
Lecture Notes on Game Theory

Theory and Examples

Xiang Sun



November 9, 2016

Contents

Acknowledgement	vii
1 Introduction	1
1.1 Timeline of the main evolution of game theory	1
1.2 Nobel prize laureates	16
1.3 Potential Nobel prize winners	19
1.4 Rational behavior	20
1.5 Common knowledge	20
2 Strategic games with complete information	23
2.1 Strategic games	23
2.2 Nash equilibrium	24
2.3 Examples	25
2.4 Existence of a Nash equilibrium	42
2.5 Symmetric games	43
2.6 Strictly competitive games (zero-sum games)	44
2.7 Contest	45
2.8 Strategic games with discontinuous payoff functions	47
3 Bayesian games	49
3.1 Bayes' rule (Bayes' theorem)	49
3.2 Bayesian games	50
3.3 Examples	54
3.4 Bilateral trading (Double auction)	66
3.5 Providing a public good under incomplete information	71

3.6	Comments on Bayesian games	71
4	Auction	75
4.1	Preliminary	75
4.2	The symmetric model	77
4.3	Second-price sealed-bid auction	77
4.4	First-price sealed-bid auction	78
4.5	Revenue comparison	86
4.6	Reserve prices	88
4.7	The revenue equivalence principle	91
4.8	All-pay auction	92
4.9	Third-price auction	93
4.10	Uncertain number of bidders	95
5	Mixed-strategy Nash equilibrium	97
5.1	Mixed-strategy Nash equilibrium	97
5.2	Examples	99
5.3	Interpretation of mixed-strategy Nash equilibrium	102
5.3.1	Purification	103
6	Correlated equilibrium	109
6.1	Motivation	109
6.2	Correlated equilibrium	110
6.3	Examples	113
7	Rationalizability	117
7.1	Rationalizability	117
7.2	Iterated elimination of never-best response	121
7.3	Iterated elimination of strictly dominated actions	122
7.4	Examples	125
7.5	Iterated elimination of weakly dominated actions	128
8	Knowledge model	131
8.1	A model of knowledge	131
8.2	Common knowledge	136
8.3	Common prior	137

8.4	“Agree to disagree” is impossible	138
8.5	No-trade theorem	140
8.6	Speculation	141
8.7	Characterization of the common prior assumption	143
8.8	Unawareness	145
9	Interactive epistemology	147
9.1	Epistemic conditions for Nash equilibrium	147
9.2	Epistemic foundation of rationalizability	150
9.3	Epistemic foundation of correlated equilibrium	151
9.4	The electronic mail game	151
10	Extensive games with perfect information	155
10.1	Extensive games with perfect information	155
10.2	Subgame perfect equilibrium	157
10.3	Examples	160
10.4	Three notable games	167
10.5	Iterated elimination of weakly dominated strategies	169
10.6	Forward induction	170
11	Bargaining games	173
11.1	A bargaining game of alternating offers	173
11.2	Bargaining games with finite horizon	174
11.3	Bargaining games with infinite horizon	175
11.4	Properties of subgame perfect equilibria in Rubinstein bargaining games	178
11.5	Bargaining games with cost	179
11.6	n -person bargaining games	179
12	Repeated games	183
12.1	Infinitely repeated games	183
12.2	Trigger strategy equilibrium	188
12.3	Tit-for-tat strategy equilibrium	193
12.4	Folk theorem	194
12.5	Nash-threats folk theorem	195
12.6	Perfect folk theorem	196
12.7	Finitely repeated games	199

13 Extensive games with imperfect information	201
13.1 Extensive games with imperfect information	201
13.2 Mixed and behavioral strategies	203
13.3 Subgame perfect equilibrium	206
13.4 Perfect Bayesian equilibrium	212
13.5 Sequential equilibrium	217
13.6 Trembling hand perfect equilibrium	221
14 Information economics 1: signaling	223
14.1 Adverse selection	223
14.2 Signaling	226
14.3 Job-market signaling	238
14.4 Cheap talk	240
14.4.1 The model with two types	242
14.4.2 The model with three types	244
14.4.3 The model with a continuum of types	244
14.4.4 Optimal communication mechanism (the commitment case)	248
14.5 Bayesian persuasion	252
14.5.1 Example	252
14.5.2 A model of Bayesian persuasion	253
14.5.3 Simplified problem	255
14.5.4 Optimal information structure	257
14.6 Optimal information disclosure	260
15 Information economics 2: screening	261
15.1 The principal-agent model	261
15.1.1 Complete information	262
15.1.2 Incomplete information	265
15.1.3 The revelation principle	269
15.1.4 General utility function for the agent	270
15.1.5 Nonresponsiveness	271
15.2 Pricing a single indivisible good	272
15.3 Nonlinear pricing	275
15.3.1 The model with two types	276
15.3.2 The model with a continuum of types	280

15.4 Insurance	282
15.5 Sequential screening	282
15.5.1 A motivating example	283
15.5.2 Optimal menu of refund contracts: two-type case	283
15.5.3 Sequential screening: general type case	287
15.5.4 Optimal selling mechanism	290
16 Information economics 3: moral hazard	291
16.1 The principle-agent problem	291
16.2 Complete information	293
16.3 Incomplete information with risk-neutral agent	294
16.4 Incomplete information with limited liability	295
16.5 Incomplete information with risk-averse agent	296
16.6 n levels of performance	298
16.6.1 Limited liability	298
16.6.2 Risk aversion	299
17 Social choice theory	301
17.1 Social choice	301
17.2 Arrow's impossibility theorem	302
17.3 Borda count, simple plurality rule, and two-round system	306
17.4 Gibbard-Satterthwaite theorem	309
18 Dominant strategy mechanism design	315
18.1 Envelope theorem	315
18.2 A general mechanism design setting	317
18.3 Revelation principle	319
18.4 Payoff equivalence	320
18.5 VCG mechanism	321
18.5.1 Pivot mechanism	323
18.5.2 Balancing the budget	326
18.6 Characterization of dominant strategy incentive compatibility	327
19 Mechanism design: Auction	329
19.1 Set-up	329
19.2 The revelation principle for Bayesian equilibrium	331

19.3 Incentive compatibility and individual rationality	332
19.4 Expected revenue maximization	335
19.5 Maximizing welfare	339
19.5.1 VCG mechanism	341
19.5.2 AGV mechanism	343
19.6 Dominant strategy mechanism	344
19.7 Bilateral trade	345
20 Mechanism design: Provision of public goods	349
20.1 Set-up	349
20.2 Ex ante and ex post budget balance	350
20.3 Welfare maximization	352
21 Implementation theory	355
21.1 Implementation	355
21.2 Implementation in dominant strategies	357
21.3 Nash implementation	360
22 Coalitional games	365
22.1 Coalitional game	365
22.2 Core	366
22.3 Shapley value	371
22.4 Nash bargaining solution	376
23 Supermodular games	387
23.1 Lattice	387
23.2 Supermodular function	391
23.3 Supermodular games	392
24 Large games 1: large strategic games	395
25 Large games 2: large distributional games	397
26 Stochastic games	399
Bibliography	401

Acknowledgement

This note was initiated in autumn of 2013 when I gave lectures at Wuhan University, and revised in autumn of 2014. I would like to thank Yi-Chun Chen, Qiang Fu, Wei He (何晔), Qian Jiao (焦倩), Bin Liu (刘斌), Xiao Luo (罗晓), Yeneng Sun (孙业能), Yifei Sun (孙一飞), Qianfeng Tang (唐前锋) and Haomiao Yu for discussion and encouragement.

I am also grateful for the following teaching assistants and students who provide lots of helpful suggestions and comments: Yiyang Cai (蔡熠阳), Yingfeng Ding (丁映峰), Siqi Fan (范思琪), Meimei Hu (胡美妹), Xiashuai Huang (黄夏帅), An Li (李安), Ming Li (黎明), Qian Li (李茜), Xiao Lin (林潇), Dawei Liu (刘大伟), Weizhao Liu (刘维钊), Yuting Liu (刘雨婷), Ji Lu (陆劼), Aixia Qu (瞿爱霞), Tianchen Song (宋天辰), Yue Teng (滕越), Deli Wang (王德利), Dexin Wang (王得鑫), Rui Wang (汪瑞), Wei Wang (王玮), Zijia Wang (汪紫珈), Ziwei Wang (王子伟), Ya Wen (文雅), Yue Wu (伍玥), Jiang Xiang (项江), Ran Xiao (肖然), Qian Xie (谢倩), Tianyang Zhang (张天洋), Wangyue Zhang (张望月), Yang Zhang (张杨).

Chapter 1

Introduction

Contents

1.1	Timeline of the main evolution of game theory	1
1.2	Nobel prize laureates	16
1.3	Potential Nobel prize winners	19
1.4	Rational behavior	20
1.5	Common knowledge	20

Game theory is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers interact. It is concerned with general analysis of strategic interaction among individuals.

1.1 Timeline of the main evolution of game theory

1.1 Reference: [A Chronology of Game Theory](#) by Paul Walker.

1.2 In 1838, the book *Researches into the Mathematical Principles of the Theory of Wealth* by [Antoine Augustin Cournot](#) (安托万·奥古斯丁·库尔诺).

In Chapter 7 of the book, “On the competition of producers”, Cournot discussed the special case of duopoly and utilises a solution concept that is a restricted version of the Nash equilibrium.

1.3 In 1913, [Zermelo’s theorem](#) by [Ernst Zermelo](#) (恩斯特·策梅洛).

✉ Ernst Zermelo, Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels, in *Proceedings of the Fifth International Congress of Mathematicians*, volume II (E. W. Hobson and A. E. H. Love, eds.), 501–504, Cambridge, Cambridge University Press, 1913.



Figure 1.1: Ernst Zermelo.

This theorem is the first theorem of game theory asserts that in any finite two-person game of perfect information in which the players move alternately and in which chance does not affect the decision making process, if the game can not end in a draw, then one of the two players must have a winning strategy. More formally, every finite extensive-form game exhibiting full information has a Nash equilibrium that is discoverable by backward induction. If every payoff is unique, for every player, this backward induction solution is unique.

When applied to chess, Zermelo's theorem states "either white can force a win, or black can force a win, or both sides can force at least a draw."

For more details of Zermelo's theorem, see [Zermelo and the early history of game theory](#) by Ulrich Schwalbe and Paul Walker.

1.4 In 1928, *Zur Theorie der Gesellschaftsspiele* (团队游戏之理论) by [John von Neumann](#) (约翰·冯·诺伊曼).

✉ John von Neumann, *Zur Theorie der Gesellschaftsspiele*, *Mathematische Annalen* **100** (1928), 295–320.



Figure 1.2: John von Neumann.

John von Neumann proved the minimax theorem in this paper. It states that every two-person zero-sum game with finitely many pure strategies for each player is determined, *i.e.* when mixed strategies are admitted, this variety of game has precisely one individually rational payoff vector. This paper also introduced the extensive form of a game.

1.5 In 1944, the book *Theory of Games and Economic Behavior* (博弈论与经济行为) by [John von Neumann](#) (约翰·

冯·诺伊曼) and [Oskar Morgenstern](#).




Figure 1.3: 60th anniversary edition (2004) of the book *Theory of Games and Economic Behavior*.

This book is considered the groundbreaking text that created the interdisciplinary research field of game theory. As well as expounding two-person zero sum theory this book is the seminal work in areas of game theory such as the notion of a cooperative game, with transferable utility, its coalitional form and its von Neumann-Morgenstern stable sets. It was also the account of axiomatic utility theory given here that led to its wide spread adoption within economics.

1.6 In 1950, Melvin Dresher and Merrill Flood carry out, at the [Rand Corporation](#), the experiment which introduced the game now known as the prisoner's dilemma. The famous story associated with this game is due to [Albert W. Tucker](#) (阿尔伯特·塔克). Howard Raiffa independently conducted, unpublished, experiments with the prisoner's dilemma.

1.7 In 1950, Nash's equilibrium points by [John Forbes Nash, Jr.](#) (约翰·福布斯·纳什).

📖 John Nash, Equilibrium points in N -person games, *Proceedings of the National Academy of Sciences of the United States of America* **36** (1950), 48–49.

📖 John Nash, Non-cooperative games, *Annals of Mathematics* **54** (1951), 286–295.




Figure 1.4: John Forbes Nash, Jr.

Nash earned a doctorate in 1950 with a 28-page dissertation on non-cooperative games. The thesis, which was

written under the supervision of doctoral advisor [Albert W. Tucker](#), contained the definition and properties of what would later be called the “Nash equilibrium”. It’s a crucial concept in non-cooperative games, and won Nash the Nobel prize in economics in 1994.

In an equilibrium no player can profitably deviate, given the other players’ equilibrium behavior.

Example: Prisoner’s dilemma. There is unique Nash equilibrium: (Confess, Confess).

	Don’t confess	Confess
Don’t confess	3, 3	0, 4
Confess	4, 0	1, 1

Figure 1.5: Prisoner’s dilemma.



Figure 1.6: Theatrical release poster of the movie “[A beautiful mind](#) (美丽心灵)”.

1.8 In 1950, [Nash bargaining solution](#) by [John Forbes Nash, Jr.](#) (约翰·福布斯·纳什).

☞ John Nash, The bargaining problem, *Econometrica* **18** (1950), 155–162.

☞ John Nash, Two person cooperative games, *Econometrica* **21** (1953), 128–140.

The Nash bargaining game is a simple two-player game used to model bargaining interactions. John Nash proposed that a solution should satisfy certain axioms (Invariant to affine transformations, Pareto optimality, Independence of irrelevant alternatives, Symmetry).

John Nash also gave a equivalent characterization for this solution. Let u and v be the utility functions of players 1 and 2, respectively. In the Nash bargaining solution, the players will seek to maximize $(u(x) - u(d)) \cdot (v(y) - v(d))$, where $u(d)$ and $v(d)$, are the *status quo* utilities (*i.e.* the utility obtained if one decides not to bargain with the other player).

Further reading:

- [John Nash’s Contribution to Economics](#), Roger B. Myerson, *Games and Economic Behavior* **14** (1996), 287–295.

1.9 1950–1953, [Harold W. Kuhn](#) provided the formulation of extensive games which is currently used, and also some basic theorems pertaining to this class of games.

☞ Harold W. Kuhn, Extensive Games, *Proceedings of the National Academy of Sciences of the United States of America* **36** (1950), 570–576.

- ☞ Harold W. Kuhn, Extensive Games and the Problem of Information, in *Contributions to the Theory of Games*, volume II (Annals of Mathematics Studies, 28) (H. W. Kuhn and A. W. Tucker, eds.), 193–216, Princeton: Princeton University Press, 1953.

Extensive games allow the modeler to specify the exact order in which players have to make their decisions and to formulate the assumptions about the information possessed by the players in all stages of the game.

1.10 In 1953, [Shapley value](#) by [Lloyd Stowell Shapley](#) (劳埃德·斯托韦尔·沙普利).

- ☞ Lloyd Shapley, A value for n -person games, in *Contributions to the Theory of Games*, volume II (Annals of Mathematics Studies, 28) (H. W. Kuhn and A. W. Tucker, eds.), *Annals of Mathematical Studies* 28, 307–317, Princeton University Press, 1953.



Figure 1.7: Lloyd Stowell Shapley.

Shapley value is a solution concept in cooperative game theory. To each cooperative game Shapley value assigns a unique distribution (among the players) of a total surplus generated by the coalition of all players.

Shapley also showed that the Shapley value is uniquely determined by a collection of desirable properties or axioms.

Further reading:

- [罗斯是沙普利的果实](#), 巫和懋, 《南方周末》, 2012 年 10 月 19 日。
- 我的导师获诺贝尔奖, 姚顺添。

1.11 In 1953, [stochastic game](#) by [Lloyd Stowell Shapley](#) (劳埃德·斯托韦尔·沙普利).

- ☞ Lloyd Shapley, Stochastic games, *Proceedings of the National Academy of Sciences of the United States of America* 39 (1953), 1095–1100.

Stochastic game is a dynamic game with probabilistic transitions played by one or more players. The game is played in a sequence of stages. At the beginning of each stage the game is in some state. The players select actions and each player receives a payoff that depends on the current state and the chosen actions. The game then moves to a new random state whose distribution depends on the previous state and the actions chosen by the players. The procedure is repeated at the new state and play continues for a finite or infinite number of stages. The total payoff to a player is often taken to be the discounted sum of the stage payoffs or the limit inferior of the averages of the stage payoffs.

Shapley showed that for the strictly competitive case, with future payoff discounted at a fixed rate, such games are determined and that they have optimal strategies that depend only on the game being played, not on the history or even on the date, *i.e.*, the strategies are stationary.

1.12 In 1960, [mechanism design](#) by [Leonid Hurwicz](#) (里奥尼德·赫维茨).

- ☞ Leonid Hurwicz, Optimality and informational efficiency in resource allocation processes, in *Mathematical Methods in the Social Sciences* (Arrow, Karlin and Suppes eds.), Stanford University Press, 1960.




Figure 1.8: Leonid Hurwicz.




Figure 1.9: The Stanley Reiter diagram.

The [Stanley Reiter](#) diagram above illustrates a game of mechanism design. The upper-left space Θ depicts the type space and the upper-right space X the space of outcomes. The social choice function $f(\theta)$ maps a type profile to an outcome. In games of mechanism design, agents send messages M in a game environment g . The equilibrium in the game $\xi(M, g, \theta)$ can be designed to implement some social choice function $f(\theta)$.

A communication system in which participants send messages to each other and/or to a “message center”, and where a pre-specified rule assigns an outcome (such as an allocation of goods and services) for every collection of received messages.

Several Chinese articles about Leonid Hurwicz by Quoqiang Tian:

- [田国强谈导师 2007 年诺贝尔经济学奖获得者赫维茨教授](#), 2007 年 10 月 16 日。
- [田国强眼中的赫维茨教授：关心中国、关注游戏规则](#), 《金融界网》, 2007 年 10 月 16 日。
- [田国强评论赫维茨教授研究成果和学术地位](#), 《金融界网》, 2007 年 10 月 16 日。
- [田国强：回忆恩师赫维茨](#), 《南方周末》, 2007 年 10 月 18 日。
- [媒体聚焦诺奖之赫维茨](#), 《第一财经日报》, 《上海证券报》, 2007 年 10 月 17 日。
- [田国强：赫维茨走了，但是他所开创的时代远未逝去](#), 《财经网》, 2008 年 7 月 4 日。

1.13 In 1961, Vickrey auction by [William Vickrey](#) (威廉·维克里) .

- ② William Vickrey, Counterspeculation, auctions, and competitive sealed tenders, *The Journal of Finance* **16** (1961), 8–37.



Figure 1.10: William Vickrey.

A Vickrey auction is a type of sealed-bid auction. Bidders submit written bids without knowing the bid of the other people in the auction. The highest bidder wins but the price paid is the second-highest bid. The auction was first described academically by William Vickrey in 1961 though it had been used by stamp collectors since 1893. This type of auction is strategically similar to an English auction and gives bidders an incentive to bid their true value.

A Vickrey–Clarke–Groves (VCG) auction is a generalization of a Vickrey auction for multiple items, which is named after William Vickrey, Edward H. Clarke, and Theodore Groves for their papers that successively generalized the idea.

1.14 In 1962, deferred-acceptance algorithm by [David Gale](#) and [Lloyd Stowell Shapley](#) (劳埃德·斯托韦尔·沙普利).

- ② David Gale and Lloyd Shapley, College admissions and the stability of marriage, *The American Mathematical Monthly* **69** (1962), 9–15.




Figure 1.11: David Gale.

Gale and Shapley asked whether it is possible to match m women with m men so that there is no pair consisting of a woman and a man who prefer each other to the partners with whom they are currently matched. They proved not only non-emptiness but also provided an algorithm for finding a point in it.

Example: 3 students $S = \{1, 2, 3\}$, 2 colleges $C = \{a, b\}$. Students' preferences: $P_1 : b, a, \emptyset$; $P_2 : a, \emptyset$; $P_3 : a, b, \emptyset$. Colleges' preferences and quotas $P_a : 1, 2, 3, q_a = 1$; $P_b : 3, 1, 2, q_b = 1$. Outcome:

	day 1	day 2	day 3
a	2, 3	2	1, 3
b	1	3 , 3	3
\emptyset	3	1	2

1.15 In 1965, [subgame perfect equilibrium](#) by [Reinhard Selten](#) (赖因哈德·泽尔腾).

☞ Reinhard Selten, Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, *Zeitschrift für die Gesamte Staatswissenschaft* 121 (1965), 301–24 and 667–89.



Figure 1.12: Reinhard Selten.

Nash equilibria that rely on non-credible threats or promises can be eliminated by the requirement of subgame perfection.

Example:




Figure 1.13

(L, RR') is a Nash equilibrium but not a subgame perfect equilibrium.

1.16 In 1967, [Bayesian games \(games with incomplete information\)](#) by [John Charles Harsanyi](#) (约翰·查理斯·海萨尼).

- ☞ John Charles Harsanyi, Games with incomplete information played by “Bayesian” players, *Management Science* 14 (1967–68) 159–182, 320–334, and 486–502, Parts I–III.




Figure 1.14: John Charles Harsanyi.

In game theory, a Bayesian game is one in which information about characteristics of the other players (*i.e.* payoffs) is incomplete. Following John C. Harsanyi's framework, a Bayesian game can be modelled by introducing Nature as a player in a game. Nature assigns a random variable to each player which could take values of types for each player and associating probabilities or a probability density function with those types (in the course of the game, nature randomly chooses a type for each player according to the probability distribution across each player's type space).

Harsanyi's approach to modelling a Bayesian game in such a way allows games of incomplete information to become games of imperfect information (in which the history of the game is not available to all players). The type of a player determines that player's payoff function and the probability associated with the type is the probability that the player for whom the type is specified is that type. In a Bayesian game, the incompleteness of information means that at least one player is unsure of the type (and so the payoff function) of another player.

Such games are called Bayesian because of the probabilistic analysis inherent in the game. Players have initial beliefs about the type of each player (where a belief is a probability distribution over the possible types for a player) and can update their beliefs according to [Bayes' rule](#) as play takes place in the game, *i.e.* the belief a player holds about another player's type might change on the basis of the actions they have played.

- 1.17 In 1967, [rent-seeking](#) by [Gordon Tullock](#) (戈登·图洛克).

- ☞ Gordon Tullock, The welfare costs of tariffs, monopolies, and theft, *Western Economic Journal* 5:3 (1967) 224–232.



Figure 1.15: Gordon Tullock.

Rent-seeking is spending wealth on political lobbying to increase one's share of existing wealth without creating wealth. The effects of rent-seeking are reduced economic efficiency through poor allocation of resources, reduced wealth creation, lost government revenue, increased income inequality, and national decline.

1.18 In 1972, [incentive compatibility](#) by [Leonid Hurwicz](#) (里奥尼德·赫维茨).

☞ Leonid Hurwicz, On informationally decentralized systems, in *Decision and Organization* (Radner and McGuire eds.), North-Holland, Amsterdam, 1972.

In mechanism design, a process is incentive compatible if all of the participants fare best when they truthfully reveal any private information asked for by the mechanism.

1.19 In 1972, the journal [International Journal of Game Theory](#) was founded by Oskar Morgenstern.

1.20 In 1970s, [revelation principle](#) by [Partha Dasgupta](#), [Allan Gibbard](#), [Peter Hammond](#), M. Harris, [Bengt R. Holmström](#), [Eric Stark Maskin](#) (埃里克·马斯金), [Roger Bruce Myerson](#) (罗杰·梅尔森), [Robert W. Rosenthal](#), R. Townsend, etc.

☞ Allan Gibbard, Manipulation of voting schemes: a general result, *Econometrica* **41** (1973), 587–602.

☞ Partha Dasgupta, Peter Hammond and Eric Maskin, The implementation of social choice rules: some general results on incentive compatibility, *Review of Economic Studies* **46** (1979), 181–216.

☞ M. Harris and R. Townsend, Resource allocation under asymmetric information, *Econometrica* **49** (1981), 33–64.

☞ Bengt R. Holmström, On incentives and control in organizations, Ph.D. dissertation, Stanford University, 1977.

☞ Roger Myerson, Incentive compatibility and the bargaining problem, *Econometrica* **47** (1979), 61–73.

☞ Roger Myerson, Optimal coordination mechanisms in generalized principal agent problems, *Journal of Mathematical Economics* **11** (1982), 67–81.

☞ Roger Myerson, Multistage games with communication, *Econometrica* **54** (1986), 323–358.

☞ Robert W. Rosenthal, Arbitration of two-party disputes under uncertainty, *Review of Economic Studies* **45** (1978), 595–604.




Figure 1.16

The revelation principle is an insight that greatly simplifies the analysis of mechanism design problems. In force of this principle, the researcher, when searching for the best possible mechanism to solve a given allocation problem, can restrict attention to a small subclass of mechanisms, so-called direct mechanisms. While direct mechanisms are not intended as descriptions of real-world institutions, their mathematical structure makes them relatively easy to analyze. Optimization over the set of all direct mechanisms for a given allocation problem is a well-defined mathematical task, and once an optimal direct mechanism has been found, the researcher can “translate back” that mechanism to a more realistic mechanism. By this seemingly roundabout method, researchers have been able to solve problems of institutional design that would otherwise have been effectively intractable. The first version of the revelation principle was formulated by Gibbard (1973). Several researchers independently extended it to the general notion of Bayesian Nash equilibrium (Dasgupta, Hammond and Maskin, 1979, Harris and Townsend, 1981, Holmstrom, 1977, Myerson, 1979, Rosenthal, 1978). Roger Myerson (1979, 1982, 1986) developed the principle in its greatest generality and pioneered its application to important areas such as regulation and auction theory.

1.21 In 1970s, [implementation theory](#) by [Eric Stark Maskin](#) (埃里克·马斯金), etc.

📖 Eric Maskin, Nash equilibrium and welfare optimality. Paper presented at the summer workshop of the Econometric Society in Paris, June 1977. Published 1999 in the *Review of Economic Studies* 66, 23–38.

The revelation principle is extremely useful. However, it does not address the issue of multiple equilibria. That is, although an optimal outcome may be achieved in one equilibrium, other, sub-optimal, equilibria may also exist. There is, then, the danger that the participants might end up playing such a sub-optimal equilibrium. Can a mechanism be designed so that all its equilibria are optimal? The first general solution to this problem was given by Eric Maskin (1977). The resulting theory, known as implementation theory, is a key part of modern mechanism design.

1.22 In 1974, [correlated equilibrium](#) by [Robert John Aumann](#) (罗伯特·约翰·奥曼).

📖 Robert John Aumann, Subjectivity and correlation in randomized strategies, *Journal of Mathematical Economics* 1 (1974), 67–96.



Example: In the following game, there are three Nash equilibria. The two pure-strategy Nash equilibria are (T, R) and (B, L) . There is also a mixed-strategy equilibrium $(\frac{2}{3} \circ T + \frac{1}{3} \circ B, \frac{2}{3} \circ L + \frac{1}{3} \circ R)$.

		Player 2	
		L	R
Player 1	T	6, 6	2, 7
	B	7, 2	0, 0

		L	R
		$p(y) = \frac{1}{3}$	$p(z) = \frac{1}{3}$
	T		
	B	$p(x) = \frac{1}{2}$	0

Since neither player has an incentive to deviate, this is a correlated equilibrium. Interestingly, the expected payoff for this equilibrium is $7 \cdot \frac{1}{3} + 2 \cdot \frac{1}{3} + 6 \cdot \frac{1}{3} = 5$ which is higher than the expected payoff of the mixed-strategy Nash equilibrium.

- Within the framework of partitional information structures, Aumann demonstrates the impossibility of agreeing to disagree: For any posteriors with a common prior, if the agents' posteriors for an event E are different (= they disagree), then the agents can not have common knowledge (= agreeing), of these posteriors.

An event is common knowledge among a set of agents if all know it and all know that they all know it and so on ad infinitum. Although the idea first appeared in the work of the philosopher D. K. Lewis in the late 1960s it was not until its formalisation in Aumann's paper that game theorists and economists came to fully appreciate its importance.

1.25 In 1982, Rubinstein bargaining game by [Ariel Rubinstein](#) (阿里埃勒·鲁宾斯坦).

📖 Ariel Rubinstein, Perfect equilibrium in a bargaining model, *Econometrica* 50 (1982), 97–110.



Figure 1.19: Ariel Rubinstein.

A Rubinstein bargaining game refers to a class of bargaining games that feature alternating offers through an infinite time horizon. Rubinstein considered a non-cooperative approach to bargaining. He considered an alternating-offer game where offers are made sequentially until one is accepted. There is no bound on the number of offers that can be made but there is a cost to delay for each player. Rubinstein showed that the subgame perfect equilibrium is unique when each player's cost of time is given by some discount factor.




Figure 1.20: A Rubinstein bargaining game.

One story for Ariel Rubinstein:

- [Sorin, Rapped](#), [economicprincipals.com](#), March 9, 2003.

- [A letter to the officers of the Game Theory Society](#), Ariel Rubinstein, December 5, 2002.

1.26 In 1982, [sequential equilibrium](#) by [David M. Kreps](#) and [Robert B. Wilson](#).

✉ David M. Kreps and Robert B. Wilson, Sequential equilibria, *Econometrica* **50** (1982), 863–894.



(a) David M. Kreps.



(b) Robert B. Wilson.

Figure 1.21

1.27 In 1984, rationalizability by [B. Douglas Bernheim](#) and D. G. Pearce.

✉ B. Douglas Bernheim, Rationalizable strategic behavior, *Econometrica* **52** (1984), 1007–1028.

✉ D. G. Pearce, Rationalizable strategic behavior and the problem of perfection, *Econometrica* **52** (1984), 1029–1050.

1.28 In 1985, construction of universal type spaces by [Jean-François Mertens](#) and [Shmuel Zamir](#).

✉ Jean-François Mertens and Shmuel Zamir, Formulation of Bayesian analysis for games with incomplete information, *International Journal of Games Theory* **14** (1985), 1–29.




Figure 1.22: Jean-François Mertens.

For a Bayesian game the question arises as to whether or not it is possible to construct a situation for which there is no sets of types large enough to contain all the private information that players are supposed to have. J.-F. Mertens and S. Zamir show that it is not possible to do so.

1.29 In 1989, the journal [Games and Economic Behavior](#) was founded.

1.30 In 1989, electronic mail game by [Ariel Rubinstein](#) (阿里埃勒·鲁宾斯坦).

✉ Ariel Rubinstein, The electronic mail game: a game with almost common knowledge, *American Economic Review* 79 (1989), 385–391.

	A	B		A	B
A	M, M	$1, -L$		$0, 0$	$1, -L$
B	$-L, 1$	$0, 0$		$-L, 1$	M, M
	G_a (probability $1 - p$)			G_b (probability p)	

Figure 1.23: The parameters satisfy $L > M > 1$ and $p < \frac{1}{2}$.

If they choose the same action but it is the “wrong” one they get 0. If they fail to coordinate, then the player who played B gets $-L$, where $L > M$. Thus, it is dangerous for a player to play B unless he is confident enough that his partner is going to play B as well.

Case 1: The true game is known initially only to player 1, but not to player 2. we can model this situation as a Bayesian game that has a unique Bayesian Nash equilibrium, in which both players always choose A .

Case 2: The game is common knowledge between two players, then it has a Nash equilibrium in which each player chooses A in state a and B in state b .

Case 3:

- The true game is known initially only to player 1, but not to player 2.
- Player 1 can communicate with player 2 via computers if the game is G_b . There is a small probability $\epsilon > 0$ that any given message does not arrive at its intended destination, however. (If a computer receives a message then it automatically sends a confirmation; this is so not only for the original message but also for the confirmation, the confirmation of the confirmation, and so on)
- If a message does not arrive then the communication stops.
- At the end of communication, each player’s screen displays the number of messages that his machine has sent.
- This game has a unique Bayesian Nash equilibrium in which both players choose A .

Rubinstein’s electronic mail game tells that players’ strategic behavior under “almost common knowledge” may be very different from that under common knowledge. Even if both players know that the game is G_b and the noise ϵ is arbitrarily small, the players act as if they had no information and play A , as they do in the absence of an electronic mail system.

1.31 In 1991, [perfect Bayesian equilibrium](#) by [Drew Fudenberg](#) (朱·弗登博格) and [Jean Tirole](#) (让·梯若尔).

✉ Drew Fudenberg and Jean Tirole, Perfect Bayesian equilibrium and sequential equilibrium, *Journal of Economic Theory* 53 (1991), 236–260.



Figure 1.24: Jean Tirole.


Further reading:

- 梯若尔：中国不应重演别国失误，人民网，2005 年。
- 专访诺奖得主梯若尔，财新《新世纪》，2014 年。


1.32 In 1999, [Game Theory Society](#) was founded.

1.2 Nobel prize laureates

1.33 In 1994, [John C. Harsanyi](#) (University of California at Berkeley), [John F. Nash Jr.](#) (Princeton University) and [Reinhard Selten](#) (University of Bonn) were awarded the Nobel Prize, “for their pioneering analysis of equilibria in the theory of non-cooperative games.”



(a) John C. Harsanyi




(b) John F. Nash Jr.



(c) Reinhard Selten

Figure 1.25

1.34 In 1996, [James Alexander Mirrlees](#) (University of Cambridge) and [William Spencer Vickrey](#) (Columbia University) were awarded the Nobel Prize, “for their fundamental contributions to the economic theory of incentives under asymmetric information.”



(a) James Alexander Mirrlees

(b) William Spencer Vickrey

Figure 1.26

- 1.35 In 2005, [Robert J. Aumann](#) (Hebrew University of Jerusalem, Stony Brook University) and [Thomas C. Schelling](#) (University of Maryland) were awarded the Nobel Prize, “for having enhanced our understanding of conflict and cooperation through game-theory analysis.”




(a) Robert J. Aumann

(b) Thomas C. Schelling

Figure 1.27

- 1.36 In 2007, [Leonid Hurwicz](#) (Minnesota University), [Eric S. Maskin](#) (Harvard University, Princeton University) and [Roger B. Myerson](#) (Northwestern University, Chicago University) were awarded the Nobel Prize, “for having laid the foundations of mechanism design theory.”



(a) Leonid Hurwicz



(b) Eric S. Maskin



(c) Roger B. Myerson

Figure 1.28

1.37 In 2012, [Alvin E. Roth](#) (Harvard University, Stanford University) and [Lloyd S. Shapley](#) (University of California at Los Angeles) were awarded the Nobel Prize, “for the theory of stable allocations and the practice of market design.”



(a) Alvin E. Roth



(b) Lloyd S. Shapley

Figure 1.29

1.38 In 2014, [Jean Tirole](#) (Toulouse 1 Capitole University) was awarded the Nobel Prize, “for his analysis of market power and regulation.”



Figure 1.30: Jean Tirole.

1.3 Potential Nobel prize winners



(a) Oliver Hart (Harvard)



(b) Bengt Holmström (MIT)



(c) David Kreps (Stanford)

Figure 1.31



(a) Paul Milgrom (Stanford)



(b) Ariel Rubinstein (Tel Aviv, NYU)



(c) Robert Wilson (Stanford)

Figure 1.32

1.4 Rational behavior

1.39 The basic assumptions that underlie game theory are that decision-makers pursue well-defined exogenous objectives (they are rational) and take into account their knowledge or expectations of other decision-makers' behavior (they are reason strategically).

1.40 A model of rational choice:

- A : set of actions, with typical element a ;
- Ω : set of states, with typical element ω ;
- C : set of outcomes;
- g : outcome function $g: A \times \Omega \rightarrow C$;
- u : utility function $u: C \rightarrow \mathbb{R}$.

1.41 A decision-maker is *rational* if the decision-maker chooses an action $a^* \in A$ that maximizes the expected value of $u(g(a, \omega))$, with respect to some probability distribution μ , i.e., a^* solves

$$\max_{a \in A} \mathbf{E}^\mu[u(g(a, \cdot))].$$

1.5 Common knowledge

1.42 E is common knowledge to players 1 and 2 if

- 1 knows E and 2 knows E ;
- 1 knows that 2 knows E and 2 knows that 1 knows E ;
- 1 knows that 2 knows that 1 knows E and 2 knows that 1 knows that 2 knows E ;
- 1 knows that 2 knows that 1 knows that 2 knows E and 2 knows that 1 knows that 2 knows that 1 knows E ;
- and so on, and so on.

1.43 For example, a handshake is common knowledge between the two persons involved. When I shake hand with you, I know you know I know you know ... that we shake hand. Neither person can convince the other that she does not know that they shake hand. So, perhaps it is not entirely random that we sometimes use a handshake to signal an agreement or a deal.

1.44 莊子與惠子游於濠梁之上。

莊子曰：儻魚出游從容，是魚之樂也。

惠子曰：子非魚，安知魚之樂？

莊子曰：子非我，安知我不知魚之樂？

惠子曰：我非子，固不知子矣；子固非魚也，子之不知魚之樂，全矣！

莊子·外篇·秋水

1.45 There are four kinds of men:

- (1) He who knows not and knows not he knows not: he is a fool—shun him;
- (2) He who knows not and knows he knows not: he is simple—teach him;
- (3) He who knows and knows not he knows: he is asleep—wake him;
- (4) He who knows and knows he knows: he is wise—follow him.

Arabian Proverb


Strategic games with complete information


Contents

2.1	Strategic games	23
2.2	Nash equilibrium	24
2.3	Examples	25
2.4	Existence of a Nash equilibrium	42
2.5	Symmetric games	43
2.6	Strictly competitive games (zero-sum games)	44
2.7	Contest	45
2.8	Strategic games with discontinuous payoff functions	47

2.1 Strategic games

2.1 A strategic game is a model of interactive decision-making in which each decision-maker chooses his plan of action once and for all, and these choices are made simultaneously.

 2.2 Definition: A strategy for a player is a complete plan of actions. It specifies a feasible action for the player in every contingency in which the player might be called on to act.

 2.3 Definition: A strategic game, denoted by $\langle N, (A_i), (\succsim_i) \rangle$, consists of

- a finite set N of players
- for each player $i \in N$ a non-empty set A_i of strategies
- for each player $i \in N$ a preference relation \succsim_i on $A = \times_{j \in N} A_j$.

A game is said to be finite if the set A_i of every player i is finite.

2.4 In a simultaneous-move game, the set of strategies is the same as the set of feasible actions.

In a dynamic game, the set of strategies may be different from the set of feasible actions.




Figure 2.1: Strategies and actions.

In this game, player 2 has 2 actions, L' and R' , but 4 strategies, $L'L'$, $L'R'$, $R'L'$ and $R'R'$.

- 2.5 The model places no restrictions on the set of strategies available to a player, which may be a huge set containing complicated plans that cover a variety of contingencies. The preference relation or utility function may not be continuous.
- 2.6 We often assume that \succsim_i can be represented by a payoff function $u_i: A \rightarrow \mathbb{R}$. In such a case we denote the game by $\langle N, (A_i), (u_i) \rangle$.
- 2.7 We may model a game in which the consequence of a profile is affected by an exogenous random variable; a profile $a \in A$ induces a lottery $g(a; \cdot)$ on outcomes. In this case, a preference relation \succsim_i over A can be defined as: $a \succsim_i b$ if and only if $g(b; \cdot)$ is at least as good as $g(a; \cdot)$, e.g., $E[u_i(g(a; \cdot))] \geq E[u_i(g(b; \cdot))]$.
- 2.8 A finite strategic game in which there are two players can be described conveniently in a payoff table.
- 2.9 When referring to the strategies/actions of the players in a strategic game as “simultaneous/static” we do not necessarily mean that these strategies/actions are taken at the same point in time.
- 2.10 A common interpretation of a strategic game is that it is a model of an event that occurs only once; each player knows the details of the game and the fact that all the players are “rational”, and the players choose their strategies/actions simultaneously and independently.

2.2 Nash equilibrium

- 2.11 Definition: A Nash equilibrium of a strategic game $\langle N, (A_i), (\succsim_i) \rangle$ is a profile $a^* \in A$ with the property that for every player $i \in N$ we have

$$(a_{-i}^*, a_i^*) \succsim_i (a_{-i}^*, a_i) \text{ for all } a_i \in A_i.$$

- 2.12 Interpretation: In an equilibrium no player can profitably deviate, given the other players' equilibrium behavior.
- 2.13 Once a player deviates, other players may want to deviate as well. But the definition does not require that a deviation be free from subsequent deviations.
- 2.14 A Nash equilibrium needs not to be Pareto optimal, for example, prisoners' dilemma. More generally, Nash equilibrium does not rule out the possibility that a subset of players can deviate jointly in a way that makes every player in the subset better off.
- 2.15 Shortcoming: The Nash equilibrium implicitly assumes that players know that each player is to play the equilibrium strategy. Given this knowledge, no player wants to deviate. So, there is a sort of circularity in this concept—the players behave in the way because they are supposed to behave in this way.

2.16 The Nash equilibrium can be justified in several ways:

- The players reach a self-enforcing agreement to play this way through pregame communication. Example: you may agree with a friend to meet at a particular restaurant for dinner.
- A steady-state convention evolved from some dynamic learning/evolutionary process. Example: we usually takes nodding your head to mean yes and shaking your head means no.
- In coordination games, certain equilibria just “stands out” as a focal point.

		Player 2	
		Mozart	Mahler
Player 1	Mozart	2, 2	0, 0
	Mahler	0, 0	1, 1

Figure 2.2

2.17 Define a correspondence (set-valued function) $B_i: A_{-i} \rightarrow A_i$ as follows:

$$B_i(a_{-i}) = \{a_i \in A_i \mid (a_{-i}, a_i) \succsim_i (a_{-i}, a'_i) \text{ for all } a'_i \in A_i\}.$$

The correspondence B_i is called the best-response correspondence of player i .

Define the correspondence $B: A \rightarrow A$ as follows:

$$B(a) = \times_{i \in N} B_i(a_{-i}).$$

2.18 Proposition: a^* is a Nash equilibrium if and only if $a^* \in B(a^*)$.

2.19 This alternative formulation of the definition points us to a method of finding Nash equilibria: first calculate the best-response correspondence of each player, then find a profile a^* for which $a_i^* \in B_i(a_{-i}^*)$ for all $i \in N$.

2.3 Examples

2.20 Example [OR Example 15.3]: Battle of the sexes.

Mary and Peter are deciding on an evening's entertainment, attending either the opera or a prize fight. Both of them would rather spend the evening together than apart, but Peter would rather they be together at the prize fight while Mary would rather they be together at the opera.

		Peter	
		Opera	Fight
Mary	Opera	2, 1	0, 0
	Fight	0, 0	1, 2

Figure 2.3: Battle of the sexes.

Answer. Two Nash equilibria: (Opera, Opera) and (Fight, Fight). □

2.21 Example [OR Example 16.1]: A two-person coordination game.

A coordination game has the property that players have a common interest in coordinating their actions. That is, two people wish to go out together, but in this case they agree on the more desirable concert.

		Player 2	
		Mozart	Mahler
Player 1	Mozart	2, 2	0, 0
	Mahler	0, 0	1, 1

Figure 2.4: A coordination game.

Answer. Two Nash equilibria: (Mozart, Mozart) and (Mahler, Mahler). □

2.22 Example [OR Example 16.2]: Prisoner's dilemma.

Two suspects in a crime are put into separate cells. If they both confess, each will be sentenced to three years in prison. If only one of them confesses, he will be freed and used as a witness against the other, who will receive a sentence of four years. If neither confesses, they will both be convicted of a minor offense and spend one year in prison.

	Player 2	
	Don't Confess	Confess
Don't Confess	3, 3	0, 4
Confess	4, 0	1, 1

Figure 2.5: Prisoner's dilemma.

Answer. This is a game in which there are gains from cooperation—the best outcome for the players is that neither confesses—but each player has an incentive to be a “free rider”. Whatever one player does, the other prefers Confess to Don't Confess, so that the game has a unique Nash equilibrium (Confess, Confess). □

2.23 Example [OR Example 16.3]: Hawk-Dove.

Two animals are fighting over some prey. Each can behave like a dove or like a hawk. The best outcome for each animal is that in which it acts like a hawk while the other acts like a dove; the worst outcome is that in which both animals act like hawks. Each animal prefers to be hawkish if its opponent is dovish and dovish if its opponent is hawkish.

	Player 2	
	Dove	Hawk
Dove	3, 3	1, 4
Hawk	4, 1	0, 0

Figure 2.6: Hawk-Dove.

Answer. Two Nash equilibria: (Dove, Hawk) and (Hawk, Dove). □

2.24 Example [OR Example 17.1]: Matching pennies.

Each of two people chooses either Head or Tail. If the choices differ, person 1 pays person 2 a dollar; if they are the same, person 2 pays person 1 a dollar. Each person cares only about the amount of money that he receives.

	Head	Tail
Head	1, -1	-1, 1
Tail	-1, 1	1, -1

Figure 2.7: Matching pennies.

Answer. No Nash equilibrium. □

- 2.25 Example: An old lady is looking for help crossing the street. Only one person is needed to help her; more are okay but no better than one. You and I are the two people in the vicinity who can help, each has to choose simultaneously whether to do so. Each of us will get pleasure worth of 3 from her success (no matter who helps her). But each one who goes to help will bear a cost of 1, this being the value of our time taken up in helping. Set this up as a game. Write the payoff table, and find all Nash equilibria.

Answer. We can formulate this game as follows:

- Two players: You (Player 1) and I (Player 2);
- Each player has 2 strategies: “Help” and “Not Help”.
- Payoffs:

		Player 2	
		Help	Not help
Player 1	Help	2, 2	2, 3
	Not help	3, 2	0, 0

Figure 2.8

There are two Nash equilibria: (Help, Not help) and (Not help, Help). □

- 2.26 Example: A game with three players.

There are three computer companies, each of which can choose to make large (L) or small (S) computers. The choice of company 1 is denoted by S_1 or L_1 , and similarly, the choices of companies 2 and 3 are denoted S_i or L_i of $i = 2$ or 3. The following table shows the profit each company would receive according to the choices which the three companies could make. Find all the Nash equilibria of the game.

	$S_2 S_3$	$S_2 L_3$	$L_2 S_3$	$L_2 L_3$
S_1	-10, -15, 20	0, -10, 60	0, 10, 10	20, 5, 15
L_1	5, -5, 0	-5, 35, 15	-5, 0, 15	-20, 10, 10

Figure 2.9: A game with three players.

Answer. Unique Nash equilibrium: (S_1, L_2, L_3) . □

- 2.27 Example: Two firms may compete for a given market of total value, V , by investing a certain amount of effort into the project through advertising, securing outlets, etc. Each firm may allocate a certain amount for this purpose. If firm 1 allocates $x \geq 0$ and firm 2 allocates $y \geq 0$, then the proportion of the market that firm 1 corners is $\frac{x}{x+y}$. The firms have different difficulties in allocating these resources. The cost per unit allocation to firm i is c_i , $i = 1, 2$. Thus the profits to the two firms are

$$\pi_1(x, y) = V \cdot \frac{x}{x+y} - c_1 x,$$

$$\pi_2(x, y) = V \cdot \frac{y}{x + y} - c_2 y.$$

If both x and y are zero, the payoffs to both are $\frac{V}{2}$.

Find the equilibrium allocations, and the equilibrium profits to the two firms, as functions of V , c_1 and c_2 .

Answer. It is natural to assume V , c_1 and c_2 are positive.

- (1) Given player 2's strategy $y = 0$, there is no best response for player 1: The payoff of player 1 is as follows

$$\pi_1(x, 0) = \begin{cases} V - c_1 x, & \text{if } x > 0; \\ \frac{V}{2}, & \text{if } x = 0. \end{cases}$$

Player 1 will try to choose $x \neq 0$ as close as possible to 0:

- We may choose x small enough, such that $\frac{V}{2} < V - c_1 x$, so $x = 0$ can not be a best response;
- For any $x > 0$, we will have $V - c_1 x < V - c_1 \frac{x}{2}$, so x can not be a best response.

Hence, the strategy profiles $(x, 0)$ and $(0, y)$ are not Nash equilibria. Therefore, we will assume that $x, y > 0$.

- (2) Given player 2's strategy $y > 0$, player 1's best response $x^*(y)$ should satisfy $\frac{\partial \pi_1}{\partial x}(x) = 0$ and $\frac{\partial^2 \pi_1}{\partial x^2}(x) \leq 0$, which implies

$$\frac{Vy}{(x^*(y) + y)^2} - c_1 = 0.$$

That is

$$\frac{y}{c_1} = \frac{(x^*(y) + y)^2}{V}. \quad (2.1)$$

Similarly, given player 1's strategy $x > 0$, we will get that player 2's best response $y^*(x)$ satisfies

$$\frac{x}{c_2} = \frac{(x + y^*(x))^2}{V}. \quad (2.2)$$

- (3) Let (x^*, y^*) be a Nash equilibrium, that is, x^* and y^* are best responses of each other, and hence (x^*, y^*) should satisfy Equations (2.1) and (2.2). From Equations (2.1) and (2.2), we will have

$$\frac{y^*}{c_1} = \frac{(x^* + y^*)^2}{V} = \frac{x^*}{c_2}.$$

Substitute this equation into Equations (2.1) and (2.2), we will obtain that

$$x^* = \frac{Vc_2}{(c_1 + c_2)^2}, \quad y^* = \frac{Vc_1}{(c_1 + c_2)^2}.$$

Notice that x^*, y^* are both positive, so they could be the solution of this problem. Hence (x^*, y^*) is the only Nash equilibrium.

Meanwhile, the equilibrium profits to the two firms are

$$\pi_1(x^*, y^*) = \frac{Vc_2^2}{(c_1 + c_2)^2}, \quad \pi_2(x^*, y^*) = \frac{Vc_1^2}{(c_1 + c_2)^2}.$$

□

2.28 Example [G Exercise 1.3]: Splitting a dollar.

Players 1 and 2 are bargaining over how to split one dollar. Both players simultaneously name shares they would like to have, s_1 and s_2 , where $0 \leq s_1, s_2 \leq 1$. If $s_1 + s_2 \leq 1$, then the players receive the shares they named; if $s_1 + s_2 > 1$, then both players receive zero. What are Nash equilibria of this game?

Answer. Given any $s_2 \in [0, 1)$, the best response for player 1 is $1 - s_2$, i.e., $B_1(s_2) = \{1 - s_2\}$.

To $s_2 = 1$, the player 1's best response is the set $[0, 1]$, because player 1's payoff is 0 no matter what she chooses.

The best-response correspondence for player 1:

$$B_1(s_2) = \begin{cases} \{1 - s_2\}, & \text{if } 0 \leq s_2 < 1, \\ [0, 1], & \text{if } s_2 = 1. \end{cases}$$

Similarly, we have the best response correspondence for player 2:

$$B_2(s_1) = \begin{cases} \{1 - s_1\}, & \text{if } 0 \leq s_1 < 1, \\ [0, 1], & \text{if } s_1 = 1. \end{cases}$$




Figure 2.10: Best-response correspondences.

From Figure 2.10, we know

$$\{(s_1, s_2) \mid s_1 + s_2 = 1, s_1, s_2 \geq 0\} \cup \{(1, 1)\}$$

is the set of all Nash equilibria. □

2.29 Example: Modified splitting a dollar.

Players 1 and 2 are bargaining over how to split one dollar. Both players simultaneously name shares they would like to have, s_1 and s_2 , where $0 \leq s_1, s_2 \leq 1$. If $s_1^2 + s_2^2 \leq 1/2$, then the players receive the shares they named; if $s_1^2 + s_2^2 > 1/2$, then both players receive zero. What are the Nash equilibria of this game?

Answer (1st method). Let $s = (s_1, s_2) \in [0, 1] \times [0, 1]$. We distinguish the following three cases:

- if $s_1^2 + s_2^2 < 1/2$, each player i can do better by choosing $s_i + \epsilon$. Thus, s is not a Nash equilibrium.
- if $s_1^2 + s_2^2 = 1/2$, no player can do better by unilaterally changing his/her strategy (because i 's payoff is 0 by choosing $s_i + \epsilon$). Thus, s is a Nash equilibrium.

- if $s_1^2 + s_2^2 > 1/2$, then we further distinguish two subcases:
 - if $s_i^2 < 1/2$, then j can do better by choosing $s_i + \epsilon$. Thus, s in this subcase is not a Nash equilibrium.
 - if $s_1^2 \geq 1/2$ and $s_2^2 \geq 1/2$, then no player can do better by unilaterally changing his/her strategy (because i 's payoff is always 0 if $s_j^2 \geq 1/2$). Thus, s in this subcase is a Nash equilibrium.

□

Answer (2nd method). Given player 2's strategy s_2 , the best response of player 1 is:

$$B_1(s_2) = \begin{cases} \left\{ \sqrt{\frac{1}{2} - s_2^2} \right\}, & \text{if } s_2 < \frac{1}{\sqrt{2}}; \\ [0, 1], & \text{if } s_2 \geq \frac{1}{\sqrt{2}}. \end{cases}$$

Note that if $s_2 < \frac{1}{\sqrt{2}}$, then player 1 should choose s_1 as much as possible, so that $s_1^2 + s_2^2 \leq \frac{1}{2}$. Hence, $\left\{ \sqrt{\frac{1}{2} - s_2^2} \right\}$ is player 1's best response to s_2 . If $s_2 \geq \frac{1}{\sqrt{2}}$, no matter what player 1 chooses, his payoff is always 0. Thus player 1 can choose any value between 0 and 1.

The graph of B_1 is showed in Figure 2.11a, and by symmetry, we can also get the best response of player 2, showed in Figure 2.11b.




Figure 2.11: Best-response correspondences.


Then the intersection of B_1 and B_2 is shown in Figure 2.12.

So the Nash equilibria are

$$\left\{ (s_1, s_2) \mid s_1 \geq 0, s_2 \geq 0, s_1^2 + s_2^2 = \frac{1}{2} \right\} \cup \left(\left[\frac{1}{\sqrt{2}}, 1 \right] \times \left[\frac{1}{\sqrt{2}}, 1 \right] \right).$$

□

Now we change the payoff rule as follows: If $s_1^2 + s_2^2 < 1/2$, then the players receive the shares they named; if $s_1^2 + s_2^2 \geq 1/2$, then both players receive zero. What are the Nash equilibria of this game?

Figure 2.12: Intersection of B_1 and B_2 .

Answer. Under the new payoff rules, the best response becomes:

$$B_i(s_j) = \begin{cases} \emptyset, & \text{if } s_j < \frac{1}{\sqrt{2}}; \\ [0, 1], & \text{if } s_j \geq \frac{1}{\sqrt{2}}, \end{cases}$$

where $(i, j) = (1, 2)$ or $(2, 1)$. Note that when $s_j < \frac{1}{\sqrt{2}}$, player i does not have the best response, because he will try to choose s_i as close as possible to $\sqrt{1/2 - s_j^2}$, but can not achieve $\sqrt{1/2 - s_j^2}$. The detailed discussion is as follows:

- For any $1 \geq s_i \geq \sqrt{1/2 - s_j^2}$, player i 's payoff is 0, which is less than the payoff when player i chooses $\frac{1}{2}\sqrt{1/2 - s_j^2}$; Hence such a s_i can not be a best response.
- For any $0 \leq s_i < \sqrt{1/2 - s_j^2}$, player i 's payoff is s_i , which is less than the payoff when player i chooses $\frac{s_i + \sqrt{1/2 - s_j^2}}{2}$; Hence such a s_i can not be a best response.

Therefore, the Nash equilibria are

$$\left[\frac{1}{\sqrt{2}}, 1 \right] \times \left[\frac{1}{\sqrt{2}}, 1 \right].$$

□

2.30 Example [G Section 1.2.A]: Cournot model of duopoly.

Suppose firms 1 and 2 produce the same product.

Let q_i be the quantity of the product produced by firm i , $i = 1, 2$. Let $Q = q_1 + q_2$, the aggregate quantity of the product.

Let the market clearing price be

$$P(Q) = \begin{cases} a - Q, & \text{if } Q < a, \\ 0, & \text{if } Q \geq a. \end{cases}$$

Let the cost of producing a unit of the product be c , where we assume $0 < c < a$.

How much shall each firm produce?

Answer. We need to translate the problem into a strategic game.

- The players of the game are the two firms.
- Each firm's strategy space is $S_i = [0, \infty)$, $i = 1, 2$. (Any value of q_i is a strategy.)
- The payoff to firm i as a function of the strategies chosen by it and by the other firm, is simply its profit function:

$$\pi_i(q_i, q_j) = P(q_i + q_j) \cdot q_i - c \cdot q_i = \begin{cases} q_i[a - (q_i + q_j) - c], & \text{if } q_i + q_j < a, \\ -cq_i, & \text{if } q_i + q_j \geq a. \end{cases}$$

We consider the following two cases:

- When $q_j \geq a$, $\pi_i(q_i, q_j) = -cq_i$, and hence $B_i(q_j) = \{0\}$.
- When $a > q_j > a - c$,
 - if $q_i \geq a - q_j (> 0)$, then $\pi_i(q_i, q_j) = -cq_i < 0$.
 - if $a - q_j > q_i > 0$, then $\pi_i(q_i, q_j) = q_i[a - (q_i + q_j) - c] < 0$.
 - if $q_i = 0$, then $\pi_i(q_i, q_j) = q_i[a - (q_i + q_j) - c] = 0$.

Therefore, $B_i(q_j) = \{0\}$.

In the following we only need to consider the case when $a - c \geq q_i, q_j \geq 0$:

- if $q_i + q_j \geq a$, then $\pi_i(q_i, q_j) = -cq_i < 0$.
- if $a > q_i + q_j \geq a - c$, then $\pi_i(q_i, q_j) = q_i[a - (q_i + q_j) - c] \leq 0$.
- if $a - c > q_i + q_j \geq 0$, then $\pi_i(q_i, q_j) \geq 0$, and in this case $\pi_i(q_i, q_j)$ achieves the maximum when $q_i = \frac{a - q_j - c}{2}$ which yields a positive payoff for i .

Therefore the best-response correspondence for i is

$$B_i(q_j) = \begin{cases} \left\{ \frac{a - q_j - c}{2} \right\}, & \text{if } q_j < a - c, \\ \{0\}, & \text{if } a - c \leq q_j. \end{cases}$$




Figure 2.13: Best-response correspondences.

From Figure 2.13, there is unique Nash equilibrium $(\frac{a-c}{3}, \frac{a-c}{3})$. □

2.31 Example [G Exercise 1.6]: Modified Cournot duopoly model.

Consider the Cournot duopoly model where inverse demand is $P(Q) = a - Q$ but firms have asymmetric marginal costs: c_1 for firm 1 and c_2 for firm 2. What is the Nash equilibrium if $0 < c_i < a/2$ for each firm? What if $c_1 < c_2 < a$ but $2c_2 > a + c_1$?

Answer. • Set of players: $\{1, 2\}$;

- For each i , player i 's strategy set: $S_i = [0, +\infty)$;
- For each i , player i 's payoff function:

$$\pi_i(q_i, q_j) = q_i(\max\{a - q_i - q_j, 0\} - c_i),$$

where $i \neq j$.

By similar method used in the previous examples, we will obtain player i 's best response:

$$B_i^*(q_j) = \begin{cases} \left\{ \frac{a - c_i - q_j}{2} \right\}, & \text{if } q_j \leq a - c_i; \\ \{0\}, & \text{if } q_j > a - c_i. \end{cases}$$




Figure 2.14: Intersection of best-response correspondences.

- If $0 < c_1, c_2 < \frac{a}{2}$, then $\frac{a-c_i}{2} < \frac{a}{2} < a - c_j$, where $i \neq j$. Hence we have the Figure 2.14a, and from it we will obtain the Nash equilibrium: $\left(\frac{a-2c_1+c_2}{3}, \frac{a-2c_2+c_1}{3}\right)$.
- If $0 < c_1 < c_2 < a$ and $2c_2 > a + c_1$, then $a - c_1 > a - c_2 > \frac{a-c_2}{2} > 0$ and $\frac{a-c_1}{2} > a - c_2 > 0$. Hence we have the Figure 2.14b, and from it we will obtain the Nash equilibrium: $\left(\frac{a-c_1}{2}, 0\right)$.

□

2.32 Example [G Exercise 1.4]: Cournot model with many firms.

Suppose there are n firms in the Cournot oligopoly model. Let q_i denote the quantity produced by firm i , and let $Q = q_1 + \dots + q_n$ denote the aggregate quantity on the market. Let P denote the market-clearing price and assume that inverse demand is given by $P(Q) = a - Q$ (assuming $Q < a$, else $P = 0$). Assume that the total cost of firm i from producing quantity q_i is $C_i(q_i) = cq_i$. That is, there are no fixed costs and the marginal cost is constant at c , where we assume $c < a$. Following Cournot, suppose that the firms choose their quantities simultaneously. What is the Nash equilibrium? What happens as n approaches infinity?

Answer. We assume $c > 0$.

- Set of players: $\{1, 2, \dots, n\}$;
- For each i , player i 's strategy set: $S_i = [0, +\infty)$;
- For each i , player i 's payoff function:

$$\begin{aligned}\pi_i(q_i, q_{-i}) &= q_i(\max\{a - q_i - q_{-i}, 0\} - c) \\ &= \begin{cases} (a - q_i - q_{-i} - c)q_i, & \text{if } q_i + q_{-i} < a; \\ -cq_i, & \text{if } q_i + q_{-i} \geq a, \end{cases}\end{aligned}$$

where $q_{-i} = \sum_{j \neq i} q_j$.

In the following, given q_{-i} , we try to find player i 's best response:

- (1) When $a \leq q_{-i}$, then we have $q_i + q_{-i} \geq a$, and hence

$$\pi_i(q_i, q_{-i}) = -cq_i \begin{cases} < 0, & \text{if } q_i > 0; \\ = 0, & \text{if } q_i = 0. \end{cases}$$

Therefore, in this case, the best response for player i is $q_i = 0$.

- (2) When $a - c \leq q_{-i} < a$, then we have

$$\pi_i(q_i, q_{-i}) = \begin{cases} 0, & \text{if } q_i = 0; \\ (a - q_i - q_{-i} - c)q_i < 0, & \text{if } 0 < q_i < a - q_{-i}; \\ -cq_i < 0, & \text{if } q_i \geq a - q_{-i}. \end{cases}$$

Therefore, in this case, the best response for player i is $q_i = 0$.

- (3) When $0 \leq q_{-i} < a - c$, then we have

$$\pi_i(q_i, q_{-i}) = \begin{cases} 0, & \text{if } q_i = 0; \\ (a - q_i - q_{-i} - c)q_i, & \text{if } 0 < q_i < a - q_{-i}; \\ -cq_i < 0, & \text{if } q_i \geq a - q_{-i}. \end{cases}$$

The function $(a - q_i - q_{-i} - c)q_i$ is concave for q_i , because its 2nd derivative is $-2 < 0$. The local maximum can be determined by the first order condition (the 1st derivative equals zero) $a - q_{-i} - c - 2q_i = 0$, thus the best response for player i is $\frac{a - c - q_{-i}}{2}$. Note that when player i chooses $\frac{a - c - q_{-i}}{2}$, his payoff is positive.

Therefore player i 's best response is

$$B_i^*(q_{-i}) = \begin{cases} \{0\}, & \text{if } a - c \leq q_{-i}; \\ \{\frac{a - c - q_{-i}}{2}\}, & \text{if } 0 \leq q_{-i} < a - c. \end{cases}$$

Remark: We can not draw graphs to find Nash equilibria, since there are more than 2 players.

Claim: There does not exist a Nash equilibrium in which some players choose 0. We will prove this claim by contradiction:

(1) Assume there is a Nash equilibrium $(q_1^*, q_2^*, \dots, q_n^*)$, where

$$J \equiv \{i: q_i^* = 0\} \neq \emptyset.$$

Let $J^c = \{1, 2, \dots, n\} - J$, then for any $j \in J^c$, $q_j^* = \frac{a-c-q_{-j}^*}{2}$.

(2) Since for any $i \in J$, $q_i^* = 0$, we will have $q_{-i}^* \geq a - c$, which implies $\sum_{j \in J^c} q_j^* \geq a - c$.

(3) Since for any $i \in J$, $q_i^* = 0$, we will have

$$q_{-j}^* = \sum_{k \in J^c, k \neq j} q_k^*,$$

for each $j \in J^c$, and hence

$$q_j^* = \frac{a - c - \sum_{k \in J^c, k \neq j} q_k^*}{2}, \quad \forall j \in J^c.$$

Summing this $|J^c|$ equations, we will have

$$\sum_{j \in J^c} q_j^* = \frac{a - c}{2} |J^c| - \frac{1}{2} (|J^c| - 1) \sum_{j \in J^c} q_j^*,$$

which implies

$$\sum_{j \in J^c} q_j^* = \frac{|J^c|}{|J^c| + 1} (a - c) < a - c.$$

Contradiction.

Assume that $(q_1^*, q_2^*, \dots, q_n^*)$ is a Nash equilibrium, then based on the claim above, we will have $q_i^* = \frac{a-c-q_{-i}^*}{2}$, for all $i = 1, 2, \dots, n$. Hence

$$q_i^* = a - c - Q^*, \quad \forall i = 1, 2, \dots, n,$$

where $Q^* = \sum_{i=1}^n q_i^*$. Summing the n equations above, we obtain

$$Q^* = \frac{n}{n+1} (a - c).$$

Substituting this into each of the above n equations, we obtain

$$q_1^* = q_2^* = \dots = q_n^* = \frac{a - c}{n + 1}.$$

As n approaches infinity, the total output $Q^* = \frac{n}{n+1} (a - c)$ approaches $a - c$ (perfect-competition output) and the price $a - Q^* = \frac{a+nc}{n+1}$ approaches c (the perfect-competition price). \square

2.33 Example [G Section 1.2.B]: Bertrand model of duopoly.

Suppose now the two firms produce different products. In this case, we can not use the aggregate quantity to determine market prices as in Cournot's model. Thus, instead of using quantities as variables, here we use prices as variables.

If firms 1 and 2 choose prices p_1 and p_2 , respectively, the quantity that consumers demand from firm i is

$$q_i(p_i, p_j) = a - p_i + bp_j$$

where $b > 0$ reflects the extent to which firm i 's product is a substitute for firm j 's product. Here we assume $c < a$. How to find the Nash equilibrium?

Answer. The strategy space of firm i consists of all possible prices, thus $S_i = [0, \infty)$, $i = 1, 2$.

The profit of firm i is

$$\pi_i(p_i, p_j) = q_i(p_i, p_j) \cdot p_i - c \cdot q_i(p_i, p_j) = [a - p_i + bp_j] \cdot (p_i - c).$$

For given p_j , $\pi_i(p_i, p_j)$ is a concave function in terms of p_i , and hence it achieves its maximum at $p_i = \frac{a+bp_j+c}{2}$.

Suppose (p_1^*, p_2^*) is a Nash equilibrium, then we have

$$p_1^* = \frac{a + bp_2^* + c}{2}, \text{ and } p_2^* = \frac{a + bp_1^* + c}{2}.$$

Thus, $p_1^* = p_2^* = \frac{a+c}{2-b}$. Note that this problem make sense only if $b < 2$. □

2.34 Example [G Exercise 1.7]: Suppose that the quantity that consumers demand from firm i is

$$q_i(p_i, p_j) = \begin{cases} a - p_i, & \text{if } p_i < p_j, \\ \frac{a-p_i}{2}, & \text{if } p_i = p_j, \\ 0, & \text{if } p_i > p_j, \end{cases}$$

that is, all customers buy the product from the firm who offers a lower price. Suppose also that there are no fixed costs and that marginal costs are constant at c , where $c < a$ and $c \leq q_1, q_2 \leq a$.

Answer. Given firm j 's price p_j , firm i 's payoff function is

$$\pi_i(p_i, p_j) = \begin{cases} (a - p_i)(p_i - c), & \text{if } p_i < p_j, \\ \frac{1}{2}(a - p_i)(p_i - c), & \text{if } p_i = p_j, \\ 0, & \text{if } p_i > p_j. \end{cases}$$

The strategy space is $S_i = [c, a]$.

We find three cases from the observation of the payoff curves.

- Case 1: Given $p_j \geq \frac{a+c}{2}$. The maximum payoff is reached at $p_i = \frac{a+c}{2}$. Thus, the best response $B_i(p_j) = \{\frac{a+c}{2}\}$.
- Case 2: Given $c < p_j \leq \frac{a+c}{2}$. It is easy to see that

$$\sup \pi_{p_i}(p_i, p_j) = (a - p_j)(p_j - c).$$

However, no $p_i \in [c, a]$ can make $\pi_i(p_i, p_j) = (a-p_j)(p_j-c)$. For $p_i \in (c, p_j)$, the function $\pi_i(p_i, p_j) = (a-p_i)(p_i-c)$ is strictly increasing. For $p_i > p_j$, $\pi_i(p_i, p_j) = 0$. For $p_i = p_j$, $\pi_i(p_i, p_j) = \frac{1}{2}(a-p_i)(p_i-c)$. Thus, there is no maximizer. This means that $B_i(p_j) = \emptyset$.

- Case 3: Given $p_j = c$. $\pi_i(p_i, c) = 0$ for any p_i . Thus any p_i is a maximizer, and $B_i(c) = [c, a]$.

The best-response correspondences are sketched below.




Figure 2.15

The only intersection of the two correspondences is (c, c) . This shows that if the firms choose prices simultaneously, then the unique Nash equilibrium is that both firms charge the price $p_i = c$. \square

2.35 Example [OR Exercise 18.2]: First-price auction.

An object is to be assigned to a player in the set $\{1, \dots, n\}$ in exchange for a payment. Player i 's valuation of the object is v_i , and $v_1 > v_2 > \dots > v_n > 0$. The mechanism used to assign the object is a (sealed-bid) auction: the players simultaneously submit bids (non-negative numbers), and the object is given to the player with the lowest index among those who submit the highest bid, in exchange for a payment.

Formulate a first-price auction as a strategic game and analyze its Nash equilibria. In particular, show that in all equilibria player 1 obtains the object.

Answer. The strategic game $\langle N, (A_i), (u_i) \rangle$ is: $N = \{1, \dots, n\}$, player i 's action set is $A_i = [0, \infty)$, and his payoff is

$$u_i(a) = \begin{cases} v_i - a_i, & \text{if } a_i > a_j \text{ (when } j < i), \text{ and } a_i \geq a_j \text{ (when } j > i), \\ 0, & \text{otherwise.} \end{cases}$$

Let $a^* = (a_1^*, a_2^*, \dots, a_n^*)$ be a Nash equilibrium.

- (1) Claim: $a_1^* \geq a_i^*$ for $i \neq 1$. Suppose that player i ($\neq 1$) submits the highest bid a_i^* and $a_i^* > a_1^*$. If $a_i^* > v_2$, then player i 's payoff is negative, so he can increase his payoff by bidding 0; if $a_i^* \leq v_2$, then player 1 can deviate to the bid a_i^* and increases his payoff. Hence, we have that $a_1^* \geq a_i^*$ for all $i \neq 1$.
- (2) Claim: $a_1^* \leq v_1$. Suppose $a_1^* > v_1$. By claim (1), we have that $a_1^* \geq a_i^*$ for all $i \neq 1$, then player 1 will win and his payoff is negative, while he can increase his payoff by bidding 0.
- (3) Claim: $a_1^* \geq v_2$. Suppose $a_1^* < v_2$. By claim (1), we have that $a_2^* < v_2$, then player 2 can increase his payoff by bidding $\frac{1}{2}(a_1^* + v_2)$.
- (4) Claim: there exists $j \in \{2, 3, \dots, n\}$, such that $a_j^* = a_1^*$. Suppose that for any $j \in \{2, 3, \dots, n\}$, $a_1^* > a_j^*$, then player 1 can choose $\max_{2 \leq j \leq n} a_j^*$.

Hence, the Nash equilibrium is (a_1^*, \dots, a_n^*) , where $a_1^* \in [v_2, v_1]$, $a_j^* \leq a_1^*$ for all $j \neq 1$, and $a_j^* = a_1^*$ for some $j \neq 1$.

Moreover, we can have that in all equilibria, player 1 will obtain the object. \square

2.36 Example [OR Exercise 18.5]: A war of attrition.

Two players are involved in a dispute over an object. The value of the object to player i is $v_i > 0$. Time is modeled as a continuous variable that starts at 0 and runs indefinitely. Each player chooses when to concede the object to the other player; if the first player to concede does so at time t , the other player obtains the object at that time. If both players concede simultaneously, the object is split equally between them, player i receiving a payoff of $v_i/2$. Time is valuable: until the first concession each player loses one unit of payoff per unit of time.

Formulate this situation as a strategic game and show that in all Nash equilibria one of the players concedes immediately.

Answer. The strategic game $\langle N, (A_i), (u_i) \rangle$ is: $N = \{1, 2\}$, player i 's action set is $A_i = [0, \infty)$, and his payoff is

$$u_i(t_1, t_2) = \begin{cases} -t_i, & \text{if } t_i < t_j, \\ v_i/2 - t_i, & \text{if } t_i = t_j, \\ v_i - t_j, & \text{if } t_i > t_j, \end{cases}$$

where $j \in \{1, 2\} \setminus \{i\}$.

Let $t^* = (t_1^*, t_2^*)$ be a Nash equilibrium.

- (1) Claim: $t_1^* \neq t_2^*$. Suppose $t_1^* = t_2^*$, then player 1 can obtain the object in its entire value instead of getting just half of it by conceding slightly later than t_1^* , so it is not a Nash equilibrium.
- (2) Claim: If $t_1^* < t_2^*$, then $t_1^* = 0$ and $t_2^* \geq v_1$. Suppose $0 < t_1^*$, then player 1 can increase his payoff to 0 by deviating to $t_1 = 0$. Suppose $0 = t_1^* < t_2^* < v_1$, then player 1 can increase his payoff by deviating to a time slightly after t_2 .
- (3) Claim: If $t_2^* < t_1^*$, then $t_2^* = 0$ and $t_1^* \geq v_2$. It is similar with the claim 2.

Hence, (t_1^*, t_2^*) is a Nash equilibrium if and only if $0 = t_1^* < t_2^*$ and $v_1 \leq t_2^*$, or $0 = t_2^* < t_1^*$ and $v_2 \leq t_1^*$. \square

2.37 Example [G Exercise 1.8]: Hotelling model.

Consider a population of voters uniformly distributed along the ideological spectrum from left ($x = 0$) to right ($x = 1$). Each of the candidates for a single office simultaneously chooses a campaign platform (*i.e.*, a point on the line between $x = 0$ and $x = 1$). The voters observe the candidates' choices, and then each voter votes for the candidate whose platform is closest to the voter's position on the spectrum. If there are two candidates and they choose platforms $x_1 = 0.3$ and $x_2 = 0.6$, for example, then all voters to the left of $x = 0.45$ vote for candidate 1, all those to the right vote for candidate 2, and candidate 2 wins the election with 55 percent of the vote. Suppose that the candidates care only about being elected—they do not really care about their platforms at all!

Question 1: If there are two candidates, what is the Nash equilibrium.

(Assume that any candidates who choose the same platform equally split the votes cast for that platform, and that ties among the leading vote-getters are resolved by coin flips.)

Answer (1st method). For player i , the strategy set is $S_i = [0, 1]$. Player i 's payoff function:

$$\pi_i(s_i, s_j) = \begin{cases} 1, & \text{if } s_j < s_i < 1 - s_j, \text{ or } 1 - s_j < s_i < s_j; \\ \frac{1}{2}, & \text{if } s_i = s_j, \text{ or } s_i = 1 - s_j; \\ 0, & \text{otherwise.} \end{cases}$$

Given player j 's strategy $s_j \neq \frac{1}{2}$, from Figures 2.16a and 2.16b, we will see that player i wins only when s_i is in the red regions.



Figure 2.16: Players' best response.

Therefore, we have player i 's best response:

$$B_i^*(s_j) = \begin{cases} (s_j, 1 - s_j), & \text{if } s_j < \frac{1}{2}; \\ \{\frac{1}{2}\}, & \text{if } s_j = \frac{1}{2}; \\ (1 - s_j, s_j), & \text{if } s_j > \frac{1}{2}. \end{cases}$$

From Figure 2.17, there is only one Nash equilibrium $(\frac{1}{2}, \frac{1}{2})$.




Figure 2.17: Intersection of the best-response correspondence.

□

Answer (2nd method). • Claim: Both candidates choose the same platform: if they choose different platforms, without loss of generality, we may assume $s_1 < \frac{1}{2} < s_2$, then 1 can do better by choosing $\frac{1}{2}$. (For any other possible case, it is similar.)

- Claim: Both candidates choose the same platform at 0.5: If both candidates choose the same platform at $x \neq 0.5$, say $x > 0.5$, then each candidate can do better by choosing $\frac{x+0.5}{2}$.

□

Question 2: If there are three candidates, exhibit a Nash equilibrium.

Answer. $(\frac{1}{3}, \frac{1}{3}, \frac{2}{3})$ is a Nash equilibrium. To see this is a Nash equilibrium,

- Player 3 has no incentive to deviate because he is the winner and obtains the maximal payoff;
- Players 1 and 2 can not do better given the other two players choose $\frac{1}{3}$ and $\frac{2}{3}$, respectively.

There are other equilibria.

□

Remark: suppose that the candidates care only about the percentage they have, in three-person game, there is no Nash equilibrium:

- If all candidates choose different platforms, say $s_1 < s_2 < s_3$, then 1 can do better by choosing $s_1 + \frac{s_2 - s_1}{2}$.
- If 1 and 2 choose same, but different with 3, say $s_1 = s_2 < s_3$, then 3 can do better by choosing $s_3 - \frac{s_3 - s_1}{2}$.
- If all of them choose same, say $s_1 = s_2 = s_3 > 0.5$, then 1 can do better by choosing $s_1 - \frac{s_1 - 0.5}{2}$.

2.38 Example [OR Exercise 19.1]: A location game.

Each of n people chooses whether or not to become a political candidate, and if so which position to take. There is a continuum of citizens, each of whom has a favorite position; the distribution of favorite positions is given by a density function f on $[0, 1]$ with $f(x) > 0$ for all $x \in [0, 1]$. A candidate attracts the votes of those citizens whose favorite positions are closer to his position than to the position of any other candidate; if k candidates choose the same position then each receives the fraction $1/k$ of the votes that the position attracts. The winner of the competition is the candidate who receives the most votes. Each person prefers to be the unique winning candidate than to tie for first place, prefers to tie for first place than to stay out of the competition, and prefers to stay out of the competition than to enter and lose.

Formulate this situation as a strategic game, find the set of Nash equilibria when $n = 2$, and show that there is no Nash equilibrium when $n = 3$.

Answer. □

2.39 Example [OR Exercise 35.1]: Guessing the average.

Let $n(n \geq 2)$ people play the following game. Simultaneously, each player i announces a number x_i in the set $\{1, 2, \dots, K\}$. A prize of \$1 is split equally between all the people whose number is closest to $\frac{2}{3} \cdot \frac{x_1 + \dots + x_n}{n}$. Find all the Nash equilibria.

Incomplete answer. Assume (x_1, x_2, \dots, x_n) is a Nash equilibrium, and x_1 is the largest number among them. We now argue as follows.

- In the equilibrium (x_1, x_2, \dots, x_n) , Player 1's payoff should be positive. Otherwise, he could be better off by choosing a number which is the closest number to $\frac{2}{3}$ of average.
- In the equilibrium (x_1, x_2, \dots, x_n) , there is some other, say Player $j(j \neq 1)$, where $x_j = x_1$. Otherwise, Player 1's payoff is 0: if $x_j < x_1$ for all $j \neq 1$, then by computation Player 1 will not win.
- In the equilibrium (x_1, x_2, \dots, x_n) , if $x_1 > 1$, then he can increase his payoff by choosing $x_1 - 1$, since by making this change he becomes the outright winner rather than tying with at least one other player.

The remaining possibility is that $x_1 = 1$: every player uses the strategy in which he announces the number 1. □

2.40 Example: Consider the following two-person game.

		Player 2	
		L	R
Player 1	U	a_1, b_1	a_2, b_2
	D	a_3, b_3	a_4, b_4

Figure 2.18

We have the following assumptions.

- (i) Neither strategy weakly dominates the other for any player.
- (ii) $a_1 \neq a_3$ and $b_1 \neq b_2$.
- (iii) It is known that this game has at least one Nash equilibrium

Prove that the game has two Nash equilibria.

Proof. Without loss of generality, we assume that $a_1 > a_3$. Then $a_2 < a_4$.

If $b_1 < b_2$, then $b_3 > b_4$, and hence there is no Nash equilibrium, contradiction. Therefore, $b_1 > b_2$ and $b_3 < b_4$. So (U, L) and (D, R) are two Nash equilibria. \square

2.41 Example: There are 100 people in a society, and two different types of mobile phones available: type A and type B .

Each of the 100 people (players in our context) chooses either A or B simultaneously. Let n_A denote the number of people who choose A and n_B denote the number of people who choose B . Note that $n_A + n_B = 100$. For each player i , his payoff is $6n_A$ if he chooses A , or $4n_B$ if he chooses B . For example, if player i chooses A and total number of people who choose B is 50, player i obtains the payoff of 300 (since $n_A = 50$, $6n_A = 300$). In this case, each of those who choose A obtains the payoff of 300 and each of those who choose B obtains the payoff of 200.

Find all the Nash equilibria. If you believe there is no Nash equilibrium, please explain.

Answer. There are two Nash equilibria:

- All players choose A ($n_A = 100$ and $n_B = 0$).
- All players choose B ($n_A = 0$ and $n_B = 100$).

It is clear that the two strategy profiles above are Nash equilibria. The following shows that there is no other Nash equilibrium:

- Any strategy profile with $0 < n_A < 40$ and $n_B = 100 - n_A$ can not be a Nash equilibrium because any player who chooses A can deviate and obtain a better payoff since $4(n_B + 1) > 6n_A$.
- Any strategy profile with $40 < n_A < 100$ and $n_B = 100 - n_A$ can not be a Nash equilibrium because any player who chooses B can deviate and obtain a better payoff since $6(n_A + 1) > 4n_B$.
- Any strategy profile with $n_A = 40$ and $n_B = 60$ can not be a Nash because any player can deviate profitably. If any player who chooses A deviates, he would obtain a better payoff since $4(n_B + 1) > 6n_A$. If any player who chooses B deviates, he would obtain a better payoff since $6(n_A + 1) > 4n_B$.

\square

2.42 Example: Each individual $i = 1, 2, \dots, 100$ must choose a number $r_i \in [0, 1]$. If an individual chooses a number that is the most closed to the value $\theta \sum_{i=1}^{100} r_i$ (where $\theta \in [0, 1]$ is a parameter), then the individual gets payoff 1; otherwise, the individual gets payoff 0. Formulate this problem as a strategic game, and find all Nash equilibria for each $\theta \in [0, 1]$.

2.43 Example (from <http://econweb.ucsd.edu/~jwatson/>): Consider a strategic setting in which two geographically distinct firms (players 1 and 2) compete by setting prices. Suppose that consumers are uniformly distributed across the interval $[0, 1]$, and each will buy either one unit or nothing. Firm 1 is located at 0 and firm 2 is located at 1. Assume that these locations are fixed. In other words, the firms cannot change their locations; they can only select prices. Simultaneously and independently, firm 1 chooses a price p_1 and firm 2 chooses a price p_2 . Assume that the firms produce at zero cost and that, due to a government regulation, they must set prices between 0 and 6.

As in the standard location game, consumers are sensitive to the distance they have to travel in order to purchase. But they are also sensitive to price. Consumers get a benefit of 6 from the good that is purchased, but they also pay a personal cost of c times the distance they have to travel to make the purchase. Assume that c is a positive constant. If the consumer at location $x \in [0, 1]$ purchases from firm 1 then this consumer's utility is $6 - cx - p_1$. If this consumer instead purchases from firm 2 then her utility is $6 - c(1 - x) - p_2$. If this consumer does not purchase the good, her utility is 0.

- (i) Suppose that, for given prices p_1 and p_2 , every consumer purchases the good. That is, ignore the case in which prices are so high that some consumers prefer not to purchase. Find an expression for the “marginal consumer” who is indifferent between purchasing from firm 1 or firm 2. Denote this consumer's location as $x^*(p_1, p_2)$.
- (ii) Continue to assume that all consumers purchase at the prices you are analyzing. Note that firm 1's payoff (profit) is $p_1 x^*(p_1, p_2)$ and firm 2's payoff is $p_2[1 - x^*(p_1, p_2)]$. Calculate the firms' best-response.
- (iii) Find the Nash equilibrium of this game for the case in which $c = 2$.
- (iv) As c converges to 0, what happens to the firms' equilibrium profits?
- (v) Find the Nash equilibrium of this game for the case in which $c = 8$.

2.44 Example (Complete-information all-pay auction): Two players, indexed by $i = 1, 2$, compete for an object of value $v > 0$. Players simultaneously place their non-negative bids x_i . The higher bidder obtains the object, while one loses his bid regardless of win or loss. Ties will be broken randomly. Their payoffs can then be written as

$$\pi_i(x_i, x_j) = \begin{cases} v - x_i, & \text{if } x_i > x_j, \\ \frac{v}{2} - x_i, & \text{if } x_i = x_j, \\ -x_i, & \text{if } x_i < x_j. \end{cases}$$

Show that there does not exist a (pure-strategy) Nash equilibrium.

2.4 Existence of a Nash equilibrium

2.45 To show that a game has a Nash equilibrium it suffices to show that there is a profile a^* such that $a^* \in B(a^*)$. Fixed-point theorems give conditions on B under which there indeed exists a value of a^* for which $a^* \in B(a^*)$.

2.46 Kakutani's fixed-point theorem: Let X be a compact convex subset of \mathbb{R}^n and let $f: X \rightrightarrows X$ be a set-valued function for which

- for all $x \in X$ the set $f(x)$ is non-empty and convex,
- the graph of f is closed, i.e. f is upper-hemicontinuous.

Then there exists $x^* \in X$ such that $x^* \in f(x^*)$.

f has a closed graph if

$$\text{Graph}(f) = \{(x, y) \mid x \in X, y \in f(x)\}$$

is closed, i.e., for all sequences $\{x_n\}$ and $\{y_n\}$ such that $y_n \in f(x_n)$ for all n , $x_n \rightarrow x$, and $y_n \rightarrow y$, we have $y \in f(x)$.

🔗 2.47 Theorem: A strategic game $\langle N, (A_i), (\succsim_i) \rangle$ has a Nash equilibrium if for all $i \in N$,

- A_i is non-empty compact convex subset of \mathbb{R}^n ,
- \succsim_i is continuous and quasi-concave on A_i .

\succsim_i is continuous if the graph of \succsim_i , $\{(a, a') \mid a \succsim_i a'\}$, is a closed set with respect to the product topology.

\succsim_i is quasi-concave on A_i if for any $a' \in A$, the upper level set for a' $\{a_i \in A_i \mid (a'_{-i}, a_i) \succsim_i a'\}$ is convex.




Figure 2.19: A quasi-concave preference.

2.48 *Proof.* (1) For each $i \in N$ and $a \in A$, $B_i(a_{-i})$ is non-empty, since \succsim_i is continuous and A_i is compact.

(2) For each $i \in N$ and $a \in A$, $B_i(a_{-i})$ is convex, since \succsim_i is quasi-concave on A_i .

$$B_i(a_{-i}) = \cap_{a'_i \in A_i} \{a_i \in A_i \mid (a_{-i}, a_i) \succsim_i (a_{-i}, a'_i)\}.$$

(3) The graph of B is closed, since each \succsim_i is continuous.

(4) By Kakutani's fixed-point theorem, B has a fixed point which is a Nash equilibrium of the strategic game. □

2.49 The existence theorem remains valid when \mathbb{R}^n is replaced by “a metric space” or “a locally convex Hausdorff topological vector space”.

2.5 Symmetric games

2.50 Symmetric games [OR Exercise 20.4]

Consider a two-person strategic game that satisfies the conditions of Theorem 2.47. Let $N = \{1, 2\}$ and assume that the game is symmetric: $A_1 = A_2$ and $(a_1, a_2) \succsim_1 (b_1, b_2)$ if and only if $(a_2, a_1) \succsim_2 (b_2, b_1)$ for all $a \in A$ and $b \in A$ (for simplicity, we could assume $u_1(a_1, a_2) = u_2(a_2, a_1)$).

Use Kakutani's fixed-point theorem to prove that there is an action $a_1^* \in A_1$ such that (a_1^*, a_1^*) is a Nash equilibrium of the game. (Such an equilibrium is called a symmetric equilibrium.) Give an example of a finite symmetric game that has only asymmetric equilibria.

Answer. Define the function $F: A_1 \rightarrow A_1$ by $F(a_1) = B_2(a_1)$ (the best response of player 2 to a_1). The function F satisfies the conditions of Theorem 2.46, and hence has a fixed point, say a_1^* . The pair of actions (a_1^*, a_1^*) is a Nash equilibrium of the game since, given the symmetry, if a_1^* is a best response of player 2 to a_1^* then it is also a best response of player 1 to a_1^* .

It is a symmetric finite game that has no symmetric equilibrium:

		Player 2	
		<i>L</i>	<i>R</i>
Player 1	<i>L</i>	3, 3	1, 4
	<i>R</i>	4, 1	0, 0

Figure 2.20

□

Remark: Consider an n -player symmetric game, and suppose we find an asymmetric Nash equilibrium, which means that not all players use the same strategy. If we find one asymmetric Nash equilibrium, then there are another $n - 1$ asymmetric Nash equilibria to be found.

2.51 Example (Fey (2012)): Symmetric games have only asymmetric equilibria.

There are two players, the strategy set is $[0, 1]$, and the payoffs are defined as follows:

$$u(a_i, a_j) = \begin{cases} (0, 0), & \text{if } a_i = a_j, \\ (1, -1), & \text{if } a_j < a_i < 1 \text{ or } 0 < a_i < a_j = 1, \\ (-1, 1), & \text{if } a_i < a_j < 1 \text{ or } 0 < a_j < a_i = 1, \\ (1, 1), & \text{if } a = (1, 0) \text{ or } a = (0, 1). \end{cases}$$

2.52 Example (Fey (2012)): Symmetric games have only asymmetric equilibria.

There are two players, the strategy set is $[0, 1]$, and the payoffs are defined as follows:

$$u_i(a_1, a_2) = \begin{cases} \max\{a_1, a_2\}, & \text{if } a \neq (1, 1), \\ 0, & \text{if } a = (1, 1). \end{cases}$$

This game clearly has a continuum of equilibria. The (pure strategy) equilibrium set is given by $\{(a_1, a_2) \mid a_1 = 1 \text{ or } a_2 = 1\} \setminus (1, 1)$.

2.6 Strictly competitive games (zero-sum games)

2.53 For an arbitrary strategic game, we can say little about the set of Nash equilibria. However, for strictly competitive games, we can say something about the qualitative character of the equilibria.

2.54 Definition: A two-person game $\langle \{1, 2\}, (A_i), (\succsim_i) \rangle$ is strictly competitive if for any $a \in A$ and $b \in A$ we have

$$a \succsim_1 b \text{ if and only if } b \succsim_2 a.$$

2.55 Without loss of generality, we may assume that a strictly competitive game can be represented as a two-person zero-sum game $\langle \{1, 2\}, (A_i), (u_i) \rangle$ in which payoff functions satisfy $u_1 + u_2 = 0$.

2.56 Definition: $a_i^* \in A_i$ is a maximinimizer for player i if

$$\min_{a_j \in A_j} u_i(a_i^*, a_j) \geq \min_{a_j \in A_j} u_i(a_i, a_j) \text{ for all } a_i \in A_i.$$

Player i maximizes if he chooses an action that is best for him on the assumption that whatever he does, player j will choose her action to hurt him as much as possible.

$$2.57 \text{ Lemma: } \max_{a_j \in A_j} \min_{a_i \in A_i} u_j(a_i, a_j) = - \min_{a_j \in A_j} \max_{a_i \in A_i} u_i(a_i, a_j).$$

$$\text{Proof. } \max_{a_j \in A_j} \min_{a_i \in A_i} u_j(a_i, a_j) = - \min_{a_j \in A_j} \max_{a_i \in A_i} -u_j(a_i, a_j) = - \min_{a_j \in A_j} \max_{a_i \in A_i} u_i(a_i, a_j). \quad \square$$

2.58 Proposition: a^* is a Nash equilibrium if and only if for $i = 1, 2$, a_i^* is i 's maximinimizer and

$$\max_{a_i \in A_i} \min_{a_j \in A_j} u_i(a_i, a_j) = u_i(a^*) = \min_{a_j \in A_j} \max_{a_i \in A_i} u_i(a_i, a_j).$$

2.59 Interpretation: a profile is a Nash equilibrium if and only if the action of each player is maximinimizer. This result provides a link between individual decision-making and the reasoning behind the notion of Nash equilibrium.

2.60 Proof. " \Rightarrow ":

- (1) Since a^* is a Nash equilibrium, $u_i(a^*) \geq u_i(a_i, a_j^*)$ for all $a_i \in A_i$. Then $u_i(a^*) \geq \min_{a_j} u_i(a_i, a_j)$ for all $a_i \in A_i$. Hence we have $u_i(a^*) \geq \max_{a_i} \min_{a_j} u_i(a_i, a_j)$.
- (2) Since $u_i + u_j = 0$, $u_i(a^*) = -u_j(a^*) \leq -u_j(a_i^*, a_j) = u_i(a_i^*, a_j)$ for all $a_j \in A_j$. Then $u_i(a^*) = \min_{a_j} u_i(a_i^*, a_j)$. Hence $u_i(a^*) \leq \max_{a_i} \min_{a_j} u_i(a_i, a_j)$.
- (3) Thus $\min_{a_j} u_i(a_i^*, a_j) = u_i(a^*) = \max_{a_i} \min_{a_j} u_i(a_i, a_j)$ and a_i^* is i 's maximinimizer.
- (4) By Lemma,

$$u_i(a^*) = -u_j(a^*) = - \max_{a_j} \min_{a_i} u_j(a_i, a_j) = \min_{a_j} \max_{a_i} u_i(a_i, a_j).$$

" \Leftarrow ": Since a_i^* is i 's maximinimizer, we have $u_i(a^*) = \min_{a_j} u_i(a_i^*, a_j) \leq u_i(a_i^*, a_j)$ for all $a_j \in A_j$. By $u_i + u_j = 0$, we have $u_j(a^*) \geq u_j(a_i^*, a_j)$ for all $a_j \in A_j$. Thus, a^* is a Nash equilibrium. \square

2.61 Proposition: The Nash equilibria of a strictly competitive game are interchangeable: if (x, y) and (x', y') are equilibria then so are (x, y') and (x', y) .

2.62 Proposition: If a two-player strictly competitive symmetric game has a Nash equilibrium, then it has a symmetric Nash equilibrium.

Proof. Suppose a two-player zero-sum symmetric game has a Nash equilibrium (σ_1^*, σ_2^*) . Then by symmetry, the strategy profile (σ_2^*, σ_1^*) is also a Nash equilibrium. As all equilibria in zero-sum games are interchangeable, it follows that (σ_1^*, σ_1^*) is a Nash equilibrium. \square

2.7 Contest

2.63 Reference: [Corchón \(2007\)](#).

2.64 Two players simultaneously commit to their rent-seeking effort outlays x_1 and x_2 to compete for an indivisible rent v . Each of them wins with a probability

$$p_i(x_1, x_2) = \begin{cases} \frac{x_i}{x_1 + x_2}, & \text{if } x_1 + x_2 > 0, \\ \frac{1}{2}, & \text{if } x_1 + x_2 = 0. \end{cases}$$

The function p_i is called the contest success function.

Assume that the effort incurs a constant marginal cost, which is normalized to one.

Player i maximizes

$$\pi_i(x_i, x_j) = \frac{x_i}{x_1 + x_2} v - x_i.$$

It is clear that π_i is a concave function of x_i . Hence, $x_i^* = \frac{v}{4}$.

2.65 Assume that the new winning probability is

$$p_i(x_1, x_2) = \frac{x_i^r}{x_1^r + x_2^r},$$

where $r \leq 2$.

The equilibrium effort is $\frac{rv}{4}$.

2.66 Now consider an n -player Tullock rent-seeking model, where the contest success function is

$$p_i(x_1, \dots, x_n) = \frac{x_i^r}{x_1^r + \dots + x_n^r}.$$

Then a pure strategy Nash equilibrium exists if and only if $r \leq \frac{n}{n-1}$, with the equilibrium effort level

$$x^* = \frac{n-1}{n^2} rv.$$

The parameter r conventionally depicts the discriminatory power of the selection mechanism. A greater r implies that one's win depends more on the amount of his bid, rather than other noisy factors.

2.67 A more general setting:

$$p_i(x_1, \dots, x_n) = \begin{cases} \frac{\phi(x_i)}{\sum_{k=1}^n \phi(x_k)}, & \text{if } \sum_{k=1}^n \phi(x_k) > 0, \\ \frac{1}{n}, & \text{if } \sum_{k=1}^n \phi(x_k) = 0. \end{cases}$$

Here $\phi(x_i)$ measures the impact of x_i in the contest. The ratio $\phi(x_i) / \sum_{k=1}^n \phi(x_k)$ measures the relative impact of i .

Assume that ϕ is strictly increasing and log-concave.

2.68 Asymmetric case: Each of them wins with a probability

$$p_i(x_1, x_2) = \begin{cases} \frac{x_i}{x_1 + x_2}, & \text{if } x_1 + x_2 > 0, \\ \frac{1}{2}, & \text{if } x_1 + x_2 = 0. \end{cases}$$

The two players have different marginal costs, with $c_1 < c_2$.

2.69 Variation: The contest organizer places a weight α on player 1's effort x_1 ; players win the contest, respectively, with probabilities

$$p_1(x_1, x_2) = \frac{\alpha x_1}{\alpha x_1 + x_2}, \quad p_2(x_1, x_2) = \frac{x_2}{\alpha x_1 + x_2}.$$

2.8 Strategic games with discontinuous payoff functions

2.70 Reference: [Reny \(1999\)](#).

2.71 Consider a strategic game $G = \langle N, (A_i), (u_i) \rangle$. Say “player i can secure a payoff \bar{u}_i at $a \in A$ ” if there exists $\bar{a}_i \in A_i$ such that $u_i(\bar{a}_i, a'_{-i}) \geq \bar{u}_i$ for all a'_{-i} close enough to a_{-i} .

2.72 The game G is better-reply secure if whenever (a^*, u^*) is in the closure of the graph of its payoff profile function and a^* is not a Nash equilibrium, some player i can secure a payoff strictly above u_i^* at a^* .

That is, for every payoff profile limit u^* resulting from strategies approaching non-equilibrium a^* , some player i has a strategy yielding a payoff strictly above u_i^* even if the others deviate slightly from a^* .

All games with continuous payoff functions are better-reply secure.

2.73 Theorem (Theorem 3.1 in [Reny \(1999\)](#)): If each A_i is non-empty, compact, convex subset of a metric space, and each u_i is quasi-concave on A_i , then the game G possesses at least one Nash equilibrium if in addition G is better-reply secure.

2.74 Example: Consider a two-person symmetric game: $G = \langle N, (A_i), (u_i) \rangle$, where $N = \{1, 2\}$, $A_1 = A_2 = [0, 1]$, and for all $a_i, a_j \in [0, 1]$, $i, j = 1, 2$, and $i \neq j$,

$$u_i(a_i, a_j) = \begin{cases} 1, & \text{if } a_i \in [\frac{1}{2}, 1] \text{ and } a_j \in [\frac{1}{2}, 1], \\ 1 + a_i, & \text{if } a_i \in [0, \frac{1}{2}) \text{ and } a_j \in (\frac{2}{3}, \frac{5}{6}), \\ a_i, & \text{otherwise.} \end{cases}$$




Figure 2.21: Payoff function $u_i(a_i, a_j)$.

Let $D = [\frac{1}{2}, 1] \times [\frac{1}{2}, 1]$. By definition, it is easy to see that the set of Nash equilibria is $\{x \in D \mid a_1, a_2 \notin (\frac{2}{3}, \frac{5}{6})\}$.

To check the better-reply secure property: let $\epsilon > 0$ be sufficiently small.

If $a^* \notin D$, then some $a_i^* < \frac{1}{2}$. Thus, i can secure payoff $a_i^* + \epsilon > u_i^* = a_i^*$ (if $a_j^* \notin (\frac{2}{3}, \frac{5}{6})$) or $a_i^* + 1 + \epsilon > u_i^* = a_i^* + 1$ (if $a_j^* \in (\frac{2}{3}, \frac{5}{6})$) by choosing a strategy $a_i^* + \epsilon$.

If $a^* \in D$, then some $a_i^* \in (\frac{2}{3}, \frac{5}{6})$ and $a_j^* \geq \frac{1}{2}$. We distinguish two subcases:

- $a_j^* > \frac{1}{2}$. As a_i^* lies in an open interval $(\frac{2}{3}, \frac{5}{6})$, j can secure payoff $1 + a_j > 1$ by choosing a strategy $a_j \in (0, \frac{1}{2})$.

- $a_j^* = \frac{1}{2}$. In this subcase, the limiting vector u^* depends on how a approaches a^* . We must distinguish two subsubcases:
 - $u^* = (1, 1)$, j can secure payoff $1 + a_j > 1$ by choosing a strategy $a_j \in (0, \frac{1}{2})$.
 - The limiting payoff vector is $u^* = (a_i^*, \frac{3}{2})$ even though the actual payoff vector at $a^* \in D$ is $(1, 1)$. Thus i can secure payoff $a_i^* + \epsilon > u_i^* = a_i^*$ by choosing a strategy $a_i^* + \epsilon$, since for any a_j that deviates slightly from $\frac{1}{2}$,

$$u_i(a_i^* + \epsilon, a_j) = \begin{cases} a_i^* + \epsilon, & \text{if } a_j < \frac{1}{2}, \\ 1, & \text{if } a_j \geq \frac{1}{2}. \end{cases}$$

Bayesian games (strategic games with incomplete information)

Contents

3.1	Bayes' rule (Bayes' theorem)	49
3.2	Bayesian games	50
3.3	Examples	54
3.4	Bilateral trading (Double auction)	66
3.5	Providing a public good under incomplete information	71
3.6	Comments on Bayesian games	71

3.1 Bayes' rule (Bayes' theorem)

3.1 Bayes' rule gives the relationship between the probabilities of A and B , $\text{Prob}(A)$ and $\text{Prob}(B)$, and the conditional probabilities of A given B and B given A , $\text{Prob}(A | B)$ and $\text{Prob}(B | A)$. In its most common form, it is:

$$\text{Prob}(A | B) = \frac{\text{Prob}(A)}{\text{Prob}(B)} \text{Prob}(B | A).$$

3.2 Interpretation:

- $\text{Prob}(A)$, the prior, is the initial degree of belief in A .
- $\text{Prob}(A | B)$, the posterior, is the degree of belief having accounted for B .

3.3 Example: A HIV test usually return a positive or a negative result (or sometimes inconclusive). Among the positive results, there are true positives and false positives. Among the negative results, there are true negatives and false negatives.

- True positive: positive test result and have the disease.
- False positive: positive test result and do not have the disease
- True negative: negative test result and do not have the disease.

- False negative: negative test result and have the disease.

The probability of having HIV is usually taken to be the prevalence (population base rate). Currently this prevalence is around 1/1000.

We also know that the probability of obtaining a positive result given that the person does not have HIV is 0.01, and the probability of obtaining a negative result given that the person has HIV is 0.05.

Question: A person obtains a positive result, then what is the probability that he has HIV?

Answer. For the purpose of this discussion, + will indicate a positive test, − will indicate a negative test, HIV will indicate having HIV and −HIV will indicate not having HIV.

$\text{Prob}(+ \mid \text{HIV})$ is the probability of obtaining a positive result, given that the person has HIV. This is known as the sensitivity. It is a measure of how good the test is at identifying individuals with HIV.

$\text{Prob}(- \mid -\text{HIV})$ is the probability of obtaining a negative test result if you do not have HIV. It is known as the specificity. It is a measure of how the test is at identifying people who do not have HIV.

$\text{Prob}(\text{HIV} \mid +)$ is the posteriori probability, that is, how likely is it that a given person has HIV after we have taken into account the base rate and updated it with the available evidence (*i.e.*, result of HIV test).

We have

$$\begin{aligned}\text{Prob}(\text{HIV} \mid +) &= \frac{\text{Prob}(+ \mid \text{HIV}) \times \text{Prob}(\text{HIV})}{\text{Prob}(+ \mid \text{HIV}) \times \text{Prob}(\text{HIV}) + \text{Prob}(+ \mid -\text{HIV}) \times \text{Prob}(-\text{HIV})} \\ &= \frac{0.95 \times 0.001}{0.95 \times 0.001 + 0.01 \times 0.999} = 8.68\%\end{aligned}$$

□

3.2 Bayesian games

3.4 We frequently wish to model situations in which some of the parties are not certain of the characteristics of some of the other parties. The model of a Bayesian game (also called strategic game with incomplete information), which is closely related to that of a strategic game, is designed for this purpose.

3.5 Example [G Section 3.1.A]: Cournot competition under asymmetric information.

Consider the Cournot duopoly model, except:

- Firm 1's cost function is $c_1(q_1) = cq_1$.
- Firm 2's cost function is

$$c_2(q_2) = \begin{cases} c_H q_2, & \text{with probability } \theta, \\ c_L q_2, & \text{with probability } 1 - \theta, \end{cases}$$

where $c_L < c_H$ are low cost and high cost respectively.

The information is asymmetric: firm 1's cost function is known by both; however, firm 2's cost function is only completely known by itself. Firm 1 knows only the marginal cost of firm 2 to be c_H with probability θ and c_L with probability $1 - \theta$.

All of the above is common knowledge. How much shall each firm produce?

Answer. Firm 2 has two payoff functions:

$$\begin{aligned}\pi_2(q_1, q_2; c_L) &= [a - q_1 - q_2 - c_L]q_2, \\ \pi_2(q_1, q_2; c_H) &= [a - q_1 - q_2 - c_H]q_2.\end{aligned}$$

Firm 1 has only one (expected) payoff function

$$\pi_1(q_1, q_2; c) = \mathbf{E}_{q_2}[a - q_1 - q_2 - c]q_1.$$

The two firms simultaneously choose $(q_1^*, q_2^*(c_H), q_2^*(c_L))$, where

- $q_2^*(c_H)$ solves $\max_{q_2} [a - q_1^* - q_2 - c_H]q_2$,
- $q_2^*(c_L)$ solves $\max_{q_2} [a - q_1^* - q_2 - c_L]q_2$,
- Firm 1 should maximize its expected payoff, i.e., q_1^* maximizes

$$\theta[a - q_1 - q_2^*(c_H) - c]q_1 + (1 - \theta)[a - q_1 - q_2^*(c_L) - c]q_1.$$

By first order condition, it is easy to obtain

$$\begin{aligned}q_2^*(c_H) &= \frac{a - q_1^* - c_H}{2}, & q_2^*(c_L) &= \frac{a - q_1^* - c_L}{2}, \\ q_1^* &= \frac{\theta[a - q_2^*(c_H) - c] + (1 - \theta)[a - q_2^*(c_L) - c]}{2}.\end{aligned}$$

By solving them, we have

$$\begin{aligned}q_1^* &= \frac{a - 2c + \theta c_H + (1 - \theta)c_L}{3} \\ q_2^*(c_H) &= \frac{a - 2c_H + c}{3} + \frac{1 - \theta}{6}(c_H - c_L) \\ q_2^*(c_L) &= \frac{a - 2c_L + c}{3} - \frac{\theta}{6}(c_H - c_L)\end{aligned}$$

□



3.6 Definition: A Bayesian game, denoted by $\langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (u_i) \rangle$, consists of

- a finite set N of players
- a set Ω of states
- a set A_i of actions available to player i
- a set T_i of signals (or types)
- a signal function $\tau_i : \Omega \rightarrow T_i$ that specifies the signal $\tau_i(\omega)$ observed by i at state ω
- a probability measure p_i on Ω (the prior belief of i) for which $p_i(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$
- a payoff function $u_i : A \times \Omega \rightarrow \mathbb{R}$ of player i .

A state $\omega \in \Omega$ contains a “complete” description of the payoff function and the beliefs of every player.

If player i receives the signal $t_i \in T_i$, then he deduces that the state is in the set $\tau_i^{-1}(t_i)$; his posterior belief about the state, denoted by $p_i(\omega \mid t_i)$ or $p_i(\omega \mid \tau_i^{-1}(t_i))$, is the probability conditional on $\tau_i^{-1}(t_i)$, i.e.,

$$p_i(\omega \mid t_i) = p_i(\omega \mid \tau_i^{-1}(t_i)) = \begin{cases} \frac{p_i(\omega)}{p_i(\tau_i^{-1}(t_i))}, & \text{if } \omega \in \tau_i^{-1}(t_i), \\ 0, & \text{otherwise.} \end{cases}$$

In a Bayesian game, player i 's strategy s_i is a function from T_i to A_i . A strategy profile can be denoted by

$$s = (s_j)_{j \in N}, \text{ or } a = \left(\left(a_{(j, t_j)} \right)_{t_j \in T_j} \right)_{j \in N} = \left(a_{(j, t_j)} \right)_{j \in N, t_j \in T_j}.$$

3.7 Belief updating: From prior belief to posterior belief:

$$\begin{array}{ccc} \text{True state } \omega & \xrightarrow{\tau_i} \text{ signal/type } \tau_i(\omega) & \xrightarrow{\text{Updating}} \tau_i^{-1}(\tau_i(\omega)) \\ \text{Prior belief } p_i & & \text{Posterior belief } p_i(\omega' \mid \tau_i(\omega)) \end{array}$$

3.8 We can model Cournot competition under asymmetric information as the Bayesian game in which

- $N = \{1, 2\}$.
- $\Omega = T_1 \times T_2 = \{(c, c_H), (c, c_L)\}$.
- $A_1 = A_2 = [0, \infty)$.
- $T_1 = \{c\}, T_2 = \{c_L, c_H\}$.
- $\tau_1(c, \cdot) = c, \tau_2(c, c_H) = c_H, \tau_2(c, c_L) = c_L$
- i 's prior belief on Ω is: (c, c_H) with probability θ , and (c, c_L) with probability $1 - \theta$.
- Profit functions.

3.9 Definition: $s^* = (s_j^*)$ is a Bayesian Nash equilibrium of $\langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (u_i) \rangle$ if for each $i \in N$ and each $\omega \in \Omega$,

$$\tilde{u}_i(s^*; \omega) \geq \tilde{u}_i(s_i, s_{-i}^*; \omega) \text{ for all } s_i: T_i \rightarrow A_i,$$

where

$$\tilde{u}_i(s; \omega) = \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i(s(\tau(\omega'))); \omega').$$

3.10 Alternative definition: A Bayesian Nash equilibrium of a Bayesian game $\langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (u_i) \rangle$ is a Nash equilibrium of its agent strategic game, denoted by $\langle \bar{N}, (\bar{A}_{(i, t_i)}), (\bar{u}_{(i, t_i)}) \rangle$, which is defined as follows:

- $\bar{N} = \{(i, t_i) \mid i \in N, t_i \in T_i\}$.
- $\bar{A}_{(i, t_i)} = A_i$, and $\bar{A} = \times_{j \in N, t_j \in T_j} \bar{A}_{(j, t_j)}$.
- $\bar{u}_{(i, t_i)}: \bar{A} \rightarrow \mathbb{R}$ is defined as follows:

$$\bar{u}_{(i, t_i)}(\bar{a}) = \sum_{\omega \in \Omega} p_i(\omega \mid t_i) \cdot u_i(a_{(j, \tau_j(\omega))}; \omega),$$

where $\bar{a} = \left(a_{(j, t_j)} \right)_{j \in N, t_j \in T_j} \in \bar{A}$.

That is, $\bar{a}^* \in \bar{A}$ is a Bayesian Nash equilibrium if and only if for each $i \in N$ and for each $t_i \in T_i$,

$$\bar{u}_{(i,t_i)}(\bar{a}^*) \geq \bar{u}_{(i,t_i)}\left(a_{(i,t_i)}, a_{-(i,t_i)}^*\right) \text{ for all } a_{(i,t_i)} \in \bar{A}_{(i,t_i)} = A_i.$$

3.11 *Proof of the equivalence.* “ \Leftarrow ”: Suppose that $s^* = (s_j^*)$ is a Bayesian Nash equilibrium. For each $j \in N$ and each $t_j \in T_j$, let $a_{(j,t_j)}^* = s_j^*(t_j) \in A_j$. So $\bar{a}^* = \left(a_{(j,t_j)}^*\right)_{j \in N, t_j \in T_j} \in \bar{A}$ is a Bayesian Nash equilibrium. For any $s_i: T_i \rightarrow A_i$, let $a_{(i,t_i)} = s_i(t_i)$ for each t_i . Therefore

$$\begin{aligned} \tilde{u}_i(s^*; \omega) &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i(s^*(\tau(\omega'))); \omega') \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(s_1^*(\tau_1(\omega')), \dots, s_i^*(\tau_i(\omega')), \dots, s_n^*(\tau_n(\omega'))\right); \omega') \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(a_{(1,\tau_1(\omega'))}^*, \dots, a_{(i,\tau_i(\omega'))}^*, \dots, a_{(n,\tau_n(\omega'))}^*\right); \omega') \\ &= \bar{u}_{(i,\tau_i(\omega))}(\bar{a}^*) \\ &\geq \bar{u}_{(i,\tau_i(\omega))}\left(a_{(i,t_i)}, a_{-(i,t_i)}^*\right) \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(a_{(1,\tau_1(\omega'))}^*, \dots, a_{(i,\tau_i(\omega'))}^*, \dots, a_{(n,\tau_n(\omega'))}^*\right); \omega') \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(s_1^*(\tau_1(\omega')), \dots, s_i(\tau_i(\omega')), \dots, s_n^*(\tau_n(\omega'))\right); \omega') \\ &= \tilde{u}_i(s_i, s_{-i}^*; \omega) \end{aligned}$$

“ \Rightarrow ”: Here we require τ_i is onto for each $i \in N$ in addition. Suppose that $\bar{a}^* = \left(a_{(j,t_j)}^*\right)$ is a Bayesian Nash equilibrium. For each $j \in N$, let $s_j^*: T_j \rightarrow A_j$ be as follows, $s_j^*(t_j) = a_{(j,t_j)}^*$. For any $a_{(i,t_i)} \in \bar{A}_{(i,t_i)}$, let $s_i: T_i \rightarrow A_i$ be as follows, $s_i(t_i) = a_{(i,t_i)}$. For any t_i , since $\tau_i: \Omega \rightarrow T_i$ is onto, there exists ω , such that $\tau_i(\omega) = t_i$.

$$\begin{aligned} \bar{u}_{(i,t_i)}(\bar{a}^*) &= \sum_{\omega' \in \Omega} p_i(\omega' \mid t_i) \cdot u_i\left(a_{(1,\tau_1(\omega'))}^*, \dots, a_{(i,\tau_i(\omega'))}^*, \dots, a_{(n,\tau_n(\omega'))}^*\right); \omega') \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(s_1^*(\tau_1(\omega')), \dots, s_i^*(\tau_i(\omega')), \dots, s_n^*(\tau_n(\omega'))\right); \omega') \\ &= \tilde{u}_i(s^*; \omega) \\ &\geq \tilde{u}_i(s_i, s_{-i}^*; \omega) \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(s_1^*(\tau_1(\omega')), \dots, s_i(\tau_i(\omega')), \dots, s_n^*(\tau_n(\omega'))\right); \omega') \\ &= \sum_{\omega' \in \Omega} p_i(\omega' \mid \tau_i(\omega)) \cdot u_i\left(a_{(1,\tau_1(\omega'))}^*, \dots, a_{(i,\tau_i(\omega'))}^*, \dots, a_{(n,\tau_n(\omega'))}^*\right); \omega') \\ &= \bar{u}_{(i,t_i)}\left(a_{(i,t_i)}, a_{-(i,t_i)}^*\right) \end{aligned}$$

□

3.12 For applications we often use the following simple version of Bayesian game, denoted by $\langle N, (A_i), (T_i), (u_i), (p_i) \rangle$:

- a set N of players.
- a set A_i of actions available to player i .

- a set T_i of signals/types.
- $u_i(a_1, a_2, \dots, a_n; t_1, t_2, \dots, t_n)$ is i 's payoff function.
- each player i has a belief $p_i(\cdot | t_i)$ on T_{-i} conditional on the signal t_i he receives.

3.13 A strategy of i is a function $s_i: T_i \rightarrow A_i$.

s^* is a Bayesian Nash equilibrium of the Bayesian game if and only if for each $i \in N$, and for each $t_i \in T_i$,

$$\tilde{u}_i(s^*; t_i) \geq \tilde{u}_i(a_i, s_{-i}^*; t_i) \text{ for all } a_i \in A_i,$$

where

$$\begin{aligned} \tilde{u}_i(s; t_i) &= \mathbf{E}_{t_{-i}} u_i(s_i(t_i), s_{-i}(t_{-i}); t_i, t_{-i}) \\ &= \sum_{t_{-i} \in T_{-i}} p_i(t_{-i} | t_i) \cdot u_i(s_1(t_1), s_2(t_2), \dots, s_n(t_n); t_i, t_{-i}). \end{aligned}$$

Moreover, if the players have a common prior p on T , then

$$p_i(t_{-i} | t_i) = \frac{p(t_{-i}, t_i)}{p(t_i)} = \frac{p(t_{-i}, t_i)}{\sum_{t'_{-i} \in T'_{-i}} p(t'_{-i}, t_i)}.$$

That is, no player wants to change his strategy, even if the change involves only one action by one type.

3.14 A Bayesian game $\langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (u_i) \rangle$ has a common prior if $p_i = p_j$ for all $i, j \in N$.

Alternatively, a posterior belief system $(p_1(\cdot | \cdot), p_2(\cdot | \cdot), \dots, p_n(\cdot | \cdot))$ is generated by a common prior if there exists a probability measure p such that for any $i \in N$, any $E \subseteq \Omega$ and any $\omega \in \Omega$,

$$p_i(E) = \sum_{\omega \in \Omega} p_i(E | \tau_i(\omega)) \cdot p(\omega).$$

Note that there may be multiple common priors.

3.15 Example [JR Exercise 7.20].

- Suppose that p is a common prior in a game of incomplete information assigning positive probability to every joint type vector. Show that if some type of some player assigns positive probability to some type, t_i , of another player i , then all players, regardless of their types, also assign positive probability to type t_i of player i .
- Provide a three-player game of incomplete information in which the players' beliefs can not be generated by a common prior that assigns positive probability to every joint vector of types.
- Provide a two-player game of incomplete information in which the players' beliefs can not be generated by a common prior that assigns positive probability to every joint vector of types and in which each player, regardless of his type, assigns positive probability to each type of the other player.

3.3 Examples

3.16 Example [G Exercise 3.2]: Cournot competition under asymmetric information.

Consider a Cournot duopoly operating in a market with inverse demand $P(q_1, q_2) = a - q_1 - q_2$, where q_i is the quantity chosen by firm i . Both firms have total costs $c_i(q_i) = cq_i$, but demand is uncertain: it is high ($a = a_H$)

with probability θ and low ($a = a_L$) with probability $1 - \theta$. (Assume $a_H > a_L > c > 0$.) Furthermore, information is asymmetric: firm 1 knows whether the demand is high or low, but firm 2 does not (however, firm 2 knows the probability θ). All of this is common knowledge. The two firms simultaneously choose quantities. Let q_{1H} denote the quantity chosen by firm 1 if it is type H (in other words, if firm 1 knows $a = a_H$), q_{1L} denote the quantity chosen by firm 1 if it is type L (in other words, if firm 1 knows $a = a_L$).

