

Advanced Microeconomics Theory

Vijay Krishna **Professor**
L^AT_EX by Rui Zhou
Penn State University

Spring 2026

Contents

Preface	i
Notation	iv
Introduction	viii
I Foundations	1
1 Representing a Game	2
1.1 Game Tree	2
1.2 Information Sets	3
1.3 Strategies	4
1.4 Normal Form	5
1.5 Dominated Strategies and Best Responses	7
1.6 Iterated Elimination of Dominated Strategies	8
2 Nash Equilibrium and Subgame Perfection	10
2.1 Nash Equilibrium and Nash's Theorem	10
2.2 Iterated Dominance and Rationalizability	12
2.3 Potential Games	14
2.4 Supermodular Games	16
2.5 Subgame Perfect Nash Equilibrium	20
II Bargaining	26
3 Two-Person Bargaining	27
3.1 Nash Bargaining Solution	27
3.2 Dividing a Dollar	31
3.3 Alternating Offer Model	32
3.4 Unequal Discounting	36
III Auctions and Mechanism Design	38
4 Private Value Auctions	39
4.1 Second Price Sealed Bid Auction (SPA)	39

4.2	First Price Sealed Bid Auction (FPA)	40
4.3	Reserve Prices	47
4.4	Revenue Equivalence Theorem	51
4.5	Risk Attitude and the Breakdown of Revenue Equivalence	54
5	Mechanism Design	57
5.1	Direct Mechanism and Revelation Principle	57
5.2	Optimal Mechanism Design (Myerson, 1981)	59
5.3	The Vickrey-Clarke-Groves (VCG) Mechanism	68
IV	Matching	80
6	Matching Theory	81
6.1	The Housing Market: Allocation with Endowments	81
6.2	The Assignment Problem: Allocation without Endowments	93
6.3	Two-Sided Matching	100
6.4	The Boston Mechanism (Immediate Acceptance)	119
V	Information and Dynamic Games	125
7	Common Knowledge	126
7.1	Motivating Example: The Corrupted Industry	126
7.2	Formal Definition	128
7.3	Aumann's Agreement Theorem	129
7.4	The Coordinated Attack Problem	130
8	Dynamic Games	134
8.1	Repeated Games	134
8.2	The Folk Theorem	142
8.3	Finitely Repeated Games	148
8.4	The Chain Store Paradox	152
9	Perfect Bayesian Equilibrium	160
9.1	From SPE to PBE	160
9.2	Definition: Strategies and Beliefs	161
9.3	Spence Signaling	163
9.4	Reputation in the Finitely Repeated PD	173
VI	Problem Sets and Solutions	178
	Problem Set 1	179
	Problem Set 2	185
	Problem Set 3	190

Problem Set 4	196
Problem Set 5	202
Problem Set 6	208
Problem Set 7	211
Problem Set 8	216
Problem Set 9	221
Problem Set 10	226
Problem Set 11	232
Problem Set 12	235
Problem Set 13	240
VII Exams and Solutions	246
First Mid-Term Exam (2025)	247
Second Mid-Term Exam (2025)	253
Final Exam (2025)	258
First Mid-Term Exam (2026)	263
Second Mid-Term Exam (2026)	269
Final Exam (2026)	275
Bibliography	281

Preface

These notes accompany the Spring 2026 offering of ECON 521 (Advanced Microeconomics Theory) at Penn State, taught by Professor Vijay Krishna. The course is the standard first-year PhD sequence on game theory, mechanism design, and matching—the strategic side of microeconomics. The notes were typeset by the student author over the course of the semester and reflect both lecture material and supplementary reading.

Audience

The primary audience is a first-year economics PhD student who has been through (or is concurrently taking) graduate consumer theory and is comfortable with measure-theoretic probability and basic real analysis. Advanced undergraduates with strong mathematical preparation should also find the material accessible, particularly Chapters 1–3. The book is self-contained in the sense that every theorem used is either proved or stated explicitly; it is not self-contained in the sense that the reader is expected to fill in routine analysis (uniform convergence, dominated convergence, basic topology) without prompting.

Structure of the Book

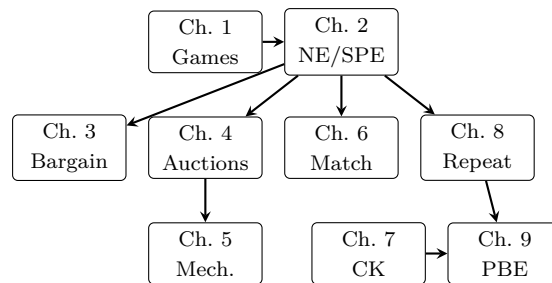
The book is organized into six parts:

- **Part I (Foundations).** Game representation in extensive and normal forms, dominance and rationalizability, Nash equilibrium and its refinements (subgame perfection).
- **Part II (Bargaining).** Nash bargaining solution, the Rubinstein alternating-offer model, and the bridge between cooperative and non-cooperative bargaining theory.
- **Part III (Auctions and Mechanism Design).** Single-object private-value auctions, revenue equivalence, optimal auctions, and the general theory of incentive compatibility (VCG, AGV, Myerson).
- **Part IV (Matching).** Two-sided matching markets, the Gale-Shapley algorithm, stability and strategy-proofness, and one-sided matching with priorities.
- **Part V (Information and Dynamic Games).** Common knowledge and its operational role; infinitely and finitely repeated games and the Folk Theorem; perfect Bayesian equilibrium and signaling.

- **Part VI (Problem Sets and Solutions)** and **Part VII (Exams and Solutions)**.
Verbatim problem statements followed by worked solutions for the thirteen homework sets and the five exams (2025 and 2026 sittings) of Econ 521.

Chapter Dependency

The chapters are not strictly linear. The following is a rough dependency map; an arrow $A \rightarrow B$ means “familiarity with A helps when reading B .”



A reader with limited time can take the following minimal paths:

- *Auctions and mechanism design only:* $1 \rightarrow 2 \rightarrow 4 \rightarrow 5$.
- *Bargaining and dynamic games only:* $1 \rightarrow 2 \rightarrow 3$, then $1 \rightarrow 2 \rightarrow 8 \rightarrow 9$.
- *Matching only:* Chapter 6 is essentially self-contained; the reader needs only the language of preferences and strategy from Chapter 1.

Conventions and Boxed Environments

Several types of coloured boxes appear throughout the book. Each type has a distinct title-bar tint that the reader should learn to recognize on first sight:

Definition

A formal definition of a term used later.

Theorem

A central mathematical statement, usually with an accompanying proof.

Lemma

An auxiliary result used inside a proof.

Corollary

An immediate consequence of a theorem or proposition.

Proposition

A formal statement, less central than a theorem.

Claim

A short assertion proved in line; one step inside a longer argument.

Example, Remark
(cyan side-bar)

A worked illustration or informal commentary; rendered with a left-side coloured bar instead of a full title block. The body of the text never depends on a remark, so remarks are skippable on a first reading.

Mathematical conventions: vectors are bold (e.g. \mathbf{x}), random variables are uppercase (e.g. X_i), realizations are lowercase (x_i), and player indices are subscripts. The notation s_{-i} refers to the strategy profile of all players *other than* i . A complete symbol list appears in the Notation Table that follows this preface.

Using the Problem Sets

Parts VI and VII reproduce the verbatim statement of every problem from the 2026 course followed by a worked solution. Each chapter starts with a *Problems* section—read this first and attempt the problems before turning to the *Solutions* section that follows. The two are separated by a page break for exactly this reason.

Acknowledgements

I thank Professor Vijay Krishna for an outstanding course; my classmates for many helpful conversations; and the open-source L^AT_EX community for the tools that made this typesetting possible. All errors, omissions, and infelicities of language are mine alone.

Rui Zhou
Spring 2026

Notation

The following symbols are used throughout the book. Specialized notation is introduced where it first appears.

Sets and numbers

$\mathbb{R}, \mathbb{R}_{\geq 0}, \mathbb{R}_{++}$	Real numbers; non-negative reals; strictly positive reals
$\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{C}$	Natural numbers, integers, rationals, complex numbers
$[n]$	The set $\{1, 2, \dots, n\}$
$\Delta(X)$	Set of probability distributions over a (measurable) set X
$\mathbf{1}_A$	Indicator function of the set A
$\mathbf{0}, \mathbf{1}$	Vector of zeros, vector of ones
$ S $	Cardinality of a finite set S
$\mathcal{P}(X)$	Power set of X

Probability and integration

X, Y, Z	Random variables (uppercase)
x, y, z	Realizations (lowercase)
$X \sim F$	X has cumulative distribution function F
$\Pr(\cdot) A$	Probability of the event A
$\mathbb{E}[X], \mathbb{E}_F[X]$	Expectation of X (under distribution F when ambiguous)
f, g	Probability density functions corresponding to CDFs F, G
Y_1	The maximum of $n - 1$ iid random values, $\max_{i \neq 1} X_i$ (Ch. 6)
$Y_k^{(n)}$	The k -th order statistic of n iid draws
$U[a, b]$	Uniform distribution on $[a, b]$

Games and players

$N = \{1, 2, \dots, n\}$	Set of players
i, j	Generic player indices
$-i$	“All players other than i ”
S_i	Pure strategy set of player i
$S = \prod_i S_i$	Set of pure strategy profiles
$s_i \in S_i$	A pure strategy of player i
$s = (s_1, \dots, s_n)$	A pure strategy profile

s_{-i}	The profile s with i 's component deleted
$\sigma_i \in \Delta(S_i)$	A mixed strategy of player i
$\sigma = (\sigma_1, \dots, \sigma_n)$	A mixed strategy profile
$u_i : S \rightarrow \mathbb{R}$	Payoff function of player i
$u_i(\sigma)$	Expected payoff under mixed profile σ
$G = (S_i, u_i)_{i=1}^n$	A finite normal-form game
\succsim_i, \succ_i	Weak / strict preference of player i
$\mu_i \in \Delta(S_{-i})$	Belief of player i over opponents' play
$B_i(\sigma_{-i})$	Best-response correspondence
BR_i	Best-response (singleton case)

Extensive-form games

X	Set of nodes in a game tree
$Z \subseteq X$	Set of terminal nodes
h	A history (path from root)
h^t	Length- t history
H_i	Information sets of player i
$A(h)$	Set of available actions at history h
\mathcal{H}	Set of all subgames
$\sigma _h$	Restriction of σ to the subgame rooted at h

Auctions and mechanism design

$X_i \sim F$	Bidder i 's private value, drawn iid from F
\bar{x}	Upper bound of the value support
$\beta(\cdot)$	Symmetric bidding strategy
b_i	Bid of player i
$Q_i^*(x)$	Allocation rule (probability i wins given report x)
$q_i(x_i)$	Interim winning probability of type x_i
$M_i^*(x)$	Ex-post payment rule
$m_i(x_i)$	Interim expected payment of type x_i
$U_i(x_i)$	Interim expected utility (information rent) of type x_i
r	Reserve price
$\varphi(x) = x - \frac{1-F(x)}{f(x)}$	Myerson's virtual value
$W(x)$	Total surplus achievable at value profile x
SPA, FPA, AP	Second-price, first-price, all-pay sealed-bid auctions

Matching

\mathcal{S}, \mathcal{C}	Sets of students / colleges (or men / women)
μ	A matching, $\mu : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$
$\mu(s), \mu^{-1}(c)$	Partner(s) of student s / college c under μ
q_c	Quota of college c

\succ_s, \succ_c	Strict preference of s over colleges / of c over students
DA	Deferred-acceptance algorithm (Gale-Shapley)
TTC	Top-trading-cycles algorithm
SD	Serial dictatorship

Repeated and dynamic games

$G(T)$	T -period repeated game
G^∞	Infinitely repeated game
$\delta \in (0, 1)$	Common discount factor
$V_i(\sigma h)$	Continuation value to i from σ at history h
\bar{v}_i	Player i 's minmax payoff
F^*	Set of feasible and individually-rational payoffs (Folk Theorem)

Bargaining

$U \subseteq \mathbb{R}^n$	Feasible utility set
$d \in U$	Disagreement (status-quo) point
$F(U, d)$	A bargaining solution
x^*, y^*	Equilibrium offers in the alternating-offer Rubinstein game

Operators and shorthand

$\arg \max, \arg \min$	Maximizer / minimizer of an expression
$f \circ g$	Composition of functions
$\nabla f, Df$	Gradient / Jacobian
$f _A$	Restriction of f to A
$\lfloor x \rfloor, \lceil x \rceil$	Floor and ceiling
$\langle \cdot, \cdot \rangle$	Inner product
$x \perp y$	x orthogonal / independent of y (context dependent)
\hookrightarrow	Inclusion / embedding
\rightrightarrows	A correspondence (multi-valued mapping)
\square	End of proof

Acronyms


NE	Nash equilibrium
SPE	Subgame-perfect equilibrium
PBE	Perfect Bayesian equilibrium
IESDS	Iterated elimination of strictly dominated strategies
IEWDS	Iterated elimination of weakly dominated strategies
IUD	Iteratively undominated set
RAT	Rationalizable set
IIA	Independence of irrelevant alternatives
IC	Incentive compatibility
IR	Individual rationality
DSIC	Dominant-strategy incentive compatibility
BNE	Bayes-Nash equilibrium
VCG	Vickrey-Clarke-Groves mechanism
AGV	Arrow-d'Aspremont-Gérard-Varet mechanism
PD	Prisoner's dilemma
CK	Common knowledge

Introduction

Remark (Reading This Chapter).

This chapter is informal motivation, not formal theory. Three short vignettes—a traffic network, a guessing game, a model of herding—preview the recurring themes of strategic interaction, information, and equilibrium that the rest of the book develops rigorously. Formal definitions begin in Chapter 1.

Example (Braess' Paradox).



The core idea to solve for equilibrium is that, in equilibrium, costs (the time spent on each possible route from A to D) must be equal on all paths. Before the shortcut is connected, the equilibrium travel time for any driver is 83 minutes. After the shortcut is built, the new equilibrium travel time becomes 92, which is actually greater than the original.

In the example, although the network capacity increased by adding the shortcut, the individual cost for every agent increased. This is due to the negative externality: when an agent switches to the shortcut, they reduce their own time marginally but impose a heavy congestion cost on the shared links AB and CD . This cost is not internalized by the selfish agents.

Remark.

- In a multi-agent situation, having more choices can make you (and even everyone) worse off.
- In a single-person decision problem, however, having more choices always makes you

at least as well off.

Example (Coin Toss with Sequential Moves).

Agent A announces her guess first; B announces hers after observing A's. If both agents make the same guess, both get 1. Otherwise, nature determines the winner by coin toss: whichever guess matches the coin toss wins 4, the other gets 0.

Coin	A	B	Payoff
H	H	H	(1, 1)
T	H	H	(1, 1)
H	T	T	(1, 1)
T	T	T	(1, 1)
H	H	T	(4, 0)
H	T	H	(0, 4)
T	H	T	(0, 4)
T	T	H	(4, 0)

If the coin toss is unavailable to both of them, B will make a guess opposite to A's as long as B is not too risk-averse. The expected payoff is (2, 2).

If instead the coin toss is observed by A and this is common knowledge, A's best strategy is to guess consistently with the actual coin toss, and B's best response is to report the same as A. Both then receive a certain payoff of (1, 1).

Remark.

Information can have negative value! In the example, making the coin toss *commonly known* to be observed by A actually hurts A. By contrast, if the revelation is A's *private* knowledge, the information must have nonnegative value.

Example (Herding via Information Cascades).

Suppose there are two restaurants in a town, A and B, with common prior

$$\Pr(A > B) = p > \frac{1}{2},$$

where $A > B$ denotes the event that restaurant A is better than B. Everyone is therefore mildly inclined toward A absent further information. Two private signals α and β are available, with

$$\begin{aligned}\Pr(\alpha|A > B) &= q, \\ \Pr(\beta|B > A) &= q,\end{aligned}$$

where $q > \frac{1}{2}$ and $q > p$.^a


The t -th agent observes all $(t - 1)$ prior agents' choices but not their private signals. Suppose the realized signal sequence is $(\alpha, \beta, \beta, \dots, \beta)$.

Agent 1. Receives α , which reinforces the prior. Since $\Pr(A > B|\alpha) > \frac{1}{2}$, she chooses A .

Agent 2. Infers from Agent 1's choice that Agent 1 must have received α . Agent 2's own signal is β . The two signals cancel out,^b so the prior dictates: Agent 2 also chooses A .

Agent 3. Infers Agent 1's α from her choice. For Agent 2: if Agent 2 had α , Agent 2 would have chosen A ; if Agent 2 had β , the two signals would cancel and Agent 2 would still choose A by the prior. So Agent 2's choice is *uninformative* about her private signal. Agent 3 therefore only learns $s_1 = \alpha$ from history; combined with her own $s_3 = \beta$, the two cancel and Agent 3 chooses A .

The same logic propagates: every subsequent agent ignores her private signal and chooses A . This is the *herding* phenomenon—an information cascade in which collective behavior locks onto the prior despite a sequence of contrary private signals.



^a $q > p$ ensures that an agent observing a signal contrary to the prior forms a posterior favoring the signal over the prior.

^bVerifiable by Bayes' rule:

$$\begin{aligned}
 \Pr(A > B|\alpha, \beta) &= \frac{\Pr(\alpha, \beta|A > B) \Pr(A > B)}{\Pr(\alpha, \beta|A > B) \Pr(A > B) + \Pr(\alpha, \beta|B > A) \Pr(B > A)} \\
 &= \frac{q(1-q) \cdot p}{q(1-q) \cdot p + (1-q)q \cdot (1-p)} \\
 &= p > \frac{1}{2}.
 \end{aligned}$$

Hence Agent 2's posterior coincides with the prior, as if she had received no informative evidence.

Remark.

- The assumption $q > p$ matters because it ensures

$$\begin{aligned}
 \Pr(B > A|\beta) &= \frac{\Pr(\beta|B > A) \Pr(B > A)}{\Pr(\beta|B > A) \Pr(B > A) + \Pr(\beta|A > B) \Pr(A > B)} \\
 &= \frac{q(1-p)}{q(1-p) + (1-q)p} > \frac{1}{2}.
 \end{aligned}$$

(Similarly, $\Pr(A > B|\alpha) > \frac{1}{2}$.) So a single private signal is strong enough to overturn the prior. Without this assumption agents would be “blinded” by the prior and ignore their signals from the start—no cascade would form because there was never any private content to suppress.

- Even with strong signals, the cascade is not perfect. Starting from Agent 2, the information aggregation breaks down: each agent's choice no longer reveals her private

signal.

- Agent 2 cares only about her own payoff. If she also cared about social welfare, she would choose according to her private signal (i.e., B if she received β) to reveal her information to subsequent agents. There is no individual incentive to do so—an information externality that leads to inefficiency.

Remark (Chapter Summary).