The strategy spaces are $\{q_{1H} \mid 0 \leq q_{1H} \leq a_H\}$, $\{q_{1L} \mid 0 \leq q_{1L} \leq a_L\}$, and $\{q_2 \mid 0 \leq q_2 \leq \theta a_H + (1 - \theta)a_L\}$. Assume $3(a_L - c) > (\theta a_H + (1 - \theta)a_L - c)$ (roughly speaking, a_H and a_L are not too far from each other).

Find all the Bayesian Nash equilibria of this game.

Answer. (i) Firm i 's action space is $\{q \mid q \geq 0\}$.

(ii) Firm 1's type space $T_1 = \{H, L\}$; Firm 2 has only one type.

(iii) Strategy space: $S_1 = \{(q_{1H}, q_{1L}) \mid q_{1H}, q_{1L} \geq 0\}$, and $S_2 = \{q_2 \mid q_2 \geq 0\}$.

(iv) Suppose that $((q_{1H}^*, q_{1L}^*), q_2^*)$ is a Bayesian Nash equilibrium, then by definition we will have:

- If the demand is high, firm 1 will choose q_{1H}^* to maximize its payoff

$$q_{1H}[a_H - c - q_2^* - q_{1H}],$$

which is a concave function, and hence

$$q_{1H}^* = \frac{a_H - c - q_2^*}{2}. \quad (3.1)$$

- If the demand is low, firm 1 will choose q_{1L}^* to maximize its payoff

$$q_{1L}[a_L - c - q_2^* - q_{1L}],$$

which is a concave function, and hence

$$q_{1L}^* = \frac{a_L - c - q_2^*}{2}. \quad (3.2)$$

- Firm 2 does not know the exact type of the demand, so it will choose q_2^* to maximize its expected payoff

$$\theta q_2[a_H - c - q_{1H}^* - q_2] + (1 - \theta)q_2[a_L - c - q_{1L}^* - q_2],$$

and hence

$$q_2^* = \frac{\theta(a_H - q_{1H}^*) + (1 - \theta)(a_L - q_{1L}^*) - c}{2}. \quad (3.3)$$

Combining Equations (3.1), (3.2) and (3.3), we get

$$\begin{aligned} q_{1H}^* &= \frac{a_H - c}{2} - \frac{\theta a_H + (1 - \theta)a_L - c}{6}, \\ q_{1L}^* &= \frac{a_L - c}{2} - \frac{\theta a_H + (1 - \theta)a_L - c}{6}, \\ q_2^* &= \frac{\theta a_H + (1 - \theta)a_L - c}{3}. \end{aligned}$$

□

3.17 Example [G Exercise 3.3]: Consider the following asymmetric-information model of Bertrand duopoly with differentiated products. Demand for firm i is $q_i(p_i, p_j) = a - p_i + b_i \cdot p_j$. Costs are zero for both firms. The sensitivity

of firm i 's demand to firm j 's price is either high or low. That is, b_i is either b_H or b_L , where $b_H > b_L > 0$. For each firm, $b_i = b_H$ with probability θ and $b_i = b_L$ with probability $1 - \theta$, independent of the realization of b_j . Each firm knows its own b_i but not its competitor's. All of this is common knowledge. What are the action spaces, type spaces, beliefs, and utility functions in this game? What are the strategy spaces? Assume that $\theta b_H + (1 - \theta)b_L < 2$. Find the pure-strategy Bayesian Nash equilibria of this game.

Answer. (i) Firm i 's action space: $A_i = [0, \infty)$.

(ii) Firm i 's type space: $T_i = \{H, L\}$.

(iii) Firm i 's beliefs: $\theta H + (1 - \theta)L$.

(iv) Firm i 's strategy space: $S_i = \{(p_{iH}, p_{iL}) \mid p_{iH}, p_{iL} \in A_i\}$.

(v) Firm i 's utility function (for type t): $(a - p_{it} + b_t(\theta p_{jH} + (1 - \theta)p_{jL}))p_{it}$.

(vi) For type $t = H, L$, firm i 's maximization problem is

$$\max_{p_{it}} \pi_{it} = (a - p_{it} + b_t(\theta p_{jH} + (1 - \theta)p_{jL}))p_{it}.$$

By first order condition, $a - 2p_{it} + b_t(\theta p_{jH} + (1 - \theta)p_{jL}) = 0$. That is, for $i = 1, 2$,

$$\begin{aligned} p_{iH} &= \frac{1}{2}a + \frac{1}{2}b_H(\theta p_{jH} + (1 - \theta)p_{jL}), \\ p_{iL} &= \frac{1}{2}a + \frac{1}{2}b_L(\theta p_{jH} + (1 - \theta)p_{jL}). \end{aligned}$$

Let $b = \theta b_H + (1 - \theta)b_L$. Then, we have

$$\begin{aligned} p_{iH} &= \frac{1}{2}a + \frac{1}{4}ab_H + \frac{1}{4}bb_H(\theta p_{iH} + (1 - \theta)p_{iL}), \\ p_{iL} &= \frac{1}{2}a + \frac{1}{4}ab_L + \frac{1}{4}bb_L(\theta p_{iH} + (1 - \theta)p_{iL}). \end{aligned}$$

Therefore, for $i = 1, 2$,

$$\begin{aligned} p_{iH} &= \frac{\frac{1}{2}a(1 + \frac{1}{2}b_H) + \frac{1-\theta}{8}ab(b_H - b_L)}{1 - \frac{1}{4}b^2}, \\ p_{iL} &= \frac{\frac{1}{2}a(1 + \frac{1}{2}b_L) - \frac{\theta}{8}ab(b_H - b_L)}{1 - \frac{1}{4}b^2} \end{aligned}$$

□

3.18 Example [G Exercise 3.4]: Find all the Bayesian Nash equilibria in the following Bayesian game:

- Nature determines whether the payoffs are as in Game 1 or as in Game 2, each game being equally likely.

	L	R
T	1, 1	0, 0
B	0, 0	0, 0

Game 1

	L	R
T	0, 0	0, 0
B	0, 0	2, 2

Game 2

- Player 1 learns whether nature has drawn Game 1 or Game 2, but player 2 does not.
- Player 1 chooses either T or B ; player 2 simultaneously chooses either L or R .

- Payoffs are given by the game drawn by nature.

Answer. • There are two players: player 1 and player 2;

- Type spaces: $T_1 = \{1, 2\}$, and $T_2 = \{\{1, 2\}\}$;
- Beliefs: player 1's belief on player 2's type is 1 on $\{T, B\}$, and player 2's belief on player 1's types is $1/2$ on 1 and $1/2$ on 2;
- Action spaces: $A_1 = \{T, B\}$, and $A_2 = \{L, R\}$;
- Strategy spaces: $S_1 = \{TT, TB, BT, BB\}$, and $S_2 = \{L, R\}$.

Now we will find the best-response correspondence for each player and each associated type: let a_1, a_2 be player 1's actions in Game 1 and Game 2, respectively, b player 2's action.

- If Game 1 is drawn by Nature, then player 1's best-response correspondence is

$$a_1^*(b) = \begin{cases} \{T\}, & \text{if } b = L; \\ \{T, B\}, & \text{if } b = R. \end{cases}$$

- If Game 2 is drawn by Nature, then player 1's best-response correspondence is

$$a_2^*(b) = \begin{cases} \{T, B\}, & \text{if } b = L; \\ \{B\}, & \text{if } b = R. \end{cases}$$

- Since player 2 does not know which game is being drawn, he will choose b to maximize his expected payoff. The following table is player 2's expected payoff table:

	L	R
TT	$1/2$	0
TB	$1/2$	1
BT	0	0
BB	0	1

Thus we get player 2's best-response correspondence:

$$b^*(a_1, a_2) = \begin{cases} \{L\}, & \text{if } a_1 a_2 = TT; \\ \{R\}, & \text{if } a_1 a_2 = TB; \\ \{L, R\}, & \text{if } a_1 a_2 = BT; \\ \{R\}, & \text{if } a_1 a_2 = BB. \end{cases}$$

Therefore, by definition, we will get all the Bayesian Nash equilibria: (TT, L) , (TB, R) and (BB, R) . The reason is as follows:

- If player 2 plays L , then player 1 must play T in Game 1 (and player 1 is indifferent between T and B in Game 2). Note that, if player 1 plays B in Game 2, then player 2 must play R . So, given that player 2 plays L , the only possible Bayesian Nash equilibrium is (TT, L) in this case.
- If player 2 plays R , then player 1 must play B in Game 2 (and player 1 is indifferent between T and B in Game 1). Note that, R is player 2's best response for TB and BB . So, given that player 2 plays R , there are two Bayesian Nash equilibria: (TB, R) and (BB, R) .

□

3.19 Example: The worker has an outside opportunity v known by himself. The firm believes that $v = 6$ and $v = 10$ with probabilities $2/3$ and $1/3$ respectively. A wage $w = 8$ is preset by the union. The firm and the worker simultaneously announce whether to accept or reject the wage. The worker will be employed by the firm if and only if both of them accept the wage. If the firm accepts the wage, its payoff is 3 if the worker is employed and -1 otherwise. If the firm rejects the wage, then its payoff is 0 regardless the worker's action. The worker's payoff is w if he is employed and v otherwise. Find all the Bayesian Nash equilibria.

Answer. Let Game 1 and Game 2 be as follows:

		Firm	
		A	R
Worker	A	8, 3	6, 0
	R	6, -1	6, 0

Game 1, $v = 6$

		Firm	
		A	R
Worker	A	8, 3	10, 0
	R	10, -1	10, 0

Game 2, $v = 10$

- There are two players: firm and worker;
- Type spaces: $T_f = \{1, 2\}$, and $T_w = \{1, 2\}$;
- Believes: work's belief on firm's type is 1 on $\{1, 2\}$, and firm's belief on work's types is $2/3$ on 1 and $1/3$ on 2;
- Action spaces: $A_w = A_f = \{A, R\}$;
- Strategy spaces: $S_f = \{A, R\}$ and $S_w = \{AA, AR, RA, RR\}$.

Now we will find the best-response correspondence for each player and each associated type: let a_1 and a_2 be worker's actions in Game 1 and Game 2, respectively, b firm's action.

- If Game 1 is drawn by Nature, then worker's best-response correspondence is

$$a_1^*(b) = \begin{cases} \{A\}, & \text{if } b = A; \\ \{A, R\}, & \text{if } b = R. \end{cases}$$

- If Game 2 is drawn by Nature, then worker's best-response correspondence is

$$a_2^*(b) = \begin{cases} \{R\}, & \text{if } b = A; \\ \{A, R\}, & \text{if } b = R. \end{cases}$$

- Since firm does not know which game is being drawn, it will choose b to maximize its expected payoff. The following table is firm's expected payoff table:

		Firm	
		A	R
Worker	AA	3	0
	AR	$5/3$	0
	RA	$1/3$	0
	RR	-1	0

Thus we get firm's best-response correspondence is

$$b^*(a_1, a_2) = \begin{cases} \{A\}, & \text{if } a_1 a_2 = AA; \\ \{A\}, & \text{if } a_1 a_2 = AR; \\ \{A\}, & \text{if } a_1 a_2 = RA; \\ \{R\}, & \text{if } a_1 a_2 = RR. \end{cases}$$

Therefore, by definition, we will get all the Bayesian Nash equilibria: (AR, A) and (RR, R) . The reason is as follows:

- If firm chooses A , then worker should choose A and R in Game 1 and Game 2, respectively. Note that, if worker chooses AR , then firm should choose A .

So, given that firm chooses A , the only possible Bayesian Nash equilibrium is (AR, A) .

- If firm chooses R , then worker can choose any strategy in each game. Note that, only when worker chooses RR , R is firm's best response. So, given that firm chooses R , the only possible Bayesian Nash equilibrium is (RR, R) .

□

3.20 Example: Consider the following Bayesian game.

- Nature selects Game 1 with probability $1/3$, Game 2 with probability $1/3$ and Game 3 with probability $1/3$.
- Player I learns whether Nature has selected Game 1 or not; player II learns whether Nature has selected Game 2 or not.
- Players I and II simultaneously choose their actions: player I either T or B , and player II either L or R .
- Payoffs are given by the game selected by Nature.

	L	R
T	0, 0	6, -1
B	-1, 6	4, 4

Game 1

	L	R
T	1, 3	0, 0
B	0, 0	3, 1

Game 2

	L	R
T	2, -2	-2, 2
B	-2, 2	2, -2

Game 3

All of this is common knowledge. Find all the Bayesian Nash equilibria.

Answer. • There are 2 players: player I and player II;

- Type spaces: $T_1 = \{\{1\}, \{2, 3\}\}$, and $T_2 = \{\{1, 3\}, \{2\}\}$;
- Beliefs: player I's belief on player II's types: $2/3$ on $\{1, 3\}$ and $1/3$ on $\{2\}$; player II's belief on player I's types: $1/3$ on $\{1\}$ and $2/3$ on $\{2, 3\}$;
- Action spaces: $A_1 = \{T, B\}$, and $A_2 = \{L, R\}$;
- Strategy spaces: $S_1 = \{TT, TB, BT, BB\}$, and $S_2 = \{LL, LR, RL, RR\}$.

Now we will find the best-response correspondence for each player and each associated type: let a_1 and a_2 be player I's actions in Game 1, and Games 2 and 3, respectively, b_1 and b_2 player II's actions in Games 1 and 3, and Game 2, respectively.

- If Game 1 is drawn, then player I's best-response correspondence is

$$a_1^*(b_1) = \begin{cases} T, & \text{if } b_1 = L; \\ B, & \text{if } b_1 = R. \end{cases}$$

- If Game 1 is not drawn, then by considering the expected payoff, player I's best-response correspondence is

$$a_2^*(b_1b_2) = \begin{cases} T, & \text{if } b_1b_2 = LL; \\ T, & \text{if } b_1b_2 = LR; \\ B, & \text{if } b_1b_2 = RL; \\ B, & \text{if } b_1b_2 = RR. \end{cases}$$

- If Game 2 is drawn, then player II's best-response correspondence is

$$b_2^*(a_2) = \begin{cases} L, & \text{if } a_2 = T; \\ R, & \text{if } a_2 = B. \end{cases}$$

- If Game 2 is not drawn, then by considering the expected payoff, player II's best-response correspondence is

$$b_1^*(a_1a_2) = \begin{cases} R, & \text{if } a_1a_2 = TT; \\ L, & \text{if } a_1a_2 = TB; \\ R, & \text{if } a_1a_2 = BT; \\ L, & \text{if } a_1a_2 = BB. \end{cases}$$

Therefore, by definition, we will get all the Bayesian Nash equilibria: (TT, RL) and (TB, LR) . The reason is as follows:

- If player I chooses TT , then player II should choose RL ; on the other hand, TT is a not best response for RL . So there is no Bayesian Nash equilibrium when player I chooses TT .
- If player I chooses TB , then player II should choose LR ; on the other hand, TB is a not best response for LR . So there is no Bayesian Nash equilibrium when player I chooses TB .
- If player I chooses BT , then player II should choose RL ; on the other hand, BT is not a best response for RL . So there is no Bayesian Nash equilibrium when player I chooses BT .
- If player I chooses BB , then player II should choose LR ; on the other hand, BB is not a best response for LR . So there is no Bayesian Nash equilibrium when player I chooses BB .

□

3.21 Example: Two individuals are involved in a synergistic relationship. If both individuals devote more effort to the relationship, they are both better off. Specifically, an effort level is a non-negative number, and player 1's payoff function is $e_1(1 + e_2 - e_1)$, where e_i is player i 's effort level. For player 2 the cost of effort is either the same as that of player 1, and hence her payoff function is given by $e_2(1 + e_1 - e_2)$, or effort is very costly for her in which case her payoff function is given by $e_2(1 + e_1 - 2e_2)$. Player 2 knows player 1's payoff function and whether the cost of effort is high for herself or not. Player 1, however, is uncertain about player 2's cost of effort. He believes that the cost of effort is low with probability p , and high with probability $1 - p$, where $0 < p < 1$. Find the Bayesian Nash equilibrium of this game as a function of p .

Answer. (i) There are two players;

(ii) Action spaces: $A_1 = A_2 = [0, \infty)$;

(iii) Type spaces: $T_1 = \{\{H, L\}\}$, and $T_2 = \{H, L\}$;

(iv) Strategy spaces: $S_1 = \{e_1 \mid e_1 \geq 0\}$, and $S_2 = \{(e_{2H}, e_{2L}) \mid e_{2H}, e_{2L} \geq 0\}$.

(v) Let $(e_1^*, e_{2H}^*, e_{2L}^*)$ be a Bayesian Nash equilibrium, then we will have:

- Player 1 does not know the exact type of the cost of effort, so he will choose e_1^* to maximize his expected payoff

$$p \times e_1(1 + e_{2L}^* - e_1) + (1 - p) \times e_1(1 + e_{2H}^* - e_1),$$

and hence

$$e_1^* = \frac{1 + pe_{2L}^* + (1 - p)e_{2H}^*}{2}. \quad (3.4)$$

- For player 2, if the cost of effort is high, then player 2 will choose e_{2H}^* to maximize his payoff

$$e_{2H}(1 + e_1^* - 2e_{2H}),$$

and hence

$$e_{2H}^* = \frac{1 + e_1^*}{4}. \quad (3.5)$$

- For player 2, if the cost of effort is low, then player 2 will choose e_{2L}^* to maximize his payoff

$$e_{2L}(1 + e_1^* - e_{2L}),$$

and hence

$$e_{2L}^* = \frac{1 + e_1^*}{2}. \quad (3.6)$$

Solving Equations (3.4), (3.5) and (3.6), we will have

$$e_1^* = \frac{5 + p}{7 - p}, \quad e_{2H}^* = \frac{3}{7 - p}, \quad e_{2L}^* = \frac{6}{7 - p}.$$

□

3.22 Example: There are 2 players who were at the scene where a crime was committed. But neither player knows whether she has been the only witness to the crime, or whether there was another witness as well. Let π be the probability with which each player believes the other player is a witness. Each player, if she is a witness, can call the police or not. The payoff to Player i is $2/3$ if she calls the police, 1 if someone else calls the police, and 0 if nobody calls.

Question 1: Write down each player's types and strategies.

Answer. Since each player knows that he is in the crime scene, each one has only one type: player 1's type is "Player 1 is a witness", and player 2's type is "player 2's type is a witness". There is no possibility that they are not in the crime scene.¹

However, they don't know whether the other person is also in the crime scene or not. Hence, what they are uncertain about is the other player's type.

Each player i has one types: $t_i = \text{"on the scene"}$. For $\pi \in [0, 1]$, each player i has two strategies C (call) and N (not call). □

Question 2: For each value of $\pi \in [0, 1]$, find the Bayesian Nash equilibria.

Answer (1st method). The story can be formulated as the following Bayesian game:

¹ Another acceptable solution is: player i 's type space is {player i is a witness, Player i is not a witness}. While there is no available action when the type is "player i is not a witness".

- $N = \{1, 2\}$.
- $\Omega = \{\omega_1 = (Y, Y), \omega_2 = (Y, N), \omega_3 = (N, Y), \omega_4 = (N, N)\}$, where, for example, $\omega_3 = (N, Y)$ means player 1 is not a witness and player 2 is a witness.
- $A_i = \{C, NC\}$, where C and NC mean “call the police” and “not call” respectively.
- $T_1 = \{Y\}$ and $T_2 = \{Y\}$, where Y means “player 1/2 is a witness”.
- $\tau_1 \equiv Y$ and $\tau_2 \equiv Y$.
- $p(\omega_1) = \pi^2, p(\omega_2) = p(\omega_3) = \pi(1 - \pi), p(\omega_4) = (1 - \pi)^2$.
-

$$u_1(a_1, a_2, \omega) = \begin{cases} \frac{2}{3}, & \text{if } a_1 = C, \\ 1, & \text{if } a_1 = NC, a_2 = C, \omega = \omega_1, \omega_3, \\ 0, & \text{otherwise,} \end{cases}$$

and

$$u_2(a_1, a_2, \omega) = \begin{cases} \frac{2}{3}, & \text{if } a_2 = C, \\ 1, & \text{if } a_1 = C, a_2 = NC, \omega = \omega_1, \omega_2, \\ 0, & \text{otherwise.} \end{cases}$$

Player i 's strategy set is identical with his action set A_i .

Player i 's payoff when he chooses C is always $\frac{2}{3}$.

Player 1's expected payoff when action profile is (NC, C) is π , and payoff when action profile is (NC, NC) is 0.

Player 2's expected payoff when action profile is (C, NC) is π , and payoff when action profile is (NC, NC) is 0.

So we have the following payoff table

		Player j	
		C	NC
Player i	C	$\frac{2}{3}, \frac{2}{3}$	$\frac{2}{3}, \pi$
	NC	$\pi, \frac{2}{3}$	$0, 0$

Thus the Bayesian Nash equilibria are as follows:

- If $\frac{2}{3} > \pi \geq 0$, then there is only one Bayesian Nash equilibrium (C, C) ;
- If $\pi = \frac{2}{3}$, then there are three Bayesian Nash equilibria (C, C) , (C, NC) and (NC, C) ;
- If $1 \geq \pi > \frac{2}{3}$, then there are two Bayesian Nash equilibria (C, NC) and (NC, C) .

□

Answer (2nd method). Each player i thinks that he is playing the following games:

- Game 1: if player i thinks that player j is also on the spot (probability π). Then player i 's payoff table is as follows:

		Player j	
		C	N
Player i	C	$\frac{2}{3}$	$\frac{2}{3}$
	N	1	0

Game 1: player j is on the scene

- Game 2: if player i thinks that player j is not on the spot (probability $1 - \pi$). Then player i think that he will get $2/3$ if he chooses C , and 0 otherwise, no matter what player j chooses.

		Player j	
		C	N
Player i	C	$2/3$	$2/3$
	N	0	0

Game 2: player j is not on the scene

Therefore, player i 's expected payoff is in the payoff table G_1 , and the game in fact can be represented by the payoff table G_2 .

		Player j	
		C	N
Player i	C	$2/3$	$2/3$
	N	π	0

G_1

		Player j	
		C	N
Player i	C	$2/3, 2/3$	$2/3, \pi$
	N	$\pi, 2/3$	$0, 0$

G_2

Thus the Bayesian Nash equilibria are as follows:

- If $2/3 > \pi \geq 0$, then there is only one Bayesian Nash equilibrium (C, C) ;
- If $\pi = 2/3$, then there are three Bayesian Nash equilibria (C, C) , (C, N) and (N, C) ;
- If $1 \geq \pi > 2/3$, then there are two Bayesian Nash equilibria (C, N) and (N, C) .

□

3.23 Example: There are $n \geq 2$ players. Each player i must simultaneously decide whether to join a team ($x_i = 1$) or not ($x_i = 0$); hence $z = \sum_{i=1}^n x_i$ is the size of the team. If player i does not join (so that $x_i = 0$) then i receives a payoff of zero. If player i joins the team (so that $x_i = 1$) then i pays a cost of c_i . If all n players join the team (so that $z = n$) then each player enjoys a benefit of v . Hence player i 's payoff is $u_i = v - c_i$ when $z = n$, and $u_i = -x_i c_i$ when $z < n$. Suppose that $v > c_i > 0$.

Question 1: Suppose that the costs c_1, \dots, c_n are common knowledge. Find all Nash equilibria.

Answer. For player i , given other players' strategies, his best-response correspondence is

$$x_i^*(x_{-i}) = \begin{cases} 0, & \text{if } x_{-i} \neq n-1 \\ 1, & \text{if } x_{-i} = n-1 \end{cases}, \quad \text{where } x_{-i} = \sum_{j \neq i} x_j.$$

It is easy to see that there are two Nash equilibria $(0, 0, \dots, 0)$ and $(1, 1, \dots, 1)$. The reason is as follows:

- If player 1 chooses 0, then each of other player should choose 0. Note that 0 is player 1's best response when each of other players chooses 0;
So, given that player 1 chooses 0, the only possible Nash equilibrium is $(0, 0, \dots, 0)$ in this case.
- If player 1 chooses 1. Note that 1 is player 1's best response only when each of other players chooses 1;
So, given that player 1 chooses 1, the only possible Nash equilibrium is $(1, 1, \dots, 1)$ in this case.

□

Question 2: Now, suppose that information is incomplete. Player i 's cost realization c_i is known only to i ; players' costs are drawn independently from the same uniform distribution: $c_i \sim U[0, \bar{c}]$. Find the symmetric Bayesian Nash equilibrium.

Answer. • There are $n \geq 2$ players;

- Type spaces: $T_i = \{c_i \mid c_i \in [0, \bar{c}]\}$;
- Action spaces: $A_i = \{0, 1\}$;
- Strategy spaces: $S_i = \{\text{functions from } T_i \text{ to } A_i\}$.

Suppose $x(c_i): T_i \rightarrow A_i$ for each player i constitutes a symmetric Bayesian Nash equilibrium. Since we know that when the cost becomes larger, the more possibility player will choose 0. So x can be characterized by $y \in [0, \bar{c}]$, that is,

$$x(c_i) = \begin{cases} 1, & \text{if } c_i \in [0, y]; \\ 0, & \text{otherwise.} \end{cases}$$

For player i , when c_i is drawn, given other players' strategies $x(c_j)$, player i 's expected payoff is

$$\begin{cases} (y/\bar{c})^{n-1}(v - c_i) + [1 - (y/\bar{c})^{n-1}](-c_i), & \text{if } x(c_i) = 1; \\ 0, & \text{if } x(c_i) = 0. \end{cases}$$

Thus player i chooses 1 if and only if $(y/\bar{c})^{n-1}(v - c_i) + [1 - (y/\bar{c})^{n-1}](-c_i) \geq 0$, that is

$$(y/\bar{c})^{n-1}v \geq c_i.$$

- If $(y/\bar{c})^{n-1}v \geq \bar{c}$, then $y = \bar{c}$.
- If $(y/\bar{c})^{n-1}v < \bar{c}$, then $(y/\bar{c})^{n-1}v = y$. We consider the following two cases.
 - If $n = 2$, then $\frac{y}{\bar{c}} \cdot v = y$. Since $v > \bar{c}$, this equation has only one solution $y = 0$.
 - If $n > 2$, then this equation has two solutions: $y = 0$ or $y = (\frac{1}{v})^{\frac{1}{n-2}} \bar{c}^{\frac{n-1}{n-2}}$. Moreover, $y = (\frac{1}{v})^{\frac{1}{n-2}} \bar{c}^{\frac{n-1}{n-2}} < \bar{c}$ if and only if $\bar{c} < v$. Thus, for both solutions the condition $(y/\bar{c})^{n-1}v = y < \bar{c}$ is satisfied.

To summarize,

- if $n = 2$, then $y = 0$ or $y = \bar{c}$;
- if $n > 2$, then $y = 0$, or $y = (\frac{1}{v})^{\frac{1}{n-2}} \bar{c}^{\frac{n-1}{n-2}}$ or $y = \bar{c}$.

Then $(x(c_i))_i$ is a symmetric Bayesian Nash equilibrium for

$$x(c_i) = \begin{cases} 1, & \text{if } c_i \in [0, y], \\ 0, & \text{otherwise,} \end{cases}$$

where y is stated above. □

3.24 Example: Exchange game.

A rich, honest, but mischievous father told his two sons that he had placed 10^n dollars and 10^{n-1} dollars in two envelopes respectively, where $n \in \{1, 2, \dots, 10\}$. The father then randomly handed each son one of the two envelopes with a probability 0.5.

After both sons opened their envelop, his father privately asked each son whether he wanted to switch his envelop with the one his brother had. If both sons agreed, then the envelops were switched. Otherwise, each son kept the original envelop he received.

Represent the sons' problem as a Bayesian game, and find all the Bayesian Nash equilibria.

Answer. We can formulate this game as follows:

- Two players: son 1 and son 2.
- $T_1 = T_2 = \{0, 1, 2, \dots, 10\}$.
- $A_i = \{Y, N\}$, where Y means that son i wants to switch.
- Payoff:

$$u_i(a_1, a_2; t_i, t_{-i}) = \begin{cases} 10^{t_{-i}}, & \text{if } a_1 = a_2 = Y, \\ 10^{t_i}, & \text{otherwise.} \end{cases}$$

- Beliefs:

$$p(t_{-i} | t_i) = \begin{cases} 1 \circ (t_i + 1), & \text{if } t_i = 0, \\ \frac{1}{2} \circ (t_i - 1) + \frac{1}{2} \circ (t_i + 1), & \text{if } t_i = 1, 2, \dots, 9, \\ 1 \circ (t_i - 1), & \text{if } t_i = 10. \end{cases}$$

- $s_i: T_i \rightarrow A_i$.

There are two kinds of Bayesian Nash equilibrium strategies:

- $s_i^*(t_i) \equiv N$.
- $s_i^*(t_i) = \begin{cases} Y, & \text{if } t_i = 0 \\ N, & \text{otherwise} \end{cases}$.

Hence, the switch of envelope would never take place.

The reason is as follows:

- (1) If a son received the envelope of $\$10^{10}$, the son would definitely say “no” (since he knew the other must receive the envelope of $\$10^{10-1}$. Therefore, at an equilibrium, $\$10^{10}$ -type player must say “no”.
- (2) Given that $\$10^{10}$ -type player says “no”, $\$10^{10-1}$ -type player would realize that he is now in the position of $\$10^{10}$ -type player and, thus, should say “no”.
- (3) Repeat the argument. For any integer $n > 0$, $\$10^n$ -type player should say “no” at an equilibrium
- (4) If $n = 0$, $\$10^0$ -type player would be indifferent between saying “no” or saying “yes”.

□

3.25 Example [OR Exercise 28.2]: Exchange game.

Each of two players receives a ticket on which there is a number in some finite subset S of the interval $[0, 1]$. The number on a player's ticket is the size of a prize that he may receive. The two prizes are identically and independently distributed, with distribution function F . Each player is asked independently and simultaneously whether he wants to exchange his prize for the other player's prize. If both players agree then the prizes are exchanged; otherwise each player receives his own prize. Each player's objective is to maximize his expected payoff. Model this situation as a Bayesian game and show that in any Nash equilibrium the highest prize that either player is willing to exchange is the smallest possible prize.

Answer. In the Bayesian game there are two players, say $N = \{1, 2\}$, the set of states is $\Omega = S \times S$, the set of actions of each player is $\{Exchange, Don'texchange\}$, the signal function of each player i is defined by $\tau_i(s_1, s_2) = s_i$, and each player's belief on Ω is that generated by two independent copies of F . Each player's preferences are represented by the payoff function $u_i((X, Y), \omega) = \omega_j$ if $X = Y = Exchange$ and $u_i((X, Y), \omega) = \omega_i$ otherwise.

Let x be the smallest possible prize and let M_i be the highest type of player i that chooses *Exchange*. If $M_i > x$ then it is optimal for type x of player j to choose *Exchange*. Thus if $M_i \geq M_j$ and $M_i > x$ then it is optimal for type M_i of player i to choose *Don'texchange*, since the expected value of the prizes of the types of player j that choose *Exchange* is less than M_i . Thus in any possible Nash equilibrium $M_i = M_j = x$: the only prizes that may be exchanged are the smallest. \square

3.4 Bilateral trading (Double auction)

3.26 Example [G Section 3.2.C]: Double auction.

There are two players: a buyer and a seller.

The buyer's valuation for the seller's good is v_b , the seller's is v_s . The valuations are private information and are drawn from certain independent distribution on $[0, 1]$.

The seller names an asking price p_s , and the buyer simultaneously names an offer price p_b . If $p_b \geq p_s$, then trade occurs at price $p = \frac{p_b + p_s}{2}$; if $p_b < p_s$, then no trade occurs.

Buyer's payoff is

$$\pi_b[p_b, p_s \mid v_b] = \begin{cases} v_b - \frac{p_b + p_s}{2}, & \text{if } p_b \geq p_s, \\ 0, & \text{if } p_b < p_s. \end{cases}$$

Seller's payoff is

$$\pi_s[p_b, p_s \mid v_s] = \begin{cases} \frac{p_b + p_s}{2} - v_s, & \text{if } p_b \geq p_s, \\ 0, & \text{if } p_b < p_s. \end{cases}$$

Question 1: Find all the linear Bayesian Nash equilibria.

Answer. Let $p_i(v_i) = a_i + c_i v_i$, $i = s, b$ be players' linear strategies, where $a_i \geq 0$ and $c_i > 0$.

Given seller's strategy $p_s(v_s)$, buyer's expected payoff is

$$\begin{aligned} E_{v_s} \pi_b[p_b, p_s(v_s) \mid v_b] &= \int_{a_s \leq p_s(v_s) \leq p_b} v_b - \frac{p_b + p_s(v_s)}{2} dv_s + \int_{p_b < p_s(v_s) \leq a_s + c_s} 0 dv_s \\ &= \int_{a_s \leq u \leq p_b} v_b - \frac{p_b + u}{2} d\frac{u}{c_s} = \frac{p_b - a_s}{c_s} \left(v_b - \frac{3}{4}p_b - \frac{1}{4}a_s \right). \end{aligned}$$

Maximizing $E_{v_s} \pi_b[p_b, p_s(v_s) \mid v_b]$ yields buyer's best response

$$p_b(v_b) = \frac{2}{3}v_b + \frac{a_s}{3},$$

which implies $c_b = \frac{2}{3}$ and $a_b = \frac{a_s}{3}$.

Analogously, given buyer's linear strategy $p_b(v_b)$, seller's expected payoff is

$$E_{v_b} \pi_s[p_s, p_b(v_b) \mid v_s] = \frac{a_b + c_b - p_s}{c_b} \left(\frac{3}{4}p_s + \frac{a_b + c_b}{4} - v_s \right).$$

Maximizing $E_{v_b} \pi_s[p_s, p_b(v_b) | v_s]$ yields seller's best response

$$p_s(v_s) = \frac{2}{3}v_s + \frac{a_b + c_b}{3},$$

which implies $c_s = \frac{2}{3}$ and $a_s = \frac{a_b + c_b}{3}$.

Therefore, the linear equilibrium strategies are

$$p_b(v_b) = \frac{2}{3}v_b + \frac{1}{12}, \quad p_s(v_s) = \frac{2}{3}v_s + \frac{1}{4}.$$

The trade occurs if and only if $p_b \geq p_s$, i.e., if and only if

$$v_b \geq v_s + \frac{1}{4}.$$

□

Remark: Myerson and Satterthwaite (*Journal of Economic Theory*, 1983) show that, for the uniform valuation distributions, the linear equilibrium yields higher expected gains for the players than any other Bayesian Nash equilibria of the double auction. This implies that there is no Bayesian Nash equilibrium of the double auction in which trade occurs if and only if it is efficient (i.e., if and only if $v_b \geq v_s$).

Question 2: The double auction above has the linear equilibrium strategies:

$$p_b(v_b) = \frac{2}{3}v_b + \frac{1}{12}, \quad p_s(v_s) = \frac{2}{3}v_s + \frac{1}{4}.$$

Note that $p_b(v_b) > v_b$ if $v_b < \frac{1}{4}$. This means that some types ($v_b < \frac{1}{4}$) of the buyer offer such prices which may probably lead to negative payoffs. Does this equilibrium look reasonable? Can you prove that actually no trade occurs with negative payoffs to any player? (You can find the similar situation for the seller.)

Answer. When they choose the following strategies

$$p_b(v_b) = \frac{2}{3}v_b + \frac{1}{12}, \quad p_s(v_s) = \frac{2}{3}v_s + \frac{1}{4},$$

then payoffs are

$$\pi_b = \begin{cases} \frac{2}{3}v_b - \frac{1}{3}v_s - \frac{1}{6}, & \text{if trade occurs;} \\ 0, & \text{otherwise.} \end{cases} \quad \text{and } \pi_s = \begin{cases} \frac{1}{3}v_b - \frac{2}{3}v_s + \frac{1}{6}, & \text{if trade occurs;} \\ 0, & \text{otherwise.} \end{cases}$$

It suffices to show $\frac{2}{3}v_b - \frac{1}{3}v_s - \frac{1}{6}$ and $\frac{1}{3}v_b - \frac{2}{3}v_s + \frac{1}{6}$ can not be negative, when trade occurs.

If $\frac{2}{3}v_b - \frac{1}{3}v_s - \frac{1}{6} < 0$, since trade occurs when $v_b \geq v_s + \frac{1}{4}$, we have $v_s < 0$, which is a contradiction.

If $\frac{1}{3}v_b - \frac{2}{3}v_s + \frac{1}{6} < 0$, since trade occurs when $v_b \geq v_s + \frac{1}{4}$, we have $v_b > 1$, which is a contradiction. □

3.27 Example: Double auction.

Consider the double auction where the seller's and buyer's valuations, v_s and v_b , are uniformly distributed on $[\alpha_s, \beta_s]$ and $[\alpha_b, \beta_b]$, respectively. Find the linear Bayesian Nash equilibrium of the game.

Answer. • There are two players: seller (s) and buyer (b);

- Type spaces: $T_s = [\alpha_s, \beta_s]$ and $T_b = [\alpha_b, \beta_b]$;
- Action spaces: $A_s = A_b = [0, \infty)$;
- Strategy spaces: $S_b = \{\text{function from } T_b \text{ to } A_b\}$, and $S_s = \{\text{function from } T_s \text{ to } A_s\}$;
- Payoff:

$$u_s(p_s, p_b; v_s, v_b) = \begin{cases} \frac{p_s + p_b}{2} - v_s, & p_b \geq p_s \\ 0, & p_b < p_s \end{cases}, \quad u_b(p_s, p_b; v_s, v_b) = \begin{cases} v_b - \frac{p_s + p_b}{2}, & p_b \geq p_s \\ 0, & p_b < p_s \end{cases}.$$

Suppose (p_s^*, p_b^*) is a linear Bayesian Nash equilibrium, where

$$p_s^*(v_s) = a_s + c_s v_s, \quad p_b^*(v_b) = a_b + c_b v_b.$$

Note that a_s, c_s, a_b, c_b are to be determined. Here we should assume $c_s, c_b > 0$.

- For seller, when v_s is drawn, given buyer's strategy p_b^* , $p_s^*(v_s)$ will maximize his expected payoff

$$\begin{aligned} & \mathbb{E}[u_s(p_s, p_b^*; v_s, v_b)] \\ &= \frac{1}{\beta_b - \alpha_b} \int_{p_s \leq p_b^*(v_b) \leq p_b^*(\beta_b)} \frac{p_s + p_b^*(v_b)}{2} - v_s \, dv_b + \frac{1}{\beta_b - \alpha_b} \int_{p_b^*(\alpha_b) \leq p_b^*(v_b) < p_s} 0 \, dv_b \\ &= \frac{1}{\beta_b - \alpha_b} \int_{\frac{p_s - a_b}{c_b}}^{\beta_b} \frac{p_s + a_b + c_b v_b}{2} - v_s \, dv_b \\ &= \frac{1}{\beta_b - \alpha_b} \left[\left(\frac{p_s + a_b}{2} - v_s \right) \left(\beta_b - \frac{p_s - a_b}{c_b} \right) + \frac{c_b}{2} \int_{\frac{p_s - a_b}{c_b}}^{\beta_b} v_b \, dv_b \right] \\ &= \frac{1}{\beta_b - \alpha_b} \left[\left(\frac{p_s + a_b}{2} - v_s \right) \left(\beta_b - \frac{p_s - a_b}{c_b} \right) + \frac{c_b}{4} \left(\beta_b - \frac{p_s - a_b}{c_b} \right) \left(\beta_b + \frac{p_s - a_b}{c_b} \right) \right] \\ &= \frac{1}{\beta_b - \alpha_b} \left(\beta_b - \frac{p_s - a_b}{c_b} \right) \left[\left(\frac{p_s + a_b}{2} - v_s \right) + \frac{c_b}{4} \left(\beta_b + \frac{p_s - a_b}{c_b} \right) \right] \\ &= \frac{c_b}{\beta_b - \alpha_b} (c_b \beta_b - p_s + a_b) \left[-v_s + \frac{3}{4} p_s + \frac{1}{4} (a_b + c_b \beta_b) \right] \end{aligned}$$

Therefore, by the first order condition,

$$p_s^*(v_s) = \frac{2}{3} v_s + \frac{1}{3} a_b + \frac{1}{3} c_b \beta_b,$$

and hence

$$c_s = \frac{2}{3}, \quad a_s = \frac{1}{3} (a_b + c_b \beta_b). \quad (3.7)$$

- For buyer, when v_b is drawn, given seller's strategy p_s^* , $p_b^*(v_b)$ will maximize his expected payoff

$$\begin{aligned} & \mathbb{E}[u_b(p_s^*, p_b; v_s, v_b)] \\ &= \frac{1}{\beta_s - \alpha_s} \int_{p_s^*(\alpha_s) \leq p_s^*(v_s) \leq p_b} v_b - \frac{p_s^*(v_s) + p_b}{2} \, dv_s + \frac{1}{\beta_s - \alpha_s} \int_{p_b < p_s^*(v_s) \leq p_s^*(\beta_s)} 0 \, dv_s \\ &= \frac{1}{\beta_s - \alpha_s} \int_{\alpha_s}^{\frac{p_b - a_s}{c_s}} v_b - \frac{a_s + c_s v_s + p_b}{2} \, dv_s \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\beta_s - \alpha_s} \left[\left(v_b - \frac{a_s + p_b}{2} \right) \left(\frac{p_b - a_s}{c_s} - \alpha_s \right) - \frac{c_s}{2} \int_{\alpha_s}^{\frac{p_b - a_s}{c_s}} v_s \, dv_s \right] \\
&= \frac{1}{\beta_s - \alpha_s} \left[\left(v_b - \frac{a_s + p_b}{2} \right) \left(\frac{p_b - a_s}{c_s} - \alpha_s \right) - \frac{c_s}{4} \left(\frac{p_b - a_s}{c_s} - \alpha_s \right) \left(\frac{p_b - a_s}{c_s} + \alpha_s \right) \right] \\
&= \frac{1}{\beta_s - \alpha_s} \left(\frac{p_b - a_s}{c_s} - \alpha_s \right) \left[\left(v_b - \frac{a_s + p_b}{2} \right) - \frac{c_s}{4} \left(\frac{p_b - a_s}{c_s} + \alpha_s \right) \right] \\
&= \frac{c_s}{\beta_s - \alpha_s} (p_b - a_s - c_s \alpha_s) \left[v_b - \frac{3}{4} p_b - \frac{1}{4} (a_s + c_s \alpha_s) \right]
\end{aligned}$$

Therefore, by the first order condition,

$$p_b^*(v_b) = \frac{2}{3} v_b + \frac{1}{3} a_s + \frac{1}{3} c_s \alpha_s,$$

and hence

$$c_b = \frac{2}{3}, \quad a_b = \frac{1}{3} (a_s + c_s \alpha_s). \quad (3.8)$$

Solving Equations (3.7) and (3.8), we will have

$$a_s = \frac{\alpha_s}{12} + \frac{\beta_b}{4}, \quad a_b = \frac{\beta_b}{12} + \frac{\alpha_s}{4}.$$

□

3.28 Example [G Exercise 3.8]: Double auction.

A firm and a worker play a double auction. The firm knows the worker's marginal product (m) and the worker knows his or her outside opportunity (v), respectively. In this context, trade means that the worker is employed by the firm. A wage w is preset by the union. If there is trade, then the firm's payoff is $m - w$ and the worker's is w ; if there is no trade then the firm's payoff is zero and the worker's is v . Suppose that m and v are independent draws from a uniform distribution on $[0, 1]$. The both players simultaneously announce either that they Accept the wage w or that they Reject that wage. The worker will be employed by the firm if and only if both of them accept the wage. Given an arbitrary value of w from $[0, 1]$, what is the Bayesian Nash equilibrium of this game? Draw a diagram showing the type-pairs that trade. Find the value of w that maximizes the sum of the players' expected payoff and compute this maximized sum.

Answer. • There are two players: firmer and worker;

- Type spaces: $T_f = \{m \mid m \in [0, 1]\}$, and $T_w = \{v \mid v \in [0, 1]\}$;
- Action spaces: $A_f = A_w = \{A, R\}$;
- Strategy spaces: $S_f = S_w = \{\text{functions from } [0, 1] \text{ to } \{A, R\}\}$;
- Payoff functions:

$$\begin{aligned}
u_f(s_f(w), s_w(v); m, v) &= \begin{cases} m - w, & \text{if } s_f(w) = s_w(v) = A; \\ 0, & \text{otherwise.} \end{cases} \\
u_w(s_f(w), s_w(v); m, v) &= \begin{cases} w, & \text{if } s_f(w) = s_w(v) = A; \\ v, & \text{otherwise.} \end{cases}
\end{aligned}$$

(i) For any $w \in [0, 1]$, it is easy to see $(s_f^*(m), s_w^*(v))$ is a Bayesian Nash equilibrium, where

$$s_f^*(m) = \begin{cases} A, & \text{if } m \geq w \\ R, & \text{otherwise} \end{cases}, \quad s_w^*(v) = \begin{cases} A, & \text{if } w \geq v \\ R, & \text{otherwise} \end{cases}.$$

(ii) There is trade when (m, v) is drawn if and only if $s_f^*(m) = s_w^*(v) = A$, and thus T is the trading area in Figure 3.1.




Figure 3.1: Trading area T

(iii) In the Bayesian Nash equilibrium, the payoff are as follows:

$$u_f(m, v) = \begin{cases} m - w, & \text{if } (m, v) \in T \\ 0, & \text{otherwise} \end{cases}, \quad u_w(m, v) = \begin{cases} w, & \text{if } (m, v) \in T \\ v, & \text{otherwise} \end{cases}.$$

Since m and v are uniformly distributed on $[0, 1]$, we have:

$$\begin{aligned} \mathbf{E}[u_f] &= \int_0^1 \int_0^1 u_f(m, v) \, dv \, dm = \iint_T (m - w) \, dv \, dm \\ \mathbf{E}[u_w] &= \int_0^1 \int_0^1 u_w(m, v) \, dv \, dm = \iint_T w \, dv \, dm + \iint_{T^c} v \, dv \, dm \\ &= \iint_T (w - v) \, dv \, dm + \int_0^1 \int_0^1 v \, dv \, dm \end{aligned}$$

and thus

$$\begin{aligned} \mathbf{E}[u_f] + \mathbf{E}[u_w] &= \iint_T (m - v) \, dv \, dm + \int_0^1 \int_0^1 v \, dv \, dm \\ &= \int_w^1 \int_0^w (m - v) \, dv \, dm + \int_0^1 \int_0^1 v \, dv \, dm = \frac{w - w^2}{2} + \frac{1}{2} \end{aligned}$$

Therefore, $w^* = \frac{1}{2}$ is the maximizer of the sum of the expected payoff.

□

3.5 Providing a public good under incomplete information

3.29 There are two players, $i = 1, 2$. Players decide simultaneously whether to contribute to the public good, and contributing is a 0-1 decision. Each player derives a benefit of 1 if at least one of them provides the public good and 0 if none does; player i 's cost of contributing is c_i .

3.30 The benefits of the public good—1 each—are common knowledge, but each player's cost is known only to that player. However, both players believe it is common knowledge that the c_i are drawn independently from the same continuous and strictly increasing cumulative distribution function p on $[\underline{c}, \bar{c}]$, where $\underline{c} < 1 < \bar{c}$. The cost c_i is player i 's "type."

3.31 A pure strategy is a function $s_i: [\underline{c}, \bar{c}] \rightarrow \{0, 1\}$, where 1 means "contribute" and 0 means "not contribute."

3.32 Player i 's payoff is

$$u_i(a_i, a_j, c_i) = \max\{a_i, a_j\} - c_i a_i.$$

3.33 A Bayesian Nash equilibrium is (s_1^*, s_2^*) such that for each player i and every c_i , the action $s_i^*(c_i)$ solves

$$\max_{a_i} \mathbf{E}_{c_j} u_i(a_i, s_j^*(c_j), c_i).$$

3.34 Let $z_j \triangleq \text{Prob}(s_j^*(c_j) = 1)$ be the equilibrium probability that player j contributes.

To maximize his expected payoff, player i will contribute if his cost c_i is less than $1(1 - z_j)$, which is his benefit from the public good times the probability that player j does not contribute. Thus

$$s_i^*(c_i) = \begin{cases} 1, & \text{if } c_i < 1 - z_j, \\ 0, & \text{if } c_i > 1 - z_j. \end{cases}$$

This shows that the types of player i who contribute lie in an interval $[\underline{c}, c_i^*]$, where $c_i^* = 1 - z_j$.

Similarly, player j contributes if and only if $c_j \in [\underline{c}, c_j^*]$ where $c_j^* = 1 - z_i$.

Clearly $c_i^* = 1 - z_j = 1 - \text{Prob}(\underline{c} \leq c_j \leq c_j^*) = 1 - p(c_j^*)$. Thus c_1^* and c_2^* must both satisfy the equation

$$c^* = 1 - p(1 - p(c^*)).$$

3.6 Comments on Bayesian games

3.35 [Harsanyi \(1967–68\)](#) argued that a situation in which the players are unsure about each other's characteristics can be modeled as a Bayesian game. Accordingly, games of incomplete information are transformed into ones with imperfect information. Harsanyi also assumed that the prior belief of every player is the same (this assumption is referred to as Harsanyi's doctrine).

By "complete information", we mean that the payoff functions are common knowledge. (applicable for strategic games and extensive games)

By "perfect information", we mean that at each move in the game, the player with the move knows the full history of the play of the game thus far. (applicable for extensive games)

3.36 A Bayesian game can be used to model not only situations in which each player is uncertain about the others' payoffs, but also situations in which each player is uncertain about the others' knowledge.

Consider a Bayesian game in which

- $N = \{1, 2\}$.
- $\Omega = \{\omega_1, \omega_2, \omega_3\}$.
- the signal functions of two players $i = 1, 2$ are given by $\tau_1(\omega_1) = \tau_1(\omega_2) = t_1$, $\tau_1(\omega_3) = t'_1$, and $\tau_2(\omega_1) = t_2$, $\tau_2(\omega_2) = \tau_2(\omega_3) = t'_2$.
- player 1's preference satisfy $(b, \omega_j) \succsim_1 (c, \omega_j)$ for $j = 1, 2$ and $(c, \omega_3) \succsim_1 (b, \omega_3)$ for some action profiles b and c .

Suppose that the true state is ω_1 . Player 2 knows that the true state is ω_1 , so he knows player 1 prefers b to c in such a game. Since in state ω_1 , player 1 does not know whether the state is ω_1 or ω_2 , and hence he does not know whether or not player 2 knows that 1 prefers b to c .

3.37 Can every situation in which the players are uncertain about each other's knowledge be modeled as a Bayesian game?

3.38 Assume that the players' payoffs depend only on a parameter $\theta \in \Theta$. Denote the set of possible beliefs of each player i by T_i . Then a belief of any player j is a probability distribution over $\Theta \times T_{-j}$. The question above is to find a collection $\{T_j\}_{j \in N}$ of sets such that for all $i \in N$,

$$T_i \sim^{\text{homeomorphism}} \Delta(\Theta \times T_{-i}),$$

where $\Delta(\Theta \times T_{-i})$ is the set of probability distributions over $\Theta \times T_{-i}$.

A function $f: X \rightarrow Y$ between two topological spaces (X, T_X) and (Y, T_Y) is called a homeomorphism (同胚) if it has the following properties:

- f is a bijection;
- f is continuous;
- f^{-1} is continuous.

3.39 If so, we can let

$$\Omega = \Theta \times (T_1 \times T_2 \times \cdots \times T_n)$$

be the state space and use the model of a Bayesian game to capture any situation in which players are uncertain not only about each other's payoffs but also about each other's beliefs.

In addition, we call $t_i \in T_i$ is a Harsanyi's type.

3.40 Consider a two-player game, where a space Θ of states of nature is primitive uncertainty faced by each player.

$$\begin{aligned} S^{[0]} &= \Theta \\ S^{[1]} &= S^{[0]} \times \Delta(S^{[0]}) \\ S^{[2]} &= S^{[1]} \times \Delta(S^{[1]}) = S^{[0]} \times \Delta(S^{[0]}) \times \Delta(S^{[1]}) \\ S^{[3]} &= S^{[2]} \times \Delta(S^{[2]}) = S^{[0]} \times \Delta(S^{[0]}) \times \Delta(S^{[1]}) \times \Delta(S^{[2]}) \\ &\dots\dots\dots \\ S^{[\ell]} &= S^{[\ell-1]} \times \Delta(S^{[\ell-1]}) = S^{[0]} \times \Delta(S^{[0]}) \times \Delta(S^{[1]}) \times \cdots \times \Delta(S^{[\ell-1]}) \end{aligned}$$

.....

Each player has 1st-beliefs, namely a distribution, about the uncertainty.

As the decisions of other players are relevant, so are their 1st-beliefs, since they affect their decisions. Thus a player must have 2nd-order beliefs about the 1st-beliefs of other players.

For the same reason, a player needs to consider 3rd-order beliefs about the 2nd-beliefs of other players about the 1st-beliefs and so on.

Let $T = \times_{\ell=0}^{\infty} \Delta(S^{[\ell]})$. We have

$$T \sim_{\text{homeomorphism}} \Delta(\Theta \times T),$$

given any one of the following conditions:

- Θ is a compact Hausdorff space: [Mertens and Zamir \(1985\)](#);
- Θ is a Polish space: [Brandenburger and Dekel \(1993\)](#).

These results are valid when there are finite players as well.

Auction

Contents

4.1 Preliminary	75
4.2 The symmetric model	77
4.3 Second-price sealed-bid auction	77
4.4 First-price sealed-bid auction	78
4.5 Revenue comparison	86
4.6 Reserve prices	88
4.7 The revenue equivalence principle	91
4.8 All-pay auction	92
4.9 Third-price auction	93
4.10 Uncertain number of bidders	95

4.1 Preliminary

4.1 An auction is a process of buying and selling goods or services by offering them up for bid, taking bids, and then selling the item to the highest bidder.

In economic theory, an auction may refer to any mechanism or set of trading rules for exchange. A common aspect of auction-like institutions is that they elicit information, in the form of bids, from potential buyers regarding their willingness to pay, and the outcome—that is, who wins what and pays how much—is determined solely on the basis of the received information.

4.2 The uncertainty regarding values facing both sellers and buyers is an inherent feature of auctions.

- The seller is unsure about the values that bidders attach to the object being sold—the maximum amount each bidder is willing to pay.
- Private value: each bidder knows the value of the object to herself at the time of bidding. Implicit in this situation is that no bidder knows with certainty the values attached by other bidder and knowledge of other bidders' value would not affect how much the object is worth to a particular bidder.
- Interdependent value: values are unknown at the time of the auction and may be affected by information available to other bidders.

- Common value: the value is unknown at the time of the auction but is the same for all bidders.

4.3 Auctions should be defined by three kinds of rules:

- rules for bidding
 - who can bid, when
 - what is the form of a bid
 - restrictions on offers, as a function of:
 - * bidder's own previous bid
 - * auction state (others' bids)
 - * eligibility (*i.e.*, budget constraints)
 - * expiration, withdrawal, replacement
- rules for what information is revealed
 - when to reveal what information to whom
- rules for clearing
 - when to clear
 - * at intervals
 - * on each bid
 - * after a period of inactivity
 - allocation (who gets what)
 - payment (who gets what)

4.4 The open ascending price or English auction is the oldest and perhaps most prevalent auction form. The word auction itself is derived from the Latin *augere*, which means “to increase” (or “augment”), via the participle *auctus* (“increasing”).

In one variant of the English auction, so-called Japanese auction, the sale is conducted by an auctioneer who begins by calling out a low price and raises it, typically in small increments, as long as there are at least two interested bidders. The auction stops when there is only one interested bidder.

One way to formally model the underlying game is to postulate that the price rises continuously and each bidder indicates an interest in purchasing at the current price in a manner apparent to all by, say, raising a hand. Once a bidder finds the price to be too high, she signals that she is no longer interested by lowering her hand. The auction ends when only a single bidder is still interested. This bidder wins the object and pays the auctioneer an amount equal to the price at which the second-last bidder dropped out.

4.5 The Dutch auction is the open descending price counterpart of the English auction. It is not commonly used in practice but is of some conceptual interest. Here, the auctioneer begins by calling out a price high enough so that presumably no bidder is interested in buying the object at that price. This price is gradually lowered until some bidder indicates her interest. The object is then sold to this bidder at the given price.

4.6 The sealed-bid first-price auction: Bidders submit bids in sealed envelopes; the person submitting the highest bid wins the object and pays what she bid.

4.7 The sealed-bid second-price auction. As its name suggests, once again bidders submit bids in sealed envelopes; the person submitting the highest bid wins the object but pays not what she bid but the second-highest bid.

- 4.8 The Dutch open descending price auction is strategically equivalent to the first-price sealed-bid auction. When values are private, the English open ascending auction is equivalent to the second-price sealed-bid auction in a weaker sense. (Exercise)

4.2 The symmetric model

- 4.9 There is a single object for sale, and N potential buyers are bidding for the object.
- 4.10 Bidder i assigns a value of X_i to the object—the maximum amount a bidder is willing to pay for the object. Each X_i is independently and identically distributed on some interval $[0, \omega]$ according to the increasing cumulative distribution function F . It is assumed that F admits a continuous density $f \equiv F'$, $F(x) = \int_0^x f(t) dt$, and has full support. It is assumed that $E[X_i] = \int_0^\omega x dF(x) = \int_0^\omega xf(x) dx < \infty$.
- 4.11 Bidder i knows the realization x_i of X_i and only that other bidders' values are independently distributed according to F .
- 4.12 Bidders are risk neutral; they seek to maximize their expected profits.
- 4.13 All components of the model other than the realized values are assumed to be commonly known to all bidders. In particular, the distribution F is common knowledge, as is the number of bidders.
- 4.14 It is also assumed that bidders are not subject to any liquidity or budget constraints. Each bidder i has sufficient resources so if necessary, she can pay the seller up to her value x_i .
- 4.15 A strategy for a bidder is a function $\beta_i: [0, \omega] \rightarrow \mathbb{R}_+$, which determines her bid for any value.
- We will typically be interested in comparing the outcomes of a symmetric equilibrium—an equilibrium in which all bidders follow the same strategy—of one auction with a symmetric equilibrium of the other.

4.3 Second-price sealed-bid auction

- 4.16 In a second-price auction, each bidder submits a sealed bid of b_i , and given these bids, the payoffs are:

$$\Pi_i(b_i, b_{-i}, x_i) = \begin{cases} x_i - \max_{j \neq i} b_j, & \text{if } b_i > \max_{j \neq i} b_j, \\ 0, & \text{if } b_i < \max_{j \neq i} b_j. \end{cases}$$

We also assume that if there is a tie, so $b_i = \max_{j \neq i} b_j$, the object goes to each winning bidder with equal probability.

- 4.17 Proposition: In a second-price sealed-bid auction, it is a weakly dominant strategy to bid according to $\beta^{\Pi}(x) = x$.

Proof. (1) Consider bidder 1, say, and suppose that $p_1 = \max_{j \neq 1} b_j$ is the highest competing bid.

- (2) By bidding x_1 , bidder 1 will win if $x_1 > p_1$ and not if $x_1 < p_1$ (if $x_1 = p_1$, bidder 1 is indifferent between winning and losing).
- (3) Suppose, however, that she bids an amount $z_1 < x_1$.
- If $p_1 > x_1 > z_1$, she still loses.
 - If $x_1 > z_1 \geq p_1$, then she still wins, and her profit is still $x_1 - p_1$.
 - If $x_1 > p_1 > z_1$, then she loses, whereas if she had bid x_1 , she would have made a positive profit.

Thus, bidding less than x_1 can never increase her profit but in some circumstances may actually decrease it.

(4) A similar argument shows that it is not profitable to bid more than x_1 .

□

4.18 Remark: It should be noted that the argument in Proposition 4.17 relied neither on the assumption that bidders' values were independently distributed nor the assumption that they were identically so. Only the assumption of private values is important, and Proposition 4.17 holds as long as this is the case.

4.19 Variation: The bidders $i = 1, 2, \dots, N$ simultaneously submit bids in \mathbb{R}_+ , and the object is given to the bidder with the lowest index among those who submit the highest bid, in exchange for a payment. Each player i knows her own valuation $x_i \in [0, \omega]$ but is uncertain of the other bidders' valuations. Assume that each bidder believes that every other bidder's valuation is drawn independently from the same distribution F over $[0, \omega]$.

The set of actions of each player i is $[0, \infty)$ (the set of possible bids) and the payoff of player i is

$$\begin{cases} x_i - \max_{j \neq i} b_j, & \text{if } b_i > b_1, \dots, b_{i-1}, \text{ and } b_i \geq b_{i+1}, \dots, b_n, \\ 0, & \text{otherwise.} \end{cases}$$

Then for any player i the bid $b_i = x_i$ is a dominant action.

Proof. To see this, let p_i be another action of player i .

- If $\max_{j \neq i} b_j \geq x_i$, then by bidding p_i player i either does not obtain the object or receives a non-positive payoff, while by bidding x_i she guarantees herself a payoff of 0.
- If $\max_{j \neq i} b_j < x_i$, then by bidding x_i player i obtains the good at the price $\max_{j \neq i} b_j$, while by bidding p_i either she wins and pays the same price or loses.

□

4.4 First-price sealed-bid auction

4.20 In a first-price auction, each bidder submits a sealed bid of b_i , and given these bids, the payoffs are

$$\Pi_i(b_i, b_{-i}, x_i) = \begin{cases} x_i - b_i, & \text{if } b_i > \max_{j \neq i} b_j \\ 0, & \text{if } b_i < \max_{j \neq i} b_j \end{cases}$$

As before, if there is more than one bidder with the highest bid, the object goes to each such bidder with equal probability.

4.21 In a first-price auction, equilibrium behavior is more complicated than in a second-price auction.

- No bidder would bid an amount equal to her value, since this would only guarantee a payoff of 0.
- Fixing the bidding behavior of others, at any bid that will neither win for sure nor lose for sure, the bidder faces a simple trade-off: an increase in the bid will increase the probability of winning while, at the same time reducing the gains from winning.

4.22 Assumption: It is expected that a bidder bids higher if she has higher private value, that is, the bidding strategy is assumed to be an increasing function.

4.23 Example [G Exercise 3.6]: N -player first-price sealed-bid auction (risk-neutral, uniform distribution, symmetric Bayesian Nash equilibrium).

Consider an N -player first-price sealed-bid auction in which the bidders' valuations are independently and uniformly distributed on $[0, 1]$. Find a symmetric Bayesian Nash equilibrium.

Part 1: Formulation.

- There are N players;
- Type spaces: $T_i = [0, 1]$, that is, each $x_i \in T_i$ is a valuation;
- Action spaces: $A_i = [0, 1]$, that is, each $b_i \in A_i$ is a bid;
- Strategy spaces: $S_i = \{\beta_i^I: T_i \rightarrow A_i\}$;
- Payoff:

$$\Pi_i(b_i, b_{-i}, x_i) = \begin{cases} x_i - b_i, & \text{if } b_i > b_j, \forall j \neq i; \\ \frac{x_i - b_i}{k}, & \text{if } b_i \text{ is one of the } k \text{ largest bids;} \\ 0, & \text{otherwise.} \end{cases}$$

- Aim: find a symmetric Bayesian Nash equilibrium $(\beta_1^I, \beta_2^I, \dots, \beta_n^I)$, where $\beta_1^I = \beta_2^I = \dots = \beta_n^I = \beta$.

□

Part 2: Heuristic derivation of symmetric equilibrium strategy.

- (1) Suppose that bidder $j \neq 1$ follow the symmetric, increasing, and differentiable equilibrium strategy β . Suppose bidder 1 receives a signal, $X_1 = x$, and bids b . We wish to determine the optimal b .
- (2) Notice that it can never be optimal to choose a bid $b > \beta(1)$, since in that case, bidder 1 would win for sure and could do better by reducing her bid slightly, so she still wins for sure but pays less.
- (3) A bidder with value 0 would never submit a positive bid, since she would make a loss if she were to win the auction. Thus, we must have $\beta(0) = 0$.
- (4) Bidder 1 wins the auction whenever she submits the highest bid—that is, whenever $\max_{i \neq 1} \beta(X_i) < b$. Her expected payoff is therefore

$$(\beta^{-1}(b))^{N-1} \times (x - b).$$

- (5) Maximizing this with respect to b yields the first-order condition:

$$(N-1)(\beta^{-1}(b))^{N-2} \frac{1}{\beta'(\beta^{-1}(b))} (x - b) - (\beta^{-1}(b))^{N-1} = 0.$$

After rearrangements, we have

$$\frac{(N-1)(x-b)}{\beta'(\beta^{-1}(b))} - \beta^{-1}(b) = 0.$$

- (6) At a symmetric equilibrium, $b = \beta(x)$, and thus we have the following differential equation

$$(N-1)x^{N-1} - (N-1)\beta(x)x^{N-2} = x^{N-1}\beta'(x),$$

or equivalently,

$$\frac{d}{dx}(x^{N-1}\beta(x)) = (N-1)x^{N-1}.$$

(7) Since $\beta(0) = 0$, we have

$$\beta(x) = \frac{N-1}{N}x.$$

Remark: The derivation of β is only heuristic: We have not formally established that if the other $N-1$ bidders follow β , then it is indeed optimal for a bidder with value x to bid $\beta(x)$. \square

Part 3: Prove β to be a symmetric equilibrium strategy.

- (1) Suppose that all but bidder 1 follow the strategy β . We will argue that in that case it is optimal for bidder 1 to follow β also.
- (2) Since β is an increasing and continuous function, in equilibrium the bidder with the highest value submits the highest bid and wins the auction.
- (3) It is not optimal for bidder 1 to bid a $b > \beta(1) = \frac{N-1}{N}$.
- (4) Suppose bidder 1 bids an amount $b \leq \beta(1) = \frac{N-1}{N}$. Denote by $z = \beta^{-1}(b)$. Then bidder 1's expected payoff from bidding $\beta(z) = b$ when her value is x is as follows:

$$\Pi(b, x) = (\beta^{-1}(b))^{N-1} \times (x - b) = z^{N-1}x - \frac{N-1}{N}z^N.$$

(5) We thus obtain that

$$\Pi(\beta(x), x) - \Pi(\beta(z), x) = (z - x)z^{N-1} - \int_x^z y^{N-1} dy \geq 0$$

regardless of whenever $z \geq x$ or $z \leq x$.

Remark: The phenomenon that $\beta^I(x) = \frac{N-1}{N}x < x$ is called bid shading. As the number of bidders increases, each bidder shades less. For each bidder, there is both an incentive to bid higher, so that she wins with higher probability ($s^{-1}(x)$ increases), and an incentive to bid lower, so that when she wins, she pays less and benefits more ($x - b$ increases). Bid shading is exactly the result of such a trade-off. \square

4.24 Heuristic derivation of symmetric equilibrium strategy for general model

- (1) Suppose that bidder $j \neq 1$ follow the symmetric, increasing, and differentiable equilibrium strategy $\beta^I \equiv \beta$. Suppose bidder 1 receives a signal, $X_1 = x$, and bids b . We wish to determine the optimal b .
- (2) Notice that it can never be optimal to choose a bid $b > \beta(\omega)$, since in that case, bidder 1 would win for sure and could do better by reducing her bid slightly, so she still wins for sure but pays less.
- (3) A bidder with value 0 would never submit a positive bid, since she would make a loss if she were to win the auction. Thus, we must have $\beta(0) = 0$.
- (4) Bidder 1 wins the auction whenever she submits the highest bid—that is, whenever $\max_{i \neq 1} \beta(X_i) < b$. Since β is increasing, $\max_{i \neq 1} \beta(X_i) = \beta(\max_{i \neq 1} X_i) = \beta(Y_1)$, where $Y_1 \equiv Y_1^{(N-1)}$, the highest of $N-1$ values. Her expected payoff is therefore

$$G(\beta^{-1}(b)) \times (x - b),$$

where G is the distribution of Y_1 .

(5) Maximizing this with respect to b yields the first-order condition:

$$\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(x - b) - G(\beta^{-1}(b)) = 0, \quad (4.1)$$

where $g = G'$ is the density of Y_1 .

(6) At a symmetric equilibrium, $b = \beta(x)$, and thus Equation (4.1) yields the differential equation

$$G(x)\beta'(x) + g(x)\beta(x) = xg(x), \quad (4.2)$$

or equivalently,

$$\frac{d}{dx}(G(x)\beta(x)) = xg(x).$$

(7) Since $\beta(0) = 0$, we have

$$\beta(x) = \frac{1}{G(x)} \int_0^x yg(y) dy = \mathbf{E}[Y_1 \mid Y_1 < x]. \quad (4.3)$$

The last equality holds due to the conditional probability.

Remark: The derivation of β is only heuristic because Equation (4.2) is merely a necessary condition: We have not formally established that if the other $N - 1$ bidders follow β , then it is indeed optimal for a bidder with value x to bid $\beta(x)$.

4.25 Proposition: Symmetric equilibrium strategies in a first-price auction are given by

$$\beta^1(x) = \mathbf{E}[Y_1 \mid Y_1 < x], \quad (4.4)$$

where Y_1 is the highest of $N - 1$ independently drawn values.

Proof. (1) Suppose that all but bidder 1 follow the strategy $\beta^1 \equiv \beta$ given in Equation (4.4). We will argue that in that case it is optimal for bidder 1 to follow β also.

(2) Since β is an increasing and continuous function, in equilibrium the bidder with the highest value submits the highest bid and wins the auction.

(3) It is not optimal for bidder 1 to bid a $b > \beta(\omega)$.

(4) Suppose bidder 1 bids an amount $b \leq \beta(\omega)$. Denote by $z = \beta^{-1}(b)$. Then bidder 1's expected payoff from bidding $\beta(z) = b$ when her value is x as follows:

$$\begin{aligned} \Pi(b, x) &= G(\beta^{-1}(b)) \times (x - b) \\ &= G(z)x - G(z) \mathbf{E}[Y_1 \mid Y_1 < z] \\ &= G(z)x - \int_0^z yg(y) dy && \text{By Equation (4.3)} \\ &= G(z)x - G(z)z + \int_0^z G(y) dy && \text{Integration by parts} \\ &= G(z)(x - z) + \int_0^z G(y) dy \end{aligned}$$

(5) We thus obtain that

$$\Pi(\beta(x), x) - \Pi(\beta(z), x) = G(z)(z - x) - \int_x^z G(y) dy \geq 0$$

regardless of whether $z \geq x$ or $z \leq x$.

To be more precise, bidding an amount $\beta(z') > \beta(x)$ rather than $\beta(x)$ results in a loss equal to the gray area in Figure 4.1: $G(z')(z' - x)$ is the area of the right rectangle, and $\int_x^{z'} G(y) dy$ is the area of the graph of G from x to z' .




Figure 4.1: Losses from over- and underbidding in a first-price auction.

Similarly, bidding an amount $\beta(z'') < \beta(x)$ results in a loss equal to the blue area.

□

4.26 Remark: From Equation (4.3), we have

$$\beta^I(x)G(x) = \int_0^x yg(y) = \int_0^x y dG(y) = xG(x) - \int_0^x G(y) dy,$$

and hence the equilibrium bid can be rewritten as

$$\beta^I(x) = x - \int_0^x \frac{G(y)}{G(x)} dy.$$

This shows that the bid is, naturally, less than the value x .

4.27 Since the degree of “shading” (the amount by which the bid is less than the value)

$$\frac{G(y)}{G(x)} = \left[\frac{F(y)}{F(x)} \right]^{N-1}$$

depends on the number of competing bidders and as N increases, approaches 0.

Thus, for fixed F , as the number of bidders increases, the equilibrium bid $\beta^I(x)$ approaches x .

4.28 Example 4.23 can be derived by letting $\omega = 1$ and $F(x) = x$.

4.29 Example: Values are exponentially distributed (the rate parameter λ is 2) on $[0, \infty)$, and there are only two bidders.

$F(x) = 1 - e^{-\lambda x} = 1 - e^{-2x}$, and $N = 2$, then

$$\beta(x) = \beta^I(x) = x - \int_0^x \frac{F(y)}{F(x)} dy = \frac{1}{2} - \frac{xe^{-2x}}{1 - e^{-2x}}.$$




Figure 4.2: Equilibrium strategy.

The equilibrium bidding strategy is depicted in Figure 4.2. The figure highlights the fact that with the exponentially distributed values, even a bidder with a very high value—say, \$1 million—will not bid more than 50 cents.

- The bidder is facing the risk of a big loss by not bidding higher.
- The probability that the bidder with a high value will lose in equilibrium is infinitesimal. Hence bidders with high values are willing to bid very small amounts.

4.30 Variation: Two-player first-price sealed-bid auction (risk-neutral, uniform distribution, linear Bayesian Nash equilibrium).

Suppose there are two bidders, $i = 1, 2$.

The bidders' valuations x_1 and x_2 for a good are independently and uniformly distributed on $[0, 1]$.

Bidders submit their bids b_1 and b_2 simultaneously. The higher bidder wins the good and pays her bidding price; the other bidder gets and pays nothing. In the case that $b_1 = b_2$, the winner is determined by a flip of a coin.

Formulate it as a Bayesian game, and find all the linear Bayesian Nash equilibria.

Answer. The formulation is as follows:

- $A_1 = A_2 = [0, \infty)$, bids $b_i \in A_i$;
- $T_1 = T_2 = [0, 1]$, valuations $x_i \in T_i$;
- $P_i(x_j)$ is the uniform distribution on $[0, 1]$;
- For any $x_i \in T_i$, player i 's payoff is

$$\Pi_i(b_1, b_2; x_1, x_2) = \begin{cases} x_i - b_i, & \text{if } b_i > b_j, \\ \frac{x_i - b_i}{2}, & \text{if } b_i = b_j, \\ 0, & \text{if } b_i < b_j. \end{cases}$$

Player i 's linear strategy is a function $\beta_i^I(x_i) = a_i + c_i x_i$ from $[0, 1]$ to $[0, \infty)$, where $1 > a_i \geq 0$ and $c_i > 0$.

$(\beta_1^I(x_1), \beta_2^I(x_2))$ is a Bayesian Nash equilibrium if for each $x_i \in [0, 1]$, $\beta_i^I(x_i)$ maximizes

$$\mathbb{E}_{x_j} u_i(b_i, \beta_j^I(x_j); x_i, x_j) = \text{Prob}(b_i > \beta_j^I(x_j)) \cdot (x_i - b_i) + \text{Prob}(b_i = \beta_j^I(x_j)) \frac{x_i - b_i}{2} = \frac{b_i - a_j}{c_j} (x_i - b_i).$$

Therefore, player i 's best response solves

$$\max_{a_j \leq b_i \leq a_j + c_j} \frac{b_i - a_j}{c_j} (x_i - b_i).$$

The unconstrained maximum is

$$\bar{b}_i = \frac{x_i + a_j}{2}.$$

If $\bar{b}_i \in [a_j, a_j + c_j]$, then the best response $b_i(x_i) = \bar{b}_i$. If $\bar{b}_i < a_j$, then $b_i(x_i) = a_j$. If $\bar{b}_i > a_j + c_j$, then $b_i(x_i) = a_j + c_j$.

Thus, player i 's best response is

$$\beta_i^I(x_i) = \begin{cases} a_j, & \text{if } x_i \leq a_j, \\ \frac{x_i + a_j}{2}, & \text{if } a_j < x_i \leq a_j + 2c_j, \\ a_j + c_j, & \text{if } x_i > a_j + 2c_j. \end{cases}$$

Since we want the strategy β_i^I to be a linear function on $[0, 1]$, there are only three cases:

$$[0, 1] \subseteq \begin{cases} (-\infty, a_j], \\ [a_j, a_j + 2c_j], \\ [a_j + 2c_j, \infty). \end{cases}$$

Case 1 violates the assumption $a_j < 1$. Case 3 violates the assumptions $a_j \geq 0$ and $c_j > 0$ which imply $a_j + 2c_j > 0$.

Therefore, $[0, 1] \subseteq [a_j, a_j + 2c_j]$, i.e., $\beta_i^I(x_i) = \frac{x_i + a_j}{2}$ for $x_i \in [0, 1]$. Hence for $i = 1, 2$ and $j \neq i$,

$$a_i = a_j/2, \quad c_i = 1/2.$$

This yields

$$a_1 = a_2 = 0, \quad c_1 = c_2 = 1/2.$$

Therefore, the unique linear Bayesian Nash equilibrium is

$$\beta_1^I(x_1) = x_1/2, \quad \beta_2^I(x_2) = x_2/2.$$

□

Remark:

- $a_i \geq 0$ reflects the fact that bids can not be negative.
- $c_i > 0$ implies high bids for high valuation.
- If $a_i \geq 1$, then, together with $c_i > 0$, it follows that $\beta_i^I(x_i) > x_i$ for each $x_i \in [0, 1]$. With such a bid, player i would always end up with negative payoffs. This bid function is certainly non-optimal. Thus we assume $a_i < 1$.
- The linear Bayesian Nash equilibrium is a symmetric equilibrium.

4.31 Variation: Two-player first-price sealed-bid auction (risk-averse, uniform distribution, linear Bayesian Nash equilibrium).

Consider the following first-price sealed-bid auction. Suppose there are two bidders, $i = 1, 2$. The bidders' valuations x_1 and x_2 for a good are independently and uniformly distributed on $[0, 1]$. The bidders have preferences represented by the utility functions $u_i(x) = x^{\alpha_i}$ where $0 < \alpha_i \leq 1$, $i = 1, 2$. Bidders submit their bids b_1 and b_2 simultaneously. The higher bidder wins the good and pays her bidding price, so that $x = x_i - b_i$; the other bidder gets and pays nothing, so that $x = 0$. In the case that $b_1 = b_2$, the winner is determined by a flip of a coin. Find a Bayesian Nash equilibrium (b_1, b_2) in which b_i is a linear function of v_i , $i = 1, 2$.

Answer. • There are two players;

- Type spaces: $T_1 = \{x_1 \mid x_1 \in [0, 1]\}$, and $T_2 = \{x_2 \mid x_2 \in [0, 1]\}$;
- Action spaces: $A_1 = A_2 = [0, 1]$;
- Strategy spaces: $S_1 = \{\text{functions from } T_1 \text{ to } A_1\}$ and $S_2 = \{\text{functions from } T_2 \text{ to } A_2\}$;
- Payoff:

$$\pi_i(b_i, b_j; x_i, x_j) = \begin{cases} (x_i - b_i)^{\alpha_i}, & \text{if } b_i > b_j; \\ 0, & \text{if } b_i < b_j. \end{cases}$$

Suppose (β_1^I, β_2^I) is a linear Bayesian Nash equilibrium, where

$$\beta_i^I(x_i) = a_i + c_i x_i, \quad i = 1, 2,$$

where a_i, c_i are to be determined. Here we should assume $c_i > 0$.

- For bidder 1, when x_1 is drawn, given bidder 2's strategy β_2^I , $\beta_1^I(x_1)$ will maximize her expected payoff

$$\begin{aligned} \mathbb{E}[\Pi_1(b_1, \beta_2^I(x_2); x_1, x_2)] &= (x_1 - b_1)^{\alpha_1} \text{Prob}(\beta_2^I(x_2) < b_1) \\ &= (x_1 - b_1)^{\alpha_1} \text{Prob}\left(x_2 < \frac{b_1 - a_2}{c_2}\right) = (x_1 - b_1)^{\alpha_1} \frac{b_1 - a_2}{c_2}. \end{aligned}$$

Note that when bidder 1 chooses b_1 , the probability that $b_1 = \beta_2^I(x_2)$ is 0, and thus we do not need to consider that.

Therefore first order condition implies

$$\beta_1(x_1) = \frac{\alpha_1}{1 + \alpha_1} a_2 + \frac{1}{1 + \alpha_1} x_1,$$

and hence

$$a_1 = \frac{\alpha_1}{1 + \alpha_1} a_2, \quad c_1 = \frac{1}{1 + \alpha_1}. \quad (4.5)$$

- For bidder 2, when x_2 is drawn, given bidder 1's strategy β_1^I , $\beta_2^I(x_2)$ will maximize her expected payoff

$$\begin{aligned} \mathbb{E}[\Pi_2(\beta_1^I(x_1), b_2; x_1, x_2)] &= (x_2 - b_2)^{\alpha_2} \text{Prob}(\beta_1^I(x_1) < b_2) \\ &= (x_2 - b_2)^{\alpha_2} \text{Prob}\left(x_1 < \frac{b_2 - a_1}{c_1}\right) = (x_2 - b_2)^{\alpha_2} \frac{b_2 - a_1}{c_1}. \end{aligned}$$

Note that when bidder 2 chooses b_2 , the probability that $b_2 = \beta_1^I(x_1)$ is 0, and thus we do not need to consider that.

Therefore first order condition implies

$$\beta_2^I(x_2) = \frac{\alpha_2}{1 + \alpha_2} a_1 + \frac{1}{1 + \alpha_2} x_2,$$

and hence

$$a_2 = \frac{\alpha_2}{1 + \alpha_2} a_1, \quad c_2 = \frac{1}{1 + \alpha_2}. \quad (4.6)$$

Solving Equations (4.5) and (4.6), we will have $a_1 = a_2 = 0$. \square

4.5 Revenue comparison

4.32 With Proposition 4.17 in hand, we can compute how much each bidder expects to pay in equilibrium in a second-price auction.

Fix a bidder—say, 1—and let the random variable $Y_1 \equiv Y_1^{(N-1)}$ denote the highest value among the $N-1$ remaining bidders. In other words, Y_1 is the highest-order statistic of X_2, X_3, \dots, X_N . Let G denote the distribution function of Y_1 . Clearly, for all y , $G(y) = F(y)^{N-1}$. In a second-price auction, the expected payment by a bidder with value x can be written as

$$\begin{aligned} m^{\text{II}}(x) &= \text{Prob}[\text{Win}] \times \mathbf{E}[\text{second highest bid} \mid x \text{ is the highest bid}] \\ &= \text{Prob}[\text{Win}] \times \mathbf{E}[\text{second highest value} \mid x \text{ is the highest value}] \\ &= G(x) \times \mathbf{E}[Y_1 \mid Y_1 < x] \end{aligned}$$

4.33 In a first-price auction, the winner pays what she bids, and thus the expected payment by a bidder with value x is

$$m^{\text{I}}(x) = \text{Prob}[\text{Win}] \times \text{Amount bid} = G(x) \times \mathbf{E}[Y_1 \mid Y_1 < x],$$

which is the same as in a second-price auction.

4.34 Figure 2.3 ?????

4.35 Proposition: With independently and identically distributed private values, the expected revenue in a first-price auction is the same as the expected revenue in a second-price auction.

Proof. (1) The *ex ante* expected payment of a particular bidder in either auction is

$$\mathbf{E}_x[m^A(x)] = \int_0^\omega m^A(x) f(x) dx = \int_0^\omega \left(\int_0^x y g(y) dy \right) f(x) dx,$$

where A is I or II.

(2) The expected revenue accruing to the seller $\mathbf{E}[R^A]$ is just N times the *ex ante* expected payment of an individual bidder, so

$$\mathbf{E}[R^A] = N \times \mathbf{E}[m^A(X)] = N \int_0^\omega y(1 - F(y)) g(y) dy.$$

\square

Furthermore, note that the density of $Y_2^{(N)}$, the second highest of N values, $f_2^{(N)}(y) = N(1 - F(y))f_1^{(N-1)}(y)$, and since $f_1^{(N-1)}(y) = g(y)$, we can write

$$\mathbb{E}[R^A] = \int_0^\omega y f_2^{(N)}(y) dy = \mathbb{E}[Y_2^{(N)}].$$

In either case, the expected revenue is just the expectation of the second-highest value.

4.36 Remark: In specific realizations of the values the price at which the object is sold may be greater in one auction or the other.

Example: There are two bidders and values are uniformly distributed. The equilibrium strategy in a first-price auction is $\beta^I(x) = \frac{1}{2}x$.

- If the realized value are such that $\frac{1}{2}x_1 > x_2$, then the revenue in a first-price auction, $\frac{1}{2}x_1$, is greater than that in a second-price auction, x_2 .
- If $\frac{1}{2}x_1 < x_2 < x_1$, the opposite is true.

4.37 Definition: Suppose X is a random variable with distribution function F . Let Z be a random variable whose distribution conditional on $X = x$, $H(\cdot | X = x)$ is such that for all x , $\mathbb{E}[Z | X = x] = 0$. Suppose $Y = X + Z$ is the random variable obtained from the first drawing X from F and then for each realization $X = x$, drawing a Z from the conditional distribution $H(\cdot | X = x)$ and adding it to X . Let G be the distribution of Y . We will say that G is a mean-preserving spread of F .

As the name suggests, while the random variables X and Y have the same mean—that is, $\mathbb{E}[X] = \mathbb{E}[Y]$ —the variable Y is “more spread out” than X since it is obtained by adding a “noise” variable Z to X .

4.38 Given two distributions F and G with the same mean, we say that F second-order stochastically dominates G if for all concave functions $U: [0, \omega] \rightarrow \mathbb{R}$,

$$\int_0^\omega U(x)f(x) dx \geq \int_0^\omega U(y)g(y) dy,$$

where f and g are density functions of F and G respectively.

4.39 Lemma: G is a mean-preserving spread of F if and only if F second-order stochastically dominates G . (Exercise)

4.40 Proposition: With independently and identically distributed private values, the distribution of equilibrium prices in a second-price auction L^{II} is a mean-preserving spread of the distribution of equilibrium prices in a first-price auction L^{I} .

Proof. (1) The revenue in a second-price auction is $R^{\text{II}} = Y_2^{(N)}$; the revenue in a first-price auction is $R^{\text{I}} = \beta(Y_1^{(N)})$, where $\beta \equiv \beta^{\text{I}}$ is the symmetric equilibrium strategy.

(2) We have

$$\mathbb{E}[R^{\text{II}} | R^{\text{I}} = p] = \mathbb{E}[Y_2^{(N)} | Y_1^{(N)} = \beta^{-1}(p)] = \mathbb{E}[Y_1^{(N-1)} | Y_1^{(N-1)} < \beta^{-1}(p)].$$

(3) By Equation (4.4), we have

$$\mathbb{E}[R^{\text{II}} | R^{\text{I}} = p] = \beta(\beta^{-1}(p)) = p.$$

(4) Therefore, there exists a random variable Z such that the distribution of R^{II} is the same as that of $R^{\text{I}} + Z$ and $\mathbb{E}[Z | R^{\text{I}} = p] = 0$. Thus, L^{II} is a mean-preserving spread of L^{I} .

□

4.41 Remark: It is clear that the revenues in a second-price auction are more variable than in its first-price counterpart. In the former, the prices can range between 0 and ω ; in the latter, they can only range between 0 and $E[Y_1]$.

From the perspective of the seller, a second-price auction is riskier than a first-price auction. Every risk-averse seller prefers the latter to the former, assuming that bidders are risk-neutral.

Figure?????

4.6 Reserve prices

4.42 In the analysis so far, the seller has played a passive role. Indeed, we have implicitly assumed that the seller parts with the object at whatever price it will fetch. In many instances, sellers reserve the right to not sell the object if the price determined in the auction is lower than some threshold amount—say, $r > 0$. Such a price is called the reserve price. We now examine what effect such a reserve price has on the expected revenue accruing to the seller.

4.43 Reserve prices in second-price auctions: Suppose that the seller sets a “small” reserve price of $r > 0$.

A reserve price makes no difference to the behavior of the bidders; it is still a weakly dominant strategy to bid one’s value.

- (1) Since the price at which the object is sold can never be lower than r , no bidder with a value $x < r$ can make a positive profit.
- (2) Consider bidder 1 with value $x_1 \geq r$, and suppose that $p_1 = \max_{j \neq i} b_j$ is the highest competing bid.
- (3) By bidding x_1 , bidder 1 will win if $x_1 > p_1$ and not if $x_1 < p_1$ (if $x_1 = p_1$, bidder 1 is indifferent between winning and losing).
- (4) Suppose, however, that she bids an amount $z_1 < x_1$.
 - If $p_1 > x_1 > z_1$, she still loses no matter what r is.
 - If $x_1 > z_1 \geq p_1$ and $z_1 \geq r$, then she still wins, and her profit is still $x_1 - \max\{p_1, r\}$.
 - If $x_1 > p_1 > z_1$, then she loses no matter what r is, whereas if she had bid x_1 , she would have made a positive profit.

Thus, bidding less than x_1 can never increase her profit but in some circumstances may actually decrease it.

- (5) A similar argument shows that it is not profitable to bid more than x_1 .

The expected payment of a bidder with value r is now just $rG(r)$, and the expected payment of a bidder with value $x \geq r$ is

$$m^{\text{II}}(x, r) = rG(r) + [G(x) - G(r)] \cdot E[Y_1 \mid Y_1 < x] = rG(r) + \int_r^x yg(y) dy,$$

since the winner pays the reserve price r whenever the second-highest bid is below r .

4.44 Reserve prices in first-price auctions: Suppose that the seller sets a “small” reserve price of $r > 0$.

- (1) Since the price at which the object is sold can never be lower than r , no bidder with a value $x < r$ can make a positive profit.
- (2) if β^{I} is a symmetric equilibrium of the first-price auction with reserve price r , it must be that $\beta^{\text{I}}(r) = r$. This is because a bidder with value r wins only if all other bidders have values less than r and, in that case, can win with a bid of r itself.

- (3) In all other respects, the analysis of a first-price auction is unaffected, and in a manner analogous to Proposition 4.25, we obtain

$$\frac{d}{dx}(G(x)\beta^I(x)) = xg(x),$$

and hence

$$G(x)\beta^I(x) - G(r)\beta^I(r) = \int_r^x yg(y) dy.$$

- (4) Therefore,

$$\beta^I(x) = r \frac{G(r)}{G(x)} + \frac{1}{G(x)} \int_r^x yg(y) dy = \mathbf{E} [\max\{Y_1^{(N-1)}, r\} \mid Y_1^{(N-1)} < x].$$

The expected payment of a bidder with value $x \geq r$ is

$$m^I(x, r) = G(x) \times \beta^I(x) = rG(r) + \int_r^x yg(y) dy,$$

which is the same as the expected payment in a second-price auction. Proposition 4.35 generalizes so as to accommodate reserve prices.

4.45 The *ex ante* expected payment of a bidder is

$$\mathbf{E}_x[m^A(X, r)] = \int_r^\omega m^A(x, r)f(x) dx = r(1 - F(r))G(r) + \int_r^\omega y(1 - F(y))g(y) dy.$$

4.46 Suppose that the seller attaches a value $x_0 \in [0, \omega)$. This means that if the object is left unsold, the seller would derive a value x_0 from its use.

- (1) The seller would not set a reserve price r that is lower than x_0 .
- (2) The expected payoff of the seller from setting a reserve price $r \geq x_0$ is

$$\Pi_0 = N \times \mathbf{E}[m^A(X, r)] + F(r)^N x_0.$$

- (3) Differentiating this with respect to r , we obtain,

$$\frac{d\Pi_0}{dr} = N[1 - F(r) - rf(r)]G(r) + NG(r)f(r)x_0.$$

Recall that the hazard rate function associated with the distribution F is defined as $\lambda(x) = f(x)/[1 - F(x)]$. Thus, we have

$$\frac{d\Pi_0}{dr} = N[1 - (r - x_0)\lambda(r)](1 - F(r))G(r).$$

4.47 A revenue maximizing seller should always set a reserve price that exceeds her value:

- If $x_0 > 0$, then the derivative of Π_0 at $r = x_0$ is positive, implying that the seller should set a reserve price $r > x_0$.
- If $x_0 = 0$, then the derivative of Π_0 at $r = 0$ is 0, but as long as $\lambda(r)$ is bounded, the expected payment attains a local minimum at 0, so a small reserve price leads to an increase in revenue.

4.48 Example: A reserve price that exceeds x_0 leads to an increase in revenue.

Consider a second-price auction with two bidders and suppose $x_0 = 0$.

- By setting a positive reserve price r , the seller runs the risk that if the highest value among the bidders, $Y_1^{(N)}$, is smaller than r , the object will remain unsold. The probability of this event is $F(r)^2$, and the loss is at most r , so for small r , the expected loss is at most $rF(r)^2$.
- This potential loss is offset by the possibility that while the highest value $Y_1^{(N)}$ exceeds r , the second highest value $Y_2^{(N)}$ is smaller than r . The application of the reserve price means that the object will be sold for r rather than $Y_2^{(N)}$. The probability of this event is $2F(r)(1 - F(r))$, and the gain is of order r , so the expected gain is of order $2rF(r)(1 - F(r))$.
- The expected gain from setting a small reserve price exceeds the expected loss. This fact is sometimes referred to as the exclusion principle, since it implies, in effect, that it is optimal for the seller to exclude some bidders—those with value below the reserve price—from the auction even though their values exceed x_0 .

4.49 The first-order condition implies that the optimal reserve price r^* must satisfy

$$(r^* - x_0)\lambda(r^*) = 1.$$

If $\lambda(\cdot)$ is increasing, this condition is also sufficient:

- If $r < r^*$, then $\frac{d\Pi_0}{dr}(r) > 0$;
- If $r > r^*$, then $\frac{d\Pi_0}{dr}(r) < 0$;
- Hence r^* is global optimal.

4.50 Note that the optimal reserve price does not depend on the number of bidders. The reason is that a reserve price comes into play only in instances when there is a single bidder with a value that exceeds the reserve price.

4.51 A positive reserve price r results in bidders with low values, lying below r , being excluded from the auction. Since their equilibrium payoffs are zero, such bidders are indifferent between participating in the auction or not. An alternative instrument that the seller can also use to exclude buyers with low values is an entry fee—a fixed and non-refundable amount that bidders must pay the seller prior to the auction in order to be able to submit bids.

A reserve price of r excludes all bidders with value $x < r$. The same set of bidders can be excluded by asking each bidder to pay an entry fee $e = G(r) \times r$. Notice that after paying e , the expected payoff of a bidder with value $x < r$ would not find it worthwhile to pay e in order to participate in the auction.

4.52 A reserve price raises the revenue to the seller but may have a detrimental effect on efficiency.

Suppose that the value that the seller attaches to the object is 0.

- In the absence of a reserve price, the object will always be sold to the highest bidder and in the symmetric model studied here, that is also the bidder with the highest value. Thus, both the first- and second-price auctions allocate efficiently in the sense that the object ends up in the hands of the person who values it the most.
- If the seller sets a reserve price $r > 0$, there is a positive probability that the object will remain in the hands of the seller and this is inefficient.

This simple observation implies that there may be a trade-off between efficiency and revenue.

4.53 Remark: We have implicitly assumed that the seller can credibly commit to not sell the object if it cannot be sold at or above the reserve price. This commitment is particularly important because by setting a reserve price the seller is giving up some gains from trade. Without such a commitment, buyers may anticipate that the object, if durable, will be offered for sale again in a later auction and perhaps with a lower reserve price. These expectations may affect

their bidding behavior in the first auction. Indeed, in the absence of a credible “no sale” commitment, the problem confronting a seller is analogous to that of a durable goods monopoly. In both, a potential future sale may cause buyers to wait for lower prices, and this may reduce demand today. In effect, potential future sales may compete with current sales. In response, the seller may have to set lower reserve prices today than would be optimal in a one-time sale or if the good were perishable.

4.54 Remark: We have assumed that the reserve price is publicly announced prior to the auction. In many situations, especially in art auctions, it is announced that there is a reserve price, but the level of the reserve price is not disclosed. In effect, the seller can opt to not sell the object after learning all the bids and thus the price. But this is rational only if the seller anticipates that in a future sale the price will be higher. Once again, buyers’ expectations regarding future sales may affect the bidding in the current auction.

4.7 The revenue equivalence principle

4.55 The auction forms we consider all have the feature that buyers are asked to submit bids—amounts of money they are willing to pay. These bids alone determine who wins the object and how much the winner pays.

4.56 We will say that an auction is standard if the rules of the auction dictate that the person who bids the highest amount is awarded the object.

- Both first- and second-price auctions are standard.
- A third-price auction, discussed later, in which the winner is the person bidding the highest amount but pays the third-highest bid is standard.
- An example of a nonstandard method is a lottery in which the chances that a particular bidder wins is the ratio of her bid to the total amount bid by all. Such a lottery is nonstandard, since the person who bids the most is not necessarily the one who is awarded the object.

4.57 Given a standard auction form, A , and a symmetric equilibrium β^A of the auction, let $m^A(x)$ be the equilibrium expected payment by a bidder with value x .

4.58 Theorem (Revenue equivalence principle): Suppose that values are independently and identically distributed and all bidders are risk neutral. Then any symmetric and increasing equilibrium of any standard auction, such that the expected payment of a bidder with value zero is zero, yields the same expected revenue to the seller.

Proof. (1) Consider a standard auction form, A , and fix a symmetric equilibrium β of A . Let $m^A(x)$ be the equilibrium expected payment in auction A by a bidder with value x . Suppose that β is such that $m^A(0) = 0$.

(2) Consider a particular bidder—say, 1—and suppose other bidders are following the equilibrium strategy β . Consider the expected payoff of bidder 1 with value x and when she bids $\beta(z)$ instead of the equilibrium bid $\beta(x)$.

(3) Bidder 1 wins when her bid $\beta(z)$ exceeds the highest competing bid $\beta(Y_1^{(N-1)})$, or equivalently, when $z > Y_1^{(N-1)}$.

(4) Her expected payoff is

$$\Pi^A(z, x) = G(z)x - m^A(z),$$

where $G(z) \equiv F(z)^{N-1}$ is the distribution of $Y_1^{(N-1)}$.

(5) Maximization results in the first-order condition,

$$\frac{\partial}{\partial z} \Pi^A(z, x) = g(z)x - \frac{d}{dz} m^A(z) = 0.$$

(6) At an equilibrium it is optimal to bid $z = x$, so we obtain that for all y ,

$$\frac{d}{dy} m^A(y) = g(y)y.$$

(7) Thus,

$$m^A(x) = m^A(0) + \int_0^x yg(y) dy = \int_0^x yg(y) dy = G(x) \times \mathbf{E} [Y_1^{(N-1)} \mid Y_1^{(N-1)} < x]. \quad (4.7)$$

Since the right-hand side does not depend on the particular auction form A , the expected revenue of the seller is constant.

□

4.59 Example: Values are uniformly distributed on $[0, 1]$.

$F(x) = x$, then $G(x) = x^{N-1}$ and for any standard auction satisfying $m^A(0) = 0$, we have

$$m^A(x) = \frac{N-1}{N} x^N,$$

and

$$\mathbf{E} [m^A(X)] = \frac{N-1}{N(N+1)},$$

while the expected revenue is

$$\mathbf{E}[R^A] = N \times \mathbf{E} [m^A(X)] = \frac{N-1}{N+1}.$$

4.8 All-pay auction

4.60 Consider an all-pay auction with the following rules. Each bidder submits a bid, and the highest bidder wins the object. The unusual aspect of an all-pay auction is that all bidders pay what they bid.

The all-pay auction is a useful model of lobbying activity. In such models, different interest groups spend money—their “bids”—in order to influence government policy and the group spending the most—the highest “bidder”—is able to tilt policy in its favored direction, thereby “winning the auction.” Since money spent on lobbying is a sunk cost borne by all groups regardless of which group is successful in obtaining its preferred policy, such situations have a natural all-pay aspect.

4.61 Suppose for the moment that there is a symmetric, increasing equilibrium of the all-pay auction such that the expected payment of a bidder with value 0 is 0. Then we know that the expected payment in such an equilibrium must be the same as in Equation (4.7).

Now in an all-pay auction, the expected payment of a bidder with value x is the same as her bid—she forfeits her bid regardless of whether she wins or not—and so if there is a symmetric, increasing equilibrium of the all-pay auction β^{AP} , it must be that

$$\beta^{\text{AP}}(x) = m^A(x) = \int_0^x yg(y) dy$$

4.62 Proposition: $\beta^{\text{AP}}(x) = m^{\text{A}}(x) = \int_0^x yg(y) dy$ is a symmetric equilibrium strategy in an all-pay auction.

Proof. (1) Suppose that all bidders except 1 are following the strategy $\beta \equiv \beta^{\text{AP}}$.

(2) If she bids an amount $\beta(z)$, the expected payoff of a bidder with value x is

$$G(z)x - \beta(z) = G(z)x - \int_0^z yg(y) dy.$$

(3) By integrating by parts, we have

$$G(z)(x - z) + \int_0^z G(y) dy,$$

which is the same as the payoff obtained in a first-price auction by bidding $\beta^{\text{I}}(z)$ against other bidders who are following β^{I} .

(4) For the same reasons as in Proposition 4.25, this is maximized by choosing $z = x$. Thus, β^{AP} is a symmetric equilibrium. □

4.9 Third-price auction

4.63 Suppose that there are at least three bidders. Consider a sealed-bid auction in which the highest bidder wins the object but pays a price equal to the third-highest bid. A third-price auction, as it is called, is a purely theoretical construct: There is no known instance of such a mechanism actually being used. It is an interesting construct nevertheless; equilibria of such an auction display some unusual properties, and it leads to a better understanding of the workings of the standard auction forms.

4.64 Suppose that there are three bidders, and for the moment that there is a symmetric, increasing equilibrium of the third-price auction—say, β^{III} —such that the expected payment of a bidder with value 0 is 0.

(1) Since the assumptions of Theorem 4.58 are satisfied, we must have that for all x , the expected payment of a bidder with value x in a third-price auction is

$$m^{\text{III}}(x) = \int_0^x yg(y) dy. \quad (4.8)$$

(2) On the other hand, consider bidder 1, and suppose that she wins in equilibrium when her value is x .

(3) Winning implies that her value x exceeds the highest of the other $N-1$ values—that is, $Y_1^{(N-1)} < x$. The price bidder 1 pays is the random variable $\beta^{\text{III}}(Y_2^{(N-1)})$, where $Y_2^{(N-1)}$ is the second highest of the $N-1$ other values.

(4) The density of $Y_2^{(N-1)}$, conditional on the event that $Y_1^{(N-1)} < x$, can be written as

$$f_2^{(N-1)}(y \mid Y_1^{(N-1)} < x) = \frac{1}{F_1^{(N-1)}(x)} \times (N-1)(F(x) - F(y)) \times f_1^{(N-2)}(y),$$

where $(N-1)(F(x) - F(y))$ is the probability that $Y_1^{(N-1)}$ exceeds $Y_2^{(N-1)} = y$ but is less than x , and $f_1^{(N-2)}(y)$ is the density of the highest of $N-2$ values.

(5) Thus, the expected payment of a bidder with value x in a third-price auction can then be written as

$$m^{\text{III}}(x) = F_1^{(N-1)}(x) \times \mathbf{E} [\beta^{\text{III}}(Y_2^{(N-1)}) \mid Y_1^{(N-1)} < x]. \quad (4.9)$$

(6) Equating Equations (4.8) and (4.9), we obtain that

$$\int_0^x yg(y) dy = F_1^{(N-1)}(x) \times \mathbf{E} [\beta^{\text{III}}(Y_2^{(N-1)}) \mid Y_1^{(N-1)} < x].$$

(7) Since $G(x) = F(x)^{N-1}$, differentiating with respect to x , we have

$$(N-1)f(x) \int_0^x \beta^{\text{III}}(y) f_1^{(N-2)}(y) dy = xg(x) = x(N-1)f(x)F(x)^{N-2}.$$

(8) This can be rewritten as

$$\int_0^x \beta^{\text{III}}(y) f_1^{(N-2)}(y) dy = xF(x)^{N-2} = xF_1^{(N-2)}(x).$$

(9) Differentiating once more with respect to x ,

$$\beta^{\text{III}}(x) f_1^{(N-2)}(x) = x f_1^{(N-2)}(x) + F_1^{(N-2)}(x),$$

and rearranging this we get

$$\beta^{\text{III}}(x) = x + \frac{F_1^{(N-2)}(x)}{f_1^{(N-2)}(x)} = x + \frac{F(x)}{(N-2)f(x)}.$$

This derivation, however, is valid only if β^{III} is increasing, and from the preceding equation it is clear that a sufficient condition for this is that the ratio F/f is increasing. This condition is the same as requiring that $\ln F$ is a concave function or equivalently that F is log-concave.

4.65 Proposition: Suppose that there are at least three bidders and F is log-concave. Symmetric equilibrium strategies in a third-price auction are given by

$$\beta^{\text{III}}(x) = x + \frac{F(x)}{(N-2)f(x)}.$$

4.66 Remark: An important feature of the equilibrium in a third-price auction is worth noting: The equilibrium bid exceeds the value.

- Notice that for much the same reason as in a second-price auction, it is dominated for a bidder to bid below her value in a third-price auction.
- Unlike in a second-price auction, however, it is not dominated for a bidder to bid above her value. Fix some equilibrium bidding strategies of the third-price auction—say, β —and suppose that all bidders except 1 follow β . Suppose bidder 1 with value x bids an amount $b > x$.
 - If $\beta(Y_2^{(N-1)}) < x < \beta(Y_1^{(N-1)}) < b$, this is better than bidding b , since it results in a profit, whereas bidding x would not.
 - If, however, $x < \beta(Y_2^{(N-1)}) < \beta(Y_1^{(N-1)}) < b$, then bidding b results in a loss.

When $b - x \equiv \epsilon$ is small, the gain in the first case is of order ϵ^2 , whereas the loss in the second case is of order ϵ^3 . Thus, it is optimal to bid higher than one's value in a third-price auction.

4.67 Remark: Comparing equilibrium bids in first-, second-, and third-price auctions in case of symmetric private values, we have seen that

$$\beta^{\text{I}}(x) < \beta^{\text{II}}(x) = x < \beta^{\text{III}}(x).$$

(assuming that the distribution of values is log-concave).

4.10 Uncertain number of bidders

4.68 In many auctions—particularly in those of the sealed-bid variety—a bidder may be uncertain about how many other interested bidders there are. In this section we show how the standard model may be amended to include this additional uncertainty.

4.69 Let $\mathcal{N} = \{1, 2, \dots, N\}$ denote the set of potential bidders and let $\mathcal{A} \subseteq \mathcal{N}$ be the set of actual bidders—that is, those that participate in the auction. All potential bidders draw their values independently from the same distribution F .

4.70 Consider an actual bidder $i \in \mathcal{A}$ and let p_n denote the probability that any participating bidder assigns to the event that she is facing n other bidders. Thus, bidder i assigns the probability p_n that the number of actual bidders is $n+1$. The exact process by which the set of actual bidders is determined from the set of potential bidders is symmetric so every actual bidder holds the same beliefs about how many other bidders she faces; the probabilities p_n do not depend on the identity of the bidder nor on her value. It is also important that the set of actual bidders does not depend on the realized values.

- 4.71 (1) Consider a standard auction A and a symmetric and increasing equilibrium β of the auction. Note that since bidders are unsure about how many rivals they face, β does not depend on n .
- (2) Consider the expected payoff of a bidder with value x who bids $\beta(z)$ instead of the equilibrium bid $\beta(x)$. The probability that she faces n other bidders is p_n . In that case, she wins if $Y_1^{(n)}$, the highest of n values drawn from F , is less than z and the probability of this event is $G^{(n)}(z) = F(z)^n$. The overall probability that she will win when she bids $\beta(z)$ is therefore

$$G(z) = \sum_{n=0}^{N-1} p_n G^{(n)}(z).$$

- (3) her expected payoff from bidding $\beta(z)$ when her value is x is then

$$\Pi^A(z, x) = G(z)x - m^A(z).$$

- (4) Suppose that the object is sold using a second-price auction. Even though the number of rival buyers that a particular bidder faces is uncertain, it is still a dominant strategy for her to bid her value. The expected payment in a second-price auction of an actual bidder with value x is therefore

$$m^{\text{II}}(x) = \sum_{n=0}^{N-1} p_n G^{(n)}(x) \mathbf{E} \left[Y_1^{(n)} \mid Y_1^{(n)} < x \right].$$

- (5) Suppose that the object is sold using a first-price auction and that β is a symmetric and increasing equilibrium. The expected payment of an actual bidder with value x is

$$m^{\text{I}}(x) = G(x)\beta(x).$$

(6) The revenue equivalence principle implies that for all x , $m^I(x) = m^II(x)$, so

$$\beta(x) = \sum_{n=0}^{N-1} \frac{p_n G^{(n)}(x)}{G(x)} \mathbf{E} \left[Y_1^{(n)} \mid Y_1^{(n)} < x \right] = \sum_{n=0}^{N-1} \frac{p_n G^{(n)}(x)}{G(x)} \beta^{(n)}(x),$$

where $\beta^{(n)}$ is the equilibrium bidding strategy in a first-price auction in which there are exactly $n + 1$ bidders for sure.

4.72 The equilibrium bid for an actual bidder with value x when she is unsure about the number of rivals she faces is a weighted average of the equilibrium bids in auctions when the number of bidders is known to all.


Mixed-strategy Nash equilibrium

Contents

5.1	Mixed-strategy Nash equilibrium	97
5.2	Examples	99
5.3	Interpretation of mixed-strategy Nash equilibrium	102
5.3.1	Purification	103

5.1 Mixed-strategy Nash equilibrium

5.1 The notion of mixed-strategy Nash equilibrium is designed to model a steady state of a game in which the participants' choices are not deterministic but are regulated by probabilistic rules.

 5.2 Consider a strategic game $G = \langle N, (A_i), (u_i) \rangle$. The mixed extension of G is defined as the strategic game


$$\langle N, (\Delta(A_i)), (U_i) \rangle,$$

in which $\Delta(A_i)$ is the set of probability distributions over A_i , and $U_i: \times_{j \in N} \Delta(A_j) \rightarrow \mathbb{R}$ that assigns to each $\alpha = (\alpha_j) \in \times_{j \in N} \Delta(A_j)$ the expected value under u_i of the lottery induced by α .

If A is finite, then

$$U_i(\alpha) = \sum_{a \in A} (\alpha_1(a_1) \alpha_2(a_2) \cdots \alpha_n(a_n)) \cdot u_i(a).$$

5.3 We refer to a member of $\Delta(A_i)$ as a mixed strategy of player i ; we refer to a member of A_i as a pure strategy. In strategic games, a pure strategy can be viewed as a degenerate mixed strategy that attaches probability one to the pure strategy.

 5.4 Definition: A mixed-strategy Nash equilibrium of a strategic game is a (pure-strategy) Nash equilibrium of its mixed extension.

5.5 Proposition: A profile a^* is a pure-strategy Nash equilibrium if and only if $1 \circ a^* = (1 \circ a_i^*, 1 \circ a_{-i}^*)$ is a mixed-strategy Nash equilibrium, where $1 \circ a_i^*$ is i 's degenerate mixed strategy that attaches probability one to a_i^* .

Proof. “ \Rightarrow ”: Since a^* is a pure-strategy Nash equilibrium, for any $\alpha_i \in \Delta(A_i)$, we have

$$U_i(1 \circ a^*) = u_i(a^*) \geq \sum_{a_i} \alpha_i(a_i) \cdot u_i(a_i, a_{-i}^*) = U_i(\alpha_i, 1 \circ a_{-i}^*).$$

“ \Leftarrow ”: Since $1 \circ a^*$ is a mixed-strategy Nash equilibrium, for any $a_i \in A_i$, we have

$$u_i(a^*) = U_i(1 \circ a^*) \geq U_i(1 \circ a_i, 1 \circ a_{-i}^*) = u_i(a_i, a_{-i}^*).$$

□

5.6 Proposition: A profile α^* is a mixed-strategy Nash equilibrium if and only if for each player i ,

$$U_i(\alpha^*) \geq U_i(1 \circ a_i, \alpha_{-i}^*) \text{ for all } a_i \in A_i.$$

5.7 Some games, e.g., matching pennies, may possess no pure-strategy Nash equilibrium. However, every finite strategic game must have at least one mixed-strategy Nash equilibrium.

Nash's theorem: Every finite strategic game $\langle N, (A_i), (u_i) \rangle$ has a mixed-strategy Nash equilibrium.

Proof. Consider the mixed extension $\langle N, (\Delta(A_i)), (U_i) \rangle$ of the strategic game. Since expected payoff is linear in probabilities, U_i is both continuous and quasi-concave on $\Delta(A_i)$. Since $\Delta(A_i)$ is a simplex in a finite-dimensional Euclidean space, there is a mixed-strategy Nash equilibrium. □

5.8 For the finite set X and $\delta \in \Delta(X)$, we denote $\delta(x)$ the probability that δ assigns to $x \in X$. Define the support of δ to be the set of elements $x \in X$ for which $\delta(x) > 0$, i.e.,

$$\text{support}(\delta) = \{x \in X \mid \delta(x) > 0\}.$$

5.9 Lemma: A profile α^* is a mixed-strategy Nash equilibrium if and only if for each i and for each $a_i^* \in \text{support}(\alpha_i^*)$,

$$U_i(1 \circ a_i^*, \alpha_{-i}^*) \geq U_i(1 \circ a_i, \alpha_{-i}^*) \text{ for all } a_i \in A_i.$$

Proof. “ \Rightarrow ”: If there is $a_i^* \in \text{support}(\alpha_i^*)$ which is not a best response to α_{-i}^* , then i can increase his payoff by transferring probability from a_i^* to a best-response action; hence α_i^* is not an equilibrium strategy.

“ \Leftarrow ”: If each $a_i^* \in \text{support}(\alpha_i^*)$ is a best response to α_{-i}^* , then i can not do better by choosing a different mixed strategy α_i ; hence α_i^* is optimal against α_{-i}^* . □

5.10 Corollary: Every action in the support of any player's equilibrium strategy yields that player the same equilibrium payoff, i.e.,

$$U_i(a_i^*, \alpha_{-i}^*) = U_i(\alpha^*) \text{ for all } a_i^* \in \text{support}(\alpha_i^*).$$

5.11 A profile α^* is a mixed-strategy Nash equilibrium if and only if

- (1) for every player i , no action in A_i yields, given α_{-i}^* , a payoff to player i that exceeds his equilibrium payoff, and
- (2) the set of actions that yield, given α_{-i}^* , a payoff less than his equilibrium payoff has α_i^* -measure zero.

5.2 Examples

5.12 Example: Matching pennies.

		Player 2	
		Head	Tail
Player 1	Head	1, -1	-1, 1
	Tail	-1, 1	1, -1

Figure 5.1: Matching pennies

Answer. Let $p_1 = (r, 1 - r)$ be a mixed strategy in which player 1 plays Head with probability r . Let $p_2 = (q, 1 - q)$ be a mixed strategy for player 2. Then given p_2 , we have

$$\begin{aligned} U_1(\text{Head}, p_2) &= q \cdot 1 + (1 - q) \cdot (-1) = 2q - 1, \\ U_1(\text{Tail}, p_2) &= q \cdot (-1) + (1 - q) \cdot 1 = 1 - 2q. \end{aligned}$$

Player 1 chooses Head if and only if $U_1(\text{Head}, p_2) \geq U_1(\text{Tail}, p_2)$ if and only if $q \geq \frac{1}{2}$. Hence

$$B_1(q) = \begin{cases} \{1\}, & \text{if } 1 \geq q > \frac{1}{2}, \\ \{0\}, & \text{if } \frac{1}{2} > q \geq 0, \\ [0, 1], & \text{if } q = \frac{1}{2}. \end{cases}$$

Similarly, we have

$$\begin{aligned} U_2(p_1, \text{Head}) &= r \cdot (-1) + (1 - r) \cdot 1, \\ U_2(p_1, \text{Tail}) &= r \cdot 1 + (1 - r) \cdot (-1), \end{aligned}$$

and

$$B_2(r) = \begin{cases} \{0\}, & \text{if } 1 \geq r > \frac{1}{2}, \\ \{1\}, & \text{if } \frac{1}{2} > r \geq 0, \\ [0, 1], & \text{if } r = \frac{1}{2}. \end{cases}$$

We draw the graphs of $B_1(q)$ and $B_2(r)$ together:

The graphs of the best-response correspondences intersect at only one point $(\frac{1}{2}, \frac{1}{2})$, and in this case $r = q = \frac{1}{2}$. Thus (p_1^*, p_2^*) is the unique mixed-strategy Nash equilibrium, where $p_1^* = (\frac{1}{2}, \frac{1}{2})$ and $p_2^* = (\frac{1}{2}, \frac{1}{2})$. \square

5.13 We may also use Corollary 5.10 to find mixed-strategy Nash equilibrium. Suppose that (p_1^*, p_2^*) is a mixed-strategy Nash equilibrium, then

$$\begin{cases} p_2^*(\text{Head}) \cdot 1 + p_2^*(\text{Tail}) \cdot (-1) = p_2^*(\text{Head}) \cdot (-1) + p_2^*(\text{Tail}) \cdot 1 \\ p_1^*(\text{Head}) \cdot (-1) + p_1^*(\text{Tail}) \cdot 1 = p_1^*(\text{Head}) \cdot 1 + p_1^*(\text{Tail}) \cdot (-1) \end{cases}$$

Thus, for $i = 1, 2$, $p_i^*(\text{Head}) = p_i^*(\text{Tail}) = \frac{1}{2}$.

5.14 Example [OR Example 34.1]: Battle of sexes.




Figure 5.2: Matching pennies

		Peter	
		Opera	Fight
Mary	Opera	2, 1	0, 0
	Fight	0, 0	1, 2

Figure 5.3: Battle of the sexes.

Answer. Suppose that (α_1, α_2) is a mixed-strategy Nash equilibrium.

If $\alpha_1(O)$ is zero or one, we obtain the two pure-strategy Nash equilibria.

If $0 < \alpha_1(O) < 1$ then, given α_2 , by Corollary 5.10 player 1's actions "Opera" and "Fight" must yield the same payoff, so that we must have $2\alpha_2(O) = \alpha_2(F)$ and thus $\alpha_2(O) = \frac{1}{3}$.

Since $0 < \alpha_2(O) < 1$ it follows from the same result that player 2's actions "Opera" and "Fight" must yield the same payoff, so that $\alpha_1(O) = 2\alpha_1(F)$, or $\alpha_1(O) = \frac{2}{3}$.

Thus the only non-degenerate mixed-strategy Nash equilibrium of the game is $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$. \square

5.15 Example: Rock-paper-scissors game.

Suppose two players play the rock-paper-scissors game. The two players simultaneously form one of their hands into the shape of either a rock, a piece of paper, or a pair of scissors. Abbreviate these shapes as R , P , and S , respectively. If the players pick the same shape, then the game ends in a tie. Otherwise, one of the players wins and the other loses. The winner is determined by the following rule: rock beats scissors, scissors beat paper, and paper beats rock. Each player obtains a payoff of 1 as a winner, -1 as a loser, and 0 if he or she ties. Find all Nash equilibria.

Answer. Unique mixed-strategy Nash equilibrium: $(\frac{1}{3}R + \frac{1}{3}P + \frac{1}{3}S, \frac{1}{3}R + \frac{1}{3}P + \frac{1}{3}S)$. \square

5.16 Example [OR Exercise 35.1]: Guessing the average.

Let n ($n \geq 2$) people play the following game. Simultaneously, each player i announces a number x_i in the set $\{1, 2, \dots, K\}$. A prize of \$1 is split equally between all the people whose number is closest to $\frac{2}{3} \cdot \frac{x_1 + \dots + x_n}{n}$. Show that the game has a unique mixed-strategy Nash equilibrium, in which each player's strategy is pure.

Answer. \square

5.17 Example [OR Exercise 35.2]: An investment race.

Two investors are involved in a competition with a prize of one dollar. Each investor can spend any amount in the interval $[0, 1]$. The winner is the investor who spends the most; in the event of a tie each investor receives half dollar. Formulate this situation as a strategic game and find its mixed-strategy Nash equilibria.

Answer.

□

5.18 Example [OR Exercise 36.1]: Guessing right.

Players 1 and 2 each choose a member of the set $\{1, 2, \dots, K\}$. If the players choose the same number then player 2 pay one dollar to player 1; otherwise no payment is made. Each player maximizes his expected monetary payoff. Find the mixed-strategy Nash equilibria of this game.

Answer.

□

5.19 Example [OR Exercise 36.2]: Air strike.

Army A has a single plane with which it can strike one of three possible targets. Army B has one anti-aircraft gun that can be assigned to one of the targets. The value of target k is v_k , with $v_1 > v_2 > v_3 > 0$. Army A can destroy a target only if the target is undefended and A attacks it. Army A wishes to maximize the expected value of the damage and army B wishes to minimize it. Formulate the situation as a (strictly competitive) strategic game and find its mixed-strategy Nash equilibria.

Answer.

□

5.20 Example: Symmetric games.

A game is symmetric if each player has the same set of pure strategies and

$$u_{\sigma(i)}(s_{\sigma(1)}, s_{\sigma(2)}, \dots, s_{\sigma(n)}) = u_i(s_1, s_2, \dots, s_n)$$

for each player i whenever the n -vector $(\sigma(1), \sigma(2), \dots, \sigma(n))$ is a permutation of $(1, 2, \dots, n)$. Prove that a finite symmetric game possesses a symmetric (mixed-strategy) Nash equilibrium—a Nash equilibrium in which every player chooses the same strategy.

Answer.

□

5.21 Example (Complete-information all-pay auction): Two players, indexed by $i = 1, 2$, compete for an object of value $v > 0$. Players simultaneously place their non-negative bids x_i . The higher bidder obtains the object, while one loses his bid regardless of win or loss. Ties will be broken randomly. Their payoffs can then be written as

$$\pi_i(x_i, x_j) = \begin{cases} v - x_i, & \text{if } x_i > x_j, \\ \frac{v}{2} - x_i, & \text{if } x_i = x_j, \\ -x_i, & \text{if } x_i < x_j. \end{cases}$$

Find the unique mixed-strategy Nash equilibrium.

5.22 Example: Contest.

n agents contest for a prize. Let agent i outlay x_i to influence the outcome of the political contest in his favor. The probability that agent i will be the successful contender is

$$p_i(x) = \begin{cases} 0, & \text{if } x_i \text{ is not a maximal element of } x_1, x_2, \dots, x_n, \\ \frac{1}{m}, & \text{if } x_i \text{ is one of } m \text{ maximal elements of } x_1, x_2, \dots, x_n. \end{cases}$$

For each i , assume agent i 's valuation is v_i , which is publicly known.

Show that

- (i) No agent will, in equilibrium, spend a positive amount β with a strictly positive probability.
- (ii) If there are only two agents, they must have the same maximum spending level.
- (iii) The minimum spending level is zero for each agent.
- (iv) At most one agent spends zero with strictly positive probability.

Given these results, if we define $1 - G_i(x_i)$ to be the probability that agent i spends more than x_i , then $G_i(x_i)$ is continuous over $(0, \infty)$.

Proof. (i) (1) Suppose that agent i does spend β with strictly positive probability.

(2) Then the probability that a rival agent j beats agent i rises discontinuously as a function of x_j at $x_j = \beta$.

(3) Therefore, there is some $\epsilon > 0$, such that agent j will bid on the interval $[\beta - \epsilon, \beta]$ with zero probability, for all $j \neq i$.

(4) But then agent i is better off spending $\beta - \epsilon$ rather than β since his probability of winning is the same, contradicting the hypothesis that $x_j = \beta$ is an equilibrium strategy.

(ii) Routine.

(iii) (1) Suppose that agent i spends less than β with zero probability, where $\beta > 0$.

(2) Then for any $j \neq i$, any spending level between zero and β yields a negative payoff since the probability of winning is zero.

(3) Since other agents can always spend zero it follows that no agent will spend in the interval $(0, \beta)$.

(4) But then agent i could reduce his spending level below β without altering the probability of winning, contradicting the hypothesis that agent i could, in equilibrium, do no better than take β as his minimal spending level.

(iv) Routine.

□

5.3 Interpretation of mixed-strategy Nash equilibrium

5.23 “Outside of Las Vegas we do not spin roulettes” by Ariel Rubinstein (*Econometrica* 1991).

5.24 A mixed strategy entails a deliberate decision by a player to introduce randomness into his behavior.

5.25 We can interpret a mixed-strategy Nash equilibrium as a stochastic steady state. The players have information about the frequencies with which actions were taken in the past; each player uses these frequencies to form his belief about the future behavior of the other players, and hence formulate his action.

5.26 Mixed strategies as pure strategies in an extended game.

Example:

		Peter	
		Opera	Fight
Mary	Opera	2, 1	0, 0
	Fight	0, 0	1, 2

This game has two pure-strategy Nash equilibria (Opera, Opera) and (Fight, Fight), and one mixed-strategy Nash equilibrium ($\frac{2}{3} \circ \text{Opera} + \frac{1}{3} \circ \text{Fight}$, $\frac{1}{3} \circ \text{Opera} + \frac{2}{3} \circ \text{Fight}$).

Now suppose that each player has three possible “moods”, determined by factors he does not understand. Each player is in each of these moods one-third of the time, independently of the other player’s mood; his mood has no effect on his payoff.

Assume that Mary chooses Opera whenever she is in moods 1 or 2 and Fight when she is in mood 3, and Peter chooses Opera when he is in mood 1 and Fight when he is in moods 2 or 3.

Viewing the situation as a Bayesian game in which the three types of each player correspond to his possible moods, this behavior defines a pure-strategy Nash equilibrium corresponding exactly to the mixed-strategy Nash equilibrium of the original game.

Note that the mood (signal) is private and independent.

5.27 Mixed strategies as beliefs.

5.3.1 Purification

5.28 Mixed strategies as pure strategies in a perturbed game, due to [Harsanyi \(1973\)](#).

A mixed-strategy Nash equilibrium in a game of complete information can almost always be interpreted as a pure-strategy Bayesian Nash equilibrium in a closely related game with a little bit of incomplete information.




Figure 5.4

5.29 Example: Battle of sexes.

		Peter	
		Opera	Fight
Mary	Opera	2, 1	0, 0
	Fight	0, 0	1, 2

This game has two pure-strategy Nash equilibria (Opera, Opera) and (Fight, Fight), and one mixed-strategy Nash equilibrium ($\frac{2}{3} \circ \text{Opera} + \frac{1}{3} \circ \text{Fight}$, $\frac{1}{3} \circ \text{Opera} + \frac{2}{3} \circ \text{Fight}$).

- (1) Suppose Mary and Peter are not completely sure each other's payoff: if both attend Opera, Mary's payoff is $2 + t_m$; if both attend Fight, Peter's payoff is $2 + t_p$, where t_m is privately known by Mary and t_p is privately known by Peter, and t_m and t_p are independently drawn from a uniform distribution on $[0, x]$.

This can be expressed as a Bayesian game $G = \langle A_m, A_p; T_m, T_p; P_m, P_p; u_m, u_p \rangle$, where

- $A_m = A_p = \{\text{Opera, Fight}\}$
- $T_m = T_p = [0, x]$
- The payoffs are

		Peter	
		Opera	Fight
Mary	Opera	$2 + t_m, 1$	$0, 0$
	Fight	$0, 0$	$1, 2 + t_p$

- (2) In general, Mary's and Peter's strategies are $s_m: [0, x] \rightarrow \{\text{Opera, Fight}\}$ and $s_p: [0, x] \rightarrow \{\text{Opera, Fight}\}$, which are defined by

$$s_m(t_m) = \begin{cases} \text{Opera,} & \text{if } t_m \geq m, \\ \text{Fight,} & \text{if } t_m < m, \end{cases} \quad s_p(t_p) = \begin{cases} \text{Fight,} & \text{if } t_p \geq p, \\ \text{Opera,} & \text{if } t_p < p. \end{cases}$$

Given $s_m(t_m)$ and $s_p(t_p)$, let

$$\delta_m = \text{Prob}(\text{Mary plays Opera}) = \frac{x - m}{x}, \quad \delta_p = \text{Prob}(\text{Peter plays Fight}) = \frac{x - p}{x}.$$

As $x \rightarrow 0$, since t_m and t_p are in $[0, x]$, the Bayesian game converges to the original game of complete information.

- (3) Given Peter's strategy $s_p(t_p)$ (or say p), Mary's expected payoff, if she plays Opera, is

$$(2 + t_m) \cdot \text{Prob}(\text{Peter chooses Opera}) + 0 \cdot \text{Prob}(\text{Peter chooses Fight}) = (2 + t_m)(1 - \delta_p) + 0 \cdot \delta_p,$$

and, if she plays Fight, is

$$0 \cdot (1 - \delta_p) + 1 \cdot \delta_p.$$

Thus Mary playing Opera is optimal if and only if

$$(2 + t_m)(1 - \delta_p) \geq \delta_p \Leftrightarrow t_m \geq \frac{\delta_p}{1 - \delta_p} - 2 = \frac{x}{p} - 3.$$

Let $m = \frac{x}{p} - 3$. Then $s_m(t_m)$ is the best response strategy to $s_p(t_p)$.

- (4) Similarly, given Mary's strategy $s_m(t_m)$ (or say m), Peter's expected payoff is $0 \cdot \delta_m + (2 + t_p)(1 - \delta_m)$ for Fight and $1 \cdot \delta_m + 0 \cdot (1 - \delta_m)$ for Opera.

Thus playing Fight is optimal if and only if

$$(2 + t_p)(1 - \delta_m) \geq \delta_m \Leftrightarrow t_p \geq \frac{\delta_m}{1 - \delta_m} - 2 = \frac{x}{m} - 3.$$

Let $p = \frac{x}{m} - 3$. Then, $s_p(t_p)$ is the best response strategy to $s_m(t_m)$.

(5) Hence, (s_m^*, s_p^*) is a Bayesian Nash equilibrium if and only if

$$\frac{x}{p^*} - 3 = m^*, \quad \frac{x}{m^*} - 3 = p^*.$$

Thus,

$$p^* = m^* = \frac{\sqrt{9+4x} - 3}{2}.$$

As $x \rightarrow 0$, the Bayesian game converges to the original game of complete information, and

$$\delta_m^* = \delta_p^* = \frac{x - p^*}{x} = 1 - \frac{\sqrt{9+4x} - 3}{2x} \rightarrow \frac{2}{3}.$$

5.30 Example: Matching pennies.

It has no pure-strategy Nash equilibrium but has one mixed-strategy Nash equilibrium: each player plays H with probability $1/2$.

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

Game G

(1) Consider the following game with incomplete information $G(x)$: where

	H	T
H	$1 + t_1, -1$	$-1, 1 - t_2$
T	$-1, 1$	$1, -1$

Game $G(x)$

- Type spaces: $T_1 = T_2 = [0, x]$, t_1 and t_2 are i.i.d. random variables and uniformly distributed on $[0, x]$.
- Action spaces: $A_1 = A_2 = \{H, T\}$.
- Strategy spaces: $S_1 = S_2 = \{s_i \text{ is a function from } [0, x] \text{ to } \{H, T\}\}$.

Note that $G(0) = G$.

In $G(x)$, suppose (s_1^*, s_2^*) is a Bayesian Nash equilibrium, $p = \text{Prob}(\{t_1 : s_1^*(t_1) = H\})$, and $q = \text{Prob}(\{t_2 : s_2^*(t_2) = H\})$.

(2) For player 1, given his type t_1 and player 2's strategy s_2^* , his expected payoff is

$$\mathbf{E}[u_1(a_1, s_2^*) \mid t_1] = \begin{cases} (1 + t_1) \cdot q - 1 \cdot (1 - q), & a_1 = H; \\ -1 \cdot q + 1 \cdot (1 - q), & a_1 = T. \end{cases}$$

Thus H is a best response if and only if $(1 + t_1) \cdot q - 1 \cdot (1 - q) \geq -1 \cdot q + 1 \cdot (1 - q)$, that is, $t_1 \geq \frac{2}{q} - 4$.

Hence, we have

$$p = \text{Prob}(\{t_1 : s_1^*(t_1) = H\}) = 1 - \frac{2/q - 4}{x} \quad (5.1)$$

(3) For player 2, given his type t_2 and player 1's strategy s_1^* , his expected payoff is

$$\mathbb{E}[u_2(a_2, s_1^*) \mid t_2] = \begin{cases} -1 \cdot p + 1 \cdot (1 - p), & a_2 = H; \\ (1 - t_2) \cdot p + (-1) \cdot (1 - p), & a_2 = T. \end{cases}$$

Thus H is a best response if and only if $-1 \cdot p + 1 \cdot (1 - p) \geq (1 - t_2) \cdot p + (-1) \cdot (1 - p)$, that is, $t_2 \geq 4 - \frac{2}{p}$.

Hence, we have

$$q = \text{Prob}(\{t_2 : s_2^*(t_2) = H\}) = 1 - \frac{4 - 2/p}{x} \quad (5.2)$$

(4) Rewriting Equations (5.1) and (5.2), we will have

$$p = \frac{2}{4 + (q - 1)x}, \quad q = \frac{2}{4 + (1 - p)x}.$$

As $x \rightarrow 0$, $p, q \rightarrow \frac{1}{2}$, that is, the Bayesian Nash equilibrium will converge to the mixed-strategy Nash equilibrium in G .

5.31 Example: Matching pennies (cont.).

(1) Consider the following game with incomplete information $G(x)$: where

	H	T
H	$1 + t_1, -1$	$-1, 1 + t_2$
T	$-1, 1$	$1, -1$

Game $G(x)$

- Type spaces: $T_1 = T_2 = [0, x]$, t_1 and t_2 are i.i.d. random variables and uniformly distributed on $[0, x]$.
- Action spaces: $A_1 = A_2 = \{H, T\}$.
- Strategy spaces: $S_1 = S_2 = \{s_i \text{ is a function from } [0, x] \text{ to } \{H, T\}\}$.

Note that $G(0) = G$.

In $G(x)$, suppose (s_1^*, s_2^*) is a Bayesian Nash equilibrium, $p = \text{Prob}(\{t_1 : s_1^*(t_1) = H\})$, and $q = \text{Prob}(\{t_2 : s_2^*(t_2) = H\})$.

(2) For player 1, given his type t_1 and player 2's strategy s_2^* , his expected payoff is

$$\mathbb{E}[u_1(a_1, s_2^*) \mid t_1] = \begin{cases} (1 + t_1) \cdot q - 1 \cdot (1 - q), & a_1 = H; \\ -1 \cdot q + 1 \cdot (1 - q), & a_1 = T. \end{cases}$$

Thus H is a best response if and only if $(1 + t_1) \cdot q - 1 \cdot (1 - q) \geq -1 \cdot q + 1 \cdot (1 - q)$, that is, $t_1 \geq \frac{2}{q} - 4$.

Hence, we have

$$p = \text{Prob}(\{t_1 : s_1^*(t_1) = H\}) = 1 - \frac{2/q - 4}{x} \quad (5.3)$$

(3) For player 2, given his type t_2 and player 1's strategy s_1^* , his expected payoff is

$$\mathbb{E}[u_2(a_2, s_1^*) \mid t_2] = \begin{cases} -1 \cdot p + 1 \cdot (1 - p), & a_2 = H; \\ (1 + t_2) \cdot p + (-1) \cdot (1 - p), & a_2 = T. \end{cases}$$

Thus H is a best response if and only if $-1 \cdot p + 1 \cdot (1 - p) \geq (1 + t_2) \cdot p + (-1) \cdot (1 - p)$, that is, $t_2 \leq \frac{2}{p} - 4$. Hence, we have

$$q = \text{Prob}(\{t_2 : s_2^*(t_2) = H\}) = \frac{2/p - 4}{x} \quad (5.4)$$

(4) Rewriting Equations (5.3) and (5.4), we will have

$$p = \frac{2}{4 + qx}, \quad q = \frac{2}{4 + (1 - p)x}.$$

As $x \rightarrow 0$, $p, q \rightarrow \frac{1}{2}$, that is, the Bayesian Nash equilibrium will converge to the mixed-strategy Nash equilibrium in G .

5.32 Harsanyi shows that any mixed strategy equilibrium can be “purified” in a similar way. For a complete proof, see [Govindan, Reny and Robson \(2003\)](#).

Correlated equilibrium

Contents

6.1	Motivation	109
6.2	Correlated equilibrium	110
6.3	Examples	113

6.1 Motivation

- 6.1 We discuss an interpretation of a mixed-strategy Nash equilibrium as a steady state in which each player's action depends on a signal that she receives from "nature". In this interpretation the signals are private and independent.
- 6.2 What happens if the signals are not private and independent?

Example:

		Peter	
		Opera	Fight
Mary	Opera	2, 1	0, 0
	Fight	0, 0	1, 2

		Opera	Fight
		$p(x) = \frac{1}{2}$	0
Opera			
Fight		0	$p(y) = \frac{1}{2}$

Suppose that both players observe a random variable that takes each of the two values x and y with probability $\frac{1}{2}$. Then there is a new equilibrium, in which both players choose Opera if the realization is x and Fight if the realization is y .

Given each player's information, her action is optimal: if the realization is x then she knows that the other player chooses Opera, so that it is optimal for him to choose Opera, and symmetrically if the realization is y .

One interpretation of this equilibrium is that the players observe the outcome of a public coin toss, which determines which of the two pure-strategy Nash equilibria they play.

在按照 p 随机抽取行动组合推荐给参与人时，每个参与人在接收到行动推荐后，给定根据 p 计算出的对方收到的行动推荐的条件概率并假定对方会服从，那么他的最优选择就是服从推荐。

- 6.3 In this example the players observe the same random variable. More generally, their information may be less than perfectly correlated.

		Player 2			
		<i>L</i>	<i>R</i>		
Player 1	<i>T</i>	6, 6	2, 7	<i>T</i>	$p(y) = \frac{1}{3}$
	<i>B</i>	7, 2	0, 0	<i>B</i>	$p(x) = \frac{1}{3}$
					$p(z) = \frac{1}{3}$
					0

Suppose that there is a random variable that takes the three values x, y , and z equally likely, and player 1 knows only that the realization is either x or that it is a member of $\{y, z\}$, while player 2 knows only that it is either a member of $\{x, y\}$ or that it is z . That is, player 1's information partition is $\{\{x\}, \{y, z\}\}$ and player 2's is $\{\{x, y\}, \{z\}\}$.

A strategy of player 1 consists of two actions: one that she uses when she knows that the realization is x and one that she uses when she knows that the realization is a member of $\{y, z\}$. Similarly, a strategy of player 2 consists of two actions, one for $\{x, y\}$ and one for z .

A player's strategy is optimal if, given the strategy of the other player, for any realization of her information she can do no better by choosing an action different from that dictated by her strategy.

Then there is a new equilibrium, they will choose (B, L) , (T, L) and (T, R) when x, y and z are realized respectively. The equilibrium payoff profile are $(5, 5)$.

Neither player has an incentive to deviate. Consider player 1. At state x , player 1 knows that player 2 plays L and thus it is optimal for player 1 to play B ; at states y and z , player 1 assigns equal probabilities to player 2 playing L and R , so that it is optimal for player 1 to play T .

6.4 Another advantage: The game above has two pure-strategy Nash equilibria (T, R) , (B, L) and one mixed-strategy Nash equilibrium $(\frac{2}{3} \circ T + \frac{1}{3} \circ B, \frac{2}{3} \circ L + \frac{1}{3} \circ R)$.

The expected payoff for the new equilibrium is $7 \cdot \frac{1}{3} + 2 \cdot \frac{1}{3} + 6 \cdot \frac{1}{3} = 5$ which is higher than the expected payoff of the mixed-strategy Nash equilibrium.




Figure 6.1

6.2 Correlated equilibrium

6.5 Definition: A correlated equilibrium, denoted by $\langle (\Omega, \pi), (\mathcal{P}_i), (\sigma_i) \rangle$, of a strategic game $\langle N, (A_i), (u_i) \rangle$ consists of

- a finite probability space (Ω, π) (Ω is a set of states and π is a probability measure on Ω)

- for each player i a partition \mathcal{P}_i of Ω (player i 's information partition)
- for each player i a function $\sigma_i: \Omega \rightarrow A_i$ with $\sigma_i(\omega) = \sigma_i(\omega')$ whenever $\omega, \omega' \in P_i$ for some $P_i \in \mathcal{P}_i$ (σ_i is player i 's strategy)

such that for every i and every function $\tau_i: \Omega \rightarrow A_i$ for which $\tau_i(\omega) = \tau_i(\omega')$ whenever $\omega, \omega' \in P_i$ for some $P_i \in \mathcal{P}_i$ (i.e. for every strategy of player i) we have

$$\sum_{\omega \in \Omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \sigma_i(\omega)) \geq \sum_{\omega \in \Omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \tau_i(\omega)).$$

Intuitively, correlated equilibrium allows us to get at preplay communication without explicit modeling the communication process.

- 6.6 The condition of correlated equilibrium can be written as: for every player i , every information cell $P_i \in \mathcal{P}_i$ with $\pi(P_i) > 0$, every $a_i \in A_i$,

$$\sum_{\omega \in P_i} \pi(\omega | P_i) \cdot u_i(\sigma(\omega)) \geq \sum_{\omega \in P_i} \pi(\omega | P_i) \cdot u_i(\sigma_{-i}(\omega), a_i).$$

- 6.7 Proposition: Let $G = \langle N, (A_i), (u_i) \rangle$ be a finite strategic game. Every probability distribution over outcomes that can be obtained in a correlated equilibrium of G can be obtained in a correlated equilibrium in which the set of states is A and for each i player i 's information partition consists of all sets of the form $\{a \in A \mid a_i = b_i\}$ for some action $b_i \in A_i$.

- 6.8 Proof. Let $\langle (\Omega, \pi), (\mathcal{P}_i), (\sigma_i) \rangle$ be a correlated equilibrium of G . Then $\langle (\Omega', \pi'), (\mathcal{P}'_i), (\sigma'_i) \rangle$ is also a correlated equilibrium, where $\Omega' = A$, $\pi'(a) = \pi(\{\omega \in \Omega \mid \sigma(\omega) = a\})$ for each $a \in A$, \mathcal{P}'_i consists of sets of the type $\{a \in A \mid a_i = b_i\}$ for some $b_i \in A_i$, and σ'_i is defined by $\sigma'_i(a) = a_i$.

For every i and every function $\tau'_i: A \rightarrow A_i$ for which $\tau'_i(a) = \tau'_i(a')$ whenever $a_i = a'_i$, we have

$$\begin{aligned} \sum_a \pi'(a) \cdot u_i(\sigma'_{-i}(a), \sigma'_i(a)) &= \sum_a \pi(\{\omega \in \Omega \mid \sigma(\omega) = a\}) \cdot u_i(a_{-i}, a_i) = \sum_a \sum_{\omega \in \sigma^{-1}(a)} \pi(\omega) \cdot u_i(a_{-i}, a_i) \\ &= \sum_{\omega} \pi(\omega) \cdot u_i(\sigma(\omega)) \geq \sum_{\omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \tau'_i(\sigma(\omega))) \\ &= \sum_a \sum_{\omega \in \sigma^{-1}(a)} \pi(\omega) \cdot u_i(a_{-i}, \tau'_i(a)) = \sum_a \pi'(a) \cdot u_i(\sigma'_{-i}(a), \tau'_i(a)). \end{aligned}$$

□

- 6.9 Remark: This result allows us to confine attention, when calculating correlated equilibrium payoffs, to equilibria in which the set of states is the set of outcomes. Note however that such equilibria may have no natural interpretation.

- 6.10 Corollary: A correlated equilibrium can be viewed simply as $\pi \in \Delta(A)$ such that, for every player i , and every function $\gamma_i: A_i \rightarrow A_i$,

$$\sum_{a \in A} \pi(a) \cdot u_i(a) \geq \sum_{a \in A} \pi(a) \cdot u_i(a_{-i}, \gamma_i(a_i)).$$

Thus, a mixed-strategy Nash equilibrium is a special form of correlated equilibrium: $\pi \in \times_{j \in N} \Delta(A_j)$.

- 6.11 (From <http://zhuanlan.zhihu.com/qianfeng/20018642>) 相关均衡是基于推荐的均衡，可以通过交通调度来理解。调度员只需要分别告诉每个司机他该走哪条路，司机收到调度指令后，根据调度分布算出

调度员给其他人的指令的概率分布是什么（从而计算拥堵程度），如果该分布是均衡，在给定其他司机都会服从调度的情况下，每个司机都会发现服从调度是最优的。

从直观上，我们已经可以看出相关均衡和纳什均衡思路上的 consistency：只要其他人服从均衡指定的推荐，那么自己服从推荐就是最优的。相关均衡更一般的地方在于，每个参与人自己得到的推荐和其他人的推荐可以是相关的，而在纳什均衡里，无论你选择什么，对方的（混合）策略都是独立于你的从而是给定的。

给定一个博弈和它的一个相关均衡 π ，由于 π 本质上是一个 **correlating device**， π 给每个参与人推荐行动的过程可以理解成一个不完全信息的生成过程：每个参与人接收到自己的信息（行动推荐），同时产生了对他人信息（行动推荐）的概率估计。

在该博弈由相关均衡 π 生成的不完全信息博弈中，每个参与人在每个信息状态下选择推荐的行动构成贝叶斯纳什均衡。

也就是说，如果我们在模拟一个博弈的时候，有一部分与博弈支付无关而在参与人之间又相关的信息没有考虑到，从而把参与人事实上面临的不完全信息博弈理解成了完全信息博弈（想象你坐着直升机在天空观察并试图理解司机之间的博弈，却没有意识到司机们背后调度的存在），那么参与人之间博弈的贝叶斯纳什均衡事实上是我们考虑的完全信息博弈的相关均衡。

6.12 Proposition: Let $G = \langle N, (A_i), (u_i) \rangle$ be a strategic game. Any convex combination of correlated equilibrium payoff profiles of G is a correlated equilibrium payoff profile of G .

6.13 Proof. Let u^1, \dots, u^K be correlated equilibrium payoff profiles and let $(\lambda^1, \dots, \lambda^K) \in \mathbb{R}^K$ with $\lambda^k \geq 0$ for all k and $\sum_k \lambda^k = 1$. For each value of k let $\langle (\Omega^k, \pi^k), (\mathcal{P}_i^k), (\sigma_i^k) \rangle$ be a correlated equilibrium that generates the payoff profile u^k ; without loss of generality assume that the sets Ω^k are disjoint. The following defines a correlated equilibrium for which the payoff profile is $\sum_k \lambda^k u^k$. Let

- $\Omega = \cup_k \Omega^k$, and for any $\omega \in \Omega$ define $\pi(\omega) = \lambda^k \pi^k(\omega)$ where $\omega \in \Omega^k$;
- For each i , let $\mathcal{P}_i = \cup_k \mathcal{P}_i^k$, where $\omega \in \Omega^k$;
- Define σ_i by $\sigma_i(\omega) = \sigma_i^k(\omega)$ where $\omega \in \Omega^k$.

For every i and every function $\tau_i: \Omega \rightarrow A_i$ for which $\tau_i(\omega) = \tau_i(\omega')$ whenever $\omega, \omega' \in P_i$ for some $P_i \in \mathcal{P}_i$, let $\tau_i^k(\omega) = \tau_i(\omega)$ where $\omega \in \Omega^k$. Then

$$\begin{aligned} \sum_{\omega \in \Omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \sigma(\omega)) &= \sum_k \sum_{\omega \in \Omega^k} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \sigma_i^k(\omega)) = \sum_k \sum_{\omega \in \Omega^k} \lambda^k \pi^k(\omega) \cdot u_i(\sigma_{-i}^k(\omega), \sigma_i^k(\omega)) \\ &\geq \sum_k \sum_{\omega \in \Omega^k} \lambda^k \pi^k(\omega) \cdot u_i(\sigma_{-i}^k(\omega), \tau_i^k(\omega)) = \sum_{\omega \in \Omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \tau_i(\omega)). \end{aligned}$$

□

6.14 Remark: The set of Nash equilibrium outcomes is generally not convex.

6.15 One of the advantages of correlated equilibria is that they are computationally less expensive than are Nash equilibria. This can be captured by the fact that computing a correlated equilibrium only requires solving a linear program whereas solving a Nash equilibrium requires finding its fixed point completely. Another way of seeing this is that it is possible for two players to respond to each other's historical plays of a game and end up converging to a correlated equilibrium. (See Foster and Vohra, *GEB*, 1996)

6.3 Examples

6.16 Example [OR Exercise 48.1]: Consider the following three-player game. Player 1 chooses one of the two rows, player 2 chooses one of the two columns, and player 3 chooses one of the three tables.

	<i>L</i>	<i>R</i>
<i>T</i>	0, 0, 3	0, 0, 0
<i>B</i>	1, 0, 0	0, 0, 0

A

	<i>L</i>	<i>R</i>
<i>T</i>	2, 2, 2	0, 0, 0
<i>B</i>	0, 0, 0	2, 2, 2

B

	<i>L</i>	<i>R</i>
<i>T</i>	0, 0, 0	0, 0, 0
<i>B</i>	0, 1, 0	0, 0, 3

C

- (i) Show that the pure-strategy equilibrium payoffs are $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 0)$.
- (ii) Show that there is a correlated equilibrium in which player 3 chooses *B* and players 1 and 2 play (T, L) and (B, R) with equal probabilities.
- (iii) Explain the sense in which player 3 prefers not to have the information that players 1 and 2 use to coordinate their actions.

6.17 Example: Find the set of correlated equilibrium payoffs of battle of sexes.

		Peter	
		Opera	Fight
Mary	Opera	2, 1	0, 0
	Fight	0, 0	1, 2

Answer. We may assume

- $\Omega = \{(O, O), (O, F), (F, O), (F, F)\}$,
- $\mathcal{P}_1 = \{\{(O, O), (O, F)\}, \{(F, O), (F, F)\}\}$, and $\mathcal{P}_2 = \{\{(O, O), (F, O)\}, \{(O, F), (F, F)\}\}$,
- $\sigma_1(a_1, a_2) = a_1$ and $\sigma_2(a_1, a_2) = a_2$.

Let the probability distribution π be as follows:

$$x = \pi((O, O)), y = \pi((O, F)), z = \pi((F, O)), w = \pi((F, F)).$$

The prior probabilities of the four outcomes are summarized by the following table:

	<i>O</i>	<i>F</i>
<i>O</i>	x	y
<i>F</i>	z	w

By Corollary 6.10, we have the following inequalities:

$$\begin{aligned}
\cancel{xu_1(O, O)} + \cancel{yu_1(O, F)} + zu_1(F, O) + wu_1(F, F) &\geq \cancel{xu_1(O, O)} + \cancel{yu_1(O, F)} + zu_1(O, O) + wu_1(O, F), \\
xu_1(O, O) + yu_1(O, F) + \cancel{zu_1(F, O)} + \cancel{wu_1(F, F)} &\geq xu_1(F, O) + yu_1(F, F) + \cancel{zu_1(F, O)} + \cancel{wu_1(F, F)}, \\
\cancel{xu_2(O, O)} + yu_2(O, F) + \cancel{zu_2(F, O)} + wu_2(F, F) &\geq \cancel{xu_2(O, O)} + yu_2(O, O) + \cancel{zu_2(F, O)} + wu_2(F, O), \\
xu_2(O, O) + \cancel{yu_2(O, F)} + zu_2(F, O) + \cancel{wu_2(F, F)} &\geq xu_2(O, F) + \cancel{yu_2(O, F)} + zu_2(F, F) + \cancel{wu_2(F, F)}.
\end{aligned}$$

That is,

$$w \geq 2z, \quad 2x \geq y, \quad 2w \geq y, \quad x \geq 2z.$$

In other words, both x and w must be greater than $2z$ and $\frac{y}{2}$. The set of correlated equilibrium payoffs is equal to

$$u_1 = 2x + w, \quad u_2 = x + 2w,$$

subject to $\min\{x, w\} \geq \max\{\frac{y}{2}, 2z\}$, $x + y + z + w = 1$, and $x, y, z, w \geq 0$.

Draw the feasible payoff set. Notice that the efficient frontier is obviously part of the correlated equilibrium payoff set because any point on the frontier can be achieved by a linear combination of the two pure-strategy Nash equilibrium. The question is how inefficient the payoff can be. To make the payoff small, we want to put as much weight on y and z as possible without violating $\min\{x, w\} \geq \max\{\frac{y}{2}, 2z\}$. This means that we want to set $x = w = \frac{y}{2} = 2z$. Let $z = \epsilon$. Then $y = 4\epsilon$, $x = w = 2\epsilon$. This means $\epsilon = \frac{1}{9}$. The corresponding payoff is $(\frac{2}{3}, \frac{2}{3})$ which is the payoff for the mixed-strategy equilibrium. In this case, the set of correlated payoffs, as shown in Figure 6.2, is

the convex hull of $\left\{ \left(\frac{2}{3}, \frac{2}{3} \right), (2, 1), (1, 2) \right\}$.




Figure 6.2

□

6.18 Example: Find the set of correlated equilibrium payoffs of the following game.

	a	b
a	2, 2	0, 3
b	3, 0	-1, -1

Answer. We may assume

- $\Omega = \{(a, a), (a, b), (b, a), (b, b)\}$
- $\mathcal{P}_1 = \{\{(a, a), (a, b)\}, \{(b, a), (b, b)\}\}$ and $\mathcal{P}_2 = \{\{(a, a), (b, a)\}, \{(a, b), (b, b)\}\}$
- $\sigma_i(a_i, a_j) = a_i$

	a	b
a	x	y
b	z	w

Let $x = \pi((a, a))$, $y = \pi((a, b))$, $z = \pi((b, a))$, and $w = \pi((b, b))$. The prior probabilities of the four outcomes are summarized by the following table:

For the strategy to be a correlated equilibrium, x, y, z and w must satisfy the following conditions:

$$\begin{aligned} 2x &\geq 3x - y \Leftrightarrow x \leq y, \\ 2z &\leq 3z - w \Leftrightarrow w \leq z, \\ 2x &\geq 3x - z \Leftrightarrow x \leq z, \\ 2y &\leq 3y - w \Leftrightarrow w \leq y. \end{aligned}$$

The set of correlated equilibrium payoffs is equal to

$$\{(U_1, U_2) \mid U_1 = 2x + 3z - w, U_2 = 2x + 3y - w\},$$

where

$$\max\{x, w\} \leq \min\{y, z\}, \quad x, y, z, w \geq 0 \text{ and } x + y + z + w = 1.$$

Result:

$$\text{the convex hull of } \left\{ \left(\frac{2}{3}, \frac{2}{3} \right), \left(\frac{5}{3}, \frac{5}{3} \right), (3, 0), (0, 3) \right\}.$$

□

6.19 Question: How to compute the set of payoff vectors of correlated equilibria? (Hint: [Fourier-Motzkin elimination](#))

Rationalizability

Contents

7.1	Rationalizability	117
7.2	Iterated elimination of never-best response	121
7.3	Iterated elimination of strictly dominated actions	122
7.4	Examples	125
7.5	Iterated elimination of weakly dominated actions	128

In equilibrium each player’s choice is optimal given his correct belief about opponents’ actions. However, it is not clear how each player can know the other players’ equilibrium actions.

“Nash behavior is neither a necessary consequence of rationality, nor a reasonable empirical proposition.”

—Berheim (*Econometrica*, 1984, p. 1007)

For this reason, game theorists have developed solution concepts that do not entail this assumption.

We study some solution concepts, in which the players’ beliefs about each other’s actions are not assumed to be correct, but are constrained by considerations of rationality: each player believes that the actions taken by every other player is a best response to some belief, and, further, each player assumes that every other player reasons in this way and hence thinks that every other player believes that every other player’s action is a best response to some belief, and so on.

7.1 Rationalizability

7.1 The basic idea behind the notion of rationalizability is that “rational” behavior must be justified by “rational” beliefs and conversely, “rational” beliefs must be based on “rational” behavior.

7.2 The idea can be illustrated explicitly by a two-person game as follows:

- A strategy of i is called 1-justifiable if it is a best response to some beliefs of i about the strategic choice of j .
- A strategy of i is called t -justifiable (where $t \geq 2$) if it is a best response to a belief of i that assigns positive probabilities only to $(t - 1)$ -justifiable strategies of j .

- A strategy of i is rationalizable if it is a best response to a belief of i that assigns positive probabilities only to t -justifiable strategies of j for all $t \geq 1$.

7.3 Example:

		Player 2	
		L	R
Player 1	T	3, 0	0, 1
	M	0, 0	3, 1
	B	1, 1	1, 0

- (1) For player 1, T and M are 1-justifiable; for player 2, L and R are 1-justifiable.
- (2) For player 1, T and M are 2-justifiable; for player 2, R is 2-justifiable.
- (3) For player 1, M is 3-justifiable; for player 2, R is 3-justifiable.
- (4) It is clear that M and R are t -justifiable for players 1 and 2 respectively for $t > 3$.

7.4 The notion aims to be weak; it determines not what actions should actually be taken, but what actions can be ruled out with confidence.



7.5 Definition: For simplicity, we restrict attention to finite games unless otherwise stated explicitly.

- A product subset $X = \times_{j \in N} X_j \subseteq A$ is rationalizable if there exists a collection $\left((X_j^t)_{j \in N} \right)_{t=0}^{\infty}$ of sets with $X_j^0 = X_j$ and $X_j^t \subseteq A_j$ for all j and $t \geq 1$ such that for each $j \in N$, each $t \geq 0$, and each $a_j \in X_j^t$, there is a belief $\mu_j^{t+1}(a_j) \in \Delta(X_{-j}^{t+1})$ such that a_j is a best response (in A_j) to the belief $\mu_j^{t+1}(a_j)$ of player j .
- We call each $a_i \in X_i$ a rationalizable action for player i .

The interpretation of X_j^{t+1} is that it is the set of all actions a_j of player j that may be used to justify some other player i 's action $a_i \in X_i^t$.

7.6 We take a belief of player i to be a probability distribution on X_{-i}^t , i.e. $\mu \in \Delta(X_{-i}^t)$, which allows each player to believe opponents' actions are correlated. In the original definition in Bernheim (*Econometrica*, 1984) and Pearce (*Econometrica*, 1984), i 's belief is a product of independent probability distributions on X_{-i}^t , one for each of the other players, i.e. $\mu \in \times_{j \neq i} \Delta(X_j^t)$. In general, $\times_{j \neq i} \Delta(X_j^t) \subsetneq \Delta(\times_{j \neq i} X_j^t)$.

7.7 Example: In the following game, there are three players; player 1 chooses one of the two rows, player 2 chooses one of the two columns, and player 3 chooses one of the four tables. All three players obtain the same payoffs, given by the numbers in the boxes.

		L	R
U	D	8	0
	D	0	0
M_1			
		L	R
U	D	4	0
	D	0	4
M_2			
		L	R
U	D	0	0
	D	0	8
M_3			
		L	R
U	D	3	3
	D	3	3
M_4			

We claim that M_2 is rationalizable if player 3's belief about his opponents' actions are correlated, but is not rationalizable if he is restricted to beliefs that are products of independent probability distributions.

Let $Z_1 = \{U, D\}$, $Z_2 = \{L, R\}$ and $Z_3 = \{M_2\}$.

- U of player 1 is a best response to a belief that assigns probability one to (L, M_2) and D is a best response to the belief that assigns probability one to (R, M_2) ;

- L of player 2 is a best response to a belief that assigns probability one to (U, M_2) and R is a best response to the belief that assigns probability one to (D, M_2) ;
- M_2 of player 3 is a best response to the belief in which players 1 and 2 play (U, L) and (D, R) with equal probabilities.

However, M_2 is not a best response to any pair of independent mixed strategies and is thus not rationalizable under the modified definition in which each player's belief is restricted to be a product of independent beliefs.

In order for M_2 to be a best response, we need

$$4pq + 4(1-p)(1-q) \geq \max\{8pq, 8(1-p)(1-q), 3\},$$

where $(p, 1-p)$ and $(q, 1-q)$ are mixed strategies of players 1 and 2 respectively. This inequality is not satisfied for any values of p and q .

7.8 Shown below are the payoffs of player 1 in a three-person game. In this game, player 1 has three pure strategies L , M and R . Player 2 chooses rows and player 3 chooses either matrix A or matrix B . For players 2 and 3, neither strategy weakly dominates the other for any player. Is L a rationalizable strategy?

	L	M	R
U	6	10	0
D	6	10	10

A

	L	M	R
U	6	10	10
D	6	0	10

B

Answer. L is a best response to the belief $\frac{1}{2} \circ (U, A) + \frac{1}{2} \circ (D, B)$. Hence, L is rationalizable. \square

 7.9 Alternative definition:

- A product subset $Z = \times_{j \in N} Z_j \subseteq A$ is rationalizable if for each $j \in N$, each $a_j \in Z_j$, there exists $\mu_j(a_j) \in \Delta(Z_{-j})$ such that a_j is a best response (in A_j) to the belief $\mu_j(a_j)$.
- We call each $z_i \in Z_i$ a rationalizable action for player i .

7.10 *Proof of the equivalence.*

“ \Rightarrow ”: Suppose that $a_i \in A_i$ is rationalizable according to Definition 7.5. Then we have a product subset X which is rationalizable according to Definition 7.5 and $a_i \in X_i$. Let

$$Z = \times_{j \in N} (\cup_{t=0}^{\infty} X_j^t).$$

For each $j \in N$, each $a_j \in Z_j = \cup_t X_j^t$, there exists t , such that $a_j \in X_j^t$, and hence there is a belief $\mu_j^{t+1}(a_j) \in \Delta(X_{-j}^{t+1})$ to which a_j is a best response of player j . Since $\Delta(X_{-j}^{t+1}) \subseteq \Delta(Z_{-j})$, we have $\mu_j^{t+1}(a_j) \in \Delta(Z_{-j})$, and hence Z is rationalizable according to Definition 7.9. Therefore $a_i \in X_i \subseteq Z_i$ is rationalizable according to Definition 7.9.

“ \Leftarrow ”: If $a_i \in A_i$ is rationalizable according to Definition 7.9. Then we have a product subset Z which is rationalizable according to Definition 7.9 and $a_i \in Z_i$. Define $X_j^t = Z_j$ for each $j \in N$ and each $t \geq 0$; hence $X = Z$ is rationalizable according to Definition 7.5, and $a_i \in X_i$ is rationalizable according to Definition 7.5. \square

7.11 Proposition:

- a^* is a Nash equilibrium if and only if the singleton $\{a^*\}$ is a rationalizable set.

- Every action used with positive probability by some player in a correlated equilibrium is rationalizable.

Proof. Denote the strategic game by $\langle N, (A_i), (u_i) \rangle$; choose a correlated equilibrium $\langle (\Omega, \pi), (\mathcal{P}_i), (\sigma_i) \rangle$. For each j , each $a_j \in A_j$, take P_j to be an information cell such that on it $\sigma_j(\omega) = a_j$. Then for every $a'_j \in A_j$, we have

$$\sum_{\omega \in P_j} \pi(\omega | P_j) \cdot u_j(\sigma_{-j}(\omega), a_j) \geq \sum_{\omega \in P_j} \pi(\omega | P_j) \cdot u_j(\sigma_{-j}(\omega), a'_j).$$

For each player $j \in N$, let Z_j be the set of actions that player j uses with positive probability in the equilibrium, i.e., $Z_j = \text{support}(\sigma_j)$. For each $k \neq j$, based on $\pi(\omega | P_j)$ and $\sigma_k(\omega)$, define a player j 's belief on player k 's actions $\sum_{\omega \in P_j} \pi(\omega | P_j) \sigma_k(\omega)$.

Then any $a_j \in Z_j$ is a best response to the belief $\left(\sum_{\omega \in P_j} \pi(\omega | P_j) \sigma_k(\omega) \right)_{k \neq j}$. The support of this belief is a subset of Z_{-j} . \square

7.12 Define $B(R) = \times_{i \in N} B_i(R_{-i})$, where

$$B_i(R_{-i}) = \{a_i \mid \text{there exists } \mu \in \Delta(R_{-i}) \text{ such that } a_i \text{ is a best response to the belief } \mu\}.$$

Then Z is rationalizable if and only if $Z \subseteq B(Z)$, i.e., Z is a fixed point of the correspondence B .

7.13 If two subsets $R = \times_j R_j$ and $R' = \times_j R'_j$ of A are rationalizable, then $\times_j (R_j \cup R'_j)$ is also rationalizable.

For a finite strategic game $\langle N, (A_i), (u_i) \rangle$, define

$$R^* = \bigcup_{R \text{ is rationalizable}} R.$$

Then R^* is the largest (with respect to the set inclusion) rationalizable set.

7.14 Example: find the largest rationalizable set of the following game.

		Player 2			
		b_1	b_2	b_3	b_4
Player 1	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, -2	0, 0	10, -1

Answer. $R^* = \{a_1, a_2, a_3\} \times \{b_1, b_2, b_3\}$:

- (a_2, b_2) is a Nash equilibrium;
- a_1 is a best response to b_3 , b_3 is a best response to a_3 , a_3 is a best response to b_1 , and b_1 is a best response to a_1 .
- b_4 is not rationalizable since if the probability that player 2's belief assigns to a_4 exceeds $\frac{1}{2}$, then b_3 yields a payoff higher than does b_4 , while if this probability is at most $\frac{1}{2}$ then b_2 yields a payoff higher than does b_4 .
- a_4 is not rationalizable since without b_4 in the support of player 1's belief, a_4 is dominated by a_2 .

\square

7.15 Example (From <http://econweb.ucsd.edu/~jwatson/>): In a strategic game, player 1 selects any real number x and player 2 selects any real number y . The payoffs are given by

$$u_1(x, y) = 2x - x^2 + 2xy, \quad u_2(x, y) = 10y - 2xy - y^2.$$

Find all pure-strategy Nash equilibria and the largest rationalizable set.

Answer. Unique Nash equilibrium $(3, 2)$, the largest rationalizable set is $\mathbb{R} \times \mathbb{R}$. □

7.2 Iterated elimination of never-best response

7.16 The notion of rationalizability determines not what actions should actually be taken, *i.e.*, the action which is a never-best response will not be taken.

7.17 Definition (Never-best response): Given a product subset $X \subseteq A$, an action a_i of player i is a never-best response given X if it is not a best response to any belief on X of player i , *i.e.*, for each $\mu \in \Delta(X_{-i})$, there is $a_i^* \in A_i$ such that

$$U_i(a_i, \mu) < U_i(a_i^*, \mu).$$

7.18 Notation: For product subsets $X, X' \subseteq A$ where $X' \subseteq X$, we use the notation

$$X \succrightarrow X'$$

to signify that for any $a_j \in X_j \setminus X'_j$ and for any $\mu \in \Delta(X_{-j})$, there is $a_j^* \in A_j$ such that

$$U_j(a_j, \mu) < U_j(a_j^*, \mu).$$

The notation $X \succrightarrow X'$ does not require that all never-best responses be eliminated; in particular, $X \succrightarrow X$.

7.19 IENBR: The product set $X \subseteq A$ of outcomes of a finite strategic game is the result of iterated elimination of never-best response (IENBR) if there is a collection $(X^t)_{t=0}^T$, where $X^0 = A$, $X^t \succrightarrow X^{t+1}$, and $X = \bigcap_{t=0}^T X^t \succrightarrow X'$ only for $X' = X$.

Sometimes, we use IENBR to denote the result of iterated elimination of never-best response.

For infinite games we may need to consider a countably/uncountably infinite number of rounds of elimination.

7.20 Claim: Let X be the result of iterated elimination of never-best response for a finite strategic game, then X is the largest rationalizable set.

Proof. Since R^* is rationalizable according to Definition 7.9, then $R^* \subseteq X^t$ for all $t \geq 0$, and hence $R^* \subseteq X$.

Since $X \succrightarrow X'$ only for $X' = X$, then X is rationalizable according to Definition 7.9, and hence $X \subseteq R^*$. □

7.21 This claim implies that the order and speed of iterated elimination of never-best response have no effect on the set of outcomes that survive.

Every Nash equilibrium survives IENBR.

7.22 Example (Order dependence): Consider the following game $G = \langle \{1\}, A_1, u_1 \rangle$, where $A_1 = (0, 1)$ and $u_1(x) = x$ for all $x \in A_1$.

Every strategy is a never-best response. For any $x \in A_1$, eliminate in round one all strategies in the set $A_1 \setminus \{x\}$, and only x survives IENBR.


7.23 Let NE denote the set of Nash equilibria, and $NE|_{R^*}$ the set of Nash equilibria in the reduced game after the IENBR procedure. We have the following $NE = NE|_{R^*}$.

Proof. Clearly, $NE \subseteq NE|_{R^*}$.

Assume, in negation, that there is $a^* \in NE|_{R^*}$ but $a^* \notin NE$. Then, for some player j , a_j^* is not a best response to $a_{-j}^* \in R_{-j}$ in A_j , although a_j^* is a best response to a_{-j}^* in R_j . Suppose that a_j is a best response to a_{-j}^* for player j . Then, a_j is eliminated by the IENBR procedure, otherwise a_j^* can not be an equilibrium action in the reduced game. A contradiction. \square

7.3 Iterated elimination of strictly dominated actions

7.24 Iterated strict dominance is one of the most basic principles in game theory. The concept of iterated strict dominance rests on the simple idea: no player would play strategies for which some alternative strategy can yield him/her a greater payoff regardless of what the other players play and this fact is common knowledge.

 7.25 Definition (Strictly dominated action): The action $a_i \in A_i$ of player i is strictly dominated if there is a mixed strategy $\alpha_i \in \Delta(A_i)$ of player i such that

$$U_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i}) \text{ for all } a_{-i} \in A_{-i}.$$

7.26 Lemma: An action is strictly dominated if and only if it is a never-best response.

7.27 *Proof.* “ \Rightarrow ”: Let a_i be strictly dominated by α_i . Then for any $\mu \in \Delta(A_{-i})$, we have

$$U_i(a_i, \mu) = \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}) \cdot u_i(a_i, a_{-i}) < \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}) \cdot U_i(\alpha_i, a_{-i}) = U_i(\alpha_i, \mu).$$

Thus, a_i is not a best response to μ , and hence a_i is never-best response.

“ \Leftarrow ”: Let a_i be not a best response to any $\mu \in \Delta(A_{-i})$. Consider a two-person zero-sum game

$$G = \langle \{i, -i\}, (\Delta(A_j)), (V_j) \rangle,$$

where $V_i(\alpha_i, \mu) = U_i(\alpha_i, \mu) - U_i(a_i, \mu)$ for all $\alpha_i \in \Delta(A_i)$ and $\mu \in \Delta(A_{-i})$.

Clearly, $\Delta(A_j)$ is non-empty, convex and compact, and the function $V_i(\alpha_i, \mu)$ is continuous and linear (hence quasi-concave) in α_i . So, there is a Nash equilibrium (α_i^*, μ^*) in game G . Since a_i is a never-best response, for any $\mu \in \Delta(A_{-i})$,

$$V_i(\alpha_i^*, \mu) \geq V_i(\alpha_i^*, \mu^*) = \max_{\alpha_i \in \Delta(A_i)} V_i(\alpha_i, \mu^*) > 0.$$

Hence, $U_i(\alpha_i^*, a_{-i}) > u_i(a_i, a_{-i})$ for all $a_{-i} \in A_{-i}$, i.e., a_i is strictly dominated by α_i^* . \square

7.28 Hyperplane separating theorem: Let A and B be convex sets in \mathbb{R}^n such that $A \cap B = \emptyset$. Then, there exists $p \in \mathbb{R}^n \setminus \{0\}$, such that $\sup_{x \in A} p \cdot x \leq \inf_{y \in B} p \cdot y$.

Alternative proof for “ \Leftarrow ” (by hyperplane separating theorem). Assume that a_i is not strictly dominated and enumerate $A_{-i} = \{a_{-i}^1, a_{-i}^2, \dots, a_{-i}^k\}$. Define $V_i = \cup_{\beta_i \in \Delta(A_i)} V_i(\beta_i)$, where

$$V_i(\beta_i) = \{(\bar{u}^m)_{m=1}^k \in \mathbb{R}^k : U_i(\beta_i, a_{-i}^m) > \bar{u}^m, \text{ for all } m = 1, 2, \dots, k\}.$$

Note that V_i is a convex set, and since a_i is not strictly dominated,

$$(u_i(a_i, a_{-i}^1), u_i(a_i, a_{-i}^2), \dots, u_i(a_i, a_{-i}^k)) \notin V_i.$$

Therefore, by hyperplane separating theorem, there is some $(p_i^m)_{m=1}^k \in \mathbb{R}^k \setminus \{0\}$ such that

$$\sum_{m=1}^k p_i^m u_i(a_i, a_{-i}^m) \geq \sum_{m=1}^k p_i^m \bar{u}^m \text{ for all } (\bar{u}^m)_{m=1}^k \in V_i. \quad (7.1)$$

Note first that $p_i^m \geq 0$ because if $p_i^m < 0$, we can pick $\bar{u}^m < 0$ so that Equation (7.1) violated. Moreover, since $(p_i^m)_{m=1}^k \neq 0$, $\sum_{m=1}^k p_i^m > 0$ and thus we can normalize $(p_i^m)_{m=1}^k$ so that $\sum_{m=1}^k p_i^m = 1$. Thus, $(p_i^m)_{m=1}^k$ is a probability distribution on A_{-i} where p_i^m is the probability that $-i$ plays a_{-i}^m . This normalization will not change Equation (7.1). Note also that for every $a'_i \in A_i$,

$$(u_i(a'_i, a_{-i}^1) - \epsilon, u_i(a'_i, a_{-i}^2) - \epsilon, \dots, u_i(a'_i, a_{-i}^k) - \epsilon) \in V_i(a'_i)$$

for any $\epsilon > 0$ and thus,

$$\sum_{m=1}^k p_i^m u_i(a_i, a_{-i}^m) \geq \sum_{m=1}^k p_i^m [u_i(a'_i, a_{-i}^m) - \epsilon] \text{ for all } a'_i \in A_i.$$

Since $\epsilon > 0$ is arbitrary, we get

$$\sum_{m=1}^k p_i^m u_i(a_i, a_{-i}^m) \geq \sum_{m=1}^k p_i^m u_i(a'_i, a_{-i}^m) \text{ for all } a'_i \in A_i.$$

Thus, a_i is a best response. A contradiction. \square

7.29 Notation: For product subsets $X, X' \subseteq A$ where $X' \subseteq X$, we use the notation

$$X \rightarrow X'$$

to signify that for any $a_j \in X_j \setminus X'_j$, there is $\alpha_j^* \in \Delta(A_j)$ such that

$$U_j(\alpha_j^*, a_{-j}) > u_j(a_j, a_{-j}) \text{ for all } a_{-j} \in X_{-j}.$$

7.30 IESDA: The product set $X \subseteq A$ of outcomes of a finite strategic game is the result of iterated elimination of strictly dominated actions (IESDA) if there is a collection $(X^t)_{t=0}^T$, where $X^0 = A$, $X^t \rightarrow X^{t+1}$, and $X = \cap_{t=0}^T X^t \rightarrow X'$ only for $X' = X$.

Sometimes, we use IESDA to denote the result of iterated elimination of strictly dominated actions.

7.31 Claim: For a finite strategic game, the result of IENBR is the same as the result of IESDA. (By Lemma 7.26)

7.32 Example:

		Player 2	
		L	R
Player 1	T	3, 0	0, 1
	M	0, 0	3, 1
	B	1, 1	1, 0

G

		Player 2	
		L	R
Player 1	T	3, 0	0, 1
	M	0, 0	3, 1

G'

		Player 2	
		R	
Player 1	T	0, 1	
	M	3, 1	

G''

- For player 1, B is strictly dominated by $\frac{1}{2} \circ T + \frac{1}{2} \circ M$, which will be eliminated;
- In reduced game G' , for player 2, L is strictly dominated by R , which will be eliminated;
- In reduced game G'' , for player 1, T is strictly dominated by M , which will be eliminated;
- Thus, (M, R) is the only outcome that survives IESDA.

7.33 Notation: For product subsets $X, X' \subseteq A$ where $X' \subseteq X$, we use the notation

$$X \rightsquigarrow X'$$

to signify that for any $a_j \in X_j \setminus X'_j$ and for any $\mu \in \Delta(X_{-j})$, there is $a_j^* \in X_j$ such that

$$U_j(a_j^*, \mu) > U_j(a_j, \mu).$$

We use the notation

$$X \hookrightarrow X'$$

to signify that for any $a_j \in X_j \setminus X'_j$, there is $\alpha_j^* \in \Delta(X_j)$ such that

$$U_j(\alpha_j^*, a_{-j}) > u_j(a_j, a_{-j}) \text{ for all } a_{-j} \in X_{-j}.$$

7.34 The product set $X \subseteq A$ of outcomes of a finite strategic game is the result of iterated elimination of strictly dominated actions' (IESDA') if there is a collection $(X^t)_{t=0}^T$, where $X^0 = A$, $X^t \hookrightarrow X^{t+1}$, and $X = \bigcap_{t=0}^T X^t \hookrightarrow X'$ only for $X' = X$.

The product set $X \subseteq A$ of outcomes of a finite strategic game is the result of iterated elimination of never-best response' (IENBR') if there is a collection $(X^t)_{t=0}^T$, where $X^0 = A$, $X^t \rightsquigarrow X^{t+1}$, and $X = \bigcap_{t=0}^T X^t \rightsquigarrow X'$ only for $X' = X$.

Similarly, the result of IESDA' is the same as the result of IENBR'.

7.35 Claim: For a finite strategic game, the result of IESDA is the same as the result of IESDA'.

Proof. Suppose that X is the result of IESDA', then it is also the result of IENBR'. Since $Z \rightsquigarrow Z'$ implies $Z \hookrightarrow Z'$, IENBR has greater elimination power than IENBR'. Therefore, $X \subseteq R^*$.

Assume, in negation, that there is $a \in X \setminus R^*$. Since $a \in X$ is not rationalizable, for any $\mu \in \Delta(X_{-j})$, there exists $a'_j \in A_j$ such that $U_j(a'_j, \mu) > U_j(a_j, \mu)$. Take a_j^* so that $U_j(a_j^*, \mu) = \max_{a'_j \in A_j} U_j(a'_j, \mu)$. Then $U_j(a_j^*, \mu) > U_j(a_j, \mu)$, and hence a_j^* is a best response to μ . It is clear that a_j^* is eliminated in the process of IENBR', otherwise a_j can not survive. A contradiction. \square

7.36 Summary: $R^* = \text{IENBR} = \text{IESDA} = \text{IENBR}' = \text{IESDA}'$

7.4 Examples

7.37 Example: Find all pure-strategy Nash equilibria and all pure rationalizable strategies in the following strategic game.

		Player 2		
		L	C	R
Player 1	U	3, 0	0, 1	1, 0
	M	0, 1	3, 0	0, 0
	D	1, 0	1, 0	0, 1

Answer. No pure-strategy Nash equilibrium.

U is a best response for L , L is a best response for M , M is a best response for C , and C is a best response. D is strictly dominated by $\frac{1}{2} \circ U + \frac{1}{2} \circ M$ and eliminated, and consequently R is strictly dominated by $\frac{1}{2} \circ L + \frac{1}{2} \circ C$. Therefore, U , L , M and C are all pure rationalizable strategies. \square

7.38 Example [OR Exercise 56.4]: Cournot duopoly.

Consider the strategic game $\langle \{1, 2\}, (A_i), (u_i) \rangle$ in which $A_i = [0, 1]$ and $u_i(a_1, a_2) = a_i(1 - a_1 - a_2)$ for $i = 1, 2$. Show that each player's only rationalizable action is his unique Nash equilibrium action.

Answer. Player i 's best response function is $B_i(a_j) = (1 - a_j)/2$; hence the only Nash equilibrium is $(\frac{1}{3}, \frac{1}{3})$.

Since the game is symmetric, the set of rationalizable actions is the same for both players; denote it by Z . Let $m = \inf Z$ and $M = \sup Z$. Any best response of player i to a belief of player j whose support is a subset of Z maximizes $E[a_i(1 - a_i - a_j)] = a_i(1 - a_i - E[a_j])$, and thus is equal to $B_i(E[a_j]) \in [B_j(M), B_j(m)] = [(1 - M)/2, (1 - m)/2]$. Hence, we need $(1 - M)/2 \leq m$ and $M \leq (1 - m)/2$, so that $M = m = \frac{1}{3}$: $\frac{1}{3}$ is the only rationalizable action of each player. \square

7.39 Example [OR Exercise 56.5]: Guessing the average.

Let n ($n \geq 2$) people play the following game. Simultaneously, each player i announces a number x_i in the set $\{1, 2, \dots, K\}$. A prize of \$1 is split equally between all the people whose number is closest to $\frac{2}{3} \cdot \frac{x_1 + \dots + x_n}{n}$. Show that each player's equilibrium action is his unique rationalizable action.

7.40 Example [OR Exercise 57.1].

7.41 Example [OR Exercise 63.1].

Consider a variant of the game in Example 2.38 in which there are two players, the distribution of the citizens' favorite positions is uniform, and each player is restricted to choose a position of the form ℓ/m for some $\ell \in \{0, \dots, m\}$, where m is even. Show that the only outcome that survives iterated elimination of weakly dominated actions is that in which both players choose the position $\frac{1}{2}$.

Answer. Only one round of elimination is needed: every action other than $\frac{1}{2}$ is weakly dominated by the action $\frac{1}{2}$. (In fact $\frac{1}{2}$ is the only action that survives iterated elimination of strictly dominated actions: on the first round Out is strictly dominated by $\frac{1}{2}$, and in every subsequent round each of the remaining most extreme actions is strictly dominated by $\frac{1}{2}$.) \square

7.42 Example [OR Exercise 63.2]: Dominance solvability.

A strategic game is dominance solvable if all players are indifferent between all outcomes that survive the iterative procedure in which all the weakly dominated actions of each player are eliminated at each stage. Give an example of a strategic game that is dominance solvable but for which it is not the case that all players are indifferent between all outcomes that survive iterated elimination of weakly dominated actions (a procedure in which not all weakly dominated actions may be eliminated at each stage).

Answer. Consider the following game. This game is dominance solvable, the only surviving outcome being (T, L) . However, if B is deleted then neither of the remaining actions of player 2 is dominated, so that both (T, L) and (T, R) survive iterated elimination of dominated actions.

		Player 2	
		L	R
Player 1	T	1, 0	0, 0
	B	0, 1	0, 0

Figure 7.1

□

7.43 Example [OR Exercise 64.1]: Announcing numbers.

Each of two players announces a non-negative integer equal to at most 100. If $a_1 + a_2 \leq 100$, where a_i is the number announced by player i , then each player i receives payoff of a_i . If $a_1 + a_2 > 100$ and $a_i < a_j$ then player i receives a_i and player j receives $100 - a_i$; if $a_1 + a_2 > 100$ and $a_i = a_j$ then each player receives 50. Show that the game is dominance solvable (see the previous exercise) and find the set of surviving outcomes.

Answer. At the first round every action $a_i \leq 50$ of each player i is weakly dominated by $a_i + 1$. No other action is weakly dominated, since 100 is a strict best response to 0 and every other action $a_i \geq 51$ is a best response to $a_i + 1$. At every subsequent round up to 50 one action is eliminated for each player: at the second round this action is 100, at the third round it is 99, and so on. After round 50 the single action pair $(51, 51)$ remains, with payoffs of $(50, 50)$. □

7.44 Consider the following game:

		Player 2		
		L	M	R
Player 1	A	4, 3	2, 5	2, 0
	B	6, 2	0, 3	1, 4
	C	3, 1	1, 0	1, 2
	D	3, 0	1, 1	3, 3

- (i) Eliminate strictly dominated strategies.
- (ii) Find all pure-strategy Nash equilibria and write down the corresponding payoffs.
- (iii) Find all mixed-strategy Nash equilibria and write down the corresponding expected payoffs.

Answer. (i) (1) C is strictly dominated by A and will be eliminated;
 (2) L is strictly dominated by M and will be eliminated;

		Player 2	
		<i>M</i>	<i>R</i>
Player 1	<i>A</i>	2, 5	2, 0
	<i>D</i>	1, 1	3, 3

G_1

		Player 2	
		<i>M</i>	<i>R</i>
Player 1	<i>A</i>	2, 5	2, 0
	<i>D</i>	1, 1	3, 3

G_2

(3) *B* is strictly dominated by *D* and will be eliminated.

Hence we will obtain the reduced game G_1 .

- (ii) From the payoff table G_2 , we obtain the pure-strategy Nash equilibria: (*A*, *M*) and (*D*, *R*) (red pairs) with payoffs (2, 5) and (3, 3), respectively.
- (iii) Let $p_1 = (r, 1 - r)$ be a mixed strategy in which player 1 plays *A* with probability r . Let $p_2 = (q, 1 - q)$ be a mixed strategy in which player 2 plays *M* with probability q . Then player 1's expected payoff is:

$$U_1(A, p_2) = 2q + 2(1 - q) = 2,$$

$$U_1(D, p_2) = q + 3(1 - q) = 3 - 2q.$$

Hence

$$r^*(q) \equiv \arg \max_{0 \leq r \leq 1} U_1(p_1, p_2) = \begin{cases} \{1\}, & \text{if } q > \frac{1}{2}; \\ \{0\}, & \text{if } q < \frac{1}{2}; \\ [0, 1], & \text{if } q = \frac{1}{2}. \end{cases}$$

Similarly, player 2's expected payoff is:

$$U_2(p_1, M) = 5r + (1 - r) = 1 + 4r,$$

$$U_2(p_1, R) = 3(1 - r).$$

Hence

$$q^*(r) \equiv \arg \max_{0 \leq q \leq 1} U_2(p_1, p_2) = \begin{cases} \{1\}, & \text{if } r > \frac{2}{7}; \\ \{0\}, & \text{if } r < \frac{2}{7}; \\ [0, 1], & \text{if } r = \frac{2}{7}. \end{cases}$$

We draw the graphs of $r^*(q)$ and $q^*(r)$ together:




Figure 7.2: Intersection of best-response correspondences.

The graphs of the best response correspondences $r^*(q)$ and $q^*(r)$ intersect at 3 points $(r = \frac{2}{7}, q = \frac{1}{2})$, $(0, 0)$ and $(1, 1)$. Hence, there are 3 mixed-strategy Nash equilibria:

- $(1 \circ A, 1 \circ M)$ with expected payoff $(2, 5)$,
- $(1 \circ D, 1 \circ R)$ with expected payoff $(3, 3)$,
- $(\frac{2}{7} \circ A + \frac{5}{7} \circ D, \frac{1}{2} \circ M + \frac{1}{2} \circ R)$ with expected payoff $(2, \frac{15}{7})$.

□

7.45 Example: Each individual $i = 1, 2, \dots, 100$ must choose a number $r_i \in [0, 1]$. If an individual chooses a number that is the most closed to the value $\theta \sum_{i=1}^{100} r_i$ (where $\theta \in [0, 1]$ is a parameter), then the individual gets payoff 1; otherwise, the individual gets payoff 0. Formulate this problem as a strategic game, and find all rationalizable strategies for each $\theta \in [0, 1]$.

7.5 Iterated elimination of weakly dominated actions

7.46 A player's action is weakly dominated if the player has another action at least as good no matter what the other players do and better for at least some vector of actions of the other players.

7.47 The action $a_i \in A_i$ of player i is weakly dominated if there is a mixed strategy $\alpha_i \in \Delta(A_i)$ of player i such that

$$\begin{aligned} U_i(\alpha_i, a_{-i}) &\geq u_i(a_i, a_{-i}) \text{ for all } a_{-i} \in A_{-i}, \\ U_i(\alpha_i, a_{-i}) &> u_i(a_i, a_{-i}) \text{ for some } a_{-i} \in A_{-i}. \end{aligned}$$

7.48 Since a weakly dominated action may be a best response to some belief, rationality does not exclude using such an action. Yet since there is no advantage to using a weakly dominated action, it seems very natural to eliminate such actions in the process of simplifying a complicated game. Indeed, a “cautious” player who holds full-support probabilistic beliefs about the opponents' behavior never uses a weakly dominated action.

7.49 Similarly to IESDA, we can define iterated elimination of weakly dominated actions (IEWDA). Unlike IESDA, IEWDA might be an order dependent procedure.

		Player 2	
		L	R
Player 1	T	1, 1	0, 0
	M	1, 1	2, 1
	B	0, 0	2, 1

7.50 Example:

- Procedure 1: T is eliminated (it is weakly dominated by M), then L is eliminated (it is weakly dominated by R). Thus, the result is $\{(M, R), (B, R)\}$.
- Procedure 2: B is eliminated (it is weakly dominated by M), then R is eliminated (it is weakly dominated by L). Thus, the result is $\{(T, L), (M, L)\}$.

7.51 Proposition [JR Exercise 7.16]: In a finite strategic game, the set of strategies surviving iterative weak dominance is non-empty.

Knowledge model

Contents

8.1	A model of knowledge	131
8.2	Common knowledge	136
8.3	Common prior	137
8.4	“Agree to disagree” is impossible	138
8.5	No-trade theorem	140
8.6	Speculation	141
8.7	Characterization of the common prior assumption	143
8.8	Unawareness	145

In game theory and economics, it is important and fundamental to account for an agent’s knowledge/beliefs about uncertainty, as well as the agent’s higher-order knowledge/beliefs about the other agents’ knowledge/beliefs. Interactive epistemology studies the logic of knowledge and belief in the context of interactions.

We present a model of knowledge and use it to formalize the idea that an event is “common knowledge”, and to express formally the assumptions on players’ knowledge that lie behind various solution concepts such as Nash equilibrium and rationalizability.

8.1 A model of knowledge

- 8.1 There are 3 players, each wearing a white hat. Each player knows that the hats are either white or black and sees the color of the hats of the other two players. The game is to guess the color of their own hats (which they don’t see).

Obviously, the hat color of the other two players contain no information about one’s own hat. Hence, initially none of the three players can tell his own hat color. Now, suppose a fourth player announces publicly to all three players that at least one hat is white. Then the fourth person asks them whether they now know the color of their hats. They said no. The fourth player asks the question again. The answer is still no. The fourth player asks the question for the third time. Now, they all answer correctly. Why?

First, do the three players learn anything that they do not know already from the fourth player? Everyone knows that some hats are white. Furthermore, everyone knows everyone else knows some hats are white. For example,

player 1 knows players 2 and 3's hats are white. So, he knows player 2 knows someone's hat (player 3's) is white. Similarly, player 1 also knows player 3 knows player 2's hat is white. But, does player 1 know player 2 knows player 3 knows someone's hat is white? The answer is no. Because player 1 does not know his own hat's color. If player 1's hat is white, then player 2 will not know player 3 knows some hats are white. So, it is not common knowledge before the fourth person's remark that some hats are white.

Making this fact common knowledge makes a whole lot of difference. Thinking from the perspective of player 1.

- (1) Player 1 knows that player 2 knows that player 3 would know immediately that his (person 3's) hat is white if both player 1's and player 2's hats were black. The fact that player 3 didn't immediately know that his hat is white should tell player 2 that either player 1's or player 2's hat must be white.
- (2) Now if player 1's hat were black, player 2 should conclude that his (person 2's) hat must be white. The fact that player 2 didn't know tells player 1 that player 1's hat must be white.

Same logic for the other two persons. The argument can extend to n persons with m white hats.

- 8.2 The standard semantic model of the knowledge of a single decision-maker is associated with Hintikka (1962, Knowledge and Belief, Cornell University Press). The model is given by

$$\langle \Omega, P \rangle,$$

- Ω is the set of states
- information function $P: \Omega \rightarrow \Omega$ such that for each $\omega \in \Omega$, $\emptyset \neq P(\omega) \subseteq \Omega$.

Interpretation of information function: when the state is ω the decision-maker knows only that the state is in the set $P(\omega)$. That is, he considers it possible that the true state could be any state in $P(\omega)$ but not any state outside $P(\omega)$.

- 8.3 Example: $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, $P(\omega_1) = P(\omega_2) = \{\omega_1, \omega_2\}$ and $P(\omega_3) = P(\omega_4) = \{\omega_3, \omega_4\}$.

In this case, the decision maker knows whether the true state lies in the set $\{\omega_1, \omega_2\}$ or $\{\omega_3, \omega_4\}$, though he can not perfectly identify the true state.

- 8.4 Example: $\Omega = \{\omega_1, \omega_2\}$, $P(\omega_1) = \{\omega_1\}$ and $P(\omega_2) = \{\omega_1, \omega_2\}$.

In this case, the decision maker knows the true state when she is at state ω_1 , but she is unsure of the true state when ω_2 occurs.

Notice that a rational decision maker (who also knows the mapping P) could reason that when she does not know the state, she must be in state ω_2 . One example of such a scenario is: decision maker is drunk in ω_2 , and not drunk in ω_1 .

- 8.5 We usually assume that $\langle \Omega, P \rangle$ satisfies the following two conditions:

- P1. $\omega \in P(\omega)$ for every $\omega \in \Omega$.
- P2. If $\omega' \in P(\omega)$ then $P(\omega') = P(\omega)$.

- P1 says that the decision-maker never excludes the true state from the set of states he regards as feasible.
- P2 says that the decision-maker uses the consistency or inconsistency with the information structure.

- 8.6 An information function P for the set Ω of states is partitional if there is a partition of Ω such that for any $\omega \in \Omega$ the set $P(\omega)$ is the element of the partition that contains ω .

8.7 Lemma: The information structure under P is partitional if and only if P satisfies P1 and P2.

Proof. “ \Rightarrow ”: If P is partitional then it clearly satisfies P1 and P2.

“ \Leftarrow ”: Suppose that P satisfies P1 and P2. If $P(\omega)$ and $P(\omega')$ intersect and $\omega'' \in P(\omega) \cap P(\omega')$ then by P2 we have $P(\omega) = P(\omega') = P(\omega'')$; by P1 we have $\cup_{\omega \in \Omega} P(\omega) = \Omega$. Thus P is partitional. \square

8.8 P2': if $\omega' \in P(\omega)$, then $P(\omega) \subseteq P(\omega')$.

8.9 Claim: P1 and P2 are equivalent to P1 and P2'.

Proof. It suffices to show that P1 and P2' imply P2. If $\omega' \in P(\omega)$, then by P2' $P(\omega) \subseteq P(\omega')$. Thus, by P1 $\omega \in P(\omega)$, $\omega \in P(\omega')$, which, again by P2', implies that $P(\omega') \subseteq P(\omega)$. \square

8.10 We refer to a subset $E \subseteq \Omega$ as an event. We say the decision-maker knows/believes E at ω if $P(\omega) \subseteq E$.

In Figure 8.1, the event E is known by player 1 at the state ω since $P_1(\omega) = P_1^2 \subseteq E$.




Figure 8.1: Knowledge

8.11 Given P , define the decision-maker's knowledge operator $K : 2^\Omega \rightarrow 2^\Omega$ by

$$K(E) = \{\omega \in \Omega : P(\omega) \subseteq E\}.$$

For any event E , the set $K(E)$ is the set of all states in which the decision-maker knows E . The statement “the decision-maker knows E ” is identified with all states in which E is known. Moreover, the set $K(K(E))$ is interpreted as “the decision-maker knows that he knows E ”.

In Figure 8.1, $K_1(E) = \{\omega : P_1(\omega) \subseteq E\} = P_1^2$.




Figure 8.2: Knowledge

In Figure 8.2, player 2 knows E at ω since $P_2(\omega) = P_2^3 \subseteq E$, and $K_2(E) = P_2^3$. However, player 1 does not know that player 2 knows E at ω , since there is no player 1's information cell in the subset P_2^3 .

8.12 The knowledge operator K satisfies the following properties:

K1. $K(\Omega) = \Omega$.

K2. If $E \subseteq F$ then $K(E) \subseteq K(F)$.

K3. $K(E) \cap K(F) = K(E \cap F)$. Moreover, $K(\cap_{\lambda} E_{\lambda}) = \cap_{\lambda} K(E_{\lambda})$.

- K1 says that in all states the decision-maker knows that some state in Ω has occurred.
- K2 says that if F occurs whenever E occurs and the decision-maker knows E then he knows F .
- K3 says that the decision-maker knows both E and F then he knows $E \cap F$.

Proof. K1: Since $P(\omega) \subseteq \Omega$, $\Omega \subseteq K(\Omega)$.

K2: If $E \subseteq F$ and $\omega \in K(E)$, then $P(\omega) \subseteq E \subseteq F$ and hence $\omega \in K(F)$.

K3: $\omega \in K(E) \cap K(F)$ if and only if $P(\omega) \subseteq E$ and $P(\omega) \subseteq F$ if and only if $P(\omega) \subseteq E \cap F$ if and only if $\omega \in K(E \cap F)$. \square

Beware $K(E) \cup K(F) \neq K(E \cup F)$.

8.13 If P satisfies P1, then the associated knowledge operator K satisfies the following additional property.

K4. (Axiom of Knowledge) $K(E) \subseteq E$.

This says that whenever the decision-maker knows E then indeed some member of E is the true state (E must be true): the decision-maker does not know anything that is false.

Replace E by $K(E)$, we have $K(K(E)) \subseteq K(E)$. If decision-maker knows that he knows that E is true, then he knows that E is true.

Proof. Let $\omega \in K(E)$. Then $P(\omega) \subseteq E$, by P1 $\omega \in P(\omega)$ and thus $\omega \in E$. \square

8.14 If P is partitional, then $K(E)$ is the union of all members of the partition that are subsets of E . In this case, the knowledge operator K satisfies the following two additional properties.

K5. (Axiom of Transparency) $K(E) \subseteq K(K(E))$.

K6. (Axiom of Wisdom) $\Omega \setminus K(E) \subseteq K(\Omega \setminus K(E))$.

- K5 says that if the decision-maker knows E then he knows that he knows E . (self awareness)
- K6 says that if the decision-maker does not know E then he knows that he does not know E .

Note that given that K satisfies K4 the properties in K5 and K6 in fact hold with equality.

Proof. K5: Let $\omega \in K(E)$. Then $P(\omega) \subseteq E$, by P2 $P(\omega') \subseteq E$ for all $\omega' \in P(\omega)$ and thus $P(\omega) \subseteq K(E)$, i.e., $\omega \in K(K(E))$.

K6: Let $\omega \notin K(E)$. Then $P(\omega) \not\subseteq E$, by P2 $P(\omega') \not\subseteq E$ for all $\omega' \in P(\omega)$ and thus $P(\omega) \subseteq \Omega \setminus K(E)$, i.e., $\omega \in K(\Omega \setminus K(E))$. \square

8.15 Alternatively we can start by defining a knowledge operator for the set Ω to be a function K that associates a subset of Ω with each event $E \subseteq \Omega$.

We can then derive from K an information function P as follows: for each state ω let

$$P(\omega) = \cap \{E \subseteq \Omega : \omega \in K(E)\}.$$

Given an information function P , let K be the knowledge operator derived from P and let P' be the information function derived from K . Then $P' = P$.

8.16 Example: Puzzle of hats

Each of n “perfectly rational” individuals, seated around a table, is wearing a hat that is either white or black. Each individual can see the hats of the other $n - 1$ individuals, but not his own. Assume that two are wearing a white hat. An observer announces: “Each of you is wearing a hat that is either white or black; at least one of the hats is white. I will start to count slowly. After each number you will have the opportunity to raise a hand. You may do so only when you know the color of your hat.” When, for the first time, will any individual raise his hand?

Answer. The two wearing white hats will raise their hands when the observer counts the number “2”.

Intuitively,

- When the observer counts the number “1”, no one knows the color of his hat. The two wearing white hats see “1 white and 98 black”, and the other see “2 white and 97 black”.
- When the observer counts the number “2”, the two wearing white hats now know the color of their hats. Each of them who is wearing a white hat can reason as follows: Since the one wearing a white hat (saw in period 1) did not raise his hand, there should be another white hat which must be on my head.

Formal reason:

Each state can be written as $c = (c_1, c_2, \dots, c_n)$, where each c_i is either W or B and at least one c_i is W . The state space is

$$\Omega = \{c \in \{W, B\}^n : |\{i : c_i = W\}| \geq 1\}.$$

Player i ’s initial information function P_i^1 is as follows:

$$P_i^1(c) = \begin{cases} \{(c_{-i}, W), (c_{-i}, B)\}, & \text{if } c \text{ is the state in which a player different from } i \text{ has a white hat,} \\ \{c\}, & \text{if } c \text{ is the state in which all the other hats are black.} \end{cases}$$

The event “ i knows the color of his hat” is

$$E_i = \{c \in \Omega : P_i(c) \subseteq \{c : c_i = B\} \text{ or } P_i(c) \subseteq \{c : c_i = W\}\}.$$

Let $F^1 = \{c : |\{i : c_i = W\}| = 1\}$, the set of states for which someone raises a hand at the first stage.

Since there are two white hats, nobody raises a hand at the first stage, and then all players obtain the additional information that the state is not in F^1 . Therefore, they will update their information: for all i and for all $c \notin F^1$, we have $P_i^2(c) = P_i^1(c) \setminus F^1$: $P_i^2(c)$ is $\{(c_{-i}, W), (c_{-i}, B)\}$ unless $c_j = W$ for exactly one player $j \neq i$, in which case $P_i^2(c_{-i}, W) = \{(c_{-i}, W)\}$ and $P_j^2(c_{-j}, W) = \{(c_{-j}, W)\}$. In other words, in any state c for which $c_j = W$ and $c_h = W$ for precisely two players j and h we have $P_j^2(c) \subseteq E_j$ and $P_h^2(c) \subseteq E_h$, and hence j and h each raises a hand at the second stage. \square

It is easy to see that if k hats are white then no one raises a hand until the observer counts k , at which point the k individuals with white hats do so.

- 8.17 孙臆，庞涓都是鬼谷子的徒弟。一天鬼谷子出了这道题目：他从 2 到 99 中选出两个不同的整数，把积告诉孙，把和告诉庞；庞说：我虽然不能确定这两个数是什么，但是我肯定你也不知道这两个数是什么。孙说：我本来的确不知道，但是听你这么一说，我现在能够确定这两个数字了。庞说：既然你这么说，我现在也知道这两个数字是什么了。请问这两个数字是什么？为什么？

8.2 Common knowledge

8.18 Let K_i be the knowledge operator of player i for each $i \in N$. For event $E \subseteq \Omega$,

- $K_i(E) = \{\omega \in \Omega : P_i(\omega) \subseteq E\}$ is the event that i knows E .
- $K(E) = \bigcap_{i \in N} K_i(E)$ is the event that E is mutually known.
- $CKE = K(E) \cap K(K(E)) \cap K(K(K(E))) \cap \dots = \bigcap_{k=1}^{\infty} K^k(E)$ is the event that E is commonly known, and for each $\omega \in CKE$, E is said to be common knowledge in the state ω .

- 8.19 Example: The state space is $\Omega = \{\omega_1, \omega_2, \omega_3\}$. The information functions are given as follows: $P_1(\omega_1) = \{\omega_1\}$, $P_1(\omega_2) = P_1(\omega_3) = \{\omega_2, \omega_3\}$; $P_2(\omega_1) = P_2(\omega_2) = \{\omega_1, \omega_2\}$, $P_2(\omega_3) = \{\omega_3\}$.

Consider $E = \{\omega_2, \omega_3\}$. We have

$$K_1(E) = \{\omega_2, \omega_3\}, K_2(E) = \{\omega_3\},$$

and hence

$$K(E) = \{\omega_3\}.$$

Then $K_1(K(E)) = \emptyset$. Therefore, $CKE = \emptyset$.

- 8.20 An event $F \subseteq \Omega$ is self-evident if $K(F) = F$, i.e., $K_i(F) = F$ for all $i \in N$. Whenever a self-evident event occurs everyone knows it occurs. The concept generalizes the condition K1.

- 8.21 Lemma: Under K4, An event E is common knowledge in the state ω if and only if it includes a self-evident event F containing ω .

Proof. " \Rightarrow ": Let $\omega \in CKE$. Define $F = CKE$, by K3, we have

$$K(F) = K(CKE) = K(\bigcap_{k=1}^{\infty} K^k(E)) = \bigcap_{k=2}^{\infty} K^k(E) \supseteq \bigcap_{k=1}^{\infty} K^k(E) = CKE = F.$$

By K4, $F \subseteq K(F)$, so $K(F) = F$. By K4 again, we have $F = CKE \subseteq E$.

" \Leftarrow ": Let $\omega \in F = K(F) \subseteq E$. By K2, $F = K(F) = K^2(F) = \dots = K^k(F) \subseteq K^k(E)$ for all $k \geq 1$. Therefore, $\omega \in F \subseteq CKE$. \square

- 8.22 Claim: Under P1 and P2, the following are equivalent:

- (1) $K(E) = E$, i.e., for all $\omega \in E$, $P_i(\omega) \subseteq E$ for all i .
- (2) E is a union of members of the partition induced by P_i for all i .

Proof. “(1) \Rightarrow (2)”: For every $\omega \in \Omega$, $P_i(\omega) \subseteq E$ for all i . Then $E = \cup_{\omega \in E} P_i(\omega)$ for all i and thus E is a union members of the partition induced by P_i .

“(2) \Rightarrow (1)”: Routine. □

8.23 Example: Let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$, let P_1 and P_2 be the partitional information functions of players 1 and 2, and let K_1 and K_2 be the associated knowledge operators. Let the partitions induced by the information functions be

$$\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4, \omega_5\}, \{\omega_6\}\}, \quad \Pi_2 = \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}, \{\omega_5\}, \{\omega_6\}\}.$$

The event $E = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ does not contain any event that is self-evident and hence in no state is E common knowledge.

The event $F = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$ is self-evident and hence is common knowledge in any state in F .

8.24 Example [OR Exercise 71.1].

8.25 Example [OR Exercise 71.2].

8.3 Common prior

8.26 We can start adding beliefs to the knowledge model. In the knowledge model, for player $i \in N$, let P_i be her information function, and p_i her prior probability. We assume p_i has full support and is positive at every state ω .

We call posterior belief the function $p_i(E | P_i(\omega))$, calculated by Bayes' rule. That is,

$$p_i(E | P_i(\omega)) = \frac{p_i(E \cap P_i(\omega))}{p_i(P_i(\omega))}.$$

8.27 Lemma: The posterior satisfies the following properties:

- $p_i(\omega' | P_i(\omega)) = 0$ if $\omega' \notin P_i(\omega)$.
- $p_i(P_i(\omega) | P_i(\omega)) = 1$.

8.28 A knowledge model has a common prior if $p_i = p_j$ for all $i, j \in N$. This can occur when players are born equal, and are Bayesian updaters.

Alternatively, posterior beliefs $p_1(\cdot | \cdot), p_2(\cdot | \cdot), \dots, p_n(\cdot | \cdot)$ are generated by a common prior, if there exists a probability p such that for any $i \in N$, and $E \subset \Omega$ and any $\omega \in \Omega$, $p(E) = \sum_{\omega \in \Omega} p_i(E | P_i(\omega)) \cdot p(\omega)$.

Note that there may be multiple common priors.

8.29 The common prior can be defined in an equivalent way:

Let Π_i be player i 's information partition, and P_i player i 's information function. For each $Q_i \in \Pi_i$, we define a probability $\mu(Q_i)$ on Q_i , which can be generalized to a probability on Ω .

A prior for i is a probability p_i is a probability distribution on Ω , such that for each information cell $Q_i \in \Pi_i$, if $p_i(Q_i) > 0$, then

$$\mu_i(Q_i)(\cdot) = p_i(\cdot | Q_i).$$

Clearly, each $\mu_i(Q_i)$ is a prior for i . The set of all priors of i , denoted by X_i , is the convex hull of $\{\mu_i(Q_i) | Q_i \in \Pi_i\}$.

A probability distribution μ on Ω is a common prior if $\mu \in \cap_{i \in N} X_i$.

8.30 Example: Suppose $\Omega = \{1, 2, 3\}$. There are two players with information partitions $\Pi_1 = \{\{1, 2\}, \{3\}\}$ and $\Pi_2 = \{\{1\}, \{2, 3\}\}$. Let $s = p_1(1 \mid \{1, 2\})$ and $t = p_2(2 \mid \{2, 3\})$. Show any $s, t \in (0, 1)$ can be generated by some common prior.

Answer. Take m such that

$$m \left(1 + \frac{1-s}{s} + \frac{1-s}{s} \frac{1-t}{t} \right) = 1.$$

Set $p_1 = m, p_2 = m \frac{1-s}{s}$, and $p_3 = m \frac{1-s}{s} \frac{1-t}{t}$. □

8.31 Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$, and the posterior probabilities are

$$p_1(\omega_1 \mid \{\omega_1, \omega_2\}) = \frac{2}{3}, p_1(\omega_3 \mid \{\omega_3\}) = 1.$$

Then we can have a continuum of priors associated with this posterior probability. Two examples are

$$p_1 = (\frac{2}{3}, \frac{1}{3}, 0), p'_1 = (\frac{1}{3}, \frac{1}{6}, \frac{1}{2}).$$

8.4 “Agree to disagree” is impossible

8.32 Reference: [Aumann \(1976\)](#).

8.33 Within the framework of partitional information structures, Aumann (1976) showed that, under the common prior assumption, it can not be common knowledge (agree) that i and j respectively assign two different posterior probabilities to the same event (disagree), even if they possess different information.

To be more clear, two individuals that have prior belief cannot have different posterior beliefs, after knowing each other's posterior beliefs (notice that it is essential that the posterior beliefs be common knowledge).

✂ 8.34 Theorem (Aumann's agree to disagree): In a two-person game with finite states, assume that there is a common prior belief μ . For an event E , let

$$E^{[\mu_i; \mu_j]} = \{\omega \in \Omega: \mu(E \mid P_i(\omega)) = \mu_i \text{ and } \mu(E \mid P_j(\omega)) = \mu_j\}.$$

Then $CKE^{[\mu_i; \mu_j]} = \emptyset$ if $\mu_i \neq \mu_j$.

8.35 *Proof.* (1) Suppose that $CKE^{[\mu_i; \mu_j]} \neq \emptyset$. Let $\omega^0 \in CKE^{[\mu_i; \mu_j]}$. Then, there is a self-evident event $F \subseteq E^{[\mu_i; \mu_j]}$ that contains ω^0 .

(2) Thus, F is a union of members of each player's information partition, i.e., $F = \cup_k A_k$ and $F = \cup_\ell B_\ell$.

(3) For any k , take an $\omega \in A_k$, then we have $P_i(\omega) = A_k$. Since $F \subseteq E^{[\mu_i; \mu_j]}$, we have

$$\mu(E \mid A_k) = \mu(E \mid P_i(\omega)) = \mu_i.$$

(4) Therefore,

$$\begin{aligned} \mu(E \cap F) &= \mu(E \cap (\cup_k A_k)) = \sum_k \mu(E \cap A_k) && \sigma\text{-additivity} \\ &= \sum_k \mu(E \mid A_k) \cdot \mu(A_k) && \text{Conditional probability} \end{aligned}$$

$$\begin{aligned}
&= \mu_i \cdot \sum_k \mu(A_k) & \mu(E \mid A_k) &= \mu_i \\
&= \mu_i \cdot \mu(F) & \sigma\text{-additivity}
\end{aligned}$$

Similarly, $\mu(E \cap F) = \mu_j \cdot \mu(F)$. Consequently, $\mu_i = \mu_j$ (Here we assume that every state is of positive probability, and hence $\mu(F) > 0$). A contradiction. □

8.36 Remark: The intuition is that if a player knows that her opponents’ beliefs are different from her own, she should revise her beliefs to take the opponents’ information into account. Of course, this intuition does not make sense if the player thinks her opponents are simply crazy; it requires that she believes that her opponents process information correctly and that the difference in the beliefs reflects some objective information. More formally, Aumann’s result requires that the players’ beliefs be derived by Bayesian updating from a common prior.

8.37 Application: For a Bayesian game, an assumption often made is that the players have identical prior beliefs. This result implies that under common prior assumption, it can not be common knowledge that the players assign different posterior probabilities to the same event.

8.38 Corollary: In a two-player game with finite states, assume that there is a common prior belief. Let $f: \Omega \rightarrow \mathbb{R}$ be a random variable. Then it can not be common knowledge that $E_1[f \mid \omega] > E_2[f \mid \omega]$ at some state ω .

Proof. (1) Suppose that it is common knowledge that $E_1[f \mid \omega] > E_2[f \mid \omega]$.

(2) Let

$$\omega_0 \in F \subseteq \{\omega \in \Omega \mid E[f \mid P_1(\omega)] > E[f \mid P_2(\omega)]\},$$

where F is a self-evident event.

(3) Then for all $\omega \in F$, we have

$$\frac{p(\omega)}{p(F)} E[f \mid P_1(\omega)] > \frac{p(\omega)}{p(F)} E[f \mid P_2(\omega)].$$

(4) Thus,

$$E[f \mid F] = \sum_{\omega \in F} \frac{p(\omega)}{p(F)} E[f \mid P_i(\omega)] > \sum_{\omega \in F} \frac{p(\omega)}{p(F)} E[f \mid P_j(\omega)] = E[f \mid F],$$

which is a contradiction. □

8.39 Remark: It must be also true that, when two individuals have the same prior beliefs and common knowledge about their posteriors, they must have the same posterior expectation over random variables.

8.40 Remark: This result fails if the players merely know each other’s posteriors, as opposed to the posteriors’ being common knowledge.

Example: Ω has four equally likely elements, a, b, c, d , player 1’s partition is $\Pi_1 = \{\{a, b\}, \{c, d\}\}$, player 2’s partition is $\Pi_2 = \{\{a, b, c\}, \{d\}\}$. Let E be the event $\{a, d\}$. That at a , player 1’s posterior of E is $q_1(E) = p[\{a, d\} \mid \{a, b\}] = \frac{1}{2}$, and player 2’s posterior of E is $q_2(E) = p[\{a, d\} \mid \{a, b, c\}] = \frac{1}{3}$.

Moreover, player 1 knows that player 2’s information is the set $\{a, b, c\}$, so player 1 knows $q_2(E)$. Player 2 knows that player 1’s information is either $\{a, b\}$ or $\{c, d\}$, and either way player 1’s posterior of E is $\frac{1}{2}$, so player 2 knows $q_1(E)$. Thus, each player knows the other player’s posterior, yet the two players’ posteriors differ.

However, the posteriors are not common knowledge. In particular, player 2 does not know what player 1 thinks $q_2(E)$ is, as $\omega = c$, is consistent with player 2's information, and in this case player 1 believes there is probability $\frac{1}{2}$ that $q_2(E) = \frac{1}{3}$ (if $\omega = c$) and probability $\frac{1}{2}$ that $q_2(E) = 1$ (if $\omega = d$).

- 8.41 Example [OR Exercise 76.1]: Common knowledge and different beliefs. Show that if two individuals with partitional information functions have the same prior then it can be common knowledge between them that they assign different probabilities to some event. Show, however, that it cannot be common knowledge that the probability assigned by individual 1 exceeds that assigned by individual 2.
- 8.42 Example [OR Exercise 76.2]: Common knowledge and beliefs about lotteries. Show that if two individuals with partitional information functions have the same prior then it cannot be common knowledge between them that individual 1 believes the expectation of some lottery to exceed some number η while individual 2 believes this expectation to be less than η . Show by an example that this result depends on the assumption that the individuals' information functions are partitional.

8.5 No-trade theorem

- 8.43 Reference: [Milgrom and Stokey \(1982\)](#).

- 8.44 When the players' posteriors are consistent with a common prior, even though they may have different expectations over a random variable in some state, the knowledge that the both sides are willing to trade reveals extra information. Incorporating this extra information into the model destroys the incentive to bet.

Example: $\Omega = \{1, 2, 3, 4\}$, $p(i) = \frac{1}{4}$. Player 1's information partition $\Pi_1 = \{\{1, 4\}, \{2, 3\}\}$, and player 2's information partition $\Pi_2 = \{\{1, 2\}, \{3, 4\}\}$. When trade occurs, payoffs are as follows: $\pi_1 = \omega - 1.9$ and $\pi_2 = 1.9 - \omega$; otherwise, each player will receive 0.

Now the true state is $\omega = 2$. Consider player 1 firstly:

- (1) [Level 0] Player 1's expectation on ω is 2.5, and player 2's expectation on ω is 1.5. Hence, both players are willing to trade: players 1 and 2's expected payoffs are 0.6 and 0.4. But to actually to carry out this trade, the players must know that each other is willing to do the trade.
- (2) [Level 1 for player 2] Player 2 knows that player 1 is always willing to trade: in player 2's opinion, ω could be 1 or 2, but player 1's expected payoff is always 0.6 no matter ω is 1 or 2. Thus, the fact that player 1 indeed is willing does not initially tell player 2 anything.
- (3) [Level 1 for player 1] If player 1 knows that player 2 is willing to trade, then player 1 knows player 2 knows $\omega \in \{1, 2\}$; otherwise player 2 is not willing to trade.
Moreover, since player 1 knows $\omega \in \{2, 3\}$, player 1 knows $\omega = 2$. So player 1 still wants to do the deal.
- (4) [Level 2 for player 2] Since player 2 knows player 1 knows player 2 initially wants to trade, player 2 knows player 1 knows player 2 knows $\omega \in \{1, 2\}$. Hence player 2 knows player 1 knows $\omega \in \{1, 2\}$.
Player 2 also knows that player 1 can distinguish 1 and 2, so player 2 knows player 1 knows $\omega = 2$, and player 2 knows $\omega = 2$ at the same time.
So player 2 in the end refuse to trade.

As long as the fact that trade is acceptable is not common knowledge, revealing the fact changes the players' beliefs in a way that destroy the incentive to trade.

8.45 We consider the case of players with a utility function defined by $u_i(a(\omega), \omega)$ where $a: \Omega \rightarrow A$ is a contract that associates each state with an action or transfer.

8.46 Definition: We say that b is *ex ante* efficient if there is no function $a: \Omega \rightarrow A$ such that for all i

$$\mathbf{E}[u_i(a(\omega), \omega) \mid P_i(\omega)] > \mathbf{E}[u_i(b(\omega), \omega) \mid P_i(\omega)],$$

or equivalently, that $a \succ_i b$ for all i .

8.47 Theorem: If b is *ex ante* efficient, then it can not be common knowledge that $a \succ_i b$ for all i .

Proof. (1) For sake of notational simplicity, write $u_i(a(\omega), \omega) = U_i(a, \omega)$ and $u_i(b(\omega), \omega) = U_i(b, \omega)$.

(2) Assume that

$$\{\omega \in \Omega \mid \mathbf{E}[U_i(a, \omega) - U_i(b, \omega) \mid P_i(\omega)] > 0 \text{ for all } i\}$$

is common knowledge at ω_0 .

(3) Let F be a self-evident event containing ω_0 . Then

$$\mathbf{E}[U_i(a, \omega) - U_i(b, \omega) \mid F] > 0 \text{ for all } i.$$

(4) With this, we can create a contract $c: \Omega \rightarrow A$ such that

$$c(\omega) = \begin{cases} b(\omega), & \text{if } \omega \notin F, \\ a(\omega), & \text{if } \omega \in F. \end{cases}$$

(5) We have therefore that b could not have been *ex ante* efficient since $c \succ_i b$ for all i :

$$\begin{aligned} \mathbf{E}[U_i(c, \omega) - U_i(b, \omega)] &= \mathbf{E}[U_i(c, \omega) - U_i(b, \omega) \mid F] + \mathbf{E}[U_i(c, \omega) - U_i(b, \omega) \mid F^c] \\ &= \mathbf{E}[U_i(a, \omega) - U_i(b, \omega) \mid F] + 0 > 0, \end{aligned}$$

which is a contradiction. □

8.48 Remark: This theorem implies that, when the common prior assumption holds and a contract b is *ex ante* efficient, there can be no trade, even after the agents receive (potentially different) new information.

Notice that the theorem can also be weakened, using a weaker concept to efficiency, in which $b \succsim_i a$ for all i and for all a , with \succ holding for at least one i , for some a .

8.6 Speculation

8.49 Reference: [Harrison and Kreps \(1978\)](#).

8.50 Consider the following scenario:

- Time is discrete: $t = 1, 2, \dots$
- A continuum of investors, with unlimit wealth.

- A single asset in the economy, which pays dividend $d_t \in \{0, 1\}$ at the end of each period.
- Two types of investors $i \in \{1, 2\}$, with different beliefs Q_i over the distribution of the dividend. Both believe that it follows a Markov process:

$$Q_1 = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} \text{ and } Q_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

- There is common knowledge over Q_i .
- Players have utility $\sum_{t=1}^{\infty} \delta^t d_t$, where $\delta = \frac{3}{4}$.

8.51 Consider the expected present value of holding an asset forever $V^i(d)$, when the most recent dividend is d :

$$\begin{aligned} V^1(0) &= \frac{\delta}{1-\delta} \cdot \frac{1}{2} = \frac{3}{2}, & V^1(1) &= \frac{\delta}{1-\delta} \cdot \frac{1}{2} = \frac{3}{2}, \\ V^2(0) &= \frac{\delta}{1-\delta} \cdot 0 = 0, & V^2(1) &= \frac{\delta}{1-\delta} \cdot 1 = 3. \end{aligned}$$

8.52 Whenever there is disagreement over the value of the asset, there must be trade, as the investors that attribute to it a lower value will wish to sell to those that attribute to it a higher one. To calculate the resulting price, we can observe that:

$$p(d_t) = \delta \cdot \max_i \mathbf{E}_i[d_{t+1} + p(d_{t+1}) \mid d_t],$$

since,

- if the price is lower than the expression on the right-hand side, investors that believe it has a higher value will compete for it, increasing its price;
- if the price is higher than the expression on the right-hand side, investors wish to sell the asset and repurchase it the next period.

8.53 When $d_t = 1$, the type 2 investor believes that the value of the asset is the maximum possible, therefore her valuation must be what drives the asset price:

$$p(1) = \delta \cdot \max_i \mathbf{E}_i[d_{t+1} + p(d_{t+1}) \mid d_t = 1] = \delta \cdot \mathbf{E}_2[d_{t+1} + p(d_{t+1}) \mid d_t = 1] = \delta[1 + p(1)].$$

Then we have

$$p(1) = \frac{\delta}{1-\delta} = 3.$$

8.54 When $d_t = 0$, the type 2 investor believes that the value of the asset is forever 0, therefore the price must be driven by type 1's valuation:

$$p(0) = \delta \cdot \max_i \mathbf{E}_i[d_{t+1} + p(d_{t+1}) \mid d_t = 0] = \delta \cdot \mathbf{E}_1[d_{t+1} + p(d_{t+1}) \mid d_t = 0] = \delta \frac{0 + p(0)}{2} + \delta \frac{1 + p(1)}{2}.$$

Then we have

$$p(0) = \frac{12}{5}.$$

8.55 Observe that the price $p(0)$ is a price resulting from speculation:

- The type 1 investor is willing to pay so much for the asset (at time t , when $d_t = 0$) only because she believes the investors of type 2 will, with some probability, purchase it back for a higher value than it actually worth.

- Her actual valuation of the asset is of only $V^1(0) = \frac{3}{2} < \frac{12}{5} = p(0)$.

8.56 In this example with a common prior, we obtain speculative trade as a result.

8.7 Characterization of the common prior assumption

8.57 Example: Suppose $\Omega = \{1, 2, 3, 4\}$. There are two players with the information structure $\Pi_1 = \{\{1, 2\}, \{3, 4\}\}$ and $\Pi_2 = \{\{1, 3\}, \{2, 4\}\}$. Let Prob_1 and Prob_2 be the posterior beliefs of players 1 and 2, respectively. Denote $s = \text{Prob}_1(1 \mid \{1, 2\})$, $t = \text{Prob}_1(3 \mid \{3, 4\})$, $y = \text{Prob}_2(1 \mid \{1, 3\})$ and $z = \text{Prob}_2(2 \mid \{2, 4\})$.

It is easy to verify that there is a common prior if and only if

$$\frac{s}{1-s} \frac{1-t}{t} = \frac{y}{1-y} \frac{1-z}{z}.$$

Interpretation: Call $\{1, 2\}$ and $\{3, 4\}$ type 1 and type 2 of player 1. Call $\{1, 3\}$ and $\{2, 4\}$ type 1 and type 2 of player 2. Type 1 of player 1 believes that it is $\frac{s}{1-s}$ times more likely that player 2 is also type 1. If $\frac{s}{1-s} / \frac{t}{1-t} > 1$, then player 1 thinks that player 2 is more likely to be type 1 when he is type 1. The equation above means that the ratio of the likelihood ratio must be the same for the two players. If one player believes that the types of the players are positively correlated then the other player will believe the same.

8.58 Reference: [Samet \(1998\)](#).

8.59 Let $f: \Omega \rightarrow \mathbb{R}$ be a random variable with $f((a, a)) = f((b, b)) = 1$ and $f((a, b)) = f((b, a)) = -1$. There exists no common prior π such that

$$E_1(f \mid \omega) > 0 > E_2(f \mid \omega) \text{ for all } \omega \in \Omega.$$

On the other hand, it is easy to find p_1 and p_2 such that $E_1(f \mid \omega) > 0$ with respect to p_1 and $E_2(f \mid \omega) < 0$ with respect to p_2 , e.g., $p_1(\{(a, a)\}) = p_1(\{(b, b)\}) = 1/2$, and $p_2(\{(b, a)\}) = p_2(\{(a, b)\}) = 1/2$.

8.60 Theorem ([Samet \(1998\)](#)): In a two-player game with finite states, there is no common prior if and only if there exists a random variable $f: \Omega \rightarrow \mathbb{R}$ such that

$$E_1(f \mid \omega) > 0 > E_2(f \mid \omega) \text{ for all } \omega \in \Omega.$$

(The “if” part is due to Aumann)

In other words, under common prior it can never be common knowledge that 1’s expectation of f is always positive when that of 2 is always negative.

8.61 *Proof.* (1) There is no common prior, then X_1 and X_2 can be strongly separated, that is, there are $g: \Omega \rightarrow \mathbb{R}$ and $c \in \mathbb{R}$, such that

$$x_1 \cdot g > c > x_2 \cdot g$$

for each $x_1 \in X_1$ and $x_2 \in X_2$.

(2) Subtracting c and write $f = g - c\mathbf{1}$, we have

$$x_1 \cdot f > 0 > x_2 \cdot f.$$

(3) Thus, for each ω , take $x_1 = \mu_1(P_1(\omega))$, and $x_2 = \mu_2(P_2(\omega))$, then

$$\mathbf{E}_1(f \mid \omega) = \mu_1(P_1(\omega)) \cdot f > 0 > \mu_2(P_2(\omega)) \cdot f = \mathbf{E}_2(f \mid \omega).$$

□

8.62 Lemma: Let K_1, K_2, \dots, K_n be convex, closed subsets of the simplex Δ^m in \mathbb{R}^m . Then $\cap_{i=1}^n K_i = \emptyset$ if and only if there are f_1, f_2, \dots, f_n in \mathbb{R}^m , such that $\sum_{i=1}^n f_i = 0$, and $x_i f_i > 0$ for each $x_i \in K_i$, for $i = 1, 2, \dots, n$.

Proof. (1) Consider the bounded, closed and convex subsets of \mathbb{R}^{mn} , $X = \times_{i=1}^n K_i$ and

$$Y = \{(p, p, \dots, p) \in \mathbb{R}^{mn} \mid p \in \Delta^m\}.$$

(2) Clearly, $\cap_{i=1}^n K_i = \emptyset$ if and only if X and Y are disjoint

(3) X and Y are disjoint if and only if there is a constant c and $g = (g_1, g_2, \dots, g_n) \in \mathbb{R}^{mn}$, where $g_i \in \mathbb{R}^m$ for each i , such that for each $x = (x_1, x_2, \dots, x_n) \in X$ and $y = (p, p, \dots, p) \in Y$,

$$x \cdot g > c > y \cdot g.$$

(4) Moreover, we may assume that $c = 0$ (by subtracting $\frac{c}{n}$ from all the components of g).

(5) Hence, $\sum_{i=1}^n x_i \cdot g_i > 0$ and $\sum_{i=1}^n p \cdot g_i < 0$.

(6) Since $\sum_{i=1}^n p \cdot g_i$ holds for all $p \in \Delta^m$ and therefore it is equivalent to $\sum_{i=1}^n g_i < 0$.

(7) Moreover, whereas the coordinates of x_i are non-negative, increasing the coordinates of the g_i does not change the first inequality, and hence the intersection of K_i s is empty if and only if there is g such that

$$\sum_{i=1}^n g_i = 0, \text{ and } \sum_{i=1}^n x_i \cdot g_i > 0.$$

(8) Let \bar{x}_i be the point that minimizes $x_i \cdot g_i$ over K_i .

(9) Whereas $\sum_{i=1}^n \bar{x}_i \cdot g_i > 0$, there are constants c_i such that

$$\bar{x}_i \cdot g_i + c_i > 0 \text{ for } i = 1, 2, \dots, n, \text{ and } \sum_{i=1}^n c_i = 0.$$

(10) Denote by $\mathbf{1}$ the vector of 1s in \mathbb{R}^m and define $f_i = g_i + c_i \mathbf{1}$.

(11) Then $\sum f_i = \sum g_i = 0$ and for each $x_i \in K_i$, $x_i \cdot f_i \geq \bar{x}_i \cdot f_i = \bar{x}_i \cdot g_i + c_i \bar{x}_i \cdot \mathbf{1} > 0$, as $\bar{x}_i \cdot \mathbf{1} = 1$.

□

8.63 Theorem (Samet (1998), Morris 1995): There exists a common prior if and only if there are no $f_1, f_2, \dots, f_n \in \mathbb{R}^\Omega$, such that $\sum_{i=1}^n f_i = 0$, and $\mathbf{E}_i f_i > 0$ for all $i \in N$.

8.64 Note that two players can make a bet that both sides expect to win in the *ex-ante* stage if and only if they have different priors. The above theorem says that if the posterior beliefs of the two players are inconsistent with a common prior, then they can still make a bet that both sides expect to win in the interim stage. The key is that all types of players 1 and 2 agree to have player 2 pay f to player 1 in the proof above. Hence, the fact the players agree to the bet does not reveal extra information.

8.8 Unawareness

8.65 Reference: [Dekel, Lipman and Rustichini \(1998\)](#).

8.66 Unawareness is a real-life phenomenon associated with an unconscious mental state directed toward, or lacking of positive knowledge about, a definite event. Unawareness can play an important role in economic implications. For example, unforeseen contingencies could prevent contracting parties from writing a complete contract with contingencies of which they were unaware at the contractual date.

8.67 Unawareness of something is related to a complete lack of positive knowledge regarding it. In particular, “knowing that not knowing an event” can not be called “being unaware of the event”.

8.68 Example: Consider the following simple version of a story from *Sherlock Holmes*.

- There are 2 states b (dog barks) and \bar{b} (dog does not bark).
- At b , Watson is aware that there is an intruder;
- At \bar{b} , Watson is not aware this.

Question: Is Watson, at \bar{b} , “unaware” b ?

Somewhat surprisingly, the concept of “unawareness” can not be modeled in any standard state space!

8.69 We consider a standard state-space model

$$\langle \Omega, K, U \rangle,$$

where Ω is a state space which has partitional information structure, $K : 2^\Omega \rightarrow 2^\Omega$ is the knowledge operator, and $U : 2^\Omega \rightarrow 2^\Omega$ is the unawareness operator.

For an event E , $U(E)$ represents the event that decision-maker is unaware of E .

8.70 [Dekel, Lipman and Rustichini \(1998\)](#) suggested three axioms for unawareness:

- Plausibility: $U(E) \subseteq \neg K(E) \cap \neg K(\neg K(E))$.
- KU introspection: $K(U(E)) = \emptyset$.
- AU introspection: $U(E) \subseteq U(U(E))$.

Plausibility means: if decision-maker is unaware of an event, then it must be the case that the event is unknown and it is not known that the event is unknown.

KU introspection means: decision-maker shouldn't know he is unaware of the event.

AU introspection means: decision-maker should be unaware he is unaware of the event.

8.71 [Dekel, Lipman and Rustichini \(1998\)](#) showed that standard state-space models preclude sensible unawareness.

Theorem: Assume (Ω, K, U) is plausible and satisfies KU introspection and AU introspection. Then

- (1) If K satisfies necessitation, i.e., $K(\Omega) = \Omega$, then for every event E , $U(E) = \emptyset$.
- (2) If K satisfies monotonicity, i.e., $K(E) \subseteq K(F)$ whenever $E \subseteq F$, then for all events E and F , $U(E) \subseteq \neg K(F)$.

Statement (1) says that “necessitation” implies that decision-maker is never unaware of anything.

Statement (2) says that “monotonicity” implies that decision-maker, being unaware of anything, knows nothing.

8.72 *Proof.* (1) By AU introspection and plausibility,

$$U(E) \subseteq U(U(E)) \subseteq \neg K(\neg K(U(E))).$$

By necessitation, $U(E) = \emptyset$ for all E .

(2) By monotonicity, $K(F) \subseteq K(\Omega)$ for all F . Hence monotonicity implies $U(E) \subseteq \neg K(F)$ for all E and F .

□

Interactive epistemology

Contents

9.1	Epistemic conditions for Nash equilibrium	147
9.2	Epistemic foundation of rationalizability	150
9.3	Epistemic foundation of correlated equilibrium	151
9.4	The electronic mail game	151

In the traditional game theory, we appeal informally to assumptions about what the players know. We here use the model of knowledge to examine formally assumptions about the players' knowledge and information that lie behind various solution concepts. This line of research can help to understand the applicability and limitations of our analysis to economic phenomena.

This approach analyzes games in terms of the rationality of the players and their epistemic state: what they know or believe about each other's rationality, actions, knowledge, and beliefs. It provides precise treatments of epistemic matters in games.

The epistemic program adds to the traditional description of a game a mathematical language for talking about the rationality or irrationality of the players, their beliefs or knowledge, and related epistemic aspects.

9.1 Epistemic conditions for Nash equilibrium

9.1 Given a fixed strategic game $G = \langle N, (A_i), (u_i) \rangle$, a model of knowledge for game G is given by

$$\mathcal{M}(G) = \langle \Omega, (P_i), (a_i), (\mu_i) \rangle,$$

where

- Ω is the set of states;
- P_i is i 's partitional information function;
- $a_i(\omega)$ is i 's action at ω ;
- $\mu_i(\omega)$ is i 's belief at ω , which is a probability measure on $A_{-i} = \times_{j \neq i} A_j$.

9.2 Say “ i is rational at ω ” if $a_i(\omega)$ is a best response of player i to his belief $\mu_i(\omega)$ in $\Delta(a_{-i}(P_i(\omega)))$, where

$$a_{-i}(P_i(\omega)) = \{a_{-i}(\omega') : \omega' \in P_i(\omega)\}.$$

In words, i is rational at ω if i 's action at ω maximizes his expected payoff with respect to the belief that i holds at ω , where the belief is required to be consistent with his knowledge ($\text{support}(\mu_i(\omega)) \subseteq a_{-i}(P_i(\omega))$).

9.3 The event that i is rational is defined as

$$R_i = \{\omega : i \text{ is rational at } \omega\}.$$

Let $R = \bigcap_{i \in N} R_i$ denote the event that everyone is rational.

9.4 We now seek sufficient epistemic conditions for Nash equilibrium that are in a sense as “sparse” as possible.

Proposition 1: Let

$$\omega \in \bigcap_{i \in N} \left(R_i \cap K_i \left(\{\omega' \mid a_{-i}(\omega') = a_{-i}(\omega)\} \right) \right).$$

Then $a(\omega)$ is a Nash equilibrium.

9.5 The condition in the Proposition can be restated as follows:

- i is rational at ω : $a_i(\omega)$ is a best response of i to his belief $\mu_i(\omega)$, which is consistent with his knowledge: $\text{support}(\mu_i(\omega)) \subseteq a_{-i}(P_i(\omega))$, i.e., $\omega \in R_i$.
- i knows the other players' actions: $P_i(\omega) \subseteq \{\omega' \mid a_{-i}(\omega') = a_{-i}(\omega)\}$, i.e., $\omega \in K_i(\{\omega' \mid a_{-i}(\omega') = a_{-i}(\omega)\})$.

9.6 *Proof.* (1) Since $\omega \in K_i(\{\omega' \mid a_{-i}(\omega') = a_{-i}(\omega)\})$, we have $P_i(\omega) \subseteq \{\omega' \mid a_{-i}(\omega') = a_{-i}(\omega)\}$, and hence $a_{-i}(P_i(\omega)) = \{a_{-i}(\omega') \mid \omega' \in P_i(\omega)\} = \{a_{-i}(\omega)\}$.

(2) Since $\omega \in R_i$, $a_i(\omega)$ is a best response of i to $\mu_i(\omega) = 1 \circ a_{-i}(\omega)$.

□

9.7 Remark: Though very simple, this proposition is significant; it calls for “mutual knowledge” of the action choices, with no need for “common knowledge/any higher order knowledge”. For rationality, not even mutual knowledge is needed; only that the players are in fact rational.

9.8 A mixed strategy of a player can be interpreted as another player's conjecture about the player's choice. The second result provides some epistemic condition for this kind of equilibrium in beliefs: In two-person games, if the rationality of the players and their “consistent” conjectures are mutual knowledge, then the conjectures constitute a mixed-strategy Nash equilibrium.

Proposition 2: Suppose that each player's belief is consistent with his knowledge. Let

$$\omega \in \bigcap_{i,j=1,2;i \neq j} \left(K_j(R_i) \cap K_j(\{\omega' \mid \mu_i(\omega') = \mu_i(\omega)\}) \right).$$

Then $(\mu_2(\omega), \mu_1(\omega))$ is a mixed-strategy Nash equilibrium.

9.9 The condition in the Proposition can be restated as follows:

- Each player's belief is consistent with his knowledge: $\text{support}(\mu_j(\omega)) \subseteq a_i(P_j(\omega))$;

- j knows that i is rational;
- j knows i 's belief: $P_j(\omega) \subseteq \{\omega' \mid \mu_i(\omega') = \mu_i(\omega)\}$.

9.10 *Proof.* (1) Let $a_i^* \in \text{support}(\mu_j(\omega))$. Since j 's belief is consistent with his knowledge, there is $\omega^* \in P_j(\omega)$ such that $a_i(\omega^*) = a_i^*$.

(2) Since $\omega \in K_j(R_i)$, we have $\omega^* \in P_j(\omega) \subseteq R_i$, and hence $a_i^* = a_i(\omega^*)$ is a best response of i to $\mu_i(\omega^*)$.

(3) By $\omega \in K_j(\{\omega' \mid \mu_i(\omega') = \mu_i(\omega)\})$, we have $\omega^* \in P_j(\omega) \subseteq \{\omega' \mid \mu_i(\omega') = \mu_i(\omega)\}$, and hence $a_i^* = a_i(\omega^*)$ a best response to $\mu_i(\omega^*) = \mu_i(\omega)$.