Three vignettes preview the recurring themes of the book. *Braess' paradox* (a routing game) shows that adding choices can hurt every player when externalities are present—the first hint that strategic interaction can invert the comparative statics of single-agent decision theory. *The sequential coin-toss game* shows that information can have negative value: making a private signal commonly observed reduces the well-informed player's equilibrium payoff. *Information cascades* show that rational individual updating can produce socially inefficient herding even when private signals would have aggregated correctly under full information. The recurring lesson—rationality plus interdependence yields surprising aggregate behavior—motivates the formal apparatus that begins in Chapter 1.

Part I

Foundations

Chapter 1

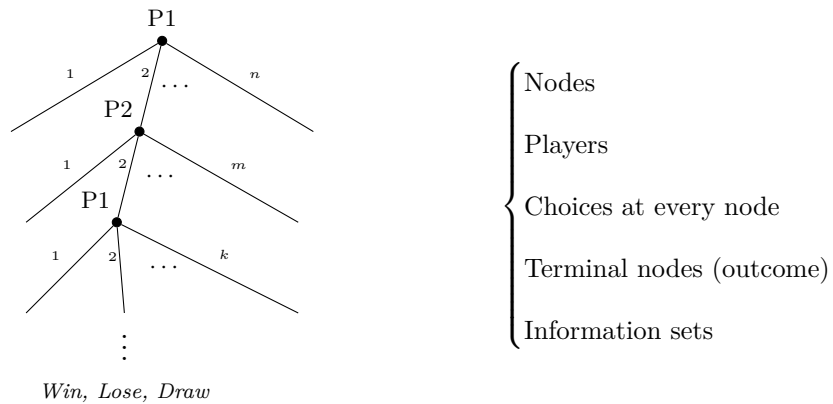
Representing a Game

Typically we have two ways to represent a game:

- Extensive form (game tree)
- Normal form (payoff matrix)

1.1 Game Tree


Example (Chess).



Definition 1.1: Game with Perfect / Imperfect Information

- A game with perfect information is a game where a player at node n knows all the moves that lead to n .
- A game with imperfect information is a game where a player at node n only knows some moves that lead to n .

Example (Matching Pennies).




1.2 Information Sets

All the nodes of a game tree are partitioned into information sets.

Rules for Information Sets

- Each node in an information set must belong to the same player.
- Each node in an information set must have the same choices.


Example.




The first tree is invalid because the two nodes in the information set have different choices. The second tree is invalid because the two nodes in the information set belong to different players.

Example (Coin Toss Revisited).

Cointoss Not Revealed to Player 1




Cointoss Revealed to Player 1




Assumption 1.2

Each path from the root to a terminal node intersects each information set at most once.

Example (Absent-Minded Driver).



Remark (“No Good Timing”).



In this game, there is no good notion of “timing” or “stages”. So the takeaway is that, the timing of players’ moves is not important. What matters is the information structure, i.e., who knows what.


1.3 Strategies

Before introducing the normal form, we need to formalize what a *strategy* is in an extensive-form game. The notion is rich: a strategy specifies what a player plans to do not just on the equilibrium path, but at every information set she might encounter.

Definition 1.3: Strategy

The strategy for player i , denoted by s_i , is a complete plan of which choices to make at every one of her information sets.

Example.



In this game, Player 1's strategies are LA , LB , RA , and RB , where LA means that Player 1 chooses L at the first information set belonging to her, and chooses A at the second. The strategies at all information sets will lead to a unique outcome.

Intuition: A player's strategy is a book with a page for each information set assigned to that player, specifying the choice there.

Definition 1.4: Mixed Strategy

A mixed (or, randomized) strategy for player i , denoted by σ_i , is a probability distribution over S_i , i.e., $\sigma_i \in \Delta(S_i)$.

Pure strategy is a special case of mixed strategy, assigning probability one to a specific action.

1.4 Normal Form

Definition 1.5: Game

A game G of n players consists of the set of strategies S_i and utility functions $u_i : S_1 \times S_2 \times \cdots \times S_n \rightarrow \mathbb{R}$ for $i = 1, 2, \dots, n$:

$$G = (S_1, S_2, \dots, S_n; u_1, u_2, \dots, u_n) \equiv (S_i, u_i)_{i=1}^n.$$

Definition 1.6: Belief

A belief of player i , denoted by $\mu_i \in \Delta(S_{-i})$, is a probability distribution over $S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n \equiv S_{-i}$.

Remark.

A belief is a probability distribution over S_{-i} , i.e., $\mu_i \in \Delta_{-i}$. But typically,

$$\Delta(S_i \times S_j) \neq \Delta(S_i) \times \Delta(S_j).$$

Example.

Consider a game of 3 players where Player 2 has strategies s_2 and s'_2 to choose from, and Player 3 has s_3 and s'_3 . From the Player 1's perspective, she can theoretically form the following two "beliefs":

	s_3	s'_3
s_2	0	1/4
s'_2	1/4	1/2

	s_3	s'_3
s_2	1/9	2/9
s'_2	4/9	4/9


Note that the first belief belongs to $\Delta(S_2 \times S_3)$ but not $\Delta(S_2) \times \Delta(S_3)$, while the second belief belongs to both $\Delta(S_2 \times S_3)$ and $\Delta(S_2) \times \Delta(S_3)$.

Consider the game where Player 1 has strategy space $S_1 = \{s_1, s'_1, s''_1\}$, and Player 2 has strategy space $S_2 = \{s_2, s'_2\}$. The normal form of this game is mapping from the strategy profiles to the payoffs of the players. In this example, it can be illustrated by a 3×2 matrix:


	s_2	s'_2
s_1	×	×
s'_1	×	×
s''_1	×	×

Remark (Relationship Between Game Trees and Normal Forms).

- Every tree leads to a unique normal form.
- However, a normal form can lead to multiple game trees that are "equivalent".



- Many game trees can lead to the same normal form.



Open Question: Do we lose some important information of the game by looking only at

the normal form?

1.5 Dominated Strategies and Best Responses

Rather than asking what is an optimal strategy, we approach it in reverse: what are the suboptimal strategies that a rational player would never choose, so that we can rule them out in the analysis of equilibrium?

Definition 1.7: Strictly Dominated Strategy

A pure strategy $s_i \in S_i$ is **strictly dominated** if there exists a (possibly mixed) strategy $\sigma_i \in \Delta(S_i)$ such that

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \quad \forall s_{-i} \in S_{-i},$$

where, by linearity, $u_i(\sigma_i, s_{-i}) = \sum_{s'_i \in S_i} \sigma_i(s'_i) u_i(s'_i, s_{-i})$.

Remark (Three Equivalent Forms).

The same concept can be stated in three equivalent ways. The third form, in terms of beliefs, is the one we will use most often when arguing about rationalizability.

- (i) *Pure-strategy form.* There exists $s'_i \in S_i$ with $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.
- (ii) *Mixed-strategy form.* The form stated in the definition above: there exists $\sigma_i \in \Delta(S_i)$ doing the dominating.
- (iii) *Belief form.* There exists $\sigma_i \in \Delta(S_i)$ such that $u_i(\sigma_i, \mu_i) > u_i(s_i, \mu_i)$ for every belief $\mu_i \in \Delta(S_{-i})$.

Form (i) is strictly weaker than (ii)—there are games (e.g. matching pennies with a third strategy) where a pure strategy is not dominated by any other pure strategy but is dominated by a mixture. Forms (ii) and (iii) are equivalent: linearity of u_i in μ_i converts a profile-by-profile inequality into a belief-by-belief one.

Definition 1.8: Best Response

A pure strategy $s_i \in S_i$ is a **best response** to the belief $\mu_i \in \Delta(S_{-i})$ if

$$u_i(s_i, \mu_i) \geq u_i(s'_i, \mu_i), \quad \forall s'_i \in S_i,$$

where $u_i(s_i, \mu_i) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \mu_i(s_{-i})$.

Definition 1.9: Never a Best Response (NBR)

A pure strategy s_i is **never a best response** if for every belief $\mu_i \in \Delta(S_{-i})$ there exists $s'_i \in S_i$ with $u_i(s'_i, \mu_i) > u_i(s_i, \mu_i)$.

Theorem 1.10: Dominance \iff Never a Best Response

In a finite game, s_i is strictly dominated if and only if s_i is never a best response.

The equivalence is the bridge between the eliminative “dominance” viewpoint and the constructive “rationalizable” viewpoint developed in Chapter 2. The proof uses a separating-hyperplane argument and is deferred to that chapter.

Definition 1.11: Weakly Dominated Strategy

A pure strategy s_i is **weakly dominated** if there exists $\sigma_i \in \Delta(S_i)$ such that

$$u_i(\sigma_i, s_{-i}) \geq u_i(s_i, s_{-i}) \quad \forall s_{-i}, \quad \text{with strict inequality for some } s_{-i}^* \in S_{-i}.$$

1.6 Iterated Elimination of Dominated Strategies

The previous theorem states that dominated strategy is a belief-free notion. This allows us to develop the concept of iterated elimination of dominated strategies (IESDS), which is a powerful tool to analyze games without having to worry about players’ beliefs.

The strategies that survive the IESDS are called *rationalizable strategies*, which are the strategies that can be rationalized by some beliefs about other players’ strategies. Rationalizability is a weaker notion than Nash equilibrium, as it does not require mutual consistency of beliefs. However, it is still a useful concept to understand the strategic behavior of players in a game.

Notably, IESDS’s outcome is independent of the order of elimination, which is a desirable property. This is because the set of dominated strategies is well-defined and does not depend on the sequence in which they are removed. As a result, we can confidently apply IESDS to simplify games and analyze the strategic interactions without worrying about the order of elimination. This is *not* the case with *IEWDS* (iterated elimination of weakly dominated strategies), where the order of elimination matters.

Example (IEWDS and Order of Elimination).

Consider the following 2-player game and IEWDS.

	<i>L</i>	<i>R</i>
<i>U</i>	1, 1	0, 0
<i>M</i>	1, 1	2, 1
<i>D</i>	0, 0	2, 1

- If Player 1 eliminates U and D , then Player 2 can eliminate nothing, and the process ends. The surviving strategies are M for Player 1, and L and R for Player 2.

- If Player 1 eliminates only U , then Player 2 can eliminate L , and the process ends. The surviving strategies are M and D for Player 1, and R for Player 2.
- If Player 1 eliminates only D , then Player 2 can eliminate R , and the process ends. The surviving strategies are U and M for Player 1, and L for Player 2.

From this example, we can clearly see the “outcome” of IEWDS depends on the order of elimination.

Remark (Chapter Summary).

This chapter introduced two complementary representations of a game and the language we will use throughout the book. The *extensive form* (game tree, with information sets capturing what each player knows when she moves) is the natural object for sequential-move games and games of imperfect information. The *normal form* (a tuple of strategy sets and payoff functions) is the natural object for simultaneous-move analysis and the home of equilibrium concepts; every extensive-form game admits a normal-form representation by enumerating strategies. A *strategy* sits between the two: it is defined on the extensive form (a complete plan for every information set) but its payoff implications live in the normal form. The chapter closes with the first nontrivial solution-concept idea—*dominated strategies and iterated elimination*—which already shows that rationality alone, without any equilibrium fixed-point reasoning, can pin down behavior in some games. Chapter 2 introduces the central equilibrium concept, Nash equilibrium, that the rest of the book builds on.

Chapter 2

Nash Equilibrium and Subgame Perfection

2.1 Nash Equilibrium and Nash's Theorem

Definition 2.1: Nash Equilibrium

A strategy profile $(\bar{\sigma}_1, \bar{\sigma}_2, \dots, \bar{\sigma}_n)$ (where $\bar{\sigma}_i \in \Delta(S_i)$) is a Nash equilibrium of the game $G = (S_i, u_i)_{i=1}^n$ if for any player i ,

$$u_i(\bar{\sigma}_i, \bar{\sigma}_{-i}) \geq u_i(s_i, \bar{\sigma}_{-i}), \quad \forall s_i \in S_i.$$

Remark.

Equivalently we can define Nash equilibrium by

$$u_i(\bar{\sigma}_i, \bar{\sigma}_{-i}) \geq u_i(\sigma_i, \bar{\sigma}_{-i}), \quad \forall \sigma_i \in \Delta(S_i).$$

But it suffices to only check pure strategies s_i because mixed strategies are just probability distributions over pure strategies, and the expected utility of a mixed strategy is a convex combination of the utilities of the pure strategies.

Theorem 2.2: Nash

Every finite game has a Nash equilibrium in mixed strategies.

Proof for Theorem

Here we apply the Kakutani's Fixed Point Theorem.

Theorem 2.3: Kakutani's Fixed Point Theorem

Suppose $F : X \rightrightarrows X$ (i.e., $F(x) \subseteq X$), and

- X is compact and convex;
- $F(x)$ is convex for all $x \in X$;
- F has a closed graph, i.e., the set $\{(x, y) | y \in F(x)\}$ is closed.

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

Let $X_i \equiv \Delta(S_i)$, and $X = X_1 \times X_2 \times \cdots \times X_n = \prod_{i=1}^n \Delta(S_i)$. Define the best response correspondence $B_i : \prod_{j \neq i} \Delta(S_j) \rightrightarrows \Delta(S_i)$ by

$$\sigma_i \in B_i(\sigma_{-i}) \iff u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}), \quad \forall \sigma'_i \in \Delta(S_i).$$

And we define $B : X \rightrightarrows X$ by

$$B(\sigma) = \begin{pmatrix} B_1 \\ B_2 \\ \vdots \\ B_n \end{pmatrix} (\sigma).$$

Three properties of B need to be verified. First, $B(\sigma)$ is convex for every σ : if $\sigma_i^a, \sigma_i^b \in B_i(\sigma_{-i})$, both attain the maximum payoff, and any convex combination $\lambda \sigma_i^a + (1 - \lambda) \sigma_i^b$ does too by linearity of u_i in player i 's mixture. Second, X is compact and convex as the product of simplices. Third, the graph of B is closed: if $\sigma^k \rightarrow \bar{\sigma}$ and $\sigma^k \in B(\sigma^k)$, continuity of u_i implies $\bar{\sigma} \in B(\bar{\sigma})$. Kakutani's theorem then delivers a fixed point $\bar{\sigma} \in B(\bar{\sigma})$, which is a Nash equilibrium. ■

Remark.

- “Mixed strategies” is important for Nash's Theorem. Finite games may not have NE in pure strategies.
- If $\bar{\sigma}$ is a NE, and $\bar{\sigma}_i(s_i) > 0$, then s_i must be a best response to $\bar{\sigma}_{-i}$, and thus $u_i(s_i, \bar{\sigma}_{-i}) = u_i(\bar{\sigma}_i, \bar{\sigma}_{-i})$. In other words, if a pure strategy is played with positive probability in a mixed strategy NE (i.e., in the *support* of the mixed strategy $\bar{\sigma}$), then it must yield the same expected payoff as the mixed strategy itself.
- Then why bother mixing? Because each player mixes to make other players mix. If a player chooses a pure strategy, then other players will have no incentive to mix, and thus the first player will have no incentive to mix either. Therefore, mixing can be a strategic move to induce other players to mix, which can lead to a more favorable outcome for the player.

Remark (Alternative Proof via Brouwer (Geanakoplos, 2003)).

Nash's theorem can also be proved using the simpler **Brouwer Fixed Point Theorem** (which only requires a continuous *function*, not a correspondence). The trick is to replace the best-response correspondence $B(\sigma)$, which is in general multi-valued, with a continuous *single-valued* map whose fixed points coincide with Nash equilibria.

For each player i , define $\phi_i : \Delta \rightarrow \Delta_i$ by

$$\phi_i(\sigma) = \arg \max_{\sigma'_i \in \Delta_i} \left\{ u_i(\sigma'_i, \sigma_{-i}) - \|\sigma'_i - \sigma_i\|_1^2 \right\}.$$

The maximand is the sum of an affine (linear in σ'_i) term and a strictly concave term, so it is strictly concave; hence $\phi_i(\sigma)$ is a singleton and, by the Theorem of the Maximum, ϕ_i is continuous in σ . The product map $\phi : \Delta \rightarrow \Delta$ is therefore continuous on the compact convex domain Δ , and Brouwer delivers a fixed point $\bar{\sigma} = \phi(\bar{\sigma})$.

To see that $\bar{\sigma}$ is a Nash equilibrium, suppose for contradiction that some player i has a strictly profitable deviation σ'_i with $u_i(\sigma'_i, \bar{\sigma}_{-i}) - u_i(\bar{\sigma}) = D > 0$. By linearity of u_i in σ_i , the convex combination $\varepsilon\sigma'_i + (1 - \varepsilon)\bar{\sigma}_i$ raises i 's payoff by exactly εD , while the penalty $\|\varepsilon\sigma'_i + (1 - \varepsilon)\bar{\sigma}_i - \bar{\sigma}_i\|_1^2 = \varepsilon^2\|\sigma'_i - \bar{\sigma}_i\|_1^2$ is of order ε^2 . For small ε , the linear gain dominates the quadratic loss, so the perturbed strategy strictly beats $\bar{\sigma}_i$ in the maximand defining ϕ_i —contradicting $\phi_i(\bar{\sigma}) = \bar{\sigma}_i$.

The Kakutani-based proof is the textbook standard, but the Brouwer route is conceptually cleaner: it makes explicit that the gap between “best-response correspondence has a fixed point” (which can be multi-valued and hence requires Kakutani) and “Nash equilibrium exists” (which only needs the existence of *some* consistent profile) can be bridged by a tiny regularization. The same trick reappears in computational game theory, where strictly concave perturbations are used to make best responses unique and approximate equilibria via gradient methods.

2.2 Iterated Dominance and Rationalizability

Chapter 1 introduced strictly dominated strategies and the iterated elimination of dominated strategies (IESDS), and informally identified the surviving set with *rationalizable* strategies. We now make that identification precise—the equivalence is a non-trivial theorem and its proof clarifies what each side of the definition is doing.