□

9.11 The following example demonstrates that Proposition 2 does not have an analog when there are more than two players.

	L	R		L	R
U	2, 3, 0	2, 0, 0	U	0, 0, 0	0, 2, 0
D	0, 3, 0	0, 0, 0	D	3, 0, 0	3, 2, 0
	A			B	

State	α	β	γ	δ	ϵ	ξ
Probability	32/63	16/63	8/63	4/63	2/63	1/63
1's action	U	D	D	D	D	D
2's action	L	L	L	L	L	L
3's action	A	B	A	B	A	B
1's partition	$\{\alpha\}$	$\{\beta\}$	$\{\gamma\}$	$\{\delta\}$	$\{\epsilon\}$	$\{\xi\}$
2's partition	$\{\alpha\}$	$\{\beta\}$	$\{\gamma\}$	$\{\delta\}$	$\{\epsilon\}$	$\{\xi\}$
3's partition	$\{\alpha\}$	$\{\beta\}$	$\{\gamma\}$	$\{\delta\}$	$\{\epsilon\}$	$\{\xi\}$

Let the set of states be $\Omega = \{\alpha, \beta, \gamma, \delta, \epsilon, \xi\}$ and let the players' action functions and information functions be those given in the table at the bottom of the figure; assume that the players' beliefs are derived from the same prior, which is given in the first row of the table.

Consider the state δ . All conditions in Proposition are satisfied:

- Since each player's belief at δ is defined from the common prior, each player has a belief that is consistent with his knowledge.
- For player 1, She knows that the state is either δ or ϵ , so that she knows that player 2's information is either $\{\gamma, \delta\}$ or $\{\epsilon, \xi\}$. In both cases player 2 believes that with probability $\frac{2}{3}$ the pair of actions chosen by players 1 and 3 is (D, A) and that with probability $\frac{1}{3}$ it is (D, B) . Given this belief, the action L is optimal for player 2. Thus player 1 knows that player 2 is rational.

Player 2 knows that player 1's information is either $\{\beta, \gamma\}$ or $\{\delta, \epsilon\}$. In both cases player 1 believes that with probability $\frac{2}{3}$ players 2 and 3 will choose (L, B) and that with probability $\frac{1}{3}$ they will choose (L, A) . Given this belief, D is optimal for player 1. Thus, player 2 knows that player 1 is rational.

Player 3 knows that player 1's information is $\{\delta, \epsilon\}$ and that player 2's information is $\{\gamma, \delta\}$. Thus, as argued above, player 3 knows that players 1 and 2 are rational.

In the three states γ, δ and ϵ , player 3's belief is that the pair of actions of players 1 and 2 is (D, L) , and thus in the state δ players 1 and 2 know player 3's belief. They also know she is rational since her payoffs are always zero.

However in the state δ the beliefs do not define a Nash equilibrium. In fact, the beliefs at δ are not common knowledge, e.g., the players' belief about each other's behavior do not even coincide: player 1 believes that player 3 chooses A with probability $\frac{1}{3}$ while player 2 believes that she does so with probability $\frac{2}{3}$. Neither of these beliefs together with the actions D and L forms a mixed-strategy Nash equilibrium of the game.

9.12 Remark: [Aumann and Brandenburger \(1995\)](#) show that if all players share a common prior and in some state rationality is mutual knowledge and the players' beliefs are common knowledge then the beliefs at that state form a mixed-strategy Nash equilibrium even if there are more than two players. The key point is that if the beliefs of players 1 and 2 about player 3's action are common knowledge and if all the players share the same prior, then the beliefs must be the same.

9.2 Epistemic foundation of rationalizability

9.13 We show that the notion of rationalizability is the logical implication of CKR .

Proposition 3: Let $\omega \in CKR$. Then $a(\omega)$ is a rationalizable strategy profile.

9.14 Proof. (1) Since $\omega \in CKR$, there exists a self-evident event $F \subseteq R$ with $\omega \in F$.

(2) For each $i \in N$, define

$$Z_i = \{a_i(\omega') \mid \omega' \in F\}.$$

Therefore, for each $\omega' \in F$, $a_{-i}(P_i(\omega')) \subseteq Z_{-i}$.

(3) Since $\omega' \in R_i$, $a_i(\omega')$ is a best response for player i to the belief $\mu_i(\omega') \in \Delta(Z_{-i})$.

□

9.15 Proposition 4: Let a^* be a rationalizable action profile. Then there is a model of knowledge $\mathcal{M}(G)$ such that $a^* = a(\omega)$ for some $\omega \in CKR$.

9.16 Proof. (1) It suffices to show that there exists $\mathcal{M}(G)$ such that $a^* = a(\omega)$ for some ω in a self-evident event in R .

(2) Since a^* is a rationalizable profile, there exists a product subset Z of action profiles that contains a^* such that for each i , each $a_i \in Z_i$ is a best response to some belief $\mu_i(a_i) \in \Delta(Z_{-i})$.

(3) Define

$$\Omega = \left\{ \omega \mid \omega = (a_i, \mu_i(a_i))_{i \in N} \right\}.$$

For any $\omega = (a_i, \mu_i(a_i))_{i \in N}$ in Ω , for each i , let $a_i(\omega) = a_i$, $\mu_i(\omega) = \mu_i(a_i)$, and

$$P_i(\omega) = \{ \omega' \in \Omega \mid a_i(\omega') = a_i(\omega) \text{ and } \mu_i(\omega') = \mu_i(\omega) \}.$$

Clearly, $a_{-i}(P_i(\omega)) = Z_{-i}$. Thus, i is rational at every ω , and hence Ω is itself a self-evident event in R .

□

9.17 In finite strategic games, we have

$$CKR = R^* = \text{IENBR} = \text{IESDA}.$$

9.18 Example [OR Exercise 81.1]: Knowledge and correlated equilibrium.

9.3 Epistemic foundation of correlated equilibrium

9.4 The electronic mail game

9.19 Example (from <http://www.zhihu.com/question/29009853>): 想象两个将军准备合力攻打一座城，一个将军在城的西边，一个在东边。假设只有两个将军同时带领各自部队进攻的时候才能攻陷城池。再假设他们事先没约好什么时候进攻，且两个将军都是极度风险厌恶者，只有百分之百确认对方会进攻时自己才会进攻。假设当时是冷兵器时代，且两个将军的唯一通讯方式是派个通信兵穿越敌人城池来到实现。这场战役会怎么收场呢？答案是两个将军根本不会进攻……

为啥呢？直觉是这样的：

城东的将军决定次日早晨进攻，他派了一个通信兵去城西通知城西的将军。此时城东的将军不会在次日早晨进攻，因为他不知城西的将军有没有收到通知（因为通信兵穿越敌人城池的时候可能被抓住）。假设城西的将军收到了通知，他派那个通信兵回去通知城东的将军自己收到通知了。此时城西的将军不会在次日早晨进攻，因为他不知道城东的将军知不知道他已经收到了通知。继续这样推理，只有当每个将军知道对方知道对方知道对方知道……（无限个对方知道）次日早晨进攻的消息才会进攻，这意味着通信兵要跑无穷次。

	<i>A</i>	<i>B</i>		<i>A</i>	<i>B</i>
<i>A</i>	M, M	$1, -L$		$0, 0$	$1, -L$
<i>B</i>	$-L, 1$	$0, 0$		$-L, 1$	M, M
G_a (probability $1 - p$)				G_b (probability p)	

Figure 9.1: The parameters satisfy $L > M > 1$ and $p < \frac{1}{2}$.

- 9.20
- In G_a (resp. G_b), the players get a positive payoff, M , if both choose the action A (resp. B).
 - If they choose the same action but it is the “wrong” one they get 0. If they fail to coordinate, then the player who played B gets $-L$, where $L > M$. Thus, it is dangerous for a player to play B unless he is confident enough that his partner is going to play B as well.
 - G_a is more likely to occur, and G_b occurs with probability $p < \frac{1}{2}$.

9.21 If the true game is common knowledge between two players, then it has a Nash equilibrium in which each player chooses A in G_a and B in G_b .

If the true game is known initially only to player 1, but not to player 2. we can model this situation as a Bayesian game that has a unique Bayesian Nash equilibrium, in which player 1 chooses A in G_a and G_b , and player 2 chooses A .

- 9.22
- The true game is known initially only to player 1, but not to player 2.
 - Player 1 can communicate with player 2 via computers if the game is G_b . There is a small probability $\epsilon > 0$ that any given message does not arrive at its intended destination, however. (If a computer receives a message then it automatically sends a confirmation; this is so not only for the original message but also for the confirmation, the confirmation of the confirmation, and so on)
 - If a message does not arrive then the communication stops.
 - At the end of communication, each player’s screen displays the number of messages that his machine has sent.




Figure 9.2

9.23 Model as a Bayesian game. Consider Figure 9.2:

Define the set of states to be

$$\Omega = \{(Q_1, Q_2) \in \mathbb{Z}_+ \times \mathbb{Z}_+ : Q_1 = Q_2 \text{ or } Q_1 = Q_2 + 1\}.$$

In state (q, q) , player 1's computer sends q messages, all of which arrive at player 2's computer, and the q th message sent by player 2's computer goes astray.

In state $(q, q + 1)$, player 1's computer sends $q + 1$ messages, and all but the last arrive at player 2's computer.

The signal function τ_i of player i is defined by

$$\tau_i(Q_1, Q_2) = Q_i.$$

Player 1's information partition is

$$\Pi_1 = \left\{ \{(0, 0)\}, \{(1, 0), (1, 1)\}, \{(2, 1), (2, 2)\}, \dots, \{(q, q-1), (q, q)\}, \dots \right\}.$$

Player 2's information partition is

$$\Pi_2 = \left\{ \{(0, 0), (1, 0)\}, \{(1, 1), (2, 1)\}, \{(2, 2), (3, 2)\}, \dots, \{(q, q), (q, q+1)\}, \dots \right\}.$$

Each player's belief on Ω is the same, derived from the technology (characterized by ϵ) and the assumption that the game is G_a with probability $1 - p$:

$$p(0, 0) = 1 - p, \quad p(q + 1, q) = p\epsilon(1 - \epsilon)^{2q}, \quad p(q + 1, q + 1) = p\epsilon(1 - \epsilon)^{2q+1}.$$

Denote by $G(Q_1, Q_2)$ the game that is played in the state (Q_1, Q_2) ; that is, $G(0, 0) = G_a$ and $G(Q_1, Q_2) = G_b$ otherwise. In each state (Q_1, Q_2) , the payoffs are determined by the game $G(Q_1, Q_2)$.

9.24 Question: If ϵ is small then with high probability each player sees a very high number on his screen. When player 1 sees "1" on her screen, she is not sure whether player 2 knows that the game is G_b , and consequently may hesitate to play B . But if the number on her screen is, for example, "17" then it seems to be "almost" common knowledge that the game is G_b , and thus it may seem that she will adhere to the more desirable equilibrium (B, B) of the game G_b .

9.25 Proposition: This game has a unique Bayesian Nash equilibrium, in which both players always choose A .

9.26 Proof. We shall prove it by induction.

- (1) In any equilibrium, player 1 must choose A when receiving the signal 0.

When player 2's signal is 0, if he chooses A , then his expected payoff is at least $\frac{M(1-p)}{1-p+p\epsilon}$; if he chooses B then his expected payoff is at most $\frac{-L(1-p)+Mp\epsilon}{1-p+p\epsilon}$. Therefore player 2 must also choose A when receiving the signal 0 in any equilibrium.

- (2) Assume inductively that when received signal is less than q , players 1 and 2 both choose A in any equilibrium.

Consider player 1's decision when receiving the signal q . Player 1 believes $(q, q-1)$ with probability $z = \frac{p\epsilon(1-\epsilon)^{2q-2}}{p\epsilon(1-\epsilon)^{2q-2}+p\epsilon(1-\epsilon)^{2q-1}} = \frac{1}{2-\epsilon} > \frac{1}{2}$ and (q, q) with probability $1-z = \frac{1-\epsilon}{2-\epsilon}$.

If player 1 chooses B , then her expected payoff is at most $-Lz + M(1-z)$ under the induction assumption; if player 1 chooses A , then her expected payoff is at least 0. Thus, player 1 should choose A .

Similarly, player 2 chooses A when receiving the signal q .

□

9.27 Rubinstein's electronic mail game tells that players' strategic behavior under "almost common knowledge" may be very different from that under common knowledge. Even if both players know that the game is G_b and the noise ϵ is arbitrarily small, the players act as if they had no information and play A , as they do in the absence of an electronic mail system.

Rubinstein's electronic mail game also tells us how even an extreme large iteration of mutual knowledge can fall short of actual common knowledge.

Chapter 10

Extensive games with perfect information

Contents

10.1 Extensive games with perfect information	155
10.2 Subgame perfect equilibrium	157
10.3 Examples	160
10.4 Three notable games	167
10.5 Iterated elimination of weakly dominated strategies	169
10.6 Forward induction	170

We investigate games with perfect information in which each player is perfectly informed about the players' previous actions at each point in the game. The standard solution concept for such games is the notion of “subgame perfect equilibrium” in which each player is required to reassess his plans as play proceeds.

10.1 Extensive games with perfect information

10.1 Definition: An extensive game with perfect information is defined as:

$$\Gamma = \langle N, H, P, (\succsim_i) \rangle.$$

- A set N of players.
- A set H of sequences that satisfies the following three properties.
 - The empty sequence \emptyset is a member of H .
 - If $(a^k)_{k=1}^K \in H$ (K may be infinite) and $L < K$ then $(a^k)_{k=1}^L \in H$.
 - If an infinite sequence $(a^k)_{k=1}^\infty$ satisfies $(a^k)_{k=1}^L \in H$ for every positive integer L then $(a^k)_{k=1}^\infty \in H$.

Each member of H is a history; each component of a history is an action.

A history $(a^k)_{k=1}^K \in H$ is terminal if it is infinite or if there is no a^{K+1} such that $(a^k)_{k=1}^{K+1} \in H$. The set of terminal histories is denoted Z .

- A function $P: H \setminus Z \rightarrow N$ that assigns to each non-terminal history a member of N .
 P is called the player function, and $P(h)$ is the player who takes an action after the history h .

- For each player i a preference relation on Z .

10.2 The game tree is a convenient representation for a extensive game.

10.3 After any non-terminal history h player $P(h)$ chooses an action from

$$A(h) = \{a: (h, a) \in H\}.$$

10.4 If the set H of possible histories is finite then the game is finite. If the length of every history is finite then the game has a finite horizon. For an extensive game Γ denote by $\ell(\Gamma)$ the length of the longest history in Γ ; we refer to $\ell(\Gamma)$ as the length of Γ .

10.5 Definition: A strategy s_i of player i in the extensive game Γ is a function that assigns an action in $A(h)$ to each non-terminal history $h \in H \setminus Z$ for which $P(h) = i$.

A strategy specifies the action chosen by a player for every history after which it is his turn to move, even for histories that, if the strategy is followed, are never reached.

10.6 For each strategy profile $s = (s_i)$ in the extensive game Γ , we define the outcome $O(s)$ of s to be the terminal history that results when each player i follows the precepts of s_i .

10.7 Definition: A Nash equilibrium of $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* such that for every player i we have

$$O(s_{-i}^*, s_i^*) \succsim_i O(s_{-i}^*, s_i) \text{ for every strategy } s_i \text{ of player } i.$$

10.8 Proposition: s^* is a Nash equilibrium of Γ if and only if it is a Nash equilibrium of the strategic game derived from Γ .

10.9 Example: Consider the following extensive game Γ .




Figure 10.1: Non-credible threat.

The strategic game derived from Γ is as follows:

- 2 players.
- $A_1 = \{L, R\}$, $A_2 = \{L', R'\}$.
- Payoffs are as follows:

		Player 2	
		L'	R'
Player 1	L	0, 0	2, 1
	R	1, 2	1, 2

There are two Nash equilibria: (L, R') and (R, L') .

10.2 Subgame perfect equilibrium

10.10 In the example above, consider the Nash equilibrium (R, L') , L' is not credible for player 2 since R' is strictly better than L' for him.

10.11 Nash equilibrium requires that each player's strategy be optimal, given the other players' strategies. Of course, what a strategy calls for at a decision node that is not reached can not matter for a player's payoff; the action assigned to a contingency matters only if one is called upon to implement it.

Nash equilibrium does not require that a strategy prescribe an optimal action for decision nodes not reached (unreached information sets) during the course of equilibrium play.

Thus, a Nash equilibrium does not require that the prescribed action be optimal for all contingencies, but rather only for those reached over the course of equilibrium play (*i.e.*, the sequence of play that occurs when players use their equilibrium strategies).

10.12 Definition: For history $h \in H$, the subgame $\Gamma(h)$ is defined as

$$\langle N, H|_h, P|_h, (\succsim_i|_h) \rangle,$$

where

- $H|_h$ is the set of sequences h' of actions for which $(h, h') \in H$.
- $P|_h$ is defined by $P|_h(h') = P(h, h')$ for each $h' \in H|_h$.
- $\succsim_i|_h$ is defined by $h' \succsim_i|_h h''$ if and only if $(h, h') \succsim_i (h, h'')$.

10.13 Definition: A subgame perfect equilibrium of Γ is a strategy profile s^* such that for every subgame $\Gamma(h)$ with $P(h) = i$ we have

$$O_h(s_{-i}^*|_h, s_i^*|_h) \succsim_i|_h O_h(s_{-i}^*|_h, s_i|_h) \text{ for every strategy } s_i \text{ of player } i,$$

where $s_i|_h$ is i 's strategy restricted to $\Gamma(h)$ and O_h is the outcome function of $\Gamma(h)$.

10.14 Proposition: s^* is a subgame perfect equilibrium if and only if $s^*|_h$ is a Nash equilibrium in every $\Gamma(h)$.

10.15 Definition (The one deviation property): The profile s^* is said to satisfy the one deviation property if for every i and every $h \in H$ with $P(h) = i$,

$$O_h(s_{-i}^*|_h, s_i^*|_h) \succsim_i|_h O_h(s_{-i}^*|_h, t_i)$$

for every strategy t_i of player i in $\Gamma(h)$ that differs from $s_i^*|_h$ only in the action it prescribes after the initial history of $\Gamma(h)$.

10.16 Interpretation of the one deviation property: For each subgame the player who makes the first move can not obtain a better outcome by changing only his initial action.

10.17 Theorem: In a finite-horizon game Γ , the profiles s^* is a subgame perfect equilibrium if and only if s^* satisfies the one deviation property.

Actually, in any perfect-information extensive game with either finite horizon or discounting, a strategy profile is a subgame perfect equilibrium if and only if it satisfies the one deviation property.

10.18 Idea of proof: If a strategy profile is not subgame perfect then some player can deviate and obtain a strictly higher payoff, say by $\epsilon > 0$, in some subgame. Now look at the last deviation. If it makes the player better off, then the strategy does not satisfy the one deviation property. If it does not makes the player better off, then the player will still be better off without the last deviation. The same argument can be repeated until we find a single beneficial deviation.

10.19 *Proof.* “ \Rightarrow ”: Trivial.

“ \Leftarrow ”:

- (1) Suppose that s^* is not a subgame perfect equilibrium. Then i can deviate profitably in $\Gamma(h')$.
- (2) Then there exists a profitable deviant strategy s_i of player i in $\Gamma(h')$ for which $s_i(h) \neq s_i^*|_{h'}(h)$ for a number of histories h not larger than the length of $\Gamma(h')$; since Γ has a finite horizon this number is finite.
- (3) From among all the profitable deviations of player i in $\Gamma(h')$ choose a strategy s_i for which the number of histories h such that $s_i(h) \neq s_i^*|_{h'}(h)$ is minimal.
- (4) Let h^* be the longest history h of $\Gamma(h')$ for which $s_i(h) \neq s_i^*|_{h'}(h)$.
- (5) Then the initial history of $\Gamma(h', h^*)$ is the only history in $\Gamma(h', h^*)$ at which the action prescribed by s_i differs from that prescribed by $s_i^*|_{h'}$.
- (6) Further, $s_i|_{h^*}$ is a profitable deviation in $\Gamma(h', h^*)$, since otherwise there would be a profitable deviation in $\Gamma(h')$ that differs from $s_i^*|_{h'}$ after fewer histories than does s_i (contradicts to the choice of s_i), i.e., without this deviation, e.g. $(s_i|_{h'-h^*}, s_i^*|_{h^*})$.
- (7) Thus $s_i|_{(h', h^*)}$ is a profitable deviation in $\Gamma(h', h^*)$ that differs from $s_i^*|_{(h', h^*)}$ only in the action that it prescribes after the initial history of $\Gamma(h', h^*)$.

□

10.20 For extensive games with infinite horizon, a strategy profile may not be a subgame perfect equilibrium although it satisfies the one deviation property.

Example: In the following one-player game, the strategy in which the player chooses S after every history satisfies the one deviation property, but is not a subgame perfect equilibrium.




Figure 10.2: In infinite-horizon games, one deviation property \nRightarrow subgame perfect equilibrium.

10.21 Note that Theorem 10.17 only works for subgame perfect equilibrium in games with perfect information. It is not true for Nash equilibrium, and it is not true for subgame perfect equilibrium in games with imperfect information.

✂ 10.22 Theorem (Kuhn's Theorem): Every finite extensive game with perfect information has a subgame perfect equilibrium.

Proof. Let $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ be a finite extensive game with perfect information. We construct a subgame perfect equilibrium of Γ by induction on $\ell(\Gamma(h))$. We also define a function R that associates a terminal history with every history $h \in H$ and show that this history is a subgame perfect equilibrium outcome of the subgame $\Gamma(h)$.

- (1) If $\ell(\Gamma(h)) = 0$, i.e., h is a terminal history of Γ , define $R(h) = h$. Since there is no feasible action, $R(h) = h$ is a subgame perfect equilibrium outcome in $\Gamma(h)$.
- (2) Suppose that $R(h)$ is defined for all $h \in H$ with $\ell(\Gamma(h)) \leq k$ for some $k \geq 0$. Let h^* be a history for which $\ell(\Gamma(h^*)) = k + 1$ and let $P(h^*) = i$. Since $\ell(\Gamma(h^*)) = k + 1$ we have $\ell(\Gamma(h^*, a)) \leq k$ for all $a \in A(h^*)$. Define

$$s_i(h^*) \in \arg \max_{a \in A(h^*)} R(h^*, a),$$

and define $R(h^*) = R(h^*, s_i(h^*))$. It is clear that $R(h^*)$ is a subgame perfect equilibrium outcome in $\Gamma(h^*)$.

By induction we have now defined a strategy profile s in Γ ; by the one deviation property, this strategy profile is a subgame perfect equilibrium of Γ . \square

10.23 The procedure used in proof of Kuhn's theorem is often referred to as backwards induction. Backwards induction will eliminate the Nash equilibria that rely on non-credible threats or promises.

10.24 Example: Backwards induction.




Figure 10.3: The procedure of backwards induction.

- (1) There are two non-terminal histories: \emptyset and L .
- (2) For L , it is player 2's turn to move. If he chooses L' , he will get 0; otherwise he will get 1. Then $R(L) = R'$.
- (3) For \emptyset , it is player 1's turn to move, and he has 2 choices: L and R . If he chooses R , he will get 1; otherwise, the subgame $\Gamma(L)$ is reached, and the equilibrium outcome therein is 2 for player 1. Thus, he will choose L .
- (4) Now we have a strategy profile: player 1 chooses L and player 2 chooses R' at histories \emptyset and L respectively, which is a subgame perfect equilibrium.

Note that we can not obtain the other Nash equilibrium (R, L') in backwards induction.

10.25 For games with finite horizon, Kuhn's theorem does not necessarily hold.

Consider the one-player game in which the player chooses a number in the interval $[0, 1)$, and prefers larger numbers to smaller ones. This game has a finite horizon (the length of the longest history is 1) but has no subgame perfect equilibrium, since $[0, 1)$ has no maximal element.

10.26 For infinite-horizon games with the requirement that after any history each player have finitely many possible actions, Kuhn's theorem does not necessarily hold.

In the infinite-horizon one-player game the beginning of which is shown in Figure 10.4 the single player chooses between two actions after every history. After any history of length k the player can choose to stop and obtain a payoff of $k + 1$ or to continue; the payoff if she continues for ever is 0. The game has no subgame perfect equilibrium.

10.27 Kuhn's theorem makes no claim of uniqueness.




Figure 10.4

10.28 Proposition: Say that a finite extensive game with perfect information satisfies the no indifferent condition if

$$z \sim_j z' \text{ for all } j \in N \text{ whenever } z \sim_i z' \text{ for some } i \in N,$$

where z and z' are terminal histories. Then, by induction on the length of subgames, every player is indifferent among all subgame perfect equilibrium outcomes of such a game. Furthermore, if s and s' are subgame perfect equilibria then so is s'' , where for each player i the strategy s''_i is equal to either s_i or s'_i , that is, the equilibria of game are interchangeable.

10.3 Examples

10.29 Example: Stackelberg model of duopoly.

Consider the Stackelberg model of duopoly where two firms produce a homogeneous product. The price for the product is $P(Q) = a - Q$ if $Q \leq a$ and 0 otherwise, where $Q = q_1 + q_2$ and q_i is the output level of firm i . The cost function of firm i is $c_i q_i$. Due to the restriction of technology, firm 1 can produce either q_{1h} or q_{1l} , where $q_{1h} \geq q_{1l} \geq 0$. Firm 2 can produce any quantity $q_2 \geq 0$. Assume $a - q_{1h} > \max\{c_2, 2c_1 - c_2\}$.

The game takes place in two stages:

- Firm 1 chooses a quantity $q_1 \in \{q_{1h}, q_{1l}\}$.
- Firm 2 observes q_1 and then chooses a quantity $q_2 \geq 0$.

For $i = 1, 2$, the payoff to firm i is given by

$$\pi_i(q_1, q_2, x) = q_i P(Q) - c_i q_i.$$

Denote $\xi = a - q_{1h} - q_{1l} - 2c_1 + c_2$.

- (i) Find the backwards induction outcomes for $\xi > 0$ and $\xi < 0$ respectively.
- (ii) For $\xi > 0$, find all the subgame perfect equilibria.
- (iii) For $\xi > 0$, find a Nash equilibrium in which firm 1's strategy is different from its strategy in the subgame perfect equilibrium.

Answer. □

10.30 Example: Each of two firms, A and B , will choose a number between 0 and 1 which represents the "location". Let x_A and x_B be the numbers chosen by A and B , respectively. The payoff function for firm A is given by

$$u_A(x_A, x_B) = \begin{cases} \frac{x_A + x_B}{2}, & \text{if } x_A \leq x_B, \\ 1 - \frac{x_A + x_B}{2}, & \text{if } x_A > x_B, \end{cases}$$

where x_A is the number chosen by A and x_B is chosen by B . The payoff function for B is $u_B(x_A, x_B) = 1 - u_A(x_A, x_B)$. Assume that A chooses the number first, and B , after observing x_A , chooses x_B . Formulate the problem as an extensive game, and find all pure-strategy subgame perfect equilibria and all pure-strategy Nash equilibria.

Answer. (i) Given x_A ,

$$u_B(x_A, x_B) = \begin{cases} 1 - \frac{x_A + x_B}{2}, & \text{if } x_A \leq x_B, \\ \frac{x_A + x_B}{2}, & \text{if } x_A > x_B. \end{cases}$$

$1 - \frac{x_A + x_B}{2}$ achieves the maximal $1 - x_A$ when $x_B = x_A$, and $\frac{x_A + x_B}{2}$ has the supremum x_A when x_B approaches to x_A .

Therefore,

- When $x_A \leq \frac{1}{2}$, $1 - x_A \geq x_A$, and hence u_B will achieve the maximal $1 - x_A$. In this case $u_A = x_A \leq \frac{1}{2}$.
- When $x_A > \frac{1}{2}$, $x_A > 1 - x_A$, and hence there is no best response for player B . In this case $u_A = 1 - x_A < \frac{1}{2}$.

Thus, player A will choose $\frac{1}{2}$ due to backwards induction. However, there is no subgame perfect equilibrium since there is no best response for player B when he faces $x_A > \frac{1}{2}$.

(ii)

□

10.31 Example: There are n lions in a clearing in the jungle, along with one dead lamb, and the lions are ranked from L_1 (highest) to L_n (lowest). The lions move sequentially, in order of rank, and they can choose to eat or not to eat. They are hungry (payoff: 0), and therefore prefer to eat (payoff: 1), but they are also cautious; they will not eat if eating will lead to their death (payoff: -1). The lions have reason to be fearful, because they are narcoleptic, cannibalistic, and cowardly: if they eat, they fall asleep immediately, at which time they will be prey to the next lion in the sequence, who will eat only sleeping lions. Finally, the lions are finicky, so they will eat only recently dead, or newly asleep, meat—they will not eat meat (*i.e.* sleeping lions or the dead lamb) that has been passed over by others. In other words, if the dead lamb is not eaten by the first lion L_1 , then no lion will choose to eat this dead lamb; if a sleeping lion L_i ($i = 1, 2, \dots, n-1$) is not eaten by the next lion L_{i+1} , then no lion will choose to eat this sleeping lion L_i .

Represent this problem as an extensive game. What is the subgame perfect outcome if there are six lions in the pride? What if there are seven lions in the pride?

Answer. (i) When there are six lions, the game tree is as follows:

By backwards induction, the subgame perfect outcome is “ L_1 does not eat”.

(ii) When there are seven lions, the game tree is as follows:

By backwards induction, the subgame perfect outcome is “ L_1 eats, and L_2 does not eat”.

□

10.32 Example: Splitting four coins/Ultimatum with a finite number of alternatives.

Players 1 and 2 are bargaining over how to split 4 coins. Player 1 proposes to take s_1 coins (s_1 should be an integer), leaving $(4 - s_1)$ coins for player 2. Then player 2 either accepts or rejects the offer. If player 2 accepts the offer, then the payoffs are s_1 coins to player 1, and $(4 - s_1)$ coins to player 2. If player 2 rejects the offer, then the payoffs are zero to both.

1	E	2	E	3	E	4	E	5	E	6	E	
	N		N		N		N		N		N	-1
												-1
												-1
												-1
0		1		-1		-1		-1		-1		-1
0		0		1		-1		-1		-1		-1
0		0		0		1		-1		-1		1
0		0		0		0		1		-1		
0		0		0		0		0		1		
0		0		0		0		0		0		

Figure 10.5

1	E	2	E	3	E	4	E	5	E	6	E	7	E	
	N		N		N		N		N		N		N	-1
														-1
														-1
														-1
0		1		-1		-1		-1		-1		-1		-1
0		0		1		-1		-1		-1		-1		-1
0		0		0		1		-1		-1		-1		-1
0		0		0		0		1		-1		-1		1
0		0		0		0		0		1		-1		
0		0		0		0		0		0		1		
0		0		0		0		0		0		0		

Figure 10.6

- (i) Find all the pure-strategy Nash equilibria.
(ii) Find all the pure-strategy subgame perfect equilibria.

Answer. Figure 10.7 is the game tree.




Figure 10.7

It is easy to see that player 1's strategy space is $S_1 = \{0, 1, 2, 3, 4\}$. Since a strategy is a complete plan of actions in every contingency when a player is called upon to make, a strategy for player 2 can be represented as a function

$$f: S_1 \rightarrow \{A, R\}.$$

For example,

$$f(s_1) = \begin{cases} A, & \text{if } s_1 = 0, 2, 4; \\ R, & \text{if } s_1 = 1, 3 \end{cases}$$

is a strategy of player 2 in which player 2 will accept if player 1 offers 0, 2 and 4, and otherwise she will reject.

Thus, the space of all strategies of player 2 is the set of all functions from S_1 to $\{A, R\}$. We denote it by S_2 .

- (i) (1) Player 1's best-response correspondence: Given a strategy f of player 2, note that for any $s_1 \in f^{-1}(A)$, player 2 will accept the offer. Hence, given f , player 1 will choose the maximum in $f^{-1}(A)$. Since $f^{-1}(A)$ is a subset of S_1 , the maximal always exists. Thus, player 1's best-response correspondence is

$$B_1^*(f) = \begin{cases} S_1, & \text{if } f^{-1}(A) = \emptyset; \\ S_1, & \text{if } 0 \text{ is the maximum of } f^{-1}(A); \\ \{s^*\}, & \text{if } f^{-1}(A) \text{ has a maximum } s^* \neq 0. \end{cases}$$

- (2) Player 2's best-response correspondence: note that player 2's strategy is a function

$$B_2^*(s_1) = \begin{cases} \{f \in S_2 : f(s_1) = A\}, & \text{if } s_1 < 4; \\ S_2, & \text{if } s_1 = 4. \end{cases}$$

That means for any $s_1 < 4$, player 2 will accept. If $s_1 = 4$, player 2 is indifferent between the two actions (accept or reject).

- (3) We can use various combinations of the conditions in the expression of B_1^* and B_2^* to construct all the Nash equilibria:

- When $f^{*-1}(A) \supsetneq \{0\}$, (s_1^*, f^*) is a Nash equilibrium if and only if $s_1^* = \max f^{*-1}(A)$;
- When $f^{*-1}(A) = \{0\}$, (s_1^*, f^*) is a Nash equilibrium if and only if $s_1^* = 0$ or 4;
- When $f^{*-1}(A) = \emptyset$, (s_1^*, f^*) is a Nash equilibrium if and only if $s_1^* = 4$.

- (ii) For each given s_1 , we need to consider a corresponding subgame, displayed in Figure 10.8. We know if f^* is




Figure 10.8

subgame perfect, $f^*(s_1) = A$ for any $s_1 < 4$. Hence, if (s_1^*, f^*) is subgame perfect, f^* should be either f_1^* or f_2^* :

$$f_1^*(s_1) = \begin{cases} A, & \text{if } s_1 = 0, 1, 2, 3; \\ R, & \text{if } s_1 = 4. \end{cases} \quad \text{or } f_2^*(s_1) \equiv A \text{ for all } s_1.$$

It is easy to check that there are 2 subgame perfect equilibria: $(s_1^* = 3, f_1^*)$ and $(s_1^* = 4, f_2^*)$.

□

Players 1 and 2 are bargaining over how to split a dollar. Player 1 proposes to take a share s_1 of the dollar, leaving $(1 - s_1)$ for player 2. The share s_1 can be any real number in the interval $[0, 1]$. Then, player 2 either accepts or rejects the offer. If player 2 accepts the offer, the payoffs are $\alpha s_1 + (1 - \alpha)(2s_1 - 1)$ to player 1, and $\beta(1 - s_1) + (1 - \beta)(1 - 2s_1)$ to player 2. If player 2 rejects the offer, the payoffs are zero to both. In this game, they may care about the difference between the shares as well as his/her own share.

- (i) Suppose $\alpha = \beta = \frac{1}{2}$. Find all pure-strategy subgame perfect equilibria.
- (ii) Suppose that α and β are real numbers in the interval $[0, 1]$ and that $\alpha + \beta > 0$. Find a pure-strategy subgame perfect equilibrium.

Answer. Figure 10.9 is the game tree.




Figure 10.9

- (i) The payoffs when player 2 accepts the offer are $\frac{3}{2}s_1 - \frac{1}{2}$ to player 1 and $1 - \frac{3}{2}s_1$ to player 2. Clearly, player 2's payoff $1 - \frac{3}{2}s_1 \geq 0$ if and only if $s_1 \leq \frac{2}{3}$. In order to be a subgame perfect equilibrium, player 2's strategy must be as follows:

$$s_2^* = \begin{cases} \text{accept} & \text{if } s_1 \leq \frac{2}{3} \\ \text{reject} & \text{if } s_1 > \frac{2}{3} \end{cases}, \quad s_2^{**} = \begin{cases} \text{accept} & \text{if } s_1 < \frac{2}{3} \\ \text{reject} & \text{if } s_1 \geq \frac{2}{3} \end{cases}.$$

Note that, if player 2 uses the strategy s_2^{**} , then there is no best choice for player 1. Therefore, the strategy profile yielding the backwards induction outcome can only be the following subgame perfect equilibrium:

$$(s_1^* = \frac{2}{3}, s_2^*).$$

- (ii) The payoffs when player 2 accepts the offer are $\alpha s_1 + (1 - \alpha)(2s_1 - 1)$ to player 1 and $\beta(1 - s_1) + (1 - \beta)(1 - 2s_1)$ to player 2. It easy to see that player 2's payoff $\beta(1 - s_1) + (1 - \beta)(1 - 2s_1) \geq 0$ if and only if $s_1 \leq \frac{1}{2 - \beta}$. In order to be a subgame perfect equilibrium, player 2's strategy must be as follows:

$$s_2^* = \begin{cases} \text{accept} & \text{if } s_1 \leq \frac{1}{2 - \beta} \\ \text{reject} & \text{if } s_1 > \frac{1}{2 - \beta} \end{cases}, \quad s_2^{**} = \begin{cases} \text{accept} & \text{if } s_1 < \frac{1}{2 - \beta} \\ \text{reject} & \text{if } s_1 \geq \frac{1}{2 - \beta} \end{cases}.$$

Note that, if player 2 uses the strategy s_2^{**} , then there is no best choice for player 1. If player 2 uses the strategy s_2^* , then player 1's optimal strategy $s_1^* = \frac{1}{2 - \beta}$ yields a payoff of $\frac{\alpha + \beta - \alpha\beta}{2 - \beta} > 0$ (because $\alpha, \beta \in [0, 1]$)

and $\alpha + \beta > 0$). Therefore, the strategy profile yielding the backwards induction outcome can only be the following subgame perfect equilibrium:

$$(s_1^* = \frac{1}{2-\beta}, s_2^*).$$

□

10.34 Example: Splitting one dollar/Ultimatum with an infinite number of alternatives.

Players 1 and 2 are bargaining over one dollar (divisible) in two periods: In the first period, Player 1 proposes s_1 for himself and $1 - s_1$ for player 2. In the second period, player 2 decides whether to accept the offer or to reject the offer. If player 2 accepts the offer, the payoff are s_1 for player 1 and $1 - s_1$ for player 2. If player 2 rejects the offer, the payoff are zero for both players.

- (i) Describe all strategies of player 1 and player 2.
- (ii) Find all Nash equilibria.
- (iii) Find a subgame perfect equilibrium of the game.
- (iv) Find a Nash equilibrium which are not subgame perfect.

Answer. Figure 10.10 is the game tree.




Figure 10.10

- (i) It is easy to see that player 1's strategy space is $S_1 = [0, 1]$. Since a strategy is a complete plan of actions in every contingency when a player is called upon to make, a strategy for player 2 can be represented as a function

$$f: [0, 1] \rightarrow \{A, R\}.$$

For example,

$$f(s_1) = \begin{cases} A, & \text{if } 0 \leq s_1 \leq \frac{1}{2}; \\ R, & \text{otherwise} \end{cases}$$

is a strategy of player 2 in which player 2 will accept if player 1 offers any $s_1 \leq \frac{1}{2}$ and otherwise she will reject.

Thus, the space of all strategies of player 2 is the set of all functions from $[0, 1]$ to $\{A, R\}$. We denote it by S_2 .

- (ii) (1) Player 1's best-response correspondence: Given a strategy f of player 2, note that for any $s_1 \in f^{-1}(A)$, player 2 will accept the offer. Hence, given f , player 1 will choose the maximum in $f^{-1}(A)$ if it exists.

Thus, player 1's best-response correspondence is

$$B_1^*(f) = \begin{cases} [0, 1], & \text{if } f^{-1}(A) = \emptyset; \\ [0, 1], & \text{if } 0 \text{ is the maximum of } f^{-1}(A); \\ \{s^*\}, & \text{if } f^{-1}(A) \text{ has a maximum } s^* \neq 0; \\ \emptyset, & \text{if } f^{-1}(A) \text{ has no maximum.} \end{cases}$$

(2) Player 2's best-response correspondence: note that player 2's strategy is a function

$$B_2^*(s_1) = \begin{cases} \{f \in S_2 : f(s_1) = A\}, & \text{if } 0 \leq s_1 < 1; \\ S_2, & \text{if } s_1 = 1. \end{cases}$$

That means for any $s_1 < 1$, player 2 will accept. If $s_1 = 1$, player 2 is indifferent between the two actions (accept or reject).

(3) We can use various combinations of the conditions in the expression of B_1^* and B_2^* to construct all the Nash equilibria:

- When $f^{*-1}(A) \supsetneq \{0\}$, (s_1^*, f^*) is a Nash equilibrium if and only if $s_1^* = \sup f^{*-1}(A) = \max f^{*-1}(A)$;
- When $f^{*-1}(A) = \{0\}$, (s_1^*, f^*) is a Nash equilibrium if and only if $s_1^* = 0$ or 1;
- When $f^{*-1}(A) = \emptyset$, (s_1^*, f^*) is a Nash equilibrium if and only if $s_1^* = 1$.

(iii) For each given s_1 , we need to consider a corresponding subgame, displayed in Figure 10.11. We know if f^*




Figure 10.11

is subgame perfect, $f^*(s_1) = A$ for any $s_1 < 1$. Hence, if (s_1^*, f^*) is subgame perfect, f^* should be either f_1^* or f_2^* :

$$f_1^*(s_1) = \begin{cases} A, & \text{if } s_1 < 1; \\ R, & \text{if } s_1 = 1. \end{cases} \quad \text{or } f_2^*(s_1) \equiv A \text{ for all } s_1.$$

It is easy to check that only $(s_1^* = 1, f_2^*)$ is the unique subgame perfect equilibrium.

(iv) $(s_1^* = 1, f^* \equiv R)$ is a Nash equilibrium but not a subgame perfect equilibrium.

□

10.35 Example [JR Exercise 7.29]: Take-it-or-leave-it game.

A referee is equipped with N dollars. He places one dollar on the table. Player 1 can either take the dollar or leave it. If he takes it, the game ends. If he leaves it, the referee places a second dollar on the table. Player two is now given the option of taking the two dollars or leaving them. If he takes them, the game ends. Otherwise the referee places a third dollar on the table and it is again player 1's turn to take or leave the three dollars. The game continues in this manner with the players alternately being given the choice to take all the money the referee has so far placed

on the table and where the referee adds a dollar to the total whenever a player leaves the money. If the last player to move chooses to leave the N dollars the game ends with neither player receiving any money. Assume that N is public information.

- (i) Without thinking too hard, how would you play this game if you were in the position of player 1? Would it make a difference if N were very large (like a million) or quite small (like 5)?
- (ii) Calculate the backward induction strategies. Do these make sense to you?
- (iii) Prove that the backward induction strategies form a Nash equilibrium.
- (iv) Prove that the outcome that results from the backward induction strategies is the unique outcome in any Nash equilibrium. Is there a unique Nash equilibrium?

10.36 Two players, A and B , take turns choosing a number between 1 and 9 (inclusive). The cumulative total of all the numbers chosen is calculated as the game progresses. The player whose choice of number takes the total to exactly 100 is the winner. Is there a first mover advantage in this game?

10.4 Three notable games

10.37 Father-Son-CEO-Manager game.

Player 1 (an entrepreneur) has to decide whether to sell the firm (action L_1) or to delegate control to his son (player 2). Player 2 can then decide to manage the firm himself (action L_2) or hire player 3 (CEO) to run the business. The CEO, in turn, may or may not delegate control to player 4 (a manager). The manager can, then, either exert effort to manage the business well (action L_4), or shirk (action R_4). Assume that the game, actions, and resulting payoffs as depicted in the following figure are all common knowledge.




Figure 10.12: Father-Son-CEO-Manager game.

There are two subgame perfect equilibria: (L_1, R_2, R_3, L_4) and (L_1, R_2, L_3, R_4) , which share the same payoff $(1, 1, 1, -1)$. However, players can do better through an course of actions (R_1, L_2, \cdot, \cdot) .

10.38 The chain-store game.

A chain-store (player CS) has branches in K cities, numbered $1, 2, \dots, K$. In each city k there is a single potential competitor, player k . In each period one of the potential competitors decides whether or not to compete with player CS ; in period k it is player k 's turn to do so. If player k decides to compete then the chain-store can either fight (F) or cooperate (C). The chain-store responds to player k 's decision before player $k + 1$ makes its decision. Thus




Figure 10.13: The chain-store game.

in period k the set of possible outcomes is $Q = \{Out, (In, C), (In, F)\}$. The structure of the players' choices and their considerations in a single period are summarized in the following figure.

There are two assumptions:

- At every point in the game all players know all the actions previously chosen. The set of histories is

$$(\cup_{k=0}^K Q^k) \cup (\cup_{k=0}^{K-1} (Q^k \times \{In\})),$$

where Q^k is the set of all sequences of k members of Q , and the player function is given by $P(h) = k + 1$ if $h \in Q^k$ and $P(h) = CS$ if $h \in Q^k \times \{In\}$, for $k = 0, 1, \dots, K - 1$.

- the payoff of the chain-store in the game is the sum of its payoffs in the K cities.

This game has a multitude of Nash equilibria: in period k , player k and chain-store choose (Out, F) or (In, C) .

This game has a unique subgame perfect equilibrium: every challenger choose In , and the chain-store always chooses C . In city K the chain-store must choose C , regardless of the history, so that in city $K - 1$ it must do the same, continuing the argument one sees that the chain-store must always choose C .

Although the chain-store's unique subgame perfect equilibrium strategy does indeed specify that it cooperate with every entrant, it seems more reasonable for a competitor who has observed the chain-store fight repeatedly to believe that its entry will be met with an aggressive response, especially if there are many cities still to be contested. If a challenger enters then it is in the myopic interest of the chain-store to be cooperative, but intuition suggests that it may be in its long-term interest to build a reputation for aggressive behavior, in order to deter future entry.

10.39 ????? Reputation

10.40 The centipede game.




Figure 10.14: The centipede game.

The set of histories in the game is

$$\{C(t) = (\underbrace{C, \dots, C}_t) : t = 0, 1, \dots, T\} \cup \{S(t) = (\underbrace{C, \dots, C}_{t-1}, S) : t = 1, \dots, T\}.$$

The player function is defined by

$$P(C(t)) = \begin{cases} 1, & \text{if } t = 0, 2, 4; \\ 2, & \text{if } t = 1, 3, 5. \end{cases}$$

The game has unique subgame perfect equilibrium; in this equilibrium each player chooses S in every period.

Now assume that there is a Nash equilibrium that ends with player i choosing S in period t . If $t \geq 2$ then player j can increase his payoff by choosing S in period $t-1$. Hence in any equilibrium player 1 chooses S in the first period. In order for this to be optimal for player 1, player 2 must choose S in period 2. The notion of Nash equilibrium imposes no restriction on the players' choices in later periods: any pair of strategies in which player 1 chooses S in period 1 and player 2 chooses S in period 2 is a Nash equilibrium.

In the unique subgame perfect equilibrium of this game each player believes that the other player will stop the game at the next opportunity, even after a history in which that player has chosen to continue many times in the past. Such a belief is not intuitively appealing.

After a history in which both a player and his opponent have chosen to continue many times in the past, the basis on which the player should form a belief about his opponent's action in the next period is far from clear.

10.5 Iterated elimination of weakly dominated strategies

10.41 Let Γ be a finite extensive game with perfect information in which no player is indifferent between any two terminal histories. Then Γ has a unique subgame perfect equilibrium.

We now define a sequence for eliminating weakly dominated actions in the induced strategic game G of Γ (weakly dominated strategies in Γ) with the property that all the action profiles of G that remain at the end of the procedure generate the unique subgame perfect equilibrium outcome of Γ .

Let h be a history of Γ with $P(h) = i$ and $\ell(\Gamma(h)) = 1$ and let $a^* \in A(h)$ be the unique action selected by the procedure of backwards induction for history h . Backwards induction eliminates every strategy of player i that chooses an action different from a_i^* after history h . Among these strategies, those consistent with h (i.e., that choose the component of h that follows h' whenever h' is a subhistory of h with $P(h') = i$) are weakly dominated actions in G . Perform this elimination for each history h with $\ell(\Gamma(h)) = 1$.

Then we turn to histories h with $\ell(\Gamma(h)) = 2$ and perform an analogous elimination; we continue back to the beginning of the game in this way.

Every strategy of player i that remains at the end of this procedure chooses the action selected by backwards induction after any history that is consistent with player i 's subgame perfect equilibrium strategy. Thus in particular the subgame perfect equilibrium remains and every strategy profile that remains generates the unique subgame perfect equilibrium outcome.

10.42 Example:

$$H = \{\emptyset, (A), (B), (A, C), (A, D), (A, C, E), (A, C, F)\}, S_1 = \{AE, AF, BE, BF\}, S_2 = \{C, D\}.$$

- (1) Consider (A, C) firstly. For player 1, E is better than F , then the strategy AF is weakly dominated by AE : if player 2 chooses D , AF and AE yield the same payoff; if player 2 chooses C , AE is strictly better than AF . AF is eliminated.
- (2) Then consider (A) , and it is player 2's turn. D is better than C given that player 1 will choose E when he choose C here. Then the strategy C is weakly dominated by D : if player 1 chooses BE or BF , then C and




Figure 10.15: Iterated elimination of weakly dominated strategies.

D yield the same payoff; if player 1 chooses AE , then D is strictly better than C . Note that AF has been eliminated, and will be not considered any longer. C is eliminated.

(3) Lastly, consider \emptyset . It is player 1's turn. AE is weakly dominated by BE , and eliminated.

(4) BE , BF and D remain.

We will see that every strategy of player i that remains at the end of this procedure chooses the action selected by backwards induction after any history that is consistent with player i 's subgame perfect equilibrium strategy.

For player 1, his subgame perfect equilibrium strategy is BE , consistent history can be \emptyset and (B) . At the history \emptyset , the outcomes BE and BF both suggest that player 1 choose B , same as the equilibrium behavior.

10.43 Note, however, that other orders of elimination may remove all subgame perfect equilibria.

Consider the strategy example above.

	C	D
AE	2, 0	1, 1
AF	0, 2	1, 1
BE	3, 3	3, 3
BF	3, 3	3, 3

(1) AE is weakly (actually strictly) dominated by BE , and eliminated;

(2) D is weakly dominated by C , and eliminated;

(3) AF is weakly dominated by BF , and eliminated;

(4) (BE, C) and (BF, C) remain, but neither of them are subgame perfect equilibria.

10.6 Forward induction

10.44 In the following game, player 1's strategy set is

$$S_1 = \{(Book, O), (Book, F), (Outside, O), (Outside, F)\},$$

and player 2's strategy set is $\{O, F\}$.

Consider its reduced strategic form:




Figure 10.16: Forward induction.

	<i>O</i>	<i>F</i>
<i>Book</i>	2, 2	2, 2
<i>O</i>	3, 1	0, 0
<i>F</i>	0, 0	1, 3

F is strictly dominated for player 1 by $Book$, and eliminated. Then F is weakly dominated for player 2 by O , and eliminated. Finally, $Book$ is strictly dominated by O for player 1. The outcome that remains is (O, O) .

This sequence of eliminations corresponds to the following argument for the extensive game:

- (1) If player 2 has to make a decision he knows that player 1 has not chosen $Book$.
- (2) Such a choice makes sense for player 1 only if she plans to choose O .
- (3) Thus player 2 should choose O also.

The logic of such an argument is referred to in the literature as “forward induction”.

- 10.45 Two individuals are going to play the battle of sexes. Before doing so player 1 can discard a dollar (take the action D) or refrain from doing so (take the action 0); her move is observed by player 2.

Player 1's strategy set is

$$S_1 = \{0OO, 0OF, 0FO, 0FF, DOO, DOF, DFO, DFF\},$$

and player 2's strategy set is $\{OO, OF, FO, FF\}$.




Figure 10.17: Forward induction.

The reduced strategic game is as follows:

	<i>OO</i>	<i>OF</i>	<i>FO</i>	<i>FF</i>
<i>OO</i>	3, 1	3, 1	0, 0	0, 0
<i>OF</i>	0, 0	0, 0	1, 3	1, 3
<i>DO</i>	2, 1	-1, 0	2, 1	-1, 0
<i>DF</i>	-1, 0	0, 3	-1, 0	0, 3

Weakly dominated actions can be eliminated iteratively as follows.

- (1) *DF* is weakly dominated for player 1 by *OO*;
- (2) *FF* is weakly dominated for player 2 by *FO*;
- (3) *OF* is weakly dominated for player 2 by *OO*;
- (4) *OF* is strictly dominated for player 1 by *DO*;
- (5) *FO* is weakly dominated for player 2 by *OO*;
- (6) *DO* is strictly dominated for player 1 by *OO*.

The single strategy pair that remains is (*OO*, *OO*): the fact that player 1 can throw away a dollar implies, under iterated elimination of weakly dominated actions, that the outcome is player 1's favorite.

An intuitive argument that corresponds to this sequence of eliminations is the following.

- (1) Player 1 must anticipate that if she chooses 0 then she will obtain an expected payoff of at least $\frac{3}{4}$, since for every belief about the behavior of player 2 she has an action that yields her at least this expected payoff.
- (2) Thus if player 2 observes that player 1 chooses *D* then he must expect that player 1 will subsequently choose *O* (since the choice of *F* can not possibly yield player 1 a payoff in excess of $\frac{3}{4}$).
- (3) Given this, player 2 should choose *O* if player 1 chooses *D*; player 1 knows this, so that she can expect to obtain a payoff of 2 if she chooses *D*.
- (4) But now player 2 can rationalize the choice 0 by player 1 only by believing that player 1 will choose *O* (since *F* can yield player 1 no more than 1), so that the best action of player 2 after observing 0 is *O*. This makes 0 the best action for player 1.

Chapter 11

Bargaining games

Contents

11.1 A bargaining game of alternating offers	173
11.2 Bargaining games with finite horizon	174
11.3 Bargaining games with infinite horizon	175
11.4 Properties of subgame perfect equilibria in Rubinstein bargaining games	178
11.5 Bargaining games with cost	179
11.6 n -person bargaining games	179

11.1 Game theory deals with situations in which people's interests conflict. The people involved may try to resolve the conflict by committing themselves voluntarily to a course of action that is beneficial to all of them. If there is more than one course of action more desirable than disagreement for all individuals and there is conflict over which course of action to pursue then some form of negotiation over how to resolve the conflict is necessary. The negotiation/bargaining process may be modeled using the tools of game theory.

11.2 Reference: Section 4.4 in [Fudenberg and Tirole \(1991\)](#), Section 2.1D in [Gibbons \(1992\)](#), Chapter 7 in [Osborne and Rubinstein \(1994\)](#), [Osborne and Rubinstein \(1990\)](#), [Rubinstein \(1982\)](#), [Ståhl \(1972\)](#).

11.1 A bargaining game of alternating offers

11.3 Players 1 and 2 are bargaining over one dollar. Let x_i denote the share of player i , $i = 1, 2$. The set of agreements is

$$X = \{(x_1, x_2) \in \mathbb{R}_+^2 \mid x_1 + x_2 = 1\}.$$

11.4 The game lasts for T periods (T could be infinite).

- (1) In period 1, player 1 makes an offer, x_1^1 for himself and x_2^1 for player 2. If player 2 accepts (A), then they split the dollar according to the offer. If player 2 rejects (R), then they move to period 2.
- (2) In period 2, they exchange roles with player 2 making an offer x_2^2 for himself and x_1^2 for player 1, and player 1 decides whether to accept.

(3) In general, player 1 makes offer in odd periods, and player 2 makes offer in even periods.

(4) The game continues until an agreement is reached or after the end of period T .

11.5 Convention: We will assume that a player accepts whenever he is indifferent between accepting and rejecting. All payoffs are evaluated from the current period.

11.6 If an agreement (x_1, x_2) is reached in period t , then player i receives payoff

$$u_i(x_i, t) = \delta_i^t x_i,$$

where $\delta_i \in (0, 1)$ is player i 's discount factor. The discount factors are assumed to be common knowledge.

11.7 If no agreement is reached after $T - 1$ periods, then an exogenous settlement $d = (s_1, s_2)$ (breakdown's payoff profile) is enforced in period T , where $s_1 + s_2 \leq 1$.

One typical breakdown's payoff profile is $(0, 0)$.

11.2 Bargaining games with finite horizon

11.8 For finite T , we can find the subgame perfect equilibria by backwards induction.

11.9 Suppose $T = 3$ and $d = (s_1, s_2)$. The unique subgame perfect equilibrium can be determined by the backwards induction.




Figure 11.1

(1) In period 2, player 1 can obtain s_1 in the next period by rejecting player 2's present offer. Thus, player 1 will reject the offer (x_1^2, x_2^2) if and only if x_1^2 is strictly worse than $\delta_1 s_1$.

Based on this observation, player 2 can obtain at least $1 - \delta_1 s_1$ in period 2.

(2) In period 1, player 2 knows that he can obtain $1 - \delta_1 s_1$ in the next period. Hence, by the same reasoning, he will accept the present offer if and only if

$$x_2^1 \geq \delta_2(1 - \delta_1 s_1) = \delta_2 - \delta_1 \delta_2 s_1,$$

$$\text{i.e., } x_1^1 \leq 1 - \delta_2 + \delta_1 \delta_2 s_1.$$

(3) Hence, in the subgame perfect equilibrium player 1 offers

$$\underbrace{(1 - \delta_2 + \delta_1 \delta_2 s_1)}_{x_1^*(3)}, \underbrace{(\delta_2 - \delta_1 \delta_2 s_1)}_{x_2^*(3)}$$

and player 2 accepts in period 1.

11.10 Suppose $T = 5$. This case is equivalent to the case $T = 3$ with the breakdown's payoff profile $(1 - \delta_2 + \delta_1 \delta_2 s_1, \delta_2 - \delta_1 \delta_2 s_1)$.

Hence player 1's equilibrium share is:

$$x_1^*(5) = 1 - \delta_2 + \delta_1 \delta_2 (1 - \delta_2 + \delta_1 \delta_2 s_1) = (1 - \delta_2)(1 + \delta_1 \delta_2) + (\delta_1 \delta_2)^2 s_1.$$

11.11 By induction, when $T = 2n + 1$, we have player 1's equilibrium share

$$x_1^*(2n + 1) = (1 - \delta_2) \sum_{i=0}^{n-1} (\delta_1 \delta_2)^i + (\delta_1 \delta_2)^n s_1.$$

11.12 When $T = 2n + 2$, we know that if the game proceeds to period 2, player 2 will obtain

$$(1 - \delta_1) \sum_{i=0}^{n-1} (\delta_1 \delta_2)^i + (\delta_1 \delta_2)^n s_2.$$

So, in this case, player 1 equilibrium share is

$$x_1^*(2n + 2) = 1 - \delta_2 (1 - \delta_1) \sum_{i=0}^{n-1} (\delta_1 \delta_2)^i - \delta_2 (\delta_1 \delta_2)^n s_2.$$

11.3 Bargaining games with infinite horizon

11.13 We will focus on the case $d = (0, 0)$.

11.14 The set of Nash equilibria of a bargaining game of alternating offers is very large, and almost any division of the dollar can be obtained as a Nash equilibrium outcome. In particular, for any $x^* \in X$ there is a Nash equilibrium in which the players immediately agree on x^* (i.e. player 1's equilibrium strategy assigns x^* in period 1 and player 2's strategy assigns A to x^*).

One such equilibrium is that in which both players always offer x^* and always accept a proposal x if and only if $x = x^*$.

11.15 For many specifications of the players' preferences there are Nash equilibria in which an agreement is not reached immediately. For example, for any agreement x and period t , there is a Nash equilibrium for which the outcome is the acceptance of x in period t .

One such equilibrium is that in which through period $t - 1$ each player demands the whole dollar and rejects all proposals, and from period t on offers x and accepts only x .

11.16 The notion of Nash equilibrium does not exclude the use of “incredible threats.” Consider the Nash equilibrium in which both players always offer $x^* = (x_1^*, x_2^*)$ and player i accepts a proposal x if and only if $x_i \geq x_i^*$.

Take a particular proposal $x = (x_1, x_2)$ such that $x_1 = x_1^*$ and $x_2^* > x_2 > \delta_2 x_2^*$. Then in the equilibrium player 2's strategy dictates that he rejects such a proposal x . This “threat” induces player 1 to offer x^* .

Player 2's threat is incredible, given player 1's strategy: the best outcome that can occur if player 2 carries out his threat to reject x is that there is agreement on x^* in the next period, an outcome that player 2 likes less than agreement on x in the previous period, which he can achieve by accepting x (since $x_2 > \delta_2 x_2^*$).

11.17 Since $T = \infty$, we can no longer find subgame perfect equilibria by backwards induction (since there is no final period). In this case, we need to use an extra trick to find the equilibrium.

11.18 Recall: A strategy profile is said to satisfy the one deviation property if for each subgame the player who makes the first move can not obtain a better outcome by changing only his initial action.

11.19 Proposition: In any perfect information extensive game with either finite horizon or discounting, a strategy profile is a subgame perfect equilibrium if and only if it satisfies the one deviation property.

This property is extremely useful in games with an infinite horizon, such as the current bargaining game or infinitely repeated games. In these games, since the players have an infinite number of strategies, it is hard to show that any particular strategy is a best response. The one deviation property says that we need only to show that at every decision node a player will not deviate in that decision node and follow the equilibrium strategy in the future.

11.20 *Proof.* OR 123.1 □

11.21 Consider Equations in 11.11 and 11.12. Let T go to infinity, then we have

$$x_1^* \equiv \lim_{T \rightarrow \infty} x_1^*(T) = \frac{1 - \delta_2}{1 - \delta_1 \delta_2}, \quad x_2^* \equiv \lim_{T \rightarrow \infty} x_2^*(T) = \delta_2 \frac{1 - \delta_1}{1 - \delta_1 \delta_2}.$$

Note that the limit is the same whether T is odd or even.

11.22 Let $(y_1^*(T), y_2^*(T))$ denote the equilibrium division when we interchange the roles of the players and let player 2 make offer in period 1. When T goes to infinity, the equilibrium share will become

$$y_1^* \equiv \lim_{T \rightarrow \infty} y_1^*(T) = \delta_1 \frac{1 - \delta_2}{1 - \delta_1 \delta_2}, \quad y_2^* \equiv \lim_{T \rightarrow \infty} y_2^*(T) = \frac{1 - \delta_1}{1 - \delta_1 \delta_2}.$$

Note that

$$y_1^* = \delta_1 x_1^* \text{ and } x_2^* = \delta_2 y_2^*.$$

✂ 11.23 Theorem (Rubinstein, 1982): In the bargaining game with infinite horizon, there is a unique subgame perfect equilibrium where

- in every odd period, player 1 offers (x_1^*, x_2^*) and player 2 accepts any $x_2 \geq x_2^*$,
- in every even period, player 2 offers (y_1^*, y_2^*) and player 1 accepts any $y_1 \geq y_1^*$.

Therefore, the subgame perfect equilibrium outcome is that player 1 offers

$$\left(\frac{1 - \delta_2}{1 - \delta_1 \delta_2}, \delta_2 \frac{1 - \delta_1}{1 - \delta_1 \delta_2} \right)$$

and player 2 accepts in period 1.

This theorem says that in each period the players will behave as if in an extremely long finite horizon game.

11.24 Note that the game is stationary: the subgame starting from any period t looks exactly like the original game. This is an extremely important property because it implies that if a strategy profile is an equilibrium in period t , it will be an equilibrium in the next period as well.

11.25 *Proof of subgame perfection.*

- (1) To show that the strategy profile is subgame perfect, we need to show that no player can gain by deviating once immediately and follow the equilibrium strategy in the future.
- (2) In all odd periods:
 - Obviously, player 1 will not be strictly better off by proposing $(x_1, 1 - x_1)$ with $x_1 < x_1^*$.
 - If player 1 offers $(x_1, 1 - x_1)$ with $x_1 > x_1^*$, then player 2 will reject the offer and player 1 will obtain

$$\delta_1 y_1^* = \delta_1^2 x_1^* < x_1^*$$

in the next period, making him worse off.

- If player 2 rejects x_2^* , then he will obtain y_2^* in the next period. Hence it is a best response to accept any $x_2 \geq \delta_2 y_2^* = x_2^*$.
- (3) The case for even periods is similar.

□

11.26 *Proof of uniqueness.*

- (1) Let \bar{x}_1 and \underline{x}_1 denote the maximal and minimal subgame perfect equilibrium payoffs for player 1 when player 1 is the proposer respectively. Let \bar{y}_2 and \underline{y}_2 denote the maximal and minimal subgame perfect equilibrium payoffs for player 2 when player 2 is the proposer respectively.
- (2) Consider an odd period. Since player 2 can get at least \underline{y}_2 in the next period by rejecting player 1's offer, in any subgame perfect equilibrium, player 1 must offer player 2 at least $\delta_2 \underline{y}_2$. Hence,

$$\bar{x}_1 \leq 1 - \delta_2 \underline{y}_2.$$

- (3) On the other hand, considering an even period, player 1 can get at most \bar{x}_1 in the next period by rejecting player 2's offer. It would not be an equilibrium for player 2 to offer more than $\delta_1 \bar{x}_1$ to player 1. Hence,

$$\underline{y}_2 \geq 1 - \delta_1 \bar{x}_1.$$

- (4) Combining the two inequalities, we have

$$\bar{x}_1 \leq 1 - \delta_2(1 - \delta_1 \bar{x}_1) = 1 - \delta_2 + \delta_1 \delta_2 \bar{x}_1,$$

which means that

$$\bar{x}_1 \leq \frac{1 - \delta_2}{1 - \delta_1 \delta_2} = x_1^*.$$

- (5) Interchanging the roles of the players, the same arguments imply that

$$\bar{y}_2 \leq 1 - \delta_1 \underline{x}_1, \text{ and } \underline{x}_1 \geq 1 - \delta_2 \bar{y}_2.$$

(6) Combining the two inequalities, we have

$$\underline{x}_1 \geq 1 - \delta_2(1 - \delta_1 \underline{x}_1) = 1 - \delta_2 + \delta_1 \delta_2 \underline{x}_1$$

which means that

$$\underline{x}_1 \geq \frac{1 - \delta_2}{1 - \delta_1 \delta_2} = x_1^*.$$

(7) Since we know $\bar{x}_1 \geq \underline{x}_1$, it follows that $x_1^* = \bar{x}_1 = \underline{x}_1$. By the same logic, we can show that $y_2^* = \bar{y}_2 = \underline{y}_2$.

(8) This shows that the subgame perfect equilibrium payoff is unique. Given this, it is obvious that the equilibrium itself must also be unique.

□

11.27 Remark: If discount factors are not assumed to be common knowledge, then bargaining can last for more than one period in the alternating offers model.

11.4 Properties of subgame perfect equilibria in Rubinstein bargaining games

11.28 Efficiency: The structure of a bargaining game of alternating offers allows negotiation to continue indefinitely. Nevertheless, in the unique subgame perfect equilibrium it terminates immediately; from an economic point of view, the bargaining process is efficient (no resources are lost in delay).

To which features of the model can we attribute this result? We saw that in a Nash equilibrium of the game, delay is possible. Thus the notion of subgame perfection plays a role in the result.

11.29 Stationarity: The subgame perfect equilibrium strategies are stationary: for any history after which it is player i 's turn to offer an agreement he offers the same agreement, and for any history after which it is his turn to respond to a proposal he uses the same criterion to choose his response.

We have not restricted players to use stationary strategies; rather, such strategies emerge as a conclusion.

11.30 First mover advantage: There is a first-mover advantage even though there are many periods of negotiation.

Suppose that $\delta_1 = \delta_2 = \delta$, then the only asymmetry in the game is that player 1 moves first. Player 1's equilibrium payoff is $\frac{1}{1+\delta}$ which exceeds $\frac{1}{2}$, but approaches $\frac{1}{2}$ as δ tends to 1. Thus if the players are equally and only slightly impatient, player 1's first mover advantage is small and the outcome is almost symmetric.

Player 1 gets the whole dollar if $\delta_2 = 0$, since a myopic player 2 will accept any positive amount this period rather than wait one period. However, even if $\delta_1 = 0$ player 2 does not get the whole dollar if $\delta_2 < 1$.

11.31 The breakdown share is irrelevant: the division is entirely driven by the discounts factor.

11.32 Effect of changes in patience: A player's share increases with his discount factor and decreases with his opponent's discount factor, and player i 's payoff converges to 1 as $\delta_i \rightarrow 1$. That is, fixing the patience of player 2, player 1's share increases as she becomes more patient.

11.33 Let $\exp(-r_i \Delta)$ be player i 's discounting factor, where r_i is player i 's discounting rate and Δ is the duration of each period. Then the subgame perfect equilibrium payoff profile converges to $\left(\frac{r_2}{r_1+r_2}, \frac{r_1}{r_1+r_2}\right)$ as $\Delta \rightarrow 0$.

11.5 Bargaining games with cost

11.34 Consider the following two-player bargaining game. Each player i incurs the cost $c_i > 0$ for every period, that is, player i 's payoff if the agreement is concluded in period t is $x_i - c_i t$.

11.35 Proposition [OR Exercise 125.2]: If $c_1 < c_2$, then the bargaining game has a unique subgame perfect equilibrium where in every odd period, player 1 offers $(1, 0)$ and player 2 accepts any $x_2 \geq 0$, and in every even period player 2 offers $(1 - c_1, c_1)$ and player 1 accepts any $y_1 \geq 1 - c_1$.

Proof. Omitted. □

11.36 Proposition [OR Exercise 125.2]: If $c_1 = c_2 = c < 1$, then the game has many subgame perfect equilibrium outcomes including, if $c < \frac{1}{3}$, equilibria in which agreement is delayed.

Proof. Omitted. □

11.6 n -person bargaining games

11.37 Reference: David P. Baron and John A. Ferejohn, Bargaining in Legislatures, *American Political Science Review* 83 (1989), 1181–1206.

11.38 Consider the following n -person bargaining game:

- n (odd) players try to allocate one dollar among them. Let

$$X = \{x \in \mathbb{R}_+^n \mid \sum_{i=1}^n x_i \leq 1\}$$

be the set of feasible allocations.

- The game is played for T periods (T could be infinite).
- In any period t , a player is chosen as a proposer with probability $\frac{1}{n}$.
- The proposer suggests how to divide the one dollar, *i.e.*, chooses some $x = (x_1, x_2, \dots, x_n)$ from X .
- The player vote publicly and sequentially (in some order). If the proposal is approved by the majority, then it is implemented and the game is over.

Otherwise the game moves on to the next period.

- If no proposal is approved by the end of the game, then everyone receives 0.
- Each player's discount factor is $\delta \in (0, 1)$, which is common knowledge.

11.39 Suppose $T = 2$.

- (1) In period 2, whoever is chosen as a proposer can request everything.
- (2) In period 1, the proposer can buy one vote by paying $\frac{\delta}{n}$. Hence the proposer can pay $\frac{\delta}{n}$ to $\frac{n-1}{2}$ players so that his proposal just gets the majority votes.

11.40 Suppose $T = \infty$. We first focus on symmetric stationary subgame perfect equilibrium, where

- (i) the distribution of proposals is same independent of histories,

- (ii) every player except for the proposer is treated symmetrically by the equilibrium proposal,
- (iii) the equilibrium voting behavior is the same across all the players.

11.41 Theorem: For any $\delta \in (0, 1)$, there exists the unique symmetric stationary subgame perfect equilibrium. In equilibrium, the proposer always proposes to distribute $\frac{\delta}{n}$ to randomly selected $\frac{n-1}{2}$ players. Player i votes for the proposal if and only if the proposal assigns player i at least $\frac{\delta}{n}$.

11.42 Proof. (1) It is easy to prove that it satisfies the one deviation property.

- (2) Take any symmetric stationary subgame perfect equilibrium. Every player's equilibrium payoff in the beginning of each period must be the same. Denote this by v .
- (3) Each proposer is guaranteed to receive at least $1 - \delta v \frac{n-1}{2}$ by paying δv to $\frac{n-1}{2}$ players.
- (4) To minimize expense, it must be the case that the proposer pays exactly δv to $\frac{n-1}{2}$ players. Note that each other player is in the coalition with probability $\frac{1}{2}$. So v must satisfy

$$v = \frac{1}{n} \left(1 - \delta v \frac{n-1}{2} \right) + \frac{n-1}{n} \frac{1}{2} \delta v.$$

Hence, $v = \frac{1}{n}$.

□

11.43 Once stationarity is dropped, then many allocations can be supported by subgame perfect equilibria.

In fact, any allocation can be supported if there are many players and the players are patient.

11.44 Theorem: Suppose that $n \geq 5$ and $\frac{n+1}{2(n-1)} < \delta < 1$. Then any $x \in X$ can be achieved by a subgame perfect equilibrium where

- every proposer proposes x if there has been no deviation by any proposer. This proposal is accepted by every player immediately.
- if player j deviates and proposes $y \neq x$, then
 - (1) it is rejected by some majority $M(y)$ that does not include j ,
 - (2) the next proposer proposes $z(y) \in X$ such that $z_j(y) = 0$ and everyone in $M(y)$ votes for $z(y)$.
- if the next proposer k proposes $w \neq y$ instead of y in the previous step, then repeat the previous step with $(z(w), k)$ instead of $(z(y), j)$.

11.45 Proof. (1) No proposer has an incentive to deviate from x because then the continuation payoff is 0.

(2) Consider (j, y) -phase. We define $M(y)$ and $z(y)$ as follows:

- $M(y)$ is a group of $\frac{n+1}{2}$ players such that $j \notin M(y)$ and $\sum_{i \in M(y)} y_i$ is minimized.
- $z_i(y) = 0$ for $i \notin M(y)$ and $z_i(y) \in \frac{y_i}{\sum_{k \in M(y)} y_k}$ for $i \in M(y)$.

- (3) No proposer in (j, y) -phase (even player j) does not have an incentive to deviate and propose something different from $z(y)$ because then the continuation payoff is 0.
- (4) Everyone votes for x and every player in $M(y)$ votes for $z(y)$ (a deviation just causes a delay).
- (5) Finally we need to make sure that $M(y)$ rejects y in favor of $z(y)$ in the next period.
 - This is trivially satisfied for $i \in M(y)$ such that $y_i = 0$.
 - For $i \in M(y)$ with $y_i > 0$, we need $\delta z_i(y) \geq y_i$, which is $\delta \geq \frac{\sum_{k \in M(y)} y_k}{\sum_{k \in M(y)} y_k}$. The least upper bound of $\sum_{k \in M(y)} y_k$ is $\frac{n+1}{2(n-1)}$, which is less than 1 if $n \geq 5$.

□

- 11.46 Remark: In this construction, if i is the pivotal voter (voting $\frac{n+1}{2}$ -th “yes”) in $M(y)$ and $z_i(y) = 0$, then i is playing a weakly dominated strategy by voting for $z_i(y)$. This can be fixed easily by considering a slightly more complicated transfer: it is possible to perturb $z_i(y)$ slightly so that $\delta z_i(y) > y_i$ holds for every $i \in M(y)$.

Chapter 12

Repeated games

Contents

12.1	Infinitely repeated games	183
12.2	Trigger strategy equilibrium	188
12.3	Tit-for-tat strategy equilibrium	193
12.4	Folk theorem	194
12.5	Nash-threats folk theorem	195
12.6	Perfect folk theorem	196
12.7	Finitely repeated games	199

A repeated game is simply a situation in which players have the same encounter over and over, *i.e.*, the same game is played repeatedly. The game that is played repeatedly is known as the stage game. The stage game itself could be a simultaneous move game or a sequential move game. If the stage game is played finitely many times then we have a finitely repeated game. If the stage game is played infinitely many times then we have an infinitely repeated game.


Repeated games generate fundamentally different incentives from isolated interactions, providing explanations to:

- collusion (oligopoly, procurement auctions, sports, *etc.*)
- cooperation (human, animals)

Repeated games often admit a simple characterization of the set of equilibrium payoffs, known as the “folk theorem”. If the players’ actions are observed at the end of each period, it becomes possible for players to condition their play on the past play of their opponents, which can lead to equilibrium outcomes that do not arise when the game is played only once.

References: Chapter 5 in [Fudenberg and Tirole \(1991\)](#), Chapters 2 and 3 in [Mailath and Samuelson \(2006\)](#), [Takahashi \(2014\)](#).

12.1 Infinitely repeated games

-  12.1 The game begins in period 0. In each period, n players play a simultaneous-move stage game.

- Player set $N = \{1, 2, \dots, n\}$.
- Let A_i denote player i 's stage-game action set, and $\Delta(A_i)$, the set of probability distributions over A_i , serves as player i 's stage-game mixed action set. Let $A = \times_{i=1}^n A_i$.
- Let $u_i: A \rightarrow \mathbb{R}$ denote player i 's stage-game payoff function. Let $u = (u_1, u_2, \dots, u_n)$.

Furthermore, we have the following notations:

- Feasible payoffs

$$F = \text{co } u(A) = \text{convex hull } \{x \in \mathbb{R}^n \mid \text{there exists } a \in A \text{ such that } u(a) = x\}.$$

- The minmax value

$$\underline{v}_i = \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

It is the minimum payoff a player can guarantee himself regardless of the other players' strategies. A payoff u_i is individually rational in the repeated game if it is not less than \underline{u}_i .

The minmax action profile

$$\begin{aligned} \underline{\alpha}_{-i}^i &= \arg \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}), \\ \underline{a}_i^i &= \arg \max_{a_i \in A_i} u_i(a_i, \underline{\alpha}_{-i}^i). \end{aligned}$$

- Feasible and weakly individually rational payoffs

$$F_+ = \{x \in F \mid \text{for each } i, x_i \geq \underline{v}_i\}.$$

- Feasible and individually rational payoffs

$$F_{++} = \{x \in F \mid \text{for each } i, x_i > \underline{v}_i\}.$$

12.2 Example: Infinitely repeated prisoners' dilemma. The stage game is as follows:

	H	C
H	1, 1	-1, 2
C	2, -1	0, 0

Figure 12.1

In this stage-game,

- F is the gray region in Figure 12.2.
- $\underline{v}_i = 0$ for $i = 1, 2$.
- $F_+ = F \cap \{(u_1, u_2) \mid u_1, u_2 \geq 0\}$.
- $F_{++} = F \cap \{(u_1, u_2) \mid u_1, u_2 > 0\}$.




Figure 12.2

- 12.3 In the beginning of each period, the players observe actions chosen in last. For each $t \in \{0, 1, 2, \dots\}$, a t -period history is a vector

$$h^t = (\emptyset, a^0, a^1, \dots, a^{t-1}),$$

where $a^s = (a_1^s, a_2^s, \dots, a_n^s)$ is the action profile chosen in period s for $s = 0, 1, \dots, t-1$. The “null” history, \emptyset , serves no real function. It is inserted for consistency in notation.

Let H^t denote the set of all t -period histories.

The set of all possible histories is $H = \cup_{t=0}^{\infty} H^t$.

- 12.4 Players choose their stage-game action based on what has happened in the past.

A period t 's (behavior) strategy for player i is

$$\sigma_i^t: H^t \rightarrow \Delta(A_i).$$

A repeated-game strategy of player i is

$$\sigma_i: H \rightarrow \Delta(A_i).$$

that describes player i 's strategy in every period.

Let Σ_i denote the set of repeated-game strategies of player i .

- 12.5 An outcome (or full history, or path) is an infinite sequence of action profiles $a = (a^0, a^1, \dots, a^t, \dots)$.

- 12.6 The players discount future payoffs by a common discount factor $\delta \in [0, 1)$.

When the outcome is $a = (a^0, a^1, \dots, a^t, \dots)$, player i 's discounted payoff is

$$\sum_{t=0}^{\infty} \delta^t u_i(a^t).$$

The discount factor can be interpreted as a preference for the present over the future. It can also be interpreted as one minus the probability that the game determinate in that period.

If a player always gets x in each period, then her discounted payoff is $\frac{x}{1-\delta}$.

- 12.7 In most applications, it is mathematically more convenient to work with the normalized discounted payoff (also

called average payoff):

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t).$$

Since the normalized discounted payoff is just a monotone transformation of the discounted payoff, they represent the same underlying preferences.

If a player receives a stage-game payoff x in each of the first t periods and y in each of all subsequent periods, then her normalized discounted payoff is a weighted average of x and y :

$$(1 - \delta) \cdot \left(\sum_{s=0}^{t-1} \delta^s x + \sum_{s=t}^{\infty} \delta^s y \right) = (1 - \delta^t) \cdot x + \delta^t y.$$

If a player always gets x in each period, then her normalized discounted payoff is also x (compare with the discounted payoff $\frac{x}{1-\delta}$).

12.8 A repeated-game strategy profile σ induces a probability distribution over the set of full histories.

If σ is a pure strategy profile, then player i 's normalized discounted payoff is

$$U_i(\sigma) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t(\sigma)),$$

where $a^t(\sigma)$ is the induced action profile in period t .

As usual, the normalized discounted payoff to player i from a profile of behavior strategies σ is the expected value of the payoffs of the realized outcomes, also denoted $U_i(\sigma)$.

Let $U = (U_1, U_2, \dots, U_n)$.

12.9 Definition: A strategy profile σ is a Nash equilibrium if $U_i(\sigma) \geq U_i(\sigma'_i, \sigma_{-i})$ for any player i and any i 's strategy σ'_i .

12.10 Proposition: If σ is a Nash equilibrium, then $U(\sigma) \in F_+$.

Proof. (1) Fix any Nash equilibrium σ .

(2) $U(\sigma)$ can be represented as a convex combination of $u(a)$ as follows:

$$U(\sigma) = \sum_{a \in A} \mu(a) U(a),$$

where $\mu(a) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \text{Prob}_{\sigma}(a^t(\sigma) = a)$ for each $a \in A$. Thus, we have $U(\sigma) \in F$.

(3) For each i , define σ'_i by $\sigma'_i(h) \in \arg \max_{a_i} u_i(a_i, \sigma_{-i}(h))$ for any h .

(4) Since σ is a Nash equilibrium, we have $U_i(\sigma) \geq U_i(\sigma'_i, \sigma_{-i})$, where

$$U_i(\sigma'_i, \sigma_{-i}) = (1 - \delta) \mathbb{E}_{\sigma_{-i}} \left[\sum_{t=0}^{\infty} \delta^t \max_{a_i} u_i(a_i, \sigma_{-i}(h)) \right] \geq (1 - \delta) \sum_{t=0}^{\infty} \delta^t v_i = v_i.$$

Thus, we have $U(\sigma) \in F_+$.

□

- 12.11 For any history $h^t \in H$, we define the continuation game to be the infinitely repeated games that begins in period t , following history h^t .

For any strategy profile σ , player i 's continuation strategy after h^t , $\sigma_i|_{h^t} : H \rightarrow \Delta(A_i)$, is given by

$$\sigma_i|_{h^t}(h') = \sigma(h^t h') \text{ for any } h' \in H,$$

where $h^t h'$ is the concatenation of the history h^t followed by the history h' .

Note that $\sigma_i|_{h^t} \in \Sigma_i$.

- 12.12 Note that the same stage game is played repeatedly and the stage-game payoffs depend only on actions taken in that period. Thus, history influences only the future actions of the players. This also means that the game is stationary—starting from any period t , any continuation game from period t onwards is exactly the same as the original game. Non-stationary repeated games are difficult to analyze.

- 12.13 Definition: A strategy profile σ is a subgame perfect equilibrium if for any player $i \in N$, any period $t \in \{0, 1, 2, \dots\}$, any history $h^t \in H^t$ and any strategy $\sigma'_i \in \Sigma_i$,

$$U_i(\sigma_i|_{h^t}, \sigma_{-i}|_{h^t}) \geq U_i(\sigma'_i, \sigma_{-i}|_{h^t}).$$

That is, σ is subgame perfect if it induces a Nash equilibrium in every continuation game.

- 12.14 There are a lot of subgames (infinitely many) and for each subgame there are a large number of possible deviations (also infinitely many)! It is impossible to verify a strategy profile is subgame perfect by brute force. In the following we shall go through two fundamental results that are crucial to the analysis of repeated games.

- 12.15 We say that a strategy profile σ does not have profitable one-shot deviations if for any player $i \in N$, any history h , and any action a_i ,

$$U_i(\sigma|_h) \geq (1 - \delta) \cdot u_i(a_i, \sigma_{-i}(h)) + \delta \sum_{a_{-i}} \sigma_{-i}(h)(a_{-i}) \cdot U_i(\sigma|_{ha_i a_{-i}}).$$

Equivalently, for any $\sigma_i \in \Sigma_i$, define

$$\Psi(\sigma_i) = \{\beta_i \in \Sigma_i \mid \beta_i^s = \sigma_i^s \text{ for all } s \geq 1\}.$$

$\Psi(\sigma_i)$ is the set of strategies that are identical to σ_i from period 1 onward.

A strategy profile σ does not have profitable one-shot deviations if for any player $i \in N$, any period $t \in \{0, 1, 2, \dots\}$, and any history $h^t \in H^t$,

$$U_i(\sigma_i|_{h^t}, \sigma_{-i}|_{h^t}) \geq U_i(\sigma'_i, \sigma_{-i}|_{h^t}) \text{ for all } \sigma'_i \in \Psi(\sigma_i|_{h^t}).$$

A strategy profile does not have profitable one-shot deviations if no player can gain by deviating only in the first period of every continuation game including those that are off the equilibrium path of σ . It is weaker than the notion of subgame perfect equilibrium because the latter concept allows players to deviate in more than one periods. There are still an infinite number of subgames to check but at least the number of first period deviations is finite (as the number of stage-game actions is finite).

- 12.16 Theorem (One-shot deviation principle): A strategy profile σ is subgame perfect if and only if it does not have

profitable one-shot deviations.

Idea of proof: It is obvious that a strategy profile has one deviation property if it is a subgame perfect equilibrium. We need to show a strategy profile is a subgame perfect equilibrium if it has one deviation property.

The formal proof is tedious, but the basic idea is quite simple. If a strategy profile is not subgame perfect then some player can obtain a strictly higher payoff, say by $\epsilon > 0$, in some continuation game by deviating. Since future payoffs are discounted, the player must be better off by deviating by only in the first T periods for some finite T . Now look at the last deviation. If it makes the player better off, then the strategy profile does not have one deviation property. If it does not make the player better off, then the player will still be better off without the last deviation. The same argument can be repeated until we find a single beneficial deviation. For a complete proof, see Page 25 in [Mailath and Samuelson \(2006\)](#).

It is probably that Bellman and Shapley in 1950s should take credit for the one-shot deviation principle in the contexts of single-agent dynamic optimization problems and of stochastic games, respectively.

12.17 Observation: If α^* is a Nash equilibrium of the stage game, then the strategy profile “each player i plays α_i^* from now on” is a subgame perfect equilibrium.

Moreover, if the game has m Nash equilibria $\{\alpha^j\}_{j=1}^m$ of the stage game, then for any map $j: \{0, 1, 2, \dots\} \rightarrow \{1, 2, \dots, m\}$, the strategy profile “play $\alpha^{j(t)}$ in period t ” is a subgame perfect equilibrium as well.

12.18 Infinitely repeated prisoners’ dilemma. The stage game is as follows:

	H	C
H	1, 1	-1, 2
C	2, -1	0, 0

Figure 12.3

H is the cooperative action, and C is the uncooperative action. When a game is repeated many times, it seems that some sort of “cooperative” behavior might be induced. We will consider two types of strategies:

- Player i ’s trigger strategy:
 - in period 1, player i plays the cooperative action;
 - in any period t , player i plays the cooperative action if no one has played the uncooperative action in the past; otherwise, plays the uncooperative action.
- Player i ’s tit-for-tat strategy:
 - in period 1, player i plays the cooperative action;
 - in any period t , player i plays the action her opponent chooses in the previous period.

12.2 Trigger strategy equilibrium

12.19 In the infinitely repeated prisoners’ dilemma, player i ’s trigger strategy is as follows:

- in period 1, player i chooses H ;
- in any period t , player i chooses H if no one has chosen C in the past; otherwise, chooses C .

12.20 To check whether the trigger strategy profile is a subgame perfect equilibrium, we need to make sure that no player can deviate once profitably in any subgame. Although there are infinite many of them, subgames all belong to one of two types:

- no player has deviated before and the continuation strategy is the same as the initial strategy;
- some player has deviated before and the continuation strategy for each player to play C in all future periods.

12.21 Consider the subgames of the first type.

- If a player follows the equilibrium strategy, then she will obtain 1 in each period, and her equilibrium payoff is therefore equal to 1.
- If she deviates and play C , she receives 2 in the current period and 0 in all subsequent periods (since according to the equilibrium strategy, they will play C after someone has deviated). The deviation payoff is $(1 - \delta) \cdot 2$.

A player will not deviate if and only if

$$(1 - \delta) \cdot 2 \leq 1,$$

that is, $\delta \geq \frac{1}{2}$.

12.22 Consider the subgames of the second type.

- The players are supposed to play C in all periods, regardless of history. The equilibrium (normalized discounted) payoff for each player is 0.
- If player 1 deviates and play H in the first period and C in all future periods, assuming that the other player is following the equilibrium strategy, her deviation payoff is $(1 - \delta) \cdot (-1)$ which is less than 0.

Hence, player 1 will be worse off if she deviates.

12.23 Hence, this strategy profile is a subgame perfect equilibrium if and only if $\delta \geq 0.5$.

12.24 This is known as a trigger strategy equilibrium. Note that both players playing C is a Nash equilibrium in the stage game. In a trigger strategy, each player begins playing a cooperate action (H in this case) and continue to do so until someone has deviated. Any deviation will “trigger” or cause a shift to the punishment phase where the players play stage-game Nash equilibrium in all future periods.

12.25 Note that in general playing a stage-game Nash equilibrium in every period regardless of history is also a subgame perfect equilibrium. Since by definition deviation from a stage-game Nash equilibrium will not be profitable in the current period, and if the strategy is history dependent, then the current deviation will not increase future payoffs. Hence, the deviation will not be profitable for the whole game. Thus, a trigger strategy profile will be an equilibrium as long as it is unprofitable to deviate from the initial cooperation phase.

12.26 Example [G Section 2.3.C]: Collusion between Cournot duopolies.

Suppose there are 2 firms in a Cournot oligopoly. Inverse demand is given by $P(Q) = a - Q$, where $Q = q_1 + q_2$ and q_i is the quantity to be produced by firm i . Each firm has a constant marginal cost of production, c , and no fixed cost. Consider the infinitely repeated game based on this stage game.

What is the lowest value of δ such that the firms can use trigger strategies to sustain the monopoly output level in a subgame perfect equilibrium?

12.27 Example [G Exercise 2.15]: Cournot model.

Suppose there are n firms in a Cournot oligopoly. Inverse demand is given by $P(Q) = a - Q$, where $Q = q_1 + \dots + q_n$ and q_i is the quantity to be produced by firm i . Each firm has a constant marginal cost of production, c , and no fixed cost. Consider the infinitely repeated game based on this stage game.

What is the lowest value of δ such that the firms can use trigger strategies to sustain the monopoly output level in a subgame perfect equilibrium?

Answer. Calculate firm i 's production and profit in the collusion, Cournot competition, and deviation from punishment cases, respectively:

- Cooperative production and profit: In the collusion, the production is $q_i^c = \frac{a-c}{2n}$, and profit is $\pi_i^c = \frac{(a-c)^2}{4n}$;
- In-cooperative production and profit: In the Cournot competition, production is $q_i^m = \frac{a-c}{n+1}$, and profit is $\pi_i^m = \frac{(a-c)^2}{(n+1)^2}$;
- Deviation production and profit: For each $j \neq i$, firm j produces $q_j^c = \frac{a-c}{2n}$, then firm i can increase its profit by producing $q_i^d = \frac{(n+1)(a-c)}{4n}$, and profit is $\pi_i^d = \frac{(n+1)^2(a-c)^2}{(4n)^2}$.

For each i , consider the following trigger strategy T_i for firm i :

- In period 1 produce q_i^c .
- In period t ($t > 1$), produce q_i^c if every firm j has produced q_j^c in each of the $t - 1$ previous stages; otherwise, produce q_i^m .

Fix firm i , and assume that each other firm $j \neq i$ chooses the trigger strategy T_j . We want to find the condition which guarantees the trigger strategy T_i to be firm i 's best response.

- If firm i does not choose the trigger strategy, then we consider the following two cases:
 - If firm i always chooses the cooperative production q_i^c in every stage game (it is a strategy for firm i , but not the trigger strategy), then the payoff is as same as the payoff when it chooses trigger strategy.
 - If firm i deviates in some period and the profit maximizer is q_i^d . Without loss of generality, we assume that period t is the first period when firm i deviates, then it can get at most π_i^d in this period. From period $(t + 1)$ on, every other firm j will produce in-cooperative production q_j^m . Thus firm i will receive at most π_i^m in each of the subsequent periods, and period t 's present value of its discounted payoff from period t onwards is at most

$$\pi_i^d + \delta\pi_i^m + \delta^2\pi_i^m + \dots = \pi_i^d + \frac{\delta\pi_i^m}{1-\delta}.$$

It is easy to understand when looking at the following table, where $*$ means we do not know exactly the action of firm i in that period.

Period	1	\dots	$t - 1$	t	$t + 1$	$t + 2$	$t + 3$	\dots
Firm $j \neq i$	q_j^c	\dots	q_j^c	q_j^c	q_j^m	q_j^m	q_j^m	\dots
Firm i	q_i^c	\dots	q_j^c	q_i^d	$*$	$*$	$*$	\dots
Firm i 's payoff	π_i^c	\dots	π_i^c	π_i^d	$\leq \pi_i^m$	$\leq \pi_i^m$	$\leq \pi_i^m$	\dots

- If firm i chooses the trigger strategy T_i , then it will receive π_i^c in each period, and period t 's present value of its discounted payoff from period t onwards is

$$\pi_i^c + \delta\pi_i^c + \delta^2\pi_i^c + \dots = \frac{\pi_i^c}{1-\delta}.$$

- In order for firm i to play trigger strategy T_i , we should have

$$\frac{\pi_i^c}{1-\delta} \geq \pi_i^d + \frac{\delta \pi_i^m}{1-\delta},$$

$$\text{that is } \delta \geq \frac{(n+1)^2}{(n+1)^2 + 4n}.$$

Since $\lim_{n \rightarrow \infty} \frac{(n+1)^2}{(n+1)^2 + 4n} = 1$, the lowest value of δ approaches 1. That is, as n increases, a larger δ is required to deter the deviation. In other words, there is more incentive to deviate the trigger strategy. \square

12.28 Example [G Exercise 2.13]: Bertrand duopoly model with homogenous products.

Find $\delta > 0$ such that the trigger strategy is a subgame perfect equilibrium for the game which infinitely repeats the stage game of Bertrand model with homogeneous products described in the lecture.

Answer. Calculate firm i 's price and profit in the collusion, Bertrand competition, and deviation from punishment cases, respectively:

- Cooperative price and profit: In the collusion, the price is $p_i^c = \frac{a+c}{2}$, and profit is $\pi_i^c = \frac{(a-c)^2}{8}$;
- In-cooperative price and profit: In the Bertrand competition, price is $p_i^m = c$, and profit is $\pi_i^m = 0$;
- Deviation price and profit: Firm j 's price is $p_j^c = \frac{a+c}{2}$, firm $i \neq j$ can increase its profit by choosing a price $p_i^d < \frac{a+c}{2}$, but as close as possible to $\frac{a+c}{2}$, and profit is almost equal to monopoly profit $\pi_i^d = \frac{(a-c)^2}{4}$.

For each i , consider the following trigger strategy T_i for firm i :

- In period 1, choose price p_i^c .
- In period t , choose p_i^c if firm j chooses price p_j^c in each of the $t - 1$ previous periods; otherwise, choose price p_i^m .

For any i , assume that firm $j \neq i$ chooses the trigger strategy T_j . We want to find the condition which guarantees the trigger strategy T_i to be firm i 's best response.

- If firm i does not choose the trigger strategy, then we consider the following two cases:
 - If firm i always chooses the cooperative production p_i^c in every stage game (it is a strategy for firm i , but not the trigger strategy), then the payoff is as same as the payoff when it chooses trigger strategy.
 - If firm i deviates in some period and the profit maximizer is p_i^d . Without loss of generality, we assume that period t is the first period when firm i deviates, then it can get at most π_i^d in this period. From period $(t + 1)$ on, firm $j \neq i$ will choose in-cooperative price p_j^m . Thus firm i will receive at most $\pi_i^m = 0$ in each of the subsequent periods, and period t 's present value of its payoff from period t onwards is at most

$$\pi_i^d.$$

It is easy to understand when looking at the following table, where $*$ means we do not know exactly the action of firm i in that period.

Period	1	...	$t - 1$	t	$t + 1$	$t + 2$	$t + 3$...
Firm $j \neq i$	p_j^c	...	p_j^c	p_j^c	p_j^m	p_j^m	p_j^m	...
Firm i	p_i^c	...	p_i^c	p_i^d	*	*	*	...
Firm i 's payoff	π_i^c	...	π_i^c	π_i^d	$\leq \pi_i^m$	$\leq \pi_i^m$	$\leq \pi_i^m$...

- If firm i chooses the trigger strategy T_i , then it will receive π_i^c in each stage, and the present value of its discounted payoff from t -th stage onwards is

$$\pi_i^c + \delta\pi_i^c + \delta^2\pi_i^c + \dots = \frac{\pi_i^c}{1-\delta}.$$

In order for firm i to play trigger strategy T_i , we should have

$$\frac{\pi_i^c}{1-\delta} \geq \pi_i^d,$$

that is $\delta \geq \frac{1}{2}$. □

12.29 Example: Reputation.

The king has borrowed 100 dollars from the lender at 10% interest. The King can repay or renege. The King may need a loan in future periods with a probability b . A loan is worth 125 dollars to the King.

- Lender's strategy: The lender initially provides a loan when needed and continues to do so as long as the king has repaid past loans. If the king ever reneges, then the lender refuses to lend to him again.
- The king's strategy: Repay the initial loan and any future loan if he has always repaid it in the past. If he ever reneges, then he reneges on all future loans.

Lender: If the king has always repaid his loans, then lending to him again yields a payoff of 10 on each loan. If, instead, the king has reneged on a past loan, then, according to the king's strategy, he'll renege on all future loans. In that case, the lender does not want to lend to him. The lender's strategy is then optimal.

King:

- If he ever reneged, then it is indeed optimal for him not to repay a loan, since, according to the lender's strategy, he won't get a future loan regardless of what he does. His payoff from repaying the loan is -110 , while it is zero from reneging.
- Suppose the king has always repaid in the past and has an outstanding loan. If he repays the loan his payoff is

$$-110 + 15b\delta + 15b\delta^2 + \dots = -110 + 15b\frac{\delta}{1-\delta}.$$

The payoff is zero in all periods from reneging.

- It is then optimal for the king to repay the loan when

$$-110 + 15b\frac{\delta}{1-\delta} \geq 0 \iff \delta \geq \frac{110}{110+15b}.$$

The king repays a loan to preserve a reputation conducive to securing more loans. As $b \rightarrow 0$, $\delta \rightarrow 1$. cooperation becomes more difficult if future probability of a loan is low.

12.30 Example: In the infinitely repeated prisoners' dilemma, we define another strategy s_i^* of player i as follows:

- in period 1, play H ;
- in period $t \geq 2$, if the outcome of the previous period is either (C, C) or (H, H) , play H ; otherwise, play C .

Find the condition on δ under which (s_1^*, s_2^*) is a subgame perfect equilibrium.

12.3 Tit-for-tat strategy equilibrium

12.31 A problem with trigger strategies is that they are too unforgiving—once someone has deviated, they will play the in-cooperative action forever. This will not work well in situations where players may deviate by mistake or due to mis-communication.

12.32 A potential alternative to trigger strategies is tit-for-tat. In tit-for-tat:

- in period 1, a player begins with playing H ;
- in all future periods, she will choose the action her opponent chooses in the previous period. That is, a player will cooperate if her opponent cooperate in the last period and will defect if her opponent defect in the last period.

12.33 Tit-for-tat strategy profile has many nice features.

- The strategy is simple, so it is easy for a player to learn that the other player is playing tit-for-tat.
- Unlike trigger strategies, it will not get stuck in the punishment phase forever. It was the winner in the famous repeated prisoners' dilemma tournament conducted by Axelrod in which strategies devised by well-known game theorists were pitched against each other. It did better than many much more complicated strategies.

12.34 Subgames all belong to one of four types:

- the last period's action profile is HH ;
- the last period's action profile is HC ;
- the last period's action profile is CH ;
- the last period's action profile is CC .

12.35 To determine whether tit-for-tat is a subgame perfect equilibrium, we need to check whether the players would want to deviate in any subgame. Since the game is symmetric, we only need to consider player 1.

12.36 Consider the subgame of the first type (the last period's action profile is HH): The equilibrium payoff is 1.

If player 1 deviates, she will get 2 in the current period. The current period outcome will be CH , so the continuation path in the next period will be: HC, CH, HC, \dots . The deviation payoff is therefore equal to

$$(1 - \delta) \cdot (2 - \delta + 2\delta^2 - \delta^3 + \dots) = \frac{2-\delta}{1+\delta}.$$

Player 1 will not deviate if and only if

$$1 \geq \frac{2-\delta}{1+\delta},$$

that is, $\delta \geq \frac{1}{2}$.

12.37 Consider the subgame of the second type (the last period's action profile is HC): The equilibrium payoff is

$$(1 - \delta) \cdot (2 - \delta + 2\delta^2 - \delta^3 + \dots) = \frac{2-\delta}{1+\delta}.$$

Player 1 is supposed to choose C , if she deviates to H , then she will get 1 in the current period. So the continuation path in the next period will be: HH, HH, HH, \dots . So the deviation payoff is 1.

Player 1 will not deviate if

$$\frac{2-\delta}{1+\delta} \geq 1,$$

that is, $\delta \leq \frac{1}{2}$.

12.38 Consider the subgame of the third type (the last period's action profile is CH): The equilibrium payoff is

$$(1 - \delta) \cdot (-1 + 2\delta - \delta^2 + 2\delta^3 - \dots) = \frac{2\delta - 1}{1 + \delta}.$$

Player 1 is supposed to choose H , if she deviates to C , then she will get 1 in the current period. So the continuation path in the next period will be: CC, CC, CC, \dots . So the deviation payoff is 0.

Player 1 will not deviate if

$$\frac{2\delta - 1}{1 + \delta} \geq 0,$$

that is, $\delta \geq \frac{1}{2}$.

12.39 Consider the subgame of the fourth type (the last period's action profile is CC): The equilibrium payoff is 0.

Player 1 is supposed to choose C , if she deviates to H , then she will get -1 in the current period. So the continuation path in the next period will be: CH, HC, CH, \dots . So the deviation payoff is

$$(1 - \delta) \cdot (-1 + 2\delta - \delta^2 + 2\delta^3 - \dots) = \frac{2\delta - 1}{1 + \delta}.$$

Player 1 will not deviate if

$$0 \geq \frac{2\delta - 1}{1 + \delta},$$

that is, $\delta \leq \frac{1}{2}$.

12.40 All four conditions can be satisfied simultaneously if $\delta = 0.5$. Hence, tit-for-tat strategy profile is subgame perfect only if δ is exactly equal to 0.5.

12.41 Intuitively, the problem of tit-for-tat strategy is that it does not distinguish whether a C is played as a deviation or as a punishment. Hence, after one player plays C as a punishment, the other player will play C in the next period, leading to a CH, HC, CH, \dots cycle. The game will not return to the cooperation path HH, HH, HH, \dots .

12.4 Folk theorem

12.42 The “folk theorems” for repeated games asserts that if the players are sufficiently patient then any feasible, individually rational payoffs can be enforced by an equilibrium.

Folk theorems refer to a collection of results concerning the set of feasible subgame perfect equilibrium payoff profiles as the discount factor converges to one. Since the equilibrium payoff set is increasing in δ , these results provide an upper bound on the set of payoffs that can be achieved through intertemporal cooperation.

This is called the “folk theorem” because it was part of game theory’s oral tradition or “folk wisdom” long before it was recorded in print.

✂ 12.43 Theorem (Folk theorem): For every payoffs $u \in U$ with $u_i > \underline{v}_i$ for all $i \in N$, there exists a $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a Nash equilibrium with payoffs u .

12.44 *Proof.* (1) For simplicity we shall assume that there is a pure-strategy profile a such that $u(a) = u$.

(2) Consider the following strategy for each player i :

- play a_i in period 1, and continue to play a_i as long as

- the action profile in the previous period is a ;
 - the action profile in the previous period differed from a in two or more components;
 - if in some period player i is the only one not to follow a , then each player $j \neq i$ plays \underline{a}_j^i for the rest of the game.
- (3) In period t , player i deviates, then she will receive at most $\max_a u_i(a)$ in period t , and at most \underline{v}_i in periods after period t . Thus, player i obtains at most

$$(1 - \delta) \max_a u_i(a) + \delta \underline{v}_i,$$

which is less than u_i as long as δ exceeds $\underline{\delta}_i$ defined by

$$u_i = (1 - \underline{\delta}_i) \max_a u_i(a) + \underline{\delta}_i \underline{v}_i.$$

- (4) Since $u_i > \underline{v}_i$, the solution $\underline{\delta}_i$ is less than 1.
- (5) Taking $\underline{\delta} = \max_i \underline{\delta}_i$ completes the proof.

□

12.45 Under the strategies used in the proof, a single deviation provokes unrelenting punishment. Now, such punishments may be very costly for the punishers to carry out. For example, in a repeated quantity-setting oligopoly, the minmax strategies require player i 's opponents to produce so much output that price falls below player i 's average cost, which may be below their own costs as well. Since minmax punishments can be costly, the question arises if player i ought to be deterred from a profitable one-shot deviation by the fear that her opponents will respond with the unrelenting punishment specified above.

More formally, the point is that the strategies we used to prove the Nash folk theorems are not subgame perfect. This raises the question of whether the conclusion of the folk theorem applies to the payoffs of subgame perfect equilibrium.

12.5 Nash-threats folk theorem

✂ 12.46 Nash-threats folk theorem (Friedman (1971)): Let a^* be a Nash equilibrium of the stage game with payoffs e . Then for any $u \in U$ such that $u_i > e_i$ for all $i \in N$, there is a $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a subgame perfect equilibrium with payoffs u .

12.47 Friedman's theorem shows that any payoff profile that strictly dominates e can be supported by subgame perfect equilibria when the players are sufficiently patient.

12.48 *Proof.* (1) For simplicity we shall assume that there is a pure-strategy profile \hat{a} such that $u(\hat{a}) = u$.

- (2) The following strategy profile, commonly known as a trigger strategy profile, is a subgame perfect equilibrium when δ converges to one.
- Cooperation phase: in period 1, player i chooses \hat{a}_i . Continue to play \hat{a}_i as long as \hat{a} has been chosen in all previous periods. Switch to the punishment phase if some player has deviated from \hat{a} .
 - Punishment phase: player i chooses a_i^* in every period regardless of history.
- (3) Since a^* is a Nash equilibrium of the stage game, to show that the strategy profile is a subgame perfect equilibrium, it is sufficient to show that no player wants to deviate from the cooperation phase.

- (4) Let \bar{u}_i be player i 's maximum stage-game payoff. If player i deviates, she gets at most \bar{u}_i in the current period and e_i in all future periods. Her deviating payoff is at most

$$(1 - \delta) \cdot (\bar{u}_i + \delta e_i + \delta^2 e_i + \cdots) = (1 - \delta)\bar{u}_i + \delta e_i$$

which is less than u_i , player i 's equilibrium payoff in the cooperation phase, for δ sufficiently close to one. \square

12.49 Note that player i 's equilibrium payoff will never be lower than \underline{v}_i . It turns out that for discount factor sufficiently close to one, the converse is also true.

12.50 Example: Infinitely repeated prisoners' dilemma. The stage game is as follows:

	H	C
H	1, 1	-1, 2
C	2, -1	0, 0

Figure 12.4

The set of feasible payoffs is the gray region in Figure 12.2, and Friedman's theorem guarantees that any point in both the gray region and the first quadrant can be achieved as the payoff in a subgame perfect equilibrium of infinitely repeated prisoners' dilemma, provided the discount factor is sufficiently close to one.

12.6 Perfect folk theorem

12.51 Two-person perfect folk theorem (Theorem 1 in [Fudenberg and Maskin \(1986\)](#)): In any two-person infinitely-repeated game, for any $u \in U$ with $u_i > \underline{v}_i$ for all $i = 1, 2$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a subgame perfect equilibrium with payoff profile u .

12.52 Since the equilibrium payoff can not go below \underline{v}_i , this result says that when any strictly individually rational payoff is a subgame perfect equilibrium payoff. Anything is possible in repeated games.

12.53 *Proof.* (1) For simplicity we shall assume that there is a pure strategy profile $a = (a_1, a_2)$ such that $u(a) = u$. Also assume that for $(i, j) = (1, 2), (2, 1)$, the minmax action \underline{a}_i^j is pure.¹ Let $\underline{a} = (\underline{a}_1^2, \underline{a}_2^1)$ denote the "mutual" minmax action profile.

- (2) Pick K sufficiently large so that for any $i = 1, 2$,

$$K(u_i - u_i(\underline{a})) > \max_a u_i(a) - u_i.$$

- (3) Consider the following strategy profile:

- Cooperation phase: play (a_1, a_2) in each period. If either player deviates, the game goes to the punishment phase.
- Punishment phase: play $(\underline{a}_1^2, \underline{a}_2^1)$ for K periods. If no player deviates during the K periods, return to the cooperation phase. If any player deviates, restart the punishment phase.

¹If not, we have to add a "reward phase" to make each player i indifferent over all actions in the support of \underline{a}_i^j .

(4) For cooperation phase, if player i deviates, she gets at most

$$\begin{aligned} & (1 - \delta) \cdot \left(\max_a u_i(a) + \delta u_i(\underline{a}) + \cdots + \delta^K u_i(\underline{a}) + \delta^{K+1} u_i + \delta^{K+2} u_i + \cdots \right) \\ &= (1 - \delta) \max_a u_i(a) + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i. \end{aligned}$$

Player i will not deviate if

$$u_i \geq (1 - \delta) \max_a u_i(a) + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i,$$

that is,

$$\max_a u_i(a) - u_i \leq \underbrace{\delta(1 + \delta + \cdots + \delta^{K-1})}_{\rightarrow K(\delta \rightarrow 1)} (u_i - u_i(\underline{a})). \quad (12.1)$$

Note that Equation (12.1) holds when δ is sufficiently close to one.

(5) For punishment phase, it is sufficient to show that the players do not have incentives to deviate in the beginning of the punishment phase. (She has the highest incentive to do so in the first period of the punishment phase)

Player i 's equilibrium payoff is

$$(1 - \delta^K) u_i(\underline{a}) + \delta^K u_i.$$

If player i deviates, her payoff is at most

$$(1 - \delta) \cdot \left(\underline{v}_i + \delta u_i(\underline{a}) + \cdots + \delta^K u_i(\underline{a}) + \delta^{K+1} u_i + \delta^{K+2} u_i + \cdots \right) = (1 - \delta) \underline{v}_i + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i.$$

Player i will not deviate if

$$(1 - \delta^K) u_i(\underline{a}) + \delta^K u_i \geq (1 - \delta) \underline{v}_i + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i,$$

that is,

$$\underline{v}_i - u_i(\underline{a}) \leq \delta^K (u_i - u_i(\underline{a})). \quad (12.2)$$

Since $u_i > \underline{v}_i \geq u_i(\underline{a})$, we can get δ such that Equation (12.2) holds.

(6) Therefore there exists $\underline{\delta} < 1$, such that for any $\delta \in (\underline{\delta}, 1)$, there is a subgame perfect equilibrium with payoff profile u .

□

12.54 Note that this proof applies only to two-player games as it is generally impossible to minmax more than two players simultaneously.

Example (Example 3 in [Fudenberg and Maskin \(1986\)](#)): In this game, player 1 chooses rows, player 2 chooses columns, and player 3 chooses matrices. Note that whatever one player gets, the others get too.

	A	B
A	1, 1, 1	0, 0, 0
B	0, 0, 0	0, 0, 0

	A	B
A	0, 0, 0	0, 0, 0
B	0, 0, 0	1, 1, 1

For any $\delta < 1$ there does not exist a subgame perfect equilibrium in which the normalized discounted payoff ϵ is less than $\frac{1}{4}$ (the mixed-strategy equilibrium payoff of the stage game).

- ⌘ 12.55 Perfect folk theorem (Theorem 2 in [Fudenberg and Maskin \(1986\)](#)): Assume that the dimensionality of U equals the number of players. Then for any $u \in U$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there exists a subgame perfect equilibrium in which player i 's normalized discounted payoff is u_i .

Proof. (1) For simplicity, assume that there exists a pure action profile $\tilde{a} \in A$ such that $u = u(\tilde{a})$. Also assume that for each i , the minmax action profile $\underline{a}^i = (\underline{a}_i^i, \underline{a}_{-i}^i)$ is pure.

- (2) Since $\dim F = n$, there exist “player-specific punishments” $u^1, u^2, \dots, u^n \in F_{++}$ such that

$$u_i > u_i^i, \text{ and } u_i^j > u_i^i \text{ for all } j \neq i.$$

Again, for simplicity, assume that for each i , there exists a pure action profile $a^i \in A$ such that $u^i = u(a^i)$.

- (3) Pick K sufficiently large so that

$$K(u_i^i - \underline{v}_i) > \max_a u_i(a) - u_i^i$$

for any $i \in N$.

- (4) Consider the following strategy for player i :

- Normal phase: play \tilde{a}_i in each period. If player j deviates, the game goes to the punishment phase j . If multiple players deviate simultaneously, we can ignore such deviations, or punish the “youngest” among all deviators.
- Punishment phase j : play \underline{a}_i^j for K periods and move to reward phase j if no one deviates; if player k deviates, move to the beginning of the punishment phase k . (This rule applies even if $k = j$.)
- Reward phase j : play a_i^j in each period. If player k deviates, move to the beginning of the punishment phase k .²

- (5) For normal phase, if player i deviates, she gets at most

$$\begin{aligned} & (1 - \delta) \cdot \left(\max_a u_i(a) + \delta u_i(\underline{a}) + \dots + \delta^K u_i(\underline{a}) + \delta^{K+1} u_i^i + \delta^{K+2} u_i^i + \dots \right) \\ &= (1 - \delta) \max_a u_i(a) + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i^i. \end{aligned}$$

Player i will not deviate if

$$u_i \geq (1 - \delta) \max_a u_i(a) + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i^i.$$

This inequality holds since RHS approaches u_i^i when δ is sufficiently close to one.

- (6) For punishment phase, it is sufficient to show that the players do not have incentives to deviate in the beginning of the punishment phase. (She has the highest incentive to do so in the first period of the punishment phase) Player i 's equilibrium payoff is

$$(1 - \delta^K) u_i(\underline{a}) + \delta^K u_i^i.$$

If player i deviates, her payoff is at most

$$(1 - \delta) (u_i(a_i, \underline{a}_{-i}^i) + \delta u_i(\underline{a}) + \dots + \delta^K u_i(\underline{a}) + \delta^{K+1} u_i^i + \dots) = (1 - \delta) u_i(a_i, \underline{a}_{-i}^i) + (\delta - \delta^{K+1}) u_i(\underline{a}) + \delta^{K+1} u_i^i.$$

²The reward phase j is not to reward player j , but to reward all other players. Without such a reward phase, player $i \neq j$ would not be willing to play \underline{a}_i^j in punishment phase j since punishing player j can be costly to player i .

Player i will not deviate if

$$(1 - \delta^K)u_i(\underline{a}) + \delta^K u_i^i \geq (1 - \delta)u_i(a_i, \underline{a}_{-i}^i) + (\delta - \delta^{K+1})u_i(\underline{a}) + \delta^{K+1}u_i^i,$$

which is equivalent to

$$\underbrace{(1 - \delta^K)u_i(\underline{a}) + \delta^K u_i^i}_{\rightarrow u_i^i(\delta \rightarrow 1)} \geq u_i(a_i, \underline{a}_{-i}^i).$$

Since $u_i(a_i, \underline{a}_{-i}^i) \leq \underline{v}_i$ and $u_i^i > \underline{v}_i$, this inequality holds when δ is sufficiently close to one.

(7) For reward phase, if player i deviates to a_i in her own reward phase, she gets

$$(1 - \delta)u_i(a_i, \underline{a}_{-i}^i) + (\delta - \delta^{K+1})u(\underline{a}) + \delta^{K+1}u_i^i.$$

Player i will not deviate if

$$u_i^i \geq (1 - \delta)u_i(a_i, \underline{a}_{-i}^i) + (\delta - \delta^{K+1})u(\underline{a}) + \delta^{K+1}u_i^i,$$

which holds when δ is sufficiently close to one due to the choice of T .

□

12.56 Reference: [Abreu, Dutta and Smith \(1994\)](#).

12.57 A game satisfies the non-equivalent utilities condition (abbreviated as “NEU”) if for any i, j there exists no c and d such that $u_i = c + du_j$.

12.58 Folk theorem (Sufficiency): If a game satisfies NEU, then any strictly individually rational payoff can be supported as a subgame perfect equilibrium as δ goes to one.

12.59 A stage game satisfies the condition of no simultaneous minimizing (abbreviated as “NSM”) if no two players can be simultaneously hold to their respective minimal attainable payoffs or below.

12.60 Folk theorem (Necessity): For any repeated game that satisfies NSM, NEU is a necessary condition for folk theorem.

12.7 Finitely repeated games

12.61 Theorem: If the stage game has a unique Nash equilibrium a^* , then for a finitely repeated game (T periods), there is a unique subgame perfect equilibrium that is a repetition of the stage game Nash equilibrium. No cooperation is sustainable.

Proof. By backwards induction, at period T , we will have that (regardless of history) $a^T = a^*$. Given this, then we have $a^{T-1} = a^*$, and continuing inductively, $a^t = a^*$ for each $t = 1, 2, \dots, T$ regardless of history. □

Chapter

13

Extensive games with imperfect information

Contents

13.1 Extensive games with imperfect information	201
13.2 Mixed and behavioral strategies	203
13.3 Subgame perfect equilibrium	206
13.4 Perfect Bayesian equilibrium	212
13.5 Sequential equilibrium	217
13.6 Trembling hand perfect equilibrium	221

13.1 Extensive games with imperfect information

13.1 In extensive games with imperfect information, some player may have only partial information about the history of the play.

We analyze this kind of games by assuming that each player, when choosing an action, forms an expectation about the unknowns.

- These expectations are not derived solely from the players' equilibrium behavior as in strategic games, since the players may face situations inconsistent with that behavior.
- These expectations relate not only to the other players' future behavior as in extensive games with perfect information but also to events that happened in the past.

13.2 Definition: An extensive game with imperfect information is defined as:

$$\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i) \rangle.$$

- A set $N = \{1, 2, \dots, n\}$ of players.
- A set H of sequences that satisfies the following three properties.
 - The empty sequence \emptyset is a member of H .
 - If $(a^k)_{k=1}^K \in H$ (K may be infinite) and $L < K$ then $(a^k)_{k=1}^L \in H$.
 - If an infinite sequence $(a^k)_{k=1}^\infty$ satisfies $(a^k)_{k=1}^L \in H$ for every positive integer L then $(a^k)_{k=1}^\infty \in H$.

Each member of H is a history; each component of a history is an action.

A history $(a^k)_{k=1}^K \in H$ is terminal if it is infinite or if there is no a^{K+1} such that $(a^k)_{k=1}^{K+1} \in H$. The set of terminal histories is denoted Z .

After any non-terminal history h player $P(h)$ chooses an action from

$$A(h) = \{a: (h, a) \in H\}.$$

- A function $P: H \setminus Z \rightarrow N \cup \{c\}$ that assigns to each non-terminal history a member of $N \cup \{c\}$.

P is called the player function, and $P(h)$ is the player who takes an action after the history h .

If $P(h) = c$ then chance determines the action taken after the history h .