Definition 2.4: Iteratively Undominated Set (IUD)

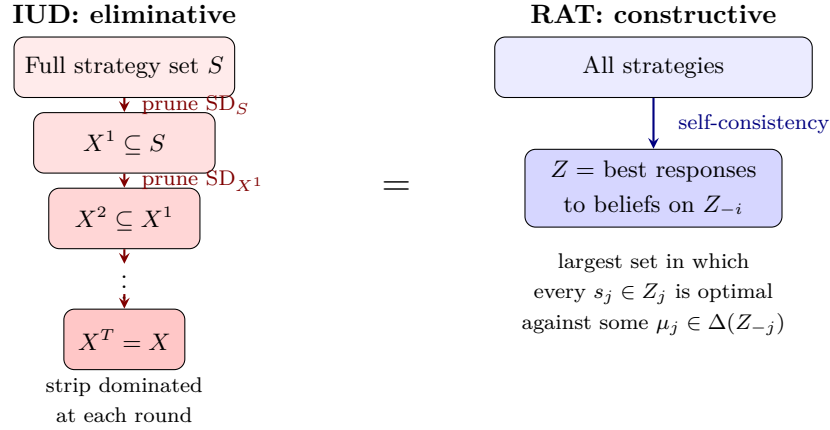
A product set $X = X_1 \times \cdots \times X_n \subseteq S$ is **iteratively undominated** if there exists a sequence $\{X^t\}_{t=0}^T$ of product sets such that

- (i) $X^0 = S$ and $X^T = X$;
- (ii) $X^{t+1} \subseteq X^t$ for every t ;
- (iii) every $s_j \in X_j^t \setminus X_j^{t+1}$ is strictly dominated in the restricted game $(X_j^t, u_j)_{j=1}^n$;
- (iv) no $s_j \in X_j^T$ is strictly dominated in the restricted game $(X_j^T, u_j)_{j=1}^n$.

Definition 2.5: Rationalizable Set (RAT)

A product set $Z = Z_1 \times \dots \times Z_n \subseteq S$ is **rationalizable** if (i) every $s_j \in Z_j$ is a best response to some belief $\mu_j \in \Delta(Z_{-j})$ whose support lies in Z_{-j} , and (ii) Z is the maximal such set: any other set Z' satisfying (i) is contained in Z .

The two definitions describe rationality from opposite directions. IUD is *eliminative*: start with everything and prune dominated strategies, justifying each removal by reference to the previous round. RAT is *constructive*: insist that every surviving strategy be optimal under *some* belief over surviving opponent strategies, and take the largest such set. The next theorem shows they pick out the same object.



Theorem 2.6: IUD = RAT

Let X be the iteratively undominated set and Z the rationalizable set of a finite game G . Then $X = Z$.

Proof for Theorem

We use the lemma from Chapter 1 (Theorem 1.5) that $s_i \in S_i$ is strictly dominated (possibly by a mixed strategy) if and only if it is never a best response (NBR) to any belief in $\Delta(S_{-i})$. The proof proceeds by mutual inclusion.

$Z \subseteq X$. We show by induction that $Z_j \subseteq X_j^t$ for every t . The base case $t = 0$ is immediate since $X_j^0 = S_j$. For the inductive step, suppose $Z_j \subseteq X_j^t$. Take any $s_j \in Z_j$. By rationalizability, s_j is a best response to some belief μ_j supported on $Z_{-j} \subseteq X_{-j}^t$. By the $SD \iff NBR$ lemma applied to the restricted game (X^t, u) , s_j is therefore not strictly dominated in (X^t, u) . Hence $s_j \in X_j^{t+1}$, completing the induction. In the limit, $Z \subseteq X^T = X$.

$X \subseteq Z$. By the termination condition (iv), no $s_j \in X_j^T = X_j$ is strictly dominated in (X, u) . By the $SD \iff NBR$ lemma, every $s_j \in X_j$ is a best response from X_j to some belief μ_j on X_{-j} . To upgrade “best response from X_j ” to “best response from S_j ” (the requirement in the definition of rationalizability), suppose for contradiction that some $s_j \in X_j$ is best-responded-to from X_j but *not* from S_j : there exists $s'_j \in S_j \setminus X_j$ that strictly improves on s_j against μ_j . Since s'_j is not in X_j , it was eliminated at some

stage t of the iterated procedure—meaning s'_j was strictly dominated in (X^t, u) . But this contradicts that s'_j is a strict best response to $\mu_j \in \Delta(X^t_{-j})$ (by the SD \iff NBR lemma in the restricted game). Hence X satisfies condition (i) of rationalizability. Since Z is by definition the maximal such set, $X \subseteq Z$. ■

Remark (Why the Equivalence Is Substantive).

IUD reaches the rationalizable set “from above,” RAT defines it “intrinsically.” Their equality is the formal statement that *rationality and common knowledge of rationality* imply nothing more than what iterated elimination of strictly dominated strategies extracts. In particular: there is no mileage to be gained by allowing players to entertain richer hierarchies of beliefs beyond what IESDS already captures. Conversely, IESDS is not a brute-force algorithm divorced from epistemic foundations—each round of elimination corresponds exactly to one more level of mutual knowledge of rationality. The order-of-elimination invariance for strict dominance, observed informally in Chapter 1, is now a corollary: every elimination order converges to the unique RAT set.

2.3 Potential Games

One of the central questions in game theory is whether a game has a *pure strategy Nash equilibrium* (PSNE). Some games, like matching penny, have no PSNE. However, certain classes of games are guaranteed to have PSNE. Two important classes are:

- Potential games
- Supermodular games

In this section we will focus on potential games, which are a broad class of games that include many interesting examples, such as congestion games, coordination games, and certain types of auctions. Potential games have a special structure that allows us to analyze them using a potential function, which captures the incentives of all players in a single function.

2.3.1 Definition

Definition 2.7: Potential Game

A finite game $G = (S_i, u_i)_{i=1}^n$ is called a *potential game* if there exists a function $P : S \rightarrow \mathbb{R}$ (where $S = S_1 \times S_2 \times \cdots \times S_n$) such that for every player i and every pair of strategies $s_i, s'_i \in S_i$:

$$u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i}) = P(s_i, s_{-i}) - P(s'_i, s_{-i}), \quad s_{-i} \in S_{-i}.$$

The function P is called the *potential function*.

Example (Congestion Game).

Consider a road network where n drivers choose routes. Let s_i be the route chosen by driver i , $n_r(s)$ the number of drivers on route r under profile s , and $c(r, k)$ the congestion cost incurred by each user of route r when k drivers use it. Driver i 's payoff is

$$u_i(s_i, s_{-i}) = -c(s_i, n_{s_i}(s)).$$

Define the potential function

$$P(s) = -\sum_r \sum_{k=1}^{n_r(s)} c(r, k).$$

When driver i switches from route r to route r' , the change in P is exactly $c(r', n_{r'}(s) + 1) - c(r, n_r(s))$, which equals the change in driver i 's own payoff (up to sign). Hence P is a potential function and the game admits a PSNE—the route assignment that minimizes total congestion cost.

Remark.

- Potential games can be thought of as games where all players have identical payoffs.
- Potential games arise naturally when:
 - Players care about a common objective (like aggregate welfare)
 - Individual incentives are aligned with some global measure
 - Strategic interactions are “symmetric” in the sense of cross-partial derivatives
- The potential function P captures the payoff differences for individual players when they deviate from one strategy to another. The key insight is that a player's incentive to deviate depends only on how the potential function changes, not on the baseline payoff levels.
- In the continuous case with smooth payoff functions, a game is a potential game if and only if:

$$\frac{\partial P}{\partial s_i} = \frac{\partial u_i}{\partial s_i} \quad \text{for all players } i$$

This leads to the following necessary condition, for all pairs of players i and j :

$$\frac{\partial^2 u_i}{\partial s_j \partial s_k} = \frac{\partial^2 u_j}{\partial s_k \partial s_j} \quad \text{whenever these derivatives exist.}$$

2.3.2 Existence of PSNE in Potential Games

Theorem 2.8: Existence of PSNE in Potential Games

Every potential game has at least one pure strategy Nash equilibrium.

Proof for Theorem

A PSNE of game G is a maximizer of the potential function P . Since S is finite (in the discrete case), P attains its maximum. At such a maximum point s^* , if player i deviates to any other strategy $s_i \neq s_i^*$, then:

$$u_i(s_i^*, s_{-i}^*) - u_i(s_i, s_{-i}^*) = P(s_i^*, s_{-i}^*) - P(s_i, s_{-i}^*) \geq 0$$

Thus no player wants to deviate, so s^* is a PSNE. ■

Remark.

Different information structures (simultaneous move vs. sequential move) can lead to different equilibria in general games. However, in potential games, the set of pure strategy Nash equilibria is often robust across different information structures. This is because the potential function captures the incentives of all players in a way that is independent of the timing of moves. As a result, the same strategy profiles that maximize the potential function will be equilibria regardless of whether players move simultaneously or sequentially.

2.4 Supermodular Games

While potential games ensure the existence of PSNE through a global potential function, supermodular games approach the problem differently: they rely on the structure of strategic complementarities between players' actions.

The key insight is *strategic complementarity*: if one player's action increases, it becomes more attractive for other players to increase their actions as well. This creates a natural coordination mechanism that leads to equilibrium. Unlike potential games, which require a global objective function, supermodular games only require local complementarity conditions on payoffs.

Examples of strategic complementarity abound:

- In a price competition setting, if a competitor raises prices, your prices become more attractive to consumers, so you want to raise your price too.
- In technology adoption, if more people adopt a technology, the benefits of adoption increase for you (network effects).
- In public goods games, if others contribute more, your optimal contribution may also increase.

2.4.1 Definitions

We need lattice-theoretic concepts to formalize strategic complementarity.

Definition 2.9: Meet & Join

Let x and y be two vectors of \mathbb{R}^n . The *meet* of x and y , denoted by $x \wedge y$, is the coordinate-wise minimum:

$$x \wedge y = \begin{pmatrix} \min(x_1, y_1) \\ \min(x_2, y_2) \\ \vdots \\ \min(x_n, y_n) \end{pmatrix}$$

The *join* of x and y , denoted by $x \vee y$, is the coordinate-wise maximum:

$$x \vee y = \begin{pmatrix} \max(x_1, y_1) \\ \max(x_2, y_2) \\ \vdots \\ \max(x_n, y_n) \end{pmatrix}$$

Remark.

Meet and join decompose a pair of strategy profiles into their “lower bound” and “upper bound” in each dimension. Later we will see, if a payoff function is supermodular, then being at both extremes (the largest and smallest actions) together is at least as good as being at the two middle ground positions.

Definition 2.10: Supermodular Function

A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *supermodular* if for all $x, y \in \mathbb{R}^n$:

$$f(x \wedge y) + f(x \vee y) \geq f(x) + f(y)$$

This is called the *supermodularity inequality*.

Intuitively, the supermodularity inequality says: having some coordinates large and others small is better than having all coordinates at intermediate levels.

In a game setting, we say player i 's payoff is supermodular if:

$$u_i(s) + u_i(s') \leq u_i(s \wedge s') + u_i(s \vee s') \quad \forall s, s' \in S$$

This means: player i 's payoff is supermodular in their own strategy relative to opponents' strategies.

Remark.

When f is twice continuously differentiable, supermodularity is equivalent to:

$$\frac{\partial^2 f}{\partial x_i \partial x_j} \geq 0 \quad \text{for all } i \neq j.$$

This means: if you increase x_i , the marginal benefit of increasing x_j also increases (*strategic complementarity*).

2.4.2 Existence of PSNE in Supermodular Games

The next two lemmas establish that best response correspondences have monotonicity properties under supermodularity, which enables us to use fixed point theorems.

Lemma 2.11: Monotonicity of Best Responses

Suppose $u : [0, 1]^p \times [0, 1]^q \rightarrow \mathbb{R}$ is supermodular. Suppose $y', y'' \in [0, 1]^q$ with $y'' \geq y'$. Let $x' \in [0, 1]^p$ be a maximizer of $u(\cdot, y')$, and $x'' \in [0, 1]^p$ be a maximizer of $u(\cdot, y'')$. Then $x' \wedge x''$ is a maximizer of $u(\cdot, y')$, and $x' \vee x''$ is a maximizer of $u(\cdot, y'')$.

Proof for Lemma

By construction, since x' maximizes $u(\cdot, y')$ and x'' maximizes $u(\cdot, y'')$:

$$u(x', y') \geq u(x' \wedge x'', y')$$

Also note that $u(x' \wedge x'', y') = u(x' \wedge x'', y' \wedge y'')$ since $y' \wedge y'' = y'$ (as $y' \leq y''$).

Similarly:

$$u(x'', y'') \geq u(x' \vee x'', y'') = u(x' \vee x'', y' \vee y'')$$

By supermodularity:

$$u(x', y') + u(x'', y'') \geq u(x' \wedge x'', y' \wedge y'') + u(x' \vee x'', y' \vee y'')$$

Suppose $u(x', y') > u(x' \wedge x'', y' \wedge y'')$. Then:

$$u(x', y') + u(x'', y'') > u(x' \wedge x'', y' \wedge y'') + u(x' \vee x'', y' \vee y'')$$

which violates supermodularity.

Therefore, we must have:

$$u(x', y') = u(x' \wedge x'', y' \wedge y'')$$

$$u(x'', y'') = u(x' \vee x'', y' \vee y'')$$

This means both $x' \wedge x''$ and $x' \vee x''$ are maximizers at the respective points. ■

Lemma 2.12: Monotonicity of Largest Maximizers

Suppose $u : [0, 1]^p \times [0, 1]^q \rightarrow \mathbb{R}$ is supermodular. Define:

$$h(y) = \max\{x : x \text{ maximizes } u(\cdot, y)\}$$

(the largest maximizer). Then h is non-decreasing: if $y'' \geq y'$, then $h(y'') \geq h(y')$.

Proof for Lemma

Suppose $y'' \geq y'$. Let $x' = h(y')$ and $x'' = h(y'')$ be the largest maximizers at y' and y'' respectively. By the previous lemma, $x' \wedge x''$ maximizes $u(\cdot, y')$ and $x' \vee x''$ maximizes $u(\cdot, y'')$.

Since $x' \vee x''$ is a maximizer at y'' and x'' is the largest maximizer at y'' , we have $x'' \geq x' \vee x''$. By definition, $x' \vee x'' \geq x'$. So $x'' \geq x'$. Therefore, $h(y'') = x'' \geq x' = h(y')$. ■

Theorem 2.13: Existence of PSNE in Supermodular Games

Every finite supermodular game has at least one pure strategy Nash equilibrium.

Proof for Theorem

This is a direct result of Tarski's Fixed Point Theorem.

Theorem 2.14: Tarski's Fixed Point Theorem

If $f : L \rightarrow L$ is a non-decreasing function on a complete lattice L , then f has a fixed point.

Why Tarski holds (sketch). *Let L be a complete lattice with top \top and bottom \perp . The set $A = \{x \in L : x \leq f(x)\}$ contains \perp (since $\perp \leq f(\perp)$) and is therefore non-empty. Let $\bar{x} = \sup A$, which exists by completeness. Monotonicity gives $f(x) \leq f(\bar{x})$ for every $x \in A$, and combining with $x \leq f(x)$ yields $x \leq f(\bar{x})$, so $f(\bar{x})$ is an upper bound of A and hence $\bar{x} \leq f(\bar{x})$. Applying f once more and using monotonicity, $f(\bar{x}) \leq f(f(\bar{x}))$, so $f(\bar{x}) \in A$ and $f(\bar{x}) \leq \bar{x}$. Together: $f(\bar{x}) = \bar{x}$. The argument is constructive on finite lattices: iterate f on \perp until the sequence stabilizes.*

For each player i , define $b_i(s_{-i})$ as the largest maximizer of $u_i(\cdot, s_{-i})$. By the lemmas above, b_i is non-decreasing in s_{-i} .

Define the best response mapping $b : S \rightarrow S$ by:

$$b(s) = (b_1(s_{-1}), b_2(s_{-2}), \dots, b_n(s_{-n}))$$

Since b is component-wise non-decreasing and S is a finite lattice, by *Tarski's Fixed Point Theorem*, b has a fixed point \bar{s} , i.e., $b(\bar{s}) = \bar{s}$. This fixed point is a Nash equilibrium. ■

Example (Bertrand Duopoly with Differentiated Products).

Consider two firms competing in prices. Let $D_i(p_1, p_2)$ be firm i 's demand, which is decreasing in its own price p_i and increasing in the competitor's price p_j . With constant marginal cost c , firm i 's profit is:

$$\pi_i(p_1, p_2) = (p_i - c)D_i(p_1, p_2)$$

The key observation is:

$$\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} = \frac{\partial}{\partial p_j} \left[\frac{\partial \pi_i}{\partial p_i} \right] > 0$$

This is because when the competitor's price increases, the marginal profit from raising your own price increases (since demand becomes more favorable). Thus, firm i 's profit is supermodular in (p_1, p_2) , and the duopoly game has a PSNE.

2.5 Subgame Perfect Nash Equilibrium

Nash equilibrium constrains only on-path play. In extensive-form games—where players move sequentially and can condition on what they have observed—this is too weak: many Nash equilibria are sustained by threats that the threatener would never actually carry out. Subgame perfection refines NE by demanding optimality in every *subgame*, ruling out such non-credible threats.

2.5.1 Definitions and Existence

Definition 2.15: Subgame

Given an extensive form T , a subgame T' is some node $r \in T$ and all its succeeding nodes such that if an information set I intersects T' , then all nodes in I are also in T' .


Definition 2.16: Subgame Perfect Nash Equilibrium

σ is a *subgame perfect equilibrium* if

- σ is a Nash equilibrium.
- σ induces a Nash equilibrium in every subgame $T' \subseteq T$.

Example (Entry Deterrence Game).

Consider the following game:



From its normal form (payoff matrix) below, we know that there are two Nash equilibria: (OUT, FIGHT) and (IN, ACCOM).

	FIGHT	ACCOM
OUT	0, 5	0, 5
IN	-2, -2	2, 2

However, only (IN, ACCOM) is a subgame perfect equilibrium, because if Player 1 chooses IN, then Player 2's best response is to choose ACCOM (so to choose FIGHT is not a credible threat). Therefore, the unique subgame perfect equilibrium is (IN, ACCOM).

Suppose now the game is repeatedly played for $T = 20$ times. At $t = 1$, Player 1 chooses IN or OUT. If Player 1 chooses OUT, the game ends. If Player 1 chooses IN, then in any period $t \geq 1$, the Player 1 can choose to stay (IN) or exit (OUT). If OUT is ever chosen by Player 1, then OUT for all the periods onwards.

Clearly, using backward induction, we know that the unique subgame perfect equilibrium is for Player 1 to choose IN and Player 2 to choose ACCOM in all periods.

However, if we add one additional constraint: Player 1 can only fight for at most 15 periods, then the analysis changes.

Claim

The unique SPE with the constraint is that

- Player 1 never enters (IN), and
- Player 2 chooses FIGHT if Player 1 chooses IN.

Proof for Claim.

Suppose that Player 1 has entered and been fought for the previous 14 periods. In period $t = 15$, if Player 1 chooses IN, Player 2 compares the payoffs of FIGHT and ACCOM:

- If Player 2 chooses FIGHT: Player 1's capacity is exhausted, forcing them to choose OUT from $t = 16$ to $t = 20$. Player 2's total remaining payoff is $-2(\text{at } t = 15) + 5 \times 4(\text{monopoly profit}) = 18$.
- If Player 2 chooses ACCOM: Player 1 stays in the market. Player 2's total remaining payoff is $2(\text{at } t = 15) + 5 \times 2(\text{duopoly profit}) = 12$.