- A function f_c that associates with every history h for which $P(h) = c$ a probability measure $f_c(\cdot | h)$ on $A(h)$, where each such probability measure is independent of every other such measure.
- For each player $i \in N$ a partition \mathcal{I}_i of $\{h \in H: P(h) = i\}$ with the property that $A(h) = A(h')$ whenever h and h' are in the same member of the partition.

For $I_i \in \mathcal{I}_i$ we denote by $A(I_i)$ the set $A(h)$ and by $P(I_i)$ the player $P(h)$ for any $h \in I_i$.

\mathcal{I}_i is the information partition of player i ; a set $I_i \in \mathcal{I}_i$ is an information set of player i .

The information set containing the history h is denoted by $I(h)$.

- For each player i a preference relation on Z .

13.3 Example 202.1

13.4 Example Myerson 1991, Section 2.1

- 13.5 Definition: A pure strategy for a player is a complete plan of actions—specifies an action for the player in every contingency; i.e.,

$$s_i: \mathcal{I}_i \rightarrow \cup_{I_i \in \mathcal{I}_i} A(I_i)$$

such that $s_i(I_i) \in A(I_i)$.

s_i can be rewritten as a vector $(s_i(I_i))_{I_i \in \mathcal{I}_i}$.

Denoted by S_i the set of player i 's strategies.

13.6 We refer to games in which at every point every player remembers whatever he knew in the past as games with perfect recall.

Let $\langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i) \rangle$ be an extensive game and let $X_i(h)$ be the record of player i 's experience along the history h : $X_i(h)$ is the sequence consisting of the information sets that the player encounters in the history h and the actions that he takes at them, in the order that these events occur.

An extensive game has perfect recall if for each player i we have $X_i(h) = X_i(h')$ whenever the histories h and h' are in the same information set of player i .

13.7 Examples of extensive games with imperfect recall.




Figure 13.1: Extensive games with imperfect recall.

13.2 Mixed and behavioral strategies

13.8 Definition: A mixed strategy of player i in an extensive game $\langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i) \rangle$ is a probability measure over S_i , usually denoted by α_i .

A behavioral strategy of player i is a collection $\beta_i = (\beta_i(I_i))_{I_i \in \mathcal{I}_i} \in \times_{I_i \in \mathcal{I}_i} \Delta A(I_i)$ of independent probability measures, where $\beta_i(I_i)$ is a probability measure over $A(I_i)$.

We denote by $\beta_i(I_i)(a)$ the probability assigned to the action $a \in A(I_i)$.

13.9 Differences between mixed strategies and behavior strategies: Consider the following game.




Figure 13.2

For player 2, there are four pure strategies: (A, C) , (A, D) , (B, C) , (B, D) . Hence, each player 2's mixed strategy can be written as $\alpha \circ (A, C) + \beta \circ (A, D) + \gamma \circ (B, C) + (1 - \alpha - \beta - \gamma) \circ (B, D)$.

Each player 2's behavioral strategy can be represented by $(x \circ A + (1 - x) \circ B, y \circ C + (1 - y) \circ D)$.

13.10 ????? Hart 1992 handbook chap 2 page 32, Remark: A useful way of viewing the difference between mixed and behavior strategies is as follows. One can think of each pure strategy as a book of instructions, where for each of the player's information sets there is on page which states what choice he should make at that information set. The player's set of pure strategies is a library of such books. A mixed strategy is a probability distribution on his library of books, so that, in playing according to a mixed strategy, the player chooses on book from his library by means of a chance device having the prescribed probability distribution. A behavior strategy is a single book of a different sort. Although each page still refers to a single information set of the player, it specifies a probability distribution over the choices at that set, not a specific choice.

13.11 For any profile $\sigma = (\sigma_i)_{i \in N}$ of either mixed or behavioral strategies in an extensive game, we define the outcome $O(\sigma)$ of σ to be the probability distribution over the terminal histories that results when each player i follows the precepts of σ_i .

13.12 Two (mixed or behavioral) strategies of any player are outcome-equivalent if for every collection of pure strategies of the other players the two strategies induce the same outcome.

13.13 Observation: For any behavioral strategy β_i of any player i in extensive games with perfect recall, the mixed strategy defined as follows is outcome-equivalent: the probability assigned to any pure strategy s_i (which specifies an action $s_i(I_i)$ for every information set $I_i \in \mathcal{I}_i$) is $\prod_{I_i \in \mathcal{I}_i} \beta_i(I_i)(s_i(I_i))$.

In Example 13.9, for the behavioral strategy $(x \circ A + (1-x) \circ B, y \circ C + (1-y) \circ D)$, the outcome-equivalent mixed strategy can be defined as:

$$xy \circ (A, C) + x(1-y) \circ (A, D) + (1-x)y \circ (B, C) + (1-x)(1-y) \circ (B, D).$$

13.14 Note that the derivation of the mixed strategy relies on the assumption that the collection $(\beta_i(I_i))_{I_i \in \mathcal{I}_i}$ is independent.

Consider the following imperfect recall game. The behavioral strategy that assigns probability $p \in (0, 1)$ to a to a generates the outcomes (a, a) , (a, b) , and b with probabilities p^2 , $p(1-p)$, and $1-p$ respectively, a distribution that can not be duplicated by any mixed strategy: the set of pure strategies are $\{a, b\}$. Any mixed strategy can be represented by $q \circ a + (1-q) \circ b$, which generates the outcomes (a, a) and b with probabilities q and $1-q$ respectively.




Figure 13.3

13.15 Proposition (Kuhn, 1950 and 1953): For any mixed strategy of a player in a finite extensive game with perfect recall there is an outcome-equivalent behavioral strategy.

13.16 Proof. (1) For any history $h = (a^1, a^2, \dots, a^k)$ define a pure strategy s_i of player i to be consistent with h if for every subhistory (a^1, a^2, \dots, a^l) of h for which $P(a^1, a^2, \dots, a^l) = i$ we have $s_i(a^1, a^2, \dots, a^l) = a^{l+1}$.

(2) Let σ_i be a mixed strategy of player i . For any history h let $\pi_i(h)$ be the sum of the probabilities according to σ_i of all the pure strategies of player i that are consistent with h .

(3) Let h and h' be two histories in the same information set I_i of player i , and let $a \in A(h)$. Since the game has perfect recall, the sets of actions of player i in h and h' are the same. Thus $\pi_i(h) = \pi_i(h')$.

(4) Since in any pure strategy of player i the action a is taken after h if and only if it is taken after h' , we also have $\pi_i(h, a) = \pi_i(h', a)$.

(5) Define a behavioral strategy β_i of player i :

$$\beta_i(I_i)(a) = \begin{cases} \frac{\pi_i(h, a)}{\pi_i(h)}, & \text{if } \pi_i(h) > 0, \\ \text{immaterial}, & \text{otherwise,} \end{cases}$$

where $h \in I_i$.

(6) Let s_{-i} be a collection of pure strategies for the players other than i . Let h be a terminal history.

- If h includes moves that are inconsistent with s_{-i} then the probability of h is zero under both α_i and β_i . Now assume that all the moves of players other than i in h are consistent with s_{-i} .

- If h includes a move after a subhistory $h' \in I_i$ of h that is inconsistent with σ_i then $\beta_i(I_i)$ assigns probability zero to this move, and thus the probability of h according to β_i is zero.
- If h is consistent with σ_i then $\pi_i(h') > 0$ for all subhistories h' of h and the probability of h according to β_i is the product of $\frac{\pi_i(h', a)}{\pi_i(h')}$ over all (h', a) that are subhistories of h ; this product is $\pi_i(h)$, the probability of h according to σ_i .

□

13.17 Remark: Kuhn's Theorem allows us to focus on behavioral strategies.

13.18 Consider Example 13.9 again. For a mixed strategy $\alpha \circ (A, C) + \beta \circ (A, D) + \gamma \circ (B, C) + (1 - \alpha - \beta - \gamma) \circ (B, D)$, the outcome-equivalent behavioral strategy can be defined as follows:


$$((\alpha + \beta) \circ A + (1 - \alpha - \beta) \circ B, (\alpha + \gamma) \circ C + (1 - \alpha - \gamma) \circ D).$$

13.19 A Nash equilibrium of an extensive game is a strategy profile σ^* with the property that for every player $i \in N$ we have

$$O(\sigma_{-i}^*, \sigma_i^*) \succsim_i O(\sigma_{-i}^*, \sigma_i) \text{ for every strategy } \sigma_i \text{ of player } i.$$

Here σ^* could refer to pure-strategy profile, mixed-strategy profile, or behavioral-strategy profile.

13.20 In the following game, the unique pure-strategy Nash equilibrium is (L, R') . When adopting the strategy profile (L, R') , player 2's information set is not reached.



		Player 2	
		L'	R'
Player 1	L	2, 2	2, 2
	M	3, 1	0, 2
	R	0, 2	1, 1

But in this case player 2's optimal action in the event that his information set is reached depends on his belief about the history that has occurred. The action R' is optimal if he assigns probability of at least $\frac{1}{2}$ to the history M , while L' is optimal if he assigns probability of at most $\frac{1}{2}$ to this history.

Thus his optimal action depends on his explanation of the cause of his having to act. His belief can not be derived from the equilibrium strategy, since this strategy assigns probability zero to his information set being reached.

13.21 The solutions for extensive games that we have studied so far have a single component: a strategy profile. We will study solutions that consist of both a strategy profile and a belief system, where a belief system specifies, for each information set, the beliefs held by the players who have to move at that information set about the history that occurred.

It is natural to include a belief system as part of the equilibrium, given our interpretation of the notion of subgame perfect equilibrium. When discussing this notion of equilibrium we argue that to describe fully the players' rea-

soning about a game we have to specify their expectations about the actions that will be taken after histories that will not occur if the players adhere to their plans, and that these expectations should be consistent with rationality.

13.3 Subgame perfect equilibrium

13.22 Definition: In an extensive game with imperfect information, the subgame after history $h \in H$

- begins at a history h that is a singleton information set (but is not the game's initial history), that is, $I(h) = \{h\}$;
- includes all the histories following h (but no histories that do not follow h);
- does not cut any information sets (*i.e.*, if h' is a history following h , and h'' is in the information set containing h' , then h'' also follows h).

13.23 Definition: A strategy profile σ is a subgame perfect equilibrium (abbreviated as "SPE") of the extensive game with imperfect information Γ if σ induces a Nash equilibrium in every subgame in Γ . Here σ^* could refer to pure-strategy profile, mixed-strategy profile, or behavioral-strategy profile.

13.24 Example (nonexistence, [Luttmer and Mariotti \(2003\)](#)): The example has five stages.

- In stage 1, player 1 chooses $a_1 \in [0, 1]$.
- In stage 2, player 2 chooses $a_2 \in [0, 1]$.
- In stage 3, Nature chooses x by randomizing uniformly over the interval $[-2 + a_1 + a_2, 2 - a_1 - a_2]$.
- After this, players 3 and 4 move sequentially.

The subgame following a history (a_1, a_2, x) and the associated payoffs for all four players are shown in the following figure. This game does not have a subgame perfect equilibrium.




Figure 13.4

13.25 Example: Consider the following game of complete information, where the three numbers below each terminal node are the payoffs to player 1, player 2, and player 3 from top to bottom.

- How many information sets does player 3 have?
- How many pure strategies does player 3 have? What are they?
- How many subgames do you find in the above game?
- Find all the pure-strategy Nash equilibria for the game.
- Identify those pure-strategy Nash equilibria which are subgame perfect or not.




Figure 13.5

- Answer.* (i) Player 3 has two information sets: One is a non-singleton information set and the other is a singleton information set.
- (ii) Player 3 has four strategies: SX , SY , TX , and TY .
- (iii) Only one subgame.
- (iv) In the following tables, one can find all the pure-strategy Nash equilibrium in the entire game. If player 1 chooses L :

		Player 3			
		SX	SY	TX	TY
Player 2	AC	0, 20, 20	0, 20, 20	0, 0, 0	0, 0, 0
	AD	0, 20, 20	0, 20, 20	0, 0, 0	0, 0, 0
	BC	8, 8, 8	8, 8, 8	12, 12, 12	12, 12, 12
	BD	8, 8, 8	8, 8, 8	12, 12, 12	12, 12, 12

If player 1 chooses R :

		Player 3			
		SX	SY	TX	TY
Player 2	AC	10, 5, 10	10, 5, 10	10, 15, 10	10, 15, 10
	AD	15, 10, 5	5, 10, 15	15, 10, 5	5, 10, 15
	BC	10, 5, 10	10, 5, 10	10, 15, 10	10, 15, 10
	BD	15, 10, 5	5, 10, 15	15, 10, 5	5, 10, 15

Subgame perfect equilibria: (L, BC, TY) , (L, BD, TY) , (R, AC, TY) , and (R, AD, SY) .

Not subgame perfect equilibria: (L, BC, TX) and (R, AC, TX) .

□

13.26 Example: Consider the following game. In the game tree, the numbers at the top are payoffs to player 1 and the numbers at the bottom are payoffs to player 2, as usual.

- (i) How many subgames are there in this game, not counting the whole game as one subgame?
- (ii) If $x \geq x_0$, then there are two pure-strategy subgame perfect equilibria. If $x < x_0$, then there is only one pure-strategy subgame perfect equilibrium. Find this threshold value x_0 . What are the two pure-strategy subgame perfect equilibria if $x \geq x_0$?

Answer. (i) 0. Thus, every Nash equilibrium is a subgame perfect equilibrium.

- (ii) Consider the following payoff table:




Figure 13.6

		Player 2	
		L'	R'
Player 1	L	$x, 1$	$-2, 0$
	M	$2, 0$	$-1, 1$
	R	$1, 1$	$1, 1$

If $x \geq 2$, then there are two pure-strategy subgame perfect equilibria (L, L') and (R, R') . Otherwise, there is unique pure-strategy subgame perfect equilibrium (R, R') . Therefore $x_0 = 2$.

□

13.27 Example [G Exercise 2.6]: Market with three oligopolists.

Three oligopolists operate in a market with inverse demand given by $P(Q) = a - Q$, where $Q = q_1 + q_2 + q_3$ and q_i is the quantity produced by firm i . Each firm has a constant marginal cost of production, c , and no fixed cost. The firms choose their quantities as follows:

- firm 1 chooses $q_1 \geq 0$;
- firms 2 and 3 observe q_1 and then simultaneously choose q_2 and q_3 , respectively.

What is the pure-strategy subgame perfect outcome?




Figure 13.7

Answer. Figure 13.7 is the game tree. Given q_1 , suppose $q_1 \leq a - c$ (otherwise firm 1's payoff is non-positive), which implies $a - q_1 \geq c$. Note that firms 2 and 3's pure strategies are both functions of q_1 .

The second stage is exactly a Cournot model of duopoly, with total demand $a' = a - q_1$, and marginal cost $c_2 = c_3 = c \leq a - q_1$. Therefore the unique Nash equilibrium is $(q_2^*(q_1), q_3^*(q_1)) = (\frac{a - q_1 - c}{3}, \frac{a - q_1 - c}{3})$.

For firm 1, consider the following optimization problem

$$\max_{q_1 \leq a - c} q_1(a - c - q_1 - q_2^*(q_1) - q_3^*(q_1)) = \max_{q_1 \leq a - c} \frac{1}{3} q_1(a - q_1 - c),$$

which has a unique maximizer $q_1^* = \frac{a-c}{2}$. Hence $q_2^* = q_3^* = \frac{a-c}{6}$.

Hence, the subgame perfect outcome is: firm 1 chooses $\frac{a-c}{2}$ in the first stage, and firms 2 and 3 choose $\frac{a-c}{6}$. \square

- 13.28 Example: Consider strategic investment in a duopoly model. Firm 1 and firm 2 currently both have a constant average cost of 2 per unit. Firm 1 can install a new technology with an average cost of 0 per unit; installing the technology costs 8. Firm 2 will observe whether or not firm 1 invests in the new technology. Once firm 1's investment decision is observed, the two firms will simultaneously choose output levels q_1 and q_2 as in Cournot model. Here let the price be $P(Q) = 14 - Q$ if $Q < 14$ and 0 otherwise. What is the pure-strategy subgame perfect outcome of the game?




Figure 13.8

Answer. Figure 13.8 is the extensive-form representation of the game. There are 2 stages:

- In the first stage, firm 1 choose “Install” or “Not install”;
 - In the second stage, firms 1 and 2 play the Cournot Duopoly Game.
- (a) If firm 1 chooses “Install” in the first stage, then $a = 14$, $c_1 = 0$, $c_2 = 2$. Since $0 \leq c_i < \frac{a}{2}$, by Example 2.31, the unique Nash equilibrium is $(q_1^*, q_2^*) = (\frac{a-2c_1+c_2}{3}, \frac{a-2c_2+c_1}{3}) = (\frac{16}{3}, \frac{10}{3})$, and firm 1's payoff is $\frac{16}{3}(14 - \frac{16}{3} - \frac{10}{3}) - 8 = 20\frac{4}{9}$.
- (b) If firm 1 chooses “Not install” in the first stage, then $a = 14$, $c_1 = c_2 = 2 < a$, and the unique Nash equilibrium is $(q_1^*, q_2^*) = (\frac{a-c}{3}, \frac{a-c}{3}) = (4, 4)$. Firm 1's payoff is $4(14 - 8 - 2) = 16$.

Since $16 < 20\frac{4}{9}$, the subgame perfect outcome is: firm 1 chooses “Install” in the first stage, and firms 1 and 2 choose $\frac{16}{3}$ and $\frac{10}{3}$, respectively in the second stage. \square

- 13.29 Example: Three pirates jointly own 6 coins. They have to decide on an allocation which exhausts the coins. They decide they should be democratic and choose the following rule: The oldest pirate proposes an allocation. If at least half approve, it is enforced. Otherwise, the oldest pirate is executed. The same procedure is then followed with the oldest pirate of the remainder proposing.

Assume that the pirates are perfectly distinguishable by seniority. Find all the pure-strategy subgame perfect outcomes.

Answer. Denote by player i the i -th oldest pirate. Let $a = (a_1, a_2, a_3)$ denote an allocation, where $a_1 + a_2 + a_3 = 6$ and $a_i = 1, 2, \dots, 6$. We assume the proposer always says “yes”. There are three subgame perfect outcomes: $(6, 0, 0)$, $(5, 0, 1)$ and $(5, 1, 0)$. The game tree is as follows:

Note that, if the allocation a is rejected in the first period, then the second oldest pirate can enforce any allocation a' in the second period (and, hence, the second oldest pirate gets all the coins and the youngest pirate gets nothing). So, in the second period, the second oldest pirate will choose the best allocation $a' = (0, 6, 0)$ that will be enforced. We need only care about the strategies in the first period.




Figure 13.9

To support the outcome $(5, 1, 0)$, for example, we may consider the following strategies by which the outcome $(5, 1, 0)$ is enforced in the first period:

$$s_1^* = (5, 1, 0), \quad s_2^* = \begin{cases} Y, & \text{if } s_1 = (0, 6, 0) \\ N, & \text{if } s_1 \neq (0, 6, 0) \end{cases}, \quad s_3^* = \begin{cases} Y, & \text{if } s_1 \neq (6, 0, 0) \\ N, & \text{if } s_1 = (6, 0, 0) \end{cases}.$$

Since $(6, 0, 0)$ can not be approved by the specified strategies, the oldest pirate can get at most 5 coins and, therefore, has no incentive to deviate from $s_1^* = (5, 1, 0)$. Since the second oldest pirate will get all the coins and the youngest pirate will get nothing if the game goes to the second period, both players 2 and 3 have no incentive to deviate, respectively, from s_2^* and s_3^* . (Intuitively, since the second oldest pirate will get all the coins if the game goes to the second period, the second oldest pirate may not accept the offer except $a = (0, 6, 0)$; since the youngest pirate will get nothing if the game goes to the second period, the youngest pirate may accept any offer.) Thus, the outcome $(5, 1, 0)$ can be supported by a subgame perfect equilibrium.

Analogous discussions also apply to the other two outcomes.

$$s_1^* = (5, 0, 1), \quad s_2^* = \begin{cases} Y, & \text{if } s_1 = (0, 6, 0) \\ N, & \text{if } s_1 \neq (0, 6, 0) \end{cases}, \quad s_3^* = \begin{cases} Y, & \text{if } s_1 \neq (6, 0, 0) \\ N, & \text{if } s_1 = (6, 0, 0) \end{cases}.$$

$$s_1^* = (6, 0, 0), \quad s_2^* = \begin{cases} Y, & \text{if } s_1 = (0, 6, 0) \\ N, & \text{if } s_1 \neq (0, 6, 0) \end{cases}, \quad s_3^* \equiv Y.$$

□

13.30 Example: Consider the following Bayesian game.

Nature selects Game 1 with probability $\frac{9}{13}$, Game 2 with probability $\frac{3}{13}$, and Game 3 with probability $\frac{1}{13}$.

Player 1 learns whether nature has selected Game 1 or not; player 2 learns whether nature has selected Game 3 or not.

Players 1 and 2 simultaneously choose their actions: player 1 chooses either T or B , and player 2 chooses either L or R .

Payoffs are given by the game selected by nature. In the cells of the tables, the first number is the payoff to player 1 and the second number is the payoff to player 2.

	<i>L</i>	<i>R</i>
<i>T</i>	2, 2	0, 0
<i>B</i>	3, 0	1, 1

Game 1

	<i>L</i>	<i>R</i>
<i>T</i>	2, 2	0, 0
<i>B</i>	0, 0	1, 1

Game 2

	<i>L</i>	<i>R</i>
<i>T</i>	2, 2	0, 0
<i>B</i>	0, 0	1, 1

Game 3

All of this is common knowledge. Notice that Game 2 and Game 3 are identical.

- (i) Represent the above Bayesian game as an extensive game.
- (ii) How many information sets does player 1 have?
- (iii) How many information sets does player 2 have?
- (iv) How many pure strategies does player 1 have? What are they?
- (v) How many pure strategies does player 2 have? What are they?
- (vi) Find the pure-strategy Bayesian Nash equilibrium. You should specify the strategy profile. Please show your work.

Answer. (i) The game tree is given in Figure 13.10.




Figure 13.10

- (ii) Two information sets for player 1.
- (iii) Two information sets for player 2.
- (iv) Four pure strategies for player 1: TT , TB , BT , and BB (for each pure strategy, the first letter is the action to take knowing that nature has selected Game 1, and the second letter is the action to take knowing that nature has selected Game 2 or 3).
- (v) Four pure strategies for player 2: LL , LR , RL , and RR (for each pure strategy, the first letter is the action to take knowing that nature has selected Game 1 or 2, and the second letter is the action to take knowing that nature has selected Game 3).
- (vi) The unique pure-strategy Bayesian Nash equilibrium is (BB, RR) :

The first number is the payoff to type $\{1\}$ of player 1, the second number is the payoff to type $\{2, 3\}$ of player 1, the third number is the expected payoff to type $\{1, 2\}$ of player 2, and the fourth number is the expected payoff to type $\{3\}$ of player 2 in each cell.

Remark: One may find the Bayesian Nash equilibrium by the following argument. Player 1 of type $\{1\}$ must choose B because it is the strictly dominant action in Game 1. Then player 2 of type $\{1, 2\}$ must choose R as his best

		Player 2			
		<i>LL</i>	<i>LR</i>	<i>RL</i>	<i>RR</i>
Player 1	<i>TT</i>	2, 2, 2, 2	2, $\frac{3}{2}$, 2, 0	0, $\frac{1}{2}$, 0, 2	0, 0, 0, 0
	<i>TB</i>	2, 0, $\frac{3}{2}$, 0	2, $\frac{1}{4}$, $\frac{3}{2}$, 1	0, $\frac{3}{4}$, $\frac{1}{4}$, 0	0, 1, $\frac{1}{4}$, 1
	<i>BT</i>	3, 2, $\frac{1}{2}$, 2	3, $\frac{3}{2}$, $\frac{1}{2}$, 0	1, $\frac{1}{2}$, $\frac{3}{4}$, 2	1, 0, $\frac{3}{4}$, 0
	<i>BB</i>	3, 0, 0, 0	3, $\frac{1}{4}$, 0, 1	1, $\frac{3}{4}$, 1, 0	1, 1, 1, 1

response regardless of player 1's choice in Game 2. Then player 1 of type $\{2, 3\}$ must choose *B* as his best response regardless of player 2's choice in Game 3. Finally, player 2 of type $\{3\}$ should choose *R* because it is his best response to player 2's choice of *B* in Game 3. \square

13.4 Perfect Bayesian equilibrium

13.31 Consider the subgame perfect equilibria of the following game.




Figure 13.11: Motivation of perfect Bayesian equilibrium

The induced strategic game is


		Player 2	
		<i>L'</i>	<i>R'</i>
Player 1	<i>L</i>	2, 1	0, 0
	<i>M</i>	0, 2	0, 1
	<i>R</i>	1, 3	1, 3

There are 2 Nash equilibria: (L, L') and (R, R') . Note that the above game has no subgames. Thus both (L, L') and (R, R') are subgame perfect equilibria.

However, (R, R') is based on a non-credible threat: if player 1 believes that player 2's threat of playing R' , then player 1 indeed should choose *R* to end the game with payoff 1 for himself and 3 for player 2 since choosing *L* or *M* will give him 0.

On the other hand, if player 1 does not believe the threat by playing *L* or *M*, then player 2 gets the move and chooses L' . Since L' strictly dominates R' for player 2. The threat of playing R' from player 2 is indeed non-credible.

13.32 For a given equilibrium in a given extensive game, an information set is on the equilibrium path if it will be reached with positive probability if the game is played according to the equilibrium strategies, and is off the equilibrium path if it is certain not to be reached if the game is played according to the equilibrium strategies, where "equilibrium" can mean Nash, subgame perfect, Bayesian, or perfect Bayesian equilibrium.

 13.33 A perfect Bayesian equilibrium (abbreviated as "PBE") is a strategy profile σ and a belief system $\mu = (\mu_1, \mu_2, \dots, \mu_n)$, where μ_i specifies i 's belief at each of his information sets, such that for every player i ,

- Sequential rationality: At each of his information sets, σ_i is a best response to σ_{-i} , given his belief μ_i at that information set.
- Belief consistency: At information sets on the equilibrium path, his belief μ_i is derived from Bayes' rule using the strategy profile σ .
- Belief consistency +: At information sets off the equilibrium path, his belief μ_i is derived from Bayes' rule using the strategy profile σ where possible.

13.34 (σ, μ) is a weak perfect Bayesian equilibrium (abbreviated as “WPBE”) if it satisfies “sequential rationality” and “belief consistency”.

13.35 Bayes' rule: Given an information set I , where this information set contains n histories: k_1, k_2, \dots, k_n , if the decision node k_i will be reached with probability p_i for each $i = 1, 2, \dots, n$, then the belief on this information set should be as follows:

- (a) If $p_1 + p_2 + \dots + p_n \neq 0$, then the player with the move should believe that the history k_i has been reached with probability

$$\frac{p_i}{p_1 + p_2 + \dots + p_n}.$$

- (b) If $p_1 + p_2 + \dots + p_n = 0$, the belief can be arbitrary.

13.36 It is clear that in a perfect Bayesian equilibrium (σ, μ) , σ is a subgame perfect equilibrium.

13.37 Theorem: There exists a (possibly mixed) perfect Bayesian equilibrium for a finite extensive game with perfect recall.

Idea of proof: Backwards induction starting from the information sets at the end ensures perfection, and one can construct a belief system supporting these strategies, so the result is a perfect Bayesian equilibrium.

13.38 Consider the game in Figure 13.11 again. We assume player 2 to believe that L has been played by player 1 with probability p , shown in Figure 13.12.




Figure 13.12

Given this belief, we can compute player 2's expected payoff:

$$\begin{cases} p \cdot 1 + (1 - p) \cdot 2 = 2 - p, & \text{if playing } L', \\ p \cdot 0 + (1 - p) \cdot 1 = 1 - p, & \text{if playing } R'. \end{cases}$$

R' is not optimal at the information set with any belief. Thus (R, R') is not a perfect Bayesian equilibrium.

13.39 Remark on “belief consistency +”: Consider the following game.

- (1) The induced strategic game is




Figure 13.13

	LL'	LR'	RL'	RR'
A	2, 0, 0	2, 0, 0	2, 0, 0	2, 0, 0
D	1, 2, 1	3, 3, 3	0, 1, 2	0, 1, 1

There are four Nash equilibria: (A, L, L') , (A, R, L') , (A, R, R') and (D, L, R') .

- (2) The game has a unique subgame: it begins at player 2's only decision node.

We can represent this subgame as the following strategic game:

		Player 3	
		L'	R'
Player 2	L	2, 1	3, 3
	R	1, 2	1, 1

The subgame has a unique Nash equilibrium: (L, R') . Therefore the unique subgame perfect equilibrium is (D, L, R') .

- (3) Consider the Nash equilibrium (A, L, L') . To support L' to be optimal for player 3, p should satisfy

$$p \cdot 1 + 2 \cdot (1 - p) \geq p \cdot 3 + (1 - p) \cdot 1,$$

that is $p \leq \frac{1}{3}$. Thus (A, L, L') and $p \in [0, \frac{1}{3}]$ satisfy the requirements “sequential rationality” and “belief consistency”.

- (4) However, requirement “belief consistency +” can rule out (A, L, L') and $p \in [0, \frac{1}{3}]$ as a perfect Bayesian equilibrium. Bayes' rule results in $p = 1$ since player 2 plays strategy L although it is off the path.

13.40 Example: Selten's horse.

- The induced strategic game is




Figure 13.14: Selten's horse.

	cL	cR	dL	dR
C	1, 1, 1	1, 1, 1	4, 4, 0	0, 0, 1
D	3, 3, 2	0, 0, 0	3, 3, 2	0, 0, 0

There are two pure-strategy Nash equilibria: (C, c, R) and (D, c, L) .

- For the Nash equilibrium (D, c, L) , belief consistency implies $p = 1$. When $p = 1$, L is a best response for player 3 given the belief $(1, 0)$ on his information set.
- To support the Nash equilibrium (C, c, R) to be a perfect Bayesian equilibrium, we have $1 - p \geq 2p$, that is, $p \leq \frac{1}{3}$.

13.41 Perfect Bayesian equilibrium is a relatively weak equilibrium concept for dynamic games of incomplete information. It is often strengthened by restricting beliefs on information sets that are not reached along the equilibrium path.

13.42 Example [G Exercise 4.10].

Two partners must dissolve their partnership. Partner 1 currently owns share s of the partnership, partner 2 owns share $1 - s$. the partners agree to play the following game: partner 1 names a price, p , for the whole partnership, and partner 2 then chooses either to buy 1's share for ps or to sell his or her share to 1 for $p(1 - s)$. Suppose it is common knowledge that the partners' valuations for owning the whole partnership are independently and uniformly distributed on $[0, 1]$, but that each partner's valuation is private information. What is the perfect Bayesian equilibrium?

Answer. Figure 13.15 is the game tree.

For $v_1 \in [0, 1]$, partner 1's maximization problem is:

$$\max_p [v_1 - p(1 - s)] \text{Prob}(v_2 - ps \leq p(1 - s)) + ps[1 - \text{Prob}(v_2 - ps \leq p(1 - s))].$$

Since $\text{Prob}(v_2 - ps \leq p(1 - s)) = p$, partner 1's maximization problem becomes:

$$\max_p [v_1 - p(1 - s)]p + ps(1 - p).$$

By first order condition, we have $p^* = \frac{v_1 + s}{2}$. Therefore, the perfect Bayesian equilibrium is: for $v_1, v_2 \in [0, 1]$,

$$s_1^*(v_1) = p^* = \frac{v_1 + s}{2}, \quad s_2^*(v_2 | p) = \begin{cases} \text{sell,} & \text{if } v_2 \leq p \\ \text{buy,} & \text{if } v_2 > p \end{cases},$$

partner 1's belief about the partner 2's valuation is a uniform distribution on $[0, 1]$, and partner 2's belief about the partner 1's valuation is a uniform distribution on $[0, 1]$. \square




Figure 13.15

13.43 Example [G Exercise 4.11].

A buyer and a seller have valuations v_b and v_s . It is common knowledge that there are gains from trade (*i.e.*, that $v_b > v_s$), but the size of the gains is private information, as follows: the seller's valuation is uniformly distributed on $[0, 1]$; the buyer's valuation $v_b = k \cdot v_s$, where $k > 1$ is common knowledge; the seller knows v_s (and hence v_b) but the buyer does not know v_b (or v_s). Suppose the buyer makes a single offer, p , which the seller either accepts or rejects. What is the perfect Bayesian equilibrium when $k < 2$? When $k > 2$?

Answer. Figure 13.16 is the game tree.




Figure 13.16

Clearly, the buyer has no incentive to offer $p > 1$, since the seller will accept $p \geq v_s$ and v_s is uniformly distributed on $[0, 1]$.

- By backwards induction, the seller's best response is

$$s_s^*(v_s | p) = \begin{cases} \text{accept,} & \text{if } v_s \leq p \\ \text{reject,} & \text{if } v_s > p \end{cases}.$$

Note that we assume seller will accept if $v_s = p$. This will not affect the our analysis of the game since the probability is zero for $v_s = p$.

- The buyer's maximization problem is:

$$\max_{0 \leq p \leq 1} E[v_b - p | v_s \leq p].$$

Since $v_b = kv_s$, the buyer's maximization problem is:

$$\max_{0 \leq p \leq 1} \int_0^p (kv_s - p) dv_s = \max_{0 \leq p \leq 1} (k/2 - 1)p^2.$$

Therefore, the maximizer is

$$p^* = \begin{cases} 1, & \text{if } k > 2 \\ 0, & \text{if } k < 2 \end{cases}.$$

- Each information set of buyer is reached, so buyer's belief is a uniform distribution on $[0, 1]$.

To summarize, the perfect Bayesian equilibrium is:

$$s_b^* = p^* = \begin{cases} 1, & \text{if } k > 2 \\ 0, & \text{if } k < 2 \end{cases},$$

and for $v_s \in [0, 1]$,

$$s_s^*(v_s | p) = \begin{cases} \text{accept,} & \text{if } v_s < p \\ \text{accept or reject,} & \text{if } v_s = p, \\ \text{reject,} & \text{if } v_s > p \end{cases}$$

the buyer's belief about the seller's valuation is a uniform distribution on $[0, 1]$. □

13.5 Sequential equilibrium

13.44 An assessment is a pair (β, μ) where β is a profile of behavioral strategies and μ is a function that assigns to every information set a probability measure on the set of histories in the information set.

☞ 13.45 Let $\beta^k \rightsquigarrow \beta$ denote a “trembling sequence” $\{\beta^k\}_{k=1}^\infty$ of completely mixed behavioral strategy profiles that converges to a behavioral strategy profile β .

☞ 13.46 An assessment (β, μ) is a sequential equilibrium (abbreviated as “SE”) of a finite extensive game with perfect recall if there is $\{\beta^k\}_{k=1}^\infty$ such that $\beta^k \rightsquigarrow \beta$, and for all $i \in N$, $I_i \in \mathcal{I}_i$, and $k \geq 1$

- Sequential consistency: $\mu^k(I_i) \rightarrow \mu(I_i)$, where $\mu^k(I_i) \in \Delta(I_i)$ is the belief on I_i which derived from β^k by Bayes' rule.

- Sequential rationality: the restriction of β_i to information sets that succeed I_i is i 's best response to the restriction of β_{-i} to these information sets, using $\mu(I_i)$.

13.47 Example:




Figure 13.17

(1) The induced strategic game is

	L'	R'
L	3, 3	3, 3
M	0, 1	0, 0
R	1, 0	5, 1

There are two pure-strategy Nash equilibria: (L, L') and (R, R') .

The set of mixed-strategy Nash equilibria is

$$\{(R, R')\} \cup \{(L, \alpha \circ L' + (1 - \alpha) \circ R') \mid \frac{1}{2} \leq \alpha \leq 1\}.$$

(2) Consider the assessment (β, μ) , where $\beta_1 = L$, $\beta_2 = L'$ and $\mu(\{M, R\})(R) = 0$.

Let $\beta_1^k = (1 - \frac{1}{k} - \frac{1}{k^2}) \circ L + \frac{1}{k} \circ M + \frac{1}{k^2} \circ R$, $\beta_2^k = (1 - \frac{1}{k}) \circ L' + \frac{1}{k} \circ R'$. Then $\beta^k \rightsquigarrow \beta$.

Given β_1^k , (M) and (R) will be reached with probabilities $\frac{1}{k}$ and $\frac{1}{k^2}$ respectively. By Bayes' rule, the belief $\mu^k(\{M, R\})(M) = \frac{1/k}{1/k + 1/k^2} = \frac{k}{k+1} \rightarrow 1 = \mu(\{M, R\})(M)$. Therefore, the sequential consistency is satisfied.

Given μ , it is clear that L' is optimal for player 2. Given player 2's strategy L' , L is optimal for player 1. Thus, the sequential rationality is satisfied.

Therefore, (β, μ) is a sequential equilibrium.

(3) Consider the assessment (β, μ) , where $\beta_1 = R$, $\beta_2 = R'$ and $\mu(\{M, R\})(R) = 1$.

Let $\beta_1^k = \frac{1}{k} \circ L + \frac{1}{k} \circ M + \frac{k-2}{k} \circ R$, $\beta_2^k = \frac{1}{k} \circ L' + \frac{k-1}{k} \circ R'$. Then $\beta^k \rightsquigarrow \beta$.

Given β_1^k , (M) and (R) will be reached with probabilities $\frac{1}{k}$ and $\frac{k-2}{k}$ respectively. By Bayes' rule, the belief $\mu^k(\{M, R\})(R) = \frac{(k-2)/k}{1/k + (k-2)/k} = \frac{k-2}{k-1} \rightarrow 1 = \mu(\{M, R\})(R)$. Therefore, the sequential consistency is satisfied.

Given μ , it is clear that R' is optimal for player 2. Given player 2's strategy R' , R is optimal for player 1. Thus, the sequential rationality is satisfied.

Therefore, (β, μ) is a sequential equilibrium.

13.48 Example: Selten's horse.

- The induced strategic game is




Figure 13.18: Selten's horse.

	cL	cR	dL	dR
C	1, 1, 1	1, 1, 1	4, 4, 0	0, 0, 1
D	3, 3, 2	0, 0, 0	3, 3, 2	0, 0, 0

There are two pure-strategy Nash equilibria: (C, c, R) and (D, c, L) .

Since the payoffs of (C, c, R) and (C, c, L) are the same. To support $(C, c, \alpha \circ R + (1 - \alpha) \circ L)$ to be a Nash equilibrium, α should satisfy:

$$1 \geq 3 \cdot (1 - \alpha) \text{ and } 1 \geq 4 \cdot (1 - \alpha),$$

that is, $\alpha \geq \frac{3}{4}$.

Since the payoffs of (D, c, L) and (D, d, L) are the same. To support $(D, \gamma \circ c + (1 - \gamma) \circ d, L)$ to be a Nash equilibrium, γ should satisfy:

$$3 \geq \gamma + 4 \cdot (1 - \gamma),$$

that is, $\gamma \geq \frac{1}{3}$.

- For any Nash equilibrium $(C, c, \alpha \circ R + (1 - \alpha) \circ L)$ with $\alpha \geq \frac{3}{4}$, there is a sequential equilibrium (β, μ) in which $\beta_1(\emptyset)(C) = 1$, $\beta_2(C)(c) = 1$, $\beta_3(I)(R) = \alpha$, and $\mu_3(I)(D) = \frac{1}{3}$, where $I = \{(D), (C, d)\}$.

Let $\beta_1^k(\emptyset)(C) = 1 - \frac{1}{k}$, $\beta_2^k(C)(d) = \frac{2/k}{1-1/k}$, and $\beta_3^k(I)(R) = \alpha - \frac{1}{k}$. Then it is clear that $\beta^k \rightsquigarrow \beta$.

Given $(\beta_i^k)_i$, (D) will be reached with probability $\frac{1}{k}$, and (C, d) will be reached with probability $(1 - \frac{1}{k}) \cdot \frac{2/k}{1-1/k} = \frac{2}{k}$. Thus, by Bayes' rule, the belief $\mu_3^k(I)(D) = \frac{1}{3} = \mu_3(I)(D)$. Therefore the sequential consistency is satisfied.

Given the belief $(\frac{1}{3}, \frac{2}{3})$ on the information set I , L and R are indifferent for player 3. Thus, sequential rationality is satisfied.

- Any Nash equilibrium $(D, \gamma \circ c + (1 - \gamma) \circ d, L)$ with $\gamma \geq \frac{1}{3}$ is not part of any sequential equilibrium: since the associated assessment violates sequential rationality at player 2's (singleton) information set (since $4 > 1 \cdot \gamma + 4(1 - \gamma)$).

13.49 Any sequential equilibrium is a perfect Bayesian equilibrium. The converse does not hold.

If (β, μ) is a sequential equilibrium, then β is a subgame perfect equilibrium.

13.50 Existence: For every finite extensive game, there exists at least one sequential equilibrium.

This result is based on the results proved later.

13.51 Consider the following game.

It is easy to see there are two pure-strategy Nash equilibria: (L, L') and (R, R') , and both of them can be supported as sequential equilibria. For the sequential equilibrium $((L, L'), \mu)$, player 2 should believe, in the event that his information set is reached, that with high probability player 1 chooses M .




Figure 13.19

However, if player 2's information set is reached then a reasonable argument for him may be that since the action M for player 1 is strictly dominated by L it is not rational for player 1 to choose M and hence she must choose R . This argument excludes any belief that supports (L, L') as a sequential equilibrium outcome.

13.52 Consider the following game.




Figure 13.20

This game has two types of sequential equilibrium, as follows.

- Both types of player 1 choose B , and player 2 fights if he observes Q and not if he observes B . If player 2 observes Q then he assigns probability of at least 0.5 that player 1 is weak.
- Both types of player 1 choose Q , and player 2 fights if he observes B and not if he observes Q . If player 2 observes B then he assigns probability of at least 0.5 that player 1 is weak.

The following argument suggests that an equilibrium of the second type is not reasonable.

- (1) If player 2 observes that player 1 chose B then he should conclude that player 1 is strong, as follows. If player 1 is weak then she should realize that the choice of B is worse for her than following the equilibrium (in which she obtains the payoff 3), whatever the response of player 2.
- (2) Further, if player 1 is strong and if player 2 concludes from player 1 choosing B that she is strong and consequently chooses N , then player 1 is indeed better off than she is in the equilibrium (in which she obtains 2).
- (3) Thus it is reasonable for a strong type of player 1 to deviate from the equilibrium, anticipating that player 2 will reason that indeed she is strong, so that player 2's belief that player 1 is weak with positive probability when she observes B is not reasonable.

13.53 Example [JR Example 7.7].

13.6 Trembling hand perfect equilibrium

13.54 A trembling hand perfect equilibrium or simply perfect equilibrium (abbreviated as “PE”) of a finite strategic game is a mixed-strategy profile σ with the property that there exists a sequence $(\sigma^k)_{k=0}^\infty$ of completely mixed-strategy profiles that converges to σ such that for each player i the strategy σ_i is a best response to σ_{-i}^k for all values of k .

☞ 13.55 A trembling hand perfect equilibrium of a finite extensive game is a behavioral strategy profile β with the property that there is a sequence $\{\beta^k\}_{k=1}^\infty$ with $\beta^k \rightsquigarrow \beta$ such that for all $i \in N$, $I_i \in \mathcal{I}_i$ and $k \geq 1$,

$$U_i(\beta_i(I_i), \beta^k(-I_i)) \geq U_i(a, \beta^k(-I_i)) \text{ for all } a \in A(I_i).$$

Observation: A behavioral strategy profile is a perfect equilibrium if and only if it is a perfect equilibrium of the agent-strategic-form of the game.

13.56 The basic idea is that each player's actions be optimal not only given his equilibrium beliefs but also given a perturbed belief that allows for the possibility of slight mistakes.

13.57 Example:

	A	B	C
A	0, 0	0, 0	0, 0
B	0, 0	1, 1	2, 0
C	0, 0	0, 2	2, 2

There are three pure-strategy Nash equilibria (A, A) , (B, B) and (C, C) . However, (B, B) is the only perfect equilibrium.

- Consider (A, A) . If player 2 chooses the mixed strategy $(1 - \epsilon_1 - \epsilon_2) \circ A + \epsilon_1 \circ B + \epsilon_2 \circ C$, then it is clear that B is strictly better than A for player 1.
- Consider (C, C) . If player 2 chooses the mixed strategy $\epsilon_1 \circ A + \epsilon_2 \circ B + (1 - \epsilon_1 - \epsilon_2) \circ C$, then it is clear that B is strictly better than C for player 1.
- Consider (B, B) . Let $\sigma_1^k = \sigma_2^k = (\frac{1}{k} \circ A + \frac{k-2}{k} \circ B + \frac{1}{k} \circ C)$. Then $\sigma^k \rightsquigarrow \sigma$. It is clear that B is always optimal for player i given σ_j^k for all $k \geq 1$. Therefore, (B, B) is a perfect equilibrium.

13.58 For every perfect equilibrium β of a finite extensive game with perfect recall there is a belief system μ such that (β, μ) is a sequential equilibrium of the game.

Proof. (1) Let (β^k) be the sequence of completely mixed behavioral strategy profiles that corresponds to the sequence of mixed-strategy profiles in the agent strategic form of the game that is associated with the equilibrium β .

(2) At each information set $I_i \in \mathcal{I}_i$, define the belief $\mu(I_i)$ to be the limit of the beliefs defined from β^k using Bayes' rule. Then (β, μ) is a consistent assessment.

(3) Since every player's information set is reached with positive probability, by the one deviation property and the continuity of payoff functions, β_i is a best response to β_{-i} when the belief at I_i is defined by $\mu(I_i)$.

(4) Thus (β, μ) is a sequential equilibrium.

The converse does not hold. Consider the game in Example 13.57. (A, A) and (C, C) can be supported as sequential equilibria, but there are not perfect equilibria. \square

13.59 Example: Selten's horse




Figure 13.21: Selten's horse.

- Any Nash equilibrium $\sigma = (C, c, \alpha \circ R + (1 - \alpha) \circ L)$ with $\alpha \geq \frac{3}{4}$ is a perfect equilibrium strategy profile. For each k , let $\sigma_1^k = (1 - \frac{1}{k}) \circ C + \frac{1}{k} \circ D$, $\sigma_2^k = \frac{1-3/k}{1-1/k} \circ c + \frac{2/k}{1-1/k} \circ d$, and

$$\sigma_3^k = \begin{cases} \alpha \circ R + (1 - \alpha) \circ L, & \text{if } \alpha < 1; \\ (1 - \frac{1}{k}) \circ R + \frac{1}{k} \circ L, & \text{if } \alpha = 1. \end{cases}$$

Then $\sigma^k \rightsquigarrow \sigma$ and σ_i is optimal given σ_{-i}^k for all $i \in N$ and $k \geq 1$.

- Any Nash equilibrium $(D, \gamma \circ c + (1 - \gamma) \circ d, L)$ with $\gamma \geq \frac{1}{3}$ is not a perfect equilibrium since it is not a sequential equilibrium.

13.60 Existence: Every finite strategic game has a perfect equilibrium. Every finite extensive game with perfect recall has a perfect equilibrium and thus also a sequential equilibrium.

Proof. It suffices to show that every finite strategic game has a perfect equilibrium.

- Define a perturbation of the game $\Gamma(\epsilon)$ by letting the set of actions of each player i be the set of mixed strategies of player i that assign probability of at least ϵ to each action of player i .
- Every such perturbed game has a pure-strategy Nash equilibrium by Nash's theorem.
- Consider a sequence of such perturbed games $\Gamma(\epsilon^k)$ and their corresponding Nash equilibria σ^k in which $\epsilon^k \rightarrow 0$. By the compactness of the set of strategy profiles, we can pick a subsequence $\{\sigma^{k_l}\}_{l=1}^\infty$ of $\{\sigma^k\}_{k=1}^\infty$ such that $\sigma^{k_l} \rightsquigarrow \sigma$.
- By the continuity of payoff functions, σ is a perfect equilibrium of the game.

\square

13.61 Summary:

$$\text{PE} \subsetneq \text{SE} \subsetneq \text{PBE} \subsetneq \text{SPE} \subsetneq \text{NE}.$$

13.62 Example [JR Exercise 7.48].

13.63 Example [JR Exercise 7.49].

Chapter

14

Information economics 1: signaling

Contents

14.1	Adverse selection	223
14.2	Signaling	226
14.3	Job-market signaling	238
14.4	Cheap talk	240
14.4.1	The model with two types	242
14.4.2	The model with three types	244
14.4.3	The model with a continuum of types	244
14.4.4	Optimal communication mechanism (the commitment case)	248
14.5	Bayesian persuasion	252
14.5.1	Example	252
14.5.2	A model of Bayesian persuasion	253
14.5.3	Simplified problem	255
14.5.4	Optimal information structure	257
14.6	Optimal information disclosure	260

14.1 In this and the following two chapters, we consider situations in which an asymmetry of information exists among market participants. Asymmetric information is usually distinguished by two types: adverse selection (逆向选择) and moral hazard (道德风险).

14.1 Adverse selection

14.2 In this section we look at problems of adverse selection where one party to a transaction knows things pertaining to the transaction that are relevant to but unknown by the second party. Adverse selection models hidden characteristics, where asymmetric information exists before the parties enter into a relationship. It refers to a market process in which undesired results occur when buyers and sellers have asymmetric information (access to different information); the “bad” products or services are more likely to be selected.

One example is the market of used cars. In the market, buyers often do not observe the quality of the cars, which is privately information of the sellers. Due to the common existence of low quality used cars (the “lemons”—没

有价值的东西, 次品), buyers will be reluctant to pay high price for a high quality car (the “peach”—(同类事物中) 极好的 [极吸引人的] 事物), since they cannot tell its quality. As a consequence of low market prices, high quality sellers are driven out of the market (they lose if they sell), and whoever sells on the market is more likely to be selling a low quality car—adverse selection arises. As a result, buyer’s willingness to pay decreases further, and eventually, the market of high quality cars disappears.

Another example is in life/health insurance. If premiums are set at actuarially fair rates for the population as a whole, insurance may be a bad deal for healthy people, who then will refuse to buy. Only the sick and dying will sign up. And premium rates then must be set to reflect this.

14.3 Simple example: There are two types of used cars: peaches and lemons. A peach, if it is known to be a peach, is worth \$3000 to a buyer and \$2500 to a seller. A lemon, on the other hand, is worth \$2000 to a buyer and \$1000 to a seller. There are twice as many lemons as peaches.

- Case 1: If buyers and sellers both had the ability to look at a car and see whether it was a peach or a lemon, there would be no problem: Peaches would sell for \$3000 and lemons for \$2000.
- Case 2: If neither buyer nor seller knew whether a particular car was a peach or a lemon, we would have no problem (at least, assuming risk neutrality, which we will to avoid complications): A seller, thinking she has a peach with probability $\frac{1}{3}$ and a lemon with probability $\frac{2}{3}$, has a car that (in expectation) is worth \$1500. A buyer, thinking that the car might be a peach with probability $\frac{1}{3}$ and a lemon with $\frac{2}{3}$, thinks that the car is worth on average \$2333.33. Assuming an inelastic supply of cars and perfectly elastic demand, the market clears at \$2333.33.
- Case 3: The seller, having lived with the car for quite a while, knows whether it is a peach or a lemon. Buyers typically can not tell. If we make the extreme assumption that the buyers can not tell at all, then the peach market breaks down.

Therefore, the expected value of the car to sellers is \$2333.33, and that would be the maximal amount she is willing to pay for the car. Given this, only sellers of lemons sell, because a peach values \$2500 to sellers. So the market attracts only sellers of lemons and the way it selects sellers is a version of adverse selection.

Moreover, if only lemons are put on the market, buyer’s beliefs update: they understand the logic behind adverse selection (sellers of peaches are not willing to sell), the actual probability that they are facing a peach is zero. As a result, we get as equilibrium: Only lemons are put on the market, at a price of \$2000.

This example says that owners of good cars will not place their cars on the used car market. This is sometimes summarized as “the bad driving out the good” (劣币驱逐良币) in the market.

14.4 Akerlof’s model (Akerlof (1970)): Assume that there are just two groups of traders: groups one and two. Each member in group 1 has a car, and each member in group 2 is a potential buyer.

(a) A buyer’s utility function is

$$u_2 = M + \frac{3}{2} \cdot q \cdot n,$$

where M is the consumption of goods other than cars, q is the quality of the car, and n is the number of cars. For sake of simplicity, we assume that n is 0 (not buy) or 1 (buy).

(b) A buyer has a budget constraint

$$y_2 = M + p \cdot n,$$

where y_2 is the income, and p is the price of the used car.

(c) A seller's utility function is

$$u_1 = M' + q \cdot n.$$

Note that the coefficient of quality in u_1 is 1 which is less than that in u_2 , $\frac{3}{2}$. It means that the car is more needed for buyers.

(d) A seller has a budget constraint

$$y_1 = M' + p \cdot n.$$

14.5 Akerlof's model with symmetric information.

If q is public information, then the trade occurs if and only if

$$q \leq p \leq \frac{3}{2}q,$$

and buyer and seller will both benefit from the trade.

14.6 Akerlof's model with asymmetric information.

(1) A buyer's expected valuation is

$$\mathbb{E}[u_2] = M + \frac{3}{2} \cdot \mathbb{E}[q] \cdot n = M + \frac{3}{2} \cdot \mu \cdot n,$$

where $\mu \triangleq \mathbb{E}[q]$ is the average quality of used cars. Therefore, the buyer's aim is to maximize

$$\mathbb{E}[u_2] = y_2 + \left(\frac{3}{2} \cdot \mu - p\right) \cdot n.$$

So, a buyer will buy ($n = 1$) if and only if

$$\frac{3}{2} \cdot \mu \geq p.$$

(2) The seller's aim is to maximize her utility (not expected utility)

$$u_1 = y_1 + (q - p) \cdot n.$$

Therefore, a seller will sell ($n = 1$) her car if and only if

$$p \geq q.$$

Assume that q is uniformly distributed on $[0, 2]$, and the seller knows q but the buyer only knows its distribution.

level 1 buyer knows the expected valuation is $\mu = 1$, and her highest buying price is $p = \frac{3}{2} \cdot \mu = \frac{3}{2}$.

level 2 seller knows that buyer's highest price is $\frac{3}{2}$. Then only the cars with quality less than $\frac{3}{2}$ will be sold.

level 3 buyer knows that only the cars with quality less than $\frac{3}{2}$ will be sold, so she believes that q is uniformly distributed on $[0, \frac{3}{2}]$. It is the first adverse selection.

Analogously, we have $\mu^2 = \frac{3}{4}$, and $p^2 = \frac{3}{2} \cdot \frac{3}{4} = \frac{9}{8}$. So the cars with quality higher than $\frac{9}{8}$ will be kicked off, and buyers believe that q is uniformly distributed on $[0, \frac{9}{8}]$.

Repeat this process, p and q will converge to zeros, that is, the good cars may be driven out of the market by the bad cars. Actually we have the bad driving out the not-so-bad driving out the medium driving out the not-so-good driving out the good in such a sequence of events that no market exists at all.

14.7 Signaling (信号传递) and screening (信号甄别) are two primary solutions to adverse selection.

- In signaling the informed party moves first. For example, the seller of a high quality car can (costly) signal to potential buyers by offering a long warranty, which is not affordable to low quality car sellers, as a way to distinguish them from low quality car sellers. The insurance company may offer life insurance on better terms if the insure is willing to accept very limited benefits for the first two or three years the policy is in effect, on the presumption that someone suffers from ill health and is about to die is unwilling to accept those limited benefits.
- In screening, the uninformed party moves first. For example, the uninformed party offers a term of exchange such that the informed party can choose to accept or reject. For example, the uninformed buyer can ask for a long enough warranty to screen out low quality sellers. The insurance market gives a better illustration, where the insurance company, who is uninformed of conditions of the clients, offers different insurance packages, and ideally, each type of clients only find one package acceptable.

14.2 Signaling

14.8 A signaling game is an extensive game of imperfect information involving two players: a Sender (S) and a Receiver (R). The timing of the game is as follows.

- (1) Nature draws a type t_i for the Sender from a set of feasible types $T = \{t_1, t_2, \dots, t_I\}$ according to a probability distribution $P(t_i)$.
- (2) The Sender observes t_i and then chooses a message m_j from a set of feasible messages $M = \{m_1, m_2, \dots, m_J\}$.
- (3) The Receiver observes m_j (but not t_i) and then chooses an action a_k from a set of feasible actions $A = \{a_1, a_2, \dots, a_K\}$.
- (4) Payoffs are given by $U_S(t_i, m_j, a_k)$ and $U_R(t_i, m_j, a_k)$.

A strategy for Receiver is a function from T to M , and a strategy for Sender is a function from M to A .

14.9 We translate the requirements for a perfect Bayesian equilibrium to the case of signaling games.

1. After observing any message m_j from M , the Receiver must have a belief about which types could have sent m_j . Denote this belief by the probability distribution $\mu(t_i | m_j)$, where $\mu(t_i | m_j) \geq 0$ for each $t_i \in T$, and $\sum_{t_i \in T} \mu(t_i | m_j) = 1$.
- 2R. For each $m_j \in M$, the Receiver's action $a^*(m_j)$ must maximize the Receiver's expected utility, given the belief $\mu(t_i | m_j)$ about which types could have sent m_j . That is, $a^*(m_j)$ solves

$$\max_{a_k \in A} \sum_{t_i \in T} \mu(t_i | m_j) U_R(t_i, m_j, a_k).$$

- 2S. For each $t_i \in T$, the Sender's message $m^*(t_i)$ must maximize the Sender's utility, given the Receiver's strategy $a^*(m_j)$. That is, $m^*(t_i)$ solves

$$\max_{m_j \in M} U_S(t_i, m_j, a^*(m_j)).$$

3. Let T_j denote the set of types that send the message m_j .

For each $m_j \in M$, if there exists $t_i \in T$ such that $m^*(t_i) = m_j$, i.e., $T_j \neq \emptyset$, then the Receiver's belief at the information set corresponding to m_j must follow from Bayes' rule and the Sender's strategy:

$$\mu(t_i | m_j) = \frac{P(t_i)}{\sum_{t_i \in T_i} P(t_i)}.$$

14.10 A simple example:




Figure 14.1

- (1) Nature has two types t_1 and t_2 with the same probability.
- (2) The Sender observes t_i and then chooses a message L or R .
- (3) The Receiver observes the message and then chooses an action u or d .
- (4) Payoffs depend on the type of Nature, the message of Sender, and the action of Receiver.

Sender's strategies are: LL, LR, RL, RR , where $m'm''$ means that Sender plays m' when facing type t_1 , and m'' when facing type t_2 .

Receiver's strategies are: uu, ud, du, dd , where $a'a''$ means that Receiver plays a' following L and a'' following R .

Receiver has two non-trivial information sets. The beliefs on them are given in the game tree: $(p, 1-p)$ for the left information set and $(q, 1-q)$ for the right information set.

We analyze the possibility of the four Sender's strategies to be perfect Bayesian equilibria.

14.11 Case 1: Pooling on L . Suppose Sender adopts the strategy LL .

- (1) By Bayes' rule, $p = 1 - p = \frac{1}{2}$.
- (2) On the left information set, given the belief $(\frac{1}{2}, \frac{1}{2})$, Receiver's expected payoff is $\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 3 = \frac{7}{2}$ for u and $\frac{1}{2} \cdot 5 + \frac{1}{2} \cdot 1 = 3$ for d . Thus, Receiver's best response for message L is u .
- (3) For the right information set, the belief $(q, 1-q)$ can not be determined by Sender's strategy, thus any belief $(q, 1-q)$ is available. Furthermore, both u and d are possible for some q respectively.
Hence we need only to see if sending L is better than sending R for both types t_1 and t_2 and for one of u and d .
- (4) If u is the Receiver's best response on the right information set, i.e., Receiver's strategy is uu , then

- For type t_1 , Sender's payoff is 2 if L is sent, and 0 if R is sent. Hence sending L is optimal.
 - For type t_2 , Sender's payoff is 3 if L is sent, and 4 if R is sent. Hence sending L is not optimal.
- (5) If d is the Receiver's best response on the right information set, *i.e.*, Receiver's strategy is ud , then
- For type t_1 , Sender's payoff is 2 if L is sent, and 3 if R is sent. Hence sending L is not optimal.
 - For type t_2 , Sender's payoff is 3 if L is sent, and 2 if R is sent. Hence sending L is optimal.

Therefore, there is no perfect Bayesian equilibrium in which Sender plays LL .

14.12 Case 2: Pooling on R . Suppose Sender adopts the strategy RR .

- (1) By Bayes' rule, $q = 1 - q = \frac{1}{2}$.
- (2) On the right information set, given the belief $(\frac{1}{2}, \frac{1}{2})$, Receiver's expected payoff is $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 1 = \frac{3}{2}$ for u and $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 2 = 2$ for d . Thus, Receiver's best response for message R is d .
- (3) For the left information set, the belief $(p, 1 - p)$ can not be determined by Sender's strategy, thus any belief $(p, 1 - p)$ is available. Furthermore, both u and d are possible for some p respectively.
Hence we need only to see if sending R is better than sending L for both types t_1 and t_2 and for one of u and d .
- (4) If u is the Receiver's best response on the left information set, *i.e.*, Receiver's strategy is ud , then
 - For type t_1 , Sender's payoff is 2 if L is sent, and 3 if R is sent. Hence sending R is optimal.
 - For type t_2 , Sender's payoff is 3 if L is sent, and 2 if R is sent. Hence sending R is not optimal.
- (5) If d is the Receiver's best response on the left information set, *i.e.*, Receiver's strategy is dd , then
 - For type t_1 , Sender's payoff is 1 if L is sent, and 3 if R is sent. Hence sending R is optimal.
 - For type t_2 , Sender's payoff is 4 if L is sent, and 2 if R is sent. Hence sending R is not optimal.

Therefore, there is no perfect Bayesian equilibrium in which Sender plays RR .

14.13 Case 3: Separation with t_1 playing L . Suppose Sender adopts the separation strategy LR .

- (1) By Bayes' rule, $p = 1$ and $q = 0$.
- (2) Based on these beliefs, d and d are Receiver's best responses on the left and right information sets respectively.
- (3) For type t_1 , Sender's payoff is 1 when sending L and 3 when sending R . Thus L is not optimal.
- (4) For type t_2 , Sender's payoff is 4 when sending L and 2 when sending R . Thus R is not optimal.

Therefore, there is no perfect Bayesian equilibrium in which Sender plays LR .

14.14 Case 4: Separation with t_1 playing R . Suppose Sender adopts the separation strategy RL .

- (1) By Bayes' rule, $p = 0$ and $q = 1$.
- (2) Based on these beliefs, u is Receiver's best response on the left information set. On the right information set, u and d are indifferent for Receiver.
- (3) If Receiver chooses u on the right information set, then for type t_1 , sending R is not optimal.
Now let receiver's best response be d .
- (4) For type t_1 , Sender's payoff is 2 when sending L and 2 when sending R . Thus R is optimal.
- (5) For type t_2 , Sender's payoff is 3 when sending L and 2 when sending R . Thus L is optimal.

Therefore, $((RL, ud), (p = 0, q = 1))$ is a perfect Bayesian equilibrium.

14.15 Alternative method.

		Receiver			
		<i>uu</i>	<i>ud</i>	<i>du</i>	<i>dd</i>
Sender	<i>LL</i>	2.5, 3.5	2.5, 3.5	2.5, 3	2.5, 3
	<i>LR</i>	3, 2.5	2, 3	2.5, 3	1.5, 3.5
	<i>RL</i>	1.5, 2.5	3, 2.5	2, 1.5	3.5, 1.5
	<i>RR</i>	2, 1.5	2.5, 2	2, 1.5	2.5, 2

The payoff profile in each cell is the expected payoff profile for two players given the prior probability of Nature. For example, when Sender and Receiver choose *RL* and *du* respectively, their payoffs are:

$$(U_S, U_R) = \underbrace{\frac{1}{2}(0, 2)}_{\text{Nature chooses } t_1} + \underbrace{\frac{1}{2}(4, 1)}_{\text{Nature chooses } t_2} = (2, 1.5).$$

There is unique Nash equilibrium (RL, ud) . Since Sender's strategy is *RL*, then $p = 0$ and $q = 1$. Based on such beliefs, Receiver's strategy is a best response. Therefore, $((RL, ud), (p = 0, q = 1))$ is a perfect Bayesian equilibrium.

14.16 Example: Beer and Quiche (Cho and Kreps (1987)).

Player 1 enters a restaurant to have breakfast, where there is a bully-player 2. Player 1 is either strong or weak; if he is strong, he prefers Beer, and if he is weak, he prefers Quiche. Player 2 would enjoy from bullying (fighting) player 1 only if player 1 is weak, but he observes only player 1's choice for breakfast (player 1's signal) but not player 1's type. This game models armed negotiation.




Figure 14.2: Beer and Quiche.

Answer. Case 1: Pooling on *B*. Suppose player 1 adopts the strategy *BB*.

- $p = 0.9$.
- On the left information set, player 2 chooses *N*. Hence, player 1 gets 3 if he is strong, and 2 if he is weak.
- On the right information set, player 2 should choose *F*; otherwise player 1 choosing *B* when he is weak is not optimal.
- $q \leq \frac{1}{2}$.

$((BB, NF), p = 0.9, q \leq \frac{1}{2})$ is a perfect Bayesian equilibrium.

Case 2: Pooling on Q . Suppose player 1 adopts the strategy QQ .

- $q = 0.9$.
- On the right information set, player 2 chooses N .
- On the left information set, player 2 should choose F ; otherwise player 1 choosing Q when he is strong is not optimal.
- $p \leq \frac{1}{2}$.

$((QQ, FN), p \leq \frac{1}{2}, q = 0.9)$ is a perfect Bayesian equilibrium.

Case 3: Separation with “strong” playing B . Suppose player 1 adopts the separation strategy BQ .

- $p = 1$ and $q = 0$.
- On the left information set, player 2 chooses N .
- On the right information set, player 2 chooses F . However, player 1 will deviate from Q to B when he is weak.

No perfect Bayesian equilibrium exists in this case.

Case 4: Separation with “strong” playing Q . Suppose player 1 adopts the separation strategy QB .

- $p = 0$ and $q = 1$.
- On the left information set, player 2 chooses F . However, player 1 will deviate from B to Q when he is weak.
- On the right information set, player 2 chooses N .

No perfect Bayesian equilibrium exists in this case.

□

14.17 Example [G Exercise 4.3]: Three-type signaling games.

Find all perfect Bayesian equilibria in the following signaling games.

Answer. Let (p_1, p_2, p_3) and (q_1, q_2, q_3) be player 2’s beliefs at the left and right information sets, respectively. Note that u dominates d at the left information set. Therefore, we only need to consider the following induced strategic game.

- $T = \{t_1, t_2, t_3\}$, $M = \{L, R\}$, and $A = \{u, d\}$.
- Payoff table:




Figure 14.3

		Receiver	
		<i>uu</i>	<i>ud</i>
Sender	<i>LLL</i>	$4/3, 1$	$4/3, 1$
	<i>LLR</i>	$1, 2/3$	$5/3, 1$
	<i>LRL</i>	$1, 1$	$1, 2/3$
	<i>LRR</i>	$2/3, 2/3$	$4/3, 2/3$
	<i>RLL</i>	$1, 1$	$1, 2/3$
	<i>RLR</i>	$2/3, 2/3$	$4/3, 2/3$
	<i>RRL</i>	$2/3, 1$	$2/3, 1/3$
	<i>RRR</i>	$1/3, 2/3$	$1, 1/3$

For example,

$$\begin{aligned}
 U(RLR, du) &= \text{Prob}(t_1) \cdot U(R, u \mid t_1) + \text{Prob}(t_2) \cdot U(L, d \mid t_2) + \text{Prob}(t_3) \cdot U(R, u \mid t_3) \\
 &= \frac{1}{3}(0, 1) + \frac{1}{3}(0, 0) + \frac{1}{3}(0, 0) = (0, 1/3)
 \end{aligned}$$

There are two pure-strategy Nash equilibria (LLL, uu) and (LLR, ud) , which are also the subgame perfect Nash equilibria since there is no subgame.

- To check whether (LLL, uu) is a perfect Bayesian equilibrium, we need only to find beliefs, satisfying Requirements 1, 2S, 2R and 3.
 - Requirement 1: Assume the probability distributions on left and right information set are (p_1, p_2, p_3) and (q_1, q_2, q_3) , respectively, displayed in the figure, where $p_1 + p_2 + p_3 = q_1 + q_2 + q_3 = 1$.
 - Requirement 2S: Holds automatically. (since (LLL, uu) is a Nash equilibrium)

- Requirement 2R: It is obvious that u is the best response for Receiver when Sender chooses L . To support u to be a best response for Receiver when Sender chooses R , we should take $q_3 \leq \frac{1}{2}$.
- Requirement 3: Since Sender chooses LLL , Bayes' rule implies $p_1 = p_2 = p_3 = \frac{1}{3}$ and q_1, q_2, q_3 could be arbitrary.

Hence, (LLL, uu) with player 2's beliefs $(p_1, p_2, p_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and (q_1, q_2, q_3) where $q_3 \leq \frac{1}{2}$.

- To check whether (LLR, ud) is a perfect Bayesian equilibrium, we need only to find beliefs, satisfying Requirements 1, 2S, 2R and 3.
 - Requirement 1: Assume the probability distributions on left and right information set are (p_1, p_2, p_3) and (q_1, q_2, q_3) , respectively, displayed in the figure, where $p_1 + p_2 + p_3 = q_1 + q_2 + q_3 = 1$.
 - Requirement 2S: Holds automatically. (since (LLR, ud) is a Nash equilibrium)
 - Requirement 2R: It is obvious that u is the best response for Receiver when Sender chooses L . To support d to be a best response for Receiver when Sender chooses R , we should take $q_3 \geq \frac{1}{2}$.
 - Requirement 3: Since Sender chooses LLR , Bayes' rule implies $p_1 = p_2 = \frac{1}{2}, p_3 = 0$ and $q_1 = q_2 = 0, q_3 = 1$.

Hence, (LLR, ud) with player 2's beliefs $(p_1, p_2, p_3) = (\frac{1}{2}, \frac{1}{2}, 0)$ and $(q_1, q_2, q_3) = (0, 0, 1)$ is a perfect Bayesian equilibrium.

□

14.18 Example: There are two players in the game: Judge and Plaintiff. The Plaintiff has been injured. Severity of the injury, denoted by v , is the Plaintiff's private information. The Judge does not know v and believes that v is uniformly distributed on $\{0, 1, \dots, 9\}$ (so that the probability that $v = i$ is $\frac{1}{10}$ for any $i \in \{0, 1, \dots, 9\}$). The Plaintiff can verifiably reveal v to the Judge without any cost, in which case the Judge will know v . The order of the events is as follows. First, the Plaintiff decides whether to reveal v or not. Then, the Judge rewards a compensation R which can be any non-negative real number. The payoff of the Plaintiff is $R - v$, and the payoff of the Judge is $-(v - R)^2$. Everything described so far is common knowledge. Find a perfect Bayesian equilibrium.

Answer. The signaling game is as follows: types $T = \{0, 1, \dots, 9\}$; signals $M = \{R, N\}$, where R is "Reveal" and N is "Not Reveal"; actions $A = \mathbb{R}_+$.

Figure 14.4 is the game tree. From the game tree, there are 10 subgames, and Judge has 11 information sets I_0, I_1, \dots, I_9 , where for $v = 0, 1, \dots, 9$, I_v denotes that Plaintiff reveals v to Judge, and I_{10} denotes the case that Plaintiff does not reveal the value.

Plaintiff's strategy space is

$$S = \{s = (s_0, s_1, \dots, s_9) \mid s_v = R \text{ or } N, v = 0, 1, \dots, 9\}.$$

For a particular strategy of Plaintiff $s = (s_0, s_1, \dots, s_9)$, s_v is the action of Plaintiff when she/he faces injury v .

Judge's strategy space is

$$Q = \{q = (x_0, x_1, \dots, x_9, x_{10}) \mid x_v \geq 0, v = 0, 1, \dots, 9, 10\}.$$

For a particular strategy of Judge $q = (x_0, x_1, \dots, x_9, x_{10})$, x_v is the action of Judge at the information set I_v .

Given any strategy s of Plaintiff, let $s^{-1}(N) = \{v : s(v) = N\}$, which denotes the set of Plaintiff's types at which the value is not revealed to Judge.




Figure 14.4

- (1) Claim 1: In any perfect Bayesian equilibrium (s^*, q^*, p^*) , if Plaintiff chooses R when $v = 0$, that is $s_0^* = R$, then Judge's action on information set I_{10} should be 0, that is, $x_{10}^* = 0$.
- (2) Proof of Claim 1: Otherwise, Plaintiff can be better off by deviating from R to N : If Plaintiff chooses R when $v = 0$, then she/he will get 0 when $v = 0$; otherwise she/he will get $x_{10}^* > 0$. Therefore, such a strategy s^* can not be a strategy in a perfect Bayesian equilibrium, which is a contradiction.
- (3) Claim 2: In a perfect Bayesian equilibrium (s^*, q^*, p^*) , if $(s^*)^{-1}(N) \neq \emptyset$, then Judge's strategy should be

$$q^* = \left(0, 1, \dots, 9, \sum_{v \in s^{-1}(N)} \frac{v}{n} \right),$$

where $n = |s^{-1}(N)|$.

- (4) Motivation of Claim 2: Based on Judge's belief p^* , her/his optimal action x_{10}^* should be weighted payoff

$$0 \cdot p_0^* + 1 \cdot p_1^* + 2 \cdot p_2^* + \dots + 9 \cdot p_9^*.$$

Given Plaintiff's strategy s^* , Judge's belief p^* on the information set I_{10} can be determined by Bayes' law.

- (5) Proof of Claim 2: (s^*, q^*) should be a subgame perfect equilibrium, and hence on the information set I_v ($v = 0, 1, \dots, 9$), Judge will choose optimal action based on her/his payoff $-(v - x_v)^2$. Therefore, Judge's action on the information set I_v should be v ($v = 0, 1, 2, \dots, 9$).

On the information set I_{10} , which is on the equilibrium path, only the branches v , where $v \in s^{-1}(N)$ can be reached. Thus, by Bayes' rule, Judge believes that these branches are reached with equal probability, $\frac{1}{n}$, where $n = |s^{-1}(N)|$. Thus, Judge will choose the optimal action based on her/his expected payoff, and the optimal action is the maximizer of the following maximization problem

$$\max_{x_{10} \geq 0} -\frac{1}{n} \sum_{v \in s^{-1}(N)} (v - x_{10})^2.$$

By first order condition, it is easy to find the unique maximizer $x_{10}^* = \frac{1}{n} \sum_{v \in s^{-1}(N)} v$.

- (6) Claim 3: In any perfect Bayesian equilibrium (s^*, q^*, p^*) , Plaintiff's strategy s^* should be

$$(R, R, \dots, R) \text{ or } (N, R, \dots, R).$$

- (7) Proof of Claim 3:

- Case 1: assume $(s^*)^{-1}(N) = \{v_0\}$, where $v_0 \neq 0$. Given such a Plaintiff's strategy s^* , that is, $s^*(v_0) = N$, and $s^*(v) = R$ for others v , by Claim 2, Judge's best response is

$$q^* = (0, 1, 2, \dots, 9, v_0).$$

However, s^* is not a best response for Plaintiff given Judge's strategy $q^*(s)$: when $v = 0$, Plaintiff can be better off if she/he chooses N rather than R : if she/he chooses R , she/he will get 0; otherwise, she/he will get $v_0 > 0$.

- Case 2: assume $(s^*)^{-1}(N)$ contains at least 2 elements. Let $v_1 = \min(s^*)^{-1}(N)$, and $v_2 = \max(s^*)^{-1}(N)$. Note that,

$$v_1 < x_{10}^* = \frac{1}{n} \sum_{v \in s^{-1}(N)} v < v_2.$$

By Claim 2, Judge's best response is

$$q^* = \left(0, 1, 2, \dots, 9, \sum_{v \in s^{-1}(N)} \frac{v}{n} \right).$$

However, s^* is not a best response for Plaintiff given Judge's strategy q^* : when the injury is v_2 , Plaintiff can get a higher amount v_2 by revealing: if she/he chooses N , she/he will get $x_{10}^* - v_2 < 0$; otherwise she/he will get 0.

Case 2 implies that there is at most 1 type at which Plaintiff chooses N in a perfect Bayesian equilibrium; and Case 1 implies that this unique type can only be $v = 0$.

- (8) Claim 4:

$$s^* = (N, R, \dots, R), \quad q^* = (0, 1, 2, \dots, 9, 0)$$

with belief $(1, 0, \dots, 0)$ on I_{10} is a perfect Bayesian equilibrium.

- (9) Proof of Claim 4: Routine.

(10) Claim 5:

$$s^* = (R, R, \dots, R), \quad q^* = (0, 1, 2, \dots, 9, 0)$$

with belief $(1, 0, \dots, 0)$ on I_{10} is a perfect Bayesian equilibrium:

(11) Proof of Claim 5: By Claims 1, 2 and 3, this strategy profile could be a strategy profile in a perfect Bayesian equilibrium. Assume Judge's belief on the information set I_{10} is $(p_0^*, p_1^*, \dots, p_9^*)$, then Judge's maximization problem is

$$\max_{x_{10} \geq 0} -p_0^*(x_{10} - 0)^2 - p_1^*(x_{10} - 1)^2 - \dots - p_9^*(x_{10} - 9)^2.$$

Then the unique maximizer is $x_{10}^* = p_0^* \cdot 0 + p_1^* \cdot 1 + \dots + p_9^* \cdot 9$. We have already known that $x_{10}^* = 0$, this implies $p_0^* = 1$ and $p_1^* = p_2^* = \dots = p_9^* = 0$, that is, Judge believes that $v = 0$ with probability 1.

□

14.19 Akerlof's model of market for "lemons".

Suppose a seller wants to sell his used car. The seller knows what is the quality of the car, but the buyer does not. The buyer knows only that the car could be a "good quality" car with probability $\frac{1}{2}$ and a "lemon" with probability $\frac{1}{2}$. If the car is good, the buyer's valuation for it is \$20,000 and the seller's is \$10,000. If it is a "lemon", both buyer's and seller's valuations are \$0.

The seller can make two offers (asking price): \$5,000 and \$15,000. Then, the buyer can accept the offer (buy the car) or reject the offer.

Find all the perfect Bayesian equilibria.

Answer. The following is the game tree.




Figure 14.5

Case 1: Separation with type H playing 5.

In this case, $p = 1$, and $q = 0$. It is easy to see that $a^*(5) = Y$ and $a^*(15) = N$. At type H , sending 5 is not optimal.

Case 2: Separation with type H playing 15.

In this case, $p = 0$, and $q = 1$. It is easy to see that $a^*(15) = Y$ and $a^*(5) = N$. At type H , sending 15 is not optimal.

Case 3: Pooling on 5.

In this case, $p = \frac{1}{2}$, and it is easy to see that $a^*(5) = Y$. If $a^*(15) = Y$ or N , then at type H , sending 5 is not optimal.

Case 4: Pooling on 15.

In this case, $q = \frac{1}{2}$, and it is easy to see that $a^*(15) = N$. When $p > \frac{1}{4}$, we have $a^*(5) = Y$. Then at type L , sending 15 is not optimal.

When $p \leq \frac{1}{4}$, we have $a^*(5) = N$. Then at both types H and L , sending 15 is optimal.

Hence, $\left((15, 15), (N, N)\right), 0 \leq p \leq \frac{1}{4}, q = \frac{1}{2}$ is a perfect Bayesian equilibrium. \square

14.20 Example: The market for “lemons” with different prior probability.

Suppose a seller wants to sell his used car. The seller knows what is the quality of the car, but the buyer does not. The buyer knows only that the car could be a “good quality” car with probability $\frac{4}{5}$ and a “lemon” with probability $\frac{1}{5}$. If the car is good, the buyer’s valuation for it is \$20,000 and the seller’s is \$10,000. If it is a “lemon”, both buyer’s and seller’s valuations are \$0.

The seller can make two offers (asking price): \$5,000 and \$15,000. Then, the buyer can accept the offer (buy the car) or reject the offer.

- (i) Would there be a separating perfect Bayesian equilibrium in this case?
- (ii) Find all the pooling perfect Bayesian equilibria.

Answer. The following is the game tree.




Figure 14.6

- (i) There is no separating perfect Bayesian equilibrium.
- (ii) The pooling perfect Bayesian equilibria are:

- $a_1^*(H) = a_1^*(L) = 15, a_2^*(5) = a_2^*(15) = Y, \beta^*(5) \in [\frac{1}{4}, 1], \beta^*(15) = \frac{4}{5}$.
- $a_1^*(H) = a_1^*(L) = 15, a_2^*(5) = N, a_2^*(15) = Y, \beta^*(5) \in [0, \frac{1}{4}], \beta^*(15) = \frac{4}{5}$.

\square

14.21 Example: The market for “lemons” with an option of passing an inspection.

In the context of previous example, consider the following variation in which the seller has an option of passing an inspection. If the inspection finds the car in good shape, the seller is charged \$200 to get the proof of inspection. Otherwise, he needs to pay \$15,200 to have his car serviced and get the proof of inspection. In this case, the buyer's valuation for the serviced car is \$20,000 and the seller's is \$10,000. Now, the seller has four choices: \$5,000 with inspection (“5i”), \$15,000 with inspection (“15i”), \$5,000 without inspection (“5”), and \$15,000 without inspection (“15”).

Show that

$$a_1^*(H) = 15i, a_1^*(L) = 5, a_2^*(5) = a_2^*(15) = N, a_2^*(5i) = a_2^*(15i) = Y, \\ \beta^*(5) = 0, \beta^*(15) \in [0, \frac{3}{4}], \beta^*(5i) \in [0, 1], \beta^*(15i) = 1,$$

is a separating perfect Bayesian equilibrium.

Answer. The following is the game tree.




Figure 14.7

□

14.3 Job-market signaling

Literature: [Spence \(1970\)](#).

14.22 Spence's model of education: A worker (the sender) knows his productive ability θ , while his employer (the receiver) does not. The timing is as follows:

- (1) Nature determines a worker's productive ability, θ , which can be either high (θ_H) with probability p_H or low ($\theta_L < \theta_H$) with probability p_L .
- (2) The worker, before entering the labor market, learns his ability and then chooses a level of education $e \in [0, +\infty)$.

- (3) A employer observes the worker's education (but not the worker's ability) and then pays the worker a wage $w \in [0, +\infty)$.
- (4) The payoffs are $w - \frac{e}{\theta}$ to the worker and $-(\theta - w)^2$ to the employer (under the assumption of perfect competition on the demand side).

14.23 The first-best outcome: suppose that θ is perfectly observable. Then:

$$w = \theta, \text{ and } e = 0.$$

14.24 The extensive form is as follows.

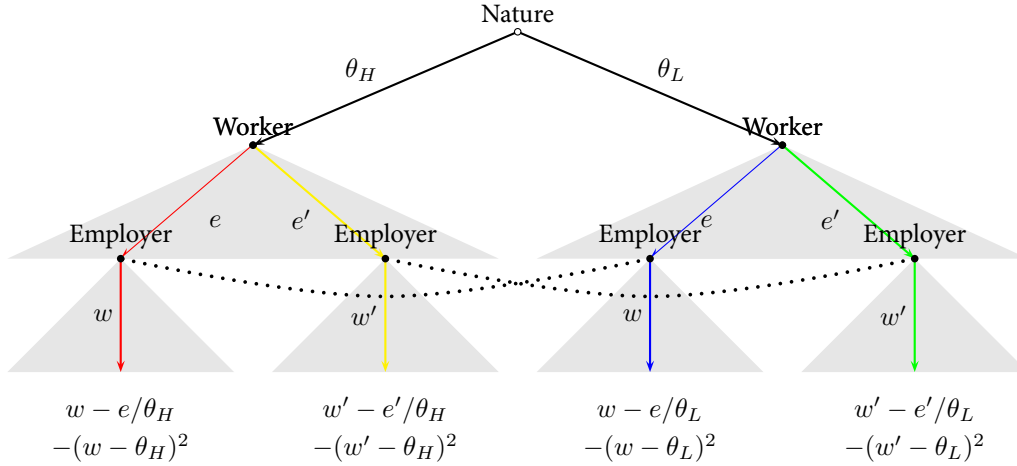


Figure 14.8

14.25 A strategy for the worker is (e_H, e_L) which specifies actions for types H and L , where $e_H, e_L \in [0, +\infty)$.

A strategy for the employer is a wage schedule $w(\cdot)$ which depends on the observed signal (*i.e.* the level of education).

14.26 $((e_H^*, e_L^*), w^*(\cdot), \mu^*)$ is a perfect Bayesian equilibrium, where $\mu^*(\cdot | e)$ specifies the belief about the worker's types when the observed signal is e .

14.27 Pooling equilibrium: Both types choose the same level of education: $e_H^* = e_L^* = e^*$.

By the definition of perfect Bayesian equilibrium, $w^*(e^*) = p_H \theta_H + p_L \theta_L$. To be a perfect Bayesian equilibrium pooling on e^* , the easiest way is to pessimistically believe that any deviation $e \neq e^*$ is from type L . Thus the wage schedule should be:

$$w^*(e) = \begin{cases} p_H \theta_H + p_L \theta_L, & \text{if } e = e^*, \\ \theta_L, & \text{if } e \neq e^*. \end{cases}$$

To be a perfect Bayesian equilibrium pooling on e^* , each type of workers does not want to deviate from e^* . Thus,

$$\left. \begin{aligned} w^*(e^*) - \frac{e^*}{\theta_H} &\geq \theta_L - \frac{e}{\theta_H}, \forall e \\ w^*(e^*) - \frac{e^*}{\theta_L} &\geq \theta_L - \frac{e}{\theta_L}, \forall e \end{aligned} \right\} \iff w^*(e^*) - \frac{e^*}{\theta_L} \geq \theta_L \iff e^* \leq p_H(\theta_H - \theta_L)\theta_L.$$

Therefore, $((e^*, e^*), w^*(\cdot), \mu^*)$ is a pooling perfect Bayesian equilibrium, where

$$\mu^*(H | e) = \begin{cases} p_H, & \text{if } e = e^* \\ 0, & \text{if } e \neq e^* \end{cases} \text{ and } 0 \leq e^* \leq p_H(\theta_H - \theta_L)\theta_L.$$

14.28 Separating equilibrium: The two types of workers choose different levels of education: $e_H^* > e_L^* = 0$ (the wage paid to L is θ_L , independent of e_L).

We consider the most pessimistic wage schedule:

$$w^*(e) = \begin{cases} \theta_H, & \text{if } e = e_H^*, \\ \theta_L, & \text{if } e \neq e_H^*. \end{cases}$$

To be a separating perfect Bayesian equilibrium, each type should have no incentive to mimic the other. Therefore,

$$\left. \begin{aligned} \theta_H - \frac{e_H^*}{\theta_H} &\geq \theta_L \\ \theta_L &\geq \theta_H - \frac{e_H^*}{\theta_L} \end{aligned} \right\} \iff (\theta_H - \theta_L)\theta_L \leq e_H^* \leq (\theta_H - \theta_L)\theta_H.$$

Therefore, $((e_H^*, 0), w^*(\cdot), \mu^*)$ is a pooling perfect Bayesian equilibrium, where

$$\mu^*(H | e) = \begin{cases} 1, & \text{if } e = e_H^* \\ 0, & \text{if } e \neq e_H^* \end{cases} \text{ and } (\theta_H - \theta_L)\theta_L \leq e_H^* \leq (\theta_H - \theta_L)\theta_H.$$

14.4 Cheap talk

Literature: [Crawford and Sobel \(1982\)](#), [Krishna and Morgan \(2008\)](#).

14.29 Cheap talk games are analogous to signaling games, but in cheap talk games the sender's messages are just talk—costless, nonbinding, nonverifiable claims.

14.30 Such talk cannot be informative in Spence's signaling game: a worker who simply announced “My ability is high” would not be believed.

It is because that all types have the same preferences over the receiver's possible actions: all workers prefer high wages, independent of ability. Therefore, a situation when two types of sender send different messages and the receiver responds differently to these messages is impossible at equilibrium: the sender-type who gets a less favorable response is better off with changing his message to the one employed by the other type.

14.31 Real examples of cheap talk include:

- Monetary mystique: a central bank is unwilling to make precise statements about its policy objectives.
- Security analyst recommendations.

14.32 It turns out that in a variety of contexts cheap talk is informative. An example is an expert advising a politician. The politician, after hearing the opinion of the expert, makes a decision which affects the payoffs of both players.

For cheap talk to be useful, the following conditions are necessary:

- Different sender-types should have different preferences over the receiver's actions.
- The receiver should prefer different actions depending on the sender's type.
- The receiver's preferences over actions should not be completely opposed to the sender's. Otherwise, the sender is worse off revealing true information about his type. Therefore, cheap talk cannot be informative in this situation: the receiver will be misled by the sender.

14.33 A decision maker (receiver) must choose some decision a . Her payoff depends on a and on an unknown state of the world θ , which is distributed uniformly on Θ . The decision maker can base her decision on the costless message m sent by an expert (sender) who knows the precise value of θ .

The decision maker's payoff is

$$u_r(\theta, a) = -(a - \theta)^2,$$

and the expert's payoff is

$$u_s(\theta, a) = -[a - (\theta + b)]^2,$$

where $b \geq 0$ is a "bias" parameter that measures how nearly agents' interests coincide. Notice that the signal m is irrelevant to the payoff functions, *i.e.*, talk is cheap. Although the message space is independent of the state space, we always let them be identical for sake of simplicity.

14.34 Because of the tractability of the "uniform-quadratic" specification, much of the cheap talk literature, restricts attention to this case.

Quadratic loss means that the marginal cost is increasing in the distance between the state and action. It also means that the players are risk-averse; they prefer a constant gap to a varying gap (depending on θ) with the same mean.

14.35 Note that under this set up, the expert consistently prefers a bigger action than the decision maker (since $b \geq 0$). A more general case, discussed in [Melumad and Shibano \(1991\)](#), is to have

$$u_s(\theta, a) = -[a - (\lambda\theta + b)]^2,$$

in which case the difference between the ideal action of the expert and decision maker depends on θ . When λ is greater than one and b is zero, the agent prefers a proportionally greater than action.

14.36 The sequence of play is as follows:

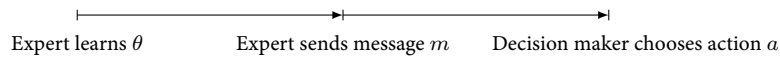


Figure 14.9

- The expert learns her type $\theta \in \Theta$;
- The expert sends a message $m \in \Theta$ to the decision maker; the message may be random, and can be viewed as a noisy estimate of θ ;
- The decision maker processes the information in the expert's message and chooses an action a , which determines players' payoffs.

14.37 In this cheap talk game, a pure-strategy PBE of this game consists of

- a strategy for the expert, denoted $m^*(\theta): \Theta \rightarrow \Theta$,
- a strategy for the decision maker, denoted $a^*(m): \Theta \rightarrow \mathbb{R}$,

- a belief system, denoted $\mu^*(\cdot | m) \in \Delta(\Theta)$,

such that

- Given the decision maker's strategy $a^*(m)$, the expert of type θ send a message $m^*(\theta)$ so that

$$m^*(\theta) \in \arg \max_{m \in \Theta} u_s(\theta, a^*(m), b).$$

- The decision maker's action $a^*(m)$ satisfies

$$a^*(m) \in \arg \max_{a \in \mathbb{R}} \int_{\Theta} u_r(\theta, a) d\mu^*(\theta | m).$$

- $\mu^*(\cdot | m)$ is derived via the Bayes' rule whenever possible.

14.4.1 The model with two types

14.38 We first consider the simplest case in which $\Theta = \{\theta_L, \theta_H\}$. The decision maker initially regards the two states as equally likely.

14.39 We now investigate the conditions under which communication can be informative. In this case, the strategy profile $m(\theta) = \theta, a(m) = m$ and the beliefs system $\mu(\theta | \theta) = 1$ consist of a PBE.

In the following, we check the incentives facing the expert. Consider Figure 14.10, which shows the the expert's payoffs in different states.




Figure 14.10: The expert's payoffs with two states

- (1) Clearly, the expert has no incentive to misrepresent the facts when the state is θ_H .
- (2) If the state instead is θ_L , a truthful (and trusted) report by the expert induces a policy $a = \theta_L$. This is smaller than the expert's ideal policy of $\theta_L + b$ in state θ_L . If the expert instead claims that the state is θ_H , the policy outcome will be $a = \theta_H$. The expert may prefer this larger policy, but it also might be too large even for her tastes. The expert will report truthfully in state θ_L if and only if

$$(\theta_L + b) - \theta_L \leq \theta_H - (\theta_L + b).$$

Notice that this inequality is satisfied for the case depicted in Figure 14.10.

(3) We can rewrite the inequality as a limitation on the size of the divergence in preferences; that is,

$$b \leq \frac{\theta_H - \theta_L}{2}. \quad (14.1)$$

14.40 When Equation (14.1) is satisfied, there exists a PBE with informative communication. In such an equilibrium, the expert educates the decision maker about the state of the world. The equilibrium that results is fully revealing, because the decision maker learns the true state for all possible values of the random variable θ .

14.41 When $\frac{\theta_H - \theta_L}{2} \geq b \geq \frac{\theta_H - \theta_L}{4}$, there also exists a PBE with uninformative communication:

$$m(\theta) = \theta_H, \quad a(\theta) = \begin{cases} \frac{\theta_L + \theta_H}{2}, & \text{if } \theta = \theta_H \\ \theta_L, & \text{if } \theta = \theta_L \end{cases}, \quad \mu(\theta_H | \theta_H) = \mu(\theta_L | \theta_H) = \frac{1}{2}, \quad \mu(\theta_L | \theta_L) = 1.$$

When $\frac{\theta_H - \theta_L}{2} \geq b$, there exists a mixed-strategy PBE with uninformative communication:

$$m(\theta) = \frac{1}{2}\theta_H + \frac{1}{2}\theta_L, \quad a(m) = \frac{1}{2}\theta_H + \frac{1}{2}\theta_L, \quad \mu(\theta_H | \theta) = \mu(\theta_L | \theta) = \frac{1}{2}.$$

They are babbling equilibria, where no information is conveyed from the expert to the decision maker.

14.42 When $\frac{\theta_H - \theta_L}{2} \geq b > \frac{\theta_H - \theta_L}{4}$, there exists a mixed-strategy PBE with partially informative communication:

$$\begin{cases} m(\theta_H) = \theta_H \\ m(\theta_L) = \frac{1-2b}{2b}\theta_H + \frac{4b-1}{2b}\theta_L \end{cases} \quad \begin{cases} a(\theta_H) = 2b \\ a(\theta_L) = \theta_L \end{cases} \quad \begin{cases} \mu(\theta_H | \theta_H) = 2b \\ \mu(\theta_L | \theta_L) = 1 \end{cases}.$$

14.43 If, in contrast, Equation (14.1) is not satisfied, the expert's message lacks credibility. The decision maker would know in such circumstances that the expert had an incentive to announce the state as θ_H no matter what the true state happened to be. For this reason, the expert's message is uninformative, and the decision maker is well justified in ignoring its content. In the event, the decision maker sets the policy $a = \frac{\theta_H + \theta_L}{2}$ that matches her prior expectation about the mean value of θ . Evidently, the transmission of information via cheap talking requires a sufficient degree of alignment between the interests of the decision maker and the expert.

14.44 In summary:

- When $\frac{\theta_H - \theta_L}{4} \geq b$, only babbling and fully revealing can be equilibrium.
- When $\frac{\theta_H - \theta_L}{2} \geq b > \frac{\theta_H - \theta_L}{4}$, (mixed) babbling, fully revealing and (mixed) partially revealing can be equilibrium.
- When $b > \frac{\theta_H - \theta_L}{2}$, the "unique" equilibrium is the babbling equilibrium.

14.45 Numerical example: Let $\theta_L = 0$, $\theta_H = 1$, $\frac{1}{2} \geq b > \frac{1}{4}$.

The comparison of the decision maker's expected utility:

$$0 \text{ for fully revealing} > -\frac{1}{2}(1 - 2b) \text{ for partially revealing} > -\frac{1}{4} \text{ for babbling.}$$

The comparison of the expert's expected utility:

$$-b^2 \text{ for fully revealing} > -\frac{1}{2}(1 - b)^2 - \frac{1}{2}b^2 \text{ for partially revealing} > -\frac{1}{2}(\frac{1}{2} - b)^2 - \frac{1}{2}(\frac{1}{2} + b)^2 \text{ for babbling.}$$

14.4.2 The model with three types

14.46 We then consider the case that $\Theta = \{\theta_L, \theta_M, \theta_H\}$.

14.47 Similar with the two-state case, the θ_L -type expert will reveal his information truthfully if and only if

$$b \leq \frac{\theta_M - \theta_L}{2},$$

and the θ_M -type expert will reveal his information truthfully if and only if

$$b \leq \frac{\theta_H - \theta_M}{2}.$$

Therefore, truth-telling is a PBE strategy if and only if

$$b \leq \min \left\{ \frac{\theta_M - \theta_L}{2}, \frac{\theta_H - \theta_M}{2} \right\}.$$

14.48 Section 18.1 in [Tadelis \(2013\)](#).

14.4.3 The model with a continuum of types

14.49 We finally turn to consider the case that $\Theta = [0, 1]$. As the number of possible states grows, full revelation becomes ever more difficult to achieve. For a sender to be able to distinguish among all possible states, b must be smaller than one-half of the distance between any two of them. But as the number of states tends to infinity—as it must, for example, when θ represents a continuous variable—this requirement becomes impossible to fulfil.

14.50 Proposition: If the expert is even slightly biased, all equilibria entail some information loss.

Proof. If the expert's message always revealed the true state and the decision maker believed him, then the expert would have the incentive to exaggerate the state: in some state θ , he would report $\theta + b$. \square

14.51 Proposition: There always exists a “babbling equilibrium” in which the sender always send the same message and the message is always ignored.

Proof. (1) Let the sender's strategy be to send a message $m_0 \in [0, 1]$ regardless of θ .

(2) This means that the message is completely uninformative and receiver believes that θ is distributed uniformly on $[0, 1]$.

(3) Conditioning on receiving the message m_0 , receiver maximizes his expected payoff

$$\max_{a \in \mathbb{R}} \mathbf{E}_\theta u_r(\theta, a) = \int_0^1 -(a - \theta)^2 d\theta = -\frac{1}{3} + a - a^2,$$

which is maximized when $a = \frac{1}{2}$.

(4) Let receiver's off-equilibrium path beliefs be

$$\text{Prob}\{\theta = \frac{1}{2} \mid m \neq m_0\} = 1$$

so that his off-equilibrium path best response to any other message is $a = \frac{1}{2}$ as well.

- (5) It is easy to see that sender is indifferent between any of his message and hence choosing $m = m_0$ is the best response.

□

14.52 Theorem (Theorem 1 in Crawford and Sobel (1982)): All the perfect Bayesian equilibria are equivalent to a partially pooling equilibrium of the following form: the type space is divided into the n intervals (steps)

$$[x_0 = 0, x_1), [x_1, x_2), \dots, [x_{n-1}, 1 = x_n];$$

all the types in a given interval send the same message, but types in different intervals send different messages.

Idea of proof. Since in equilibrium $a^*(m)$ is weakly increasing, every points in between must send the same message.

□


We will refer to the message sent when $\theta \in [x_{i-1}, x_i]$ as m_i for $i = 1, 2, \dots, n$.

14.53 A two-step ($n = 2$) equilibrium (m^*, a^*, μ^*):

- (1) Suppose all the types in the interval $[0, x_1)$ send one message m_1 while those in $[x_1, 1]$ send another message m_2 .
- (2) After receiving the message m_1 from the types in $[0, x_1)$, the decision maker will believe that the expert's type is uniformly distributed on $[0, x_1)$, so the decision maker's optimal action will be $\frac{x_1}{2}$; likewise, after receiving the message m_2 from the types in $[x_1, 1]$, the decision maker's optimal action will be $\frac{x_1+1}{2}$.
- (3) For the types in $[0, x_1)$ to be willing to send their message m_1 , it must be that all these types prefer the action $\frac{x_1}{2}$ to the action $\frac{x_1+1}{2}$; likewise, all the types above x_1 must prefer $\frac{x_1+1}{2}$ to $\frac{x_1}{2}$.
- (4) Since the expert's utility is symmetric around her optimal action $\theta + b$, the type- θ expert prefers $\frac{x_1}{2}$ to $\frac{x_1+1}{2}$ if the midpoint between these two actions exceeds that type's optimal action $\theta + b$, but prefers $\frac{x_1+1}{2}$ to $\frac{x_1}{2}$ if $\theta + b$ exceeds the midpoint.
- (5) Thus, for a two-step equilibrium to exist, the x_1 -type expert must be indifferent between $\frac{x_1+1}{2}$ and $\frac{x_1}{2}$:

$$x_1 + b = \frac{1}{2} \left[\frac{x_1}{2} + \frac{x_1+1}{2} \right],$$

that is, $x_1 = \frac{1}{2} - 2b$.



- (6) Since the type space is $\Theta = [0, 1]$, x_1 must be positive, so a two-step equilibrium exists only if $b < \frac{1}{4}$; for $b \geq \frac{1}{4}$ the players' preferences are too dissimilar to allow even the limited communication.
- (7) To complete the characterization of this two-step equilibrium, we address the issue of messages that are off the equilibrium path. For example, let the expert's strategy be that all types $\theta < x_1$ send the message $m_1 \in [0, x_1)$ and all types $\theta \geq x_1$ send the message $m_2 \in [x_1, 1]$. Then we may let the decision maker's off-path belief after observing any message from $[0, x_1) \setminus \{m_1\}$ be that θ is uniformly distributed on $[0, x_1)$, and after observing any message from $[x_1, 1] \setminus \{m_2\}$ be that θ is uniformly distributed on $[x_1, 1]$.

14.54 An n -step equilibrium (m^*, a^*, μ^*) :

- (1) By Bayes' rule, $\mu^*((a, b) \mid m_i) = \frac{|(a, b) \cap [x_{i-1}, x_i]|}{x_i - x_{i-1}}$.
- (2) Sequential rationality implies that

$$a^*(m_i) = \frac{x_i + x_{i-1}}{2}.$$

- (3) In equilibrium, the x_i -type expert must be indifferent between m_i and m_{i+1} for $i = 1, 2, \dots, n-1$. Given the quadratic-loss utility function, it must be that

$$(x_i + b) - \frac{x_{i-1} + x_i}{2} = \frac{x_i + x_{i+1}}{2} - (x_i + b),$$

equivalently,

$$(x_{i+1} - x_i) = (x_i - x_{i-1}) + 4b.$$

The width of each step increases by $4b$.

- (4) If the first step is of length d , then the boundary condition must imply

$$d + (d + 4b) + \dots + [d + (n-1)4b] = 1,$$

equivalently,

$$nd + n(n-1)2b = 1.$$

- (5) Hence, given any n such that $n(n-1)2b < 1$, there exists a value of d such that $nd + n(n-1)2b = 1$. That is, there is an n -step equilibrium as long as $n(n-1)2b < 1$.


- (6) Since the length of the first step must be positive, the largest possible number of steps in such an equilibrium, $n^*(b)$, is the largest value of n such that $n(n-1)2b < 1$, i.e., $n^*(b)$ is the largest integer less than

$$\frac{1}{2} \left[1 + \sqrt{1 + \frac{2}{b}} \right].$$

Imprecise messages can still be credible when the interests of the expert and the decision maker do not align completely.

- 14.55 Remark: More communication can occur through cheap talk when the players' preferences are more closely aligned. But perfect communication cannot occur unless the players' preferences are perfectly aligned.

- 14.56 Numerical example: $b = \frac{1}{32}$. Then $n^*(b) = 4$, $d = \frac{1}{16}$, $x_1 = \frac{1}{16}$, $x_2 = \frac{4}{16}$, $x_3 = \frac{9}{16}$.



- 14.57 Remark: Theorem 14.52 crucially depends on the quadratic utility assumption. By this assumption, expert at different states has different points. This makes the expert has incentives to reveal some information.

However, if we assume that the expert's utility function is monotonic in a , for example $u_s(\theta, a) = a - \theta$, then the unique equilibrium is the babbling equilibrium.

- 14.58 $n^*(b)$ decreases in b but approaches infinity only as b approaches zero: more communication can occur through cheap talk when the players' preferences are more closely aligned, but perfect communication cannot occur unless the players' preferences are perfectly aligned.

If there exists an equilibrium with n messages, there must be other equilibria with less than n messages. It always includes the babbling equilibrium in which the decision maker never listen the expert and the expert never convey the true information.

- 14.59 Expected welfare analysis: let us rank the equilibria by evaluating the expected welfare of the decision maker and the expert in each possible equilibrium.

- (1) Since the decision maker's utility function is $u_r(\theta, a) = -(\theta - a)^2$ and she sets $a = \frac{x_{i-1} + x_i}{2}$ if the decision maker heard θ is in $[x_{i-1}, x_i]$.

(2) Hence, the decision maker's expected welfare is

$$U_r(n) = - \sum_{i=1}^n \int_{x_{i-1}}^{x_i} \left(\frac{x_{i-1} + x_i}{2} - \theta \right)^2 d\theta = - \frac{1}{12} \sum_{i=1}^n (x_i - x_{i-1})^3 = \frac{1}{12n^2} + \frac{b^2(n^2 - 1)}{3}.$$

(3) Likewise, the expert's expected welfare is

$$U_s(n) = - \sum_{i=1}^n \int_{x_{i-1}}^{x_i} \left(\frac{x_{i-1} + x_i}{2} - \theta - b \right)^2 d\theta = U_r(n) - b^2.$$

(4) Since $U_r(n)$ is an increasing function of n . Thus, we can conclude that in the *ex ante* sense, the more n we get, the better equilibrium we achieve.

14.4.4 Optimal communication mechanism (the commitment case)

14.60 Literature: [Melumad and Shibano \(1991\)](#).

14.61 We consider the case where decision maker can commit to a decision rule (a strategy for the decision maker) which is a mapping from the message space to the action space: $a: \Theta \rightarrow \mathbb{R}$.

14.62 We still assume that $\Theta = [0, 1]$ and θ is uniformly distributed on Θ .

14.63 By the revelation principle we need to consider truth-telling mechanism in which the sender's strategy is $m(\theta) = \theta$ for all θ . This requires the IC condition:

$$u_s(\theta, a(\theta)) \geq u_s(\theta, a(\theta'))$$

for any $\theta, \theta' \in \Theta$.

14.64 Lemma: $a: \Theta \rightarrow \mathbb{R}$ is increasing.

Proof. Consider any two distinct θ and θ' . The IC condition implies:

$$-(a(\theta) - \theta - b)^2 \geq -(a(\theta') - \theta - b)^2 \text{ and } -(a(\theta') - \theta' - b)^2 \geq -(a(\theta) - \theta' - b)^2.$$

These two inequalities yield

$$-(a(\theta) - \theta - b)^2 - (a(\theta') - \theta' - b)^2 \geq -(a(\theta') - \theta - b)^2 - (a(\theta) - \theta' - b)^2.$$

That is,

$$(a(\theta) - a(\theta'))(\theta - \theta') \geq 0.$$

Therefore, a is increasing. □

14.65 Lemma: If a is continuous and strictly increasing on (θ_1, θ_2) , then $a(\theta) = \theta + b$ on (θ_1, θ_2) .

Proof. (1) Suppose that there exists θ such that $a(\theta) \neq \theta + b$.

(2) Without loss of generality, assume $a(\theta) > \theta + b$.

(3) Since a is continuous and strictly increasing, there exists $\theta' < \theta$ such that $\theta + b < a(\theta') < a(\theta)$.

(4) Then it is optimal for the sender under θ to misreport θ' . Contradiction.

□

14.66 Lemma: If a is discontinuous at θ_0 , the discontinuity must be a jump discontinuity that satisfies

(i) $u_s(\theta_0, a^-(\theta_0)) = u_s(\theta_0, a^+(\theta_0))$, where $a^-(\theta_0) = \lim_{\theta \uparrow \theta_0} a(\theta)$ and $a^+(\theta_0) = \lim_{\theta \downarrow \theta_0} a(\theta)$.

(ii) $a(\theta) = \begin{cases} a^-(\theta_0), & \text{if } \theta \in [a^-(\theta_0) - b, \theta_0), \\ a^+(\theta_0), & \text{if } \theta \in (\theta_0, a^+(\theta_0) - b]. \end{cases}$

(iii) $a(\theta_0) \in \{a^-(\theta_0), a^+(\theta_0)\}$.

Proof. (i) Assume that $u_s(\theta_0, a^-(\theta_0)) < u_s(\theta_0, a^+(\theta_0))$. Then the sender will strictly prefer $a^+(\theta_0)$ to $a^-(\theta_0) \approx a(\theta')$ at θ' slightly less than θ_0 . Contradiction.

(ii) Since the sender's favorite action when $\theta = a^-(\theta_0) - b$ is $a^-(\theta_0)$, the IC condition requires that

$$a(a^-(\theta_0) - b) = a^-(\theta_0).$$

Since a is increasing, a must be flat between $[a^-(\theta_0) - b, \theta_0)$.

(iii) Assume that $a(\theta_0) \in (a^-(\theta_0), a^+(\theta_0))$. Then the sender will strictly prefer $a(\theta_0)$ to either $a^+(\theta_0)$ or $a^-(\theta_0)$. So will the sender when θ' is near θ_0 since u_s is continuous in θ . Contradiction.

□

14.67 The following figure depicts a general IC decision rule according to the previous lemmas.




Figure 14.11

14.68 Lemma: The decision rule a is everywhere continuous.

Proof. (1) Suppose that a is not continuous at θ_0 .

(2) Then we have

$$a(\theta) = \begin{cases} a^-(\theta_0), & \text{if } \theta \in [a^-(\theta_0) - b, \theta_0), \\ a^+(\theta_0), & \text{if } \theta \in (\theta_0, a^+(\theta_0) - b]. \end{cases}$$

(3) Note that replacing any segment of an IC decision rule by some other IC segment does not affect the IC property of the decision rule.

(4) We will argue that (i) it is incentive compatible to replace this part of the rule by $a'(\theta) = \theta + b$ (the sender's favorite action) and (ii) that the receiver will benefit from the change in contradiction to the assumed optimality of $a(\theta)$.

(5) When $\theta < a^-(\theta_0) - b$, we have $\theta + b < a^-(\theta_0)$, and hence

$$-(a^-(\theta_0) - \theta - b)^2 \geq -(a - \theta - b)^2$$

for any $a \in (a^-(\theta_0), a^+(\theta_0))$.




Figure 14.12

(6) Therefore, the sender prefers $a^-(\theta_0)$ to any $a \in (a^-(\theta_0), a^+(\theta_0))$ when $\theta < a^-(\theta_0) - b$.

(7) Similarly, the sender prefers $a^+(\theta_0)$ to any $a \in (a^-(\theta_0), a^+(\theta_0))$ when $\theta > a^+(\theta_0) - b$.

(8) So, the change will not create incentives for the sender to deviate when $\theta \notin (a^-(\theta_0), a^+(\theta_0))$.

(9) When $\theta \in (a^-(\theta_0), a^+(\theta_0))$, the sender will obtain his best action by telling the truth, so there is no incentive to deviate.

(10) The receiver's expected utility is, then,

$$-\int_{a^-(\theta_0)-b}^{a^+(\theta_0)-b} (a'(\theta) - \theta)^2 d\theta = -\int_{a^-(\theta_0)-b}^{a^+(\theta_0)-b} (\theta + b - \theta)^2 d\theta = -b^2 (a^+(\theta_0) - a^-(\theta_0)).$$

(11) The receiver's original expected utility is

$$-\int_{a^-(\theta_0)-b}^{\theta_0} (a^-(\theta_0) - \theta)^2 d\theta - \int_{\theta_0}^{a^+(\theta_0)-b} (a^+(\theta_0) - \theta)^2 d\theta = -\int_{\theta_0 - a^+(\theta_0)}^{\theta_0 - a^-(\theta_0)} x^2 dx,$$

which is strictly less than $-b^2 (a^+(\theta_0) - a^-(\theta_0))$ due to the Jensen's inequality.

□

14.69 Theorem: The optimal decision rule is

$$a^*(\theta) = \begin{cases} \theta + b, & \text{if } \theta \leq 1 - 2b, \\ 1 - b, & \text{if } \theta > 1 - 2b. \end{cases}$$

Proof. (1) We have already shown that the optimal a must be continuous and equal $\theta + b$ when it is strictly increasing. This means that the optimal rule must be of the form

$$a^*(\theta) = \begin{cases} \underline{\theta} + b, & \text{if } \theta \leq \underline{\theta}, \\ \theta + b, & \text{if } \theta \in (\underline{\theta}, \bar{\theta}), \\ \bar{\theta} + b, & \text{if } \theta \geq \bar{\theta}. \end{cases}$$

(2) The receiver chooses $\underline{\theta} \geq 0$ and $\bar{\theta} \in [\underline{\theta}, 1]$ to maximize:

$$U_r(\underline{\theta}, \bar{\theta}) = - \int_0^{\underline{\theta}} (\underline{\theta} + b - \theta)^2 d\theta - \int_{\underline{\theta}}^{\bar{\theta}} b^2 d\theta - \int_{\bar{\theta}}^1 (\bar{\theta} + b - \theta) d\theta.$$

(3) It is easy to show that

$$\frac{\partial U_r}{\partial \underline{\theta}} = -2 \int_0^{\underline{\theta}} (\underline{\theta} + b - \theta) d\theta < 0,$$

which implies that $\underline{\theta}^* = 0$.

(4) Similarly,

$$\frac{\partial U_r}{\partial \bar{\theta}} = -2 \int_{\bar{\theta}}^1 (\bar{\theta} + b - \theta) d\theta.$$

(5) By letting $\frac{\partial U_r}{\partial \bar{\theta}} = 0$, we have $\bar{\theta} = 1 - 2b$ or $\bar{\theta} = 1$. However, the second derivative is negative at $1 - 2b$ but positive at 1. Therefore, U_r is maximized at $\bar{\theta}^* = 1 - 2b$.

□

14.70 The optimal rule is to set a limit on the highest action the sender (who is biased for higher actions) can take and let the sender picks the action.




Figure 14.13

14.71 Further readings: [Krishna and Morgan \(2004\)](#) on multi-stage communication, [Battaglini \(2002\)](#) and [Ambrus and Takahashi \(2008\)](#) on multiple senders, [Chen, Kartik and Sobel \(2008\)](#) on equilibrium selection.

14.5 Bayesian persuasion

14.72 Literature: [Kamenica and Gentzkow \(2011\)](#).

Reading: [Gentzkow and Kamenica \(2016\)](#), [Gentzkow and Kamenica \(2016\)](#), [Li and Norman \(2015\)](#), [Wang \(2013\)](#).

14.73 Consider an example of a prosecutor (sender) trying to convince a judge (receiver) that a defendant is guilty.

When the defendant is indeed guilty, revealing the facts of the case will tend to help the prosecutor's case. When the defendant is innocent, revealing facts will tend to hurt the prosecutor's case.

Can the prosecutor structure his arguments, selection of evidence, *etc.* so as to increase the probability of conviction by a rational judge on average?

14.5.1 Example

14.74 There are two states of the world: the defendant is either guilty or innocent. The judge must choose one of two actions: to acquit or innocent.

The judge gets utility 1 for choosing the just action (convict when guilty and acquit when innocent) and utility 0 for choosing the unjust action (convict when innocent and acquit when guilty).

The prosecutor gets utility 1 if the judge convicts and utility 0 if the judge acquits, regardless of the state.

The prosecutor and the judge share a prior belief $\mu_0(\text{guilty}) = 0.3$.

14.75 The prosecutor conducts an investigation and is required by law to report its full outcome.

We can think of the choice of the investigation as consisting of the decisions on whom to subpoena, what forensic tests to conduct, what questions to ask an expert witness, *etc.*

We formalize an investigation as distributions $q(\cdot \mid \text{guilty})$ and $q(\cdot \mid \text{innocent})$ on some set of signal realizations. The prosecutor chooses q and must honestly report the signal realization to the judge.

14.76 If the prosecutor chooses a fully informative investigation, one that leaves no uncertainty about the state, the judge convicts 30 percent of the time.

14.77 However, the prosecutor can do better. For example, he can choose the following binary signal $\{i, g\}$ such that

$$q(g \mid \text{guilty}) = 1 \text{ and } q(g \mid \text{innocent}) = \frac{3}{7}.$$

The first message g leads to a posterior

$$\mu_g = \frac{\mu_0(\cdot) \cdot q(g \mid \cdot)}{\mu_0(\text{guilty}) \cdot q(g \mid \text{guilty}) + \mu_0(\text{innocent}) \cdot q(g \mid \text{innocent})} = \frac{1}{2}\delta_{\text{innocent}} + \frac{1}{2}\delta_{\text{guilty}}.$$

That is, the sender says “this person is guilty enough to convict”. The second message i leads to a posterior

$$\mu_i = \frac{\mu_0(\cdot) \cdot q(i \mid \cdot)}{\mu_0(\text{guilty}) \cdot q(i \mid \text{guilty}) + \mu_0(\text{innocent}) \cdot q(i \mid \text{innocent})} = \delta_{\text{innocent}}.$$

That is, the sender says “this person is innocent”.

This leads the judge to convict with (*ex ante* level) probability 60 percent ($0.7 \times \frac{3}{7} + 0.3 \times 1 = 0.6$).

The judge knows 70 percent of defendants are innocent, yet she convicts 60 percent of them! She does so even though she is fully aware that the investigation was designed to maximize the probability of conviction.

14.78 The posteriors:

$$\begin{aligned} \mu(\text{guilty} \mid g) &= \frac{\mu_0(\text{guilty}) \cdot q(g \mid \text{guilty})}{\mu_0(\text{guilty}) \cdot q(g \mid \text{guilty}) + \mu_0(\text{innocent}) \cdot q(g \mid \text{innocent})} = \frac{1}{2} \\ \mu(\text{innocent} \mid i) &= \frac{\mu_0(\text{innocent}) \cdot q(i \mid \text{innocent})}{\mu_0(\text{innocent}) \cdot q(i \mid \text{innocent}) + \mu_0(\text{guilty}) \cdot q(i \mid \text{guilty})} = 1 \end{aligned}$$

Clearly,

$$\mu_0(\theta) = \text{Prob}(g) \cdot \mu(\theta \mid g) + \text{Prob}(i) \cdot \mu(\theta \mid i)$$

for $\theta = \text{guilty}$ or innocent .

14.5.2 A model of Bayesian persuasion

14.79 There is a sender and a receiver. The state space Θ is finite and the action space is compact. The sender and the receiver share the (full support) common prior $\mu_0(\cdot)$ on Θ .

14.80 Receiver has a continuous utility function $u_r(a, \theta)$, and sender has a continuous utility function $u_s(a, \theta)$.

14.81 The special feature of the model is that the sender’s strategy is to pick an information structure (or information disclosure rule), which consists of a finite set M of messages and a mapping $q: \Theta \rightarrow \Delta(M)$. Here, $q(m \mid \theta)$ describes the probability that the receiver hears the message m when θ is the true state.

Note that the mapping q can be induced by a joint distribution on states and messages.

14.82 The persuasion game works as follows.

- (1) The sender selects the information structure (M, q) .
- (2) Given the information structure, the receiver selects an action rule (what to do following any message).

The sender is potentially informed but can choose to commit to public information acquisition and disclosure.




Figure 14.14

14.83 In equilibrium:

- The receiver's action rule is a best response to the information structure (for each realized message).

After observing the sender's choice of information structure (M, q) and a message realization $m \in M$ (drawn by nature), the receiver forms the posterior $\mu_m(\cdot) \in \Delta(\Theta)$ via Bayes' rule:

$$\mu_m(\theta) = \frac{\mu_0(\theta) \cdot q(m | \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot q(m | \theta')},$$

and takes an action $\hat{a}(\mu_m) \in A$ to maximize her expected utility

$$\mathbf{E}_{\mu_m}[u_r(a, \theta)] = \sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_r(a, \theta).$$

If there are multiple best actions, assume that the receiver picks one of the optimizers that maximizes the sender's expected utility $\mathbf{E}_{\mu_m}[u_s(a, \theta)] = \sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_s(a, \theta)$.

- The sender selects the information structure (M, q) to maximize her expected utility given the receiver's response.

Given an information structure (M, q) , let

$$P(m) = \sum_{\theta \in \Theta} q(m | \theta) \cdot \mu_0(\theta)$$

be the probability that message m is heard.

The sender's *ex ante* expected utility is

$$U_s(M, q) = \mathbf{E}_P \left[\mathbf{E}_{\mu_m} \left[u_s(\hat{a}(\mu_m), \theta) \right] \right] = \sum_{m \in M} P(m) \left[\sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_s(\hat{a}(\mu_m), \theta) \right].$$

Given the receiver's action rule, the sender chooses the information structure (M, q) to maximize his expected utility.

14.84 Key assumption: The sender cannot distort or conceal information once the message is realized.

14.85 We simplify the analysis further by noting that, without loss of generality, we can restrict our attention to a particular class of information structures. Say that an information structure is straightforward if $M \subseteq A$ and the receiver's equilibrium action equals the message realization.

In other words, a straightforward information structure produces a "recommended action" and the receiver always follows the recommendation.

14.86 Lemma: The following are equivalent:

- (i) There exists an information structure with the expected utility u_s^* .
- (ii) There exists a straightforward information structure with expected utility u_s^* .

Proof. (1) By definition, (ii) implies (i).

- (2) Given an information structure (M, q) with the expected utility u_s^* , let

$$M^a = \{m \in M \mid \hat{a}(\mu_m) = a\} \text{ for each } a \in A.$$

- (3) Consider an information structure (M', q') with $M' = A$ and

$$q'(a \mid \theta) = \sum_{m \in M^a} q(m \mid \theta).$$

- (4) For each $a \in A$, we have

$$\begin{aligned} \sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_r(a, \theta) &= \sum_{\theta \in \Theta} \frac{\mu_0(\theta) \cdot \sum_{m \in M^a} q(m \mid \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \cdot u_r(a, \theta) \\ &= \frac{1}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) \cdot u_r(a, \theta) \\ &= \frac{1}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) \cdot u_r(\hat{a}(\mu_m), \theta) \\ &\geq \frac{1}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot \sum_{m' \in M^a} q(m' \mid \theta')} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) \cdot u_r(a', \theta) \\ &= \sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_r(a', \theta) \end{aligned}$$

for each $a' \in A$. That is, a is an optimal response for the receiver to the realization a from q' .

- (5) Sender's expected utility under the new information structure is

$$\begin{aligned} U_s(M', q') &= \sum_{a \in A} P(a) \left[\sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_s(\hat{a}(\mu'_a), \theta) \right] \\ &= \sum_{a \in A} P(a) \left[\sum_{\theta \in \Theta} \mu'_a(\theta) \cdot u_s(a, \theta) \right] \\ &= \sum_{a \in A} \left(\sum_{\theta' \in \Theta} \sum_{m \in M^a} q(m \mid \theta') \cdot \mu_0(\theta') \right) \left[\sum_{\theta \in \Theta} \frac{\mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta)}{\sum_{\theta'' \in \Theta} \mu_0(\theta'') \sum_{m \in M^a} q(m \mid \theta'')} u_s(a, \theta) \right] \\ &= \sum_{a \in A} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M^a} q(m \mid \theta) u_s(\hat{a}(\mu_m), \theta) \\ &= \sum_{m \in M} P(m) \left[\sum_{\theta \in \Theta} \mu_m(\theta) \cdot u_s(\hat{a}(\mu_m), \theta) \right] = U_s(M, q) \end{aligned}$$

□

14.5.3 Simplified problem

14.87 An information structure (M, q) naturally induces a distribution over posterior beliefs.

Conversely, a distribution over posterior beliefs $\tau \in \Delta(\Delta(\Theta))$ is induced by an information structure (M, q) if

$$\begin{aligned}\mu_m(\theta) &= \frac{\mu_0(\theta) \cdot q(m \mid \theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta') \cdot q(m \mid \theta')} \text{ for all } \theta \text{ and } m, \\ \text{support}(\tau) &= \{\mu_m\}_{m \in M}, \\ \tau(\mu) &= \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \text{ for all } \mu.\end{aligned}$$

14.88 When both sender and receiver hold some posterior μ , the sender's expected utility is equal to

$$U_s(\mu) = \mathbf{E}_\mu [u_s(\hat{a}(\mu), \theta)] = \sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta).$$

14.89 Since sender's and receiver's beliefs coincide, the sender's utility from any distribution of posteriors τ (induced by some information structure) is

$$U_s(\tau) = \mathbf{E}_\tau U_s(\mu) = \mathbf{E}_\tau \left[\mathbf{E}_\mu [u_s(\hat{a}(\mu), \theta)] \right] = \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right].$$

14.90 A distribution of posteriors τ is Bayesian plausible if the expected posterior probability equals the prior:

$$\mathbf{E}_\tau \mu = \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \cdot \mu = \mu_0.$$

14.91 Clearly, if τ is induced by an information structure (M, q) , then τ is Bayesian plausible.

For each θ ,

$$\begin{aligned}\sum_{\mu \in \text{support}(\tau)} \tau(\mu) \cdot \mu(\theta) &= \sum_{\mu \in \text{support}(\tau)} \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \cdot \mu(\theta) \\ &= \sum_{\mu \in \text{support}(\tau)} \sum_{m: \mu_m = \mu} q(m \mid \theta) \cdot \mu_0(\theta) = \mu_0(\theta)\end{aligned}$$

14.92 Proposition: The following are equivalent:

- (i) There exists an information structure (M, q) such that the expected utility $U_s(M, q)$ is u_s^* .
- (ii) There exists a Bayesian plausible distribution of posteriors τ such that the expected utility $U_s(\tau)$ is u_s^* .

Proof. (1) Given an information structure (M, q) , let τ be the distribution of posteriors induced by (M, q) . Then

$$\begin{aligned}U_s(\tau) &= \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right] \\ &= \sum_{\mu \in \text{support}(\tau)} \left[\sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \right] \cdot \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right] \\ &= \sum_{\mu \in \text{support}(\tau)} \sum_{\theta \in \Theta} \sum_{m: \mu_m = \mu} \sum_{\theta' \in \Theta} q(m \mid \theta') \cdot \mu_0(\theta') \cdot \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \\ &= \sum_{\mu \in \text{support}(\tau)} \sum_{\theta \in \Theta} \sum_{m: \mu_m = \mu} \mu_0(\theta) \cdot q(m \mid \theta) \cdot u_s(\hat{a}(\mu), \theta)\end{aligned}$$

$$= \sum_{m \in M} \sum_{\theta \in \Theta} \mu_0(\theta) \cdot q(m \mid \theta) \cdot u_s(\hat{a}(\mu_m), \theta) = U_s(M, q).$$

(2) Assume that τ is a Bayesian plausible distribution of posteriors with the expected utility u_s^* .

(3) Then we have

$$u_s^* = \sum_{\mu \in \text{support}(\tau)} \tau(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right].$$

(4) Since Θ is finite, Carathéodory's theorem¹ implies that there exists a Bayesian plausible τ^* with finite support such that

$$u_s^* = \sum_{\mu \in \text{support}(\tau^*)} \tau^*(\mu) \left[\sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) \right].$$

(5) Define M so that $\text{support}(\tau^*) = \{\mu_m\}_{m \in M}$ and let $q(m \mid \theta) = \tau^*(\mu_m) \frac{\mu_m(\theta)}{\mu_0(\theta)}$.

(6) Clearly, $U_s(M, q) = U_s(\tau^*)$.

□

14.93 The key implication is that to evaluate whether the sender benefits from persuasion and to determine the value of an optimal information structure we need only ask how $U_s(\tau)$ varies over the space of Bayesian plausible distributions of posteriors.

14.94 Simplified problem:

$$\begin{aligned} & \underset{\tau}{\text{maximize}} && U_s(\tau) = \mathbf{E}_{\tau} U_s(\mu) \\ & \text{subject to} && \mathbf{E}_{\tau} \mu = \mu_0. \end{aligned}$$

14.5.4 Optimal information structure

14.95 The focus on sender-preferred equilibria implies that $U_s(\mu)$ is upper semicontinuous which in turn ensures the existence of an optimal information structure.

14.96 Let \hat{U}_s be the concave closure of the function $U_s(\mu): \Delta(\Theta) \rightarrow \mathbb{R}$:

$$\hat{U}_s(\mu) = \sup \{v \mid (\mu, v) \in \text{co}(U_s)\},$$

where $\text{co}(U_s)$ denotes the convex hull of the graph of U_s .

14.97 Clearly, $\hat{U}_s: \Delta(\Theta) \rightarrow \mathbb{R}$ is a concave function by construction. (check by yourself)

Actually, it is the smallest concave function that is everywhere weakly greater than U_s .

The following figure shows an example of the construction of \hat{U}_s . In the figure, the state space is binary, and we identify a distribution μ with the probability of one of the states.

¹In convex geometry Carathéodory's theorem states that if a point x of \mathbb{R}^d lies in the convex hull of a set P , there is a subset P' of P consisting of $d + 1$ or fewer points such that x lies in the convex hull of P' .




Figure 14.15

14.98 Proposition (Corollary 2 in [Kamenica and Gentzkow \(2011\)](#)): The expected utility of an optimal information structure is $\hat{U}_s(\mu_0)$, and revealing information structure is better than non-revealing if and only if $\hat{U}_s(\mu_0) > U_s(\mu_0)$.

Proof. (1) For any $(\mu', v) \in \text{co}(U_s)$, there exists a distribution of posteriors τ such that

$$(\mu', v) = \mathbf{E}_\tau (\mu, U_s(\mu)) .$$

- (2) Thus, $\text{co}(U_s)$ is the set of (μ', v) such that if the prior is μ' , there exists an information structure with the expected utility v .
- (3) Hence, given the prior μ_0 , $\hat{U}_s(\mu_0)$ is the largest expected utility the sender can achieve with any information structure.

□

14.99 The following figure shows the function $U_s(\mu)$, the concave closure \hat{U}_s , and the optimal information structure for Example 14.74. In the figure, μ denotes the probability that the state is *guilty*.




Figure 14.16

U_s is a step function: the prosecutor's expected utility is 0 whenever $\mu < \frac{1}{2}$ (since the judge will choose *acquit*) and 1 whenever $\mu \geq \frac{1}{2}$ (since the judge will choose *convict*).

The concave closure \hat{U}_s is

$$\hat{U}_s = \begin{cases} 2\mu, & \text{if } \mu \geq \frac{1}{2}, \\ 1, & \text{if } \mu > \frac{1}{2}. \end{cases}$$

It is clear $\hat{U}_s(\mu_0) > U_s(\mu_0)$ and the expected utility of the optimal information structure is $\hat{U}_s(\mu_0)$. The prosecutor can benefit from persuasion if and only if $\mu_0 < \frac{1}{2}$.

By simple calculation, the optimal information structure induces the distribution of posteriors τ^* :

$$\tau^* = \frac{2}{5}\delta_0 + \frac{3}{5}\delta_{\frac{1}{2}}.$$

Let $M = \{i, g\}$ such that $\mu_i = \delta_0$ and $\mu_g = \delta_{\frac{1}{2}}$. Then

$$\begin{aligned} q(g \mid \text{guilty}) &= \tau^*(\mu_g) \frac{\mu_g(\text{guilty})}{\mu_0(\text{guilty})} = \frac{3}{5} \frac{1/2}{0.3} = 1, \\ q(g \mid \text{innocent}) &= \tau^*(\mu_g) \frac{\mu_g(\text{innocent})}{\mu_0(\text{innocent})} = \frac{3}{5} \frac{1/2}{0.7} = \frac{3}{7}. \end{aligned}$$

14.100 Corollary (Remark 1 in [Kamenica and Gentzkow \(2011\)](#)): If $U_s(\mu)$ is concave, the sender does not benefit from persuasion for any prior. If $U_s(\mu)$ is convex and not concave, the sender benefits from persuasion for every prior.

Proof. The sender benefits from persuasion is and only if there exists a τ such that $\mathbf{E}_\tau[U_s(\mu)] > U_s(\mathbf{E}_\tau[\mu])$. \square

14.101 We say “there is information the sender would like to share” if there is a posterior that is better for the sender than the prior, that is, there exists μ such that

$$U_s(\mu) = \sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu), \theta) > \sum_{\theta \in \Theta} \mu(\theta) \cdot u_s(\hat{a}(\mu_0), \theta).$$

In other words, there is a μ such that, if the sender had private information that led him to believe μ , he would prefer to share this information with the receiver rather than have the receiver act based on μ_0 .

We say the receiver’s preference is discrete at belief μ if the receiver’s expected utility from her preferred action $\hat{a}(\mu)$ is bounded away from her expected utility from any other action, *i.e.*, if there is an $\epsilon > 0$ such that for any $a \neq \hat{a}(\mu)$, $\mathbf{E}_\mu u_r(\hat{a}(\mu), \theta) > \mathbf{E}_\mu u_r(a, \theta) + \epsilon$.

14.102 Proposition (Proposition 2 in [Kamenica and Gentzkow \(2011\)](#)): If there is no information the sender would share, the sender does not benefit from persuasion. If there is information the sender would share and the receiver’s preference is discrete at the prior, the sender benefits from persuasion.

Proof. (1) If there is no information the sender would share, then for any information structure τ which induces a distribution of posteriors τ ,

$$\mathbf{E}_\tau U_s(\mu) \leq \underbrace{\mathbf{E}_\tau \mathbf{E}_\mu u_s(\hat{a}(\mu_0), \theta)}_{\text{Bayesian plausibility}} = U_s(\mu_0).$$

Informally, any realization of message m leads the receiver to take an action $\hat{a}(\mu_m)$ the sender weakly dislikes relative to the default action $\hat{a}(\mu_0)$. Hence, a completely noninformative information structure is optimal.

- (2) Since the receiver’s preference is discrete at the prior μ_0 , there exists an $\epsilon > 0$ such that for any $a \neq \hat{a}(\mu_0)$, $\sum_\theta \mu_0(\theta) \cdot u_r(\hat{a}(\mu_0), \theta) > \sum_\theta \mu_0(\theta) \cdot u_r(a, \theta) + \epsilon$.
- (3) Since $u_r(a, \theta)$ is continuous in θ , $\sum_\theta \mu(\theta) \cdot u_r(a, \theta)$ is continuous in μ .
- (4) Thus, there is a $\delta > 0$ such that for any $\mu \in B_\delta(\mu_0)$ and for any $a \neq \hat{a}(\mu_0)$, $\sum_\theta \mu(\theta) \cdot u_r(\hat{a}(\mu_0), \theta) > \sum_\theta \mu(\theta) \cdot u_r(a, \theta)$.
- (5) Hence, $\hat{a}(\mu) = \hat{a}(\mu_0)$ for any $\mu \in B_\delta(\mu_0)$.

- (6) Since there is information the sender would share, there exists μ_h such that $\sum_{\theta} \mu_h(\theta) \cdot u_s(\hat{a}(\mu_h), \theta) > \sum_{\theta} \mu_h(\theta) \cdot u_s(\hat{a}(\mu_0), \theta)$.
- (7) Consider a ray from μ_h through μ_0 . Since μ_0 is not on the boundary of $\Delta(\Theta)$, there exists a belief μ_ℓ on that ray such that $\mu_\ell \in B_\delta(\mu_0)$ and $\mu_0 = \gamma\mu_\ell + (1 - \gamma)\mu_h$ for some $\gamma \in (0, 1)$.
- (8) Consider the Bayesian plausible distribution of posteriors $\tau = \gamma\mu_\ell + (1 - \gamma)\mu_h$.
- (9) We have

$$\begin{aligned} \mathbf{E}_\tau \mathbf{E}_\mu u_s(\hat{a}(\mu), \theta) &= \gamma U_s(\mu_\ell) + (1 - \gamma) U_s(\mu_h) \\ &> \gamma \sum_{\theta} \mu_\ell(\theta) \cdot u_s(\hat{a}(\mu_0), \theta) + (1 - \gamma) \sum_{\theta} \mu_h(\theta) \cdot u_s(\hat{a}(\mu_0), \theta) = U_s(\mu_0). \end{aligned}$$

Therefore, the sender benefits from persuasion. □

14.103 Lemma: If A is finite, the receiver's preference is discrete at the prior generically. □

Proof. Omitted. □

14.104 Application: Lobbying.

Consider a setting where a lobbying group commissions a study with the goal of influencing a benevolent, but nonetheless rational, politician. The politician (Receiver) chooses a unidimensional policy $a \in [0, 1]$. The state $\theta \in [0, 1]$ is the socially optimal policy. The lobbyist (Sender) is employed by the interest group whose preferred action is $a_0 = \alpha\theta + (1 - \alpha)\theta_0$ with $\alpha \in [0, 1]$ and $\theta_0 > 1$. Politician's payoff $-(a - \theta)^2$ and lobbyist's payoff $-(a - a_0)^2$.

Since politician's payoff is $-(a - \theta)^2$, $\hat{a}(\mu) = \mathbf{E}_\mu[\theta]$. Given this \hat{a} , we have

$$U_s(\mu) = -(1 - \alpha)^2 \theta_0^2 + 2(1 - \alpha)^2 \theta_0 \mathbf{E}_\mu[\theta] - \alpha^2 \mathbf{E}_\mu[\theta^2] + (2\alpha - 1)(\mathbf{E}_\mu[\theta])^2.$$

U_s is linear in μ when $\alpha = \frac{1}{2}$, strictly convex when $\alpha > \frac{1}{2}$, and strictly concave when $\alpha < \frac{1}{2}$.

Therefore we have full disclosure if $\alpha > \frac{1}{2}$ and no disclosure if $\alpha < \frac{1}{2}$. There is thus a natural sense in which some alignment of preferences is necessary for information to be communicated in equilibrium even when Sender has the ability to commit.

Note that the lobbyist either commissions a fully revealing study or no study at all.

The optimal information structure is independent of θ_0 . This is important because θ_0 also captures a form of disagreement between the lobbyist and the politician. We might have expected communication to be difficult when θ_0 is much greater than one. Unlike α , however, θ_0 does not affect the way the lobbyist's payoff varies across realizations of a message. The loss the lobbyist suffers from high values of θ_0 is thus a sunk cost and does not affect the decision of how best to persuade.

14.6 Optimal information disclosure

14.105 Literature: [Eső and Szentes \(2007\)](#), [Rayo and Segal \(2010\)](#). [Zhang and Zhou \(forthcoming\)](#)

Chapter 15

Information economics 2: screening

Contents

15.1 The principal-agent model	261
15.1.1 Complete information	262
15.1.2 Incomplete information	265
15.1.3 The revelation principle	269
15.1.4 General utility function for the agent	270
15.1.5 Nonresponsiveness	271
15.2 Pricing a single indivisible good	272
15.3 Nonlinear pricing	275
15.3.1 The model with two types	276
15.3.2 The model with a continuum of types	280
15.4 Insurance	282
15.5 Sequential screening	282
15.5.1 A motivating example	283
15.5.2 Optimal menu of refund contracts: two-type case	283
15.5.3 Sequential screening: general type case	287
15.5.4 Optimal selling mechanism	290

15.1 Signaling (信号传递) and screening (信号甄别) are two primary solutions to adverse selection. In screening, the uninformed party moves first. For example, the uninformed party offers a term of exchange such that the informed party can choose to accept or reject. For example, the uninformed buyer can ask for a long enough warranty to screen out low quality sellers. The insurance market gives a better illustration, where the insurance company, who is uninformed of conditions of the clients, offers different insurance packages, and ideally, each type of clients only find one package acceptable.

15.1 The principal-agent model

15.2 Reference: Chapters 2 and 3 in [Laffont and Martimort \(2002\)](#).

15.3 Consider a consumer (the principal) who wants to delegate to an agent the production of q units of a good.

The value for the principal of these q units is $S(q)$ where $S' > 0$, $S'' < 0$ and $S(0) = 0$. The marginal value of the good is thus positive and strictly decreasing with the number of units bought by the principal.

The production cost of the agent is unobservable to the principal, but it is common knowledge that the marginal cost θ belongs to the set $\Theta = \{\theta_H, \theta_L\}$. The agent can be either efficient (θ_L) or inefficient (θ_H) with respective probabilities β and $1 - \beta$. In other words, he has the cost function

$$c(q, \theta_L) = \theta_L q \text{ with probability } \beta$$

or

$$c(q, \theta_H) = \theta_H q \text{ with probability } 1 - \beta.$$

We denote by $\Delta\theta = \theta_H - \theta_L > 0$ the spread of uncertainty on the agent's marginal cost.

The principal's utility if she purchases q units of the good and pays a monetary transfer t to the agent is

$$S(q) - t,$$

and at this case the agent's utility is

$$t - c(q, \theta).$$

15.4 The economic variables are quantity produced q and the transfer t received by the agent. These variables are both observable and verifiable by a third party such as a benevolent court of law.

Let \mathcal{A} be the set of all feasible contract, that is,

$$\mathcal{A} = \{(q, t) \mid q \in \mathbb{R}_+, t \in \mathbb{R}\}.$$

15.5 The sequence of play is as follows:




Figure 15.1

15.1.1 Complete information

15.6 First suppose that there is no asymmetry of information between the principal and the agent.

Then the principal will try to maximize her utility subject to inducing the agent to accept the proposed contract. Clearly, the agent obtains 0 if he does not take the principal's contract. So the principal will solve the following problem:

$$\begin{aligned} & \underset{(q_i, t_i) \in \mathcal{A}}{\text{maximize}} && S(q) - t \\ & \text{subject to} && t - c(q, \theta_i) \geq 0. \end{aligned}$$

The solution to this problem will be the first-best contract menu $\{(q_i^*, t_i^*)\}_{i=H,L}$ such that

$$S'(q_i^*) = \theta_i.$$

That is, the efficient production levels are obtained by equating the principal's marginal value and the agent's marginal cost.

15.7 Since $S'' < 0$ and $\theta_H > \theta_L$, we have

$$q_L^* > q_H^*,$$

i.e., the optimal production of an efficient agent is greater than that of an inefficient agent.

15.8 The complete information efficient production levels q_H^* and q_L^* should be both carried out if their social values, respectively $W_H^* = S(q_H^*) - \theta_H q_H^*$ and $W_L^* = S(q_L^*) - \theta_L q_L^*$, are non-negative.

The social value of production when the agent is efficient, W_L^* , is greater than when he is inefficient, namely W_H^* :

$$W_L^* = \overbrace{S(q_L^*) - \theta_L q_L^*}^{q_L^* \text{ maximizes } S(q) - \theta_L q} \geq \underbrace{S(q_H^*) - \theta_L q_H^*}_{\theta_L < \theta_H} \geq S(q_H^*) - \theta_H q_H^* = W_H^*.$$

For trade to be always carried out, it is thus enough that production be socially valuable for the least efficient type, *i.e.*, the following condition must be satisfied

$$W_H^* = S(q_H^*) - \theta_H q_H^* \geq 0,$$

a hypothesis that we will maintain throughout this section.

15.9 For a successful delegation of the production, the principal must offer the agent a utility level that is at least as high as the utility level that the agent obtains from outside opportunity. We refer to these constraints as the agent's individual rationality constraints or participation constraints.

If we normalize to zero the agent's outside opportunity utility level (*i.e.*, his status quo utility level), these conditions are written as

$$t_L - \theta_L q_L \geq 0 \text{ and } t_H - \theta_H q_H \geq 0.$$

Obviously, the first-best contract menu $\{(q_i^*, t_i^*)\}_{i=H,L}$ satisfies these conditions, if we let $t_i^* = \theta_i q_i^*$.

15.10 To implement the first-best production levels q_L^* and q_H^* , the principal can make the following take-it-or-leave-it offers to the agent: If $\theta = \theta_L$ (resp. θ_H), the principal offers the transfer t_L^* (resp. t_H^*) for the production level q_L^* (resp. q_H^*) with $t_L^* = \theta_L q_L^*$ (resp. $t_H^* = \theta_H q_H^*$).

Whatever his type, agent accepts the offer and makes zero utility. The complete information optimal contracts are thus (t_L^*, q_L^*) if $\theta = \theta_L$ and (t_H^*, q_H^*) if $\theta = \theta_H$.

15.11 Under complete information delegation is costless for the principal, who achieves the same utility level that he could get if he was carrying out the task himself (with the same cost function as the agent).

15.12 In Figure 15.2, we draw the indifference curves of a θ_L -agent (heavy curves) and of a θ_H -agent (light curves) in the (q, t) space. The isoutility curves of both types correspond to increasing levels of utility when one moves in the northwest direction. These indifference curves are straight lines with a slope θ corresponding to the agent's type.




Figure 15.2: Indifference curves of both types

Since $\theta_H > \theta_L$, the isoutility curves of the inefficient agent θ_H have a greater slope than those of the efficient agent. Thus, the isoutility curves for different types cross only once. This property is called the *single-crossing* or *Spence-Mirrlees* property.

- 15.13 The complete information optimal contract is finally represented in Figure 15.3 by the pair of points (A^*, B^*) . For each of those two points, the strictly concave indifference curve of the principal is tangent to the zero rent isoutility curve of the corresponding type.

Note that the isoutility curves of the principal correspond to increasing levels of utility when one moves in the southeast direction. Thus the principal reaches a higher profit when dealing with the efficient type.




Figure 15.3: First-best contracts

We denote by V_H^* (resp. V_L^*) the principal's level of utility when he faces the θ_H - (resp. θ_L -) type. Because the principal has all the bargaining power (complete information) in designing the contract, we have $V_H^* = W_H^*$ (resp. $V_L^* = W_L^*$) under complete information.

15.14 In Figure 15.3, the payment t_L^* is greater than t_H^* , but we note that t_L^* can be greater or smaller than t_H^* depending on the curvature of the function $S(\cdot)$, as it can be easily seen graphically.

15.1.2 Incomplete information

15.15 Suppose that the marginal cost θ is the agent's private information.

15.16 Consider the case where the principal offers the menu of contracts $\{(t_L^*, q_L^*), (t_H^*, q_H^*)\}$ hoping that an agent with type θ_L will select (t_L^*, q_L^*) and an agent with type θ_H will select instead (t_H^*, q_H^*) .

From Figure 15.3, we see that B^* is preferred to A^* by both types of agents:

- The θ_L -agent's isoutility curve that passes through B^* corresponds to a positive utility level instead of a zero utility level at A^* .
- The θ_H -agent's isoutility curve that passes through A^* corresponds to a negative utility level, which is less than the zero utility level this type gets by choosing B^* .

Offering the menu (A^*, B^*) fails to have the agents self-selecting properly within this menu. The efficient type mimics the inefficient one and selects also contract B^* . The complete information optimal contracts can no longer be implemented under asymmetric information.

15.17 Definition: A menu of contracts $\{(t_L, q_L), (t_H, q_H)\}$ is incentive compatible when (t_L, q_L) is weakly preferred to (t_H, q_H) by the type- θ_L agent and (t_H, q_H) is weakly preferred to (t_L, q_L) by the type- θ_H agent.

Mathematically,

$$t_L - \theta_L q_L \geq t_H - \theta_L q_H, \quad (\text{IC}_L)$$

$$t_H - \theta_H q_H \geq t_L - \theta_H q_L. \quad (\text{IC}_H)$$

15.18 If a menu of contracts $\{(t_L, q_L), (t_H, q_H)\}$ is incentive compatible, then $q_L \geq q_H$, which is called the monotonicity constraint. Indeed,

$$\overbrace{\theta_L(q_H - q_L) \geq t_H - t_L}^{\text{By Equation (IC}_L\text{)}} \geq \underbrace{\theta_H(q_H - q_L)}_{\text{By Equation (IC}_H\text{)}},$$

and hence $q_H - q_L \leq 0$.

15.19 Definition: A menu of contracts $\{(t_L, q_L), (t_H, q_H)\}$ is individually rational if

$$t_L - \theta_L q_L \geq 0, \quad (\text{IR}_L)$$

$$t_H - \theta_H q_H \geq 0. \quad (\text{IR}_H)$$

15.20 Under complete information, the principal is able to maintain all types of agents at their zero status quo utility level. Their respective utility levels U_L^* and U_H^* at the first-best contracts satisfy

$$U_L^* = t_L^* - \theta_L q_L^* = 0 \text{ and } U_H^* = t_H^* - \theta_H q_H^* = 0.$$

Generally this will not be possible anymore under incomplete information, at least when the principal wants both types of agents to be active.

Let $U_L = t_L - \theta_L q_L$ and $U_H = t_H - \theta_H q_H$ denote the respective information rent (the utility in excess of the reservation utility) of each type.

15.21 The principal's problem is to solve

$$\begin{aligned} & \underset{(q_L, t_L), (q_H, t_H)}{\text{maximize}} && \beta(S(q_L) - t_L) + (1 - \beta)(S(q_H) - t_H) \\ & \text{subject to} && \text{Equations (IC}_L\text{)} - (\text{IR}_H\text{)}. \end{aligned}$$

15.22 Since $U_L = t_L - \theta_L q_L$ and $U_H = t_H - \theta_H q_H$, the principal's objective function can then be rewritten as

$$\underbrace{\beta(S(q_L) - \theta_L q_L) + (1 - \beta)(S(q_H) - \theta_H q_H)}_{\text{Expected social value/allocative efficiency}} - \underbrace{(\beta U_L + (1 - \beta)U_H)}_{\text{Expected information rent}}.$$

The incentive constraints and individual rationality constraints are rewritten as

$$U_L \geq U_H + \Delta\theta q_H, \quad (\text{IC}'_L)$$

$$U_H \geq U_L - \Delta\theta q_L, \quad (\text{IC}'_H)$$

$$U_L \geq 0, \quad (\text{IR}'_L)$$

$$U_H \geq 0. \quad (\text{IR}'_H)$$

15.23 Optimal contract without shutdown: Under asymmetric information, the optimal menu of contracts entails:

- No output distortion for the efficient type with respect to the first-best, $q_L^{\text{SB}} = q_L^*$. A downward output distortion for the inefficient type, $q_H^{\text{SB}} < q_H^*$ with

$$S'(q_H^{\text{SB}}) = \theta_H + \frac{\beta}{1 - \beta} \Delta\theta.$$

Here we assume that the equation above has positive solution. Otherwise q_H^{SB} should be set at zero, and we are in the special case of a contract with shutdown.

- Only the efficient type gets a positive information rent given by

$$U_L^{\text{SB}} = \Delta\theta q_H^{\text{SB}}.$$

- The second-best transfers are respectively given by

$$t_L^{\text{SB}} = \theta_L q_L^* + \Delta\theta q_H^{\text{SB}} \text{ and } t_H^{\text{SB}} = \theta_H q_H^{\text{SB}}.$$

15.24 Proof:

- (1) Step 1: If $q_H > 0$, then the constraint (IR'_L) is always strictly satisfied due to constraints (IC'_L) and (IR'_H) .
- (2) Step 2: The constraint (IR'_H) is binding at the optimal, i.e., $U_H = 0$. Suppose that $U_H = \epsilon > 0$. Then the principal can decrease U_H by ϵ and consequently also U_L by ϵ and gain ϵ .
- (3) Step 3: The constraint (IC'_L) is binding at the optimal, i.e., $U_L = \Delta\theta q_H$. Suppose that $U_L - \Delta\theta q_H = \epsilon > 0$. Then the principal can decrease U_L by ϵ and gain $\beta\epsilon$.

(4) Step 4: By Steps 2 and 3, we obtain a reduced program

$$\underset{q_L, q_H}{\text{maximize}} \quad \beta(S(q_L) - \theta_L q_L) + (1 - \beta)(S(q_H) - \theta_H q_H) - \beta\Delta\theta q_H.$$

Compared with the full information setting, asymmetric information alters the principal's optimization simply by the subtraction of the expected rent that has to be given up to the efficient type. The inefficient type gets no rent, but the efficient type θ_L gets the information rent that he could obtain by mimicking the inefficient type θ_H . This rent depends only on the level of production requested from this inefficient type.

(5) Step 5: The first order condition on q_L implies

$$S'(q_L^{\text{SB}}) = \theta_L, \text{ that is, } q_L^{\text{SB}} = q_L^*.$$

Hence, there is no distortion away from the first-best for the efficient type's output.

The first order condition on q_H implies

$$(1 - \beta)(S'(q_H^{\text{SB}}) - \theta_H) = \beta\Delta\theta.$$

This equation expresses the important trade-off between efficiency and rent extraction which arises under asymmetric information. The expected marginal efficiency gain (resp. cost) and the expected marginal cost (resp. gain) of the rent brought about by an infinitesimal increase (resp. decrease) of the inefficient type's output are equated. Thus, the principal is neither willing to increase nor to decrease the inefficient agent's output.

(6) Step 6: We have the following inequality

$$q_L^{\text{SB}} = q_L^* > \underbrace{q_H^*}_{S'' < 0} > q_H^{\text{SB}},$$

and hence

$$U_H = 0 > \Delta\theta q_H - \Delta\theta q_L = U_L - \Delta\theta q_L.$$

That is, the constraint (IC'_H) is strictly satisfied.

15.25 Starting from the complete information optimal contract (A^*, B^*) that is not incentive compatible, we can construct an incentive compatible contract (B^*, C) with the same production levels by giving a higher transfer to the agent producing q_L^* (Figure 15.4).




Figure 15.4: Rent needed to implement the first-best outputs

The contract C is on the θ_L -agent's indifference curve passing through B^* . Hence, the θ_L -agent is now indifferent between B^* and C . (B^*, C) becomes an incentive-compatible menu of contracts. The rent that is given up to the θ_L -agent is now $\Delta\theta q_H^*$.

15.26 Rather than insisting on the first-best production level q_H^* for an inefficient type, the principal prefers to slightly decrease q_H by an amount dq .

- By doing so, expected efficiency is just diminished by a second-order term $\frac{1}{2}|S''(q_H^*)|(dq)^2$ since q_H^* is the first-best output that maximizes efficiency when the agent is inefficient.
- Instead, the information rent left to the efficient type diminishes to the first-order $\Delta\theta dq$.

Of course, the principal stops reducing the inefficient type's output when a further decrease would have a greater efficiency cost than the gain in reducing the information rent it would bring about. The optimal trade-off finally occurs at (A^{SB}, B^{SB}) as shown in Figure 15.5.




Figure 15.5: Second-best contracts

15.27 Optimal contract with shutdown.

- (1) If the equation $S'(q_H^{SB}) = \theta_H + \frac{\beta}{1-\beta}\Delta\theta$ has no positive solution, q_H^{SB} should be set at zero.
- (2) Then B^{SB} coincides with O and A^{SB} with A^* in Figure 15.5.
- (3) No rent is given up to the θ_L -agent by the unique non-null contract (t_L^*, q_L^*) offered and selected only by agent θ_L .
- (4) The shutdown of the agent occurs when $\theta = \theta_L$.

With such a contract, a significant inefficiency emerges because the inefficient type does not produce. The benefit of such a contract is that no rent is given up to the efficient type.

15.28 More generally, such a shutdown contract is optimal when

$$\beta(S(q_L^*) - \theta_L q_L^*) \geq \beta(S(q_L^{SB}) - \theta_L q_L^{SB} - \Delta\theta q_H^{SB}) + (1 - \beta)(S(q_H^{SB}) - \theta_H q_H^{SB})$$

or, noting that $q_L^* = q_L^{SB}$, when

$$\beta\Delta\theta q_H^{SB} \geq (1 - \beta)(S(q_H^{SB}) - \theta_H q_H^{SB}).$$

The left-hand side represents the expected cost of the efficient type's rent due to the presence of the inefficient one when the latter produces a positive amount q_H^{SB} . The right-hand side represents the expected benefit from transacting with the inefficient type at the second-best level of output. Thus, shutdown for the inefficient type is optimal when this expected benefit is lower than the expected cost.

15.29 When Inada condition $S'(0) = +\infty$ is satisfied and $\lim_{q \rightarrow 0} S'(q)q = 0$, the shutdown is never desirable.

- (1) q_H^{SB} defined by $S'(q_H^{SB}) = \theta_H + \frac{\beta}{1-\beta}\Delta\theta$ is necessarily strictly positive since $S'(0) = +\infty$.
- (2)

$$S(q_H^{SB}) - (\theta_H + \frac{\beta}{1-\beta}\Delta\theta)q_H^{SB} = S(q_H^{SB}) - S'(q_H^{SB})q_H^{SB}$$

is strictly positive since $S(q) - S'(q)q$ is strictly increasing with q and is equal to zero for $q = 0$. Hence,

$$\beta\Delta\theta q_H^{SB} < (1 - \beta)(S(q_H^{SB}) - \theta_H q_H^{SB})$$

and the shutdown of the least efficient type does not occur.

15.1.3 The revelation principle

15.30 In the above analysis, we have restricted the principal to offer a menu of contracts, one for each possible type.

- First, one may wonder if a better outcome could be achieved with a more complex contract allowing the agent possibly to choose among more options.
- Second, one may also wonder whether some sort of communication device between the agent and the principal could be used to transmit information to the principal so that the latter can recommend outputs and payments as a function of transmitted information.

The revelation principle ensures that there is no loss of generality in restricting the principal to offer simple menus having at most as many options as the cardinality of the type space. Those simple menus are actually examples of direct revelation mechanisms for which we now give a couple of definitions.

15.31 Definition: A direct mechanism is a mapping $g: \Theta \rightarrow \mathcal{A}$.

The principal commits to offer the transfer $t(\theta)$ and the production level $q(\theta)$ if the agent announces the value θ for any $\theta \in \Theta$.

15.32 Definition: A direct mechanism g is truthful if it is incentive compatible for the agent to announce his true type for any type, *i.e.*, if the direct mechanism satisfies the following incentive compatibility constraints:

$$t(\theta_L) - \theta_L q(\theta_L) \geq t(\theta_H) - \theta_L q(\theta_H) \text{ and } t(\theta_H) - \theta_H q(\theta_H) \geq t(\theta_L) - \theta_H q(\theta_L).$$

15.33 A more general mechanism can be obtained when communication between the principal and the agent is more complex than simply having the agent report his type to the principal. Let \mathcal{M} be the message space offered to the agent by a more general mechanism. This message space can be very complex. Conditionally, on a given message m received from the agent, the principal requests a production level $\tilde{q}(m)$ and provides a corresponding payment $\tilde{t}(m)$.

A mechanism is a message space \mathcal{M} and a mapping $\tilde{g}: \mathcal{M} \rightarrow \mathcal{A}$.

15.34 When facing such a mechanism, the agent with type θ chooses a best message $m^*(\theta)$ that is implicitly defined as

$$\tilde{t}(m^*(\theta)) - \theta \tilde{q}(m^*(\theta)) \geq \tilde{t}(\tilde{m}) - \theta \tilde{q}(\tilde{m}) \text{ for all } \tilde{m} \text{ in } \mathcal{M}.$$

The mechanism (\mathcal{M}, \tilde{g}) induces therefore an allocation rule $a(\theta) = (\tilde{q}(m^*(\theta)), \tilde{t}(m^*(\theta)))$ mapping the set of types Θ into the set of allocations \mathcal{A} .

15.35 Revelation principle: Any allocation rule $a(\theta)$ obtained with a mechanism (\mathcal{M}, \tilde{g}) can also be implemented with a truthful direct mechanism.

Proof. Let $g = \tilde{g} \circ m^*$. □

15.1.4 General utility function for the agent

15.36 Consider a general cost function $C(q, \theta)$ with the assumption

$$C_q > 0, C_\theta > 0, C_{qq} > 0, C_{q\theta} > 0.$$

15.37 The generalization of the Spence-Mirrlees property used so far is now

$$C_{q\theta} > 0.$$

This condition still ensures that the different types of the agent have indifference curves which cross each other at most once. This condition simply says that a more efficient type is also efficient at the margin.

15.38 Optimal contract entails:

- No output distortion with respect to the first-best outcome for the efficient type, $q_L^{\text{SB}} = q_L^*$ with

$$S'(q_L^*) = C_q(q_L^*, \theta_L).$$

A downward output distortion for the inefficient type, $q_H^{\text{SB}} < q_H^*$ with

$$S'(q_H^*) = C_q(q_H^*, \theta_H)$$

and

$$S'(q_H^{\text{SB}}) = C_q(q_H^{\text{SB}}, \theta_H) + \frac{\beta}{1-\beta} \Phi'(q_H^{\text{SB}}),$$

where $\Phi(q) = C(q, \theta_H) - C(q, \theta_L)$ with $\Phi' > 0$ and $\Phi'' > 0$ from the assumptions made on C .

- Only the efficient type gets a positive information rent given by $U_L^{\text{SB}} = \Phi(q_H^{\text{SB}})$.
- The second-best transfers are respectively given by $t_L^{\text{SB}} = C(q_L^*, \theta_L) + \Phi(q_H^{\text{SB}})$ and $t_H^{\text{SB}} = C(q_H^{\text{SB}}, \theta_H)$.

15.39 We still have

$$q_L^{\text{SB}} = q_L^* > q_H^* > q_H^{\text{SB}}.$$

- $S'(q_L^*) = C_q(q_L^*, \theta_L) < C_q(q_L^*, \theta_H)$ because $C_{q\theta} > 0$. Hence, using the fact that $S(q) - C(q, \theta_H)$ is concave in q and maximum for q_H^* , we have $q_L^* > q_H^*$.
- $\Phi' > 0$ implies that $S'(q_H^{\text{SB}}) > C_q(q_H^{\text{SB}}, \theta_H)$. Thus, $q_H^{\text{SB}} < q_H^*$.

15.1.5 Nonresponsiveness

15.40 Assume that the principal's value function S depends directly on θ and is written as $S(q, \theta)$.

We still assume that $S_q > 0$ and $S_{qq} < 0$. Furthermore, we assume that $S_{q\theta} > 1$. This assumption simply means that the marginal gross value of trade for the principal increases quickly with the agent's type.

15.41 The first-best contracts are the solutions to the following problem.

$$\begin{aligned} & \underset{(q_i, t_i) \in \mathcal{A}}{\text{maximize}} && S(q, \theta_i) - t \\ & \text{subject to} && t - c(q, \theta_i) \geq 0. \end{aligned}$$

Thus, the first-best productions are now defined by

$$S_q(q_L^*, \theta_L) = \theta_L \text{ and } S_q(q_H^*, \theta_H) = \theta_H.$$

15.42 Since $S_q(q^*(\theta), \theta) = \theta$, by implicit function theorem, we have

$$\frac{dq^*}{d\theta}(\theta) = -\frac{S_{q\theta} - 1}{S_{qq}} > 0.$$

Thus, $q_L^* = q^*(\theta_L) < q^*(\theta_H) = q_H^*$.

However, the monotonicity condition implies that $q_L^* \geq q_H^*$.

That is, there exists a strong conflict between the principal's desire to have the θ_H -type produce more than the θ_L -agent for pure efficiency reasons and the monotonicity condition imposed by asymmetric information. This is what Guesenerie and Laffont (1984) call a phenomenon of nonresponsiveness.

15.43 Nonresponsiveness makes screening of types quite difficult:

- (1) the second-best optimum induces screening only when $q_L^{SB} = q_L^*$ and q_H^{SB} defined by

$$S_q(q_H^{SB}, \theta_H) = \theta_H + \frac{\beta}{1 - \beta} \Delta\theta$$

satisfy the monotonicity condition $q_L^{SB} \geq q_H^{SB}$.

- (2) When β is small enough, q_H^{SB} defined above is close to the first-best outcome q_H^* .
 (3) Thus, we have $q_H^{SB} > q_L^{SB}$, and monotonicity condition is violated.

Nonresponsiveness forces the principal to use a pooling allocation.

15.2 Pricing a single indivisible good

15.44 Literature:

15.45 A seller seeks to sell a single indivisible good. The seller herself does not attach any value to the good. Her objective is to maximize the expected revenue from selling the good, that is, she is risk-neutral.

15.46 There is just one potential buyer. The buyer's von Neumann-Morgenstern utility if he purchases the good and pays a monetary transfer t to the seller is $\theta - t$.

Here θ is the number that we can interpret as the buyer's valuation of the good. We shall refer below to θ as the buyer's type.

We have assumed that the buyer's utility is additively separable—the sum of the utility derived from the good and the disutility resulting from the money payment.

We have also assumed that the buyer is risk-neutral with respect to money, that is, his utility is linear in money.

Utility functions that satisfy additive separability of utility and linearity of utility in money are referred to as “quasi-linear” utility functions.

15.47 We assume that the value of θ is known to the buyer, but it is not known to the seller. The seller has a subjective probability distribution over possible values of θ . To be more precise, θ is assumed to be a random variable with the cumulative distribution function F and the strictly positive density function f on $[\underline{\theta}, \bar{\theta}]$.

15.48 Our interest is in procedures for selling the good which the seller should adopt to maximize expected profits.

One obvious way would be to pick a price p and to say to the buyer that he can have the good if and only if he is willing to pay p . In this procedure, the probability that the buyer's value is above p , and hence he accepts a price offer p , is $1 - F(p)$. Seller's expected revenue is therefore $p \cdot (1 - F(p))$, and the optimal strategy for the seller is to pick some price that maximizes $p \cdot (1 - F(p))$.

15.49 Is picking a price p really the best the seller can do? What else could the seller do? The seller could, for example, negotiate with the buyer. The seller could offer the buyer a lottery where in return for higher or lower payments the buyer could be given a larger or smaller chance of getting the object.

Our objective is thus to study the optimality problem in which the seller's choices are a lottery (design) and a strategy in the game associated with the lottery and in which the seller's objective function is expected revenue, with the constraint that the buyer will choose an expected utility maximizing strategy and several other constraints introduced below.

15.50 A direct mechanism consists of functions q and t where

$$q: [\underline{\theta}, \bar{\theta}] \rightarrow [0, 1], \quad t: [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}.$$

The interpretation is that in a direct mechanism the buyer is asked to report θ . The seller commits to transferring the good to the buyer with the probability $q(\theta)$ if the buyer reports that his type is θ , and the buyer has to pay the seller $t(\theta)$ if he reports that his type is θ .

15.51 Given a direct mechanism (q, t) , we define the buyer's expected utility $u(\theta)$ conditional on his type being θ by $u(\theta) = \theta q(\theta) - t(\theta)$.

15.52 A direct mechanism (q, t) is incentive compatible (IC) if truth telling is optimal for every type $\theta \in [\underline{\theta}, \bar{\theta}]$, that is, if

$$u(\theta) \geq \theta q(\theta') - t(\theta') \text{ for all } \theta, \theta' \in [\underline{\theta}, \bar{\theta}].$$

15.53 Lemma: If a direct mechanism (q, t) is IC, then q is increasing in θ .

Proof. Straightforward. □

15.54 Lemma: If a direct mechanism (q, t) is IC, then

- u is increasing,
- u is convex, and hence differentiable except in at most countably many points,
- for all θ for which it is differentiable, it satisfies $u'(\theta) = q(\theta)$.

Proof. (1) IC means that for all θ we have

$$u(\theta) = \max_{\theta' \in [\underline{\theta}, \bar{\theta}]} \theta q(\theta') - t(\theta').$$

(2) Given any value of θ' , $\theta q(\theta') - t(\theta')$ is an increasing and affine (and hence convex) function of θ .

(3) The maximum of increasing functions is increasing, and the maximum of convex functions is convex. Therefore, u is increasing and convex.

(4) Convex functions are not differentiable in at most countably many points.

(5) Consider any θ for which u is differentiable. By IC we have

$$\lim_{\epsilon \downarrow 0} \frac{u(\theta + \epsilon) - u(\theta)}{\epsilon} \geq \lim_{\epsilon \downarrow 0} \frac{[(\theta + \epsilon)q(\theta) - t(\theta)] - [\theta q(\theta) - t(\theta)]}{\epsilon} = q(\theta).$$

Similarly,

$$\lim_{\epsilon \downarrow 0} \frac{u(\theta) - u(\theta - \epsilon)}{\epsilon} \leq \lim_{\epsilon \downarrow 0} \frac{[\theta q(\theta) - t(\theta)] - [(\theta - \epsilon)q(\theta) - t(\theta)]}{\epsilon} = q(\theta).$$

(6) Putting the two inequalities together, we obtain $u'(\theta) = q(\theta)$ whenever u is differentiable. □

15.55 Proposition (Payoff equivalence): Consider an IC direct mechanism (q, t) . Then for all $\theta \in [\underline{\theta}, \bar{\theta}]$ we have

$$u(\theta) = u(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(x) \, dx.$$

Proof. (1) u is convex, then u is absolutely continuous; see Corollary 17 in Royden (1988, Chapter 5).

(2) By Theorem 10 in Royden (1988, Chapter 5), it is the integral of its derivative. □

15.56 Proposition (Revenue equivalence): Consider an IC direct mechanism (q, t) . Then for all $\theta \in [\underline{\theta}, \bar{\theta}]$ we have

$$t(\theta) = t(\underline{\theta}) + [\theta q(\theta) - \underline{\theta} q(\underline{\theta})] - \int_{\underline{\theta}}^{\theta} q(x) dx.$$

15.57 Proposition: A direct mechanism (q, t) is IC if and only if

- q is increasing.
- For every $\theta \in [\underline{\theta}, \bar{\theta}]$ we have

$$t(\theta) = t(\underline{\theta}) + [\theta q(\theta) - \underline{\theta} q(\underline{\theta})] - \int_{\underline{\theta}}^{\theta} q(x) dx.$$

Proof. (1) The “only if” part is due to Lemma 15.53 and Proposition 15.56.

(2) For any θ and θ' , we have

$$[\theta q(\theta) - t(\theta)] - [\theta' q(\theta') - t(\theta')] = \int_{\theta'}^{\theta} q(x) dx - q(\theta')(\theta - \theta') \geq 0.$$

□

15.58 A direct mechanism is individually rational (IR) if the buyer, conditional on his type, is willing to participate, that is, if

$$u(\theta) \geq 0 \text{ for all } \theta \in [\underline{\theta}, \bar{\theta}].$$

15.59 Proposition: An IC direct mechanism (q, t) is IR if and only if $u(\underline{\theta}) \geq 0$, or equivalently $\underline{\theta} q(\underline{\theta}) \geq t(\underline{\theta})$.

Proof. Due to Lemma 15.54. □

15.60 Proposition: If an IC and IR direct mechanism (q, t) maximizes the seller's expected revenue, then

$$\underline{\theta} q(\underline{\theta}) = t(\underline{\theta}).$$

Proof. If $\underline{\theta} q(\underline{\theta}) > t(\underline{\theta})$, then the seller could increase expected revenue by choosing a direct mechanism with the same q , but a higher $t(\underline{\theta})$. □

15.61 Remark: If the buyer has only two possible types, the revenue equivalence principle does not necessarily hold.

Let θ_H and θ_L be buyer's types. Firstly, IR for the low type should be binding, that is, $\theta_L q(\theta_L) = t(\theta_L)$.

Next, IC for the high type should be binding. Otherwise, we have

$$\theta_H q(\theta_H) - t(\theta_H) > \theta_H q(\theta_L) - t(\theta_L) \geq \theta_L q(\theta_L) - t(\theta_L) = 0.$$

Then the seller can increase $t(\theta_H)$ without breaking IC and IR.

Thus, we have

$$t(\theta_H) = t(\theta_L) + [\theta_H q(\theta_H) - \theta_H q(\theta_L)].$$

Since $\int_{\theta_L}^{\theta_H} q(x) dx$ may not be zero, the revenue equivalence principle does not necessarily hold.

In the discrete case, the distance between two types allows the variety of the transfers.

For example, let $\theta_L = 1$, $\theta_H = 2$, $\text{Prob}(\theta_H) = \text{Prob}(\theta_L) = \frac{1}{2}$, $q(\theta_L) = 0$, $q(\theta_H) = 1$, $t(\theta_L) = 0$, $t(\theta_H) = 2$.

15.62 Let \mathcal{M} be the set of all increasing functions $q: [\underline{\theta}, \bar{\theta}] \rightarrow [0, 1]$.

Lemma:

- \mathcal{M} is compact and convex.
- Let X be a compact, convex subset of a vector space, $f: X \rightarrow \mathbb{R}$ be a continuous linear function. Then the set E of extreme points of X is nonempty, and there exists an $e \in E$ such that $f(e) \geq f(x)$.
- A function $q \in \mathcal{M}$ is an extreme point of \mathcal{M} if and only if $q(\theta) \in \{0, 1\}$ for almost all θ .

15.63 The lemma above implies that the seller can restrict her attention to nonstochastic mechanisms. But a nonstochastic mechanism is monotonic if and only if there is some $p^* \in [\underline{\theta}, \bar{\theta}]$ such that $q(\theta) = 0$ if $\theta < p^*$ and $q(\theta) = 1$ if $\theta > p^*$.

Proposition: The following direct mechanism maximizes the seller's expected revenues among all IC and IR direct mechanisms:

$$p^* \in \arg \max_{p \in [\underline{\theta}, \bar{\theta}]} p(1 - F(p)), \quad q(\theta) = \begin{cases} 0, & \text{if } \theta < p^* \\ 1, & \text{if } \theta > p^* \end{cases}, \quad t(\theta) = \begin{cases} 0, & \text{if } \theta < p^* \\ p^*, & \text{if } \theta > p^* \end{cases}.$$

15.64 It may seem that we have gone to considerable lengths to derive a disappointing result, namely, a result that does not offer the seller any more sophisticated selling mechanisms than we are familiar with from elementary microeconomics.

15.3 Nonlinear pricing

15.65 Literature:

15.66 Consider a transaction between a seller and a buyer, where the seller seeks to sell a good to the buyer.

The buyer's utility if she purchases q units of the good and pays a monetary transfer t to the seller is

$$u_b(q, t, \theta) = \theta v(q) - t,$$

where $v(0) = 0$, $v'(q) > 0$ and $v''(q) < 0$ for all q . The characteristics $\theta > 0$ is a number that we can interpret as the buyer's valuation of the good.

Assuming that the seller's unit production costs are given by $c > 0$, her profit from selling q units against a sum of money t is given by

$$u_s(q, t) = t - cq.$$

15.67 We assume that the value of θ is known to the buyer, but it is not known to the seller. This seems plausible in many contexts. Buyers often know better than sellers how well some particular product meets their preferences. We shall refer below to θ as the buyer's type. While the seller does not know the buyer's type, she does have a subjective probability distribution over possible values of θ .

15.68 The question of interest here is, what is the best, that is, the profit maximizing, contract (q, t) that the seller will be able to induce the buyer to choose?

Let \mathcal{A} be the set of all feasible contracts, that is,

$$\mathcal{A} = \{(q, t) \mid q \geq 0, t \in \mathbb{R}\}.$$

15.3.1 The model with two types

15.69 We suppose that set of possible characteristics

$$\Theta = \{\theta_H, \theta_L\}, \text{ with } \theta_H > \theta_L.$$

The first-best contract: perfect discrimination

15.70 If the seller can observe the types θ of the buyer, she can then treat each type of buyer separately and offer her a type-specific contract, that is, (q_i, t_i) for type θ_i ($i = H, L$).

15.71 The seller will try to maximize her profits subject to inducing the buyer to accept the proposed contract. Assume the buyer obtains 0 if she does not take the seller's offer. So the seller will solve the following problem:

$$\begin{aligned} & \underset{(q_i, t_i) \in \mathcal{A}}{\text{maximize}} && t_i - cq_i \\ & \text{subject to} && \theta_i v(q_i) - t_i \geq 0. \end{aligned}$$

The solution to this program will be the first-best contract menu $\{(q_i^*, t_i^*)\}_{i=H,L}$ such that

$$\theta_i v'(q_i^*) = c \text{ and } \theta_i v(q_i^*) = t_i^* \text{ for } i = H, L.$$

15.72 Figure 15.6 illustrates the two first-best contracts when $v(q) = \sqrt{q}$, $c = \frac{2}{3}$, $\theta_L = 1$ and $\theta_H = 2$. The two curves shown are the indifference curves corresponding to zero utility for the two types of the buyer. The lines tangent to them are isoprofit curves, with equation $t = cq + K$, where $K \in \mathbb{R}$. Note that the utility of the buyer increases when going southeast, while the profit of the seller increases when going northwest.

Since $\theta_H > \theta_L$ and v' is decreasing, we have

$$q_H^* > q_L^*.$$




Figure 15.6: Illustration of a first-best contract, where $v(q) = \sqrt{q}$, $c = \frac{2}{3}$, $\theta_L = 1$ and $\theta_H = 2$.

The second-best: optimal nonlinear pricing

15.73 We now consider the case where the seller cannot observe directly the buyer's type. We assume that

$$\text{Prob}(\theta = \theta_L) = \beta \text{ and } \text{Prob}(\theta = \theta_H) = 1 - \beta.$$

If the seller proposes the first-best contracts (q_i^*, t_i^*) , the θ_H -type buyer will choose (q_L^*, t_L^*) , the contract designed for the θ_L -type:

$$\begin{aligned} u(q_L^*, t_L^*, \theta_H) &= \theta_H v(q_L^*) - t_L^* = (\theta_H - \theta_L)v(q_L^*) + [\theta_L v(q_L^*) - t_L^*] \\ &= (\theta_H - \theta_L)v(q_L^*) \geq 0 = u(q_H^*, t_H^*, \theta_H), \end{aligned}$$

where the inequality holds strictly whenever $q_L^* > 0$. Thus, the two types cannot be treated separately any more. Both will choose (q_L^*, t_L^*) ; see Figure 15.7.




Figure 15.7: The θ_H -type buyer will choose (q_L^*, t_L^*) , where $v(q) = \sqrt{q}$, $c = \frac{2}{3}$, $\theta_L = 1$ and $\theta_H = 2$.

15.74 Our interest is in the best pair of contracts (the second-best optimum). At first sight, this looks like a hard problem, as the contract set that the seller can use is potentially large. However, a simple, yet crucial result enables us to get a handle on this optimization problem. The result says that we can restrict our attention to a small set of contracts.

15.75 A menu of contracts $\{(q_L, t_L), (q_H, t_H)\}$ is incentive compatible when (q_L, t_L) is weakly preferred to (q_H, t_H) by type- θ_L buyer, and (q_H, t_H) is weakly preferred to (q_L, t_L) by type- θ_H buyer; that is,

$$\theta_L v(q_L) - t_L \geq \theta_L v(q_H) - t_H, \quad (\text{IC}_L)$$

$$\theta_H v(q_H) - t_H \geq \theta_H v(q_L) - t_L, \quad (\text{IC}_H)$$

15.76 A menu of contracts $\{(q_L, t_L), (q_H, t_H)\}$ is individual rational if

$$\theta_L v(q_L) - t_L \geq 0, \quad (\text{IR}_L)$$

$$\theta_H v(q_H) - t_H \geq 0, \quad (\text{IR}_H)$$

15.77 The problem of the seller is therefore to solve

$$\max_{(q_L, t_L), (q_H, t_H)} \beta[t_L - cq_L] + (1 - \beta)[t_H - cq_H],$$

subject to

$$\theta_L v(q_L) - t_L \geq \theta_L v(q_H) - t_H, \quad (\text{IC}_L)$$

$$\theta_H v(q_H) - t_H \geq \theta_H v(q_L) - t_L, \quad (\text{IC}_H)$$

$$\theta_L v(q_L) - t_L \geq 0, \quad (\text{IR}_L)$$

$$\theta_H v(q_H) - t_H \geq 0, \quad (\text{IR}_H)$$

15.78 Step 1: Equation (IR_H) is not bind at the optimum. Indeed Equation (IR_H) will be satisfied automatically because of Equations (IR_L) and (IC_H):

$$\theta_H v(q_H) - t_H \geq \theta_H v(q_L) - t_L \geq \theta_L v(q_L) - t_L \geq 0. \quad (15.1)$$

15.79 Step 2: Equation (IR_L) is bind at the optimum, so $t_L = \theta_L v(q_L)$. If Equation (IR_L) was not bind, so would be Equation (IR_H) by Equation (15.1). We then could increase t_L and t_H by the same amount. This would increase the seller's profit without any effect on incentive compatibility.

15.80 Step 3: Equation (IC_H) is bind at the optimum, so $t_H - t_L = \theta_H[v(q_H) - v(q_L)]$. Assume not, then

$$\theta_H v(q_H) - t_H > \theta_H v(q_L) - t_L \geq \theta_L v(q_L) - t_L \geq 0.$$

We can therefore increase t_H without breaking incentive compatibility or individual rationality. This obviously increases the seller's profit, and hence the original contract cannot be optimal.

15.81 Step 4: $q_H \geq q_L$ at the optimum. To see it, we add Equations (IC_L) and (IC_H) and get

$$\theta_H[v(q_H) - v(q_L)] \geq \theta_L[v(q_H) - v(q_L)].$$

Since $\theta_H > \theta_L$ and v is increasing, we have $q_H \geq q_L$.

15.82 Step 5: Equation (IC_L) is not bind at the optimum. This because

$$\begin{aligned} \theta_L v(q_L) - t_L &= \theta_H v(q_L) - t_L - (\theta_H - \theta_L)v(q_L) \\ &= \theta_H v(q_H) - t_H - (\theta_H - \theta_L)v(q_L) && \text{By Equation (IC}_H\text{)} \\ &= \theta_L v(q_H) - t_H + (\theta_H - \theta_L)[v(q_H) - v(q_L)] \\ &\geq \theta_L v(q_H) - t_H && \text{By Step 4} \end{aligned}$$

15.83 Step 6: The seller's problem now reduces to

$$\max_{(q_L, t_L), (q_H, t_H)} \beta[t_L - cq_L] + (1 - \beta)[t_H - cq_H],$$

subject to

$$\begin{aligned} \theta_H v(q_H) - t_H &= \theta_H v(q_L) - t_L, \\ \theta_L v(q_L) - t_L &= 0. \end{aligned}$$

Substituting for the values of t_L and t_H in the seller's objective function, we obtain the following unconstrained optimization problem:

$$\max_{q_L, q_H} \beta[\theta_L v(q_L) - cq_L] + (1 - \beta)[\theta_H v(q_H) - (\theta_H - \theta_L)v(q_L) - cq_H].$$

The following first-order conditions characterize the unique interior solution $(q_L^{\text{SB}}, q_H^{\text{SB}})$ to the above program, if this solution exists:

$$\theta_H v'(q_H^{\text{SB}}) = c \text{ and } \theta_L v'(q_L^{\text{SB}}) = \frac{c}{1 - \frac{1-\beta}{\beta} \frac{\theta_H - \theta_L}{\theta_L}} > c.$$

15.84 Figure 15.8 illustrates the second-best contract when $v(q) = \sqrt{q}$, $c = \frac{2}{3}$, $\theta_L = 1$, $\theta_H = 2$ and $\beta = \frac{2}{3}$.




Figure 15.8: Illustration of a second-best contract, where $v(q) = \sqrt{q}$, $c = \frac{2}{3}$, $\theta_L = 1$, $\theta_H = 2$ and $\beta = \frac{2}{3}$.

15.85 $q_H^{\text{SB}} = q_H^*$ and $q_L^{\text{SB}} < q_L^*$.

The θ_L -type buyer has no surplus, while the other type has a strictly positive information rent:

$$\theta_H v(q_H^{\text{SB}}) - t_H^{\text{SB}} = \theta_H v(q_L^{\text{SB}}) - t_L^{\text{SB}} = (\theta_H - \theta_L) v(q_L^{\text{SB}}).$$

15.3.2 The model with a continuum of types

15.86 Suppose that θ is distributed according to F on $[\underline{\theta}, \bar{\theta}]$ (with density $f > 0$).

15.87 We assume that $\bar{\theta}v'(0) > c$. This means that the largest marginal willingness to pay that the buyer might possibly have is above the marginal cost of production. This assumption ensures that the seller and the buyer have an incentive to trade at least for the largest type of the buyer.

15.88 A final assumption is that $\lim_{q \rightarrow \infty} \bar{\theta}v'(q) < c$. This means that even the highest type's marginal willingness to pay falls below c as q gets large. This assumption ensures that the quantity that the seller supplies to the buyer is finite for all possible types of the buyer.

15.89 A direct mechanism consists of functions q and t where

$$q: \Theta \rightarrow \mathbb{R}_+, \quad t: \Theta \rightarrow \mathbb{R}.$$

The interpretation is that the buyer is asked to report θ , and the seller commits to selling quantity $q(\theta)$ to the buyer and the buyer commits to paying $t(\theta)$ to the seller.

15.90 Given a direct mechanism, we define the buyer's expected utility $u(\theta)$ conditional on her type being θ by

$$u(\theta) = \theta v(q(\theta)) - t(\theta).$$

15.91 A direct mechanism (q, t) is incentive compatible (IC) if truth telling is optimal for every type, that is, if

$$\theta v(q(\theta)) - t(\theta) \geq \theta v(q(\theta')) - t(\theta') \text{ for all } \theta, \theta' \in \Theta.$$

15.92 A direct mechanism is individually rational (IR) if the buyer, conditional on her type, is voluntarily willing to participate, that is, if

$$u(\theta) \geq 0 \text{ for all } \theta \in \Theta.$$

15.93 Proposition: A direct mechanism (q, t) is incentive compatible if and only if q is increasing, and

$$t(\theta) = t(\underline{\theta}) + \left[\theta v(q(\theta)) - \underline{\theta} v(q(\underline{\theta})) \right] + \int_{\underline{\theta}}^{\theta} v(q(x)) \, dx \text{ for all } \theta \in \Theta.$$

15.94 Proposition: An IC direct mechanism (q, t) is IR if and only if $u(\underline{\theta}) \geq 0$.

15.95 The seller's problem reduces to

$$\max_{\substack{q: [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+ \\ q \text{ is increasing}}} U(q) = \int_{\underline{\theta}}^{\bar{\theta}} \left[\theta v(q(\theta)) - \int_{\underline{\theta}}^{\theta} v(q(x)) \, dx - cq(\theta) \right] f(\theta) \, d\theta.$$

15.96 Since

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} v(q(x)) \, dx f(\theta) \, d\theta &= \int_{\underline{\theta}}^{\bar{\theta}} \int_x^{\bar{\theta}} v(q(x)) f(\theta) \, d\theta \, dx = \int_{\underline{\theta}}^{\bar{\theta}} v(q(x)) \int_x^{\bar{\theta}} f(\theta) \, d\theta \, dx \\ &= \int_{\underline{\theta}}^{\bar{\theta}} v(q(x)) (1 - F(x)) \, dx = \int_{\underline{\theta}}^{\bar{\theta}} v(q(\theta)) (1 - F(\theta)) \, d\theta, \end{aligned}$$

we have

$$\begin{aligned} U(q) &= \int_{\underline{\theta}}^{\bar{\theta}} \theta v(q(\theta)) f(\theta) \, d\theta - \int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{\theta}}^{\theta} v(q(x)) \, dx f(\theta) \, d\theta - \int_{\underline{\theta}}^{\bar{\theta}} cq(\theta) f(\theta) \, d\theta \\ &= \int_{\underline{\theta}}^{\bar{\theta}} \left[v(q(\theta)) \left(\theta - \frac{1 - F(\theta)}{f(\theta)} \right) - cq(\theta) \right] f(\theta) \, d\theta. \end{aligned}$$

15.97 If the seller can observe θ , then the seller can get the surplus $\theta v(q(\theta)) - cq(\theta)$.

When θ is not observable, the surplus $\psi(\theta) \triangleq v(q(\theta)) \left(\theta - \frac{1 - F(\theta)}{f(\theta)} \right) - cq(\theta)$ (called virtual surplus) generated by the transaction with buyer type θ that the seller can recover from the buyer is lower than it would be if θ were observable. This can be seen from the fact that we subtract a positive term from θ in $\psi(\theta)$. The subtracted term represents the reduction in surplus recovered by the seller that is due to the constraint that she must offer incentives to reveal θ truthfully. This term is called the buyer's information rent. $\theta - \frac{1 - F(\theta)}{f(\theta)}$ is called the virtual type of θ .

15.98 Suppose we ignore for the moment the constraint that q must be increasing. Then the seller can choose $q(\theta)$ for each θ separately to maximize the expression in the square bracket.

The first-order condition for maximizing $\psi(\theta)$ for given θ is

$$v'(q)\psi(\theta) = c.$$

- If $\psi(\theta) \leq 0$, then there is obviously no solution, and hence the optimal choice is $q(\theta) = 0$.

- If $\psi(\theta) > 0$ and $v'(0)\psi(\theta) \leq c$, since v' is decreasing, $q(\theta) = 0$.
- If $\psi(\theta) > 0$ and $v'(0)\psi(\theta) > c$, there is a unique solution.

15.99 Regular assumption: $\psi(\theta)$ is increasing in θ .

The optimal q is the intersection point of $v'(q)\psi(\theta)$ with c or 0, whatever is greater. It is then easy to see that the optimal q is increasing.

15.100 A sufficient condition for the regular assumption is that $\frac{f(\theta)}{1-F(\theta)}$ is increasing in θ . This condition is often referred to as the “increasing hazard rate” condition. Think of $F(\theta)$ as the probability that an individual dies before time θ . Then $1 - F(\theta)$ is the probability that the individual survives until time θ , and $\frac{f(\theta)}{1-F(\theta)}$ can be thought of as the conditional probability of dying at time θ of an individual that has survived until time θ .

15.101 Proposition: Suppose that F is regular. Then an expected profit maximizing choice of q is given by

- If $v'(0)\psi(\theta) \leq c$, then $q(\theta) = 0$.
- Otherwise, $q(\theta)$ is determined by $v'(q)\psi(\theta) = c$.

15.102 Numerical example: Let $\Theta = [0, 1]$, $c = 1$, $v(q) = \sqrt{q}$ and $F(\theta) = \theta$.

- (1) Regular assumption is satisfied since $\psi(\theta) = 2\theta - 1$ is increasing in θ .
- (2) Next we determine for which values of θ the optimal quantity $q(\theta)$ equals zero:

$$v'(0)\psi(\theta) \leq 0 \Leftrightarrow \psi(\theta) \leq 0 \Leftrightarrow \theta \leq 0.5.$$

- (3) If $\theta > 0.5$, the optimal $q(\theta)$ is given by $v'(q)\psi(\theta) = c$, and so $q(\theta) = (\theta - \frac{1}{2})^2$.
- (4) The corresponding transfer $t(\theta)$ is zero if $\theta \leq 0.5$, and if $\theta > 0.5$ it is given by

$$t(\theta) = \theta v(q(\theta)) - \int_{\underline{\theta}}^{\theta} v(q(x)) \, dx = \frac{\theta^2}{2} - \frac{1}{8}.$$

- (5) We can express the transfer t as a function of q :

$$t(\theta) = \frac{1}{2}q + \frac{1}{2}\sqrt{q},$$

where $q \in [0, 1/4]$.

15.4 Insurance

Literature: [Rothschild and Stiglitz \(1976\)](#).

15.5 Sequential screening

15.103 Literature: [Courty and Li \(2000\)](#) and Chapter 11.2.1 in [Borgers \(2015\)](#).

15.104 The classic mechanism design approach considers a static framework, where agents observe private information only once, and the mechanism implements a single allocation. The model of sequential screening developed by [Courty and Li \(2000\)](#) analyzes a framework in which the allocation itself—the sale of a single indivisible good—is still static, but the agents’ private information changes over time.

15.5.1 A motivating example

15.105 Consider the demand for plane tickets. Travellers typically do not know their valuations for tickets until just before departure, but they know in advance their likelihood to have high and low valuations.

A monopolist can wait until the travellers learn their valuations and charge the monopoly price, but more consumer surplus can be extracted by requiring them to reveal their private information sequentially. (To be shown later.)

15.106 Suppose that one-third of all potential buyers are leisure travellers whose valuation is uniformly distributed on $[1, 2]$, and two-thirds are business travellers whose valuation is uniformly distributed on $[0, 1] \cup [2, 3]$.

Intuitively, business travellers face greater valuation uncertainty than leisure travellers.

Suppose that cost of flying an additional traveller is 1.

15.107 If the seller waits until travellers have privately learned their valuations, she faces a valuation distribution distribution that is uniform on $[0, 3]$. By solving

$$\max_p \frac{1}{3}(3-p)(p-1),$$

the optimal monopoly price is 2 with expected profit of $\frac{1}{3}$, thus excluding all leisure travellers as well as half of business travellers who turned out to have low valuations.

15.108 Suppose instead that the seller offers two contracts before the travellers learn their valuation, one with an advance payment of 1.5 and no refund, and the other with an advance payment of 1.75 and a partial refund of 1.

Leisure travellers strictly prefer the contract with no refund: since the valuation of a leisure traveller is at least 1 which is higher than the partial refund in the second contract, leisure travellers will never get refund when they choose the second contract.

Business travellers are indifferent between the two contracts so we assume that they choose the contract with refund.

The monopolist separates the two types and earns an expected profit of $\frac{2}{3}$:

$$\underbrace{\frac{1}{3}(1.75 - 1)}_{[0,1]} + \underbrace{\frac{1}{3}(1.5 - 1)}_{[1,2]} + \underbrace{\frac{1}{3}(1.75 - 1)}_{[2,3]} = \frac{2}{3},$$

twice as much as the profits of charging the monopoly price after travellers have learned their valuations.

15.5.2 Optimal menu of refund contracts: two-type case

15.109 Consider a monopoly seller of airplane tickets facing two types of travellers, B and L , with proportions p_B and p_L respectively. Throughout this subsection, we will think of type B as the “business traveller,” and type L as the “leisure traveller.”

Suppose that the range of valuations is $[\underline{v}, \bar{v}]$.

15.110 There are two periods.

- In the beginning of period one, the traveller privately learns his type which determines the probability distribution of his valuation for the ticket. The seller and the traveller contract at the end of period one.
- In the beginning of period two, the traveller privately learns his actual valuation v for the ticket, and then decides whether to travel.




Figure 15.9

- 15.111 Each ticket costs the seller c . The seller and the traveller are risk-neutral, and do not discount. The reservation utility of each type of traveller is normalized to zero.
- 15.112 The business type may value the ticket more in the sense of first-order stochastic dominance (FSD): type B 's distribution of valuation F_B first-order stochastically dominates the leisure type's distribution F_L if

$$F_B(v) \leq F_L(v) \text{ for all } v \in [\underline{v}, \bar{v}].$$

It means that the business type has greater mean valuation.

Alternatively, the business type may face greater valuation uncertainty in the sense of mean-preserving-spread (MPS): F_B dominates F_L by MPS if they have the same mean and

$$\int_{\underline{v}}^v (F_B(u) - F_L(u)) \, du \geq 0 \text{ for all } v \in [\underline{v}, \bar{v}].$$

It means that the business type has greater uncertainty.

Let f_B and f_L be the density functions of two types.

- 15.113 A refund contract consists of an advance payment a at the end of period one and a refund k that can be claimed at the end of period two after the traveller learns his valuation.

Clearly, regardless of the payment a , the consumer travels only if he values the ticket more than k .

- 15.114 The seller offers a menu of refund contracts $\{(a_B, k_B), (a_L, k_L)\}$ hoping that a type- B traveller will select (a_B, k_B) and a type- L traveller will select instead (a_L, k_L) .

- 15.115 A menu of refund contracts $\{(a_B, k_B), (a_L, k_L)\}$ is incentive compatible if

$$-a_B + k_B F_B(k_B) + \int_{k_B}^{\bar{v}} v \, dF_B(v) \geq -a_L + k_L F_B(k_L) + \int_{k_L}^{\bar{v}} v \, dF_B(v) \quad (\text{IC}_{B,L})$$

$$-a_L + k_L F_L(k_L) + \int_{k_L}^{\bar{v}} v \, dF_L(v) \geq -a_B + k_B F_L(k_B) + \int_{k_B}^{\bar{v}} v \, dF_L(v) \quad (\text{IC}_{L,B})$$

- 15.116 A menu of refund contracts $\{(a_B, k_B), (a_L, k_L)\}$ is individually rational if

$$-a_B + k_B F_B(k_B) + \int_{k_B}^{\bar{v}} v \, dF_B(v) \geq 0 \quad (\text{IR}_B)$$

$$-a_L + k_L F_L(k_L) + \int_{k_L}^{\bar{v}} v \, dF_L(v) \geq 0 \quad (\text{IR}_L)$$

15.117 The profit maximization problem can be written as

$$\begin{aligned} & \underset{\substack{(a_B, k_B) \\ (a_L, k_L)}}{\text{maximize}} && \sum_{t=B, L} p_t \cdot \left(a_t - k_t F_t(k_t) - c(1 - F_t(k_t)) \right) \\ & \text{subject to} && \text{Equations (IC}_{B,L}\text{)} - (\text{IR}_L). \end{aligned}$$

15.118 Lemma: Under either FSD or MPS, IR_L and $\text{IC}_{B,L}$ imply IR_B .

Proof. (1) $\text{IC}_{B,L}$ implies

$$-a_B + \int_{\underline{v}}^{\bar{v}} \max\{k_B, v\} dF_B(v) \geq -a_L + \int_{\underline{v}}^{\bar{v}} \max\{k_L, v\} dF_B(v).$$

(2) Case 1: Since $\max\{k_L, v\}$ is an increasing function of v , if F_B dominates F_L by FSD,

$$-a_L + \int_{\underline{v}}^{\bar{v}} \max\{k_L, v\} dF_B(v) \geq -a_L + \int_{\underline{v}}^{\bar{v}} \max\{k_L, v\} dF_G(v).$$

Case 2: Since $\max\{k_L, v\}$ is a convex function, if F_B dominates F_L by MPS, we also have

$$-a_L + \int_{\underline{v}}^{\bar{v}} \max\{k_L, v\} dF_B(v) \geq -a_L + \int_{\underline{v}}^{\bar{v}} \max\{k_L, v\} dF_G(v).$$

(3) IR_L implies

$$-a_L + \int_{\underline{v}}^{\bar{v}} \max\{k_L, v\} dF_L(v) \geq 0.$$

(4) Thus,

$$-a_B + \int_{\underline{v}}^{\bar{v}} \max\{k_B, v\} dF_B(v) \geq 0.$$

That is, IR_B holds. □

The business type gets more utility than the leisure type from any refund contract, whether it is defined by greater mean valuation or by greater uncertainty.

15.119 The lemma above implies that IR_L binds in the optimal menu of refund contracts, otherwise increasing both a_L and a_B by the same amount would increase profits.

Also, $\text{IC}_{B,L}$ binds in the optimal menu of refund contracts, otherwise profit could be increased by increasing a_B .

15.120 Substituting IR_L and $\text{IC}_{B,L}$ into the objective function and ignoring $\text{IC}_{L,B}$, we obtain the following “relaxed” problem:

$$\begin{aligned} & \max_{(a_B, k_B), (a_L, k_L)} p_B \cdot \left(a_B - k_B F_B(k_B) - c(1 - F_B(k_B)) \right) + p_L \cdot \left(a_L - k_L F_L(k_L) - c(1 - F_L(k_L)) \right) \\ &= \max_{k_B, k_L} p_B \cdot \left(\int_{k_B}^{\bar{v}} v dF_B(v) - \int_{k_L}^{\bar{v}} v dF_B(v) - k_L F_B(k_L) + k_L F_L(k_L) + \int_{k_L}^{\bar{v}} v dF_L(v) - c(1 - F_B(k_B)) \right) \\ & \quad + p_L \cdot \left(\int_{k_L}^{\bar{v}} v dF_L(v) - c(1 - F_L(k_L)) \right) \\ &= \max_{k_B, k_L} p_B \int_{k_B}^{\bar{v}} v dF_B(v) - p_B c \int_{k_B}^{\bar{v}} 1 dF_B(v) + p_L \int_{k_L}^{\bar{v}} v dF_L(v) - p_L c \int_{k_L}^{\bar{v}} 1 dF_L(v) \end{aligned}$$

$$\begin{aligned}
& -p_B \int_{k_L}^{\bar{v}} v \, dF_B(v) + p_B \int_{k_L}^{\bar{v}} v \, dF_L(v) - p_B k_L F_B(k_L) + p_B k_L F_L(k_L) \\
& = \max_{k_B, k_L} \int_{k_B}^{\bar{v}} p_B(v - c) f_B(v) \, dv + \int_{k_L}^{\bar{v}} p_L(v - c) f_L(v) \, dv - p_B \left(v F_B(v) \Big|_{k_L}^{\bar{v}} - \int_{k_L}^{\bar{v}} F_B(v) \, dv \right) - p_B k_L F_B(k_L) \\
& \quad + p_B \left(v F_L(v) \Big|_{k_L}^{\bar{v}} - \int_{k_L}^{\bar{v}} F_L(v) \, dv \right) + p_B k_L F_L(k_L) \\
& = \max_{k_B, k_L} \int_{k_B}^{\bar{v}} p_B(v - c) f_B(v) \, dv + \underbrace{\int_{k_L}^{\bar{v}} p_L(v - c) f_L(v) \, dv}_{S(k_L)} - \underbrace{\int_{k_L}^{\bar{v}} p_B(F_L(v) - F_B(v)) \, dv}_{R(k_L)}.
\end{aligned}$$

15.121 Since the choice of refund for the business type is unconstrained, it should be equal to c to maximize the surplus from the business type.

15.122 Let

$$S(k_L) = \int_{k_L}^{\bar{v}} (v - c) f_L(v) \, dv$$

be the surplus from the leisure type, and

$$R(k_L) = \int_{k_L}^{\bar{v}} (F_L(v) - F_B(v)) \, dv$$

be the rent to the business type, both as function of the refund to the leisure type.

The solution to the relaxed problem is given by

$$\arg \max_k p_L S(k) - p_B R(k).$$

15.123 Proposition: Under either FSD or MPS, in the optimal menu of refund contracts, $k_B = c$ and $k_L \in \arg \max_k p_L S(k) - p_B R(k)$.

Proof. (1) It suffices to show that $IC_{L,B}$ is satisfied by the given $\{k_B, k_L\}$.

(2) Since $IC_{B,L}$ binds, we have

$$\begin{aligned}
a_L - a_B &= \int_{k_L}^{k_B} v \, dF_B(v) + k_L F_B(k_L) - k_B F_B(k_B) \\
&= v F_B(v) \Big|_{k_L}^{k_B} - \int_{k_L}^{k_B} F_B(v) \, dv + k_L F_B(k_L) - k_B F_B(k_B) = \int_{k_B}^{k_L} F_B(v) \, dv.
\end{aligned}$$

(3) Then we have

$$-a_L + k_L F_L(k_L) + \int_{k_L}^{\bar{v}} v \, dF_L(v) = -a_B + k_B F_L(k_B) + \int_{k_B}^{\bar{v}} v \, dF_L(v) - \int_{k_B}^{k_L} (F_B(v) - F_L(v)) \, dv.$$

(4) Thus, $IC_{L,B}$ is satisfied if and only if

$$\int_{k_B}^{k_L} (F_B(v) - F_L(v)) \, dv \leq 0.$$

(5) Case 1: If the solution to the relaxed problem has $k_L = c$, then $\int_{k_B}^{k_L} (F_B(v) - F_L(v)) \, dv = 0$. The proposition follows immediately.

(6) Case 2: Suppose that the solution $k_L \neq c$, and $\int_{k_B}^{k_L} (F_B(v) - F_L(v)) dv > 0$.

(i) Now consider an alternative menu where the two types have the same refund c .

(ii) The surplus $S(c)$ from type L is greater in the alternative menu.

(iii) The rent $R(c)$ to type B is smaller in the alternative menu:

$$R(c) = \int_{k_L}^{\bar{v}} (F_L(v) - F_B(v)) dv + \int_{k_B}^{k_L} (F_L(v) - F_B(v)) dv < \int_{k_L}^{\bar{v}} (F_L(v) - F_B(v)) dv.$$

(iv) This contradicts the assumption that $\{k_L, k_B\}$ solves the relaxed problem.


□

15.124 Consider the example in last subsection. By definition, the surplus function is

$$S(k) = \begin{cases} \frac{1}{2}, & \text{if } k \in [0, 1), \\ k - \frac{1}{2}k^2, & \text{if } k \in [1, 2], \\ 0, & \text{if } k \in (2, 3], \end{cases}$$

and the rent function is

$$R(k) = \begin{cases} \frac{1}{4}k^2, & \text{if } k \in [0, 1), \\ \frac{1}{4} - \frac{1}{2}(k-1)(k-2), & \text{if } k \in [1, 2], \\ \frac{1}{4}(k-3)^2, & \text{if } k \in (2, 3]. \end{cases}$$



According to the proposition above, $k_B = c = 1$ and $k_L = \arg \max_k \frac{1}{3}S(k) - \frac{2}{3}R(k) = 0$. Since IR_L binds, we have $a_L = 1.5$. Since $IC_{B,L}$ binds, we have $a_B = 1.75$. This verifies the optimality of the refund contracts given in the last subsection.

15.5.3 Sequential screening: general type case

15.125 We assume that types are continuously distributed over $\Theta = [\underline{\theta}, \bar{\theta}]$ ($\underline{\theta} \geq 0$), with a density function $g(\theta) > 0$ and distribution $G(\theta)$.

Consumer type is privately learned before contracting.

15.126 Each type θ is represented by a distribution of valuations over $[\underline{v}, \bar{v}]$, with a differentiable density function $f(v | \theta) > 0$ and cumulative distribution function $F(v | \theta)$. We assume that the conditional distributions $f(v | \theta)$ have the same support for all $\theta \in \Theta$.

15.127 We assume that the family $F(v \mid \theta)$ ($\theta \in \Theta$) is ordered in the sense of FSD so that

$$F(v \mid \theta) \leq F(v \mid \theta') \text{ for } \theta \geq \theta' \text{ and } v \in (\underline{v}, \bar{v}).$$

FSD captures the idea that a higher type θ is good news about the buyer's valuation v .

15.128 A (dynamic) direct mechanism consists of functions q and t where

$$q: [\underline{\theta}, \bar{\theta}] \times [\underline{v}, \bar{v}] \rightarrow [0, 1], \quad t: [\underline{\theta}, \bar{\theta}] \times [\underline{v}, \bar{v}] \rightarrow \mathbb{R}.$$

A dynamic direct mechanism requires the buyer to report her private information and then seems a straightforward extension of static mechanisms.

15.129 For dynamic revelation principle, see Proposition 11.1 in [Borgers \(2015\)](#).

15.130 Given a direct mechanism, let $u(\theta, v) = vq(\theta, v) - t(\theta, v)$, and

$$U(\theta' \mid \theta) = \int_{\underline{v}}^{\bar{v}} u(\theta', v) f(v \mid \theta) dv,$$

and $U(\theta) = U(\theta \mid \theta)$.

15.131 A direct mechanism is incentive compatible if

- (a) it is incentive compatible with respect to the *ex post* valuation v , meaning that truth telling about v is optimal for every *ex ante* type θ and every *ex post* valuation v :

$$u(\theta, v) \geq vq(\theta, v') - t(\theta, v') \text{ for all } \theta \in [\underline{\theta}, \bar{\theta}] \text{ and } v, v' \in [\underline{v}, \bar{v}].$$

- (b) it is incentive compatible with respect to the *ex ante* type θ , meaning that truth telling about θ is optimal for every *ex ante* type θ and any subsequent report about v :

$$U(\theta) \geq \int_{\underline{v}}^{\bar{v}} [vq(\theta', v^r(v)) - t(\theta', v^r(v))] f(v \mid \theta) dv$$

for all $\theta, \theta' \in [\underline{\theta}, \bar{\theta}]$ and all $v^r: [\underline{v}, \bar{v}] \rightarrow [\underline{v}, \bar{v}]$.

15.132 A direct mechanism is individually rational if

$$U(\theta) \geq 0 \text{ for all } \theta \in \Theta.$$

15.133 Proposition: A direct mechanism is incentive compatible if and only if it satisfies

$$u(\theta, v) \geq vq(\theta, v') - t(\theta, v') \text{ for all } \theta \in [\underline{\theta}, \bar{\theta}] \text{ and } v, v' \in [\underline{v}, \bar{v}],$$

and

$$U(\theta) \geq U(\theta' \mid \theta) \text{ for all } \theta, \theta' \in \Theta.$$

Proof. (1) If a direct mechanism is incentive compatible, then it should satisfy the two equations above.

(2) Assume that the two equations above hold.

(3) Since $u(\theta, v) \geq vq(\theta, v') - t(\theta, v')$ for all θ, v and v' , we have

$$u(\theta', v) \geq vq(\theta', v^r(v)) - t(\theta', v^r(v))$$

for all θ', v and $v^r: [\underline{v}, \bar{v}] \rightarrow [\underline{v}, \bar{v}]$.

(4) Since $U(\theta) \geq U(\theta' | \theta) = \int_{\underline{v}}^{\bar{v}} u(\theta', v) f(v | \theta) dv$, we have

$$U(\theta) \geq \int_{\underline{v}}^{\bar{v}} u(\theta', v) f(v | \theta) dv \geq \int_{\underline{v}}^{\bar{v}} [vq(\theta', v^r(v)) - t(\theta', v^r(v))] f(v | \theta) dv.$$

□

15.134 Proposition: A direct mechanism is incentive compatible with respect to the *ex post* valuation v if and only if

- for every *ex ante* type θ , the function $q(\theta, v)$ is increasing in v .
- for every *ex ante* type θ and *ex post* valuation v : $\partial u(\theta, v) / \partial v = q(\theta, v)$.
- for every *ex ante* type θ and *ex post* valuation v :

$$t(\theta, v) = t(\theta, \underline{v}) + (vq(\theta, v) - \underline{v}q(\theta, \underline{v})) - \int_{\underline{v}}^v q(\theta, v) dv.$$

The proof is the same as that of Proposition 15.57.

15.135 Proposition: If a direct mechanism is incentive compatible with respect to the *ex ante* type θ , then

- for every *ex ante* type θ :

$$U'(\theta) = - \int_{\underline{v}}^{\bar{v}} q(\theta, v) \frac{\partial F(v | \theta)}{\partial \theta} dv;$$

- for every *ex ante* type θ :

$$\begin{aligned} & \int_{\underline{v}}^{\bar{v}} t(\theta, v) f(v | \theta) dv \\ &= \int_{\underline{v}}^{\bar{v}} vq(\theta, v) f(v | \theta) dv + \int_{\underline{v}}^{\bar{v}} [t(\underline{\theta}) - vq(\underline{\theta}, v)] f(v | \underline{\theta}) dv + \int_{\underline{\theta}}^{\theta} \int_{\underline{v}}^{\bar{v}} q(\theta', v) \frac{\partial F(v | \theta')}{\partial \theta} d\theta' dv. \end{aligned}$$

Proof. (1) Incentive compatibility implies

$$U(\theta + \delta) \geq U(\theta | \theta + \delta) = \int_{\underline{v}}^{\bar{v}} u(\theta, v) f(v | \theta + \delta) dv,$$

for all $\theta \in \Theta$ and $\delta > 0$.

(2) Thus, we have

$$\begin{aligned} \lim_{\delta \downarrow 0} \frac{U(\theta + \delta) - U(\theta)}{\delta} &\geq \lim_{\delta \downarrow 0} \int_{\underline{v}}^{\bar{v}} u(\theta, v) \frac{f(v | \theta + \delta) - f(v | \theta)}{\delta} dv \\ &= \int_{\underline{v}}^{\bar{v}} u(\theta, v) \lim_{\delta \downarrow 0} \frac{f(v | \theta + \delta) - f(v | \theta)}{\delta} dv = \int_{\underline{v}}^{\bar{v}} u(\theta, v) \frac{\partial f(v | \theta)}{\partial \theta} dv, \end{aligned}$$

where we exchange the order of taking the limit and the integral by Lebesgue's dominated convergence theorem.

(3) Likewise,

$$\lim_{\delta \downarrow 0} \frac{U(\theta) - U(\theta - \delta)}{\delta} \leq \lim_{\delta \downarrow 0} \int_{\underline{v}}^{\bar{v}} u(\theta, v) \frac{\partial f(v | \theta)}{\partial \theta} dv.$$

(4) Putting the two inequalities together we obtain

$$U'(\theta) = \int_{\underline{v}}^{\bar{v}} u(\theta, v) \frac{\partial f(v | \theta)}{\partial \theta} dv = - \int_{\underline{v}}^{\bar{v}} q(\theta, v) \frac{\partial F(v | \theta)}{\partial \theta} dv.$$

(5) Hence,

$$U(\theta) = U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} \int_{\underline{v}}^{\bar{v}} q(\theta', v) \frac{\partial F(v | \theta')}{\partial \theta'} dv d\theta',$$

which completes the proof. □

15.136 Proposition: If a direct mechanism is incentive compatible, then

$$t(\theta, v) = t_0(\theta) + vq(\theta, v) - \int_{\underline{v}}^v q(\theta, v') dv',$$

where $t_0(\theta)$ is

$$= t(\underline{\theta}, \underline{v}) - \underline{v}q(\underline{\theta}, \underline{v}) + \int_{\underline{\theta}}^{\theta} \int_{\underline{v}}^{\bar{v}} q(\theta', v) \frac{\partial F(v | \theta')}{\partial \theta'} d\theta' dv + \int_{\underline{v}}^{\bar{v}} \int_{\underline{v}}^v [q(\theta, v')f(v | \theta) - q(\underline{\theta}, v')f(v | \underline{\theta})] dv' dv.$$

15.137 Proposition: If $q(\theta, v)$ is increasing in θ and v , then there exists a transfer schedule $t(\theta, v)$ such that the direct mechanism (q, t) is incentive compatible.

15.138 Proposition: An incentive compatible direct mechanism is individually rational if and only if $U(\underline{\theta}) \geq 0$.

15.5.4 Optimal selling mechanism

15.139

Chapter 16

Information economics 3: moral hazard

Contents

16.1 The principle-agent problem	291
16.2 Complete information	293
16.3 Incomplete information with risk-neutral agent	294
16.4 Incomplete information with limited liability	295
16.5 Incomplete information with risk-averse agent	296
16.6 n levels of performance	298
16.6.1 Limited liability	298
16.6.2 Risk aversion	299

16.1 Moral hazard models hidden action, where asymmetric information forms after the parties enter into a relationship. A moral hazard is a situation in which a party is more likely to take risks because the costs that could result will not be borne by the party taking the risk. Moral hazard arises because an individual or institution does not take the full consequences and responsibilities of its actions, and therefore has a tendency to act less carefully than it otherwise would, leaving another party to hold some responsibility for the consequences of those actions.

In a principal–agent problem, one party, called an agent, acts on behalf of another party, called the principal. The agent usually has more information about his or her actions or intentions than the principal does, because the principal usually cannot completely monitor the agent. The agent may have an incentive to act inappropriately (from the viewpoint of the principal) if the interests of the agent and the principal are not aligned.

In particular, consider that a firm (the principal) hires a worker (the agent) to work on a project, which succeeds with probability p if the worker exerts effort. The firm may only observe the outcome of the project but not the agent's effort level. In such a situation, the firm's payment contract can only depend on the outcome, which is an imperfect indicator of the worker's effort level. If the worker is paid fixed wage or if the payment conditional on success is not high enough, since effort is costly, the worker will shirk—moral hazard arises.

16.1 The principle-agent problem

16.2 A principal (employer) hires an agent (employee) for production. The agent can exert a costly effort $e \in \{0, 1\}$. Exerting effort e implies a disutility for the agent that is equal to $\psi(e)$ with the normalizations $\psi(0) = 0$ and

$\psi(1) = \psi$. The agent receives a transfer t from the principal.

The agent's utility is assumed to be

$$U = u(t) - \psi(e),$$

where u is increasing and concave, and $u(0) = 0$. We also assume that $h = u^{-1}$, which is increasing and convex.

- 16.3 Production is stochastic, and effort affects the production level as follows: the stochastic production level q can only take two values $\{\underline{q}, \bar{q}\}$ with $\Delta q \triangleq \bar{q} - \underline{q} > 0$, and the stochastic influence of effort on production is characterized by the probabilities

$$\text{Prob}(q = \bar{q} \mid e = 0) = \pi_0 \text{ and } \text{Prob}(q = \bar{q} \mid e = 1) = \pi_1,$$

with $\Delta\pi \triangleq \pi_1 - \pi_0 > 0$.

- 16.4 Effort improves production in the sense of first-order stochastic dominance, *i.e.*, $\text{Prob}(q \leq q^* \mid e)$ is decreasing with e for any given production q^* . That is, we have

$$\begin{aligned} \text{Prob}(q \leq \underline{q} \mid e = 1) &= 1 - \pi_1 < 1 - \pi_0 = \text{Prob}(q \leq \underline{q} \mid e = 0), \\ \text{Prob}(q \leq \bar{q} \mid e = 1) &= 1 = \text{Prob}(q \leq \bar{q} \mid e = 0). \end{aligned}$$

- 16.5 The value for the principal of these q units is $S(q)$ where $S' > 0$, $S'' < 0$ and $S(0) = 0$. The marginal value of the good is thus positive and strictly decreasing with the number of units bought by the principal.

The risk-neutral principal's expected utility is

$$V_1 = \pi_1(S(\bar{q}) - \bar{t}) + (1 - \pi_1)(S(\underline{q}) - \underline{t})$$

if the agent makes a positive effort $e = 1$, and

$$V_0 = \pi_0(S(\bar{q}) - \bar{t}) + (1 - \pi_0)(S(\underline{q}) - \underline{t})$$

if the agent makes no effort $e = 0$.

- 16.6 The problem of the principal is to decide whether to induce the agent to exert effort or not and, if he chooses to do so, then to decide which incentive contract should be used.

- 16.7 Incentive compatibility: The principal wishes to induce the agent to exert effort, that is,

$$\pi_1 u(\bar{t}) + (1 - \pi_1) u(\underline{t}) - \psi \geq \pi_0 u(\bar{t}) + (1 - \pi_0) u(\underline{t}).$$

- 16.8 Individual rationality: The agent is willing to exert effort, that is,

$$\pi_1 u(\bar{t}) + (1 - \pi_1) u(\underline{t}) - \psi \geq 0,$$

where we normalize the agent's reservation utility at zero.

- 16.9 The timing is as follows:




Figure 16.1

16.2 Complete information

16.10 First assume that the principal can observe effort. Then the principal's problem becomes:

$$\begin{aligned} & \underset{(\bar{t}, \underline{t})}{\text{maximize}} && \pi_1(S(\bar{q}) - \bar{t}) + (1 - \pi_1)(S(\underline{q}) - \underline{t}) \\ & \text{subject to} && \pi_1 u(\bar{t}) + (1 - \pi_1)u(\underline{t}) - \psi \geq 0. \end{aligned}$$

Indeed, only the agent's individual rationality matters for the principal, because the agent can be forced to exert a positive level of effort. If the agent were not choosing positive effort, his action could be perfectly detected by the principal, and hence the agent could be heavily punished.

16.11 Denoting the multiplier of the individual rationality constraint by μ and optimizing with respect to \bar{t} and \underline{t} yields, respectively, the following first-order conditions:

$$\begin{aligned} -\pi_1 + \mu \pi_1 u'(\bar{t}^*) &= 0, \\ -(1 - \pi_1) + \mu(1 - \pi_1)u'(\underline{t}^*) &= 0, \end{aligned}$$

where \bar{t}^* and \underline{t}^* are the first-best transfers.

We immediately derive that $\mu = \frac{1}{u'(\underline{t}^*)} = \frac{1}{u'(\bar{t}^*)} > 0$, and finally that $t^* = \bar{t}^* = \underline{t}^*$.

16.12 Remark:

- The transfer t^* the agent receives is the same whatever the state of nature.
- Because the IR constraint is binding we also obtain the value of this transfer, which is just enough to cover the disutility of effort, namely $t^* = h(\psi)$. It is called the first-best cost C^* of implementing the positive effort level.

16.13 For the principal, inducing effort yields an expected payoff equal to

$$V_1 = \pi_1 S(\bar{q}) + (1 - \pi_1)S(\underline{q}) - h(\psi).$$

Had the principal decided to let the agent exert no effort, $e = 0$, he would make a zero payment to the agent whatever the realization of output. In this scenario, the principal would instead obtain a payoff equal to

$$V_0 = \pi_0 S(\bar{q}) + (1 - \pi_0)S(\underline{q}).$$

Inducing effort is thus optimal from the principal's point of view when $V_1 \geq V_0$, i.e.,

$$\Delta\pi\Delta S \geq h(\psi), \tag{16.1}$$

where $\Delta S = S(\bar{q}) - S(\underline{q})$.

16.14 The left-hand side of Equation (16.1) captures the gain of increasing effort from $e = 0$ to $e = 1$. This gain comes from the fact that the return $S(\bar{q})$, which is greater than $S(\underline{q})$, arises more often when a positive effort is exerted. The right-hand side of Equation (16.1) is instead the first-best cost of inducing the agent's acceptance when he exerts a positive effort.

16.15 The first-best outcome calls for $e^* = 1$ if and only if $\Delta\pi\Delta S \geq h(\psi)$.

16.3 Incomplete information with risk-neutral agent

16.16 If the agent is risk-neutral, we can assume that (up to an affine transformation) $u(t) = t$ for all t and $h(u) = u$ for all u .

16.17 The principal's problem is

$$\begin{aligned} & \underset{(\bar{t}, \underline{t})}{\text{maximize}} && \pi_1(S(\bar{q}) - \bar{t}) + (1 - \pi_1)(S(\underline{q}) - \underline{t}) \\ & \text{subject to} && \pi_1\bar{t} + (1 - \pi_1)\underline{t} - \psi \geq \pi_0\bar{t} + (1 - \pi_0)\underline{t} \\ & && \pi_1\bar{t} + (1 - \pi_1)\underline{t} - \psi \geq 0. \end{aligned}$$

16.18 IR condition should be binding; otherwise the principal can decrease \bar{t} or \underline{t} with breaking IR condition.

16.19 Let IC condition be binding. Then we have

$$\bar{t}^* = \frac{1 - \pi_0}{\Delta\pi}\psi \text{ and } \underline{t}^* = -\frac{\pi_0}{\Delta\pi}\psi. \quad (16.2)$$

The agent is rewarded if production is high, and his utility is $\bar{t}^* - \psi = \frac{1 - \pi_1}{\Delta\pi}\psi > 0$.

The agent is punished if production is low, and his utility is $\underline{t}^* - \psi = -\frac{\pi_1}{\Delta\pi}\psi < 0$.

The principal makes an expected payment

$$\pi_1\bar{t}^* + (1 - \pi_1)\underline{t}^* = \psi,$$

which is equal to the disutility of effort he would incur if he could control the effort level perfectly or if he was carrying the agent's task himself.

The principal can costlessly structure the agent's payment so that the latter has the right incentives to exert effort. Indeed, by increasing effort from $e = 0$ to $e = 1$, the agent receives the transfer \bar{t}^* more often than the transfer \underline{t}^* . His expected gain from exerting effort is thus $\Delta\pi(\bar{t}^* - \underline{t}^*) = \psi$, i.e., it exactly compensates the agent for the extra disutility of effort that he incurs when increasing his effort from $e = 0$ to $e = 1$.

Therefore, we have shown: Moral hazard is not an issue with a risk-neutral agent despite the nonobservability of effort. The first-best level of effort is still implemented.

16.20 Let us consider another pair of transfers

$$\bar{t}^{*'} = S(\bar{q}) - T^* \text{ and } \underline{t}^{*'} = S(\underline{q}) - T^*,$$

where T^* is an up-front payment made by the agent before output realizes.

These transfers satisfy the agent's IC constraint since:

$$\Delta(\bar{t}^{*'} - \underline{t}^{*'}) = \Delta\pi\Delta S > h(\psi) = \psi,$$

where the inequality comes from the fact that effort is socially optimal in a first-best world (Equation (16.1)).

The up-front payment T^* can be adjusted by the principal to have the agent's IR constraint be binding:

$$T^* = \pi_1 S(\bar{q}) + (1 - \pi_1) S(\underline{q}) - \psi.$$

With the transfers $\bar{t}^{*'}$ and $\underline{t}^{*'}$, the agent becomes residual claimant for the profit of the firm. The up-front payment T^* is precisely equal to this expected profit. The principal chooses this *ex ante* payment to reap all gains from delegation.

16.4 Incomplete information with limited liability

16.21 Let us consider a risk-neutral agent. Let us also assume that the agent's transfer must always be greater than some exogenous level $-l$, with $l \geq 0$. Limited liability in both states are thus written as

$$\bar{t} \geq -l \text{ and } \underline{t} \geq -l.$$

16.22 The principal's problem is

$$\begin{aligned} \underset{(\bar{t}, \underline{t})}{\text{maximize}} \quad & \pi_1(S(\bar{q}) - \bar{t}) + (1 - \pi_1)(S(\underline{q}) - \underline{t}) \\ \text{subject to} \quad & \pi_1 \bar{t} + (1 - \pi_1) \underline{t} - \psi \geq \pi_0 \bar{t} + (1 - \pi_0) \underline{t} \\ & \pi_1 \bar{t} + (1 - \pi_1) \underline{t} - \psi \geq 0 \\ & \bar{t} \geq -l \\ & \underline{t} \geq -l \end{aligned}$$

16.23 For $l > \frac{\pi_0}{\Delta\pi} \psi$, the first-best outcome can be implemented, and one optimal transfers are given by Equation (16.2). In this case, the agent has no expected limited liability rent.

16.24 For $0 \leq l \leq \frac{\pi_0}{\Delta\pi} \psi$, we conjecture that the IC condition and the limited liability condition for low production level are only relevant constraints.

- (1) The limited liability condition for high production level is obviously irrelevant (IC implies $\bar{t} \geq \frac{\psi}{\Delta\pi} + \underline{t}$).
- (2) The IR condition is also irrelevant:

$$\pi_1 \bar{t} + (1 - \pi_1) \underline{t} - \psi \geq \pi_1(-l + \frac{\psi}{\Delta\pi}) + (1 - \pi_1)(-l) - \psi = \frac{\pi_0}{\Delta\pi} \psi - l \geq 0.$$

- (3) Since the principal is willing to minimize the transfers made to the agent, both constraints must be binding.
- (4) Therefore,

$$\bar{t}^{**} = -l + \frac{\psi}{\Delta\pi} \text{ and } \underline{t}^{**} = -l.$$

In this case, the agent's expected limited liability rent is non-negative:

$$\pi_1 \bar{t}^{**} + (1 - \pi_1) \underline{t}^{**} - \psi = -l + \frac{\pi_0}{\Delta\pi} \psi \geq 0.$$

16.25 Remark:

- Only the limited liability constraint for the bad state may be binding.
- When the limited liability constraint for the bad state is binding, the principal is limited in his punishments to induce effort. The risk-neutral agent does not have enough assets to cover the punishment requested by the principal if \underline{q} is realized in order to induce effort provision. As a result, the agent receives a non-negative *ex ante* limited liability rent. Compared with the case without limited liability, this rent is actually the additional payment that the principal must incur because of the conjunction of moral hazard and limited liability.
- As the agent is endowed with more assets, *i.e.*, as l gets larger, the conflict between moral hazard and limited liability diminishes and then disappears whenever l is large enough. In this case, the agent avoids bankruptcy even when he has to pay the optimal penalty to the principal in the bad state of nature.

16.26 For the sake of simplicity, we assume $l = 0$.

When the principal induces positive effort from the agent, his expected utility is

$$V_1^{**} = \pi_1 S(\bar{q}) + (1 - \pi_1) S(\underline{q}) - \frac{\pi_1}{\Delta\pi} \psi.$$

When the principal gives up the goal of inducing effort from the agent, he can choose $\bar{t} = \underline{t} = 0$ and instead obtain the expected utility level

$$V_0 = \pi_0 S(\bar{q}) + (1 - \pi_0) S(\underline{q}).$$

It is worth inducing effort if $V_1^{**} \geq V_0$, *i.e.*, when

$$\Delta\pi \Delta S \geq \psi + \frac{\pi_0}{\Delta\pi} \psi.$$

The left-hand side is the gain of inducing effort, *i.e.*, the gain of increasing the probability of a high production level. The right-hand side is instead the second-best cost C^{**} of inducing effort, which is the disutility of effort ψ plus the limited liability rent $\frac{\pi_0}{\Delta\pi} \psi$. This second-best cost of implementing effort obviously exceeds the first-best cost. It is clear that the limited liability and moral hazard together make it more costly to induce effort.

16.5 Incomplete information with risk-averse agent

16.27 When the agent is risk-averse, the principal's program is written as:

$$\begin{aligned} & \underset{(\bar{t}, \underline{t})}{\text{maximize}} && \pi_1 (S(\bar{q}) - \bar{t}) + (1 - \pi_1) (S(\underline{q}) - \underline{t}) \\ & \text{subject to} && \pi_1 u(\bar{t}) + (1 - \pi_1) u(\underline{t}) - \psi \geq \pi_0 u(\bar{t}) + (1 - \pi_0) u(\underline{t}) \\ & && \pi_1 u(\bar{t}) + (1 - \pi_1) u(\underline{t}) - \psi \geq 0. \end{aligned}$$

16.28 Let $\bar{u} = u(\bar{t})$ and $\underline{u} = u(\underline{t})$. Then the principal's program can be written as:

$$\begin{aligned} & \underset{(\bar{u}, \underline{u})}{\text{maximize}} && \pi_1(S(\bar{q}) - h(\bar{u})) + (1 - \pi_1)(S(\underline{q}) - h(\underline{u})) \\ & \text{subject to} && \pi_1 \bar{u} + (1 - \pi_1) \underline{u} - \psi \geq \pi_0 \bar{u} + (1 - \pi_0) \underline{u} \\ & && \pi_1 \bar{u} + (1 - \pi_1) \underline{u} - \psi \geq 0. \end{aligned}$$

Note that the principal's objective function is now strictly concave in (\bar{u}, \underline{u}) because h is strictly convex. The constraints are now linear and the interior of the constrained set is obviously nonempty, and therefore it is a concave problem, with the Kuhn and Tucker conditions being sufficient and necessary for characterizing optimality.

16.29 Letting λ and μ be the non-negative multipliers associated respectively with the constraints, the first-order conditions of this program can be expressed as

$$\begin{aligned} -\pi_1 h'(\bar{u}^{**}) + \lambda \Delta \pi + \mu \pi_1 &= -\frac{\pi_1}{u'(\bar{t}^{**})} + \lambda \Delta \pi + \mu \pi_1 = 0 \\ -(1 - \pi_1) h'(\underline{u}^{**}) - \lambda \Delta \pi + \mu(1 - \pi_1) &= -\frac{1 - \pi_1}{u'(\underline{t}^{**})} - \lambda \Delta \pi + \mu(1 - \pi_1) = 0, \end{aligned}$$

where \bar{t}^{**} and \underline{t}^{**} are the second-best optimal transfers.

16.30 Rearranging terms, we get

$$\frac{1}{u'(\bar{t}^{**})} = \mu + \lambda \frac{\Delta \pi}{\pi_1} \text{ and } \frac{1}{u'(\underline{t}^{**})} = \mu - \lambda \frac{\Delta \pi}{1 - \pi_1}.$$

Multiplying the left equation by π_1 and the right equation by $1 - \pi_1$, and then adding those two modified equations, we obtain

$$\mu = \frac{\pi_1}{u'(\bar{t}^{**})} + \frac{1 - \pi_1}{u'(\underline{t}^{**})} > 0.$$

Hence, the IR condition is binding.

16.31 The IC condition implies

$$\bar{u}^{**} - \underline{u}^{**} \geq \frac{\psi}{\Delta \pi} > 0,$$

and thus $\bar{t}^{**} > \underline{t}^{**}$.

Therefore,

$$\lambda = \frac{\pi_1(1 - \pi_1)}{\Delta \pi} \left(\frac{1}{u'(\bar{t}^{**})} - \frac{1}{u'(\underline{t}^{**})} \right) > 0,$$

and hence the IC condition is also binding.

16.32 Since the IR and IC conditions are binding, we have

$$\bar{u}^{**} = \psi + \frac{1 - \pi_1}{\Delta \pi} \psi \text{ and } \underline{u}^{**} = \psi - \frac{\pi_1}{\Delta \pi} \psi,$$

and hence

$$\bar{t}^{**} = h \left(\psi + \frac{1 - \pi_1}{\Delta \pi} \psi \right) \text{ and } \underline{t}^{**} = h \left(\psi - \frac{\pi_1}{\Delta \pi} \psi \right).$$

16.33 The agent receives more than the complete information transfer when a high output is realized, $\bar{t}^{**} > h(\psi)$. When a low output is realized, the agent instead receives less than the complete information transfer, $\underline{t}^{**} < h(\psi)$.

A risk premium must be paid to the risk-averse agent to induce his participation since he now incurs a risk by the fact that $\underline{t}^{**} < \bar{t}^{**}$. Indeed, we have

$$\psi = \pi_1 u(\bar{t}^{**}) + (1 - \pi_1) u(\underline{t}^{**}) < u(\pi_1 \bar{t}^{**} + (1 - \pi_1) \underline{t}^{**}),$$

where the inequality follows from Jensen's inequality. That is, the expected payment $\pi_1 \bar{t}^{**} + (1 - \pi_1) \underline{t}^{**}$ given by the principal is thus larger than the first-best cost $h(\psi)$, which is incurred by the principal when effort is observable.

16.34 The second-best cost of inducing effort under moral hazard is the expected payment made to the agent

$$C^{**} = \pi_1 \bar{t}^{**} + (1 - \pi_1) \underline{t}^{**} = \pi_1 h\left(\psi + \frac{1 - \pi_1}{\Delta\pi} \psi\right) + (1 - \pi_1) h\left(\psi - \frac{\pi_1}{\Delta\pi} \psi\right) > h(\psi) = C^*,$$

where the inequality follows from Jensen's inequality (h is strictly convex).

16.35 The benefit of inducing effort is still $\Delta\pi\Delta S$, and a positive effort $e^* = 1$ is the optimal choice of the principal whenever

$$\Delta\pi\Delta S \geq C^{**} > C^*.$$

16.6 n levels of performance

16.36 We consider a production process where n possible outcomes can be realized. Those performances can be ordered so that

$$q_1 < q_2 < \dots < q_i < \dots < q_n.$$

We denote the principal's return in each of those states of nature by $S_i = S(q_i)$.

Let π_{ik} be the probability that production q_i takes place when the effort level is e_k . We assume that $\pi_{ik} > 0$ for all pairs (i, k) with $\sum_{i=1}^n \pi_{ik} = 1$. We still denote $\Delta\pi_i = \pi_{i1} - \pi_{i0}$.

Only two levels of effort are feasible, i.e., $e_k \in \{0, 1\}$.

16.37 In this context, a contract is a n -uple of payments (t_1, \dots, t_n) .

16.6.1 Limited liability

16.38 The principal's program is

$$\begin{aligned} & \underset{(t_1, \dots, t_n)}{\text{maximize}} && \sum_{i=1}^n \pi_{i1} (S_i - t_i) \\ & \text{subject to} && \sum_{i=1}^n \pi_{i1} t_i - \psi \geq \sum_{i=1}^n \pi_{i0} t_i \\ & && \sum_{i=1}^n \pi_{i1} t_i - \psi \geq 0 \\ & && t_i \geq 0 \text{ for } i = 1, \dots, n. \end{aligned}$$

16.39 The IR condition is implied by the IC condition and LL conditions:

$$\sum_{i=1}^n \pi_{i1} t_i - \psi = \sum_{i=1}^n (\pi_{i1} - \pi_{i0}) t_i - \psi + \sum_{i=1}^n \pi_{i0} t_i \geq 0 + 0 = 0.$$

16.40 Denoting the multiplier of the IC condition by λ and the respective multipliers of the LL conditions by ξ_i , the first-order conditions leads to

$$-\pi_{i1} + \lambda \Delta \pi_i + \xi_i = 0,$$

with the slackness condition $\xi_i t_i = 0$ for each $i = 1, \dots, n$.

- (1) For i such that $t_i^{**} > 0$, we have $\xi_i = 0$, and hence $\lambda = \frac{\pi_{i1}}{\pi_{i1} - \pi_{i0}}$ for any such i .
- (2) If the ratios $\frac{\pi_{i1}}{\pi_{i1} - \pi_{i0}}$ are all different, there exists a single index j such that $\frac{\pi_{j1}}{\pi_{j1} - \pi_{j0}}$ is the highest ratio.
- (3) Then the agent receives a strictly positive transfer only in this particular state of nature j , and this payment is such that the IC condition is binding, i.e., $t_j^{**} = \frac{\psi}{\pi_{j1} - \pi_{j0}}$. In all other states, the agent receives no transfer and $t_i^{**} = 0$ for all $i \neq j$.
- (4) The agent gets a strictly positive *ex ante* limited liability rent $\pi_{j1} \frac{\psi}{\pi_{j1} - \pi_{j0}} - \psi = \frac{\pi_{j0}}{\pi_{j1} - \pi_{j0}} \psi$.

The important point here is that the agent is rewarded in the state of nature that is the most informative about the fact that he has exerted a positive effort. Indeed, $\frac{\pi_{i1} - \pi_{i0}}{\pi_{i1}}$ can be interpreted as a likelihood ratio. The principal therefore uses a maximum likelihood ratio criterion to reward the agent. The agent is only rewarded when this likelihood ratio is maximum.

16.41 The probabilities of success satisfy the monotone likelihood ratio property (MLRP) if $\frac{\pi_{i1} - \pi_{i0}}{\pi_{i1}}$ is nondecreasing in i .

A higher effort level increases the likelihood of a high production level more than the likelihood of a low production level.

16.42 If the probability of success satisfies MLRP, the second-best payment t_i^{**} received by the agent may be chosen to be nondecreasing with the level of production q_i .

16.6.2 Risk aversion

16.43 Suppose now that the agent is strictly risk-averse.

16.44 The principal's problem is

$$\begin{aligned} & \underset{(t_1, \dots, t_n)}{\text{maximize}} && \sum_{i=1}^n \pi_{i1} (S_i - t_i) \\ & \text{subject to} && \sum_{i=1}^n \pi_{i1} u(t_i) - \psi \geq \sum_{i=1}^n \pi_{i0} u(t_i) \\ & && \sum_{i=1}^n \pi_{i1} u(t_i) - \psi \geq 0. \end{aligned}$$

16.45 The first-order condition (with multipliers λ and μ) leads to

$$\frac{1}{u'(t_i^{**})} = \mu + \lambda \left(\frac{\pi_{i1} - \pi_{i0}}{\pi_{i1}} \right) \text{ for all } i. \quad (16.3)$$

16.46 Multiplying each of these equations by π_{i1} and summing over i yields

$$\mu = \mathbf{E}_q \left(\frac{1}{u'(t_i^{**})} \right) > 0,$$

where E_q denotes the expectation operator with respect to the distribution of outputs induced by effort $e = 1$.