Since $18 > 12$, FIGHT becomes a *credible threat* at $t = 15$. Anticipating this, Player 1 will choose OUT at $t = 15$ to avoid the -2 payoff.

Consider period $t = 14$. Player 1 anticipates that choosing IN will result in Player 2 choosing FIGHT (to trigger the exit at $t = 15$ and secure future monopoly profits). Choosing IN at $t = 14$ followed by exit at $t = 15$ yields a lower payoff than choosing OUT immediately at $t = 14$.

Following this logic backward to $t = 1$, consequently, the only rational choice for Player 1 is to never enter (OUT), and for Player 2 to threaten to FIGHT at any entry node. ■

Theorem 2.17

Every finite game in extensive form has at least one subgame perfect equilibrium.

Proof for Theorem

The proof is to find all subgames that do not have any proper subgames. In each such subgame, find a Nash equilibrium of the subgame, and replace each node by expected payoffs from any Nash equilibrium of the subgame. ■

Remark.

The notion of subgame perfect equilibrium is to generalize backward induction to games with imperfect information.

2.5.2 The Backward Induction Algorithm

The proof of the existence theorem is constructive: it gives an explicit *algorithm* for computing an SPE in any finite extensive-form game. The procedure—**backward induction**—is the workhorse of finite-horizon analysis.

Backward Induction Algorithm

Let T be a finite extensive-form game.

1. **Identify minimal subgames.** A subgame is *minimal* if it contains no proper subgame. Equivalently, every information set inside it is a singleton in which

only one player moves before reaching a terminal node.

2. **Solve each minimal subgame.** Find a Nash equilibrium of the minimal subgame (in finite games this exists, by Nash’s theorem applied to the subgame’s normal form).
3. **Substitute continuation payoffs.** Replace every minimal subgame by a single terminal node carrying the equilibrium payoff vector. This yields a smaller extensive-form game T' with one less “layer.”
4. **Recurse.** Apply Steps 1–3 to T' . Terminate when only the root remains; the strategies recorded along the way constitute an SPE.

Remark (Why Backward Induction Produces an SPE).


By construction, the strategies prescribed in each minimal subgame are a Nash equilibrium of that subgame. After substitution, the next iteration treats those subgames as terminal nodes with fixed payoffs, so any Nash equilibrium of the reduced game T' remains a Nash equilibrium when expanded back. Iterating, the prescribed strategies form a Nash equilibrium in every subgame—which is exactly the SPE condition. The argument is the game-theoretic analogue of **dynamic programming**: a globally optimal plan is built by stitching together locally optimal continuations.

Remark (Verification via the One-Deviation Property).

A practical corollary: to check that a candidate strategy profile σ^* is an SPE, it suffices to verify, at every history (or every information set in the perfect-information case), that no player can profitably deviate by changing her action at that single node and reverting to σ^* thereafter. This is the **one-shot deviation principle**, a finite-horizon analogue of the result we will state and prove for infinitely repeated games in the chapter on dynamic games. The principle reduces SPE verification from a global check across all alternative strategies to a local, period-by-period inequality at each node—precisely what the entry-deterrence example above (with Player 1’s 15-period capacity constraint) implicitly relied on.

The backward-induction algorithm settles existence and provides a constructive solution method, but it does *not* guarantee uniqueness: if a minimal subgame has multiple Nash equilibria, choosing different ones leads to different SPEs. This multiplicity is what powers the finite-horizon Folk Theorem (in the chapter on dynamic games): when a stage game has several payoff-distinct NE, the choice of which NE to play in late periods can be used as a reward or punishment to sustain non-NE play in early periods.


Example (Centipede Game).



The unique subgame perfect equilibrium is for each player to exit immediately at their first decision node, yielding payoff $(1, 0)$. However, the backward induction outcome $(1, 0)$ is Pareto dominated by the outcome $(64, 64)$ at the end. Both players would be better off if the game continued to the end, but the equilibrium logic prevents this.

2.5.3 Trembling-Hand Perfect Equilibrium

Example (Selten's Horse).



There are two pure strategy Nash equilibria: (d, A, l) and (a, A, r) . Since this game has no proper subgames, both equilibria are subgame perfect.

However, the SPE (d, A, l) appears problematic: if Player 1 accidentally chooses a instead of d , Player 2 would prefer to choose D rather than A . This suggests the equilibrium is not robust to small mistakes.

Recall that σ is a Nash equilibrium if and only if for each player i and any pure strategy s'_i : $u_i(\sigma_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i})$, and for all the pure strategies that are in the support of σ_i , the expected payoff is the same: $u_i(s_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i})$ for all s_i such that $\sigma_i(s_i) > 0$. Hence, if for some pure strategy s'_i , we have $u_i(\sigma_i, \sigma_{-i}) > u_i(s'_i, \sigma_{-i})$, then $\sigma_i(s'_i) = 0$.

To formalize the intuition from the previous example of Selten's Horse, suppose all players' choices are subject to small trembles (mistakes).

Definition 2.18: ε -Equilibrium

A mixed strategy profile σ^ε is an ε -equilibrium (for small $\varepsilon > 0$) if:

- Every pure strategy has positive probability: $\sigma_i^\varepsilon(s_i) > 0$ for all i and s_i .
- Strategies that are not a best response are played with low probability: if $u_i(s'_i, \sigma_{-i}^\varepsilon) > u_i(s_i, \sigma_{-i}^\varepsilon)$, then $\sigma_i^\varepsilon(s'_i) < \varepsilon$.

Definition 2.19: Trembling Hand Perfect Equilibrium

A strategy profile σ is a *trembling hand perfect equilibrium* (THPE) if it is the limit of a sequence of ε -equilibria σ^ε as $\varepsilon \rightarrow 0$.

Remark (Chapter Summary).

This chapter built the central solution concept of the book. Nash's theorem (Theorem 2.1) guarantees that every finite game has at least one mixed-strategy equilibrium; the proof, via Kakutani or—more directly—Brouwer, identifies an equilibrium with a fixed point of a best-response correspondence. We then mapped out two families of refinements and existence results. *Iterated dominance and rationalizability* (IUD = RAT) describe what rationality plus common knowledge of rationality alone can deliver, before equilibrium fixed-point reasoning enters. *Potential and supermodular games* carve out structural classes in which pure-strategy equilibria are guaranteed and have nice comparative-statics properties (Topkis-Milgrom-Roberts). For dynamic games, *subgame perfection* (Definition 2.5.1) sharpens Nash equilibrium by requiring optimality at every history; *trembling-hand perfection* sharpens it again by ruling out equilibria sustained by zero-probability mistakes. Each refinement narrows the prediction at the cost of stronger assumptions about behavior off the equilibrium path—a tradeoff that recurs in every later chapter.

Part II

Bargaining

Chapter 3

Two-Person Bargaining

3.1 Nash Bargaining Solution

Definition 3.1: Bargaining Problem

A **(two-person) bargaining problem** is a pair (U, d) , where $U \subseteq \mathbb{R}^2$ is the *feasible utility set*—a compact, convex set of expected utility profiles attainable by the two parties—and $d \in U$ is the *disagreement point*, the outcome that arises if no agreement is reached. We assume non-degeneracy: there exists some $u \in U$ with $u \gg d$ (i.e., $u_i > d_i$ for both i).

Definition 3.2: Bargaining Solution

A **bargaining solution** is a function F that assigns to every bargaining problem (U, d) a point $F(U, d) \in U$ representing the agreed-upon utility profile.

We now ask: what are the natural axioms a bargaining solution should satisfy, and is there a unique one that does?

Axiom: Scale Invariance

Consider the transformation $(U, d) \mapsto (\bar{U}, \bar{d})$ such that

$$\bar{U} = \left\{ (\bar{u}_1, \bar{u}_2) \mid \bar{u}_i = \alpha_i u_i + \beta_i \text{ for some } \alpha_i > 0, \quad (u_1, u_2) \in U \right\},$$
$$\bar{d}_i = \alpha_i d_i + \beta_i.$$

Then,

$$F_i(\bar{U}, \bar{d}) = \alpha_i F_i(U, d) + \beta_i.$$

Axiom: Efficiency

$\nexists u \in U$ such that $u \gg F(U, d)$.

Axiom: Symmetry

If U is symmetric (i.e., $(u_1, u_2) \in U \iff (u_2, u_1) \in U$) and $d_1 = d_2$, then

$$F_1(U, d) = F_2(U, d).$$

Axiom: Independence of Irrelevant Alternatives (IIA)

Consider (U, d) and (\bar{U}, d) such that $U \subseteq \bar{U}$. If $F(\bar{U}, d) \in U$, then

$$F(\bar{U}, d) = F(U, d).$$

Theorem 3.3


There is one and only one F that satisfies the above requirements. F is given by:

$$F^N(U, d) = \arg \max_{u \in U, u \geq d} (u_1 - d_1)(u_2 - d_2).$$

Proof for Theorem

Firstly, it is easy to check that F^N is well-defined in the sense that F^N satisfies the above requirements.

Let $(u_1^*, u_2^*) = F^N(U, d)$. We can always rescale the utilities such that $(u_1^*, u_2^*) \mapsto (1, 1)$ and $(d_1, d_2) \mapsto (0, 0)$, without loss of generality. Denote the new space as (\bar{U}, \bar{d}) .

**Claim**

If we can show that $F = F^N$ for any F in the scaled space (\bar{U}, \bar{d}) , then we must have $F = F^N$ for any (U, d) .

Proof for Claim.


This is true by scale invariance. ■

Claim


$$F^N(\bar{U}, d) = (1, 1).$$

Proof for Claim.

Define the set $H = \{(u_1, u_2) \mid u_1 + u_2 \leq 2\}$. Note that the boundary of H is the tangent line to the hyperbola $u_1 u_2 = 1$ through $(1, 1)$ with slope -1 . Since $(1, 1) \in \bar{U}$ and $(1, 1) \in H$, if we can show that $\bar{U} \subseteq H$, then we must have $(1, 1)$ as the solution.



To show that $\bar{U} \subseteq H$, suppose for contradiction that \bar{U} contains some point (u_1, u_2) such that $u_1 + u_2 > 2$. The line segment connecting $(1, 1)$ and (u_1, u_2) must intersect the hyperbola $u_1 u_2 = 1$ at some point $(\tilde{u}_1, \tilde{u}_2)$. So the line segment joining $(\tilde{u}_1, \tilde{u}_2)$ and $(1, 1)$ must be above the hyperbola $u_1 u_2 = 1$. Since \bar{U} is convex, we can find a convex combination of (u_1, u_2) and $(1, 1)$ that falls within the line segment joining $(\tilde{u}_1, \tilde{u}_2)$ and $(1, 1)$, denoted by (\hat{u}_1, \hat{u}_2) such that $(\hat{u}_1, \hat{u}_2) \in \bar{U}$ and $\hat{u}_1 \hat{u}_2 > 1$. This contradicts the assumption that $F^N(\bar{U}, d) = (1, 1)$.




Claim

For any F , the solution must be $(1, 1)$.

Proof for Claim.

We can always include the \bar{U} using a symmetric right triangle whose hypotenuse is on the boundary of H and whose right angle sits on the line of $u_1 = u_2$. Denote the new triangle by $\bar{\bar{U}}$.



By efficiency of F , the solution must lie on the boundary of $\bar{\bar{U}}$. By symmetry of F , the solution must be $(1, 1)$. Additionally, note that $\bar{U} \subseteq \bar{\bar{U}}$, and $F(\bar{\bar{U}}, \bar{d}) \in \bar{U}$. By IIA, we must have $F(\bar{\bar{U}}, \bar{d}) = F(\bar{U}, \bar{d})$. Therefore, $F(\bar{\bar{U}}, \bar{d}) = (1, 1)$. ■

From all the results above, we conclude that $F = F^N$ for any (U, d) . ■

Remark.

This is actually a *normative* result in the sense that we require the bargaining solution to satisfy certain axioms. The Nash bargaining solution is the unique solution that satisfies these axioms. However, it is not necessarily a positive result in the sense that we do not have a game-theoretic model that can generate the Nash bargaining solution as an equilibrium outcome.

3.2 Dividing a Dollar

Let $X = [0, 1]$. Suppose utility functions for the two players are $u_1 : X \rightarrow \mathbb{R}$ and $u_2 : X \rightarrow \mathbb{R}$, which are continuous, strictly increasing, and strictly concave. The problem is given by

$$\max_{x \in [0, 1]} u_1(x)u_2(1 - x).$$

The FOC is:

$$\begin{aligned} u_1'(x^*)u_2(1-x^*) - u_1(x^*)u_2'(1-x^*) &= 0 \\ \implies \frac{u_1'(x^*)}{u_1(x^*)} &= \frac{u_2'(1-x^*)}{u_2(1-x^*)}. \end{aligned}$$

Suppose in addition that Player 1 becomes more risk averse than Player 2:

$$\begin{aligned} v_1 : X &\rightarrow \mathbb{R}, \text{ where } v_1 = \phi(u_1), \quad \phi' > 0, \quad \phi'' < 0, \\ u_2 : X &\rightarrow \mathbb{R}. \end{aligned}$$

So the problem becomes:

$$\max_{x \in [0,1]} v_1(x)u_2(1-x) = \max_{x \in [0,1]} \phi(u_1(x))u_2(1-x).$$

The FOC is:

$$\begin{aligned} \phi'(u_1(x^{**}))u_1'(x^{**})u_2(1-x^{**}) - \phi(u_1(x^{**}))u_2'(1-x^{**}) &= 0 \\ \implies \frac{\phi'(u_1(x^{**}))u_1'(x^{**})}{\phi(u_1(x^{**}))} &= \frac{u_2'(1-x^{**})}{u_2(1-x^{**})}. \end{aligned}$$

Since ϕ is a concave function,

$$\phi'(t) < \frac{\phi(t)}{t} \implies \frac{\phi'(t)}{\phi(t)} < \frac{1}{t}.$$

Thus,

$$\frac{u_1'(x^*)}{u_1(x^*)} > \frac{\phi'(u_1(x^{**}))u_1'(x^{**})}{\phi(u_1(x^{**}))} = \frac{u_2'(1-x^{**})}{u_2(1-x^{**})}.$$

So in comparison,

$$\begin{cases} \frac{u_1'(x^*)}{u_1(x^*)} = \frac{u_2'(1-x^*)}{u_2(1-x^*)} \\ \frac{u_1'(x^{**})}{u_1(x^{**})} > \frac{u_2'(1-x^{**})}{u_2(1-x^{**})} \end{cases} \implies x^{**} < x^*.$$

3.3 Alternating Offer Model

Consider the problem of dividing a dollar: Two players will propose a division of a dollar in turns. Without loss of generality, suppose Player 1 proposes first. If Player 1 proposes to give x to Player 2, then Player 2 can either accept or reject. If Player 2 accepts, then the payoffs are $(1-x, x)$. If Player 2 rejects, then the game moves to the next period, where Player 2 proposes a division of a dollar and Player 1 can accept or reject. If Player 1 accepts, then the payoffs are $(x', 1-x')$. If Player 1 rejects, then the game continues to the next period, and so on. The game ends when one proposal is accepted. The discount factor for both players is $\delta \in (0, 1)$.

Claim

Any split $(x, 1-x)$ where $x \in (0, 1)$ is a NE outcome.

Proof for Claim.

This split can be supported by the strategies where Player 1 only accepts proposals that give them at least x , and Player 2 only accepts proposals that give them at least $1 - x$.


Theorem 3.4: Rubinstein

Let (x^*, y^*) be the solution to

$$\begin{aligned} u_1(y) &= \delta u_1(x) \\ u_2(1 - x) &= \delta u_2(1 - y) \end{aligned}$$

Then there is a unique subgame perfect equilibrium in which in G_1 , Player 1 asks for x^* and accepts iff $y \geq x^*$, and in G_2 , Player 2 offers y^* and Player 1 accepts iff $x \leq y^*$.

Proof for Theorem



Let (x^*, y^*) be the solution to

$$\begin{aligned} u_1(y) &= \delta u_1(x) \\ u_2(1 - x) &= \delta u_2(1 - y) \end{aligned}$$

Player 1 always offers x^* , and rejects Player 2's offer if $y < x^*$. Player 2 always offers y^* , and rejects Player 1's offer if $x > y^*$.

(z, t) is an outcome if a split of $(z, 1 - z)$ is accepted in period t .

Define $v_1(z, t)$ as the dollar amount such that Player 1 is indifferent between receiving $v_1(z, t)$ today and receiving z in period t .

$$u_1(v_1(z, t)) = \delta^{t-1} u_1(z).$$

Define $v_2(z, t)$ as the dollar amount such that Player 2 is indifferent between giving Player 1 $v_2(z, t)$ today and giving Player 1 z in period t .

$$u_2(1 - v_2(z, t)) = \delta^{t-1} u_2(1 - z).$$

Define

$$M_i = \sup \{v_i(z, t) : (z, t) \text{ is a SPE of } G_i\},$$

$$m_i = \inf \{v_i(z, t) : (z, t) \text{ is a SPE of } G_i\},$$

where G_i is the subgame starting with Player i 's proposal.

Claim

$$u_1(M_2) \leq \delta u_1(M_1).$$

Proof for Claim.

If Player 2 offers y such that $u_1(y) > \delta u_1(M_1)$, then Player 1 would accept the offer, because the right hand side $\delta u_1(M_1)$ is the most that Player 1 can get in the subsequent G_1 . So Player 2 would never offer such y . Hence, we must have $u_1(M_2) \leq \delta u_1(M_1)$. ■

Claim

$$u_1(M_2) \geq \delta u_1(M_1).$$

Proof for Claim.

If Player 2 offers y such that $u_1(y) < \delta u_1(M_1)$, then Player 1 will reject the offer, because by rejecting and entering G_1 , Player 1 can get at most M_1 (with current payoff equivalent $\delta u_1(M_1)$). Therefore, any successful offer in G_2 must satisfy $u_1(y) \geq \delta u_1(M_1)$. Taking the supremum over all y in G_2 equilibria, we have $u_1(M_2) \geq \delta u_1(M_1)$. ■

Claim

$$u_2(1 - M_1) \geq \delta u_2(1 - M_2).$$

Proof for Claim.

In G_1 , if Player 1 demands more than $v_2(M_2, 2)$, then Player 2 will reject. This implies

$$M_1 \leq v_2(M_2, 2).$$

And this implies

$$u_2(1 - M_1) \geq u_2(1 - v_2(M_2, 2)) = \delta u_2(1 - M_2).$$

Combining the three claims: $u_1(M_2) = \delta u_1(M_1)$ and $u_2(1 - M_1) \geq \delta u_2(1 - M_2)$. Together with $x^* \leq M_1$ (by definition of supremum) and the fact that the system in (1) admits a unique solution, we obtain $(M_1, M_2) = (x^*, y^*)$.