Thus, the IR condition is binding, that is, $E_q(u(t_i^{**})) = \psi$.

16.47 Multiplying the equation above by $\pi_{i1}u(t_i^{**})$, summing all these equations over i , and taking into account the expression of μ obtained above yields

$$\lambda \left(\sum_{i=1}^n (\pi_{i1} - \pi_{i0}) u(t_i^{**}) \right) = \sum_{i=1}^n \pi_{i1} u(t_i^{**}) \left(\frac{1}{u'(t_i^{**})} - \mu \right) = \sum_{i=1}^n \pi_{i1} [u(t_i^{**}) - \psi] \left(\frac{1}{u'(t_i^{**})} - \mu \right)$$

16.48 Using the slackness condition $\lambda(\sum_{i=1}^n (\pi_{i1} - \pi_{i0}) u(t_i^{**}) - \psi) = 0$, we have

$$\lambda \psi = \text{cov} \left(u(t_i^{**}), \frac{1}{u'(t_i^{**})} \right).$$

By assumption, u and u' covary in opposite directions. Moreover, a constant wage $t_i^{**} = t^{**}$ for all i does not satisfy the incentive constraint, and thus t_i^{**} cannot be constant everywhere. Hence, the right-hand side is necessarily strictly positive. Thus we have $\lambda > 0$, and the IC condition is binding.

16.49 Coming back to Equation (16.3), we observe that the left-hand side is increasing in t_i^{**} since u is concave. For t_i^{**} to be nondecreasing with i , MLRP must again hold. Then higher outputs are also those that are the more informative ones about the realization of a high effort. Hence, the agent should be more rewarded as output increases.

Chapter 17

Social choice theory

Contents

17.1 Social choice	301
17.2 Arrow's impossibility theorem	302
17.3 Borda count, simple plurality rule, and two-round system	306
17.4 Gibbard-Satterthwaite theorem	309

17.1 Social choice

17.1 A society, denoted by $\langle X, N, (\succsim_i) \rangle$, consists of

- Non-empty set X of mutually exclusive social states (or alternatives). Although X could be infinite, we focus on the case that X is finite.
- Set $N = \{1, 2, \dots, n\}$ of individuals, where $n \geq 2$.
- Each individual i has his own preference \succsim_i over the set of social states. Let \mathcal{L} denote the set of preferences on X .

17.2 To determine the social choice, we will need some ranking of the social states in X that reflects society's preferences. Ideally, we would like to be able to compare any two alternatives in X from a social point of view, and we would like those binary comparisons to be consistent in the usual way.

A social preference relation, \succsim , is a complete and transitive binary relation on the set X of social states. For x and y in X , we read $x \succsim y$ as the statement “ x is socially at least as good as y .”

17.3 Issue 1: How can we go from the often divergent, but individually consistent, personal views of society's members to a single and consistent social view?

17.4 Condorcet's paradox: When we insist on transitivity as a criterion for consistency in social choice, certain well-known difficulties can easily arise.

A society of three individuals and three alternatives x, y and z . The preferences of the individuals are as follows:
The outcome is determined by majority voting.

Individual 1	Individual 2	Individual 3
x	y	z
y	z	x
z	x	y

Table 17.1


- (1) In a choice between x and y , x would get two votes and y would get one, so the social preference under majority rule would be $x \succ y$.
- (2) In a choice between y and z , majority voting gives $y \succ z$.
- (3) Because $x \succ y$ and $y \succ z$, transitivity of social preferences would require that $x \succ z$.
- (4) In a choice between y and z , majority voting gives $z \succ x$, which violates transitivity.

Hence, no single best alternative can be determined by majority rule.

17.2 Arrow's impossibility theorem

17.5 Issue 2: How can we go from consistent individual views to a social view that is consistent and that also respects certain basic values on matters of social choice that are shared by members of the community?

That is, we can imagine our problem as one of finding a “rule,” or function, capable of aggregating and reconciling the different individual views represented by the individual preference relations \succsim_i into a single social preference relation \succsim satisfying certain ethical principles.

 17.6 Definition: A social welfare function F is a function from \mathcal{L}^n to \mathcal{L} .

17.7 Arrow has proposed a set of four conditions that might be considered minimal properties the social welfare function, F , should possess. They are as follows.

U. Unrestricted domain. The domain of F must include all possible combinations of individual preferences on X .

PE. Pareto efficiency. For any pair of alternatives x and y in X , if $x \succ_i y$ for all i , then $x \succ y$.

IIA. Independent of irrelevant alternatives. Let

$$\succsim = F(\succsim_1, \succsim_2, \dots, \succsim_n), \quad \succsim' = F(\succsim'_1, \succsim'_2, \dots, \succsim'_n),$$

and let x and y be any two alternatives in X . If each individual i ranks x versus y under \succsim_i the same way that he does under \succsim'_i , then the social ranking of x versus y is the same under \succsim and \succsim' .

D. Non-dictatorship. There is no individual i such that for all x and y in X , $x \succ_i y$ implies $x \succ y$ regardless of the preferences \succsim_j of all other individuals $j \neq i$.

17.8 Remark:

- Condition U says that F is able to generate a social preference ordering regardless of what the individuals' preference relations happen to be.

As we have seen before, this condition, together with the transitivity requirement on \succsim , rules out majority voting as an appropriate mechanism because it sometimes fails to produce a transitive social ordering when there are more than three alternatives to consider.

- Condition PE says society should prefer x to y if every single member of society prefers x to y .
Notice that this is a weak Pareto requirement because it does not specifically require the social preference to be for x if, say, all but one strictly prefer x to y , yet one person is indifferent between x and y .
- Condition IIA says that the social ranking of x and y should depend only on the individual rankings of x and y .
- Condition D says there should be no single individual who “gets his way” on every single social choice, regardless of the views of everyone else in society.

17.9 Arrow's impossibility theorem: If there are at least three social states, then there is no social welfare function F that simultaneously satisfies Conditions U, PE, IIA, and D.

17.10 *Proof.* The strategy of the proof is to show that conditions U, PE, and IIA imply the existence of a dictator.

- (1) Consider any social state, c . Suppose each individual places state c at the bottom of his ranking. By PE, the social ranking must place c at the bottom as well. See Table 17.2.

\succsim_1	\succsim_2	\cdots	\succsim_n	\succsim
*	*	\cdots	*	*
\vdots	\vdots		\vdots	\vdots
\vdots	\vdots		\vdots	\vdots
c	c	\cdots	c	c

Table 17.2

- (2) (i) Imagine now moving c to the top of individual 1's ranking, leaving the ranking of all other states unchanged.
(ii) Next, do the same with individual 2: move c to the top of 2's ranking.
(iii) Continue doing this one individual at a time, keeping in mind that each of these changes in individual preferences might have an effect on the social ranking.
(iv) Eventually, c will be at the top of every individual's ranking, and so it must then also be at the top of the social ranking by Condition PE.

Consequently, there must be a first time during this process that the social ranking of c increases. Let individual m be the first such that raising c to the top of his ranking causes the social ranking of c to increase.

- (3) We claim that, as shown in Table 17.3, when c moves to the top of individual m 's ranking, the social ranking of c not only increases but c also moves to the top of the social ranking.

\succsim_1	\succsim_2	\cdots	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	\succsim
c	c	\cdots	c	*	\cdots	*	c
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
*	*	\cdots	*	c	\cdots	c	*

Table 17.3

To see this, assume by way of contradiction that the social ranking of c increases, but not to the top; *i.e.*, $\alpha \succsim c$ and $c \succsim \beta$ for some states $\alpha, \beta \neq c$. See Table 17.4.

\succsim_1	\succsim_2	\cdots	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	\succsim
c	c	\cdots	c	$*$	\cdots	$*$	$*$
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
\vdots	\vdots		\vdots	\vdots		\vdots	α
α, β	α, β		α, β	α, β		α, β	c
\vdots	\vdots		\vdots	\vdots		\vdots	β
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
$*$	$*$	\cdots	$*$	c	\cdots	c	$*$

Table 17.4

Now, because c is either at the bottom or at the top of every individual's ranking, we can change each individual i 's preferences so that $\beta \succ_i \alpha$, while leaving the position of c unchanged for that individual. See Table 17.5.

\succsim_1	\succsim_2	\cdots	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	\succsim
c	c	\cdots	c	$*$	\cdots	$*$	$*$
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
β	β		β	β		β	α
\vdots	\vdots		\vdots	\vdots		\vdots	c
α	α		α	α		α	β
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
$*$	$*$	\cdots	$*$	c	\cdots	c	$*$

Table 17.5

On one hand, $\beta \succ_i \alpha$ for every individual implies by PE that β must be strictly preferred to α according to the social ranking; *i.e.*, $\beta \succ \alpha$.

On the other hand, because the rankings of c relative to α and of c relative to β have not changed in any individual's ranking (see Tables 17.4 and 17.5), IIA implies that the social rankings of c relative to α and of c relative to β must be unchanged; *i.e.*, as initially assumed, we must have $\alpha \succsim c$ and $c \succsim \beta$. But transitivity then implies $\alpha \succsim \beta$, contradicting $\beta \succ \alpha$. This establishes our claim that c must have moved to the top of the social ranking.

- (4) Consider now any two distinct social states a and b , each distinct from c . In Table 17.3, change the profile of preferences as follows: change individual m 's ranking so that $a \succ_m c \succ_m b$, and for every other individual rank a and b in any way so long as the position of c is unchanged for that individual. See Table 17.6.

\succsim_1	\succsim_2	\dots	\succsim_m	\succsim_{m+1}	\dots	\succsim_n	\succsim
c	c	\dots	$*$	$*$	\dots	$*$	\dots
\vdots	\vdots		\vdots	\vdots		\vdots	a
\vdots	\vdots		a	\vdots		\vdots	\vdots
a, b	a, b		c	a, b		a, b	c
\vdots	\vdots		b	\vdots		\vdots	\vdots
\vdots	\vdots		\vdots	\vdots		\vdots	b
$*$	$*$	\dots	$*$	c	\dots	c	\vdots

Table 17.6

Note that in the new profile of preferences the ranking of a to c is the same for every individual as it was just before raising c to the top of individual m 's ranking in Step (2). Therefore, by IIA, the social ranking of a and c must be the same as it was at that moment, see Table 17.7. But this means that $a \succ c$ because at that moment c was still at the bottom of the social ranking.

\succsim_1	\succsim_2	\dots	\succsim_{m-1}	\succsim_m	\dots	\succsim_n	\succsim
c	c	\dots	c	$*$	\dots	$*$	$*$
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
a	a		a	a		a	a
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
$*$	$*$	\dots	$*$	c	\dots	c	c

Table 17.7

Similarly, in the new profile of preferences, the ranking of c to b is the same for every individual as it was just after raising c to the top of individual m 's ranking in Step (2), see Table 17.8. Therefore by IIA, the social ranking of c and b must be the same as it was at that moment. But this means that $c \succ b$ because at that moment c had just risen to the top of the social ranking.

\succsim_1	\succsim_2	\dots	\succsim_m	\succsim_{m+1}	\dots	\succsim_n	\succsim
c	c	\dots	c	$*$	\dots	$*$	c
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
b	b		b	b		b	b
\vdots	\vdots		\vdots	\vdots		\vdots	\vdots
$*$	$*$	\dots	$*$	c	\dots	c	$*$

Table 17.8

Therefore, as in Table 17.6, because $a \succ c$ and $c \succ b$, we may conclude by transitivity that $a \succ b$. Note then that no matter how the others rank a and b , the social ranking agrees with individual m 's ranking. By IIA, and

because a and b were arbitrary, we may therefore conclude that for all social states a and b distinct from c

$$a \succ_m b \text{ implies } a \succ b.$$

That is, individual m is a dictator on all pairs of social states not involving c .

(5) The final step shows that individual m is in fact a dictator.

Let a be distinct from c . We may repeat the above steps with a playing the role of c to conclude that some individual is a dictator on all pairs not involving a .

However, recall that individual m 's ranking of c (bottom or top) in Table 17.3 affects the social ranking of c (bottom or top). Hence, it must be individual m who is the dictator on all pairs not involving a . Because a was an arbitrary state distinct from c , and together with our previous conclusion about individual m , this implies that m is a dictator.

□

17.11 Although Arrow's theorem is a mathematical result, it is often expressed in a non-mathematical way with a statement such as "No voting method is fair", "Every ranked voting method is flawed", or "The only voting method that isn't flawed is a dictatorship".

More importantly, Arrow's theorem says that a deterministic preferential voting mechanism—that is, one where a preference order is the only information in a vote, and any possible set of votes gives a unique result—can not comply with all of the conditions given above simultaneously.

17.3 Borda count, simple plurality rule, and two-round system

17.12 Borda count.

The Borda count is commonly used for making collective choices. Individual i assigns a Borda count, $B_i(x)$, to every alternative x , where $B_i(x)$ is the number of alternatives in X to which x is preferred by agent i . Alternatives are then ranked according to their total Borda count as follows:

$$x \succ y \iff \sum_{i=1}^n B_i(x) \geq \sum_{i=1}^n B_i(y).$$

17.13 Example.

Individual 1	Individual 2	Individual 3
x	y	x
y	z	z
z	x	y

Table 17.9

x , y and z have 4, 3 and 2 points respectively.

17.14 Example: Consider 100 individuals who can be broken down into three groups based on their preferences over three alternatives, x , y and z .

40%	24%	36%
x	y	z
y	z	y
z	x	x

Table 17.10

x gets $40 \times 2 = 80$ points, y gets $24 \times 2 + 76 \times 1 = 124$ points and z gets $36 \times 2 + 24 \times 1 = 96$ points. With this procedure y wins with z in the second place and x in the third place.

17.15 Borda count satisfies U, PE and D, but does not satisfy IIA.

17.16 Because it sometimes elects broadly acceptable candidates, rather than those preferred by the majority, the Borda count is often described as a consensus-based electoral system, rather than a majoritarian one.

17.17 Borda count was developed independently several times, but is named for the 18-th century French mathematician and political scientist Jean-Charles de Borda, who devised the system in 1770.

17.18 Reversal paradox.

Consider seven individuals and four alternatives $\{x, y, z, w\}$. The preferences are:

Individual 1	Individual 2	Individual 3	Individual 4	Individual 5	Individual 6	Individual 7
x	z	x	y	z	z	y
y	x	y	w	x	x	w
w	y	w	z	y	y	z
z	w	z	x	w	w	x

Table 17.11

Total points: $x : 12$, $y : 13$, $z : 11$ and $w : 6$. So, y is the winner and the social ranking is $y \succ x \succ z \succ w$.

If the worst alternative w is eliminated, then the rankings are:

Individual 1	Individual 2	Individual 3	Individual 4	Individual 5	Individual 6	Individual 7
x	z	x	y	z	z	y
y	x	y	z	x	x	z
z	y	z	x	y	y	x

Table 17.12

Total points: $x : 7$, $y : 6$ and $z : 8$. So, the winner is z and the ranking is $z \succ x \succ y$.

The social ranking is completely reversed! (Reason: failure of IIA.)

17.19 Simple plurality rule.

In this system the single winner is the person with the most votes (plurality); there is no requirement that the winner gain an absolute majority of votes, but rather only a plurality, sometimes called a relative/simple majority.

17.20 Example: Consider 100 individuals who can be broken down into three groups based on their preferences over three alternatives, x , y and z .

40%	24%	36%
x	y	z
y	z	y
z	x	x

Table 17.13

With simple plurality rule, x gets 40%, y gets 25% and z gets 36%. So, x is the winner, although 60% rank it the lowest!

17.21 Two-round system.

The two-round system (also known as the second ballot, runoff voting or ballotage) is a voting system used to elect a single winner where the voter casts a single vote for their chosen candidate. However, if no candidate receives the required number of votes (usually an absolute majority or 40–45% with a winning margin of 5–15%), then those candidates having less than a certain proportion of the votes, or all but the two candidates receiving the most votes, are eliminated, and a second round of voting occurs.

17.22 Example.

40%	24%	36%
x	y	z
y	z	y
z	x	x

Table 17.14

In the first round, x gets 40%, y gets 24% and z gets 36%. If y is eliminated, in the second round x gets 40% and z gets 60% of the votes and z is the winner.

17.23 Outcomes under different methods.

40%	24%	36%
x	y	z
y	z	y
z	x	x

Table 17.15

- Simple plurality rule: x is the best.
- Borda count: y is the best.
- Two-round system: z is the best.

17.4 Gibbard-Satterthwaite theorem

17.24 We have focused solely on the task of aggregating the preference profile into a single preference for society. This task, as we have seen, is a formidable one. Indeed, it can not be carried out if we insist on all of Arrow's conditions.

17.25 Issue 3: Maybe Arrow's impossibility theorem held because we required a whole preference ordering. So social choice functions might be easier to find. Firstly We will need to redefine our criteria for the social choice function setting; PE and IIA discussed the ordering.

17.26 Issue 4: Implicit in our analysis has been the assumption that the true preferences of each individual can be obtained and that society's preferences are then determined according to its social welfare function. But how, exactly, does society find out the preferences of its individual members?

One possibility, of course, is to simply ask each individual to report his ranking of the social states. But this introduces a serious difficulty. Individuals would be better off lying about their preferences than reporting them truthfully if a false report leads to a better social state for them.

Thus, in addition to the problem of coherently aggregating individual rankings into a social ranking, there is the problem of finding out individual preferences in the first place.

17.27 Definition: A social choice function is a function $f: \mathcal{L}^n \rightarrow X$. Specifically, for each preference profile $(\succsim_1, \succsim_2, \dots, \succsim_n)$, $f(\succsim_1, \succsim_2, \dots, \succsim_n)$ is the society's choice from X .

17.28 Definition: A social choice function f is dictatorial if there is an individual i such that whenever $f(\succsim_1, \succsim_2, \dots, \succsim_n) = x$ it is the case that $x \succsim_i y$ for every $y \in X$.

17.29 Definition: A social choice function f is Pareto efficient if $f(\succsim_1, \succsim_2, \dots, \succsim_n) = x$ whenever $x \succ_i y$ for every individual i and every $y \in X$ distinct from x .

17.30 Definition: A social choice function f is monotonic if whenever $f(\succsim_1, \succsim_2, \dots, \succsim_n) = x$ and for every individual i and every alternative y the preference \succsim'_i ranks x above y if \succsim_i does, then $f(\succsim'_1, \succsim'_2, \dots, \succsim'_n) = x$.

An alternative x must remain the winner whenever the support for it is increased in a preference profile under which x was already winning.

17.31 Definition: A social choice function f is strategy-proof when, for every individual, i , for every pair \succsim_i and \succsim'_i of his preferences, and for every profile \succsim_{-i} of others' preferences, we have

$$f(\succsim_i, \succsim_{-i}) \succsim_i f(\succsim'_i, \succsim_{-i}).$$

17.32 Example.

Individual 1	Individual 2	Individual 3	Individual 4
x	z	w	x
y	x	y	z
z	w	z	w
w	y	x	y

Table 17.16

Apply Borda count. Total points: $x : 8$, $y : 4$, $z : 7$ and $w : 5$. So, x is the best.

Individuals 1 and 4 will tell the truth since the winner x is their top choices.

If individual 2 reports the preference $z \succ_2 w \succ_2 y \succ_2 x$ as shown in the following table, the total points are: $x : 6$, $y : 5$, $z : 7$ and $w : 6$, and hence z is the best. That is, individual 2 prefers z to x in the original ranking and would like to lie.

Individual 1	Individual 2	Individual 3	Individual 4
x	z	w	x
y	w	y	z
z	y	z	w
w	x	x	y

Table 17.17

If individual 3 reports the preference $z \succ_3 w \succ_3 y \succ_3 x$ as shown in the following table, the total points are: $x : 8$, $y : 3$, $z : 9$ and $w : 4$, and hence z is the best. That is, individual 3 prefers z to x in the original ranking and would like to lie.

Individual 1	Individual 2	Individual 3	Individual 4
x	z	z	x
y	x	w	z
z	w	y	w
w	y	x	y

Table 17.18

17.33 Lemma: Suppose that f is a monotonic social choice function and that $f(\succsim_1, \dots, \succsim_n) = x$, where $\succsim_1, \dots, \succsim_n$ are each strict rankings of the social states in X .

- (i) Suppose that for some individual i , \succsim_i ranks y just below x , and let \succsim'_i be identical to \succsim_i except that y is ranked just above x , i.e., the ranking of x and y is reversed. Then either $f(\succsim'_i, \succsim_{-i}) = x$ or $f(\succsim'_i, \succsim_{-i}) = y$.
- (ii) Suppose that $\succsim'_1, \dots, \succsim'_n$ are strict rankings such that for every individual i , the ranking of x versus any other social state is the same under \succsim'_i as it is under \succsim_i . Then $f(\succsim'_1, \dots, \succsim'_n) = x$.

Proof. (i) Suppose that $f(\succsim'_i, \succsim_{-i}) = z \neq x, y$, then we have for every $w \in X$ $z \succ_i w$ whenever $z \succ'_i w$, and for every $j \neq i$ and every $w \in X$ $z \succ_j w$ whenever $z \succ'_j w$. Then by monotonicity $f(\succsim_i, \succsim_{-i}) = z \neq x$, a contradiction.

(ii) Routine. □

17.34 Lemma: Let f be a monotonic social choice function and suppose that the social choice must be x whenever all individual rankings are strict and x is at the top of individual m 's ranking. Then the social choice must be at least as good as x for individual m when the individual rankings are not necessarily strict and x is at least as good for individual m as any other social states.

Proof. By Lemma 17.33. Argue by contradiction and change preferences monotonically so that all preferences are strict and x is at the top of m 's ranking. □

17.35 Lemma: Let x and y be distinct social states. Suppose that the social choice is at least as good as x for individual i whenever x is at least as good as every other social state for i . Suppose also that the social choice is at least as good as y for individual j whenever y is at least as good as every other social state for j . Then $i = j$.

Proof. □

📌 17.36 **Gibbard-Satterthwaite theorem:** If there are at least three social states, then every onto strategy-proof social choice function is dictatorial.

This theorem is named after Allan Gibbard and Mark Satterthwaite.

17.37 *Part 1: strategy-proofness implies monotonicity.* Let $(\succsim_1, \succsim_2, \dots, \succsim_n)$ be an arbitrary preference profile and suppose that $f(\succsim_1, \dots, \succsim_n) = x$. Fix an individual, i say, and let \succsim'_i be a preference for i such that for every $y \in X$ distinct from x , $x \succ'_i y$ if $x \succsim_i y$. We shall show that $f(\succsim'_i, \succsim_{-i}) = x$.

- (1) Suppose, by way of contradiction, that $f(\succsim'_i, \succsim_{-i}) = y \neq x$.
- (2) Given that the others report \succsim_{-i} , individual i , when his preferences are \succsim_i can report truthfully and obtain the social state x or he can lie by reporting \succsim'_i and obtain the social state y . Strategy-proofness requires that lying can not be strictly better than telling the truth. Hence we must have $x \succsim_i y$.
- (3) According to the definition of \succsim'_i , we then have $x \succ'_i y$.
- (4) Consequently, when individual i 's preferences are \succsim'_i he strictly prefers x to y and so, given that the others report \succsim_{-i} , individual i strictly prefers lying (reporting \succsim_i and obtaining x) to telling the truth (reporting \succsim'_i and obtaining y), contradicting strategy-proofness.
- (5) We conclude that $f(\succsim'_i, \succsim_{-i}) = x$.

□

17.38 *Part 2: onto and monotonicity implies Pareto efficiency.* Let x be an arbitrary social state and let $(\succsim_i)_i$ be a preference profile with x at the top of each individual's ranking \succsim_i . We must show that $f(\succsim_1, \dots, \succsim_n) = x$.

- (1) Because f is onto, $f(\succsim'_1, \dots, \succsim'_n) = x$ for some $(\succsim'_1, \dots, \succsim'_n) \in \mathcal{L}^n$.
- (2) Obtain the preference profile $(\succsim''_i)_i$ from $(\succsim'_i)_i$ by moving x to the top of every individual's ranking \succsim'_i .
- (3) By monotonicity, $f(\succsim''_1, \dots, \succsim''_n) = x$.
- (4) Because $(\succsim_i)_i$ places x at the top of every individual ranking \succsim_i and $f(\succsim''_1, \dots, \succsim''_n) = x$, we can again apply monotonicity and conclude that $f(\succsim_1, \dots, \succsim_n) = x$, as desired.

□

17.39 *Part 3: $|X| \geq 3$, monotonicity and Pareto efficiency imply dictatorship.* (1) Consider any two distinct social states $x, y \in X$ and a profile of strict rankings in which x is ranked highest and y lowest for every individual i . Pareto efficiency implies that the social choice at this profile is x .

- (2) Consider now changing individual 1's ranking by strictly raising y in it one position at a time. By monotonicity, the social choice remains equal to x so long as y is below x in 1's ranking.
- (3) When y finally does rise above x , Lemma 17.33 implies that the social choice either changes to y or remains equal to x .

- (4) If the social choice is x , then begin the same process with individual 2, then 3, *etc.* until for some individual m , the social choice does change from x to y when y rises above x in m 's ranking. There must be such an individual m because y will eventually be at the top of every individual's ranking and by Pareto efficiency the social choice will then be y . Tables 17.19 and 17.20 depict the situations just before and just after individual m 's ranking of y is raised above x .

\succsim_1	\cdots	\succsim_{m-1}	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	Social choice
y	\cdots	y	x	x	\cdots	x	x
x	\cdots	x	y	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
$*$	\cdots	$*$	$*$	y	\cdots	y	

Table 17.19

\succsim_1	\cdots	\succsim_{m-1}	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	Social choice
y	\cdots	y	y	x	\cdots	x	y
x	\cdots	x	x	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
$*$	\cdots	$*$	$*$	y	\cdots	y	

Table 17.20

- (5) Consider Tables 17.21 and 17.22 below. Table 17.21 is derived from Table 17.19 (and Table 17.22 from Table 17.20) by moving x to the bottom of individual i 's ranking for $i < m$ and moving x to the second last position in i 's ranking for $i > m$.

\succsim_1	\cdots	\succsim_{m-1}	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	Social choice
y	\cdots	y	x	$*$	\cdots	$*$	x
\vdots		\vdots	y	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
\vdots		\vdots	\vdots	x	\cdots	x	
x	\cdots	x	$*$	y	\cdots	y	

Table 17.21

\succsim_1	\cdots	\succsim_{m-1}	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	Social choice
y	\cdots	y	y	$*$	\cdots	$*$	y
\vdots		\vdots	x	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
\vdots		\vdots	\vdots	x	\cdots	x	
x	\cdots	x	$*$	y	\cdots	y	

Table 17.22

- (6) We wish to argue that these changes do not affect the social choices, *i.e.*, that the social choices are as indicated in the tables.

Note that the social choice in Table 17.22 must, by Lemma 17.33, be y because the social choice in Table 17.20 is y and no individual's ranking of y versus any other social state changes in the move from Table 17.20 to Table 17.22.

Note that the preference profiles in Tables 17.21 and 17.22 differ only in individual m 's ranking of x and y . By Lemma 17.33, the social choice in Table 17.21 must be either x or y because the social choice in Table 17.22 is y . When the social choice in Table 17.21 is y , by Lemma 17.33, the social choice in Table 17.19 must be y , a contradiction.

- (7) Because there are at least three social states, we may consider a social state $z \in X$ distinct from x and y . Since the (otherwise arbitrary) profile of strict rankings in Table 17.23 can be obtained from the Table 17.21 profile without changing the ranking of x versus any other social state in any individual's ranking, the social choice in Table 17.23 must, by Lemma 17.33, be x .

\succsim_1	\cdots	\succsim_{m-1}	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	Social choice
$*$	\cdots	$*$	x	$*$	\cdots	$*$	x
\vdots		\vdots	z	\vdots		\vdots	
\vdots		\vdots	y	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
z	\cdots	z	\vdots	z	\cdots	z	
y	\cdots	y	\vdots	x	\cdots	x	
x	\cdots	x	$*$	y	\cdots	y	

Table 17.23

- (8) Consider the profile of rankings in Table 17.24 derived from the Table 17.23 profile by interchanging the ranking of x and y for individuals $i > m$.

Because this is the only difference between the profiles in Tables 17.23 and 17.24, and because the social choice in Table 17.23 is x , the social choice in Table 17.24 must, by Lemma 17.33, be either x or y .

But the social choice in Table 17.24 can not be y because z is ranked above y in every individual's Table 17.24 ranking, and monotonicity would then imply that the social choice would remain y even if z were raised to the top of every individual's ranking, contradicting Pareto efficiency.

Hence the social choice in Table 17.24 is x .

\succsim_1	\cdots	\succsim_{m-1}	\succsim_m	\succsim_{m+1}	\cdots	\succsim_n	Social choice
*	\cdots	*	x	*	\cdots	*	x
\vdots		\vdots	z	\vdots		\vdots	
\vdots		\vdots	y	\vdots		\vdots	
\vdots		\vdots	\vdots	\vdots		\vdots	
z	\cdots	z	\vdots	z	\cdots	z	
y	\cdots	y	\vdots	y	\cdots	y	
x	\cdots	x	*	x	\cdots	x	

Table 17.24

- (9) Note that an arbitrary profile of strict rankings with x at the top of individual m 's ranking can be obtained from the profile in Table 17.24 without reducing the ranking of x versus any other social state in any individual's ranking. Hence, Lemma 17.33 implies that the social choice must be x whenever individual rankings are strict and x is at the top of individual m 's ranking.

Lemma 17.34 implies that even when individual rankings are not strict and indifferences are present, the social choice must be at least as good as x for individual m whenever x is at least as good as every other social state for individual m .

- (10) So, we may say that individual m is a dictator for the social state x . Because x was arbitrary, we have shown that for each social state $x \in X$, there is a dictator for x . But Lemma 17.35 implies there can not be distinct dictators for distinct social states. Hence there is a single dictator for all social states and therefore the social choice function is dictatorial.

□

17.40 Proposition: A social choice function f is strongly monotonic if whenever $f(\succsim_1, \dots, \succsim_n) = x$ and for every individual i and every alternative y the preference \succsim'_i ranks x above y if \succsim_i does, then $f(\succsim'_1, \dots, \succsim'_n) = x$.

Suppose there are two individuals, 1 and 2, and three social states, x , y , and z . Define the social choice function f to choose individual 1's top-ranked social state unless it is not unique, in which case the social choice is individual 2's top-ranked social state among those that are top-ranked for individual 1, unless this too is not unique, in which case, among those that are top-ranked for both individuals, choose x if it is among them, otherwise choose y .

(i) f is strategy-proof.

(ii) Show by example that f is not strongly monotonic. (Hence, strategy-proofness does not imply strong monotonicity, even though it implies monotonicity.)

17.41 Proposition: Show that if f is an onto monotonic social choice function and the finite set of social states is X , then for every $x \in X$ there is a profile, $(\succsim_i)_i$, of strict rankings such that $f(\succsim_1, \dots, \succsim_n) = x$.

17.42 Proposition: Show that when there are just two alternatives and an odd number of individuals, the majority rule social choice function (*i.e.*, that which chooses the outcome that is the top ranked choice for the majority of individuals) is Pareto efficient, strategy-proof and non-dictatorial.

Dominant strategy mechanism design

Contents

18.1 Envelope theorem	315
18.2 A general mechanism design setting	317
18.3 Revelation principle	319
18.4 Payoff equivalence	320
18.5 VCG mechanism	321
18.5.1 Pivot mechanism	323
18.5.2 Balancing the budget	326
18.6 Characterization of dominant strategy incentive compatibility	327

18.1 The theory of mechanism design can be thought of as the “engineering” side of economic theory. Much theoretical work focuses on existing economic institutions. The theorist wants to explain or forecast the economic or social outcomes that these institutions generate.

But in mechanism design theory the direction of inquiry is reversed. We begin by identifying our desired social goal. We then ask whether or not an appropriate institution (mechanism) could be designed to attain that goal.

Almost any kind of market institution or economic organization can be viewed, in principle, as a mechanism. Examples include: school choice, auction, kidney exchange, tax codes, contract design, *etc.*

18.2 Leonid Hurwicz defined a mechanism as a communication system in which participants send messages to each other and/or to a “message center”, and where a pre-specified rule assigns an outcome (such as an allocation of goods and services) for every collection of received messages.

The difficulty in mechanism design is that the individuals have private information and different objectives, and so may not have the incentive to behave in a way that reveals what they know. The key point is how to design “incentive compatible” mechanisms that can generate the information needed as they are executed.

18.1 Envelope theorem

18.3 Consider a one-agent decision problem

$$V(\theta) = \max_{a \in A} h(a, \theta),$$

where a is the agent's chosen action and $\theta \in \Theta$ an exogenous parameter. In an auction, θ could be bidder's valuation, and a bidder's choice of bid.

18.4 A could be either discrete or continuous, but Θ is an interval.

18.5 Let $a^*(\theta)$ be the set of optimal choices, that is,

$$a^*(\theta) = \arg \max_{a \in A} h(a, \theta).$$

Let h_a and h_θ denote partial derivatives of h .

18.6 Theorem (Envelope theorem): Suppose for all $\theta \in \Theta$, $a^*(\theta)$ is non-empty, and for all a and θ , h_θ exists. Let $a(\theta)$ be any selection from $a^*(\theta)$.

(i) If V is differentiable at θ , then

$$V'(\theta) = h_\theta(a(\theta), \theta).$$

(ii) If V is absolutely continuous, then for any $\theta' > \theta$,

$$V(\theta') - V(\theta) = \int_{\theta}^{\theta'} h_\theta(a(t), t) dt.$$

Proof of (i). (1) If V is differentiable at θ , then

$$V'(\theta) = \lim_{\epsilon \downarrow 0} \frac{V(\theta + \epsilon) - V(\theta)}{\epsilon} = \lim_{\epsilon \downarrow 0} \frac{V(\theta) - V(\theta - \epsilon)}{\epsilon}.$$

(2) Take $a(\theta) \in a^*(\theta)$, then $V(\theta) = h(a(\theta), \theta)$, and

$$V(\theta + \epsilon) = \max_a h(a, \theta + \epsilon) \geq h(a(\theta), \theta + \epsilon).$$

(3) Then we have

$$V'(\theta) = \lim_{\epsilon \downarrow 0} \frac{V(\theta + \epsilon) - V(\theta)}{\epsilon} \geq \lim_{\epsilon \downarrow 0} \frac{h(a(\theta), \theta + \epsilon) - h(a(\theta), \theta)}{\epsilon} = h_\theta(a(\theta), \theta).$$

(4) For the same number $a(\theta)$, $V(\theta - \epsilon) = \max_a h(a, \theta - \epsilon) \geq h(a(\theta), \theta - \epsilon)$, and hence

$$V'(\theta) = \lim_{\epsilon \downarrow 0} \frac{V(\theta) - V(\theta - \epsilon)}{\epsilon} \leq \lim_{\epsilon \downarrow 0} \frac{h(a(\theta), \theta) - h(a(\theta), \theta - \epsilon)}{\epsilon} = h_\theta(a(\theta), \theta).$$

(5) So

$$h_\theta(a(\theta), \theta) \leq V'(\theta) \leq h_\theta(a(\theta), \theta).$$

□

Proof of (ii). (1) Absolute continuity: for all $\epsilon > 0$, there exists $\delta > 0$ such that for any finite, disjoint set of intervals $\{[x_k, y_k]\}_{k=1,2,\dots,M}$ with $\sum_k |y_k - x_k| < \delta$,

$$\sum_k |V(y_k) - V(x_k)| < \epsilon.$$

- (2) Absolute continuity is equivalent to V being differentiable almost everywhere and being the integral of its derivative, so the second part follows directly from the first part.

□

18.7 Remark: The derivative of the value function is the derivative of the objective function, evaluated at the maximizer.

18.8 Corollary: Assume that

- for each $a \in A$, $h(a, \cdot)$ is differentiable,
- there exists $B > 0$, such that for all $a \in A$ and almost all $\theta \in \Theta$

$$|h_\theta(a, \theta)| \leq B,$$

- $a^*(\theta) = \arg \max_{a \in A} h(a, \theta) \neq \emptyset$.

Then V is Lipschitz continuous with Lipschitz constant 1, and hence absolutely continuous and almost everywhere differentiable. Therefore the two formulas in Theorem 18.6 still hold.

Proof. For any two distinct θ and θ' , we have

$$|V(\theta) - V(\theta')| = \left| \max_{a \in A} h(a, \theta) - \max_{a \in A} h(a, \theta') \right| \leq \max_{a \in A} |h(a, \theta) - h(a, \theta')| \leq \max_{a \in A} B \cdot |\theta - \theta'| = B \cdot |\theta - \theta'|.$$

□

18.2 A general mechanism design setting

18.9 Mechanism design theory distinguishes sharply between the apparatus under the control of the designer, which we call a mechanism, and the world of things that are beyond the designer's control, which we call the environment.

18.10 An environment comprises three lists:

- a list of participants or potential participants,
- a list of the possible outcomes,
- a list of the participants' possible types—that is, their capabilities, preferences, information, and beliefs.

A mechanism consists of rules that govern what the participants are permitted to do, and how these permitted actions determine outcomes.

18.11 Mechanism theory evaluates alternative designs based on their comparative performance. Formally, performance is the function that maps environments into outcomes.

The goal of mechanism design analysis is to determine what performance is possible and how mechanism can best be designed to achieve the designer's goals. Mechanism design addresses three common questions:

- Is it possible to achieve a certain kind of performance, for instance a map that picks an efficient allocation for every possible environment in some class?
- What is the complete set of performance functions that are implementable by some mechanism?
- What mechanism optimizes performance according to the mechanism designer's performance criterion?

18.12 Setup:

- There are N agents. The set of agents is denoted by $\mathcal{N} = \{1, 2, \dots, N\}$.
- The set of potential social decisions is denoted by D .
- Agent i 's information is represented by a type θ_i which lies in a set Θ_i . Let $\theta = (\theta_1, \theta_2, \dots, \theta_N)$, and $\Theta = \Theta_1 \times \Theta_2 \times \dots \times \Theta_N$.
- Agents have preferences over decisions that are represented by a utility function. Agent i 's utility if decision d is chosen, and agent i pays transfer t_i is:

$$v_i(d, \theta_i) - t_i.$$

18.13 A decision rule is a mapping $d: \Theta \rightarrow D$.

A decision rule $d(\cdot)$ is efficient if

$$\sum_i v_i(d(\theta), \theta_i) \geq \sum_i v_i(d', \theta_i) \text{ for all } \theta \in \Theta \text{ and } d' \in D,$$

that is

$$d(\theta) \in \arg \max_{d' \in D} \sum_i v_i(d', \theta_i) \text{ for all } \theta \in \Theta.$$

18.14 Agent i 's transfer function is a mapping $t_i: \Theta \rightarrow \mathbb{R}$. $t_i(\theta)$ represents the payment that i receives based on the announcement of types θ . Let $t(\theta) = (t_1(\theta), t_2(\theta), \dots, t_N(\theta))$.

A transfer function t is said to be feasible if $\sum_i t_i(\theta) \geq 0$ for all θ .

A transfer function t is said to be balanced if $\sum_i t_i(\theta) = 0$ for all θ . (d, t) satisfies budget balance if the transfer function is balanced.

18.15 A pair (d, t) will be referred to as a social choice function.

18.16 The utility that i receives, if θ' is the announced vector of types, and i 's true type is θ_i , is

$$v_i(d(\theta'), \theta_i) - t_i(\theta').$$

18.17 A mechanism is a pair (M, g) , where

- $M = M_1 \times M_2 \times \dots \times M_N$ is a cross product of message spaces.
- $g: M \rightarrow D \times \mathbb{R}^N$ is an outcome function.

18.18 Note that a social choice function (d, t) can be viewed as a mechanism, where $M_i = \Theta_i$ and $g = (d, t)$. This is referred to as a direct mechanism.

18.19 A social choice function (d, t) is implemented by a mechanism (M, g) if there exists a function $m = (m_1, m_2, \dots, m_N)$, where each $m_i: \Theta_i \rightarrow M_i$, such that

$$(d(\theta), t(\theta)) = (g \circ m)(\theta) \text{ for all } \theta.$$




Figure 18.1

We always require that m satisfies some equilibrium condition, for example m is a dominant strategy profile or a Bayesian Nash equilibrium.

18.3 Revelation principle

18.20 A strategy $m_i \in M_i$ is a dominant strategy at θ_i , if

$$v_i(g_d(m_i, m_{-i}), \theta_i) - g_{t_i}(m_i, m_{-i}) \geq v_i(g_d(m'_i, m_{-i}), \theta_i) - g_{t_i}(m'_i, m_{-i})$$

for all m_{-i} and m'_i .

18.21 A social choice function (d, t) is implemented in dominant strategies by the mechanism (M, g) if there exists functions $m_i: \Theta_i \rightarrow M_i$ such that

- $m_i(\theta_i)$ is a dominant strategy for each i and $\theta_i \in \Theta_i$,
- $g(m(\theta)) = (d, t)(\theta)$ for all $\theta \in \Theta$.

18.22 A direct mechanism $f = (d, t)$ is dominant strategy incentive compatible if θ_i is a dominant strategy at θ_i for each i and $\theta_i \in \Theta_i$.

That is, for all θ_i, θ'_i and θ_{-i} ,

$$v_i(d(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) \geq v_i(d(\theta'_i, \theta_{-i}), \theta_i) - t_i(\theta'_i, \theta_{-i}).$$

18.23 Theorem (Revelation principle for dominant strategies): If a mechanism (M, g) implements a social choice function (d, t) in dominant strategies, then the direct mechanism (d, t) is dominant strategy incentive compatible.

Proof. Note that $(d, t)(\theta) = g(m(\theta))$ for each θ .




Figure 18.2: Revelation principle

□

The powerful implication of the revelation principle is that if we wish to find out the social choice functions can implemented in dominant strategies, we can restrict our attention to the set of direct mechanisms.

18.4 Payoff equivalence

§ 18.24 Theorem (Payoff equivalence theorem): Suppose $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$, $v_i(d, \cdot)$ is differentiable, and there exists $B > 0$ such that for all d and θ , $|\frac{dv_i}{d\theta_i}(d, \theta_i)| \leq B$. If the direct mechanism $(d(\cdot), t(\cdot))$ is dominant strategy incentive compatible. Then for every θ ,

$$v_i(d(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) = v_i(d(\underline{\theta}_i, \theta_{-i}), \underline{\theta}_i) - t_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} \frac{dv_i}{d\theta_i}(d(s, \theta_{-i}), s) ds.$$

18.25 Proof.

- (1) Fixed θ_{-i} , let $h(\theta'_i, \theta_i) = v_i(d(\theta'_i, \theta_{-i}), \theta_i) - t_i(\theta'_i, \theta_{-i})$ and $V(\theta_i) = \max_{\theta'_i} h(\theta'_i, \theta_i)$.
- (2) Then we have $\frac{dh}{d\theta_i} = \frac{dv_i}{d\theta_i}$ for all θ_i .
- (3) Since (d, t) is dominant strategy incentive compatible, we have

$$\theta_i \in \arg \max_{\theta'_i} h(\theta'_i, \theta_i),$$

and hence $V(\theta_i) = h(\theta_i, \theta_i)$.

- (4) By Corollary 18.8, we have

$$V(\theta_i) - V(\underline{\theta}_i) = \int_{\underline{\theta}_i}^{\theta_i} \frac{dh}{d\theta_i}(t, t) dt,$$

that is,

$$v_i(d(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) = v_i(d(\underline{\theta}_i, \theta_{-i}), \underline{\theta}_i) - t_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} \frac{dv_i}{d\theta_i}(d(s, \theta_{-i}), s) ds.$$

□

18.26 Corollary: Suppose $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$, $v_i(d, \cdot)$ is differentiable, and there exists $B > 0$ such that for all d and θ , $|\frac{dv_i}{d\theta_i}(d, \theta_i)| \leq B$. If (d, t) and (d, t') are two dominant strategy incentive compatible, then for each i and each θ_{-i} , there exists $c_i(\theta_{-i}) \in \mathbb{R}$, such that for all θ_i

$$t_i(\theta_i, \theta_{-i}) - t'_i(\theta_i, \theta_{-i}) = c_i(\theta_{-i}).$$

Proof. By Theorem 18.24, we have

$$v_i(d(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) = v_i(d(\underline{\theta}_i, \theta_{-i}), \underline{\theta}_i) - t_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} \frac{dv_i}{d\theta_i}(d(s, \theta_{-i}), s) ds.$$

$$v_i(d(\theta_i, \theta_{-i}), \theta_i) - t'_i(\theta_i, \theta_{-i}) = v_i(d(\underline{\theta}_i, \theta_{-i}), \underline{\theta}_i) - t'_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} \frac{dv_i}{d\theta_i}(d(s, \theta_{-i}), s) ds.$$

Therefore,

$$t_i(\theta_i, \theta_{-i}) - t'_i(\theta_i, \theta_{-i}) = t_i(\underline{\theta}_i, \theta_{-i}) - t'_i(\underline{\theta}_i, \theta_{-i}) \triangleq c_i(\theta_{-i}).$$

□

18.27 Remark: Suppose $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$, $v_i(d, \cdot)$ is differentiable, and there exists $B > 0$ such that for all d and θ , $|\frac{dv_i}{d\theta_i}(d, \theta_i)| \leq B$. Suppose that (d, t) is a dominant strategy incentive compatible mechanism. Then a direct mechanism with the same decision rule, (d, t') , is dominant strategy incentive compatible if and only if for every $i \in \mathcal{N}$ and every $\theta_{-i} \in \Theta_{-i}$ there is a number $c_i(\theta_{-i}) \in \mathbb{R}$ such that

$$t'_i(\theta) = t_i(\theta) + c_i(\theta_{-i}) \text{ for all } \theta \in \Theta.$$

18.28 Theorem: Suppose that for every $i \in \mathcal{N}$ the set Θ_i is a convex subset of a finite-dimensional Euclidean space. Moreover, assume that for every $i \in \mathcal{N}$ the function $v_i(d, \theta_i)$ is a convex function of θ_i . Suppose that (d, t) is a dominant strategy incentive compatible mechanism. Then a direct mechanism with the same decision rule, (d, t') , is dominant strategy incentive compatible if and only if for every $i \in \mathcal{N}$ and every $\theta_{-i} \in \Theta_{-i}$ there is a number $c_i(\theta_{-i}) \in \mathbb{R}$ such that

$$t'_i(\theta) = t_i(\theta) + c_i(\theta_{-i}) \text{ for all } \theta \in \Theta.$$

Proof. The necessity part is obvious. See [Krishna and Maenner \(2001\)](#) for the sufficiency part. \square

18.5 VCG mechanism

18.29 Definition: A direct mechanism (d, t^{VCG}) is called a Vickrey-Clarke-Groves mechanism if d is an efficient decision rule, and if for every i there is a function

$$h_i: \Theta_{-i} \rightarrow \mathbb{R},$$

such that

$$t_i^{\text{VCG}}(\theta) = - \sum_{j \neq i} v_j(d(\theta), \theta_j) + h_i(\theta_{-i}) \text{ for all } \theta \in \Theta.$$

18.30 Exercise: How to derive transfer function of VCG mechanism?

18.31 In a VCG mechanism each agent i is paid the sum of the other agents' utility from the implemented alternative whereby utilities are calculated using the agents' reported types. This is the first term in the formula. This term aligned agent i 's interests with utilitarian welfare. The second term is a constant that depends on the other agents' reported types, and that does not affect agent i 's incentives. This constant can be used to raise the overall revenue from the mechanism.

18.32 Proposition: VCG mechanisms are dominant strategy incentive compatible.

Proof. (1) Consider any agent i and take θ_{-i} as given.

(2) If agent i is of type θ_i , and reports that she is of type θ'_i , then her utility is:

$$v_i(d(\theta'_i, \theta_{-i}), \theta_i) + \sum_{j \neq i} v_j(d(\theta'_i, \theta_{-i}), \theta_j) - h_i(\theta_{-i}) = \sum_{j \in \mathcal{N}} v_j(d(\theta'_i, \theta_{-i}), \theta_j) - h_i(\theta_{-i}).$$

(3) Note that $h_i(\theta_{-i})$ is not changed by agent i 's report. Only the first expression matters for i 's incentives.

(4) Since d is efficient, we have

$$\sum_{j \in \mathcal{N}} v_j(d(\theta_i, \theta_{-i}), \theta_j) \geq \sum_{j \in \mathcal{N}} v_j(d(\theta'_i, \theta_{-i}), \theta_j)$$

for all θ'_i .

(5) Therefore, it is optimal for agent i to report her true type. □

18.33 Remark: Every efficient social choice function can be truthfully implemented in a dominant strategy by a VCG mechanism.

18.34 Proposition: Suppose $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$, $v_i(d, \cdot)$ is differentiable, and there exists $B > 0$ such that for all d and θ , $|\frac{dv_i}{d\theta_i}(d, \theta_i)| \leq B$. Suppose that (d, t) is a dominant strategy incentive compatible mechanism, and suppose that d is efficient. Then (d, t) is a VCG mechanism.

Proof. Every dominant strategy incentive compatible mechanism that implements an efficient decision rule d must involve the same transfers as the VCG mechanism up to additive constants $c_i(\theta_{-i})$ that may be added to any agent i 's transfers. But adding such constants to a VCG mechanism yields by the definition of VCG mechanisms another VCG mechanism. □

18.35 Proposition: Suppose that for every i , the set Θ_i is a convex subset of a finite-dimensional Euclidean space. Moreover, assume that for every i the function $v_i(d, \theta_i)$ is a convex function of θ_i . Suppose that (d, t) is a dominant strategy incentive compatible mechanism, and suppose that d is efficient. Then (d, t) is a VCG mechanism.

18.36 Proposition: If d is an efficient decision rule, (d, t) is dominant strategy incentive compatible, and the type spaces are complete in the sense that

$$\{v_i(\cdot, \theta_i) \mid \theta_i \in \Theta_i\} = \{v: D \rightarrow \mathbb{R}\} \text{ for each } i,$$

then for each i there exists a function $h_i: \Theta_{-i} \rightarrow \mathbb{R}$ such that the transfer function t satisfies

$$t_i(\theta) = h_i(\theta_{-i}) - \sum_{j \neq i} v_j(d(\theta), \theta_j).$$

Proof. (1) Let d be efficient, (d, t) dominant strategy incentive compatible, and the type spaces complete.

(2) Note that for each i there exists a function $h_i: \Theta \rightarrow \mathbb{R}$ such that

$$t_i(\theta) = h_i(\theta) - \sum_{j \neq i} v_j(d(\theta), \theta_j).$$

We need only show that h_i is independent of θ_i .

(3) Suppose to the contrary, that there exists i , θ and θ'_i such that $h_i(\theta) < h_i(\theta_{-i}, \theta'_i)$.

(4) Let $\epsilon = \frac{1}{2}[h_i(\theta_{-i}, \theta'_i) - h_i(\theta)]$.

(5) By dominant strategy incentive compatibility, it follows that

$$v_i(d(\theta_{-i}, \theta_i), \theta_i) - t(\theta_{-i}, \theta_i) \geq v_i(d(\theta_{-i}, \theta'_i), \theta_i) - t(\theta_{-i}, \theta'_i),$$

that is,

$$\begin{aligned} & v_i(d(\theta_{-i}, \theta_i), \theta_i) + \sum_{j \neq i} v_j(d(\theta_{-i}, \theta_i), \theta_j) - h_i(\theta_{-i}, \theta_i) \\ & \geq v_i(d(\theta_{-i}, \theta'_i), \theta_i) + \sum_{j \neq i} v_j(d(\theta_{-i}, \theta'_i), \theta_j) - h_i(\theta_{-i}, \theta'_i). \end{aligned}$$

If $d(\theta) = d(\theta_{-i}, \theta'_i)$, then we have

$$h_i(\theta_{-i}, \theta'_i) \geq h_i(\theta),$$

which is a contradiction. Hence $d(\theta) \neq d(\theta_{-i}, \theta'_i)$.

(6) Given the completeness of type spaces, there exists $\theta''_i \in \Theta_i$ such that

$$v_i(d(\theta_{-i}, \theta'_i), \theta''_i) + \sum_{j \neq i} v_j(d(\theta_{-i}, \theta'_i), \theta_j) = \epsilon,$$

and

$$v_i(d, \theta''_i) + \sum_{j \neq i} v_j(d, \theta_j) = 0 \text{ for any } d \neq d(\theta_{-i}, \theta'_i).$$

(7) If $d(\theta_{-i}, \theta''_i) \neq d(\theta_{-i}, \theta'_i)$, then by efficiency of d we have

$$\begin{aligned} 0 &= v_i(d(\theta_{-i}, \theta''_i), \theta''_i) + \sum_{j \neq i} v_j(d(\theta_{-i}, \theta''_i), \theta_j) \\ &\geq v_i(d(\theta_{-i}, \theta'_i), \theta''_i) + \sum_{j \neq i} v_j(d(\theta_{-i}, \theta'_i), \theta_j) = \epsilon > 0, \end{aligned}$$

which is a contradiction. Hence, $d(\theta_{-i}, \theta''_i) = d(\theta_{-i}, \theta'_i)$.

(8) By dominant strategy incentive compatibility, we have

$$\begin{aligned} v_i(d(\theta_{-i}, \theta'_i), \theta'_i) - t_i(\theta_{-i}, \theta'_i) &\geq v_i(d(\theta_{-i}, \theta''_i), \theta'_i) - t_i(\theta_{-i}, \theta'_i), \\ v_i(d(\theta_{-i}, \theta''_i), \theta''_i) - t_i(\theta_{-i}, \theta''_i) &\geq v_i(d(\theta_{-i}, \theta'_i), \theta''_i) - t_i(\theta_{-i}, \theta'_i). \end{aligned}$$

Then $t_i(\theta_{-i}, \theta'_i) = t_i(\theta_{-i}, \theta''_i)$.

(9) Thus, the utility to i from truthful announcement at θ''_i is

$$v_i(d(\theta_{-i}, \theta''_i), \theta''_i) - t_i(\theta_{-i}, \theta''_i) = \epsilon - h_i(\theta_{-i}, \theta'_i),$$

and by lying and reporting θ_i at θ''_i , i gets $-h_i(\theta)$.

(10) This contradicts dominant strategy incentive compatibility since $h_i(\theta) < h_i(\theta_{-i}, \theta'_i) - \epsilon$.

□

18.5.1 Pivot mechanism

18.37 One version of VCG mechanism is called the pivot mechanism, where

$$h_i(\theta_{-i}) = \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

18.38 In the pivot mechanism, i 's transfer becomes

$$t_i^{\text{pivot}}(\theta) = - \sum_{j \neq i} v_j(d(\theta), \theta_j) + \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

This transfer is always non-negative, and so the pivot mechanism is always feasible.

- 18.39
- The term $\max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j)$ maximizes the sum of everyone else's utility if i were ignored.
 - The term $\sum_{j \neq i} v_j(d(\theta), \theta_j)$ is the maximum sum of other agents' utility when i is taken into account.

Agent i get paid everyone else's utility under the allocation that is actually chosen, i.e., $\sum_{j \neq i} v_j(d(\theta), \theta_j)$, and get charged everyone's utility in the world where you do not participate. That is, agent i pays her social cost.

- 18.40
- If i 's presence makes no difference in maximizing choice of d in two cases, then $t_i(\theta) = 0$, that is, agents who do not affect the outcome pay 0.
 - Otherwise, we can think of i as being pivotal, and then t_i represents the loss in value that is imposed on the other agents due to the change in decision that results from i 's presence in society.

18.41 Definition: A social choice function (d, t) is (ex post) individually rational if for each agent i , for each θ_i and θ_{-i} ,

$$v_i(d(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) \geq 0.$$

18.42 Proposition: If the function v_i is always nonnegative, the pivot mechanism is individually rational.

Proof. Routine. □

18.43 Proposition (Uniqueness of VCG transfers): Suppose $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$, $v(d, \cdot)$ is differentiable, and there exists $B > 0$ such that for all d and θ $|v_\theta(d, \theta)| \leq B$. If (d, t) is dominant strategy incentive compatible and d is efficient, then there exists $h_i: \Theta_{-i} \rightarrow \mathbb{R}$, such that

$$t_i(\theta) = t_i^{\text{pivot}}(\theta) + h_i(\theta_{-i}) \text{ for all } \theta \in \Theta.$$

Proof. By Corollary 18.26, for two dominant strategy incentive compatible mechanisms (d, t) and (d, t^{pivot}) , we have

$$t_i(\theta) = t_i^{\text{pivot}}(\theta) + h_i(\theta_{-i}).$$

□

18.44 Proposition: Among all dominant strategy incentive compatible and individually rational mechanism, the pivot mechanism has the largest expected budget surplus.

Proof. ????? □

18.45 Example: Three agents run the pivot mechanism to decide whether or not to build an airport. Their reports are:

Agent	Utility (build)	Utility (not build)	Payment
1	0	60	?
2	45	15	?
3	45	5	?

What outcome will pivot mechanism given the above reports? If the outcome is to build, what are agents 2 and 3's payments?

Answer. Since $0 + 45 + 45 = 90 > 80 = 60 + 15 + 5$, the outcome is to build the airport.

If agent 2 were not present, then the airport would not be built, and the other agents would get $60 + 5 = 65$. The other agents get utility $0 + 45$ from the airport's being built. So agent 2's payment is $65 - 45 = 20$.

If agent 3 were not present, then the airport would not be built, and the other agents would get $60 + 15 = 75$. The other agents get utility $0 + 45$ from the airport's being built. So agent 3's payment is $75 - 45 = 30$. □

18.46 Example: A social planner wishes to build a road connecting A and F . There are several agents who can build the sub-road with some cost. The costs are presented in Figure 18.3; e.g., agent AB can build the part AB with cost 3. The cost is the agents' private information, and they could lie about their cost. The planner's goal is to find the relevant agents to work together to build the entire road and minimize the total cost. What outcome will be selected by the planner? In the outcome, what are the costs for agents AC , AB , BE and BF ?




Figure 18.3: Building a road.

Solution. (1) Note that minimizing the social cost is equivalent to maximizing the negative of total cost, which goes back to the familiar expression in VCG mechanism that the goal is to maximize something.

(2) It is clear that the path $ABEF$ will be selected.

(3) For agent AC : The smallest cost taking AC 's declaration into account is 5, and imposes cost 5 on agents other than AC . The smallest cost without AC 's declaration is also 5. So $t_{AC} = (-5) - (-5) = 0$.

(4) For agent AB : The smallest cost taking AB 's declaration into account is 5, and imposes cost 2 on agents other than AB . The smallest cost without AB 's declaration is 6. So $t_{AB} = (-6) - (-2) = -4$.

(5) For agent BE : $t_{BE} = (-6) - (-4) = -2$.

(6) For agent BF : $t_{BF} = (-7) - (-4) = -3$.

□

Note that EF has more market power: for the other agents, the situation without EF is worse than the situation without BE .

18.47 The pivot mechanism is susceptible to collusion.

Agent	Utility (build)	Utility (not build)	Payment
1	200	0	150
2	100	0	50
3	0	250	0

If agents 1 and 2 both increase their declared valuations by 50, we have

Agent	Utility (build)	Utility (not build)	Payment
1	250	0	100
2	150	0	0
3	0	250	0

The choice is unchanged, but both of their payments are reduced. Thus, while no agent can gain by changing her reporting, groups can.

18.48 Revenue monotonicity: The revenue that a mechanism would obtain would always weakly increase as adding agents to the mechanism.

The pivot mechanism may violate revenue monotonicity.

Agent	Utility (build)	Utility (not build)	Payment
1	0	90	0
2	100	0	90

Agent	Utility (build)	Utility (not build)	Payment
1	0	90	0
2	100	0	0
3	100	0	0

Adding agent 3 causes the pivot mechanism to make the same choice but to collect zero revenue.

18.5.2 Balancing the budget

18.49 Theorem: There exists a VCG mechanism that satisfies budget balance if and only if for every i there is a function $g_i: \Theta_{-i} \rightarrow \mathbb{R}$ such that

$$\sum_{i=1}^N v_i(d(\theta), \theta_i) = \sum_{i=1}^N g_i(\theta_{-i}) \text{ for all } \theta \in \Theta.$$

18.50 *Proof of necessity.*

(1) Suppose that a VCG mechanism $(d(\cdot), t^{\text{VCG}}(\cdot))$ is budget balanced, then we have

$$\sum_{i=1}^N \left[h_i(\theta_{-i}) - \sum_{j \neq i} v_j(d(\theta), \theta_j) \right] = \sum_{i=1}^N t_i^{\text{VCG}}(\theta) = 0.$$

(2) This equality is equivalent to

$$\sum_{i=1}^N h_i(\theta_{-i}) = \sum_{i=1}^N \sum_{j \neq i} v_j(d(\theta), \theta_j) = (N-1) \sum_{i=1}^N v_i(d(\theta), \theta_i).$$

(3) Hence, if we set for every i and θ_{-i} ,

$$g_i(\theta_{-i}) = \frac{h_i(\theta_{-i})}{N-1},$$

we have obtained the desired form for the function $\sum_{i=1}^N v_i(d(\theta), \theta_i)$.

□

18.51 *Proof of sufficiency.*

(1) Suppose that $\sum_{i=1}^N v_i(d(\theta), \theta_i)$ has the form described in the statement.

(2) For every i and every θ_{-i} we consider the VCG mechanism with

$$h_i(\theta_{-i}) \triangleq (N-1)g_i(\theta_{-i}).$$

(3) Then for every θ , the sum of agents' payments is

$$\sum_{i=1}^N t_i^{\text{VCG}}(\theta) = \sum_{i=1}^N \left[h_i(\theta_{-i}) - \sum_{j \neq i} v_j(d(\theta), \theta_j) \right] = (N-1) \left[\sum_{i=1}^N g_i(\theta_{-i}) - \sum_{i=1}^N v_i(d(\theta), \theta_i) \right] = 0.$$

□

18.6 Characterization of dominant strategy incentive compatibility

18.52 In this section, we focus on the environment:

- only one agent, whose utility is $v(d, \theta) - t$ when decision d is chosen and transfer is t .
- D is an interval in \mathbb{R} .
- $\Theta = [\underline{\theta}, \bar{\theta}]$.

In this case, the strict dominant incentive compatibility is equivalent to the Bayesian incentive compatibility.

18.53 Definition: $v: D \times \Theta \rightarrow \mathbb{R}$ satisfies strictly increasing difference (abbreviated as “SID”) if for each $\theta, \theta' \in D$, for each $\theta, \theta' \in \Theta$, $d' > d$ and $\theta' > \theta$ implies

$$v(d', \theta') - v(d, \theta') > v(d', \theta) - v(d, \theta).$$

18.54 Lemma: If $v_{d\theta}(d, \theta) \triangleq \frac{\partial^2 v}{\partial d \partial \theta}$ exists everywhere, then v has SID if and only if $v_{d\theta} > 0$ for all $(d, \theta) \in D \times \Theta$.

18.55 Lemma: If v has SID, and $(d(\cdot), t(\cdot))$ is incentive compatible, then $d(\cdot)$ is non-decreasing.

Proof.

(1) Fix $\theta' > \theta$. Suppose $(d(\cdot), t(\cdot))$ is incentive compatible.

(2) Then we have

$$v(d(\theta), \theta) - t(\theta) \geq v(d(\theta'), \theta) - t(\theta'),$$

thus,

$$v(d(\theta), \theta) - v(d(\theta'), \theta) \geq t(\theta) - t(\theta').$$

(3) Likewise,

$$v(d(\theta'), \theta') - t(\theta') \geq v(d(\theta), \theta') - t(\theta),$$

thus,

$$t(\theta) - t(\theta') \geq v(d(\theta), \theta') - v(d(\theta'), \theta').$$

(4) Therefore,

$$v(d(\theta), \theta) - v(d(\theta'), \theta) \geq v(d(\theta), \theta') - v(d(\theta'), \theta').$$

(5) If we assume $d(\theta) > d(\theta')$, then the equation above contradicts SID. Therefore, we must have $d(\theta') \geq d(\theta)$.

□

18.56 Lemma: Suppose v has SID. If $(d(\cdot), t(\cdot))$ satisfies

$$v(d(\theta), \theta) - t(\theta) = v(d(\underline{\theta}), \underline{\theta}) - t(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} v_{\theta}(d(s), s) \, ds \text{ for all } \theta \in \Theta,$$

and $d(\cdot)$ is non-decreasing, then $(d(\cdot), t(\cdot))$ is IC.

Proof.

(1) Fix $\theta, \theta' \in \Theta$.

(2) We have

$$\begin{aligned} & v(d(\theta), \theta) - t(\theta) - v(d(\theta'), \theta) + t(\theta') \\ &= v(d(\theta), \theta) - t(\theta) - v(d(\theta'), \theta') + t(\theta') + v(d(\theta'), \theta') - v(d(\theta'), \theta) \\ &= \int_{\theta'}^{\theta} v_{\theta}(d(s), s) \, ds + \int_{\theta}^{\theta'} v_{\theta}(d(\theta'), s) \, ds \\ &= \int_{\theta'}^{\theta} v_{\theta}(d(s), s) - v_{\theta}(d(\theta'), s) \, ds \end{aligned}$$

(3) Since $d(\cdot)$ is non-decreasing and v has SID,

$$v(d(\theta), \theta) - t(\theta) - v(d(\theta'), \theta) + t(\theta') \geq 0.$$

□

✂ 18.57 Theorem (IC characterization): If v has SID, then $(d(\cdot), t(\cdot))$ is IC if and only if envelope condition and monotonicity condition hold.

Chapter

19

Mechanism design: Auction

Contents

19.1 Set-up	329
19.2 The revelation principle for Bayesian equilibrium	331
19.3 Incentive compatibility and individual rationality	332
19.4 Expected revenue maximization	335
19.5 Maximizing welfare	339
19.5.1 VCG mechanism	341
19.5.2 AGV mechanism	343
19.6 Dominant strategy mechanism	344
19.7 Bilateral trade	345

19.1 Set-up

19.1 A seller seeks to sell a single indivisible good. There are $N \geq 2$ potential buyers. We denote the set of potential buyers by $\mathcal{N} = \{1, 2, \dots, N\}$.

19.2 Buyer i 's utility if he purchases the good and pays a transfer t_i to the seller is: $\theta_i - t_i$.

19.3 The seller's utility if she obtains transfers t_i from buyers $i = 1, 2, \dots, N$ is $\sum_{i \in \mathcal{N}} t_i$.

19.4 Buyer i knows θ_i , but that neither the seller nor any other buyer $j \neq i$ knows θ_i .

Buyer's value θ_i is distributed over $[\underline{\theta}_i, \bar{\theta}_i]$ according to the cumulative distribution function F_i with associated density f_i , where $0 < \underline{\theta}_i < \bar{\theta}_i$. Although we do not assume that the random variables θ_i have the same distribution for different i , we do assume that they have the same support. This is for convenience only. For technical convenience, we also assume that $f_i(\theta_i) > 0$ for all $i \in \mathcal{N}$ and all $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$.

We assume that for $i, j \in \mathcal{N}$ with $i \neq j$, the random variables θ_i and θ_j are independent.

The distributions $\{F_i\}_{i \in \mathcal{N}}$ are common knowledge among the buyers and the seller.

19.5 We will be interested in procedures that the seller can use to sell her good for revenue maximization.

19.6 General criteria for the selling mechanism designer (*i.e.*, seller):

- The mechanism designer can not force people to play—they have to be willing to play.
- The mechanism designer need to assume people will play an equilibrium within whatever game you define.

19.7 We assume that the seller has full commitment power—once she defines the rules of the game, the players have complete confidence that she will honor those rules. This is important—you had bid differently in an auction if you thought that, even if you won, the seller might demand a higher price or mess with you some other way.

19.8 Broadly speaking, mechanism design takes the environment as given—the players, their value distributions, and their preferences over the different possible outcomes—and designs a game for the players to play in order to select one of the outcomes. Outcomes can be different legislative proposals, different allocations of one or more objects, *etc.*

Here we will focus on the auction problem—designing a mechanism to sell a single object, and try to maximize the expected revenue or expected welfare. So the set of possible outcomes consists of who (if anyone) gets the object, and how much each person pays.

19.9 A selling mechanism (\mathcal{B}, π, μ) has the following components:

- a set of possible messages (or “bids”) \mathcal{B}_i for each buyer;
- an allocation rule $\pi: \mathcal{B} \rightarrow \Delta$, where $\Delta = \{(\pi_1, \pi_2, \dots, \pi_N) \mid 0 \leq \pi_i \leq 1 \text{ for all } i \text{ and } \sum_{i \in \mathcal{N}} \pi_i \leq 1\}$ is the set of probability distributions over the set of buyers and seller $\mathcal{N} \cup \{0\}$;
- a payment rule $\mu: \mathcal{B} \rightarrow \mathbb{R}^N$.

An allocation rule determines, as a function of all N messages, the probability $\mu_i(b)$ that i will get the object.

A payment rule determines, as a function of all N messages, for each buyer i , the expected payment $\mu_i(b)$ that i must make.

19.10 Example: Both first- and second-price auctions are mechanisms.

- The set of possible bids \mathcal{B}_i in both can be safely assumed to be \mathcal{X}_i .
- There is no reservation price, and

$$\pi_i(b) = \begin{cases} 1, & \text{if } b_i > \max_{j \neq i} b_j, \\ 0, & \text{if } b_i < \max_{j \neq i} b_j. \end{cases}$$

- For a first-price auction,

$$\mu_i^I(b) = \begin{cases} b_i, & \text{if } b_i > \max_{j \neq i} b_j, \\ 0, & \text{if } b_i < \max_{j \neq i} b_j. \end{cases}$$

For a second-price auction,

$$\mu_i^{II}(b) = \begin{cases} \max_{j \neq i} b_j, & \text{if } b_i > \max_{j \neq i} b_j, \\ 0, & \text{if } b_i < \max_{j \neq i} b_j. \end{cases}$$

19.11 Every mechanism defines a game of incomplete information among the buyers. An N -tuple of strategies $\beta_i: \Theta_i \rightarrow \mathcal{B}_i$ is a Bayesian Nash equilibrium of a mechanism if for all i and for all θ_i , given the strategies β_{-i} of other buyers, $\beta_i(\theta_i)$ maximizes i 's expected payoff.

19.2 The revelation principle for Bayesian equilibrium

19.12 A mechanism could, in principle, be quite complicated since we have made no assumptions on the sets i of “bids” or “messages.” A smaller and simpler class consists of those mechanisms for which the set of messages is the same as the set of values—that is, for all i , $\mathcal{B}_i = \Theta_i$. Such mechanisms are called direct, since every buyer is asked to directly report a value.

19.13 Formally, a direct mechanism (q, t) consists of a pair of functions

$$q: \Theta \rightarrow \Delta, \quad t: \Theta \rightarrow \mathbb{R}^N,$$

where $q_i(\theta)$ is the probability that i will get the object, $1 - \sum_{i \in \mathcal{N}} q_i(\theta)$ is the probability that the good is not sold, and $t_i(\theta)$ is the expected payment by i .

We refer to the pair $(q(\theta), t(\theta))$ as the outcome of the mechanism at θ .

19.14 Theorem (Revelation principle for Bayesian equilibrium): Given a mechanism and a Bayesian Nash equilibrium for that mechanism, there exists a direct mechanism in which

- (i) it is a Bayesian Nash equilibrium for each buyer to report his or her value truthfully;
- (ii) the outcomes are the same as in the given equilibrium of the original mechanism.

Proof. (1) Suppose that (\mathcal{B}, π, μ) is a mechanism and β is a Bayesian Nash equilibrium of this mechanism.

(2) Let $q: \Theta \rightarrow \Delta$ and $t: \Theta \rightarrow \mathbb{R}^N$ be defined as follows:

$$q(\theta) = \pi(\beta(\theta)), \quad t(\theta) = \mu(\beta(\theta)).$$




Figure 19.1: Revelation principle for Bayesian equilibrium

(3) It is clear that

$$\begin{aligned}
 & \mathbf{E}_{\theta_{-i}} [q_i(\theta_i, \theta_{-i}) \cdot \theta_i - t_i(\theta_i, \theta_{-i})] \\
 &= \mathbf{E}_{\theta_{-i}} [\pi_i(\beta_i(\theta_i), \beta_{-i}(\theta_{-i})) \cdot \theta_i - \mu_i(\beta_i(\theta_i), \beta_{-i}(\theta_{-i}))] \\
 &> \mathbf{E}_{\theta_{-i}} [\pi_i(\beta_i(\theta'_i), \beta_{-i}(\theta_{-i})) \cdot \theta_i - \mu_i(\beta_i(\theta'_i), \beta_{-i}(\theta_{-i}))] \\
 &= \mathbf{E}_{\theta_{-i}} [q_i(\theta'_i, \theta_{-i}) \cdot \theta_i - t_i(\theta'_i, \theta_{-i})]
 \end{aligned}$$

for all $\theta'_i \in \Theta_i$, which implies that truthfully reporting is an equilibrium strategy for each buyer in the direct mechanism (q, t) , and its outcome is the same as in β .

□

- 19.15 Remark: This result shows that the outcomes resulting from any equilibrium of any mechanism can be replicated by a truthful equilibrium of some direct mechanism. In this sense, there is no loss of generality in restricting attention to direct mechanisms.

19.3 Incentive compatibility and individual rationality

- 19.16 Given a direct mechanism (q, t) , define

$$Q_i(\theta'_i) = \int_{\Theta_{-i}} q_i(\theta'_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

to be the probability that i will get the object when she reports her value to be θ'_i and all other buyers report their values truthfully.

Similarly, define

$$T_i(\theta'_i) = \int_{\Theta_{-i}} t_i(\theta'_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

to be the expected payment of i when his report is θ'_i and all other buyers tell the truth.

Note that both the probability of getting the object and the expected payment depend only on the reported value θ'_i and not on the true value θ_i .

- 19.17 The expected payoff of buyer i when his true value is θ_i and he reports θ'_i , again assuming that all other buyers tell the truth, can then be written as

$$Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i).$$

- 19.18 The direct mechanism (q, t) is said to be Bayesian incentive compatible if truth telling is a Bayesian Nash equilibrium; that is, if for all i , for all θ_i and for all θ'_i ,

$$U_i(\theta_i) \triangleq Q_i(\theta_i) \cdot \theta_i - T_i(\theta_i) \geq Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i).$$

We will refer to U_i as the equilibrium payoff function.

- 19.19 Lemma: If a direct mechanism (q, t) is Bayesian incentive compatible, then for every i , the function Q_i is increasing.

Proof. (1) Consider θ_i and θ'_i with $\theta_i > \theta'_i$.

(2) Bayesian incentive compatibility requires

$$Q_i(\theta_i) \cdot \theta_i - T_i(\theta_i) \geq Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i), \text{ and } Q_i(\theta'_i) \cdot \theta'_i - T_i(\theta'_i) \geq Q_i(\theta_i) \cdot \theta'_i - T_i(\theta_i).$$

(3) Then we have

$$[Q_i(\theta_i) - Q_i(\theta'_i)] \cdot (\theta_i - \theta'_i) \geq 0,$$

and hence $Q_i(\theta_i) \geq Q_i(\theta'_i)$.

□

- 19.20 Lemma: If a direct mechanism (q, t) is Bayesian incentive compatible, then for every i , the function U_i is increasing. It is also convex, and hence differentiable except in at most countably many points. For all θ_i for which it is differentiable, it satisfies

$$U'_i(\theta_i) = Q_i(\theta_i).$$

Proof. (1) Bayesian incentive compatibility implies that for all θ_i ,

$$U_i(\theta_i) = \max_{\theta'_i \in \Theta_i} (Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i)).$$

(2) Given any value of θ'_i , $Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i)$ is an increasing and affine (and hence convex) function.

(3) The maximum of increasing functions is increasing, and the maximum of convex functions is convex. Therefore, U_i is increasing and convex.

(4) Convex functions are not differentiable in at most countably many points.

(5) Then, by envelope theorem (Theorem 18.6), we have $U'_i(\theta_i) = Q_i(\theta_i)$ whenever U_i is differentiable. □

Remark: Bayesian incentive compatibility is equivalent to the requirement that for all θ_i and θ'_i ,

$$U_i(\theta'_i) \geq U_i(\theta_i) + Q_i(\theta_i) \cdot (\theta'_i - \theta_i).$$

This implies that for all θ_i , $Q_i(\theta_i)$ is the slope of a line that supports the function U_i at the point θ_i .

19.21 Proposition (Payoff equivalence): Consider a direct Bayesian incentive compatible mechanism. Then for all i and all θ_i , we have

$$U_i(\theta_i) = U_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

Proof. Since U_i is convex, it is absolutely continuous. Since for all θ_i for which U_i is differentiable, it satisfies $U'_i(\theta_i) = Q_i(\theta_i)$, we have

$$U_i(\theta_i) - U_i(\underline{\theta}_i) = \int_{\underline{\theta}_i}^{\theta_i} U'_i(\theta'_i) d\theta'_i = \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

□

19.22 Remark: Proposition 19.21 shows that the interim expected payoffs of the different buyer values are pinned down by the functions Q_i and the expected payoff of the lowest value. That is, Proposition 19.21 implies that up to an additive constant, the interim expected payoff to a buyer in a Bayesian incentive compatible direct mechanism (q, t) depends only on the allocation rule q .

If (q, t) and (q, t') are two Bayesian incentive compatible mechanisms with the same allocation rule q but different payment rules, then the expected payoff functions associated with the two mechanisms, U_i and U'_i , respectively, differ by at most a constant; the two mechanisms are payoff equivalent. Put another way, the “shape” of the expected payoff function is completely determined by the allocation rule q alone. The payment rule t only serves to determine the constants $U_i(\underline{\theta}_i)$.

19.23 Proposition: A direct mechanism (q, t) is Bayesian incentive compatible if and only if for every i

- (i) Q_i is increasing;
- (ii) For every $\theta_i \in \Theta_i$,

$$U_i(\theta_i) = U_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

Proof. (1) Let $\theta_i > \theta'_i \in \Theta_i$.

(2) Since Q_i is increasing, we have

$$\int_{\theta'_i}^{\theta_i} Q_i(\theta''_i) d\theta''_i \geq \int_{\theta'_i}^{\theta_i} Q_i(\theta'_i) d\theta''_i = Q_i(\theta'_i) \cdot (\theta_i - \theta'_i).$$

(3) Since

$$\int_{\theta'_i}^{\theta_i} Q_i(\theta''_i) d\theta''_i = \left[\int_{\underline{\theta}_i}^{\theta_i} - \int_{\underline{\theta}_i}^{\theta'_i} \right] Q_i(\theta''_i) d\theta''_i = U_i(\theta_i) - U_i(\theta'_i),$$

we have

$$U_i(\theta_i) - U_i(\theta'_i) \geq Q_i(\theta'_i) \cdot (\theta_i - \theta'_i).$$

(4) Then

$$U_i(\theta_i) \geq Q_i(\theta'_i) \cdot (\theta_i - \theta'_i) + U_i(\theta'_i) = Q_i(\theta'_i) \cdot (\theta_i - \theta'_i) + Q_i(\theta'_i) \cdot \theta'_i - T_i(\theta'_i) = Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i).$$

(5) If $\theta_i < \theta'_i$, the argument is analogous. □

19.24 Proposition (Revenue equivalence): Consider a direct Bayesian incentive compatible mechanism. Then for all i and all θ_i , we have

$$T_i(\theta_i) = T_i(\underline{\theta}_i) + Q_i(\theta_i) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

Proof. Since $U_i(\theta_i) = Q_i(\theta_i) \cdot \theta_i - T_i(\theta_i)$ and $U_i(\underline{\theta}_i) = -T_i(\underline{\theta}_i)$, by Proposition 19.21, we have

$$Q_i(\theta_i) \cdot \theta_i - T_i(\theta_i) = -T_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

□

19.25 Proposition 19.24 shows similarly that the interim expected payments of the different buyer values are pinned down by the functions Q_i and the expected payment of the lowest value. Note that this does not mean that the *ex post* payment functions t_i are uniquely determined. Different functions t_i might give rise to the same interim expected payments T_i .

19.26 Proposition 19.24 generalizes the revenue equivalence principle in Theorem 4.58, to situations where buyers may be asymmetric.

19.27 Example: Consider the symmetric case in which $F_1 = F_2 = \dots = F_N = F$. Suppose we wanted to compare the auctioneer's expected revenue from the second-price auction with minimum bid 0 to the expected revenue from the first-price auction with minimum bid 0.

In the second-price auction it is a weakly dominant strategy, and hence a Bayesian Nash equilibrium, to bid one's true value.

A symmetric Bayesian Nash equilibrium for the first-price auction is constructed in Proposition 4.25. This equilibrium is in strictly increasing strategies. Hence this equilibrium shares with the equilibrium of the second-price auction that the expected payment of the lowest value is zero (because this value's probability of winning is zero), and that the highest value wins with probability 1. Therefore, the equilibria imply the same values for $T_i(0)$ and $Q_i(0)$ for all i and θ_i . The revenue equivalence theorem implies that the expected revenue from the equilibria of the two different auction formats is the same.

19.28 Proposition: A direct mechanism (q, t) is Bayesian incentive compatible if and only if for every i

- (i) Q_i is increasing;
- (ii) For every $\theta_i \in \Theta_i$,

$$T_i(\theta_i) = T_i(\underline{\theta}_i) + Q_i(\theta_i) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i. \quad (19.1)$$

Proof. Similar with proof of Proposition 19.23. □

19.29 Remark: We have now obtained a complete understanding of the implications of Bayesian incentive compatibility for the seller's choice. The seller can focus on two choice variables: firstly the allocation rule q , and secondly the interim expected payment by a buyer with the lowest type: $T_i(\underline{\theta}_i)$.

As long as the seller picks an allocation rule q such that the functions $\{Q_i\}_{i \in \mathcal{N}}$ are increasing, she can pick the interim expected payments by the lowest values in any arbitrary way, and be assured that there will be some transfer scheme that makes the allocation rule Bayesian incentive compatible and that implies the given interim expected payments by the lowest values. Moreover, any such transfer scheme will give her the same expected revenue, and therefore the seller does not have to worry about the details of this transfer scheme.

19.30 A direct mechanism is individually rational if each agent, conditional on her type, is willing to participate, that is, if

$$U_i(\theta_i) \geq 0 \text{ for all } i \text{ and } \theta_i.$$

We are implicitly assuming here that by not participating, a buyer can guarantee herself a payoff of zero.

19.31 Proposition: A Bayesian incentive compatible direct mechanism is individually rational if and only if for every i , we have

$$U_i(\underline{\theta}_i) \geq 0.$$

Proof. U_i is increasing for Bayesian incentive compatible direct mechanisms. Therefore, $U_i(\theta_i)$ is non-negative for all θ_i if and only if it is non-negative for the lowest θ_i , which is zero. □

Since $U_i(\underline{\theta}_i) = -T_i(\underline{\theta}_i)$, this is equivalent to the requirement that $T_i(\underline{\theta}_i) \leq 0$.

19.4 Expected revenue maximization

19.32 In this section we view the seller as the designer of the mechanism and examine mechanisms that maximize the expected revenue—the sum of the expected payments of the buyers—among all mechanisms that are Bayesian incentive compatible and individually rational. We reiterate that when carrying out this exercise, the revelation principle guarantees that there is no loss of generality in restricting attention to direct mechanisms. Suppose that the seller uses the direct mechanism (q, t) .

We will refer to a mechanism that maximizes expected revenue, subject to the Bayesian incentive compatibility and individual rationality constraints, as an optimal mechanism.

19.33 Expected revenue:

- (1) The expected revenue of the seller is

$$\sum_{i \in \mathcal{N}} \mathbf{E}_{\theta_i} [T_i(\theta_i)],$$

where the *ex ante* expected payment of buyer i is

$$\mathbf{E}_{\theta_i} [T_i(\theta_i)] = \int_{\underline{\theta}_i}^{\bar{\theta}_i} T_i(\theta_i) f_i(\theta_i) d\theta_i.$$

- (2) By substituting Equation (19.1), we have

$$\mathbf{E}_{\theta_i} [T_i(\theta_i)] = T_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\bar{\theta}_i} Q_i(\theta_i) \cdot \theta_i \cdot f_i(\theta_i) d\theta_i - \int_{\underline{\theta}_i}^{\bar{\theta}_i} \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i f_i(\theta_i) d\theta_i.$$

- (3) Interchanging the order of integration in the last term results in

$$\int_{\underline{\theta}_i}^{\bar{\theta}_i} \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) f_i(\theta_i) d\theta'_i d\theta_i = \int_{\underline{\theta}_i}^{\bar{\theta}_i} \int_{\theta'_i}^{\bar{\theta}_i} Q_i(\theta'_i) f_i(\theta_i) d\theta_i d\theta'_i = \int_{\underline{\theta}_i}^{\bar{\theta}_i} (1 - F_i(\theta'_i)) Q_i(\theta'_i) d\theta'_i.$$

- (4) Thus, we can write

$$\begin{aligned} \mathbf{E}_{\theta_i} [T_i(\theta_i)] &= T_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\bar{\theta}_i} \left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) Q_i(\theta_i) f_i(\theta_i) d\theta_i \\ &= T_i(\underline{\theta}_i) + \int_{\Theta} \left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) q_i(\theta) f(\theta) d\theta. \end{aligned}$$

- (5) The seller's objective therefore is to find a mechanism that maximizes

$$\sum_{i \in \mathcal{N}} T_i(\underline{\theta}_i) + \sum_{i \in \mathcal{N}} \int_{\Theta} \left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) q_i(\theta) f(\theta) d\theta,$$

subject to the constraint that the mechanism is

IC Bayesian incentive compatible, which is equivalent to the requirement that Q_i be increasing and that Equation (19.1) be satisfied;

IR individually rational, which is equivalent to the requirement that $T_i(\underline{\theta}_i) \leq 0$.

19.34 We first ask which function q the seller would choose if she did not have to make sure that the functions Q_i are increasing. In a second step, we introduce an assumption that makes sure that the optimal q from the first step implies increasing functions Q_i .

- (1) Let

$$\psi_i(\theta_i) = \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \text{ for all } i \text{ and } \theta_i,$$

which is referred to as virtual valuation of a buyer with value θ_i .

- (2) Focus on

$$\max_q \int_{\Theta} \left[\sum_{i \in \mathcal{N}} \psi_i(\theta_i) q_i(\theta) f(\theta) d\theta \right],$$

then the optimal choice of q without constraints is: for all i and θ ,

$$q_i(\theta) \begin{cases} > 0, & \text{if } \psi_i(\theta_i) = \max_{j \in \mathcal{N}} \psi_j(\theta_j) \geq 0, \\ = 0, & \text{otherwise.} \end{cases}$$

19.35 Condition of regularity: for every i , the function $\psi_i(\theta_i)$ is strictly increasing.

Since

$$\psi_i(\theta_i) = \theta_i - \frac{1}{\lambda_i(\theta_i)},$$

where $\lambda_i = f_i/(1 - F_i)$ is the hazard rate function associated with F_i , a sufficient condition for regularity is that for all i , λ_i is increasing.

19.36 Lemma: If $\psi_i(\theta_i)$ is strictly increasing, then Q_i is increasing.

Proof. (1) Suppose $\theta'_i < \theta_i$. Then by the regularity condition, $\psi_i(\theta'_i) < \psi_i(\theta_i)$.

(2) For any θ_{-i} , if $q_i(\theta'_i, \theta_{-i}) \geq 0$, then

$$\psi_i(\theta'_i) = \max_{j \in \mathcal{N}} \psi_j(\theta_j) \geq 0,$$

and hence

$$\psi_i(\theta'_i) > \max_{j \neq i} \psi_j(\theta_j) \geq 0.$$

Therefore, $q_i(\theta_i, \theta_{-i}) = 1 \geq q_i(\theta'_i, \theta_{-i})$.

(3) If $q_i(\theta'_i, \theta_{-i}) = 0$, it means the virtual value of bidder i with value θ'_i is not the highest. Now when her value is θ_i , the virtual value is either still not the highest, which gives zero, or the virtual value becomes the highest among all bidders which give strictly positive number. Thus $q_i(\theta_i, \theta_{-i}) \geq q_i(\theta'_i, \theta_{-i})$.

(4) Therefore, Q_i is an increasing function. □

19.37 Lemma: If a Bayesian incentive compatible and individually rational direct mechanism maximizes the seller's expected revenue, then for every i , $T_i(\underline{\theta}_i) = 0$.

Proof. It is clear that $T_i(\underline{\theta}_i) \leq 0$. If $T_i(\underline{\theta}_i) < 0$, then the seller could increase expected revenue by choosing a direct mechanism with the same q , but a higher $T_i(\underline{\theta}_i)$. By the formula for payments in Proposition 19.28, all values' payments would increase. □

When $T_i(\underline{\theta}_i) = 0$, we have

$$T_i(\theta_i) = Q_i(\theta_i) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

19.38 Theorem: Suppose the design problem is regular. Then the following is an optimal mechanism:

$$q_i(\theta) \begin{cases} \geq 0, & \text{if } \psi_i(\theta_i) = \max_{j \in \mathcal{N}} \psi_j(\theta_j) \geq 0, \\ = 0, & \text{otherwise,} \end{cases}$$

and

$$T_i(\theta) = Q_i(\theta_i) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

Proof. It is clear that (q, t) is Bayesian incentive compatible and individually rational. It is optimal, since it separately maximizes

$$\sum_{i \in \mathcal{N}} T_i(\theta_i) \text{ and } \sum_{i \in \mathcal{N}} \int_{\Theta} \psi_i(\theta_i) q_i(\theta) f(\theta) d\theta$$

over all $q: \Theta \rightarrow \Delta$. In particular, it gives positive weight only to non-negative and maximal terms in

$$\sum_{i \in \mathcal{N}} \psi_i(\theta_i) q_i(\theta).$$

□

19.39 Remark: We have characterized the optimal choice of the allocation rule q and of the interim expected payments. We have not described the actual transfer schemes that make these choices Bayesian incentive compatible and individually rational.

19.40 In the optimal mechanism, the maximized value of the expected revenue is

$$\mathbb{E}_{\theta}[\max\{\psi_1(\theta_1), \psi_2(\theta_2), \dots, \psi_N(\theta_N), 0\}].$$

In other words, it is the expectation of the highest virtual valuation, provided it is non-negative.

19.41 One possible payment rule t is

$$t_i(\theta) = q_i(\theta) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} q_i(\theta'_i, \theta_{-i}) d\theta'_i.$$

It is clear that $t_i(\underline{\theta}_i, \theta_{-i}) = 0$ for all θ_{-i} , and hence

$$T_i(\underline{\theta}_i) = \int_{\Theta_{-i}} t_i(\underline{\theta}_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i} = 0.$$

19.42 Let

$$\vartheta_i(\theta_{-i}) = \inf\{\theta_i \mid \psi_i(\theta_i) \geq 0 \text{ and } \psi_i(\theta_i) \geq \psi_j(\theta_j) \text{ for all } j \neq i\}$$

as the smallest value for i that “wins” against θ_{-i} .

Then

$$q_i(\theta_i, \theta_{-i}) = \begin{cases} 1, & \text{if } \theta_i > \vartheta_i(\theta_{-i}), \\ 0, & \text{otherwise,} \end{cases}$$

which results in

$$\int_{\underline{\theta}_i}^{\theta_i} q_i(\theta'_i, \theta_{-i}) d\theta'_i = \begin{cases} \theta_i - \vartheta_i(\theta_{-i}), & \text{if } \theta_i > \vartheta_i(\theta_{-i}), \\ 0, & \text{if } \theta_i < \vartheta_i(\theta_{-i}), \end{cases}$$

and so payment rule becomes

$$t_i(\theta) = \begin{cases} \vartheta_i(\theta_{-i}), & \text{if } q_i(\theta) = 1, \\ 0, & \text{if } q_i(\theta) = 0. \end{cases}$$

That is, only the winning buyer pays anything: she pays the smallest value that would results in her winning.

19.43 Corollary: Suppose the design problem is regular. Then the following is an optimal mechanism:

$$q_i(\theta) = \begin{cases} 1, & \text{if } \psi_i(\theta_i) > \max_{j \neq i} \psi_j(\theta_j) \text{ and } \psi_i(\theta_i) \geq 0, \\ 0, & \text{if } \psi_i(\theta_i) < \max_{j \neq i} \psi_j(\theta_j), \end{cases}$$

and

$$t_i(\theta) = \begin{cases} \vartheta_i(\theta_{-i}), & \text{if } q_i(\theta) = 1, \\ 0, & \text{if } q_i(\theta) = 0. \end{cases}$$

19.44 Example: Suppose that we have a symmetric problem so the distributions of values are identical across buyers. In other words, for all i , $f_i = f$, and hence for all i , $\psi_i = \psi$. Now we have that,

$$\vartheta_i(\theta_{-i}) = \max \left\{ \psi^{-1}(0), \max_{j \neq i} \theta_j \right\}.$$

Thus, the optimal mechanism is a second-price auction with a reserve price $r^* = \psi^{-1}(0)$.

19.45 Remark: Note that in the case with asymmetric buyers, the optimal mechanism may sometimes give the good to a buyer who does not have the highest value.

19.5 Maximizing welfare

19.46 Suppose that the seller is not maximizing expected revenue but expected welfare. So the seller uses the following utilitarian welfare function, where each agent has equal weight:

$$\sum_{i \in \mathcal{N}} q_i(\theta) \cdot \theta_i.$$

Note that this seller is no longer concerned with transfer payments, and expected welfare depends only on the allocation rule q .

19.47 By Lemma 19.19, the seller can choose any rule q that is such that the functions Q_i are increasing. By Proposition 19.31, she can choose any transfer payments such that $T_i(\underline{\theta}_i) \leq 0$ for all i .

19.48 If values were known, maximization of the welfare function would require that the object be allocated to the potential buyer for whom θ_i is largest.

Because transferring to the buyer for whom θ_i is largest maximizes welfare for every type vector, it also maximizes expected welfare.

In this case, it is clear that Q_i is increasing.

19.49 Proposition: Among all Bayesian incentive compatible, individually rational direct mechanisms, a mechanism maximizes expected welfare if and only if for all i and all θ :

(i)

$$q_i(\theta) = \begin{cases} 1, & \text{if } \theta_i > \theta_j \text{ for all } j \neq i, \\ 0, & \text{otherwise.} \end{cases}$$

(ii)

$$T_i(\theta_i) \leq Q_i(\theta_i) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

19.50 Remark: Note that this result does not rely on regularity condition.

19.51 Differences between welfare maximizing and revenue maximizing mechanisms in the case that regularity condition holds.

- Revenue maximizing mechanism allocates the object to the highest virtual type whereas the welfare maximizing mechanism allocates the object to the highest actual type. In the symmetric case, the functions ψ_i are the same for all i and there is no difference between these two rules. But in the asymmetric case the revenue maximizing mechanism might allocate the object inefficiently.
- Revenue maximizing mechanism sometimes does not sell the object at all, whereas the welfare maximizing mechanism always sells the object. This is an instance of the well-known inefficiency that monopoly sellers make goods artificially scarce.

19.52 Example: Suppose that $\Theta_1 = \Theta_2 = [0, 1]$, and that θ_1 and θ_2 are independently and uniformly distributed on $[0, 1]$. Then $F_i(\theta_i) = \theta_i$, and

$$\psi_i(\theta_i) = \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} = 2\theta_i - 1.$$

Note that the regularity condition is satisfied.

In an expected revenue maximizing auction, the good is sold to neither bidder if

$$\psi_1(\theta_1), \psi_2(\theta_2) < 0,$$

that is,

$$\theta_1, \theta_2 < \frac{1}{2}.$$

If the good is sold, it is sold to bidder 1 if and only if

$$\psi_1(\theta_1) > \psi_2(\theta_2) \Leftrightarrow \theta_1 > \theta_2.$$

The expected revenue maximizing auction will allocate the object to the buyer with the highest value provided that this value is larger than $\frac{1}{2}$. A first- or second-price auction with reserve price $\frac{1}{2}$ will implement this mechanism.

19.53 Example: Suppose that $\Theta_1 = \Theta_2 = [0, 1]$, and that $F_1(\theta_1) = \theta_1^2$ and $F_2(\theta_2) = 2\theta_2 - \theta_2^2$. Thus, buyer 1 is more likely to have high values than buyer 2.

$$\psi_1(\theta_1) = \theta_1 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} = \frac{3}{2}\theta_1 - \frac{1}{2\theta_1},$$

and

$$\psi_2(\theta_2) = \theta_2 - \frac{1 - F_2(\theta_2)}{f_2(\theta_2)} = \frac{3}{2}\theta_2 - \frac{1}{2}.$$

Note that the regularity condition is satisfied.

In an expected revenue maximizing auction, the good is sold to neither bidder if


$$\psi_1(\theta_1), \psi_2(\theta_2) < 0,$$

that is,

$$\theta_1 < \frac{1}{\sqrt{3}}, \text{ and } \theta_2 < \frac{1}{3}.$$

If the good is sold, it is sold to buyer 1 if and only if

$$\psi_1(\theta_1) > \psi_2(\theta_2) \Leftrightarrow \theta_2 < \theta_1 - \frac{1}{3\theta_1} + \frac{1}{3}.$$



19.5.1 VCG mechanism

- 19.54 • Recall that an allocation rule $q^*: \Theta \rightarrow \Delta$ is said to be efficient if it maximizes “social welfare”—that is, for all $\theta \in \Theta$,

$$q^*(\theta) \in \arg \max_{q \in \Delta} \sum_{j \in \mathcal{N}} q_j \theta_j.$$

When there are no ties, an efficient rule allocates the object to the buyer who values it the most.

- A mechanism with an efficient allocation rule is said to be efficient.

- 19.55 Given an efficient allocation rule q^* , define the maximized value of social welfare by

$$w(\theta) \triangleq \sum_{j \in \mathcal{N}} q_j^*(\theta) \theta_j,$$

when the values are x .

Similarly, define

$$w_{-i}(\theta) \triangleq \sum_{j \neq i} q_j^*(\theta) \theta_j$$

as the welfare of agents other than i .

- 19.56 The Vickrey-Clarke-Groves (VCG) mechanism (q^*, t^{VCG}) is an efficient mechanism with the payment rule $t^{\text{VCG}}: \Theta \rightarrow \mathbb{R}^N$ given by

$$t_i^{\text{VCG}}(\theta) = -w_{-i}(\theta_i, \theta_{-i}) + h_i(\theta_{-i}).$$

- 19.57 Proposition: VCG mechanism is Bayesian incentive compatible: If the other buyer report values θ_{-i} , then by reporting a value of θ'_i , agent i 's payoff is

$$\begin{aligned} q_i^*(\theta'_i, \theta_{-i})\theta_i - t_i^{\text{VCG}}(\theta'_i, \theta_{-i}) &= q_i^*(\theta'_i, \theta_{-i})\theta_i - h_i(\theta_{-i}) + \sum_{j \neq i} q_j^*(\theta'_i, \theta_{-i})\theta_j \\ &= \sum_{j \in \mathcal{N}} q_j^*(\theta'_i, \theta_{-i})\theta_j - h_i(\theta_{-i}). \end{aligned}$$

The definition of q^* implies that for all θ_{-i} , the first term is maximized by choosing $\theta'_i = \theta_i$; and since the second term does not depend on θ'_i , it is optimal to report $\theta'_i = \theta_i$.

Actually, the VCG mechanism is dominant strategy incentive compatible.

- 19.58 In the following, we consider a special VCG mechanism by taking $h_i(\theta_{-i})$ to be $w(\underline{\theta}_i, \theta_{-i})$.

In this case, $t_i^{\text{VCG}}(\theta)$ is thus the difference between social welfare at i 's lowest possible value $\underline{\theta}_i$ and the welfare of other agents at i 's reported value θ_i , assuming that in both cases that the efficient allocation rule q^* is employed.

- 19.59 i 's equilibrium payoff when the values are θ is

$$q_i^*(\theta)\theta_i - t_i^{\text{VCG}}(\theta) = q_i^*(\theta)\theta_i - w(\underline{\theta}_i, \theta_{-i}) + \sum_{j \neq i} q_j^*(\theta)\theta_j = w(\theta) - w(\underline{\theta}_i, \theta_{-i}),$$

which is just the difference in social welfare induced by i when she reports her true value θ_i as opposed to her lowest possible value $\underline{\theta}_i$.

- 19.60 Proposition: VCG mechanism is individually rational.

Proof. Since the VCG mechanism is incentive compatible, the equilibrium expected payoff function U_i^{VCG} associated with the VCG mechanism,

$$U_i^{\text{VCG}}(\theta_i) = \mathbf{E}_{\theta_{-i}}[w(\theta) - w(\underline{\theta}_i, \theta_{-i})]$$

is convex and increasing.

Clearly, $U_i^{\text{VCG}}(\underline{\theta}_i) = 0$, and the monotonicity of U_i^{VCG} implies that VCG mechanism is individually rational. \square

- 19.61 Proposition: Among all mechanisms for allocating a single object that are efficient, incentive compatible, and individually rational, the VCG mechanism maximizes the expected payment of each agent.

Proof. If (q^*, t) is some other efficient mechanism that is also incentive compatible, then by the revenue equivalence principle we know that for all i , the expected payoff functions for this mechanism, say U_i , differ from U_i^{VCG} by at most an additive constant, say $c_i = U_i(\theta_i) - U_i^{\text{VCG}}(\theta_i)$.

If (q^*, t) is also individually rational, then this constant must be non-negative—that is, $c_i = U_i(\theta_i) - U_i^{\text{VCG}}(\theta_i) \geq 0$. This is because otherwise we would have $U_i(\underline{\theta}_i) < U_i^{\text{VCG}}(\underline{\theta}_i) = 0$, contradicting that (q^*, t) was individually rational.

Since the expected payoffs in (q^*, t) are greater than in the VCG mechanism, and the two have the same allocation rule, the expected payments must be lower. \square

- 19.62 Remark: Taking $\underline{\theta}_i = 0$, then it is routine to see that the VCG mechanism is the same as a second-price auction:

$$t_i^{\text{VCG}}(\theta) = w(0, \theta_{-i}) - w_{-i}(\theta_i, \theta_{-i}) = w_{-i}(0, \theta_{-i}) - w_{-i}(\theta_i, \theta_{-i}),$$

$$= \sum_{j \neq i} \left[q_j^*(0, \theta_{-i}) - q_j^*(\theta_i, \theta_{-i}) \right] \theta_j,$$

and this is positive if and only if $\theta_i \geq \max_{j \neq i} \theta_j$.

In that case, $t_i^{\text{VCG}}(\theta)$ is equal to $\max_{j \neq i} \theta_j$, the second-highest value.

19.5.2 AGV mechanism

19.63 A mechanism is said to balance the budget if for every realization of values, the net payments from agents sum to zero—that is, for all θ ,

$$\sum_{i \in \mathcal{N}} t_i(\theta) = 0.$$

19.64 The Arrow-d'Aspremont-Gérard-Varet or AGV mechanism (also called the “expected externality” mechanism) (q^*, t^{AGV}) is defined by

$$t_i^{\text{AGV}}(\theta) = \frac{1}{N-1} \sum_{j \neq i} \mathbf{E}_{\theta_{-j}} [w_{-j}(\theta_j, \theta_{-j})] - \mathbf{E}_{\theta_{-i}} [w_{-i}(\theta_i, \theta_{-i})].$$

So that for all θ ,

$$\sum_{i \in \mathcal{N}} t_i^{\text{AGV}}(\theta) = 0.$$

19.65 It is easy to see that the AGV mechanism may not satisfy the individual rationality constraint.

19.66 Proposition: There exists an efficient, incentive compatible, and individually rational mechanism that balances the budget if and only if the VCG mechanism results in an expected surplus.

Proof. By Proposition 19.61, if the VCG mechanism runs a deficit, that is

$$\mathbf{E}_{\theta} \left[\sum_i t_i^{\text{VCG}}(\theta) \right] < 0,$$

then all efficient, incentive compatible, and individually rational mechanisms must run a deficit, which can not balance the budget.

We now show that the condition is sufficient by explicitly constructing an efficient, incentive compatible mechanism that balances the budget and is individually rational.

(1) Since we know

$$U_i^{\text{VCG}}(\theta_i) = \mathbf{E}_{\theta_{-i}} [w(\theta_i, \theta_{-i}) - w(\underline{\theta}_i, \theta_{-i})] = \mathbf{E}_{\theta_{-i}} [w(\theta_i, \theta_{-i})] - c_i^{\text{VCG}},$$

where $c_i^{\text{VCG}} = \mathbf{E}_{\theta_{-i}} [w(\underline{\theta}_i, \theta_{-i})]$.

(2) Since (q^*, t^{VCG}) and (q^*, t^{AGV}) are incentive compatible direct mechanisms, by revenue equivalence principle, we have

$$U_i^{\text{AGV}}(\theta_i) = \mathbf{E}_{\theta_{-i}} [w(\theta_i, \theta_{-i})] - c_i^{\text{AGV}},$$

for some constant c_i^{AGV} .

(3) Suppose that the VCG mechanism runs an expected surplus, that is,

$$\mathbf{E}_\theta \left[\sum_{i \in \mathcal{N}} t_i^{\text{VCG}}(\theta) \right] \geq 0.$$

(4) Then

$$\mathbf{E}_\theta \left[\sum_{i \in \mathcal{N}} t_i^{\text{VCG}}(\theta) \right] \geq \mathbf{E}_\theta \left[\sum_{i \in \mathcal{N}} t_i^{\text{AGV}}(\theta) \right],$$

where the right-hand side is exactly 0 due to the budget balance constraint. Equivalently,

$$\sum_{i \in \mathcal{N}} c_i^{\text{VCG}} \geq \sum_{i \in \mathcal{N}} c_i^{\text{AGV}}.$$

(5) For all $i > 1$, define $d_i = c_i^{\text{AGV}} - c_i^{\text{VCG}}$, and let $d_1 = -\sum_{i=2}^N d_i$. Consider the mechanism \bar{t} defined by

$$\bar{t}_i(\theta) = t_i^{\text{AGV}}(\theta) - d_i.$$

(6) Clearly, \bar{t} balances the budget, and is incentive compatible, since the payoff to each agent in the mechanism \bar{t} differs from the payoff from an incentive compatible mechanism, t^{AGV} , by an additive constant.

(7) For all $i > 1$,

$$\bar{U}_i(\theta_i) = U_i^{\text{AGV}}(\theta_i) + d_i = U_i^{\text{AGV}}(\theta_i) + c_i^{\text{AGV}} - c_i^{\text{VCG}} = U_i^{\text{VCG}}(\theta_i) \geq 0.$$

Since $\sum_i d_i = 0$ and $\sum_i c_i^{\text{VCG}} \geq \sum_i c_i^{\text{AGV}}$, we have

$$d_1 = -\sum_{i>1} d_i = \sum_{i>1} (c_i^{\text{VCG}} - c_i^{\text{AGV}}) \geq c_1^{\text{AGV}} - c_1^{\text{VCG}}.$$

Thus

$$\bar{U}_1(\theta_1) = U_1^{\text{AGV}}(\theta_1) + d_1 \geq U_1^{\text{AGV}}(\theta_1) + c_1^{\text{AGV}} - c_1^{\text{VCG}} = U_1^{\text{VCG}}(\theta_1) \geq 0.$$

Therefore, \bar{t} is individually rational.

□

19.6 Dominant strategy mechanism

19.67 A direct mechanism (q, t) is dominant strategy incentive compatible if truth telling is a dominant strategy for each type of each buyer, that is, for all $i \in \mathcal{N}$, all θ, θ' and all θ_{-i} ,

$$\theta_i q_i(\theta_i, \theta_{-i}) - t_i(\theta_i, \theta_{-i}) \geq \theta_i q_i(\theta'_i, \theta_{-i}) - t_i(\theta'_i, \theta_{-i}).$$

19.68 A direct mechanism (q, t) is ex post individually rational if for each type of each buyer participation is a dominant strategy, that is, if for all $i \in I$, all θ_i and all θ_{-i} ,

$$\theta_i q_i(\theta_i, \theta_{-i}) - t_i(\theta_i, \theta_{-i}) \geq 0.$$

19.69 Proposition: A direct mechanism (q, t) is dominant strategy incentive compatible if and only if for every i and

every θ_{-i} ,

- (i) $q_i(\theta_i, \theta_{-i})$ is increasing in θ_i .
- (ii) for every θ_i ,

$$t_i(\theta_i, \theta_{-i}) = t_i(\underline{\theta}_i, \theta_{-i}) + \theta_i q_i(\theta_i, \theta_{-i}) - \underline{\theta}_i q_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} q_i(\theta'_i, \theta_{-i}) d\theta'_i.$$

19.70 Proposition: A dominant strategy incentive compatible direct mechanism (q, t) is ex post individually rational if and only if for every i and every θ_{-i} we have

$$t_i(\theta_i, \theta_{-i}) \leq \underline{\theta}_i q_i(\underline{\theta}_i, \theta_{-i}).$$

19.71 A direct mechanism (q, t) is called a canonical auction if there are strictly increasing and continuous functions $\psi_i: \Theta_i \rightarrow \mathbb{R}$ for i such that for all θ and i ,

$$q_i(\theta) = \begin{cases} \frac{1}{n}, & \text{if } \psi_i(\theta_i) \geq 0 \text{ and } \psi_i(\theta_i) \geq \psi_j(\theta_j) \text{ for all } j \neq i, \\ 0, & \text{otherwise,} \end{cases}$$

where n is the number of agents k such that $\psi_k(\theta_k) = \psi_i(\theta_i)$, and

$$t_i(\theta) = \begin{cases} \frac{1}{n} \min\{\theta'_i \mid q_i(\theta'_i, \theta_{-i}) > 0\}, & \text{if } q_i(\theta) > 0, \\ 0, & \text{otherwise,} \end{cases}$$

for all θ .

19.72 Proposition: Every canonical auction (q, t) is DSIC and ex post IR. Moreover, for every i , $u_i(\underline{\theta}_i, \theta_{-i}) = 0$ for all θ_{-i} .

19.7 Bilateral trade

19.73 A seller S owns a single indivisible good, and there is one potential buyer B .

19.74 The seller's utility if her valuation is θ_S and she receives a payment t is

$$(1 - d) \cdot \theta_S - t_S,$$

where $d = 0$ means “no trade” and $d = 1$ means “trade takes place”. Assume θ_S is a random variable on $\Theta_S = [0, 1]$ with cumulative distribution function F_S and density f_S .

19.75 The buyer's utility if her valuation is θ_B and she pays t is

$$d \cdot \theta_B - t_B,$$

where $d = 0$ means “no trade” and $d = 1$ means “trade takes place”. Assume θ_B is a random variable on $\Theta_B = [0, 1]$ with cumulative distribution function F_B and density f_B .

19.76 A direct mechanism (d, t) is efficient if

$$d(\theta) = d^*(\theta) \triangleq \begin{cases} 1, & \text{if } \theta_B \geq \theta_S, \\ 0, & \text{otherwise.} \end{cases}$$

Note that the efficiency is a special case in Definition 18.13.

19.77 A direct mechanism (d, t) is (ex post) budget balanced if for each θ

$$t_B(\theta) + t_S(\theta) = 0.$$

19.78 Let $D_i: \Theta_i \rightarrow [0, 1]$ be the probability that trade takes place conditioning on agent i 's type.

Let $T_i: \Theta_i \rightarrow \mathbb{R}$ be agent i 's expected transfer conditional on her type.

Let $U_i: \Theta_i \rightarrow \mathbb{R}$ be agent i 's expected utility conditional on her type:

$$\begin{aligned} U_S(\theta_S) &= (1 - D_S(\theta_S)) \cdot \theta_S - T_S(\theta_S), \\ U_B(\theta_B) &= D_B(\theta_B) \cdot \theta_B - T_B(\theta_B). \end{aligned}$$

19.79 A direct mechanism (d, t) is interim individually rational if

$$U_S(\theta_S) \geq \theta_S, \text{ and } U_B(\theta_B) \geq 0.$$

✂ 19.80 Theorem (Myerson and Satterthwaite (1983)): In the bilateral trade problem, there is no mechanism that is efficient, Bayesian incentive compatible, individual rational and budget balanced.

19.81 *Proof.*

(1) Consider an efficient and Bayesian incentive compatible mechanism (d^*, t) which is interim individually rational.

(2) By payoff equivalence theorem on interim expected payoff (Theorem 18.24), we have

$$\begin{aligned} T_S(\theta_S) &= T_S^{\text{pivot}}(\theta_S) - T_S^{\text{pivot}}(1) + T_S(1), \\ T_B(\theta_B) &= T_B^{\text{pivot}}(\theta_B) - T_B^{\text{pivot}}(0) + T_B(0). \end{aligned}$$

(3) By the definition of pivot mechanism, we have $U_S^{\text{pivot}}(1) = U_B^{\text{pivot}}(0) = 0$.

(4) Since (d^*, t) is interim individually rational, we have $U_S(1) \geq 1$ and $U_B(0) \geq 0$. Then

$$1 + T_S(1) = U_S(1) \geq 1, \text{ and } T_B(0) \geq 0.$$

(5) Thus,

$$T_S(\theta_S) \geq T_S^{\text{pivot}}(\theta_S), \text{ and } T_B(\theta_B) \geq T_B^{\text{pivot}}(\theta_B).$$

(6) By the definition of pivot mechanism, we have

$$t_S^{\text{pivot}}(\theta) = \begin{cases} -\theta_B, & \text{if } \theta_B \geq \theta_S \\ 0, & \text{otherwise} \end{cases}, \quad t_B^{\text{pivot}}(\theta) = \begin{cases} \theta_S, & \text{if } \theta_B \geq \theta_S \\ 0, & \text{otherwise} \end{cases}.$$

(7) This pivot mechanism runs an expected deficit of

$$\mathbf{E}_\theta [t_S^{\text{pivot}}(\theta_S, \theta_B) + t_B^{\text{pivot}}(\theta_S, \theta_B)] = \mathbf{E}_\theta [\max\{\theta_S - \theta_B, 0\}] > 0.$$

(8) Therefore, we have

$$\mathbf{E}_\theta [t_S(\theta) + t_B(\theta)] \geq \mathbf{E}_\theta [t_S^{\text{pivot}}(\theta_S, \theta_B) + t_B^{\text{pivot}}(\theta_S, \theta_B)] > 0,$$

which contradicts the budget balanced condition.

□

Mechanism design: Provision of public goods

Contents

20.1 Set-up	349
20.2 Ex ante and ex post budget balance	350
20.3 Welfare maximization	352

20.1 Set-up

20.1 A community consists of N agents: $\mathcal{N} = \{1, 2, \dots, N\}$ where $N \geq 2$. They have to choose whether to produce some indivisible, non-excludable public good. We denote this decision by $d \in \{0, 1\}$. If the public good is produced, then $d = 1$; otherwise $d = 0$.

20.2 Agent i 's utility is $\theta_i d - t_i$ if the decision is d and she pays a transfer t_i to the community. Here, θ_i is agent i 's valuation of the public good.

20.3 Agent i knows θ_i , but that neither the designer nor any other agent $j \neq i$ knows θ_i .

Agent's value θ_i is distributed over $[\underline{\theta}_i, \bar{\theta}_i]$ according to the cumulative distribution function F_i with associated density f_i , where $0 < \underline{\theta}_i < \bar{\theta}_i$. Although we do not assume that the random variables θ_i have the same distribution for different i , we do assume that they have the same support. This is for convenience only. For technical convenience, we also assume that $f_i(\theta_i) > 0$ for all $i \in \mathcal{N}$ and all $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$.

We assume that for $i, j \in \mathcal{N}$ with $i \neq j$, the random variables θ_i and θ_j are independent.

The distributions $\{F_i\}_{i \in \mathcal{N}}$ are common knowledge among the agents and the designer.

20.4 The cost of producing the public good is assumed to be $c > 0$, so that a decision d implies cost cd .

20.5 A direct mechanism consists of functions q and t

$$q: \Theta \rightarrow \{0, 1\} \text{ and } t: \Theta \rightarrow \mathbb{R}^N.$$

20.6 Given a direct mechanism (q, t) , define

$$Q_i(\theta_i) = \int_{\Theta_{-i}} q_i(\theta_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

to be the interim conditional probability that the public good is produced, conditioning on agent i 's type being θ_i .

Similarly, define

$$T_i(\theta_i) = \int_{\Theta_{-i}} t_i(\theta_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

to be the interim conditional expected value of the transfer that agent i makes to the community, conditioning on agent i 's type being θ_i .

Agent i 's expected utility is defined as

$$U_i(\theta_i) = Q_i(\theta_i) \cdot \theta_i - T_i(\theta_i),$$

conditioning on agent i 's type being θ_i .

20.7 The direct mechanism (q, t) is said to be Bayesian incentive compatible if truth telling is a Bayesian Nash equilibrium; that is, if for all i , for all θ_i and for all θ'_i ,

$$Q_i(\theta_i) \cdot \theta_i - T_i(\theta_i) \geq Q_i(\theta'_i) \cdot \theta_i - T_i(\theta'_i).$$

20.8 Proposition: A direct mechanism (q, t) is Bayesian incentive compatible if and only if for every i

- (i) Q_i is increasing;
- (ii) For every $\theta_i \in \Theta_i$,

$$T_i(\theta_i) = T_i(\underline{\theta}_i) + Q_i(\theta_i) \cdot \theta_i - \int_{\underline{\theta}_i}^{\theta_i} Q_i(\theta'_i) d\theta'_i.$$

20.9 A direct mechanism is individually rational if each agent, conditional on her type, is willing to participate, that is, if

$$U_i(\theta_i) \geq 0 \text{ for all } i \text{ and } \theta_i.$$

20.10 Proposition: A Bayesian incentive compatible direct mechanism (q, t) is individually rational if and only if for every i , we have

$$U_i(\underline{\theta}_i) \geq 0.$$

20.2 Ex ante and ex post budget balance

20.11 A direct mechanism is ex post budget balanced if for every θ we have

$$\sum_{i \in \mathcal{N}} t_i(\theta) \geq cq(\theta).$$

A direct mechanism is ex ante budget balanced if

$$\int_{\Theta} \sum_{i \in \mathcal{N}} t_i(\theta) f(\theta) d\theta \geq \int_{\Theta} cq(\theta) f(\theta) d\theta.$$

20.12 Clearly, ex post budget balance implies ex ante budget balance.

20.13 Two direct mechanisms are equivalent if they have the same decision rule, and if for all agents i and for all types θ_i, θ'_i , agent i 's expected transfers, conditional on agent i 's type being θ_i and agent i reporting to be type θ'_i , is the same in the two mechanisms.

20.14 Proposition: For every direct mechanism that is ex ante budget balanced, there is an equivalent direct mechanism that is ex post budget balanced.

Proof. Suppose first that we have a mechanism for which the ex ante budget balance condition holds with equality. Let t be the payments in the ex ante budget balanced mechanism, and $T_j(\theta_i)$ the expected value of agent j 's transfer, conditioning on agent i 's type being θ_i .