An entirely symmetric argument applied to the infimum pair (m_1, m_2) yields $(m_1, m_2) = (x^*, y^*)$ as well: if Player 1's worst SPE outcome in G_1 were below x^* , Player 2 could mimic that outcome, and a parallel chain of inequalities (with sup and inf reversed) contradicts (1). Since sup and inf coincide, every SPE of G_1 delivers exactly x^* to Player 1, and every SPE of G_2 delivers exactly y^* . The proposed strategies form the unique SPE.

Proposition 3.5: Convergence to the Nash Bargaining Solution

Let $(z^*, 1 - z^*)$ be the Nash Bargaining Solution to $\max_{0 \leq z \leq 1} u_1(z) u_2(1 - z)$. Then as $\delta \rightarrow 1$, both $x^*(\delta)$ and $y^*(\delta)$ converge to z^* .

Proof for Proposition.

From the Rubinstein system $u_1(y^*) = \delta u_1(x^*)$, the multiplier $\delta < 1$ forces $y^* < x^*$ for every fixed δ , and as $\delta \rightarrow 1$ the multiplier vanishes so $x^* - y^* \rightarrow 0$.

Multiply the two equations in (1):

$$u_1(y^*) u_2(1 - x^*) = \delta^2 u_1(x^*) u_2(1 - y^*),$$

which after rearrangement gives

$$u_1(x^*) u_2(1 - x^*) = u_1(y^*) u_2(1 - y^*).$$

So the two payoff vectors $(u_1(x^*), u_2(1 - x^*))$ and $(u_1(y^*), u_2(1 - y^*))$ lie on the *same* hyperbola $u_1 \cdot u_2 = \text{const}$. They are both efficient (they exhaust the unit pie), and their distance shrinks to zero as $\delta \rightarrow 1$. The unique limit point on the efficient frontier of this hyperbola family is exactly the NBS, which by definition maximizes $u_1 \cdot u_2$.

Remark (Strategic Foundation of an Axiomatic Solution).

This is the celebrated **Nash program** at work. The NBS was originally introduced as the unique axiomatic answer to the bargaining problem (scale invariance, efficiency, symmetry, IIA). The Rubinstein game gives it a *strategic* foundation: the unique SPE of an explicit non-cooperative bargaining protocol converges, as friction (impatience)

vanishes, to the same solution the axioms picked out. The two perspectives—“what *should* a fair split be?” and “what *will* two patient bargainers actually agree on?”—turn out to coincide in the limit, lending each weight to the other.

Example (Rubinstein with Square-Root Utilities).

Let $u_1(x) = u_2(x) = \sqrt{x}$ and $\delta = 0.9$. The Rubinstein system reads

$$\sqrt{y^*} = \delta\sqrt{x^*}, \quad \sqrt{1-x^*} = \delta\sqrt{1-y^*}.$$

Squaring both equations,

$$y^* = \delta^2 x^*, \quad 1 - x^* = \delta^2(1 - y^*).$$

Substituting the first into the second:

$$1 - x^* = \delta^2(1 - \delta^2 x^*) = \delta^2 - \delta^4 x^*,$$

so $x^*(1 - \delta^4) = 1 - \delta^2$, giving

$$x^* = \frac{1 - \delta^2}{1 - \delta^4} = \frac{1}{1 + \delta^2}.$$

Plugging $\delta = 0.9$ gives

$$x^* = \frac{1}{1 + 0.81} \approx 0.5525, \quad y^* = \delta^2 x^* \approx 0.4475.$$

Player 1 (who proposes first) keeps about $1 - x^* \approx 0.4475$ for herself and offers $x^* \approx 0.5525$ to Player 2; Player 2, when proposing, would offer Player 1 only $y^* \approx 0.4475$. Notice the **first-mover advantage**: as $\delta \rightarrow 1$, both x^* and y^* tend to $1/2$ and the advantage vanishes; for finite δ , the proposer captures the strictly larger share.

Remark (The Symmetric-Discount, Symmetric-Utility Formula).

The example generalizes: whenever both players share the same utility function u and the same discount δ , the Rubinstein system always reduces to a single algebraic equation in x^* , with $y^* = u^{-1}(\delta u(x^*))$. The asymmetric case, treated next, requires reducing to a common discount factor by reweighting one player’s utility—the clever transformation in the next section.

3.4 Unequal Discounting

Suppose now Player 1 has utility function w_1 and discount factor δ_1 , and Player 2 has utility function w_2 and discount factor δ_2 , with $\delta_1 < \delta_2$. We make the following transformation of the original problem:

- Player 2 has $u_2 = w_2$ and $\delta = \delta_2$.

- Player 1 has $u_1 = [w_1(x)]^{\frac{\ln \delta_2}{\ln \delta_1}}$ and $\delta = \delta_2$. (Note that $\frac{\ln \delta_2}{\ln \delta_1} < 1$ so that u_1 is concave, and is even “more concave” than w_1 .)

Without loss of generality, assume that

$$\begin{aligned}
& \delta_1^{t-1} w_1(z) > \delta_1^{t'-1} w_1(z') \\
\iff & [\delta_1^{t-1} w_1(z)]^{\frac{\ln \delta_2}{\ln \delta_1}} > [\delta_1^{t'-1} w_1(z')]^{\frac{\ln \delta_2}{\ln \delta_1}} \\
\iff & \frac{\ln \delta_2}{\ln \delta_1} [(t-1) \ln \delta_1 + \ln w_1(z)] > \frac{\ln \delta_2}{\ln \delta_1} [(t'-1) \ln \delta_1 + \ln w_1(z')] \\
\iff & (t-1) \ln \delta_2 + \frac{\ln \delta_2}{\ln \delta_1} \ln w_1(z) > (t'-1) \ln \delta_2 + \frac{\ln \delta_2}{\ln \delta_1} \ln w_1(z') \\
\iff & \ln \delta_2^{t-1} + \ln [w_1(z)]^{\frac{\ln \delta_2}{\ln \delta_1}} > \ln \delta_2^{t'-1} + \ln [w_1(z')]^{\frac{\ln \delta_2}{\ln \delta_1}} \\
\iff & \delta_2^{t-1} u_1(z) > \delta_2^{t'-1} u_1(z').
\end{aligned}$$

Remark (Chapter Summary).

Bargaining theory has two complementary strands. The *cooperative* strand axiomatizes a solution function $F(U, d)$ on the space of bargaining problems; the four Nash axioms (Pareto, symmetry, scale invariance, IIA) pin down the unique product-maximizing solution. The *non-cooperative* strand specifies a strategic protocol—Rubinstein’s infinite-horizon alternating-offer game (Theorem 3.3)—and derives a unique SPE in stationary strategies. The connection between the two is the *Nash program*: as the time between offers shrinks (or, equivalently, as $\delta \rightarrow 1$), the unique SPE outcome of the Rubinstein game converges to the Nash bargaining solution of the underlying problem. Cooperative axioms and non-cooperative protocols thus pick out the same point. The chapter also showed how to handle asymmetric discount rates by reweighting one player’s utility—a clever change of variables that reduces an asymmetric problem to a symmetric one and preserves the convergence result.

Part III

Auctions and Mechanism Design

Chapter 4

Private Value Auctions

Definition 4.1: Private-Value Auction Setting

A **private-value auction** consists of n risk-neutral bidders competing for a single indivisible object. Each bidder i independently draws a valuation $X_i \sim F$, where F is a continuous cumulative distribution function on $[0, \bar{x}]$ that is common knowledge. The realization x_i is bidder i 's *private information*—no other bidder, nor the seller, observes it. Each bidder submits a bid $b_i \in \mathbb{R}_{\geq 0}$, and the auctioneer determines the winner and the payment from the bid profile (b_1, \dots, b_n) according to the auction's rules.

4.1 Second Price Sealed Bid Auction (SPA)

Definition 4.2: Second-Price Sealed-Bid Auction

Each bidder simultaneously submits a sealed bid. The bidder with the highest bid wins the object and pays the *second*-highest bid; ties are broken uniformly at random.

Claim

It is a weakly dominant strategy for each bidder to bid their true value in a SPA, i.e., $b_i = x_i$.

Proof for Claim.

Graphical proof:

Suppose bidder 1 has value x_1 . We show that bidding $b_1 = x_1$ is a weakly dominant strategy for bidder 1. Let x_{-1} denote the highest competing bid. Bidder 1's profit is

$$\pi_1(b_1, x_{-1}) = \begin{cases} x_1 - x_{-1} & \text{if } b_1 > x_{-1}, \\ 0 & \text{if } b_1 < x_{-1}. \end{cases}$$

This corresponds to the **black** line in the figure below.

If instead bidder 1 bids $b_1^+ > x_1$, the profit becomes


$$\pi_1(b_1^+, x_{-1}) = \begin{cases} x_1 - x_{-1} & \text{if } b_1^+ > x_{-1}, \\ 0 & \text{if } b_1^+ < x_{-1}. \end{cases}$$

This corresponds to the **red** line in the figure below.

If bidder 1 bids $b_1^- < x_1$, the profit becomes

$$\pi_1(b_1^-, x_{-1}) = \begin{cases} x_1 - x_{-1} & \text{if } b_1^- > x_{-1}, \\ 0 & \text{if } b_1^- < x_{-1}. \end{cases}$$

This corresponds to the **blue** line in the figure below.



The figure shows that the black profit profile (truthful bid $b_1 = x_1$) weakly dominates both the red one (b_1^+ , which can yield negative payoffs) and the blue one (b_1^- , which forfeits positive surplus when $b_1^- < x_{-1} < x_1$) at every realization of x_{-1} . Hence $b_1 = x_1$ is weakly dominant. ■

4.2 First Price Sealed Bid Auction (FPA)

Definition 4.3: First-Price Sealed-Bid Auction

Each bidder simultaneously submits a sealed bid. The highest bidder wins and pays *her own* bid; ties are broken uniformly at random.

In a first price sealed bid auction, each bidder submits a bid without knowing the bids of others. The highest bidder wins the item and pays their own bid.

Example (Motivating Example of Two Bidders).

Suppose the true value X_i for bidder i is drawn from a uniform distribution on $[0, 1]$. Player 1 believes that Player 2 will bid $B_2 = kX_2$ for some $\frac{1}{2} \leq k < 1$. If Player 1 bids $b_1 \leq k$, her expected profit is

$$\begin{aligned}\mathbb{E}(\pi_1(b_1)) &= \Pr(kX_2 < b_1)(x_1 - b_1) \\ &= \Pr\left(X_2 < \frac{b_1}{k}\right)(x_1 - b_1) \\ &= \frac{b_1}{k}(x_1 - b_1).\end{aligned}$$

Then the best response of Player 1 to her belief $B_2 = kX_2$ with $\frac{1}{2} \leq k < 1$ is to bid $b_1 = \frac{x_1}{2}$.

Similarly, we can show that it is a symmetric equilibrium for both players to bid $b_i = \frac{x_i}{2}$ in a FPA.

Remark (Why $k \geq 1/2$?).

The condition $k \geq 1/2$ is strictly necessary to ensure that the unconstrained interior maximizer $b_1^* = \frac{x_1}{2}$ is globally valid for all possible valuations $x_1 \in [0, 1]$.

In the derivation, Player 1's expected profit function $\mathbb{E}(\pi_1(b_1)) = \frac{b_1}{k}(x_1 - b_1)$ is predicated on the assumption that $b_1 \leq k$, since Player 2's maximum possible bid is essentially $k \cdot 1 = k$. To guarantee that the optimal bid always falls within this valid domain (i.e., $\frac{x_1}{2} \leq k$), this inequality must hold for the supremum of x_1 . Since $x_1 < 1$, this immediately requires $k \geq \frac{1}{2}$.

If we were to relax this condition such that $k < 1/2$, a bidder with a sufficiently high valuation ($x_1 > 2k$) would find her unconstrained best response to be $\frac{x_1}{2} > k$. However, since bidding exactly k (or marginally above it) already guarantees a 100% probability of winning, bidding $\frac{x_1}{2}$ would strictly decrease her profit without any marginal gain in win rate. Consequently, the best response for high-valuation types would collapse into a corner solution $b_1 = k$. Thus, $k \geq 1/2$ rules out this corner solution and preserves the symmetric linear strategy over the entire support.

4.2.1 Symmetric Equilibrium of FPA

Now we extend the above example to the case with n bidders whose private values are drawn from an arbitrary distribution function F . We start the analysis by assuming there is a symmetric equilibrium where each bidder bids $b_i = \beta(X_i)$ for some strictly increasing function β .

Without loss of generality, suppose bidder 1 has value x_1 and bids b_1 , while all other

bidders follow the symmetric strategy β . The probability that bidder 1 wins is

$$\begin{aligned}
\Pr\left(b_1 > \max_{i \neq 1} b_i\right) &= \Pr\left(b_1 > \max_{i \neq 1} \beta(X_i)\right) \\
&= \Pr\left(b_1 > \beta\left(\max_{i \neq 1} X_i\right)\right) \\
&= \Pr\left(\beta^{-1}(b_1) > \max_{i \neq 1} X_i\right) \\
&= \Pr(X_i < \beta^{-1}(b_1), \forall i \neq 1) \\
&= [F(\beta^{-1}(b_1))]^{n-1}.
\end{aligned}$$

For simplicity, we define

$$Y_1 = \max_{i \neq 1} X_i.$$

Let G be the distribution function of Y_1 . Then we have

$$\begin{aligned}
G(y) &= \Pr(Y_1 \leq y) \\
&= \Pr\left(\max_{i \neq 1} X_i \leq y\right) \\
&= \Pr(X_i \leq y, \forall i \neq 1) \\
&= [F(y)]^{n-1}.
\end{aligned}$$

Then the probability that bidder 1 wins the auction can be rewritten as

$$\Pr\left(b_1 > \max_{i \neq 1} \beta(X_i)\right) = \Pr(Y_1 < \beta^{-1}(b_1)) = G(\beta^{-1}(b_1)).$$

The expected profit for bidder 1 is then given by

$$\begin{aligned}
\mathbb{E}(\pi_1(b_1)) &= \Pr\left(b_1 > \max_{i \neq 1} \beta(X_i)\right) (x_1 - b_1) \\
&= G(\beta^{-1}(b_1))(x_1 - b_1).
\end{aligned}$$

Since b_1 is the best response to the symmetric strategy β , we have

$$\begin{aligned}
\frac{d}{db_1} \mathbb{E}(\pi_1(b_1)) &= \frac{d}{db_1} G(\beta^{-1}(b_1))(x_1 - b_1) \\
&= G'(\beta^{-1}(b_1)) \cdot \frac{d}{db_1} \beta^{-1}(b_1) \cdot (x_1 - b_1) - G(\beta^{-1}(b_1)) \\
&= 0.
\end{aligned}$$

Note that $\frac{d}{db_1} \beta^{-1}(b_1) = \frac{1}{\beta'(\beta^{-1}(b_1))}$.¹ And by the optimality of β , we have $b_1 = \beta(x_1)$

¹This is true because $f(f^{-1}(x)) = x$ for any strictly increasing function f . Differentiating both sides with respect to x , we have $f'(f^{-1}(x)) \cdot \frac{d}{dx} f^{-1}(x) = 1$. Hence, $\frac{d}{dx} f^{-1}(x) = \frac{1}{f'(f^{-1}(x))}$.

and thus $\beta^{-1}(b_1) = x_1$. So the FOC can be rewritten as

$$\begin{aligned}
 & \frac{g(x_1)}{\beta'(x_1)} (x_1 - \beta(x_1)) - G(x_1) = 0 \\
 \implies & \beta'(x_1)G(x_1) + \beta(x_1)g(x_1) = x_1g(x_1) \\
 \implies & \frac{d}{dx_1} (\beta(x_1)G(x_1)) = x_1g(x_1) \\
 \implies & \beta(x_1)G(x_1) - \beta(0)G(0) = \int_0^{x_1} yg(y)dy \\
 \implies & \beta(x_1)G(x_1) = \int_0^{x_1} yg(y)dy \\
 \implies & \beta(x_1) = \frac{\int_0^{x_1} yg(y)dy}{G(x_1)} \\
 \implies & \beta(x_1) = \mathbb{E}[Y_1|Y_1 < x_1].
 \end{aligned}$$

A bidder with the lowest possible valuation $x = 0$ wins with probability zero, and any positive bid would yield a strictly negative expected payoff; hence $\beta(0) = 0$, which fixes the constant of integration.

Remark.

- **Why** $\frac{\int_0^{x_1} yg(y)dy}{G(x_1)} = \mathbb{E}[Y_1|Y_1 < x_1]$?

This equality follows directly from the definition of conditional expectation.

First, determine the conditional density of the highest competing valuation Y_1 given the event that the bidder wins (i.e., $Y_1 < x_1$). Restricting Y_1 to $[0, x_1]$, we normalize $g(y)$ by the total probability $\Pr(Y_1 < x_1) = G(x_1)$ to obtain

$$f(y | Y_1 < x_1) = \frac{g(y)}{G(x_1)} \quad \text{for } 0 \leq y < x_1.$$

By the standard definition of expected value, we integrate the variable y multiplied by its conditional PDF over the restricted support:

$$\begin{aligned}
 \mathbb{E}[Y_1|Y_1 < x_1] &= \int_0^{x_1} y \cdot f(y | Y_1 < x_1) dy \\
 &= \int_0^{x_1} y \left(\frac{g(y)}{G(x_1)} \right) dy.
 \end{aligned}$$

Since $G(x_1)$ only depends on x_1 and acts as a constant with respect to the integration variable y , we can factor it out of the integral:

$$\mathbb{E}[Y_1|Y_1 < x_1] = \frac{1}{G(x_1)} \int_0^{x_1} yg(y)dy.$$

- **Economic Intuition of FPA's Bidding Strategy:** The result $\beta(x) = \mathbb{E}[Y_1|Y_1 < x]$ reveals a profound economic intuition that elegantly connects the First-Price Auction (FPA) with the Second-Price Auction (SPA).

In an SPA, a winning bidder pays exactly the highest competing valuation, Y_1 . This

payment is determined *ex-post* (after the bids are submitted).

In an FPA, however, a winning bidder pays her own bid. Because the bid must be chosen *ex ante* (without knowing Y_1), the bidder effectively “simulates” the SPA in her mind: conditional on winning (which implies $Y_1 < x$), her best unbiased estimate of what she *would* have paid in an SPA is the expected highest competing valuation, $\mathbb{E}[Y_1 | Y_1 < x]$. The optimal FPA bid is therefore the bidder’s expected SPA payment.

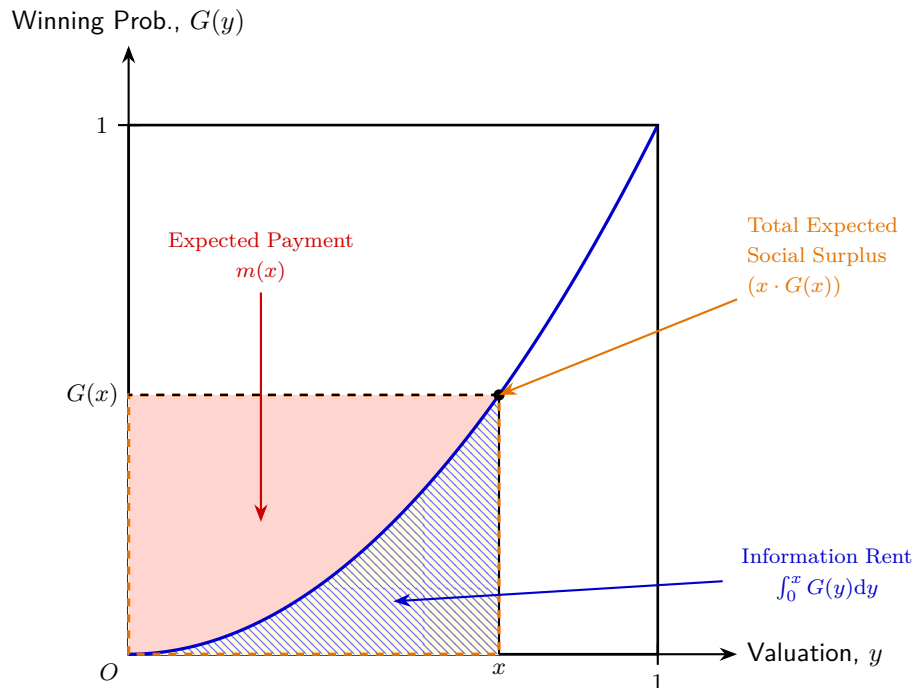
In an FPA, a bidder with value x bids $\beta(x) = \mathbb{E}[Y_1 | Y_1 < x]$ and pays her own bid upon winning. She wins iff $Y_1 < x$, which happens with probability $G(x)$. The expected payment $m(x)$ is therefore

$$m(x) = G(x) \cdot \mathbb{E}[Y_1 | Y_1 < x] = \int_0^x yg(y)dy.$$

Furthermore, by integration by parts,

$$\begin{aligned} m(x) &= \beta(x)G(x) \\ &= \int_0^x yg(y)dy \\ &= \int_0^x ydG(y) \\ &= [yG(y)] \Big|_0^x - \int_0^x G(y)dy \\ &= xG(x) - \int_0^x G(y)dy. \end{aligned}$$

This derived identity, $m(x) = xG(x) - \int_0^x G(y)dy$, allows for a beautiful geometric decomposition of the total social surplus generated by the auction. This surplus-splitting perspective is visualized in the figure below:



Remark (Surplus Splitting and Information Rent).

The identity $m(x) = xG(x) - \int_0^x G(y)dy$ provides a powerful geometric and economic interpretation of the auction's outcome, as illustrated in the figure above:

- **Total Expected Surplus ($xG(x)$):** Represented by the entire rectangle defined by $(x, G(x))$, this is the total social value created when an object is allocated to a bidder with valuation x (who wins with probability $G(x)$).
- **Information Rent ($\int_0^x G(y)dy$):** This corresponds to the area *under* the curve $G(y)$. Because the seller cannot observe the bidder's private valuation x , the bidder extracts a surplus from the mechanism. This "rent" is the cost the seller must pay to induce truthful revelation. It increases with valuation, reflecting the greater bargaining power of high-value types.
- **Expected Payment ($m(x)$):** This is the area *above* the curve but within the $x \times G(x)$ rectangle. It is the portion of the total surplus that the seller successfully captures as revenue.

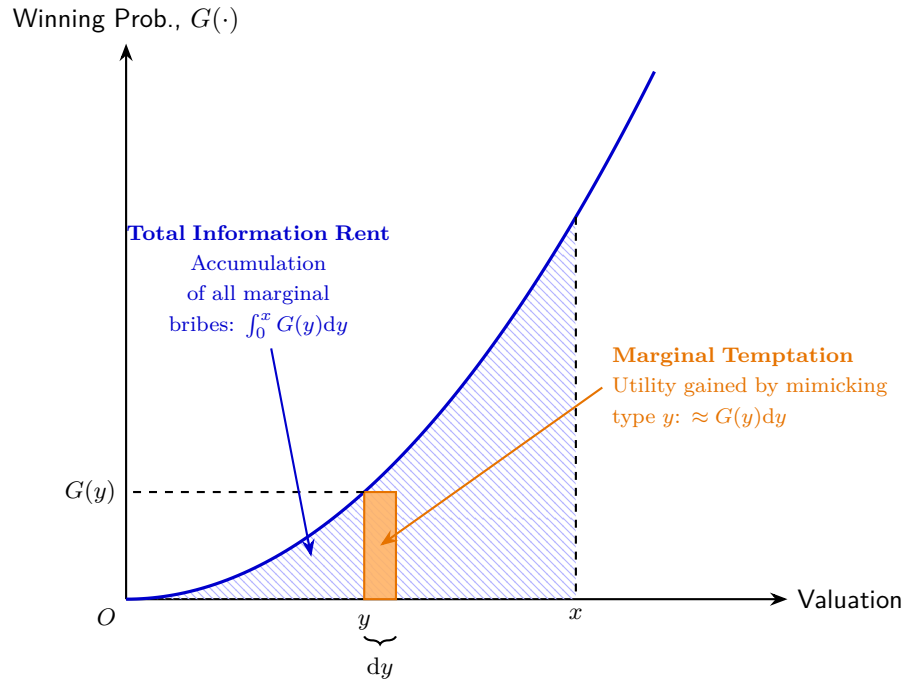
From this perspective, the Revenue Equivalence Theorem (shown later) holds because, in any efficient auction, the *information rent* is uniquely determined by the allocation rule $G(x)$.

Remark (Intuition and Derivation of the Information Rent).

- **Why the Information Rent is Equal to $\int_0^x G(y)dy$?**

In any auction, a high-value bidder wants to under-report her valuation to secure a lower price. Consider a bidder of true type $y + dy$ who instead reports type y : she wins with probability $G(y)$, and because her actual valuation is higher by dy , she gains a marginal utility of $G(y)dy$ from the misreport.

To achieve *Incentive Compatibility* (IC), the mechanism must compensate her for not lying. To prevent a type- x bidder from mimicking *any* lower type all the way down to 0, the mechanism must compensate her for all these accumulated marginal temptations. Integrating from 0 to x yields the total required information rent $\int_0^x G(y)dy$.



• **Mathematical Derivation (The Envelope Theorem):**

Let $U(z|x) = xG(z) - m(z)$ be the expected utility of a bidder with true value x reporting value z . Under truth-telling (which is optimal by design), the equilibrium utility—which is exactly the information rent—is defined as $V(x) = U(x|x)$.

By the *Envelope Theorem*, the total derivative of $V(x)$ with respect to x is simply the partial derivative of the objective function $U(z|x)$ with respect to the parameter x , evaluated at the optimal choice $z = x$:

$$V'(x) = \left. \frac{\partial U(z|x)}{\partial x} \right|_{z=x} = G(x)$$

This differential equation reveals that the marginal growth rate of the information rent is exactly the winning probability. Integrating both sides from 0 to x , and assuming the lowest type gets zero rent to satisfy Individual Rationality ($V(0) = 0$), we immediately obtain:

$$V(x) = \int_0^x G(y)dy$$

Remark (An Alternative, Intuitive Form of the Bidding Strategy in a FPA).

Recall that for a bidder with true value x , the bidding strategy is

$$\beta(x) = \frac{\int_0^x yg(y)dy}{G(x)}.$$

By integration by parts,

$$\int_0^x yg(y)dy = xG(x) - 0 \cdot G(0) - \int_0^x G(y)dy.$$

We can rewrite the bidding strategy in a more intuitive form:

$$\beta(x) = x - \frac{\int_0^x G(y)dy}{G(x)}.$$

This formulation perfectly captures the essence of strategic behavior in a FPA. A rational bidder will never bid their true valuation x (which would yield zero surplus); instead, they “shade” their bid downwards. This formula tells us exactly *how much* they shade:

- x is the bidder’s true valuation (their absolute maximum willingness to pay).
- $\int_0^x G(y)dy$ is the expected *information rent* the bidder is theoretically entitled to because their valuation is private.
- Dividing this expected rent by $G(x)$ converts it into a *conditional rent*—the actual rent the bidder secures *conditional on winning the auction*.

In an FPA, a winning bidder pays exactly her bid $\beta(x)$, so her realized profit conditional on winning is the deterministic amount $x - \beta(x)$. Since she wins only with probability $G(x)$, achieving the unconditional expected information rent $\int_0^x G(t)dt$ requires the realized profit to be scaled up by $1/G(x)$. Equating realized profit to the conditional rent yields $x - \beta(x) = \int_0^x G(t)dt/G(x)$, the standard bid-shading identity.

4.3 Reserve Prices

Definition 4.4: Reserve Price



A **reserve price** $r \geq 0$ is a minimum acceptable bid: the auction sells the object only if at least one submitted bid satisfies $b_i \geq r$. If all bids fall below r , the seller retains the object and no payment is made.

4.3.1 SPA with $n = 2$

In a SPA, the seller can set a reserve price r , meaning the item is only sold if the highest bid is no less than r . With the reserve price, the payment is then the maximum between the payment determined by the auction rule (namely the second-highest bid) and the reserve price. A natural question arises: *Is introducing a reserve price strictly profitable for the seller?*

To answer this intuitively, let’s assume there are $n = 2$ bidders with valuations drawn from a standard Uniform distribution, $F \sim U[0, 1]$. We can perform a marginal analysis by comparing the seller’s revenue with no reserve price ($r = 0$) to a scenario with a very small positive reserve price ($r \approx 0$). We can divide the possible valuation realizations into three

distinct cases:



- **Case 1: Both valuations are greater than r .** The reserve price has no effect on the outcome. The bidder with the highest valuation wins and pays the second-highest valuation. The expected revenue is exactly the same as if there were no reserve price.
- **Case 2: Both valuations are below the reserve price r .** In this case, the item is not sold, yielding a revenue of zero. If there were no reserve price, the seller would have received the lower of the two valuations. Since both are drawn from $U[0, r]$, the conditional expected payment would be $r/3$. This case happens with probability r^2 . Therefore, the unconditional expected loss for the seller is:

$$\text{Expected Loss} = r^2 \cdot \frac{r}{3} = \frac{r^3}{3}$$

- **Case 3: One valuation is above r , and the other is below r .** The bidder with the higher valuation wins. Without a reserve price, they would only pay the losing bidder's valuation (expected to be $r/2$). With the reserve price, the winner is forced to pay r . The probability of exactly one bidder drawing a value above r and one below r is $2r(1 - r)$. Therefore, the expected gain for the seller is:

$$\text{Expected Gain} = 2r(1 - r) \cdot \left(r - \frac{r}{2}\right) = 2r(1 - r) \cdot \frac{r}{2} = r^2 - r^3$$

Claim

$r = 0$ is never optimal for a revenue-maximizing seller.

Proof for Claim.

We evaluate the net change in expected revenue when introducing a small positive reserve price ($r > 0$) compared to not having one ($r = 0$). The net change is the expected gain from Case 3 minus the expected loss from Case 2:

$$\Delta \text{Revenue} = \text{Expected Gain} - \text{Expected Loss} = (r^2 - r^3) - \frac{r^3}{3} = r^2 - \frac{4}{3}r^3$$

For a sufficiently small r (specifically, $0 < r < \frac{3}{4}$), the r^2 term strictly dominates the r^3 term, meaning $r^2 > \frac{4}{3}r^3$. Consequently, $\Delta\text{Revenue} > 0$. Since setting a tiny positive reserve price strictly increases expected revenue compared to setting $r = 0$, a reserve price of zero is never optimal. ■

4.3.2 Optimal Reserve Price in a FPA

Claim

In an FPA with reserve price r , the symmetric bidding strategy for a bidder with valuation x is

$$\beta^r(x) = \begin{cases} \frac{rG(r) + \int_r^x tg(t)dt}{G(x)} & \text{if } x > r, \\ 0 & \text{if } x \leq r. \end{cases}$$

where $G(x) = F(x)^{n-1}$ is the distribution function of the highest valuation among the $n-1$ opposing bidders, and $g(x)$ is its corresponding probability density function.

Proof for Claim.

- **The Conditional Expectation (Intuitive Approach)**

Without a reserve price, a bidder's optimal strategy in a standard FPA is to bid the expected value of the highest competing valuation, conditional on winning: $\beta(x) = \mathbb{E}[Y_1 | Y_1 < x]$, where Y_1 is the highest valuation among opponents.

When a reserve price r is introduced, the seller acts as a "ghost bidder" who steadfastly bids r . To win, a bidder must defeat both the actual opponents ($Y_1 < x$) and the ghost bidder. Thus, the *effective* highest competing valuation becomes $\max(Y_1, r)$. The optimal bid is the conditional expectation of this effective opponent:

$$\beta^r(x) = \mathbb{E}[\max(Y_1, r) | Y_1 < x]$$

By definition of conditional expectation, we integrate over the distribution of Y_1 up to x :

$$\beta^r(x) = \int_0^x \max(y, r) \left(\frac{g(y)}{G(x)} \right) dy$$

We can split this integral at $t = r$:

$$\beta^r(x) = \frac{1}{G(x)} \left(\int_0^r rg(y)dy + \int_r^x yg(y)dy \right)$$

Since $\int_0^r g(y)dy = G(r)$, the first term evaluates to $rG(r)$. This leads directly to our familiar result:

$$\beta^r(x) = \frac{rG(r) + \int_r^x yg(y)dy}{G(x)}$$

- **First Order Approach**

Suppose a bidder has a true valuation $x > r$, but considers reporting a valuation z (and bidding $\beta(z)$). The bidder's expected utility is the probability of winning

multiplied by the surplus:

$$U(z, x) = (x - \beta(z))G(z)$$

To find the symmetric equilibrium, truth-telling must be optimal, meaning the first-order derivative with respect to z evaluated at $z = x$ must be zero:

$$\left. \frac{\partial U(z, x)}{\partial z} \right|_{z=x} = -\beta'(x)G(x) + (x - \beta(x))g(x) = 0$$

Rearranging terms yields a linear ordinary differential equation (ODE):

$$\beta'(x)G(x) + \beta(x)g(x) = xg(x) \implies \frac{d}{dx}[\beta(x)G(x)] = xg(x)$$

Integrating both sides from the reserve price r to x :

$$\beta(x)G(x) - \beta(r)G(r) = \int_r^x yg(y)dy$$

Consider the marginal buyer whose valuation is exactly $x = r$. To acquire the item without strictly losing money, this bidder must bid exactly the reserve price. Thus, $\beta(r) = r$. Substituting this boundary condition into the equation gives:

$$\beta(x)G(x) - rG(r) = \int_r^x yg(y)dy \implies \beta^r(x) = \frac{rG(r) + \int_r^x yg(y)dy}{G(x)}$$

Proposition 4.5: Optimal Reserve Price For a FPA

The optimal reserve price r^* for a revenue-maximizing seller in a First-Price Auction with n bidders and valuations drawn from F satisfies

$$r^* = \frac{1 - F(r^*)}{f(r^*)}.$$

Proof for Proposition.

For a FPA with a reserve price r , the symmetric bidding strategy is then

$$\beta^r(x) = \begin{cases} \frac{rG(r) + \int_r^x yg(y)dy}{G(x)} & \text{if } x > r, \\ 0 & \text{if } x \leq r. \end{cases}$$

The expected payment for a bidder with $x > r$ is therefore

$$m^r(x) = rG(r) + \int_r^x tg(t)dt.$$

And the expected payment for a bidder with value $x \leq r$ is zero: $m^r(x) = 0$.

The *ex-ante* expected payment for a bidder with valuation drawn from F is then

$$m^r = \int_0^1 m^r(x) f(x) dx = r [1 - F(r)] G(r) + \int_r^1 t [1 - F(t)] g(t) dt.$$

Hence, the optimal reserve price satisfies

$$\frac{dm^r}{dr} = [1 - F(r)] G(r) - rG(r)f(r) + r [1 - F(r)] g(r) - r [1 - F(r)] g(r) = 0.$$

Therefore, the optimal reserve price r^* is

$$r^* = \frac{1 - F(r^*)}{f(r^*)}.$$

Interestingly, the optimal reserve price does not depend on the number of bidders n .

4.4 Revenue Equivalence Theorem

We make a comparison of the expected payment in a FPA and a SPA with private value x .

In an SPA, the bidder bids x and, upon winning, pays the second-highest bid Y_1 . Conditional on winning, expected payment is $\mathbb{E}[Y_1 | Y_1 < x]$; the unconditional expected payment, weighting by the win probability $G(x)$, is $m(x) = G(x) \cdot \mathbb{E}[Y_1 | Y_1 < x]$.

In an FPA, the bidder bids $\beta(x) = \mathbb{E}[Y_1 | Y_1 < x]$ and pays her own bid upon winning, which occurs iff $Y_1 < x$ (probability $G(x)$).

The expected payment in both mechanisms for a bidder with valuation x is:

$$m(x) = G(x) \cdot \mathbb{E}[Y_1 | Y_1 < x] = \int_0^x yg(y) dy.$$

Ex ante, the expected payment for each player is

$$m = \mathbb{E}_X [m(X)] = \int_0^1 m(x) f(x) dx = \int_0^1 \left(\int_0^x yg(y) dy \right) f(x) dx.$$

Hence, the expected revenue (oftentimes called the expected selling price) is the same in a FPA and a SPA, which is just $n \times m$.

This result is more general than just comparing the FPA and SPA. It actually applies to any two auction mechanisms that satisfy certain conditions, as stated in the following theorems.

Theorem 4.6: Revenue Equivalence Theorem (RET)

Assume there are n risk-neutral bidders with independent private values drawn from a strictly increasing cumulative distribution function F on an interval $[0, \bar{x}]$. Any auction mechanism that satisfies the following two conditions:

1. **Allocative Efficiency:** The object is always allocated to the bidder with the highest valuation.
2. **Zero Payment at the Bottom:** Any bidder with the lowest possible valuation ($x = 0$) has an expected payment of zero.

will yield the exact same expected payment $m(x)$ for every bidder type x , and consequently, generate the exact same expected revenue for the seller.

Corollary 4.7

Standard auction formats such as the FPA and the SPA all satisfy these conditions and are therefore revenue equivalent.

Remark (The “Pie-Splitting” Economic Intuition).