- (1) We construct the payments in the ex post budget balanced mechanism by modifying the payments in the ex post budget balanced mechanism.
- (2) In the modified mechanism, if the type profile is θ , then agent 1's payment is

$$t_1(\theta) + \underbrace{\left(cq(\theta) - \sum_{i \in \mathcal{N}} t_i(\theta) \right)}_{\text{coverage of ex post deficit}} - \underbrace{\left(cQ_1(\theta_1) - \sum_{i \in \mathcal{N}} T_i(\theta_1) \right)}_{\text{expected deficit conditional on her own type}}.$$

Then agent 1's expected transfer, conditional on agent 1's type being θ_1 and agent 1 reporting to be type θ'_1 is

$$\int_{\Theta_{-1}} t_1(\theta'_1, \theta_{-1}) d\theta_{-1} + \int_{\Theta_{-1}} cq(\theta'_1, \theta_{-1}) - \sum_{i \in \mathcal{N}} t_i(\theta'_1, \theta_{-1}) d\theta_{-1} - \int_{\Theta_{-1}} cQ_1(\theta'_1) - \sum_{i \in \mathcal{N}} T_i(\theta'_1) d\theta_{-1} = T_1(\theta'_1),$$

the same as the expected transfer in the original mechanism, conditional on agent 1's type being θ_1 and agent 1 reporting to be type θ'_1 .

- (3) In the modified mechanism, if the type profile is θ , then agent 2's payment is

$$t_2(\theta) + \underbrace{cQ_1(\theta_1) - \sum_{i \in \mathcal{N}} T_i(\theta_1)}_{\text{expected deficit conditional on agent 1's type}}.$$

Since the original mechanism is ex ante budget balanced,

$$\int_{\Theta_{-2}} cQ_1(\theta_1) - \sum_{i \in \mathcal{N}} T_i(\theta_1) d\theta_{-2} = \int_{\Theta} cq(\theta) - \sum_{i \in \mathcal{N}} t_i(\theta) d\theta = 0,$$

and hence agent 2's expected transfer, conditional on agent 2's type being θ_2 and agent 2 reporting to be type θ'_2 , is the same in the original mechanism.

- (4) The modified transfer for agent $i \neq 1, 2$ is the same as before, $t_i(\theta)$.
- (5) It is clear the modified transfers are ex post budget balanced.

If there is an ex ante budget surplus, we subtract from some agent's payments a constant until the mechanism is exactly budget balanced. Then we conduct the transformation described above. Then we add the constant again to this agent's payments. The mechanism that we obtain has the required properties. \square

20.15 Intuitively, the ex post budget balanced mechanism in the above proof is constructed so that agent 1 covers the ex post deficit, and agent 2 compensates agent 1 by paying to agent 1 a "fair" insurance premium.

20.16 We shall assume that the mechanism designer in this section considers only ex post budget balanced mechanisms. The above result makes clear that this is equivalent to requiring ex ante budget balance. We shall work with either condition, whichever is more convenient.

20.3 Welfare maximization

20.17 We attribute to the mechanism designer a utilitarian welfare function with equal welfare weights for all agents:

$$\sum_{i \in \mathcal{N}} \theta_i d - \sum_{i \in \mathcal{N}} t_i.$$

The mechanism designer's objective is to maximize the expected welfare.

20.18 As the mechanism designer subtracts transfer payments in her welfare function, she would never raise transfers larger than what is required to cover cost.

20.19 It is clear that the optimal decision rule is

$$q^*(\theta) = \begin{cases} 1, & \text{if } \sum_{i \in \mathcal{N}} \theta_i \geq c, \\ 0, & \text{otherwise.} \end{cases}$$

20.20 As the utilitarian designer with equal welfare weights for all agents does not care how costs are distributed among agents. Therefore, any transfer rules that satisfy

$$\sum_{i \in \mathcal{N}} t_i^*(\theta) = \begin{cases} c, & \text{if } q^*(\theta) = 1, \\ 0, & \text{otherwise,} \end{cases}$$

are optimal. We call these direct mechanisms “first best.” They maximize welfare for any type vector θ , and therefore they also maximize ex ante expected welfare.

20.21 Theorem: An incentive-compatible and individually rational first best mechanism exists if and only if either $\sum_{i \in \mathcal{N}} \underline{\theta}_i \geq c$ or $\sum_{i \in \mathcal{N}} \bar{\theta}_i \leq c$.

Proof of sufficiency. If $\sum_{i \in \mathcal{N}} \underline{\theta}_i \geq c$ then a mechanism where the public good is always produced and it is not difficult to find a first best, incentive-compatible and individually rational mechanism. If $\sum_{i \in \mathcal{N}} \bar{\theta}_i \leq c$ then a mechanism where the public good is never produced and no agent ever pays anything is first best and incentive-compatible and individually rational. \square

20.22 The pivot mechanism is the mechanism that is given by the first best decision rule q^* and by the following transfer scheme:

$$t_i(\theta) = \underline{\theta}_i q^*(\underline{\theta}_i, \theta_{-i}) + (q^*(\theta) - q^*(\underline{\theta}_i, \theta_{-i})) \left(c - \sum_{j \neq i} \theta_j \right)$$

for all i and θ .

20.23 The second term equals the change to the social welfare of all other agents caused by agent i 's report, and hence agent i pays only if her report is “pivotal” for the collective decision. Here we compare the actual outcome to the outcome that would have occurred had agent i reported the lowest type $\underline{\theta}_i$. Agent i 's report changes the collective

decision if $q^*(\theta) - q^*(\underline{\theta}_i, \theta_{-i}) = 1$. In that case, agent i pays for the difference between the costs of the project, and the sum of all other agents' valuations of the project.

20.24 Lemma: The pivot mechanism is incentive compatible and individually rational.

Proof. (1) We are going to show that truthful reporting is optimal whatever the other agents' types θ_{-i} are.

(2) If we leave out terms that do not depend on agent i 's report, then agent i 's utility if reporting θ'_i is:

$$\theta_i q^*(\theta'_i, \theta_{-i}) - q^*(\theta'_i, \theta_{-i}) \left(c - \sum_{j \neq i} \theta_j \right) = q^*(\theta'_i, \theta_{-i}) \left(\sum_{i \in \mathcal{N}} \theta_i - c \right).$$

(3) Because q^* is first best, agent i 's utility is maximized if she reports truthfully θ_i . This proves incentive compatibility.

(4) To verify individual rationality note that the expected utility of agent i obviously equals zero if her type is $\underline{\theta}_i$.

(5) By a result analogous to Lemma 19.20, all types' interim expected utility is at least zero, and the mechanism is individually rational. □

20.25 Lemma: No incentive compatible and individually rational mechanism that implements the first best decision rule q^* has larger ex ante expected budget surplus (ex ante expected transfers as well), that is revenue minus cost, than the pivot mechanism.

Proof. (1) The ex ante expected budget surplus from an incentive compatible mechanism is

$$\sum_{i \in \mathcal{N}} T_i(\underline{\theta}_i) - \sum_{i \in \mathcal{N}} \underline{\theta}_i Q_i(\underline{\theta}_i) + \int_{\Theta} q(\theta) \left[\sum_{i \in \mathcal{N}} \left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) - c \right] f(\theta) d\theta.$$

(2) Thus, the ex ante expected budget surplus from an incentive compatible mechanism that implements the first best decision rule q^* equals the interim expected payments of the lowest types plus a term that is the same for all such rules.

(3) If a mechanism is individually rational then the interim expected payments of the lowest types can be at most such that the expected utility of the lowest types are zero, that is, $\underline{\theta}_i Q_i(\underline{\theta}_i)$.

(4) For the pivot mechanism the expected utilities of the lowest types are exactly equal to zero.

(5) Therefore, no incentive-compatible, individually rational direct mechanism can have higher expected surplus than the pivot mechanism. □

20.26 *Proof of necessity.* (1) Assume that $\sum_{i \in \mathcal{N}} \underline{\theta}_i < c < \sum_{i \in \mathcal{N}} \bar{\theta}_i$. We show that the ex post budget surplus of the pivot mechanism is always non-positive and with positive probability negative. This implies that the ex ante expected budget surplus is negative.

(2) Consider first θ such that $q^*(\theta) = 0$. In this case, there are no costs and no agent pays any transfer. Hence the deficit is zero.

(3) Consider next θ such that $q^*(\theta) = 1$ and $q^*(\underline{\theta}_i, \theta_{-i}) = 1$ for every $i \in \mathcal{N}$. In this case each agent pays $\underline{\theta}_i$. However, $\sum_{i \in \mathcal{N}} \underline{\theta}_i < c$, and hence there is a deficit.

- (4) Consider finally θ such that $q^*(\theta) = 1$, and $q^*(\underline{\theta}_i, \theta_{-i}) = 0$ for some i . Let J be the set of all i for which this holds, and call these agents “pivotal.” Then the total transfer is

$$\sum_{i \in J} \left(c - \sum_{j \neq i} \theta_j \right) + \sum_{i \in J^c} \underline{\theta}_i = |J|c - (|J| - 1) \sum_{i \in \mathcal{N}} \theta_i - \sum_{i \in J^c} (\theta_i - \underline{\theta}_i).$$

Since $q^*(\theta) = 1$, $c \leq \sum_{i \in \mathcal{N}} \theta_i$, and hence the total transfer is at most

$$|J|c - (|J| - 1)c - \sum_{i \in J^c} (\theta_i - \underline{\theta}_i) = c - \sum_{i \in J^c} (\theta_i - \underline{\theta}_i).$$

Therefore, the total transfer is not more than c .

- (5) Since the pivot mechanism has larger ex ante expected budget surplus among all IC, IR and first best mechanisms, and the ex ante expected budget surplus is negative, no incentive compatible and individually rational first best mechanism exists when $\sum_{i \in \mathcal{N}} \underline{\theta}_i < c < \sum_{i \in \mathcal{N}} \bar{\theta}_i$.

□

Implementation theory

Contents

21.1 Implementation	355
21.2 Implementation in dominant strategies	357
21.3 Nash implementation	360

21.1 Implementation

21.1 The essential part of mechanism design is implementation theory which, given a social goal, characterizes when we can design a mechanism whose predicted outcomes (*e.g.* equilibrium outcomes) coincide with the desirable outcomes, according to that goal.

21.2 An example: Consider a society consisting of two consumers of energy, Alice and Bob. An energy authority is charged with choosing the type of energy to be used by Alice and Bob. The options—from which the authority must make a single selection—are gas, oil, nuclear power, and coal.

Let us suppose that there are two possible states of the world. In state 1, the consumers place relatively little weight on the future, *i.e.*, they have comparatively high temporal discount rates. In state 2, by contrast, they attach a great deal of importance to the future, meaning that their rates of discount are correspondingly low.

In each state, the consumers’ rankings in the two states are given in Table 21.1.

Assume that the energy authority is interested in selecting an energy source that both consumers are reasonably happy with. If we interpret “reasonably happy” as getting one’s first or second choice, then oil is the optimal choice in state 1, whereas gas is the best outcome in state 2. In the language of implementation theory, we say that the authority’s social choice rule prescribes oil in state 1 and gas in state 2. Thus, if f is the choice rule, it is given in Table 21.1.

Suppose, however, that the authority does not know the state (although Alice and Bob do). This means that it does not know which alternative the choice rule prescribes, *i.e.*, whether oil or gas is the optimum.

- Probably the most straightforward mechanism would be for the authority to ask each consumer to announce $\tau_A, \tau_B \in \{\theta_1, \theta_2\}$ respectively. After receiving their reports, the following mechanism will produce the out-

state	$\theta_1 = \succsim$		$\theta_2 = \succsim'$	
social goal	$f(\theta_1) = \text{oil}$		$f(\theta_2) = \text{gas}$	
preference	Alice	Bob	Alice	Bob
	gas	nuclear	nuclear	oil
	oil	oil	gas	gas
	coal	coal	coal	coal
	nuclear	gas	oil	nuclear

Table 21.1

comes as follows:

$$g(\tau_A, \tau_B) = \begin{cases} \text{oil}, & \tau_A = \tau_B = \theta_1, \\ \text{gas}, & \tau_A = \tau_B = \theta_2, \\ \frac{1}{2} \circ \text{oil} + \frac{1}{2} \circ \text{gas}, & \text{otherwise.} \end{cases}$$

Note: Alice/Bob would always report states θ_2/θ_1 (respectively).

- If Alice reports θ_1 , she may obtain “oil”, and if she reports θ_2 , she may get “gas”. “gas” is better than “oil” for Alice no matter the true state is θ_1 or θ_2 .
- If Bob reports θ_1 , he may obtain “oil”, and if he reports θ_2 , he may get “gas”. “oil” is better than “gas” for Bob no matter the true state is θ_1 or θ_2 .

Thus this mechanism implements the social goal only with a 50 percent chance.

- The following mechanism, which specifies set of players, set of actions for each player and outcome function, can implement the social goal.

		Bob				Bob				Bob	
		L	R			L	R			L	R
Alice	U	oil	coal	Alice	U	3, 3	2, 2	Alice	U	1, 4	2, 2
	D	nuclear	gas		D	1, 4	4, 1		D	4, 1	3, 3
Outcome function				Preferences in θ_1				Preferences in θ_2			

Figure 21.1

It is clear the Nash equilibrium outcome is “oil” when the true state is θ_1 and “gas” when the true state is θ_2 .

21.3 Definition: An environment $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ consists of

- a finite set N of players, with $|N| \geq 2$,
- a set C of outcomes,
- a set \mathcal{P} of preference profiles over C , with typical profile $\succsim \in \mathcal{P}$,
- a set \mathcal{G} of game forms/mechanisms with consequences in C .
 - A strategic game form is a triple $G = \langle N, (A_i), g \rangle$, where $g: A \rightarrow C$ is an outcome function.
 - An extensive game form is a tuple $\langle N, H, P, g \rangle$, where $g: Z \rightarrow C$ is an outcome function.

A strategic game form $G = \langle N, (A_i), g \rangle$ with a preference profile \succsim induce a strategic game $\langle N, (A_i), \succsim' \rangle$, where $a \succsim'_i b$ if and only if $g(a) \succsim_i g(b)$.

Similarly, an extensive game form with a preference profile induce an extensive game.

The game form can be regarded as the rules of the game.

21.4 Definition: A choice rule $f: \mathcal{P} \rightarrow C$ is a function that assigns a subset of C to each preference profile in \mathcal{P} .

We refer to a choice rule that is a singleton-valued as a choice function.

21.5 Notation: For a game form $G \in \mathcal{G}$ and a preference profile $\succsim \in \mathcal{P}$, let $\mathcal{S}(G, \succsim)$ denote the set of solutions under the solution concept \mathcal{S} of the game induced by G and \succsim .

21.6 Definition: Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment and \mathcal{S} a solution concept.

The game form $G \in \mathcal{G}$ with outcome function g is said to \mathcal{S} -implement the choice rule $f: \mathcal{P} \rightarrow C$ if

$$g(\mathcal{S}(G, \succsim)) = f(\succsim) \text{ for all } \succsim \in \mathcal{P}.$$




Figure 21.2

We say the choice rule f is \mathcal{S} -implementable in $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ if there exists $G \in \mathcal{G}$ with outcome function g which \mathcal{S} -implements f .

21.7 Definition: Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which \mathcal{G} is a set of strategic game forms for which the set of actions of each player i is the set \mathcal{P} of preference profiles, and \mathcal{S} a solution concept.

The strategic game form $G = \langle N, (A_i), g \rangle \in \mathcal{G}$ truthfully \mathcal{S} -implements the choice rule $f: \mathcal{P} \rightarrow C$ if for every preference profile $\succsim \in \mathcal{P}$ we have

- every player reporting the “true” preference profile is a solution of the game:

$$a^* \in \mathcal{S}(G, \succsim), \quad a_i^* = \succsim \text{ for each } i \in N.$$

- the outcome if every player reports the true preference profile is a member of $f(\succsim)$:

$$g(a^*) \in f(\succsim), \quad a_i^* = \succsim \text{ for each } i \in N.$$

The choice rule f is truthfully \mathcal{S} -implementable in $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ if there is a strategic game form $G \in \mathcal{G}$ with outcome function g , such that G truthfully \mathcal{S} -implements f .

21.8 Remark: One important aspect of implementation theory is the requirement that all the solution outcomes lie in the given choice rule.

The mechanism design literature focuses on incentive compatibility issues, asking whether a given outcome can be induced as an equilibrium of some mechanism and generally ignores whether there are other undesired equilibrium outcomes.

21.2 Implementation in dominant strategies

21.9 Definition: A profile $a^* \in A$ in strategic game $\langle N, (A_i), (\succsim_i) \rangle$ is a dominant-strategy equilibrium (abbreviated as “DSE”) if, for every player i , $(a_{-i}, a_i^*) \succsim_i (a_{-i}, a_i)$ for all $a \in A$.

21.10 Lemma (revelation principle for DSE-implementation): Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which \mathcal{G} is the set of strategic game forms. If $f: \mathcal{P} \rightarrow C$ is DSE-implementable, then

- f is truthfully DSE-implementable.
- there is a strategic game form $G^* = \langle N, (A_i^*), g^* \rangle \in \mathcal{G}$ in which A_i^* is \mathcal{P}_i —the set of all preferences for player i , such that for all $\succsim \in \mathcal{P}$ the action profile \succsim is a dominant-strategy equilibrium of $\langle G^*, \succsim \rangle$ and $g^*(\succsim) \in f(\succsim)$.

21.11 *Proof of the second statement.* (1) Let the strategic game form $G = \langle N, (A_i), g \rangle$ DSE-implement f . Then for all $\succsim \in \mathcal{P}$, $g(\text{DSE}(G, \succsim)) = f(\succsim)$.

(2) For each $i \in N$, the set of dominant strategies of player i depends only on \succsim_i , so that for any $\succsim \in \mathcal{P}$ we can define $a_i(\succsim_i)$ to be a dominant strategy for player i in the game $\langle G, \succsim \rangle$ induced by G and \succsim .

(3) Clearly, $(a_i(\succsim_i))_{i \in N}$ is a dominant-strategy equilibrium in the game $\langle G, \succsim \rangle$, and

$$g\left((a_i(\succsim_i))_{i \in N}\right) \in f(\succsim).$$

(4) For each $i \in N$, let $A_i^* = \mathcal{P}_i$ and $g^*: A^* \rightarrow C$,

$$g^*(\succsim) = g\left((a_i(\succsim_i))_{i \in N}\right).$$

It is clear that $g^*(\succsim) \in f(\succsim)$.




Figure 21.3

(5) Now suppose that there is a preference profile \succsim for which \succsim_j is not a dominant strategy for player j in the game $\langle G^*, \succsim \rangle$ induced by G^* and \succsim . Then there is a preference profile \succsim' such that

$$g^*(\succsim'_{-j}, \succsim'_j) \succ_j g^*(\succsim'_{-j}, \succsim_j),$$

and hence

$$g(a_{-j}(\succsim'_{-j}), a_j(\succsim'_j)) \succ_j g(a_{-j}(\succsim'_{-j}), a_j(\succsim_j)).$$

That is, $a_j(\succsim_j)$ is not a dominant strategy of player j in $\langle G, \succsim \rangle$, a contradiction.

□

21.12 *Proof of the first statement.* (1) Define a new strategic game form $G' = \langle N, (A'_i), g' \rangle$ with $A'_i = \mathcal{P}$ and

$$g'((\succsim^i)_{i \in N}) = g^*((\succsim_i)_{i \in N}),$$

where $\succsim^i \in A'_i = \mathcal{P}$.




Figure 21.4

- (2) Fix a preference profile \succsim . For any $j \in N$ and any strategy profile $(\succsim^i)_{i \in N} \in A' = \mathcal{P}^{|N|}$, since $((\succsim_k^k)_{k \neq j}, \succsim_j)$ is a dominant-strategy equilibrium of $\langle G^*, ((\succsim_k^k)_{k \neq j}, \succsim_j) \rangle$, we have

$$g'((\succsim^k)_{k \neq j}, \succsim) = g^*((\succsim_k^k)_{k \neq j}, \succsim_j) \succsim_j g^*((\succsim_k^k)_{k \neq j}, \succsim_j^j) = g'((\succsim^k)_{k \neq j}, \succsim^j),$$

that is, $(\underbrace{\succsim, \dots, \succsim}_{|N| \text{ terms}})$ is a dominant-strategy equilibrium of $\langle G', \succsim \rangle$.

- (3) Moreover, it is trivial that $g'(\underbrace{(\succsim, \dots, \succsim)}_{|N| \text{ terms}}) = g^*((\succsim_i)_{i \in N}) \in f(\succsim)$.

- (4) Therefore, f is truthfully DSE-implementable. □

21.13 Definition: We say that a choice rule $f: \mathcal{P} \rightarrow C$ is dictatorial if there is a player $j \in N$ such that for any preference profile $\succsim \in \mathcal{P}$ and outcome $c \in f(\succsim)$ we have $c \succsim_j c'$ for all $c' \in C$.

21.14 Gibbard-Satterthwaite theorem: Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which C contains at least 3 elements, \mathcal{P} is the set of all possible preference profiles, and \mathcal{G} is the set of strategic game forms. Let $f: \mathcal{P} \rightarrow C$ be DSE-implementable and satisfy the condition:

$$\text{for every } c \in C, \text{ there exists } \succsim \in \mathcal{P}, \text{ such that } f(\succsim) = \{c\}. \quad (21.1)$$

Then f is dictatorial.

21.15 Lemma (Gibbard-Satterthwaite theorem in social choice theory): Let C be a set that contains at least three members and let \mathcal{P} be the set of all possible preference profiles. If a choice function $f: \mathcal{P} \rightarrow C$ satisfies Equation (21.1) and for every preference profile $\succsim \in \mathcal{P}$ we have $f(\succsim_{-j}, \succsim_j) \succsim_j f(\succsim_{-j}, \succsim'_j)$ for every preference relation \succsim'_j then f is dictatorial.

21.16 Proof. Since f is DSE-implementable by the game form G , any selection g^* of f (i.e., $g^*(\succsim) \in f(\succsim)$ for all $\succsim \in \mathcal{P}$) has the property that for every preference profile \succsim we have

$$g^*(\succsim_{-j}, \succsim_j) \succsim_j g^*(\succsim_{-j}, \succsim'_j)$$

for every preference relation \succsim'_j .

Since f satisfies Equation (21.1), g^* does also. Consequently by the lemma above g^* is dictatorial, so that f is also. □

21.3 Nash implementation

21.17 Nash-implementable: $g(\text{NE}(G, \succsim)) = f(\succsim)$ for all \succsim .

Truthfully Nash-implementable: $g^*(a^*) \in f(\succsim) \cap g^*(\text{NE}(G^*, \succsim))$ for all \succsim , where $a_i^* = \succsim_i$ for all i .

21.18 Lemma (revelation principle for Nash implementation): Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which \mathcal{G} is the set of strategic game forms. If a choice rule is Nash-implementable then it is truthfully Nash-implementable.

21.19 *Proof.* (1) Let $G = \langle N, (A_i), g \rangle$ Nash-implement f . Let $(a_i(\succsim))_{i \in N}$ be a Nash equilibrium in the game $\langle G, \succsim \rangle$.
 (2) Define $G^* = \langle N, (A_i^*), g^* \rangle$ in which $A_i^* = \mathcal{P}$ and $g^*((\succsim^i)_{i \in N}) = g((a_i(\succsim^i))_{i \in N})$ where $\succsim^i \in A_i^*$.




Figure 21.5

(3) Clearly, $a^* = (\succsim, \dots, \succsim)$ is a Nash equilibrium of the game $\langle G^*, \succsim \rangle$:

$$g^*(\succsim, \dots, \succsim) = g(a_{-i}(\succsim), a_i(\succsim)) \succsim_i g(a_{-i}(\succsim), a_i(\succsim')) = g^*(\succsim^{-i}, \succsim') \text{ for all } \succsim' \in A_i^* = \mathcal{P}.$$

(4) We also have $g^*(a^*) = g((a_i(\succsim))_{i \in N}) \in f(\succsim)$, and hence f is truthfully Nash-implementable. □

21.20 Definition (Maskin's monotonicity): A choice rule $f: \mathcal{P} \rightarrow C$ is monotonic if whenever $c \in f(\succsim)$ and $c \notin f(\succsim')$ there is some player $i \in N$ and some outcome $b \in C$ such that $c \succsim_i b$ but $b \succ'_i c$. (Compare with the monotonicity in social choice theory)

In other words, if $c \in f(\succsim)$ does not fall in anyone's ranking relative to any other alternative in going from \succsim to \succsim' , monotonicity requires that $c \in f(\succsim')$.

21.21 Proposition: Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which \mathcal{G} is the set of strategic game forms. If a choice rule is Nash-implementable then it is monotonic.

Proof. (1) Let the strategic game form $G = \langle N, (A_i), g \rangle$ Nash-implement f .

(2) Let $c \in f(\succsim)$ and $c \notin f(\succsim')$.

(3) Since $g(\text{NE}(G, \succsim)) = f(\succsim)$ for all \succsim , there is a strategy profile a , such that

- $g(a) = c$,
- a is a Nash equilibrium of the game $\langle G, \succsim \rangle$,
- a is not a Nash equilibrium of the game $\langle G, \succsim' \rangle$.

(4) Therefore, there are a player i and some strategy a'_i such that

$$b = g(a_{-i}, a'_i) \succ'_i g(a) = c \text{ and } c = g(a) \succsim_i g(a_{-i}, a'_i) = b.$$

□

state	$\theta_1 = \succsim$		$\theta_2 = \succsim'$	
social goal	$f(\theta_1) = \text{oil}$		$f(\theta_2) = \text{nuclear}$	
preference	Alice	Bob	Alice	Bob
	gas	nuclear	gas	nuclear
	oil	oil	oil	oil
	coal	coal	nuclear	coal
	nuclear	gas	coal	gas

Table 21.2

21.22 Example:

This choice rule f does not satisfy Maskin's monotonicity:

- $\text{oil} = f(\theta_1)$ and $\text{oil} \neq f(\theta_2)$;
- for player 1, "oil" is always better than "nuclear" and "coal", and worse than "gas" in both states;
- for player 2, "oil" is always better than "coal" and "gas", and worse than "nuclear" in both states.

Hence no mechanism NE-implements f .

21.23 Example:

state	$\theta_1 = \succsim$		$\theta_2 = \succsim'$	
social goal	$f(\theta_1) = \text{oil}$		$f(\theta_2) = \text{gas}$	
preference	Alice	Bob	Alice	Bob
	gas	nuclear	nuclear	oil
	oil	oil	gas	gas
	coal	coal	coal	coal
	nuclear	gas	oil	nuclear

Table 21.3

This choice rule f satisfies Maskin's monotonicity.

21.24 Example: Solomon's predicament

"If it please you, my lord," one of the women said, "this woman and I live in the same house, and while she was in the house I gave birth to a child. Now it happened on the third day after my delivery that this woman also gave birth to a child. We were alone together; there was no one else in the house with us....Now one night this woman's son died....And in the middle of the night she got up and took my son from beside me while I was asleep; she put him to her breast and put her own dead son to mine. While I got up to suckle my child, there he was, dead. But in the morning I looked at him carefully, and he was not the child I had borne at all." Then the other woman spoke. "That is not true! My son is the live one, yours is the dead one"; and the first retorted, "That is not true! Your son is the dead one, mine is the live one." And so they wrangled before the king...."Bring me a sword" said the king; and a sword was brought into the king's presence. "Cut the living child in two," the king said "and give half to one, half to the other." At this the woman who was the mother of the living child addressed the king, for she burned with pity for her son. "If it please you, my lord," she said "let them give her the child; only do not let them think of killing it!" But the other said, "He shall belong to neither of us. Cut him up." Then the king gave his decision. "Give the child to the first woman," he said "and do not kill him. She is his mother." All Israel came to hear of the judgement the king had pronounced, and held the king in awe, recognising that he possessed divine wisdom for dispensing justice.

Each of two women claims a baby; each knows who is the true mother, but neither can prove her motherhood. Solomon tries to seduce the truth by threatening to cut the baby in two, relying on the fact that the false mother prefers this outcome to that in which the true mother obtains the baby while the true mother prefers to give the baby away than to see it cut in two.

Let a , b and d denote the outcomes “the baby is given to mother 1”, “the baby is given to mother 2” and “the baby is cut in two” respectively.

Let θ_1 and θ_2 denote states “mother 1 is the true mother” and “mother 2 is the true mother” respectively.

state	$\theta_1 = \succsim$		$\theta_2 = \succsim'$	
social goal	$f(\theta_1) = a$		$f(\theta_2) = b$	
preference	mother 1	mother 2	mother 1	mother 2
	a	b	a	b
	b	d	d	a
	d	a	b	d

Table 21.4

This choice rule f does not satisfy Maskin’s monotonicity:

- $a = f(\theta_1) \neq f(\theta_2)$;
- For mother 1, a is always better than b and d no matter the state is θ_1 or θ_2 ;
- For mother 2, a is worse than b and d when state is θ_1 .

21.25 Remark on Solomon’s predicament. In the biblical story Solomon succeeds in assigning the baby to the true mother: he gives it to the only woman to announce that she prefers that it be given to the other woman than be cut in two (*i.e.*, one says “don’t cut”). But, Solomon’s idea does not work from a game-theoretic view.

Indeed, the following Solomon’s mechanism g can not Nash implement f .

		mother 2	
		cut	don’t
mother 1	cut	d	b
	don’t	a	d

Figure 21.6

Clearly, $g[\text{NE}(\theta_1)] = b \neq f(\theta_1)$ and $g[\text{NE}(\theta_2)] = a \neq f(\theta_2)$.

21.26 Definition: A choice rule $f: \mathcal{P} \rightarrow C$ has no veto power (abbreviated as “NVP”) if $c \in f(\succsim)$ whenever for at least $|N| - 1$ players we have $c \succsim_i y$ for all $y \in C$.

21.27 Proposition (sufficient condition for Nash implementation): Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which \mathcal{G} is the set of strategic game forms. If $|N| \geq 3$ then any choice rule that is monotonic and has no veto power is Nash-implementable.

21.28 Proof. (1) Let $f: \mathcal{P} \rightarrow C$ be a monotonic choice rule that has no veto power.

(2) Define a game form $G = \langle N, (A_i), g \rangle$ where $A_i = \mathcal{P} \times C \times \mathbb{N}$ and outcome function $g((p_i, c_i, m_i)_{i \in N})$ as follows:

If for some $j \in N$, we have $(p_i, c_i, m_i) = (\succsim, c, m)$ with $c \in f(\succsim)$ for all $i \neq j$ then

$$g((p_i, c_i, m_i)_{i \in N}) = \begin{cases} c_j, & \text{if } c \succsim_j c_j, \\ c, & \text{otherwise.} \end{cases}$$

Otherwise $g((p_i, c_i, m_i)_{i \in N}) = c_k$ where k is such that $m_k \geq m_i$ for all $i \in N$.

- (3) Let $c \in f(\succsim)$. Define $a_i^* = (\succsim, c, 1)$ for all $i \in N$. Then a^* is a Nash equilibrium of $\langle G, \succsim \rangle$ with $g(a^*) = c$: since any deviation by j , say (\succsim', c', m') , affects the outcome only if the outcome is $c \succ_j c'$.

Thus $f(\succsim) \subseteq g(\text{NE}(G, \succsim))$.

- (4) Let $a^* \in \text{NE}(G, \succsim)$. We shall show $c^* = g(a^*) \in f(\succsim)$.

- (5) Case 1: Suppose $a_i^* = (\succsim', c^*, m')$ for all $i \in N$ and $c^* \in f(\succsim')$.

If $c^* \notin f(\succsim)$ then the monotonicity of f implies that there is some player j and some outcome $b \in C$ such that $c^* \succsim'_j b$ and $b \succ_j c^*$.

Hence, j would have a profitable deviation $(\succsim, b, 1)$ from the Nash equilibrium a^* :

- j will get c^* when he chooses a_j^* .
- j will get b when he chooses $(\succsim, b, 1)$ since $c^* \succsim'_j b$.
- For player j , $b \succ_j c^*$.

- (6) Case 2: Suppose $a_i^* = (\succsim', c^*, m')$ for all $i \in N$ but $c^* \notin f(\succsim')$.

If there is some j and $b \in C$ such that $b \succ_j c^*$ then j would like to deviate to (\succsim, b, m'') with $m'' > m'$:

- j will get c^* when he chooses a_j^* .
- j will get b when he chooses (\succsim, b, m'') since $c^* \notin f(\succsim)$ and $m'' > m'$.
- For j , $b \succ_j c^*$.

Thus $c^* \succsim_i y$ for all $i \in N$ and $y \in C$, and hence by NVP $c^* \in f(\succsim)$.

- (7) Case 3: Suppose $a_i^* \neq a_j^*$ for some distinct i and j .

Since $|N| \geq 3$, if there is an outcome b such that $b \succ_k c^*$ for some $k \neq i$, then k would have a profitable deviation (\succsim', b, m'') with $m'' > m_l$ for all $l \neq k$:

- k will get c^* if he chooses a_k^* .
- k will get b if he chooses (\succsim', b, m'') since $m'' > m_l$ for all $l \neq k$.
- $b \succ_k c^*$.

Thus for all $k \neq i$ we have $c^* \succsim_k b$ for all $b \in C$ and, by NVP $c^* \in f(\succsim)$.

□

Chapter

22

Coalitional games

Contents

22.1 Coalitional game	365
22.2 Core	366
22.3 Shapley value	371
22.4 Nash bargaining solution	376

22.1 Coalitional game

22.1 A coalitional game with transferable payoff (henceforth “coalitional game”) $\langle N, v \rangle$ consists of

- a finite set N of players,
- a function $v: 2^N \setminus \{\emptyset\} \rightarrow \mathbb{R}$.

Every member in $2^N \setminus \{\emptyset\}$ is called a coalition, and $v(S)$ is called the worth of the coalition S . The function is called the characteristic function.

22.2 In a coalitional game each coalition S is characterized by a single number $v(S)$, with the interpretation that $v(S)$ is a payoff that may be distributed in any way among the members of S .

There is a more general concept, in which each coalition can not necessarily achieve all distributions of some fixed payoff; rather, each coalition S is characterized by an arbitrary set $V(S)$ of consequences. This concept is called a coalitional game without transferable payoff.

22.3 Convention: If $\{i_1, i_2, \dots, i_j\}$ is a set of players, we will sometimes write $v(i_1, i_2, \dots, i_j)$ rather than $v(\{i_1, i_2, \dots, i_j\})$ for the worth of $\{i_1, i_2, \dots, i_j\}$.

22.4 A coalitional game $\langle N, v \rangle$ is

- monotonic if $T \subseteq S$ implies $v(S) \geq v(T)$;
- cohesive if

$$v(N) \geq \sum_{k=1}^K v(S_k) \text{ for every partition } \{S_1, S_2, \dots, S_K\} \text{ of } N;$$

- super-additive if $S \cap T = \emptyset$ implies $v(S \cup T) \geq v(S) + v(T)$.

It is clear that $\langle N, v \rangle$ is cohesive if it is super-additive.

We will assume that the coalitional games are cohesive or super-additive.

22.5 A coalitional game $\langle N, v \rangle$ is simple if $v(S)$ is either 0 or 1 for any coalition S .

S is a winning coalition in a simple game if $v(S) = 1$; a veto player in such a game is a player who is a member of every winning coalition.

22.6 A coalitional game $\langle N, v \rangle$ is 0-normalized if $v(i) = 0$ for all $i \in N$; it is 0-1 normalized if it is 0-normalized and $v(N) = 1$.

22.7 i and j , elements of N , are substitutes in v if for all S containing neither i nor j , $v(S \cup \{i\}) = v(S \cup \{j\})$.

22.8 $i \in N$ is called a null player if $v(S \cup \{i\}) = v(S)$ for all $S \subseteq N$.

22.2 Core

22.9 The core is a solution concept for coalitional games that requires that no set of players be able to break away and take a joint action that makes all of them better off.

22.10 Let $\langle N, v \rangle$ be a coalitional game.

A vector $(x_i)_{i \in S}$ of real numbers is an S -feasible payoff vector if $v(S) = \sum_{i \in S} x_i$.

We refer to an N -feasible payoff vector as a feasible payoff profile.

22.11 The core of the coalitional game $\langle N, v \rangle$ is the set of feasible payoff profiles $(x_i)_{i \in N}$ for which there is no coalition S and S -feasible payoff vector $(y_i)_{i \in S}$ for which $y_i > x_i$ for all $i \in S$.

22.12 If x is in the core, then x satisfies

- (individual rational) $x_i \geq v(i)$ for all $i \in N$,
- (group rational) $\sum_{i \in N} x_i = v(N)$.

22.13 A definition that is obviously equivalent is that the core is the set of feasible payoff profiles $(x_i)_{i \in N}$ for which $\sum_{i \in S} x_i \geq v(S)$ for every coalition S .

Proof. “ \Leftarrow ”: Suppose that $x = (x_i)_{i \in N}$ satisfies

$$\sum_{i \in N} x_i = v(N), \text{ and } \sum_{i \in S} x_i \geq v(S) \text{ for all coalition } S.$$

Assume x is not in the core, that is, there exist a coalition S and $y = (y_i)_{i \in S}$, such that $\sum_{i \in S} y_i = v(S)$ and $y_i > x_i$ for all $i \in S$. Then we have $\sum_{i \in S} y_i > \sum_{i \in S} x_i \geq v(S)$, a contradiction.

“ \Rightarrow ”: Suppose that $x = (x_i)_{i \in N}$ does not satisfy

$$\sum_{i \in N} x_i = v(N), \text{ and } \sum_{i \in S} x_i \geq v(S) \text{ for all coalition } S.$$

If $\sum_{i \in N} x_i \neq v(N)$, x can not be in the core.

Suppose, then, that there is a coalition S such that

$$\sum_{i \in S} x_i = v(S) - \epsilon,$$

where $\epsilon > 0$. For $i \in S$, define

$$z_i = x_i + \frac{\epsilon}{|S|}.$$

It is easily seen that $\sum_{i \in S} z_i = v(S)$ and $z_i > x_i$ for all $i \in S$. Hence x is not in the core. \square

22.14 The core is the set of payoff profiles satisfying a system of weak linear inequalities and hence is closed and convex.

22.15 Example: Two-person bargaining game.

$$N = \{1, 2\}, v(N) = 1, \text{ and } v(1) = v(2) = 0.$$

Answer. (x_1, x_2) is in the core if and only if

$$x_1 \geq 0, x_2 \geq 0, \text{ and } x_1 + x_2 = 1.$$

\square

22.16 Example: Three-person bargaining game.

$$N = \{1, 2, 3\}, v(N) = 1 \text{ and } v(S) = 0 \text{ for all } S \subsetneq N.$$

Answer. (x_1, x_2, x_3) is in the core if and only if

$$x_1 + x_2 + x_3 = v(N) = 1, \text{ and } \sum_{i \in S} x_i \geq v(S) = 0 \text{ for all } S \subsetneq N.$$

The core is therefore the set

$$\{(x_1, x_2, x_3) \mid x_1, x_2, x_3 \geq 0, x_1 + x_2 + x_3 = 1\}.$$

\square

22.17 Example: Market with two sellers and a buyer.

$$N = \{1, 2, 3\}, v(N) = v(1, 2) = v(1, 3) = 1, \text{ and } v(S) = 0 \text{ for all other } S \subseteq N.$$

Answer. x is in the core if and only if

$$x_1 + x_2 + x_3 = 1, x_1 + x_2 \geq 1, x_1 + x_3 \geq 1, x_1 \geq 0, x_2 \geq 0, x_3 \geq 0.$$

Hence the core is $\{(1, 0, 0)\}$. \square

22.18 Example: Three-person majority game.

Suppose that three players can obtain one unit of payoff, any two of them can obtain 1 independently of the actions of the third, and each player alone can obtain nothing, independently of the actions of the remaining two players.

$$N = \{1, 2, 3\}, v(N) = v(1, 2) = v(1, 3) = v(2, 3) = 1 \text{ and } v(i) = 0 \text{ for all } i \in N.$$

Answer. For x to be in the core, we need $x_1 + x_2 + x_3 = 1$, $x_i \geq 0$ for all $i \in N$, $x_1 + x_2 \geq 1$, $x_1 + x_3 \geq 1$ and $x_2 + x_3 \geq 1$. There exists no x satisfying these condition, so the core is empty. \square

22.19 Example: Modified three-person majority game.

$N = \{1, 2, 3\}$, $v(N) = 1$, $v(S) = \alpha$ whenever $|S| = 2$ and $v(i) = 0$ for all $i \in N$.

Answer. The core of this game is the set of all non-negative payoff profiles x for which $x_1 + x_2 + x_3 = 1$ and $\sum_{i \in S} x_i \geq \alpha$ for every two-person coalition S . Hence the core is non-empty if and only if $\alpha \leq \frac{2}{3}$. \square

22.20 Example: A majority game.

A group of n players, where $n \geq 3$ is odd, has one unit to divide among its members. A coalition consisting of a majority of the players can divide the unit among its members as it wishes. This situation is modeled by the coalitional game $\langle N, v \rangle$ in which $|N| = n$ and

$$v(S) = \begin{cases} 1, & \text{if } |S| \geq \frac{n}{2}, \\ 0, & \text{otherwise.} \end{cases}$$

Answer. The game has an empty core by the following argument. Assume that x is in the core. If $|S| = n - 1$ then $v(S) = 1$ so that $\sum_{i \in S} x_i \geq 1$. Since there are n coalitions of size $n - 1$ we thus have

$$\sum_{\{S: |S|=n-1\}} \sum_{i \in S} x_i \geq n.$$

On the other hand, we have

$$\sum_{\{S: |S|=n-1\}} \sum_{i \in S} x_i = \sum_{i \in N} \sum_{\{S: |S|=n-1, S \ni i\}} x_i = \sum_{i \in N} (n-1)x_i = n-1,$$

a contradiction. \square

22.21 Example: The drug game.

Joe Willie has invented a new drug. Joe can not manufacture the drug itself. He can sell the drug formula to company 2 or company 3, but can not sell it to both companies. Company 2 can make a profit of 2 millions if it buys the formula. Company 3 can make a profit of 3 millions if it buys the formula.

Let Joe, companies 2 and 3 be players 1, 2 and 3. Characteristic function v can be defined as

$$v(1) = v(2) = v(3) = 0, v(1, 2) = 2, v(1, 3) = 3, v(2, 3) = 0, v(1, 2, 3) = 3.$$

Answer. $x = (x_1, x_2, x_3)$ is in the core if and only if x satisfies

$$x_1 \geq 0 \text{ (1)}, x_2 \geq 0 \text{ (2)}, x_3 \geq 0 \text{ (3)}, x_1 + x_3 \geq 2 \text{ (4)}, x_1 + x_2 \geq 3 \text{ (5)}, x_2 + x_3 \geq 0 \text{ (6)}, x_1 + x_2 + x_3 = 3 \text{ (7)}.$$

(2), (5) and (7) imply

$$x_2 = 0 \text{ (8)}, x_1 + x_3 = 3 \text{ (9)}.$$

(4), (8) and (9) imply

$$2 \leq x_1 \leq 3, x_3 = 3 - x_1.$$

Conversely, any $x \in X = \{(x_1, x_2, x_3) \mid x_2 = 0, x_3 = 3 - x_1, 2 \leq x_1 \leq 3\}$ satisfies (1)–(7). Hence X is the core. \square

22.22 Example: An expedition of n people has discovered treasure in the mountains; each pair of them can carry out one piece. A coalitional game that models this situation is $\langle N, v \rangle$, where

$$v(S) = \begin{cases} \frac{|S|}{2}, & \text{if } |S| \text{ is even,} \\ \frac{|S|-1}{2}, & \text{if } |S| \text{ is odd.} \end{cases}$$

Answer. If $|N| \geq 4$ is even then the core consists of the single payoff profile $(\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$.

If $|N| \geq 3$ is odd then the core is empty. \square

22.23 Proposition: Let $\delta_i = v(N) - v(N \setminus \{i\})$, $i = 1, 2, \dots, n$, for a coalitional game $\langle N, v \rangle$. Then the core is empty if $\sum_{i=1}^n \delta_i < v(N)$.

Proof.

$$v(N) > \sum_{i=1}^n \delta_i \geq \sum_{i=1}^n v(i),$$

which is impossible. \square

22.24 Denote by \mathcal{C} the set of all coalitions, and for any coalition S denote by $\mathbb{R}^{|S|}$ the $|S|$ -dimensional Euclidian space in which the dimensions are indexed by the members of S .

22.25 Denote by $\mathbf{1}_S \in \mathbb{R}^{|N|}$ the characteristic vector of S given by

$$(\mathbf{1}_S)_i = \begin{cases} 1, & \text{if } i \in S, \\ 0, & \text{otherwise.} \end{cases}$$

22.26 A collection $(\lambda_S)_{S \in \mathcal{C}}$ of numbers in $[0, 1]$ is a balanced collection of weights if for every player i the sum of λ_S over all the coalitions that contain i is 1:

$$\sum_{S \in \mathcal{C}} \lambda_S \mathbf{1}_S = \mathbf{1}_N.$$

Example 1: the collection (λ_S) in which $\lambda_N = 1$ and $\lambda_S = 0$ for all other S is a balanced collection of weights.

Example 2: let $|N| = 3$. Then the collection (λ_S) in which $\lambda_S = \frac{1}{2}$ if $|S| = 2$ and $\lambda_S = 0$ otherwise is a balanced collection of weights; so too is the collection (λ_S) in which $\lambda_S = 1$ if $|S| = 1$ and $\lambda_S = 0$ otherwise.

\Rightarrow 22.27 A game $\langle N, v \rangle$ is balanced if

$$\sum_{S \in \mathcal{C}} \lambda_S v(S) \leq v(N) \text{ for every balanced collection of weights.}$$

22.28 One interpretation of the notion of a balanced game is the following. Each player has one unit of time, which he must distribute among all the coalitions of which he is a member. In order for a coalition S to be active for the fraction of time λ_S , all its members must be active in S for this fraction of time, in which case the coalition yields the payoff $\lambda_S v(S)$. In this interpretation the condition that the collection of weights be balanced is a feasibility condition on the players' allocation of time, and a game is balanced if there is no feasible allocation of time that yields the players more than $v(N)$.

Each player is endowed with one unit of time that he allocates among the coalitions S ; λ_S is the fraction of his time that each member of S allocates to the coalition S ; the condition $\sum_{S \in \mathcal{C}} \lambda_S \mathbf{1}_S = \mathbf{1}_N$ is a feasibility condition (for every individual the sum of its amounts of his time he spends with each coalition must equal exactly the amount of time he is endowed with).

22.29 Bondareva-Shapley theorem: A coalitional game has a non-empty core if and only if it is balanced.

Proof. Let $\langle N, v \rangle$ be a coalitional game.

“ \Rightarrow ”: Let x be a payoff profile in the core of $\langle N, v \rangle$ and $(\lambda_S)_{S \in \mathcal{C}}$ a balanced collection of weights. Then

$$\sum_{S \in \mathcal{C}} \lambda_S v(S) \leq \sum_{S \in \mathcal{C}} \lambda_S \sum_{i \in S} x_i = \sum_{i \in N} x_i \sum_{S \ni i} \lambda_S = \sum_{i \in N} x_i = v(N),$$

so that $\langle N, v \rangle$ is balanced.

“ \Leftarrow ”: Assume that $\langle N, v \rangle$ is balanced. Then there is no balanced collection $(\lambda_S)_{S \in \mathcal{C}}$ of weights for which

$$\sum_{S \in \mathcal{C}} \lambda_S v(S) > v(N).$$

Therefore the convex set

$$\{(\mathbf{1}_N, v(N) + \epsilon) \in \mathbb{R}^{|N|+1} : \epsilon > 0\}$$

is disjoint from the convex cone

$$\left\{ y \in \mathbb{R}^{|N|+1} \mid y = \sum_{S \in \mathcal{C}} \lambda_S (\mathbf{1}_S, v(S)) \text{ where } \lambda_S \geq 0 \text{ for all } S \in \mathcal{C} \right\},$$

since if not then $\mathbf{1}_N = \sum_{S \in \mathcal{C}} \lambda_S \mathbf{1}_S$, so that $(\lambda_S)_{S \in \mathcal{C}}$ is a balanced collection of weights and $\sum_{S \in \mathcal{C}} \lambda_S v(S) > v(N)$. Thus by hyperplane separating theorem there is a non-zero vector $(\alpha_N, \alpha) \in \mathbb{R}^{|N|} \times \mathbb{R}$ such that

$$(\alpha_N, \alpha) \cdot y \geq 0 > (\alpha_N, \alpha) \cdot (\mathbf{1}_N, v(N) + \epsilon)$$

for all y in the cone and all $\epsilon > 0$. Since $(\mathbf{1}_N, v(N))$ is in the cone, we have $\alpha < 0$.

Now let $x = \alpha_N / (-\alpha)$. Since $(\mathbf{1}_S, v(S))$ is in the cone for all $S \in \mathcal{C}$, we have $x(S) = x \cdot \mathbf{1}_S \geq v(S)$ for all $S \in \mathcal{C}$, and $v(N) \geq \mathbf{1}_N x = \sum_{i \in N} x_i$. Thus $v(N) = \sum_{i \in N} x_i$, so that x is in the core of $\langle N, v \rangle$. \square

22.30 Example: n -person weighted majority game with weights $(w_i)_{i \in N}$ and quota q is defined by

$$v(S) = \begin{cases} 1, & \text{if } \sum_{i \in S} w_i \geq q, \\ 0, & \text{if } \sum_{i \in S} w_i < q. \end{cases}$$

A 0-1-normalized weighted majority game has a non-empty core if and only if there is at least one veto player.

Answer. “ \Rightarrow ”: Number the players in such a way that $w_1 \geq w_2 \geq \dots \geq w_n$. Then if there is at least one veto player, player 1 is such a player, i.e., $v(S) = 0$ if $1 \notin S$. Hence $x = (1, 0, \dots, 0)$ is in the core: $\sum_{i \in N} x_i \geq 1$, $x_i \geq 0$ for all $i \in N$, and

$$\sum_{i \in S} x_i = \begin{cases} 1, & \text{if } 1 \in S, \text{ in which case } v(S) \leq 1, \\ 0, & \text{if } 1 \notin S, \text{ in which case } v(S) = 0. \end{cases}$$

This establishes sufficiency.

“ \Leftarrow ”: Suppose there are no veto players. Consider the collection of coalitions $\mathcal{S} = \{N \setminus \{1\}, N \setminus \{2\}, \dots, N \setminus \{n\}\}$ with balancing weights $\lambda_S = \frac{1}{n-1}$ for all $S \in \mathcal{S}$. We have

$$\sum_{S \in \mathcal{S}} \lambda_S \mathbf{1}_S = \frac{1}{n-1} \sum_{S \in \mathcal{S}} \mathbf{1}_S = \frac{1}{n-1} \begin{pmatrix} n-1 \\ n-1 \\ \vdots \\ n-1 \end{pmatrix} = \mathbf{1}_N,$$

so \mathcal{S} is a balanced collection. Since there are no veto players $v(S) = 1$ for all $S \in \mathcal{S}$, so $\sum_{S \in \mathcal{S}} \lambda_S v(S) = \frac{n-1}{n} > 1 = v(N)$. Hence by the Bondareva-Shapley theorem the core is empty. This establishes necessity. \square

22.31 The core of a 0-1-normalized weighted majority game with veto players $1, 2, \dots, p$ is

$$\left\{ x = (a_1, a_2, \dots, a_p, 0, \dots, 0) \mid a_i \geq 0 \text{ for all } i \in N \text{ and } \sum_{i=1}^p a_i = 1 \right\}.$$

Let S be such that $v(S) = 1$, then $\{1, 2, \dots, p\} \subseteq S$, so $\sum_{i \in S} x_i = 1 = v(S)$. Let S be such that $v(S) = 0$, then $\sum_{i \in S} x_i \geq 0 = v(S)$. Hence any $x = (a_1, a_2, \dots, a_p, 0, \dots, 0)$ with $a_i \geq 0$ for all $i \in N$ and $\sum_{i=1}^p a_i = 1$ is in the core of the game, and the core consists of solely such points.

22.32 Proposition: Suppose player 1 is a null player, and (x_1, x_2, \dots, x_n) is in the core. Then $x_1 = 0$.

Proof. By definition of core, we have

$$\sum_{i=2}^n x_i \geq v(\{2, 3, \dots, n\}) = v(\{1, 2, \dots, n\}) = \sum_{i=1}^n x_i,$$


that is, $x_1 \leq 0$.

On the other hand, we also have

$$x_1 \geq v(\{1\}) = v(\emptyset) = 0,¹$$

and hence $x_1 = 0$. \square

22.3 Shapley value

 22.33 Given a coalitional game $\langle N, v \rangle$ where $N = \{1, 2, \dots, n\}$, the Shapley value is an n -vector, denoted by $\phi(v) = (\phi_1(v), \phi_2(v), \dots, \phi_n(v))$, satisfying the following conditions:

S1. Symmetry condition: if i and j are substitutes in v , then $\phi_i(v) = \phi_j(v)$.

S2. Null player condition: if i is a null player, then $\phi_i(v) = 0$.

S3. Efficiency condition: $\sum_{i \in N} \phi_i(v) = v(N)$.

S4. Additivity condition: $\phi_i(v + w) = \phi_i(v) + \phi_i(w)$.

¹Technically, we should assume that coalition could be empty.

22.34 $\phi_i(v)$ is interpreted as the power of player i in the coalitional game $\langle N, v \rangle$, or what it is worth to i to participate in the game $\langle N, v \rangle$.

22.35 Conditions S1, S2 and S4 are weak restrictions which are easy to accept as “reasonable”, while S3 is much stronger.

22.36 Shapley theorem: Shapley value is uniquely determined:

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} [v(S \cup \{i\}) - v(S)] = \frac{1}{n} \sum_{s=0}^{n-1} \frac{1}{\binom{n-1}{s}} \sum_{S \subseteq N \setminus \{i\}, |S|=s} [v(S \cup \{i\}) - v(S)].$$

Proof. It suffices to prove the uniqueness. Let ϕ be a Shapley value of $\langle N, v \rangle$. For each coalition S , define a characteristic function v_S by

$$v_S(T) = \begin{cases} 1, & \text{if } S \subseteq T, \\ 0, & \text{otherwise.} \end{cases}$$

Note that for any real number α , members of $N \setminus S$ are null players in the game $\langle N, \alpha v_S \rangle$, and members of S are substitutes for each other in the game $\langle N, \alpha v_S \rangle$.

Hence by the null player condition, $\phi_i(\alpha v_S) = 0$ when $i \notin S$, and by the symmetry condition $\phi_i(\alpha v_S) = \phi_j(\alpha v_S)$ when $i, j \in S$. Hence, by the efficiency condition

$$\sum_{i \in N} \phi_i(\alpha v_S) = \alpha v_S(N) = \alpha.$$

Thus $\alpha = \sum_{i \in S} \phi_i(\alpha v_S) = |S| \phi_i(\alpha v_S)$ for any $i \in S$. Hence,

$$\phi_i(\alpha v_S) = \begin{cases} \frac{\alpha}{|S|}, & \text{if } i \in S, \\ 0, & \text{if } i \notin S. \end{cases}$$

Now, each characteristic function can be regarded as a $(2^{|N|} - 1)$ -vector, and there are $2^{|N|} - 1$ coalitions. We know $\phi(\alpha v_S)$ for all α and S , so by additivity we know $\phi(\sum_{i=1}^k \alpha_i v_{S_i})$ for all linear combinations $\sum_{i=1}^k \alpha_i v_{S_i}$ of the v_S 's. Hence if we prove that the v_S 's are linearly independent, we will have shown that ϕ is uniquely determined by v_S 's.

Suppose they are not linearly independent; then we may write

$$v_S = \sum_{i=1}^j \beta_i v_{S_i},$$

where $|S| \leq |S_i|$ for all i and S_i 's are different from each other and from S . Then

$$1 = v_S(S) = \sum_{i=1}^j \beta_i v_{S_i}(S) = \sum_{i=1}^j \beta_i \cdot 0 = 0,$$

a contradiction. □

22.37 Notation: Let $\gamma(s) = \frac{s!(n-s-1)!}{n!}$. Then we have

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \gamma(|S|) [v(S \cup \{i\}) - v(S)].$$

22.38 Example: Two-person bargaining game.

$$N = \{1, 2\}, v(1, 2) = 1, v(1) = v(2) = 0.$$

Since $n = 2$, we have $\gamma(0) = \gamma(1) = \frac{1}{2}$.

For player 1, we have

S	\emptyset	$\{2\}$
$v(S \cup \{1\}) - v(S)$	0	1

Table 22.1

$$\text{Hence, } \phi_1(v) = 0\frac{1}{2} + 1\frac{1}{2} = \frac{1}{2}.$$

For player 2, we have

S	\emptyset	$\{1\}$
$v(S \cup \{1\}) - v(S)$	0	1

Table 22.2

$$\text{Hence, } \phi_2(v) = 0\frac{1}{2} + 1\frac{1}{2} = \frac{1}{2}.$$

For player 2, we can get $\phi_2(v) = \frac{1}{2}$ by efficiency condition directly.

22.39 Example: Three-person majority game.

$$N = \{1, 2, 3\}, v(1) = v(2) = v(3) = 0, v(1, 2) = v(1, 3) = v(2, 3) = v(N) = 1$$

Since $n = 3$, we have $\gamma(0) = \gamma(2) = \frac{1}{3}$, and $\gamma(1) = \frac{1}{6}$.

For player 1, we have

S	\emptyset	$\{2\}$	$\{3\}$	$\{2,3\}$
$v(S \cup \{1\}) - v(S)$	0	1	1	0

Table 22.3

$$\text{Hence, } \phi_1(v) = 0\frac{1}{3} + 1\frac{1}{6} + 1\frac{1}{6} + 0\frac{1}{3} = \frac{1}{3}.$$

For player 2, we have

S	\emptyset	$\{1\}$	$\{3\}$	$\{1,3\}$
$v(S \cup \{2\}) - v(S)$	0	1	1	0

Table 22.4

Hence, $\phi_2(v) = 0\frac{1}{3} + 1\frac{1}{6} + 1\frac{1}{6} + 0\frac{1}{3} = \frac{1}{3}$.

For player 3, we can get $\phi_3(v) = \frac{1}{3}$ by efficiency condition directly.

22.40 Example: Market with two sellers and one buyer.

$N = \{1, 2, 3\}$, $v(1, 2, 3) = v(1, 2) = v(1, 3) = 1$, and $v(S) = 0$ for all other $S \subseteq N$.

Since $n = 3$, we have $\gamma(0) = \gamma(2) = \frac{1}{3}$, and $\gamma(1) = \frac{1}{6}$.

For player 1, we have

S	\emptyset	$\{2\}$	$\{3\}$	$\{2,3\}$
$v(S \cup \{1\}) - v(S)$	0	1	1	1

Table 22.5

Hence, $\phi_1(v) = 0\frac{1}{3} + 1\frac{1}{6} + 1\frac{1}{6} + 1\frac{1}{3} = \frac{2}{3}$.

For player 2, we have

S	\emptyset	$\{1\}$	$\{3\}$	$\{1,3\}$
$v(S \cup \{2\}) - v(S)$	0	1	0	0

Table 22.6

Hence, $\phi_2(v) = 0\frac{1}{3} + 1\frac{1}{6} + 0\frac{1}{6} + 0\frac{1}{3} = \frac{1}{6}$.

For player 3, we can get $\phi_3(v) = \frac{1}{6}$ by efficiency condition directly.

22.41 Note that the core allocation in the example above $(1, 0, 0)$ differs considerably from the Shapley value $(\frac{2}{3}, \frac{1}{6}, \frac{1}{6})$. One can interpret that zero payoff to players 2 and 3 in the core allocation as the result of cutthroat competition between them.

22.42 Example: Consider an n -person game in which the only winning coalitions are those coalitions containing player 1 and at least one other player. If a winning coalition receives a reward of \$1, find the core and the Shapley value of the game.

Answer. When $n = 2$, the solution is quite easy. (Exercise)

In the following, we assume $n \geq 3$. The characteristic function is

$$v(S) = \begin{cases} 1, & \text{if } |S| \geq 2, \text{ player 1 belongs to } S, \\ 0, & \text{otherwise.} \end{cases}$$

(i) Core: Suppose (x_1, x_2, \dots, x_n) is in the core. Then we have

$$\sum_{i=1}^n x_i = 1, \quad x_i \geq 0,$$

$$x_1 + \sum_{i \in S} x_i \geq 1, \quad \text{for all } S \subset \{2, 3, \dots, n\} \text{ and } S \text{ is non-empty.}$$

It is easy to see that the only solution is $x_1 = 1, x_2 = \cdots = x_n = 0$: take $S = \{3, 4, \dots, n\}$, we have $x_1 + x_3 + \cdots + x_n \geq 1$, and hence $x_2 = 0$. Similarly $x_3 = x_4 = \cdots = x_n = 0$. Therefore the core is $\{(1, 0, 0, \dots, 0)\}$.

- (ii) Shapley value: For player $i \neq 1$, $v(S \cup \{i\}) - v(S) = 1$ if and only if $S = \{1\}$. Otherwise it is zero. Since $\gamma(1) = \frac{1}{n(n-1)}$, player i 's Shapley value is

$$\phi_i(v) = \frac{1}{n(n-1)}.$$

For player 1, we have

$$\phi_1(v) = 1 - \sum_{i=2}^n \phi_i(v) = \frac{n-1}{n}.$$

Therefore, the Shapley value is

$$\left(\frac{n-1}{n}, \frac{1}{n(n-1)}, \dots, \frac{1}{n(n-1)} \right).$$

□

22.43 Consider the following cost allocation problem. Building an airfield will benefit n players. Player j requires an airfield that costs c_j to build, so to accommodate all the players, the field will be built at a cost of $\max_{1 \leq j \leq n} c_j$. How should this cost be split among the players? Suppose all the costs are distinct and let $0 < c_1 < c_2 < \cdots < c_n$. Take the characteristic function of the game to be $v(S) = -\max_{j \in S} c_j$ for $S \subset \{1, 2, \dots, n\}$.

- (i) Let $R_k = \{k, k+1, \dots, n\}$ for $k = 1, 2, \dots, n$, and define the characteristic function v_k through the equation

$$v_k(S) = \begin{cases} -(c_k - c_{k-1}), & \text{if } S \cap R_k \neq \emptyset \\ 0, & \text{if } S \cap R_k = \emptyset \end{cases}$$

For convenience, let $c_0 = 0$. Show that $v = \sum_{k=1}^n v_k$.

- (ii) Find the Shapley value of the game v in the form of $\phi_i(v) = \sum_{k=1}^i \alpha_{ik}(c_k - c_{k-1})$, $i = 1, 2, \dots, n$, where the coefficients α_{ik} are independent of c_1, c_2, \dots, c_n .

Answer. (i) For every coalition S , we have

$$\sum_{k=1}^n v_k(S) = \sum_{k=1}^{\max(S)} v_k(S) + \sum_{k=\max(S)+1}^n v_k(S) = \sum_{k=1}^{\max(S)} v_k(S) = - \sum_{k=1}^{\max(S)} (c_k - c_{k-1}) = -c_{\max(S)} = v(S).$$

- (ii) Since

$$v_k(S) = \begin{cases} -(c_k - c_{k-1}), & \text{if } \max(S) \geq k \\ 0, & \text{if } \max(S) < k \end{cases}$$

we have

$$v_k(S \cup \{i\}) - v_k(S) = \begin{cases} -(c_k - c_{k-1}), & \text{if } \max(S) < k \leq i \\ 0, & \text{otherwise} \end{cases}$$

and hence $\phi_1(v_k) = \cdots = \phi_{k-1}(v_k) = 0$, and

$$\phi_k(v_k) = \cdots = \phi_n(v_k)$$

$$\begin{aligned}
&= \frac{1}{n} \sum_{s=0}^{n-1} \frac{1}{\binom{n-1}{s}} \sum_{|S|=s} [v_k(S \cup \{i\}) - v_k(S)] \\
&= \frac{1}{n} \sum_{s=0}^{n-1} \frac{1}{\binom{n-1}{s}} \left\{ \sum_{|S|=s, \max(S) < k} [v_k(S \cup \{i\}) - v_k(S)] + \sum_{|S|=s, \max(S) \geq k} [v_k(S \cup \{i\}) - v_k(S)] \right\} \\
&= \frac{1}{n} \sum_{s=0}^{n-1} \frac{1}{\binom{n-1}{s}} \binom{k-1}{s} [-(c_k - c_{k-1})]
\end{aligned}$$

Therefore, we have

$$\phi_i(v) = \sum_{k=1}^n \phi_i(v_k) = \sum_{k=1}^i \phi_i(v_k) = -\frac{1}{n} \sum_{k=1}^i \sum_{s=0}^{k-1} \frac{\binom{k-1}{s}}{\binom{n-1}{s}} (c_k - c_{k-1}).$$

□

22.4 Nash bargaining solution

22.44 A two-person bargaining problem, denoted by $\langle U, d \rangle$, consists of

- U is the set of possible agreements in terms of utilities that they yield to 1 and 2. An element of U is a pair $u = (u_1, u_2)$.
- d is a pair (d_1, d_2) , called the disagreement point or threat point.

If agreement $u = (u_1, u_2) \in U$ is reached, then 1 gets utility u_1 and 2 gets utility u_2 . If no agreement is reached then 1 gets utility d_1 and 2 gets utility d_2 .

The set of two-person bargaining games is denoted by W .

22.45 Convention: Assume that

- U is compact and convex.
- U contains a point y for which $y_i > d_i$ for $i = 1, 2$, that is, bargaining is worthwhile for both the players.

22.46 The Nash bargaining solution is a mapping $f: W \rightarrow \mathbb{R}^2$ that associates a unique element $f(U, d)$ with the game $\langle U, d \rangle$, satisfying the following axioms:

N1. Feasibility: $f(U, d) \in U$.

N2. Individual rationality: $f(U, d) \geq d$ for all $\langle U, d \rangle \in W$.

N3. Pareto optimality: $f(U, d)$ is Pareto optimal. That is, there does not exist a point $(u_1, u_2) \in U$ such that

$$u_1 \geq f_1(U, d), u_2 \geq f_2(U, d), (u_1, u_2) \neq f(U, d).$$

N4. Symmetry: If $\langle U, d \rangle \in W$ satisfies $d_1 = d_2$ and $(x_1, x_2) \in U$ implies $(x_2, x_1) \in U$, then $f_1(U, d) = f_2(U, d)$.

N5. Invariance under linear transformations: Let $a_1, a_2 > 0$, $b_1, b_2 \in \mathbb{R}$, and $\langle U, d \rangle, \langle U', d' \rangle \in W$ where $d'_i = a_i d_i + b_i$, $i = 1, 2$, and $U' = \{x \in \mathbb{R}^2 \mid x_i = a_i y_i + b_i, i = 1, 2, y \in U\}$. Then $f_i(U'_i, d'_i) = a_i f_i(U, d) + b_i$, $i = 1, 2$.

N6. Independence of irrelevant alternatives: If $\langle U, d \rangle, \langle U', d' \rangle \in W$, $d = d'$, $U \subseteq U'$, and $f(U', d') \in U$, then $f(U, d) = f(U', d')$.

The interpretation is that, given any bargaining problem $\langle U, d \rangle$, the solution function tells us that the agreement $u = f(U, d)$ will be reached.

22.47 Pareto optimality: there are no points in U that are “North-East” of $f(U, d)$. See Figure 22.1.




Figure 22.1: Pareto optimality.

22.48 Symmetry: suppose that $\langle U, d \rangle$ is such that U is symmetric around the 45° line and $d_1 = d_2$, then $f_1(U, d) = f_2(U, d)$, that is, when everything in $\langle U, d \rangle$ is symmetric, the point $f(U, d)$ is itself on the 45° line. See Figure 22.2.




Figure 22.2: Symmetry.

22.49 Invariance under linear transformations: suppose we have two bargaining problems $\langle U, d \rangle$ and $\langle U', d' \rangle$ with the following property. For some vector $b = (b_1, b_2)$,

$$d' = d + b, \quad U' = U + b.$$

Then invariance under linear transformations imposes that

$$f(U', d') = f(U, d) + b,$$

see Figure 22.3.




Figure 22.3: Independence of utility origins.

Suppose we have two bargaining problems $\langle U, d \rangle$ and $\langle U', d' \rangle$ with $d = (0, 0)$ and the following property.

$$U'_1 = k_1 U_1, \quad U'_2 = k_2 U_2.$$

Then invariance under linear transformations imposes that

$$f_1(U', d) = k_1 f_1(U, d), \quad f_2(U', d) = k_2 f_2(U, d).$$

In Figure 22.4, we depict a change for u_2 only with $k_2 = 2$.




Figure 22.4: Independence of utility units.

22.50 Independence of irrelevant alternatives:




Figure 22.5: Independence of irrelevant alternatives.

22.51 Theorem: A game $\langle U, d \rangle \in W$ has a unique Nash solution $u^* = f(U, d)$ satisfying Conditions N1 to N6. Furthermore, the solution u^* satisfies Conditions N1 to N6 if and only if

$$(u_1^* - d_1)(u_2^* - d_2) > (u_1 - d_1)(u_2 - d_2)$$

for all $u \in U$, $u \geq d$, and $u \neq u^*$.

22.52 Remark:

- Existence of an optimal solution: Since the set U is compact and the objective function is continuous, there exists an optimal solution.
- Uniqueness of the optimal solution: The objective function is strictly quasi-concave. Therefore, maximization problem has a unique optimal solution.

22.53 Part 1. We first prove that Nash bargaining solution satisfies the six axioms. The first two are automatically satisfied, and we focus on other four.

N3. Pareto optimality: This follows immediately from the fact that the objective function $(u_1 - d_1)(u_2 - d_2)$ is increasing in u_1 and u_2 .

N4. Symmetry: Assume $d_1 = d_2$. Let $u^* = (u_1^*, u_2^*)$ be the Nash bargaining solution. Then, it can be seen that (u_2^*, u_1^*) is also an optimal solution. By the uniqueness of the optimal solution, we must have $(u_1^*, u_2^*) = (u_2^*, u_1^*)$, that is, $u_1^* = u_2^*$.

N5. Invariance under linear transformation: By definition, $f(U', d')$ is an optimal solution of the problem

$$\begin{aligned} &\text{maximize} && (u_1 - a_1 d_1 - b_1)(u_2 - a_2 d_2 - b_2) \\ &\text{subject to} && (u_1, u_2) \in U' \end{aligned}$$

Performing the change of variables $u'_1 = a_1 u_1 + b_1$, $u'_2 = a_2 u_2 + b_2$, it follows immediately that $f_i(U', d') = a_i f_i(U, d) + b_i$ for $i = 1, 2$.

N6. Independence of irrelevant alternatives: Let $U \subseteq U'$. From the optimization problem characterization of the Nash bargaining solution, it follows that the objective function value at the solution $f(U', d)$ is greater than

or equal to that at $f(U, d)$. If $f(U', d) \in U$, then the objective function values must be equal, i.e., $f(U', d)$ is optimal for U and by uniqueness of the solution $f(U, d) = f(U', d)$.

□

22.54 *Part 2.* We then show that if a bargaining solution satisfies the six axioms, it must be equal to f .

- (1) Let g be a bargaining solution satisfying the six axioms. We prove that $g(U, d) = f(U, d)$ for every bargaining problem $\langle U, d \rangle$.
- (2) Given a bargaining problem $\langle U, d \rangle$, let $z = f(U, d)$, and define the set

$$U' = \{(au + b) \mid u \in U, az + b = (\tfrac{1}{2}, \tfrac{1}{2}), ad + b = (0, 0)\},$$

that is, we map the point z to $(\frac{1}{2}, \frac{1}{2})$ and the point d to $(0, 0)$.

- (3) Since $g(U, d)$ and $f(U, d)$ both satisfy axiom N5 (invariance under linear transformation), we have $g(U, d) = f(U, d)$ if and only if $g(U', \mathbf{0}) = f(U', \mathbf{0}) = (\frac{1}{2}, \frac{1}{2})$. Hence, to establish the desired claim, it is sufficient to prove $g(U', \mathbf{0}) = (\frac{1}{2}, \frac{1}{2})$.
- (4) Let us show that there is no $u \in U'$ such that $u_1 + u_2 > 1$: Assume that there is a $u \in U'$ such that $u_1 + u_2 > 1$. Let $t = (1 - \lambda)(\frac{1}{2}, \frac{1}{2}) + \lambda(u_1, u_2)$ for some $\lambda \in (0, 1)$. Since U' is convex, we have $t \in U'$. We can choose λ sufficiently small so that $t_1 t_2 > \frac{1}{4} = f(U', \mathbf{0})$, but this contradicts the optimality of $f(U', \mathbf{0})$, showing that for all $u \in U'$, we have $u_1 + u_2 \leq 1$.
- (5) Let $U'' = \{(u_1, u_2) \mid u_1 + u_2 \leq 1, u_1 \geq 0, u_2 \geq 0\}$. Then $U' \subseteq U''$ and $(\frac{1}{2}, \frac{1}{2})$ is on the boundary of U'' .
- (6) By axiom N3 (Pareto optimality) and N4 (symmetry), $g(U'', \mathbf{0}) = (\frac{1}{2}, \frac{1}{2})$.
- (7) By axiom N6 (Invariance under linear transformation), since $U' \subseteq U''$, we have $g(U', \mathbf{0}) = (\frac{1}{2}, \frac{1}{2})$, completing the proof.

□

22.55 *Example:* Find the Nash bargaining solution of the following problem.

$$U = \{(u_1, u_2) \mid u_1 + u_2 \leq 10, u_1 \geq 0, u_2 \geq 0\},$$

$$d = (2, 4).$$

We need to maximize $(u_1 - d_1) \cdot (u_2 - d_2) = (u_1 - 2) \cdot (u_2 - 4)$.

The Nash bargaining solution lies on the frontier, so, we can assume that $u_1 + u_2 = 10$. Hence, it suffices maximize

$$(u_1 - 2)(10 - u_1 - 4) = (u_1 - 2)(6 - u_1).$$

By first order condition, $u_1^* = 4$ and $u_2^* = 6$.

22.56 *Example:* Two bargaining problems have identical threat points $d = (0, 0)$. In one case, the $U_1 = \{(u_1, u_2) \mid u_1 + u_2 \leq 6, u_1 \geq 0, u_2 \geq 0\}$. In the other case, the $U_2 = \{(u_1, u_2) \mid u_1 + u_2 \leq 6, u_1 \geq 0, 0 \leq u_2 \leq 4\}$. Are the Nash solution to these two bargaining problems the same?

22.57 *Example:* Dividing one dollar.

Two individuals can divide a dollar in any way they wish. If they fail to agree on a division, the dollar is forfeited. The individuals may, if they wish, discard some of the dollar. Thus, the set of outcomes is all possible divisions of the dollar.

$$X = \{(x_1, x_2) \mid x_1 + x_2 \leq 1, x_1 \geq 0, x_2 \geq 0\}.$$

Neither player receives anything in the event of disagreement. So, the disagreement outcome is $(0, 0)$. Each player is concerned only about his/her share of the dollar and prefers more to less.

The utility functions of the two players are u_1 and u_2 .

$$U = \{(u_1(x_1), u_2(x_2)) \mid x_1 + x_2 \leq 1, x_1 \geq 0, x_2 \geq 0\},$$

$$d = (u_1(0), u_2(0)).$$

Players are either risk neutral or risk averse, then U is convex.

- First suppose that the players' preferences are the same, so that they can be represented by the same utility function, u . Then we have a symmetric bargaining problem. In this case, we know that the Nash solution is the unique symmetric efficient utility pair $(u(1/2), u(1/2))$, which corresponds to the physical outcome in which the dollar is shared equally between the players.
- If the players have different preferences, then equal division of the dollar may no longer be the agreement given by the Nash solution. Rather, the solution depends on the nature of the players' preferences.

To investigate this dependence, consider a simple example. Suppose that player 1 is risk neutral so that her payoff function is $u_1(x_1) = x_1$ and player 2 is risk averse so that his payoff function is $u_2(x_2) = \sqrt{x_2}$.

Maximize $[u_1(x_1) - u_1(0)] \cdot [u_2(x_2) - u_2(0)]$ subject to $x_1 + x_2 \leq 1$.

$u_1(x_1) = x_1$ and $u_2(x_2) = \sqrt{x_2}$ give: maximize $x_1\sqrt{x_2}$ subject to $x_1 + x_2 \leq 1$.

When this is maximized, it must be the case that $x_1 + x_2 = 1$. So, we need to maximize

$$(1 - x_2)\sqrt{x_2}.$$

The solution is $x_1^* = \frac{2}{3}$ and $x_2^* = \frac{1}{3}$.

So, player 2's share goes down. More generally, we can have the following intuition. If player 2 becomes more risk averse, then player 1's share of the dollar in the Nash solution increases. If player 2 is more risk averse than player 1, then player 1's share of the dollar in the Nash solution exceeds $1/2$.

22.58 Example: Dividing one dollar. (cont.)

General analysis. • $u_1 = u_2 = u$, and u is concave and $u(0) = 0$. It is a symmetric bargaining problem, and hence $f(U, d) = (\frac{1}{2}, \frac{1}{2})$, that is, the dollar is shared equally.

The bargaining problem is the optimal solution of the following problem

$$\max_{0 \leq z \leq 1} u_1(z)u_2(1-z) = \max_{0 \leq z \leq 1} u(z)u(1-z).$$

We denote the optimal solution of this problem by z_u . By first order condition, we have

$$u'(z)u(1-z) = u(z)u'(1-z),$$

implying that $\frac{u'(z_u)}{u(z_u)} = \frac{u'(1-z_u)}{u(1-z_u)}$.

- Player 2 is more risk averse, i.e., $u_1 = u$ and $u_2 = h \circ u$, where $h: \mathbb{R} \rightarrow \mathbb{R}$ is an increasing concave function with $h(0) = 0$. The Nash bargaining solution is the optimal solution of the following problem

$$\max_{0 \leq z \leq 1} u_1(z)u_2(1-z) = \max_{0 \leq z \leq 1} u(z)h(u(1-z)).$$

We denote the optimal solution of this problem by z_v . By the first order condition, we have

$$u'(z)h(u(1-z)) = u(z)h'(u(1-z))u'(1-z),$$

$$\text{implying } \frac{u'(z_v)}{u(z_v)} = \frac{h'(u(1-z_v))u'(1-z_v)}{h(u(1-z_v))}.$$

Since h is an increasing concave function and $h(0) = 0$, we have

$$h'(t) \leq \frac{h(t)}{t} \text{ for all } t \geq 0.$$

This implies that

$$\frac{u'(z_v)}{u(z_v)} \leq \frac{u'(1-z_v)}{u(1-z_v)},$$

and therefore $z_u \leq z_v$. This shows that when player 2 is more risk averse, player 1's share increases. □

22.59 Example: Imagine a firm that is a monopolist in the market for its output and in the labor market (hiring) as well. At the same time, a labor union is a monopolist in the labor market (supplying). Letting L denote the level of employment and w , the wage rate, suppose the union has a utility function $u(L, w) = \sqrt{Lw}$. The firm's utility is its profit $\pi = L(100 - L) - wL$.

In this situation, the payoff set

$$U = \{(u, \pi) \mid u = \sqrt{Lw}, \pi = L(100 - L) - wL, L \geq 0, w \geq 0\}.$$

The most natural threat on the part of the firm is to cease production, which means the union members will all be without jobs. Similarly, the union can refuse to supply any labor to the firm, and, again, there will be neither production nor jobs. That is, $(0, 0)$ is the disagreement point.

The Nash bargaining solution maximizes the function

$$(u - d_1)(\pi - d_2) = \sqrt{Lw}(L(100 - L) - wL)$$

for $L \geq 0$ and $w \geq 0$. Optimality conditions are both partial derivatives with respect to L and w being 0:

$$300 - 5L - 3w = 0, \quad 100 - L - 3w = 0$$

which results in

$$L = 50, \quad w = \frac{50}{3}.$$

Thus, the Nash bargaining solution is

$$u^* = \frac{50}{\sqrt{3}}, \quad \pi^* = \frac{5000}{3}.$$

22.60 Example: Suppose the set U consists of the points lying on and within a circle of radius 2, having a center at $(2, 2)$.

If the threat point, d , is at $(2, 2)$, what is the Nash bargaining solution? If the threat point, d , is at $(0, 2)$, what is the Nash bargaining solution?

Answer. $U = \{(u_1, u_2) : (u_1 - 2)^2 + (u_2 - 2)^2 \leq 4\}$.

(i) $d = (2, 2)$. Consider the following problem:

$$\text{maximize } (u_1 - 2)(u_2 - 2) \quad (22.1)$$

$$\text{subject to } (u_1 - 2)^2 + (u_2 - 2)^2 \leq 4 \quad (22.2)$$

$$u_1 \geq 2, u_2 \geq 2 \quad (22.3)$$

Consider Equations (22.1) and (22.2), and apply the method of Lagrange multipliers, we will have

$$f(u_1, u_2, \lambda) = (u_1 - 2)(u_2 - 2) - \lambda[(u_1 - 2)^2 + (u_2 - 2)^2 - 4]$$

$$\frac{\partial f}{\partial u_1} = 0 \Rightarrow (u_2 - 2) = 2\lambda(u_1 - 2)$$

$$\frac{\partial f}{\partial u_2} = 0 \Rightarrow (u_1 - 2) = 2\lambda(u_2 - 2)$$

$$\frac{\partial f}{\partial \lambda} = 0 \Rightarrow (u_1 - 2)^2 + (u_2 - 2)^2 = 4$$

The solutions are: $(2 + \sqrt{2}, 2 + \sqrt{2})$ and $(2 - \sqrt{2}, 2 - \sqrt{2})$. Note that only $(2 + \sqrt{2}, 2 + \sqrt{2})$ satisfies (22.3). Therefore, $(2 + \sqrt{2}, 2 + \sqrt{2})$ is the unique Nash bargaining solution.

(ii) $d = (0, 2)$. Consider the following problem:

$$\text{maximize } (u_1 - 0)(u_2 - 2) \quad (22.4)$$

$$\text{subject to } (u_1 - 2)^2 + (u_2 - 2)^2 \leq 4 \quad (22.5)$$

$$u_1 \geq 0, u_2 \geq 2 \quad (22.6)$$

Consider Equations (22.4) and (22.5), and apply the method of Lagrange multipliers, we will have

$$f(u_1, u_2, \lambda) = u_1(u_2 - 2) - \lambda[(u_1 - 2)^2 + (u_2 - 2)^2 - 4]$$

$$\frac{\partial f}{\partial u_1} = 0 \Rightarrow (u_2 - 2) = 2\lambda(u_1 - 2)$$

$$\frac{\partial f}{\partial u_2} = 0 \Rightarrow u_1 = 2\lambda(u_2 - 2)$$

$$\frac{\partial f}{\partial \lambda} = 0 \Rightarrow (u_1 - 2)^2 + (u_2 - 2)^2 = 4$$

The solutions are: $(0, 2)$ and $(3, 2 + \sqrt{3})$, where the former is not Pareto optimal. $(3, 2 + \sqrt{3})$ is the unique Nash bargaining solution.

□

22.61 Example: Player 1 and player 2 have been willed equal shares of an estate consisting of \$200,000 cash and 100 acres of farmland. Player 1 has a sentimental attachment to the land and values it at $v_1 = \$3,000$ per acre, whereas player 2 has no such attachment and values it at $v_2 = \$1,000$ per acre. Assume that their payoff functions are linear in money and land at these rates: $u_i = x_i + v_i y_i$ if player i receives x_i dollars of cash and y_i acres of land.

The players may reach an agreement on dividing the land and money so as to maximize their payoffs. If they fail to reach agreement they divide the land and money equally.

- (i) Carefully draw the bargaining set and label the disagreement point.
- (ii) Find the Nash bargaining solution.

Answer. (i) Assume in an agreement, the outcome is (x_1, x_2) and (y_1, y_2) , where

$$x_1 + x_2 = 200000, y_1 + y_2 = 100, x_1, x_2, y_1, y_2 \geq 0,$$

and corresponding payoffs are

$$u_1 = x_1 + 3000y_1, u_2 = x_2 + 1000y_2.$$

Hence, we have

$$u_1 + u_2 = 300000 + 2000y_1, u_1 + 3u_2 = 500000 + 2000x_1,$$

and hence

$$300000 \leq u_1 + u_2 \leq 500000, 500000 \leq u_1 + 3u_2 \leq 900000.$$

Disagreement outcome is $x_1 = x_2 = 100000$, and $y_1 = y_2 = 50$, and hence $u_1 = 250000$ and $u_2 = 150000$, which is a threat point in

$$U = \{(u_1, u_2) : 300000 \leq u_1 + u_2 \leq 500000, 500000 \leq u_1 + 3u_2 \leq 900000\}.$$

- (ii) Consider the following problem:

$$\text{maximize } (u_1 - 250000)(u_2 - 150000) \quad (22.7)$$

$$\text{subject to } u_1 + u_2 \leq 500000 \quad (22.8)$$

$$u_1 + 3u_2 \leq 900000 \quad (22.9)$$

$$300000 \leq u_1 + u_2 \quad (22.10)$$

$$500000 \leq u_1 + 3u_2 \quad (22.11)$$

$$u_1 \geq 0, u_2 \geq 0 \quad (22.12)$$

Consider Equations (22.7), (22.8) and (22.9), and apply the method of Lagrange multipliers, we will have

$$f(u_1, u_2, \lambda) = (u_1 - 250000)(u_2 - 150000) - \lambda_1[u_1 + u_2 - 500000] - \lambda_2[u_1 + 3u_2 - 900000]$$

$$\frac{\partial f}{\partial u_1} = 0 \Rightarrow u_2 - 150000 = \lambda_1 + \lambda_2$$

$$\frac{\partial f}{\partial u_2} = 0 \Rightarrow u_1 - 250000 = \lambda_1 + 3\lambda_2$$

$$\frac{\partial f}{\partial \lambda_1} = 0 \Rightarrow u_1 + u_2 = 500000$$

$$\frac{\partial f}{\partial \lambda_2} = 0 \Rightarrow u_1 + 3u_2 = 900000$$

The solution is: $(300000, 200000)$. Note it satisfies Equations (22.10), (22.11) and (22.12). Therefore, it is the unique Nash bargaining solution.

S	\emptyset	$\{j\}$
$v(S \cup \{i\}) - v(S)$	v_i	$v - v_j$

Table 22.7

□

22.62 Find the Shapley values of the game with $N = \{1, 2\}$ and the characteristic function v . Now consider the bargaining game where $U = \{(u_1, u_2) \mid u_1 + u_2 = v(N), u_1 \geq v(\{1\}), u_2 \geq v(\{2\})\}$ and $d = (v(\{1\}), v(\{2\}))$. Find the bargaining solution of the game (U, d) .

Answer. Since $n = 2$, we have $\gamma(0) = \gamma(1) = \frac{1}{2}$. Denote $v = v(N)$, $v_1 = v(\{1\})$ and $v_2 = v(\{2\})$.

- (i) Shapley value. For player i , Hence the Shapley value for player i is $\frac{v_i + v - v_j}{2}$.
- (ii) To get the Nash bargaining solution, we solve the following problem

$$\max_{u_1 + u_2 = v, u_1 \geq v_1, u_2 \geq v_2} (u_1 - v_1)(u_2 - v_2).$$

The solution is $u_i^* = \frac{v_i + v - v_j}{2}$. Note that we need to check whether $u_i^* \geq v_i$.

Hence, both Nash bargaining solution and the Shapley value give the same result. □

22.63 Relation of Nash bargaining model to Rubinstein bargaining model:

Consider the variant of the bargaining game with alternating offers with exogenous probabilistic breakdown. Assume there is an exogenous probability α of breaking down.

We can assume without loss of generality that $\delta \rightarrow 1$, since the possibility of a breakdown puts pressure to reach an agreement.

It can be seen that this game has a unique subgame perfect equilibrium in which,

- Player 1 proposes x^* and accepts an offer y if and only if $y_1 \geq y_1^*$,
- Player 2 proposes y^* and accepts an offer x if and only if $x_1 \geq x_1^*$,

where

$$x_1^* = \frac{1 - d_2 + (1 - \alpha)d_1}{2 - \alpha}, \quad y_1^* = \frac{(1 - \alpha)(1 - d_2) + d_1}{2 - \alpha}.$$

Letting $\alpha \rightarrow 0$, we have $x_1^* \rightarrow \frac{1}{2} + \frac{1}{2}(d_1 - d_2)$.

The optimization problem

$$\max(x - d_1)(1 - x - d_2),$$

gives Nash bargaining solution $(\frac{1}{2} + \frac{1}{2}(d_1 - d_2), \frac{1}{2} - \frac{1}{2}(d_1 - d_2))$, which coincides with the subgame perfect equilibrium payoffs.

That is, the variant of the bargaining game with alternating offers with exogenous probabilistic breakdown and Nash's axiomatic model, though built entirely of different components, yield the same outcome.

Chapter

23

Supermodular games

Contents

23.1 Lattice	387
23.2 Supermodular function	391
23.3 Supermodular games	392

23.1 Lattice

23.1 Definition: A partial order is a binary relation \geq over a nonempty set S which is reflexive, transitive and antisymmetric, i.e., for all x, y and z in S , we have that:

- $x \geq x$;
- if $x \geq y$ and $y \geq z$, then $x \geq z$;
- if $x \geq y$ and $y \geq x$, then $x = y$.

A set S with a partial order \geq is called a partially ordered set, abbreviated as poset, denoted by (S, \geq) .

23.2 Definition: Let (S, \geq) be a poset, T a nonempty subset of S , and x_0 an element of S , then

- x_0 is a maximal element of T if $x_0 \in T$ and $(\forall x \in T) x \not\geq x_0$;
- x_0 is a minimal element of T if $x_0 \in T$ and $(\forall x \in T) x_0 \not\geq x$;
- x_0 is a largest element of T if $x_0 \in T$ and $(\forall x \in T) x_0 \geq x$;
- x_0 is a smallest element of T if $x_0 \in T$ and $(\forall x \in T) x \geq x_0$;
- x_0 is a upper bound of T if $(\forall x \in T) x_0 \geq x$;
- x_0 is a lower bound of T if $(\forall x \in T) x \geq x_0$;
- x_0 is a supremum of T if x_0 is the smallest upper bound of T ;
- x_0 is a infimum of T if x_0 is the largest lower bound of T .

23.3 Definition: Let (S, \geq) be a poset. (S, \geq) is a lattice if for each two-point set $\{x, y\} \subset S$, there is a supremum for $\{x, y\}$ (denoted by $x \vee^S y$ and called the join of x and y) in S and an infimum (denoted by $x \wedge^S y$ and called the meet of x and y) in S .

The lattice (S, \geq) is complete if every nonempty subset T of S has both infimum (denoted by $\inf_S T$ or $\wedge_{t \in T}^S t$) and supremum (denoted by $\sup_S T$ or $\vee_{t \in T}^S t$) in S .

23.4 Definition: A sublattice (T, \geq) of a lattice (S, \geq) is a nonempty subset of (S, \geq) such that for every pair of elements x and y in T , both $x \vee^S y$ and $x \wedge^S y$ are in T .

A sublattice T of a complete lattice (S, \geq) is called a closed sublattice if for every subset U of T the elements $\vee_{t \in U}^S t$ and $\wedge_{t \in U}^S t$, as defined in S , are in T . (The term “closed” is used here in the lattice theoretical sense)

23.5 Remark: A subset that is a lattice in its own right may not be a sublattice of a larger lattice. An example is as follows:

$$T = \{(0, 0), (1, 0), (0, 1), (2, 2)\} \text{ and } S = \mathbb{R}^2 \text{ with the product order.}$$

23.6 Remark: A subset that is a complete lattice in its own right may not be a closed sublattice of a larger lattice. An example is as follows:

$$T = [0, 1) \cup \{2\} \text{ and } S = [0, 2] \text{ with the natural order.}$$

23.7 Lemma: Let (S, \geq) be a complete lattice, and T a closed sublattice of S . Then T is a complete lattice.

Proof. For any nonempty subset U of T , let $x = \vee_{t \in U}^S t$, then $x \in T$.

- x is an upper bound of U in T .
- For any $y \in T$, if for every $t \in U$, $t \leq y$, then $x \leq y$.

Based on the above discussion, x is smallest upper bound of U in T , i.e., $\vee_{t \in U}^T t = x \in T$. Similarly, we have $\wedge_{t \in U}^T t \in T$. □

23.8 Lemma: Let (S, \geq) be a complete lattice, S' a closed sublattice of S , and T a closed sublattice of S . If T is a subset of S' , then T is a closed sublattice of S' .

Proof. (1) For any nonempty subset U of T , we would like to show $\vee_{t \in U}^{S'} t \in T$ and $\wedge_{t \in U}^{S'} t \in T$.

(2) By Proposition 23.7, we know that S' is a complete lattice, and hence $y \triangleq \vee_{t \in U}^{S'} t \in S'$.

(3) Since T is a closed sublattice of S , we have $x \triangleq \vee_{t \in U}^S t \in T \subseteq S'$.

(4) Since $S' \subseteq S$, we have $y \geq x$.

(5) On the other hand, since $x \in S'$ is an upper bound of U , and $y \in S'$ is the smallest upper bound of U in S' , we have $x \geq y$.

(6) Therefore, $\vee_{t \in U}^{S'} t = y = x = \vee_{t \in U}^S t \in T$.

(7) Similarly, we have $\wedge_{t \in U}^{S'} t \in T$. □

23.9 Definition: Let (S, \geq) be a lattice. For subsets X and Y of S , we have $X \geq_+ Y$ if for any $x \in X$, $y \in Y$, it follows that $x \wedge y \in Y$ and $x \vee y \in X$. We call \geq_+ the strong set order.

23.10 Lemma: If S is a complete lattice, X and Y are in 2^S (the power set of S), and $X \geq_+ Y$, then $\inf X \geq \inf Y$ and $\sup X \geq \sup Y$.

Proof. For any $x \in X$ and $y \in Y$, $x \wedge^S y \in Y$ and $x \vee^S y \in X$ since $X \geq_+ Y$. Therefore, $x \geq x \wedge^S y \geq \inf Y$ and $\sup X \geq x \vee^S y \geq y$, and hence $\inf X \geq \inf Y$ and $\sup X \geq \sup Y$. \square

23.11 Tarski's fixed-point Theorem (Theorem 1 in [Tarski \(1955\)](#)): Let (S, \geq) be a nonempty complete lattice and $f: S \rightarrow S$ an increasing function, if E be the set of all fixed points of f , then E is nonempty and (E, \geq) is a complete lattice.

Step 1. Let $u = \sup\{x \in S \mid f(x) \geq x\}$.

- (1) For every x with $f(x) \geq x$, we have $u \geq x$.
- (2) Since f is increasing, we have $f(u) \geq f(x) \geq x$.
- (3) Since u is the supremum, we have $f(u) \geq u$.
- (4) Since f is increasing, we have $f(f(u)) \geq f(u)$, and hence $f(u) \in \{x \in S \mid f(x) \geq x\}$.
- (5) Since u is supremum, we have $u \geq f(u)$, and hence $f(u) = u$. That is, u is a fixed point of f , and E is nonempty.
- (6) Since $E \subseteq \{x \in S \mid f(x) \geq x\}$, and $u = \sup\{x \in S \mid f(x) \geq x\}$, we have u is the largest element of E .

Similarly we have $\inf\{x \in S \mid f(x) \leq x\}$ is the smallest element of E . \square

Step 2. Let Y be any subset of E , and denote $\{x \in S: \sup Y \geq x \geq \sup Y\}$ by $[\sup Y, \sup S]$.

- (1) Since (S, \geq) is a complete lattice, we have $([\sup Y, \sup S], \geq)$ is a complete lattice.
- (2) For any $x \in Y$, we have $\sup Y \geq x$ and hence $f(\sup Y) \geq f(x) = x$. Therefore $f(\sup Y) \geq \sup Y$.
- (3) For any z with $z \geq \sup Y$ implies $f(z) \geq f(\sup Y) \geq \sup Y$.
- (4) Consider the restriction of f on $[\sup Y, \sup S]$:

$$f': [\sup Y, \sup S] \rightarrow [\sup Y, \sup S],$$

which is an increasing function on the complete lattice $([\sup Y, \sup S], \geq)$.

- (5) Then the infimum v of all fixed points of f' is itself a fixed point of f' .
- (6) Obviously, v is a fixed point of f and in fact the smallest fixed point of f which is an upper bound of all elements of Y ; in other words, v is the supremum of Y in (E, \geq) .

Similarly the infimum of Y in (E, \geq) exists in S . \square

23.12 Theorem (Theorem 1 in [Zhou \(1994\)](#)): Let (S, \geq) be a complete lattice, Ψ is a correspondence from S to S , and E the set of fixed points of Ψ . If $\Psi(s)$ is a nonempty closed sublattice of S for every $s \in S$, and Ψ is increasing in s , then E is a nonempty complete lattice.

Step 1. To show E is nonempty.

- (1) Consider the set

$$C = \{c \in S \mid \text{there exists } x_c \in \Psi(c) \text{ such that } x_c \leq c\}.$$

C is nonempty since $\sup S \in C$. Let $a = \bigwedge_{c \in C}^S c$.

- (2) For any $c \in C$, there exists $x_c \in \Psi(c)$ such that $x_c \leq c$.
- (3) Since Ψ is increasing and $a \leq c$, we have $\Psi(a) \leq_+ \Psi(c)$, and hence there exists $y_c \in \Psi(a)$ such that $y_c \leq x_c (\leq c)$.
- (4) Let $y = \bigwedge_{c \in C}^S y_c$. Since $\Psi(a)$ is a subcomplete sublattice of S , we have $y \in \Psi(a)$.
- (5) Clearly $y = \bigwedge_{c \in C}^S y_c \leq \bigwedge_{c \in C}^S c = a$. Since Ψ is increasing and $y \in \Psi(a)$, there is $z \in \Psi(y)$, such that $z \leq y$. Hence $y \in C$.
- (6) Since $a = \bigwedge_{c \in C}^S a \leq y$. Therefore, $a = y \in \Psi(a)$, i.e., $a \in E$.

□

Step 2: To show $\bigvee_{e \in E}^S e \in E$ and $\bigwedge_{e \in E}^S e \in E$.

Since $E \subseteq C$, based on Step 1, $\bigvee_{e \in E}^S e = a \in E$. Similarly, we can show that $\bigwedge_{e \in E}^S e \in E$.

□

Step 3: To show E is complete lattice.

It suffices to show that $\bigvee_{e \in U}^E e$ and $\bigwedge_{e \in U}^E e$ exist for any $U \subseteq E$. Note that $\bigvee_{e \in U}^E e$ and $\bigwedge_{e \in U}^E e$ are respectively the largest and smallest elements of U in E instead of S .

- (1) Let $b = \bigvee_{e \in U}^S e$, the largest element of U in S .
- (2) For any $e \in U \subseteq E$, since $e \leq b$ and Ψ is increasing, we have $\Psi(e) \leq_+ \Psi(b)$. Moreover, since $e \in \Psi(e)$, there is $x_e \in \Psi(b)$, such that $e \leq x_e$.
- (3) Let $x = \bigvee_{e \in U}^S x_e$, then $x \geq \bigvee_{e \in U}^S e = b$. Since $\Psi(b)$ is a subcomplete sublattice, $x \in \Psi(b)$.
- (4) Let $S' = [b, \sup S]$, Φ a correspondence from S' to S' defined by $\Phi(s) = \Psi(s) \cap S'$ for all $s \in S'$.
For any $s \geq b$, since Ψ is increasing and $x \in \Psi(b)$, there is $x_s \in \Psi(s)$ with $x_s \geq x \geq b$, and hence $\Phi(s)$ is nonempty for every $s \in S'$.
Since both $\Psi(s)$ and S' are subcomplete sublattices of S for every $s \in S'$, $\Phi(s)$ must be a subcomplete sublattice of S' .
Since both Ψ and Ω , which assigns each $s \in S'$ the constant interval S' , are increasing on S , $\Phi = \Psi \cap \Omega$ is increasing on S' .
Hence, S' and Φ satisfy the assumptions of the theorem. Therefore, if we let $b' = \bigwedge_{e \in E'}^{S'} e$, where E' is the set of fixed points of Φ on S' , then $b' \in E'$ according to Steps 1 and 2. Since $E' = E \cap S'$, b' is indeed the smallest fixed point that is larger than or equal to b .
- (5) Since $b = \bigvee_{e \in U}^S e$, for any $e \in U$, $b \geq e$ and hence $b' \geq b \geq e$. That is, b' is an upper bound of U in E .
For any upper bound b'' of U in E , b'' is also an upper bound of U in S , and hence we have $b'' \geq b$. Note that b' is the smallest fixed point that is larger than or equal to b , then $b'' \geq b'$.
Therefore b' is the smallest upper bound of U in E , i.e., $\bigvee_{t \in U}^E t = b' \in E$.

□

☞ 23.13 Let (S, \geq) be a lattice. An interval is a set of the form $[x, y] \triangleq \{z \in S \mid x \leq z \leq y\}$.

By the interval topology of a lattice (S, \geq) , we mean that defined by taking the sets of the type $(-\infty, z]$ and $[z, \infty)$ to form a sub-basis for closed sets.

☞ 23.14 Theorem: A lattice is complete if and only if it is compact in its interval topology.

Proof. See Theorem 20 in Birkhoff (1967).

□

A sublattice of a complete lattice is closed in the lattice theoretical sense if and only if it is topologically closed in the interval topology.

23.2 Supermodular function

23.15 Suppose f is a real-valued function on a lattice (S, \geq) .

- If

$$f(x \vee y) + f(x \wedge y) \geq f(x) + f(y)$$

for all x and y in S , then f is called to be supermodular on S .

- If

$$f(x \vee y) + f(x \wedge y) > f(x) + f(y)$$

for all noncomparable elements x and y in S , then f is called to be strictly supermodular on S .

23.16 Let S and T be two posets.

- A function $f(x, t): S \times T \rightarrow \mathbb{R}$ has increasing differences in (x, t) if for all $x \geq x'$, the difference $f(x, t) - f(x', t)$ is increasing in t .
- A function $f(x, t): S \times T \rightarrow \mathbb{R}$ has strictly increasing differences in (x, t) if for all $x \geq x'$, the difference $f(x, t) - f(x', t)$ is strictly increasing in t .

23.17 Proposition:

- (Theorem 4.1 in [Topkis \(1978\)](#)) Let (S, \geq) be a nonempty lattice, and $f: S \rightarrow \mathbb{R}$ a supermodular function on S . Then $\arg \max_{x \in S} f(x)$ is a sublattice of S .
- (Corollary 4.1 in [Topkis \(1978\)](#)) Let (S, \geq) be a nonempty complete lattice, and $f: S \rightarrow \mathbb{R}$ a supermodular function on S which is upper semi-continuous. Then $\arg \max_{x \in S} f(x)$ is a nonempty closed sublattice of S .

Proof. (a) For any y and z in $\arg \max_{x \in S} f(x)$, since f is supermodular and y and z are in $\arg \max_{x \in S} f(x)$, we have $0 \geq f(y \wedge z) - f(z) \geq f(y) - f(y \vee z) \geq 0$. Hence, $y \wedge z$ and $y \vee z$ are in $\arg \max_{x \in S} f(x)$. Therefore, $\arg \max_{x \in S} f(x)$ is a sublattice of S . (Hence itself a lattice)

- By the previous part, $\arg \max_{x \in S} f(x)$ is a sublattice of S .

By Theorem 23.14, S is compact with the interval topology. Since f is upper semi-continuous, we have that $\arg \max_{x \in S} f(x)$ is nonempty and compact. By Theorem 23.14 again, $\arg \max_{x \in S} f(x)$ is closed.

□

23.18 Topkis's monotonicity theorem: Let (S, \geq) be a lattice, and T a poset. Suppose $f(x, t): S \times T \rightarrow \mathbb{R}$ is supermodular in $x \in S$ for fixed $t \in T$ and has increasing differences in (x, t) .

- (Lemma 6.1 in [Topkis \(1978\)](#)) Suppose that $t \geq t'$ and $X \geq_+ X'$, then $\arg \max_{x \in X} f(x, t) \geq_+ \arg \max_{x \in X'} f(x, t')$.
- (Theorem 6.1 in [Topkis \(1978\)](#), Theorem 5 in [Milgrom and Shannon \(1994\)](#)) In particular (when $t = t'$ and $X =_+ X'$), the set of maximizers of f is a sublattice.

Proof. (a) By Theorem 23.17, $\arg \max_{x \in X} f(x, t)$ is a sublattice for any sublattice X and $t \in T$. Choose $y \in \arg \max_{x \in X} f(x, t)$ and $z \in \arg \max_{x \in X'} f(x, t')$. Since $X \geq_+ X'$, we have $y \wedge z \in X'$ and $y \vee z \in X$. Since $f(x, t)$ is supermodular in $x \in S$ for fixed $t \in T$, we have

$$0 \geq f(y \vee z, t) - f(y, t) \geq f(z, t) - f(y \wedge z, t).$$

Since $f(x, t)$ has increasing differences in (x, t) ,

$$f(z, t) - f(y \wedge z, t) \geq f(z, t') - f(y \wedge z, t') \geq 0.$$

Hence, $f(y \vee z, t) = f(y, t)$ and $f(z, t') = f(y \wedge z, t')$, i.e., $y \vee z \in \arg \max_{x \in X} f(x, t)$ and $y \wedge z \in \arg \max_{x \in X'} f(x, t')$. Therefore, $\arg \max_{x \in X} f(x, t) \geq_+ \arg \max_{x \in X'} f(x, t')$.

(b) Trivial. □

23.19 Theorem: Let (S, \geq) be a complete lattice, and T a poset. Suppose $f(x, t): S \times T \rightarrow \mathbb{R}$ is supermodular in $x \in S$ for fixed $t \in T$ and upper semi-continuous, and has increasing differences in (x, t) . Then $\arg \max_{x \in S} f(x, t)$ has a smallest element $\underline{s}(t) \triangleq \inf \arg \max_{x \in S} f(x, t)$, and a largest element $\bar{s}(t) \triangleq \sup \arg \max_{x \in S} f(x, t)$, and $\underline{s}(t)$ and $\bar{s}(t)$ are both increasing in $t \in T$.

Proof. By Theorem 23.17, for each $t \in T$, $\arg \max_{x \in S} f(x, t)$ is a nonempty complete lattice, and hence has a smallest element $\underline{s}(t)$ and a largest element $\bar{s}(t)$.

If $t \geq t'$, then $\arg \max_{x \in S} f(x, t) \geq_+ \arg \max_{x \in S} f(x, t')$ by Theorem 23.18. By Lemma 23.10, $\underline{s}(t)$ and $\bar{s}(t)$ are increasing. □

23.3 Supermodular games

23.20 Model:

- $N = \{1, 2, \dots, n\}$ is the set of players.
- Each player $i \in N$ has a strategy set S_i with typical element x_i . Each strategy set S_i comes with a partial order \geq_i .
- $S = \times_i S_i$ endowed with product order, i.e., $x \geq x'$ if and only if $x_i \geq_i x'_i$ for all $i \in N$.
- Player i 's payoff function is $f_i(x_i, x_{-i})$.
- Player i 's best-response correspondence $\Psi_i: S_{-i} \rightarrow S_i$,

$$\Psi_i(x_{-i}) = \arg \max_{x_i \in S_i} f_i(x_i, x_{-i}).$$

$$\Psi = \times_i \Psi_i: S \rightarrow S.$$

- The set of Nash equilibria $E = \{x \in S \mid x \in \Psi(x)\}$.

23.21 Theorem (Theorem 3.1 in Topkis (1979)): Suppose n is finite, and for each $i \in N$,

- (S_i, \geq_i) is a complete lattice;
- f_i is supermodular in x_i for fixed x_{-i} ;
- f_i has increasing differences in (x_i, x_{-i}) ;
- f_i is upper semi-continuous in x_i for fixed x_{-i} .

Then the set of pure-strategy Nash equilibria of the game

$$\Gamma = \langle N, (S_i, \geq_i)_{i \in N}, (f_i)_{i \in N} \rangle$$

is nonempty and has a largest equilibrium \bar{x} and a smallest equilibrium \underline{x} .

Proof. By Theorem 23.19, $\sup \Psi_i$ and $\inf \Psi_i$ are increasing selections of Ψ_i , and therefore $\inf \Psi = (\inf \Psi_1, \dots, \inf \Psi_n)$ is an increasing selection of Ψ . From Tarski's fixed-point theorem (Theorem 23.11), we have that $\underline{x} \triangleq \inf\{x \in A \mid \inf \Psi(x) \leq x\}$ is a fixed point of $\inf \Psi_i$, i.e., $\underline{x} = \inf \Psi(\underline{x}) \in \Psi(\underline{x})$, and hence \underline{x} is a fixed point of Ψ .

In the following, we will show $\underline{x} = \inf E$: for any $y \in E$, we have $y \in \Psi(y)$ and $y \geq \inf \Psi(y)$, and hence $y \geq \inf\{x \in S \mid x \geq \inf \Psi(x)\} = \underline{x}$. \square

23.22 Theorem (Theorem 2 in Zhou (1994)): Suppose n is finite, and for each $i \in N$,

- (S_i, \geq_i) is a complete lattice;
- f_i is supermodular in $x_i \in S_i$;
- f_i has increasing differences in (x_i, x_{-i}) ;
- f_i is upper semi-continuous in x_i for fixed x_{-i} .

Then the set of pure-strategy equilibrium of the game

$$\Gamma = \langle N, (S_i, \geq_i)_{i \in N}, (f_i)_{i \in N} \rangle$$

is nonempty complete lattice.

Proof. By Theorem 23.17, we have $\Psi_i(x)$ is a closed sublattice in S_i for every $x \in S$.

By Theorem 23.18, we have $\Psi_i(x)$ is an increasing correspondence.

Then apply Theorem 23.12. \square

Large games 1: large strategic games

24.1 Example (Example 1 in [Noguchi and Zame \(2006\)](#)): We consider a large number of agents living on the banks of a long river. We identify locations with the interval $T = [0, 1]$, and assume the population is uniformly distributed along the river, so the population measure τ is the Lebesgue measure.

There are two commodities, each perfectly divisible. Each agent is endowed with one unit of each good: $e(t) \equiv (1, 1)$. Agents derive utility from their own consumption, but suffer a pollution externality from the consumption of others who live upstream from them. If we choose directions so that upstream from s means to the left of s , and write $f(t) = (f_1(t), f_2(t))$ for the consumption of an agent located at t , then the externality experienced by an agent located at s is

$$\eta(s, f) = \int_0^s f(t) \, dt.$$

Note that η is two-dimensional; write $\eta = (\eta_1, \eta_2)$ for the components of η . The utility of an agent located at s who consumes the bundle (x_1, x_2) when the consumption of all agents is described by f is

$$u_s(x_1, x_2, f) = (2 - \eta_1(s, f))x_1^2 + (2 - \eta_2(s, f))x_2^2.$$

In the individualistic framework, an equilibrium consists of prices p_1, p_2 (without loss, normalize so that $p_1 + p_2 = 1$) and a consumption allocation $f = (f_1, f_2): T \rightarrow \mathbb{R}_+^2$ so that almost every agent optimizes in his/her budget set and the market clears.

We claim that no such equilibrium exists.

Chapter 25

Large games 2: large distributional games

Chapter 26

Stochastic games

Bibliography

- Dilip Abreu, Prajit Dutta and Lones Smith, The folk theorem for repeated games: a NEU condition, *Econometrica* **62** (1994), 939–948. [199](#)
- George A. Akerlof, The market for “lemons”: quality uncertainty and the market mechanism, *The Quarterly Journal of Economics* **84** (1970), 488–500. [224](#)
- Attila Ambrus and Satoru Takahashi, Multi-sender cheap talk with restricted state spaces, *Theoretical Economics* **3:1** (2008), 1–27. [252](#)
- Robert Aumann, Agreeing to disagree, *Annals of Statistics* **4** (1976), 1236–1239. [138](#)
- Robert Aumann and Adam Brandenburger, Epistemic conditions for Nash equilibrium, *Econometrica* **63** (1995), 1161–1180. [150](#)
- Marco Battaglini, Multiple referrals and multidimensional cheap talk, *Econometrica* **70:4** (2002), 1379–1401. [252](#)
- Garrett Birkhoff, *Lattice Theory*, 3rd ed., American Mathematical Society (Colloquium Publications, Providence, RI). [390](#)
- Tilman Borgers, *An Introduction to the Theory of Mechanism Design*, Oxford University Press, 2015. [282](#), [288](#)
- Adam Brandenburger and Eddie Dekel, Hierarchies of beliefs and common knowledge, *Journal of Economic Theory* **59** (1993), 189–198. [73](#)
- Ying Chen, Navin Kartik and Joel Sobel, Selecting cheap-talk equilibria, *Econometrica* **76:1** (2008), 117–136. [252](#)
- In-Koo Cho and David M. Kreps, Signaling games and stable equilibria, *The Quarterly Journal of Economics* **102:2** (1987), 179–222. [229](#)
- Luis C. Corchón, The theory of contests: a survey, *Review of Economic Design* **11** (2007), 69–100. [45](#)
- Pascal Courty and Hao Li, Sequential screening, *The Review of Economic Studies* **67:4** (2000), 697–717. [282](#)
- Vincent P. Crawford and Joel Sobel, Strategic information transmission, *Econometrica* **50:6** (1982), 1431–1451. [240](#), [245](#)
- Eddie Dekel, Barton L. Lipman and Aldo Rustichini, Standard state-space models preclude unawareness, *Econometrica* **66** (1998), 159–173. [145](#)
- Péter Eső and Balázs Szentes, Optimal information disclosure in auctions and the hangicao auction, *Review of Economics Studies* **74** (2007), 705–731. [260](#)

- Mark Fey, Symmetric games with only asymmetric equilibria, *Games and Economic Behavior* 75 (2012), 424–427. [44](#)
- James W. Friedman, A non-cooperative equilibrium For supergames, *Review of Economic Studies* 38:1 (1971), 1–12. [195](#)
- Drew Fudenberg and David Levine.
- Drew Fudenberg and Eric Maskin, The folk theorem in repeated games with discounting or with incomplete information, *Econometrica* 54:3 (1986), 533–554. [196](#), [197](#), [198](#)
- Drew Fudenberg and Jean Tirole, *Game Theory*, The MIT Press, 1991. [173](#), [183](#)
- Matthew Gentzkow and Emir Kamenica, Competition in persuasion, *Review of Economic Studies* forthcoming. [252](#)
- Matthew Gentzkow and Emir Kamenica, Bayesian persuasion with multiple senders and rich signal spaces, working paper. [252](#)
- Robert Gibbons, *A Primer in Game Theory*, 1992. [173](#)
- Srihari Govindan, Philip J. Reny and Arthur J. Robson, A short proof of Harsanyi's purification theorem, *Games and Economics Behavior* 45 (2003), 363–374. [107](#)
- Michael Harrison and David Kreps, Speculative investor behavior in a stock market with heterogeneous expectations, *Quarterly Journal of Economics* (1978), 323–336. [141](#)
- John C. Harsanyi, Games with incomplete information played by 'Bayesian' players. *Management Science* 14 (1967–68) 159–182, 320–334, and 486–502, Parts I–III. [71](#)
- John C. Harsanyi, Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points, *International Journal of Game Theory* 2:1 (1973), 1–23. [103](#)
- Emir Kamenica and Matthew Gentzkow, Bayesian persuasion, *The American Economic Review* 101:6, 2590–2615. [252](#), [258](#), [259](#)
- Vijay Krishna and Eliot Maenner, Convex potentials with an application to mechanism design, *Econometrica* 69:4 (2001), 1113–1119. [321](#)
- Vijay Krishna and John Morgan, The art of conversation: eliciting information from experts through multi-stage communication, *Journal of Economic Theory* 117 (2004), 147–179. [252](#)
- Vijay Krishna and John Morgan, Cheap Talk, in *The New Palgrave Dictionary of Economics* (Steven Durlauf and Lawrence Blume eds.), New York: Palgrave Macmillan, 2nd edition. [240](#)
- Jean-Jacques Laffont and David Martimort, *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press, 2002. [261](#)
- Fei Li and Peter Norman, On Bayesian persuasion with multiple senders, working paper. [252](#)
- Erzo G. J. Luttmer and Thomas Mariotti, The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A comment, *Econometrica* 71:6 (2003), 1909–1911. [206](#)
- George J. Mailath and Larry Samuelson, *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, 2006. [183](#), [188](#)
- Nahum D. Melumad and Toshiyuki Shibano, Communication in settings with no transfers, *RAND Journal of Economics* 22:2 (1991), 173–198. [241](#), [248](#)

- Jean-François Mertens and Shmuel Zamir, Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory* **14:1**, 1–29. [73](#)
- Paul R. Milgrom and Chris Shannon, Monotone comparative statics, *Econometrica*, **62** (1994), 157–180. [391](#)
- Paul Milgrom and Nancy Stokey, Information, trade and common knowledge, *Journal of Economic Theory* **26:1** (1982), 17–27. [140](#)
- Roger B. Myerson and Mark A. Satterthwaite, Efficient mechanisms for bilateral trading, *Econometrica* **29** (1983), 265–281. [346](#)
- Mitsunori Noguchi and William R. Zame, Competitive markets with externalities, *Theoretical Economics* **1** (2006), 143–166. [395](#)
- Martin J. Osborne and Ariel Rubinstein, *Bargaining and Markets*, San Diego: Academic Press, 1990. [173](#)
- Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, The MIT Press, 1994. [173](#)
- Luis Rayo and Ilya Segal, Optimal information disclosure, *Journal of Political Economy* **118** (2010), 949–987. [260](#)
- Philip J. Reny, On the existence of pure and mixed strategy Nash equilibria in discontinuous games, *Econometrica* **67** (1999), 1029–1056. [47](#)
- Michael Rothschild and Joseph Stiglitz, Equilibrium in competitive insurance markets: an essay on the economics of imperfect information, *The Quarterly Journal of Economics* **90:4** (1976), 629–649. [282](#)
- Halsey L. Royden, *Real Analysis*, Third Edition, Macmillan, New York, NY, 1988. [274](#)
- Ariel Rubinstein, Perfect equilibrium in a bargaining model, *Econometrica* **50:1** (1982), 97–109. [173](#)
- Dov Samet, Common priors and separation of convex sets, *Games and Economic Behavior* **24** (1998), 172–174. [143](#), [144](#)
- Michael Spence, Job market signaling, *The Quarterly Journal of Economics* **87** (1970), 355–374. [238](#)
- Ingolf Ståhl, *Bargaining Theory*, Stockholm: Economics Research Institute at the Stockholm School of Economics, 1972. [173](#)
- Steven Tadelis, *Game Theory: An Introduction*, Princeton University Press, 2013. [244](#)
- Satoru Takahashi, Lecture notes on Repeated Games, National University of Singapore, 2014. Available at https://ap4.fas.nus.edu.sg/fass/ecsst/teach/lecture_notes_on_repeated_games.pdf. [183](#)
- Alfred Tarski, A lattice-theoretical fixpoint theorem and its applications, *Pacific Journal of Mathematics*, **5** (1955), 285–309. [389](#)
- Donald M. Topkis, Minimizing a submodular function on a lattice, *Operations Research*, **26** (1978), 305–321. [391](#)
- Donald M. Topkis, Equilibrium points in nonzero-sum n -person submodular games, *SIAM Journal of Control and Optimization*, **17** (1979), 773–787. [392](#)
- Xavier Vives, Nash equilibrium with strategic complementarities, *Journal of Mathematical Economics*, **19** (1990), 305–321.
- Yun Wang, Bayesian persuasion with multiple receivers, working paper. [252](#)
- Jun Zhang and Junjie Zhou, Information disclosure in contests: A Bayesian persuasion approach, *Economic Journal* forthcoming. [260](#)
- Lin Zhou, The set of Nash equilibria of a supermodular game is a complete lattice, *Game and Economic Behavior*, **7** (1994), 295–300. [389](#), [393](#)