- **Baking the Pie (Total Social Surplus):** The allocation rule decides who ultimately receives the object, which dictates the total value created by the mechanism (the size of the pie). Because standard auctions always allocate the object to the bidder with the highest valuation, they are all efficient and thus “bake” the exact same size of pie.
- **The Bidder’s Slice (Information Rent):** Because valuations are private, the seller must leave some surplus to the bidders to induce truth-telling. By the Envelope Theorem, this information rent is dictated *entirely* by the allocation probability ($V(x) = \int_0^x G(y)dy$). Since the allocation rules are identical across these standard auctions, the winning probabilities are identical, meaning the bidders carve out the exact same slice of the pie.
- **The Seller’s Slice (Expected Revenue):** The fundamental accounting identity of any mechanism is:

$$\text{Expected Revenue} = \text{Total Social Surplus} - \text{Information Rent}$$

If two auctions bake the exact same total pie, and the bidders are guaranteed to take the exact same size slice in both formats, then the remaining portion—the revenue left on the table for the seller—must logically be absolutely identical.

Having established the revenue equivalence of standard auctions, a natural question arises for a profit-maximizing seller: *Are standard auctions the best we can do?* To answer this, we must step out of the rigid, predefined rules of specific auction formats (like FPA or SPA)

and adopt the “God’s-eye view” of a mechanism designer. Instead of asking “what happens in an auction,” we ask: “over the space of *all possible mechanisms*, which one maximizes the seller’s expected revenue?” This requires a more fundamental, generalized mathematical law that governs *any* allocation rule, not just the efficient ones. The general framework—which we develop in the next chapter via the Revelation Principle and virtual values—reveals that the optimal mechanism is itself a slightly modified auction with a carefully chosen reserve price, and that VCG (covered in the second half of the next chapter) is the corresponding object on the *efficiency*-maximizing side rather than the revenue-maximizing one.

Theorem 4.8: General RET (Payoff Equivalence)

Assume there are n risk-neutral bidders with independent private values drawn from a strictly increasing cumulative distribution function F on an interval $[0, \bar{x}]$. Consider any two incentive-compatible mechanisms. If both mechanisms satisfy the following two conditions:

1. **Identical Allocation Rule:** The interim probability of receiving the object for a bidder with valuation x , denoted by $q(x)$, is exactly the same across both mechanisms.
2. **Identical Payment at the Bottom:** The expected payment of a bidder with the lowest possible valuation ($x = 0$) is the same across both mechanisms (typically zero).

Then, both mechanisms yield the exact same expected payment $m(x)$ for every bidder type x , and consequently, they generate the exact same expected revenue for the seller.

Remark (Why is this theorem more fundamental than the Standard RET?).

The Standard RET rigorously requires **Allocative Efficiency**—meaning the object *must* always go to the bidder with the highest valuation. The General RET, however, removes this straitjacket. It mathematically proves that for *any* arbitrary, monotonic allocation rule $q(x)$ (even inefficient ones), the payment function is completely and uniquely pinned down by the Envelope Theorem:

$$m(x) = m(0) + xq(x) - \int_0^x q(t)dt$$

This “Payoff Equivalence” is the ultimate theoretical weapon for the seller. It implies that the seller’s revenue maximization problem is no longer a daunting search through complex, multidimensional payment rules, but can be *entirely reduced to designing the optimal one-dimensional allocation rule $q(x)$* .

Armed with the General RET, the seller arrives at a somewhat counterintuitive economic truth: *To strictly maximize revenue, the seller must be willing to destroy social efficiency.* If the seller simply uses a standard efficient auction, they leave money on the table in the form of excessive information rents conceded to high-value bidders. To extract more surplus, for

example, the seller can strategically introduce a *reserve price* (r). With a reserve price, if the highest bidder's valuation is positive but falls below r , the seller retains the object. This explicitly violates the “*Allocative Efficiency*” condition of the Standard RET (because mutually beneficial trade is artificially blocked). However, the *General RET still applies perfectly!* We merely apply the generalized math to a newly shaped allocation rule $q^r(x)$, where the probability of winning is squashed to zero for all types below r . This strategic manipulation of the allocation rule to squeeze information rents leads directly to the seminal framework of Optimal Mechanism Design.

4.5 Risk Attitude and the Breakdown of Revenue Equivalence

All the analysis before was built on the assumption that bidders are risk-neutral. What happens when we relax this assumption? Suppose now that the bidders are strictly *risk averse*, with a utility function $u(\cdot)$ such that $u' > 0$, $u'' < 0$, and $u(0) = 0$.

First, for the SPA, risk aversion *has no effect* on the equilibrium of an SPA. Truth-telling ($\beta(x) = x$) remains a weakly dominant strategy.

However, the equilibrium in an FPA changes dramatically. Suppose the other $n - 1$ bidders follow a strictly increasing and differentiable symmetric strategy $\gamma(x)$ with $\gamma(0) = 0$. A bidder with true value x who submits bid b has expected utility

$$\begin{aligned} U(b, x) &= u(x - b) \cdot \Pr\left(b > \max_{j \neq i} \gamma(X_j)\right) \\ &= u(x - b) \cdot \Pr\left(\gamma^{-1}(b) > \max_{j \neq i} X_j\right) \\ &= u(x - b) \cdot G(\gamma^{-1}(b)), \end{aligned}$$

where $G(\cdot)$ is the distribution function of the highest opposing valuation.

To find the optimal bid, we take the FOC with respect to b :

$$\frac{\partial U(b, x)}{\partial b} = -u'(x - b) \cdot G(\gamma^{-1}(b)) + u(x - b) \cdot \frac{g(\gamma^{-1}(b))}{\gamma'(\gamma^{-1}(b))} = 0$$

In a symmetric equilibrium, the optimal bid must be $b = \gamma(x)$, which implies $\gamma^{-1}(b) = x$. Substituting this into the FOC gives:

$$-u'(x - \gamma(x)) \cdot G(x) + u(x - \gamma(x)) \cdot \frac{g(x)}{\gamma'(x)} = 0$$

Rearranging this equation yields the differential equation for the bidding strategy under risk aversion:


$$\gamma'(x) = \frac{u(x - \gamma(x))}{u'(x - \gamma(x))} \cdot \frac{g(x)}{G(x)}$$

To see how this compares to the risk-neutral case, recall that the risk-neutral bidding strategy $\beta(x)$ satisfies:

$$\beta'(x) = (x - \beta(x)) \cdot \frac{g(x)}{G(x)}$$

Note that because $u(\cdot)$ is a strictly concave function with $u(0) = 0$, its graph lies strictly below its tangent lines. Geometrically, for any point $t > 0$, the slope of the secant line connecting the origin to $(t, u(t))$ is strictly steeper than the slope of the tangent line at that exact point. Therefore, we directly have the slope inequality:

$$\frac{u(t)}{t} > u'(t) \implies \frac{u(t)}{u'(t)} > t$$



Proposition 4.9: Overbidding under Risk Aversion

For any strictly positive valuation $x > 0$, a risk-averse bidder bids strictly higher than a risk-neutral bidder:

$$\gamma(x) > \beta(x).$$

Proof for Proposition.

Note that a bidder with the lowest possible valuation ($x = 0$) bids zero regardless of their risk attitude, so we have the boundary condition $\gamma(0) = \beta(0) = 0$.

Suppose, for the sake of contradiction, that there exists some $x > 0$ where $\gamma(x) \leq \beta(x)$. Let $t = x - \gamma(x) > 0$. By the concavity property discussed above, $\frac{u(t)}{u'(t)} > t$, meaning:

$$\frac{u(x - \gamma(x))}{u'(x - \gamma(x))} > x - \gamma(x) \geq x - \beta(x).$$

Substituting this strict inequality into our differential equation for $\gamma'(x)$:

$$\gamma'(x) = \frac{u(x - \gamma(x))}{u'(x - \gamma(x))} \cdot \frac{g(x)}{G(x)} > (x - \beta(x)) \cdot \frac{g(x)}{G(x)} = \beta'(x).$$

This implies that at any point where $\gamma(x)$ is equal to or below $\beta(x)$, the slope of γ is strictly steeper than the slope of β . Since both functions start at the same origin ($\gamma(0) = \beta(0) = 0$), γ must immediately rise above β . Furthermore, γ can never cross β

from above, because at any intersection, γ must be steeper. Thus, we conclude $\gamma(x) > \beta(x)$ for all $x > 0$. ■

Remark (The “Insurance Premium” Intuition).

Why does a risk-averse bidder bid higher? In a First-Price Auction, a bidder faces a trade-off: bidding lower increases the profit *if* they win, but increases the probability of losing and getting exactly zero. A risk-averse individual deeply dislikes the downside risk of losing the item. To protect themselves against this zero-payoff outcome, they are willing to sacrifice some potential profit by bidding more aggressively. In essence, the difference $\gamma(x) - \beta(x)$ is an **insurance premium** the bidder pays to increase their probability of winning.

Because risk aversion causes bidders to bid more aggressively in the FPA, while leaving the dominant strategy in the SPA unchanged, the expected payments in the FPA are strictly higher. Therefore, under risk aversion, the First-Price Auction generates strictly higher expected revenue for the seller than the Second-Price Auction.

Remark (Chapter Summary).

Single-object private-value auction theory rests on three pillars. *Strategic equivalence within the standard family.* The second-price (Vickrey) auction has a weakly dominant truthful strategy, and the first-price auction has a unique symmetric equilibrium $\beta(x) = \mathbb{E}[Y_1 \mid Y_1 < x]$ in which bidders shade their bids by exactly their conditional information rent. *Revenue equivalence* (Theorem 4.4). All standard auctions that allocate the object to the highest-value bidder and give zero rent to the lowest type generate the same expected revenue—the auction format is irrelevant for revenue under risk neutrality, independent private values, and symmetric bidders. The three assumptions matter: risk aversion breaks revenue equivalence in favor of the FPA, asymmetric distributions break it ambiguously, and correlated values can be exploited by full-surplus extraction mechanisms. *Optimal reserves.* Setting $r^* = \varphi^{-1}(0)$, the root of the virtual value, strictly improves revenue at the cost of efficiency—the seller deliberately sometimes refuses to sell to extract higher payments from high-value bidders.

Chapter 5

Mechanism Design

5.1 Direct Mechanism and Revelation Principle

In exploring the optimal way to sell an object to n buyers, we establish a rigorous mathematical framework. Suppose the seller's reservation value for the object is x_0 , and the buyers' true valuations X_i are drawn independently from a distribution F_i .

Any complex selling procedure in reality can be abstracted as a *general mechanism*. This mechanism consists of three mathematical components, denoted as (\mathcal{B}, π, μ) :

- **Message Space (\mathcal{B}):** The set of messages (or “bids”) that buyers can submit. For n buyers, the global message space is $\mathcal{B} = \mathcal{B}_1 \times \mathcal{B}_2 \times \cdots \times \mathcal{B}_n$.
- **Allocation Rule ($\pi(b)$):** Given a profile of submitted messages b , $\pi_i(b)$ determines the probability that buyer i receives the object.
- **Payment Rule ($\mu(b)$):** Given a profile of submitted messages b , $\mu_i(b)$ determines the amount that buyer i must pay to the seller.

Under this mechanism, suppose the buyers play a game and reach an equilibrium. Each buyer employs an *equilibrium strategy* $\beta_i(x_i)$, which maps their true valuation x_i to an optimal message b_i that maximizes their expected utility.

Analyzing general mechanisms can be mathematically daunting due to the potentially infinite complexity of the message space. To drastically simplify this problem, we introduce the concept of a *direct mechanism*.

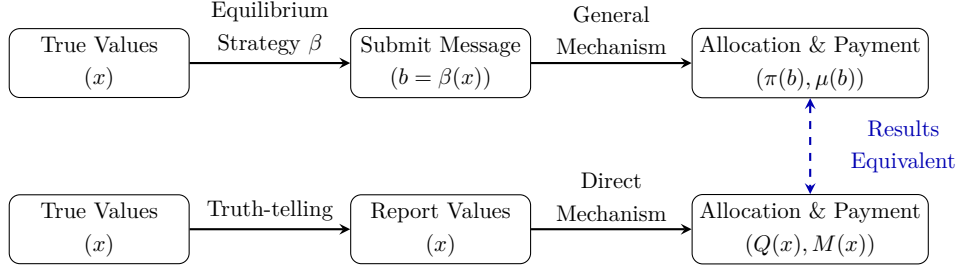
In a direct mechanism $([0, 1]^n, Q, M)$, the seller simply asks the buyers to report their true valuations directly. Thus, the message space is restricted to the valuation space itself (i.e., $\mathcal{B}_i = [0, 1]$). The mechanism's allocation and payment rules are implemented by applying the equilibrium strategies $\beta(x)$ from the original mechanism “behind the scenes”:

- **Direct Allocation Rule:** $Q_i(x) = \pi_i(\beta(x))$
- **Direct Payment Rule:** $M_i(x) = \mu_i(\beta(x))$

The Revelation Principle states a profound result: If there exists an equilibrium outcome in some complex general mechanism, we can always construct a corresponding direct mechanism where truth-telling (reporting true valuations honestly) constitutes an equilibrium

for all buyers. Furthermore, this truth-telling equilibrium yields the exact same allocation probabilities and expected payments as the original mechanism.

This principle is the cornerstone of mechanism design. It greatly simplifies the search for an “optimal auction”: instead of exhaustively searching through all possible auction formats and bidding rules, the mechanism designer only needs to optimize over the set of direct mechanisms subject to the constraint that buyers are willing to report truthfully (the *incentive compatibility constraint*).



Theorem 5.1: Revelation Principle

If β is an equilibrium of a general mechanism (\mathcal{B}, π, μ) , then truth-telling (reporting x_i honestly) is an equilibrium of the corresponding direct mechanism $([0, 1]^n, Q, M)$.

Proof for Theorem

Intuitively, if a buyer found it optimal to play the strategy $\beta(x_i)$ in the original complex mechanism, they will find it optimal to simply report their true type x_i in the direct mechanism, because the direct mechanism automatically plays their optimal strategy $\beta(x_i)$ on their behalf. Any deviation from truth-telling would result in the mechanism executing a sub-optimal strategy for that buyer. ■

Corollary 5.2: Expected Payment Identity

In any incentive-compatible mechanism with a reserve price r , the expected payment $m(x)$ of a bidder with valuation $x \geq r$ is uniquely determined by the allocation rule $G(\cdot)$ and the expected payment of the marginal type $m(r)$. Specifically, it is given by:

$$m(x) = m(r) + \int_r^x yg(y)dy,$$

where $g(y) = G'(y)$ is the density function associated with the allocation rule.

Proof for Corollary.

Let $U(z|x) = xG(z) - m(z)$ denote the expected utility of a bidder with true valuation x who reports z . By the Revelation Principle, we focus on the truth-telling equilibrium where the value function is:

$$V(x) = U(x|x) = \max_z (xG(z) - m(z)).$$

By the Envelope Theorem, the derivative of the value function with respect to x is simply

$V'(x) = G(x)$. Integrating this from the reserve price r to x yields:

$$V(x) = V(r) + \int_r^x G(y)dy.$$

Using the definition $V(x) = xG(x) - m(x)$, we can isolate the expected payment:

$$m(x) = xG(x) - V(r) - \int_r^x G(y)dy.$$

Applying integration by parts to the term $\int_r^x yg(y)dy = xG(x) - rG(r) - \int_r^x G(y)dy$, we can rewrite the payment equation as:

$$m(x) = rG(r) + \int_r^x yg(y)dy - V(r).$$

Recall that the equilibrium utility for the marginal bidder is $V(r) = rG(r) - m(r)$. Substituting this definition into the equation above immediately yields the desired result:

$$m(x) = m(r) + \int_r^x yg(y)dy.$$

5.2 Optimal Mechanism Design (Myerson, 1981)

5.2.1 Motivating Example: Monopoly Pricing

What is the best way to sell an object to n buyers? Let us start with the simplest case: $n = 1$.

Suppose the single buyer's valuation X is drawn from a distribution with cumulative distribution function $F(x)$ and probability density function $f(x)$. The seller wants to post a take-it-or-leave-it price p .

Importantly, let x_0 be the *seller's own valuation* (or reservation value) of the object. This represents the utility the seller retains if the transaction fails and they keep the item.

The seller's expected profit $\mathbb{E}[\pi(p)]$ from posting price p consists of two mutually exclusive scenarios:

- **Transaction succeeds (Probability $1 - F(p)$):** The buyer's valuation $X \geq p$. The seller collects the payment p .
- **Transaction fails (Probability $F(p)$):** The buyer's valuation $X < p$. The seller keeps the object and retains a value of x_0 .

Thus, the expected profit is given by:

$$\mathbb{E}[\pi(p)] = p \cdot \Pr(X \geq p) + x_0 \cdot \Pr(X < p) = p(1 - F(p)) + x_0F(p)$$

To find the revenue-maximizing optimal price p^* , we take the first derivative of the

expected profit with respect to p and set it to zero:

$$\frac{d\mathbb{E}[\pi(p)]}{dp} = (1 - F(p)) - pf(p) + x_0f(p) = 0$$

By dividing by $f(p)$ and rearranging the terms, we arrive at a profound economic condition:

$$p^* - \frac{1 - F(p^*)}{f(p^*)} = x_0$$

The expression $p - \frac{1-F(p)}{f(p)}$ is a fundamental concept in mechanism design known as the *virtual value*. The term $\frac{f(p)}{1-F(p)}$ is defined as the *hazard rate* of the valuation distribution. This equation dictates that a revenue-maximizing seller should set the optimal posted price p^* exactly at the point where the buyer's *virtual value* equals the seller's *true valuation* x_0 .

In many standard textbook problems, we assume the object has zero use value to the seller, i.e., $x_0 = 0$. In such cases, the FOC simplifies to $p^ - \frac{1-F(p^*)}{f(p^*)} = 0$.*

5.2.2 Framing Seller's Problem

By the Revelation Principle, instead of searching through an infinite space of possible game rules and message spaces, the seller *only needs to find the optimal direct mechanism* (Q, M) subject to incentive compatibility (truth-telling) constraints. There is no loss of generality in restricting our attention to direct mechanisms where the message space is simply the type space, $\mathcal{B}_i = [0, 1]$.

The mechanism designer (the seller) optimizes over two functions: the *allocation rule* Q and the *payment rule* M :

$$Q : [0, 1]^n \rightarrow \Delta, \quad M : [0, 1]^n \rightarrow \mathbb{R}^n,$$

where Δ is the probability simplex over the n buyers (i.e., $Q_i(x)$ is the probability that buyer i receives the object given the profile of reported valuations x).

Definition 5.3: Interim Expected Probability of Winning & Interim Expected Payment

For any bidder i , suppose they report their valuation as z_i while all other $n - 1$ bidders report truthfully. The *interim expected probability of winning* $q_i(z_i)$ and the *interim expected payment* $m_i(z_i)$ are defined as:

$$q_i(z_i) = \int_{[0,1]^{n-1}} Q_i(z_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i}$$

$$m_i(z_i) = \int_{[0,1]^{n-1}} M_i(z_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i}$$

where $f_{-i}(x_{-i}) = \prod_{j \neq i} f_j(x_j)$ is the joint density of all other bidders' true valuations.

The expected payoff for bidder i with true valuation x_i who reports z_i is therefore:

$$\pi_i(x_i, z_i) = q_i(z_i) \cdot x_i - m_i(z_i)$$

To ensure that truth-telling is a Bayesian Nash Equilibrium,¹ the mechanism must satisfy the *incentive compatibility constraint*:

Incentive Compatibility Constraint (IC)

A mechanism (Q, M) is *incentive compatible* if the following holds:

$$\pi_i(x_i, z_i) \geq \pi_i(x_i, x_i), \quad \forall i, \quad \forall x_i \in [0, 1], \quad \forall z_i \in [0, 1].$$

If the IC constraint is satisfied, every bidder will tell the truth. In this case, we can express the equilibrium payoff of bidder i as a function of their true type x_i :

$$u_i(x_i) \equiv \pi_i(x_i, x_i) = \max_{z_i} \{q_i(z_i) \cdot x_i - m_i(z_i)\}$$

Claim: Properties of the IC Mechanism

From the definition of the equilibrium payoff function $u_i(x_i)$,

$$u_i(x_i) \equiv \pi_i(x_i, x_i) = \max_{z_i} \{q_i(z_i) \cdot x_i - m_i(z_i)\},$$

we immediately obtain three fundamental results:

1. $u_i(x_i)$ is a convex function of x_i .
2. $u'_i(x_i) = q_i(x_i)$ almost everywhere.
3. $q_i(x_i)$ is non-decreasing in x_i .

Proof for Claim.

- **Convexity:** The value function $u_i(x_i)$ is defined as the upper envelope (the point-wise maximum) of a family of functions in x_i . Since $q_i(z_i)x_i - m_i(z_i)$ is an affine (linear) function of x_i , their upper envelope is globally convex.^a
- **Envelope Theorem:** By the Envelope Theorem, the total derivative of the maximized value function with respect to the parameter x_i equals the partial derivative of the objective function evaluated at the optimal choice ($z_i = x_i$). Thus, $u'_i(x_i) = q_i(x_i)$.
- **Monotonicity:** Because $q_i(x_i) = u'_i(x_i)$ and thus $q'_i(x_i) = u''_i(x_i)$, and u_i is convex ($u''_i(x_i) \geq 0$), $q_i(x_i)$ is naturally non-decreasing in x_i .

^a**Linear vs. Affine:** Strictly speaking, a *linear* function must pass through the origin ($f(x) = ax$), whereas an *affine* function is a linear function plus a constant translation ($f(x) = ax + b$). Sometimes economists sometimes use the terms interchangeably.

¹**Nash Equilibrium (NE) vs. Bayesian Nash Equilibrium (BNE):** In a traditional NE, players have complete information about the game, including the exact payoffs of all other players. In a mechanism design setting, information is incomplete: a bidder knows their own valuation but only knows the *distribution* of others' valuations. A BNE is the extension of NE to games of incomplete information, where each player's strategy maximizes their *expected* payoff, taking expectations over the possible types of other players based on a common prior distribution.

Proposition 5.4: Allocation Rule Uniquely Determines Payment Rule

In any incentive-compatible (IC) and individually rational (IR) mechanism designed by a profit-maximizing seller, the equilibrium expected payoff $u_i(x_i)$ and the expected payment $m_i(x_i)$ for any bidder i are completely and uniquely determined by the allocation rule Q . Specifically, the lowest type extracts zero surplus ($u_i(0) = 0$), yielding the exact payoff function:

$$u_i(x_i) = \int_0^{x_i} q_i(t) dt$$

Proof for Proposition.

Since $u_i'(x_i) = q_i(x_i)$, we can express the payoff function as an integral of the interim winning probability:

$$u_i(x_i) = u_i(0) + \int_0^{x_i} q_i(t) dt$$

This implies that once the allocation rule Q (and thus the interim probability q_i) is fixed, the equilibrium payoff function u_i is entirely pinned down up to a constant $u_i(0)$. Loosely speaking, the IC constraint determines the *shape* of the payoff function. Consequently, the shape of the expected payment $m_i(x_i) = q_i(x_i)x_i - u_i(x_i)$ is also uniquely determined by Q .

Furthermore, to maximize revenue, the seller should extract as much surplus as possible without violating the *individual rationality constraint*:

Individual Rationality Constraint (IR)

A mechanism (Q, M) is *individually rational* if the following holds:

$$u_i(x_i) \geq 0, \quad \forall i.$$

Since $u_i(x_i)$ is non-decreasing, the IR constraint is satisfied everywhere if and only if $u_i(0) \geq 0$. Note that by definition, $u_i(0) = q_i(0) \cdot 0 - m_i(0) = -m_i(0)$, so:

$$u_i(0) \geq 0 \implies m_i(0) \leq 0.$$

Because the seller is a profit-maximizer, $m_i(0) \leq 0$ immediately implies they will set $m_i(0) = 0$, which means $u_i(0) = 0$. Thus:

$$u_i(x_i) = \int_0^{x_i} q_i(t) dt.$$

That is, the IR constraint further pins down the constant term in the payoff function, leaving *zero degrees of freedom* for the payoff function (and thus the payment rule) once the allocation rule Q is fixed. ■

The seller’s ultimate problem is to maximize total expected revenue:


$$\max_{Q, M} \sum_{i=1}^n \int_0^1 m_i(x_i) f_i(x_i) dx_i$$

Substitute the expected payment identity $m_i(x_i) = q_i(x_i)x_i - \int_0^{x_i} q_i(t)dt$ into the integral for bidder i :

$$\int_0^1 m_i(x_i) f_i(x_i) dx_i = \int_0^1 q_i(x_i) x_i f_i(x_i) dx_i - \int_0^1 \left(\int_0^{x_i} q_i(t) dt \right) f_i(x_i) dx_i$$

We swap the order of integration for the second term (visual proof below):

$$\int_0^1 \int_0^{x_i} q_i(t) dt f_i(x_i) dx_i = \int_0^1 \int_t^1 f_i(x_i) dx_i q_i(t) dt = \int_0^1 (1 - F_i(x_i)) q_i(x_i) dx_i$$



Substituting this back into the revenue equation gives us the expected revenue purely as a function of the allocation rule:

$$\int_0^1 m_i(x_i) f_i(x_i) dx_i = \int_0^1 \left[x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} \right] q_i(x_i) f_i(x_i) dx_i$$

We define the bracketed term as the *virtual value*:

Definition 5.5: Virtual Value

The *virtual value* of a bidder with true type x_i , denoted by $\phi_i(x_i)$, is defined as

$$\phi_i(x_i) = x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}.$$

Why? To understand the logical leap to “money”, we must look at the mechanics of the Incentive Compatibility (IC) constraint: $u_i(x_i) = \int_0^{x_i} q_i(t)dt$.

Suppose the seller decides to increase the allocation probability $q_i(x)$ for a specific type x . Because the expected payoff u_i is the integral of q_i from 0 up to the true type, increasing $q_i(x)$ does not just increase the payoff for type x —it automatically increases the payoff for every single type strictly greater than x !

If the seller makes it attractive for type x to win the object, mimicking type x suddenly becomes very tempting for everyone with a higher valuation. To prevent all higher types from lying and claiming to be type x to get a cheaper price, the seller is forced to “bribe” them by leaving more surplus (money) on the table.

How many higher types are there? The probability mass of bidders with a valuation $> x$ is exactly $1 - F_i(x)$. Therefore, $1 - F_i(x)$ represents the total volume of “bribes” (information rent) the seller must pay out to the upper tail of the distribution. Dividing this by the local density $f_i(x)$ simply normalizes this total cost into a *marginal* cost evaluated at type x .

Thus, $\frac{1 - F_i(x_i)}{f_i(x_i)}$ is exactly the expected dollar amount of surplus the mechanism must concede *locally* at type x_i : allocating to x_i forces the seller to surrender rent to the $1 - F_i(x_i)$ mass of people above them!

Remark (Understanding the Information Rent).

The virtual value $\phi_i(x_i)$ is strictly less than the true valuation x_i . The subtracted term, $\frac{1 - F_i(x_i)}{f_i(x_i)}$ (the inverse hazard rate), represents the *information rent*. This is because of the IC constraint: If the seller allocates the object to a bidder with valuation x_i , they cannot simply charge them x_i . To prevent all types strictly higher than x_i from lying and mimicking type x_i to get a cheaper price, the seller must leave some surplus (rent) on the table for those higher types.

By the Envelope Theorem, a bidder of type x_i earns a payoff of $u_i(x_i) = \int_0^{x_i} q_i(t) dt$. When we take the expectation of this payoff across all possible types, the mathematical result is $\mathbb{E}[u_i] = \mathbb{E}\left[q_i(x_i) \frac{1 - F_i(x_i)}{f_i(x_i)}\right]$. Thus, $\frac{1 - F_i(x_i)}{f_i(x_i)}$ is exactly the expected dollar amount of surplus (information rent) the mechanism must concede *locally* at type x_i to ensure all higher types truthfully reveal themselves.

Therefore, $\phi_i(x_i)$ is the *actual marginal revenue* the seller extracts when allocating the object to a bidder with value x_i .

Recall that by definition, the interim allocation probability is the expected value of the ex-post allocation rule over all other bidders’ types:

$$q_i(x_i) = \int_{[0,1]^{n-1}} Q_i(x_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i}$$

By expanding $q_i(x_i)$ back into the joint integral over all n bidders, the seller’s optimization problem simplifies remarkably to maximizing the expected virtual surplus:

$$\max_Q \int_{[0,1]^n} \left[\sum_{i=1}^n \phi_i(x_i) Q_i(x) \right] f(x) dx$$

subject to the physical constraints that $Q_i(x) \in [0, 1]$, $\sum Q_i(x) \leq 1$, and the IC monotonicity constraint that $q_i(\cdot)$ is non-decreasing.

Assumption 5.6: Regularity Condition on ϕ_i

The virtual value $\phi_i(x_i)$ is strictly increasing in x_i .

Remark (Why Regularity Guarantees Monotonicity).

To understand why this condition is critical, recall that the seller’s objective is to maximize the expected virtual surplus. For any given profile of reported types x , the seller

wants to choose an allocation rule $q(x)$ to maximize:

$$\max_q \sum_{i=1}^n q_i(x) \phi_i(x_i)$$

subject to the feasibility constraint $\sum q_i(x) \leq 1$.

Ignoring the IC constraint for a moment, the unconstrained optimal way to maximize this sum is simply to allocate the object to the bidder with the highest strictly positive virtual value. That is, $q_i(x) = 1$ if $\phi_i(x_i) > \max_{j \neq i} \phi_j(x_j)$ and $\phi_i(x_i) > 0$.

Recall from the Envelope Theorem that for a mechanism to be Incentive Compatible (IC), the interim allocation probability $q_i(x_i)$ must be non-decreasing in x_i .

If the Regularity Condition holds (i.e., ϕ_i is strictly increasing), then a higher true type x_i directly translates to a strictly higher virtual value $\phi_i(x_i)$. A higher virtual value makes it strictly more likely that bidder i outbids all competitors and clears the seller's reserve price. Therefore, the allocation probability $q_i(x_i)$ naturally goes up as x_i goes up.

In short, the regularity condition ensures that the greedy, unconstrained pointwise maximization algorithm *automatically* satisfies the global IC monotonicity constraint! (If regularity fails and ϕ_i decreases in some regions, higher types might be assigned lower probabilities, violating IC.)

This regularity condition ensures that pointwise maximization naturally satisfies the monotonicity constraint of q_i .

Proposition 5.7: Optimal Allocation Rule and Reserve Prices

Under the regularity assumption, we can solve the integral via *pointwise maximization*. For any realized profile of valuations $x = (x_1, \dots, x_n)$, the seller should simply examine the virtual values and assign the object to the bidder with the highest virtual value, provided it is non-negative.

Mathematically, the optimal allocation rule is:

$$Q_i^*(x) = \begin{cases} 1 & \text{if } \phi_i(x_i) = \max_j \phi_j(x_j) \text{ and } \phi_i(x_i) \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

From this, the seller can define an optimal reserve price for each bidder as follows:

Definition 5.8: Individual Reserve Price

The optimal reserve price r_i^* for bidder i is exactly the root of their virtual value function:

$$\phi_i(r_i^*) = 0 \implies r_i^* - \frac{1 - F_i(r_i^*)}{f_i(r_i^*)} = 0$$

Notably, to maximize revenue, the seller does *not* necessarily allocate the object to the person with the highest true value, but to the person with the highest *virtual value*. Furthermore, if no bidder's virtual value exceeds zero (i.e., everyone's valuation is below their respective r_i^*), the seller optimally retains the object to avoid paying excessive information

rents. From this perspective, the optimal mechanism is to allocate the object to the buyer with the highest valuation *strictly above their respective reserve price*.

5.2.3 The Monopoly Pricing Isomorphism: What is Virtual Value?

While the derivation of the optimal mechanism using the Envelope Theorem and integration by parts is mathematically rigorous, the economic intuition behind the *virtual valuation* can be beautifully illuminated by a simple parallel: *Standard Monopoly Pricing*.

Imagine a monopolist selling a single object to a market of buyers whose valuations are distributed according to $F(\cdot)$ with density $f(\cdot)$.

Claim: Virtual Valuation is Marginal Revenue

If we view the probability of sale as the “quantity” demanded, the virtual valuation $\psi(p)$ is exactly the monopolist’s marginal revenue evaluated at price p .

Proof for Claim.

Let p be the price set by the monopolist. A buyer will purchase the object if their valuation is greater than or equal to p . Thus, the quantity demanded q at price p is exactly the survival function:

$$q(p) = 1 - F(p)$$

To find the marginal revenue, we first need the inverse demand curve $P(q)$. By inverting the demand function, we get:

$$P(q) = F^{-1}(1 - q)$$

The monopolist’s total revenue $R(q)$ as a function of quantity is price times quantity:

$$R(q) = q \cdot P(q) = q \cdot F^{-1}(1 - q)$$

Now, we differentiate $R(q)$ with respect to q to find the MR :

$$\begin{aligned} MR(q) &= \frac{dR(q)}{dq} \\ &= F^{-1}(1 - q) + q \left[\frac{1}{f(F^{-1}(1 - q))} \cdot (-1) \right] \\ &= F^{-1}(1 - q) - \frac{q}{f(F^{-1}(1 - q))} \end{aligned}$$


Finally, we substitute the original price variable $p = F^{-1}(1 - q)$ and quantity $q = 1 - F(p)$ back into the MR equation to express marginal revenue as a function of price:

$$MR(p) = p - \frac{1 - F(p)}{f(p)}$$

This expression is exactly the definition of the *virtual valuation* ■

This profound mathematical equivalence means we can analyze optimal auction mechanisms using the familiar geometric tools of intermediate microeconomics. The optimal mech-

anism simply allocates the good to the buyer with the highest MR , provided $MR \geq MC$ (where MC could be the seller's own value c).



Remark (Information Rent is the Inframarginal Loss).

Why does the inverse hazard rate $\frac{1-F(p)}{f(p)}$ appear in both mechanism design and standard monopoly pricing?

In intermediate microeconomics, to sell one additional unit (i.e., to induce a buyer with a slightly lower valuation to buy), the monopolist must lower the price. However, because price discrimination is impossible, the monopolist must lower the price not just for this marginal buyer, but for *all inframarginal buyers* who would have been willing to pay a higher price. This loss in revenue on the inframarginal units drives a wedge between the price p and the marginal revenue MR .

In mechanism design, this exact same dynamic is what we call the **Information Rent**. To convince a high-type buyer to reveal their true valuation rather than mimicking a lower type, the mechanism designer must leave them some surplus. The “loss of revenue on inframarginal buyers” in the monopoly model is mathematically identical to the “information rent paid to higher types” to satisfy the Incentive Compatibility (IC) constraint!

The story so far has developed the apparatus for designing mechanisms that maximize the seller's expected *revenue*: the Revelation Principle reduces the search to direct mechanisms, IC + IR pin down each bidder's interim utility up to a constant, and pointwise maximization of virtual values delivers the optimal allocation rule with reserve $r^* = \phi^{-1}(0)$. The construction is asymmetric—the seller exploits the bidders' private information for her own gain, paying out information rent only as much as IC strictly requires.

A complementary question takes the social planner's perspective: how should the object be allocated to maximize *efficiency* (total surplus), regardless of how the surplus is split between seller and buyers? The same Revelation Principle still applies, but the design

objective changes. The remainder of this chapter introduces the **Vickrey-Clarke-Groves (VCG) mechanism**, which solves the efficient design problem and reveals a striking parallel to the optimal auction: while the optimal mechanism uses virtual values to extract revenue, VCG uses *actual* values to maximize surplus, and each bidder is charged the externality her presence imposes on the others. The second-price auction of the previous chapter is recovered as the special case of VCG with a single object.

5.3 The Vickrey-Clarke-Groves (VCG) Mechanism

Motivating Example: Public Good Provision

Suppose there are N people in a society. Each person has a private value $x_i \in [\alpha, \omega]$ for a new public project (e.g., building a bridge), where α is the minimum possible value (might be negative) and ω is the maximum. The cost of building the bridge is c .

From a social planner's perspective, the efficient decision is simple:

- Build the bridge if $\sum_{i=1}^N x_i \geq c$.
- Do not build it if $\sum_{i=1}^N x_i < c$.

However, the planner does not know the true x_i . If the planner just asks people for their values, they have an incentive to exaggerate (if they want the bridge but don't pay) or understate (if they want to freeride). We need a mechanism to elicit the truth.

5.3.1 The General VCG Framework

Consider a general direct mechanism (Q, M) .

Definition 5.9: Efficient Allocation Rule

An allocation rule Q^* is *efficient* if it maximizes the total reported social surplus:

$$Q^*(x) \in \arg \max_Q \sum_{j=1}^N Q_j(x) \cdot x_j,$$

subject to feasibility constraints.

- If the object is a private good, feasibility requires $\sum_{j=1}^N Q_j(x) \leq 1$.
- If it is a pure public good, then everyone consumes it equally:

$$Q_1(x) = \dots = Q_N(x) \in \{0, 1\}.$$

Definition 5.10: Social Surplus and VCG Payment Rule

Suppose Q^* is an efficient allocation rule. We define two surplus functions based on the reported types x :

- **Total Maximized Surplus:**

$$W(x) = \sum_{j=1}^N Q_j^*(x) \cdot x_j.$$

- **Surplus of Players Other than i :**

$$W_{-i}(x) = \sum_{j \neq i} Q_j^*(x) \cdot x_j.$$

The **VCG Mechanism** pairs the efficient allocation Q^* with the following payment rule M_i^* :

$$M_i^*(x) = W(\alpha_i, x_{-i}) - W_{-i}(x)$$

where α_i is the lowest possible type.

Remark.

- **Intuition: Pricing the Externality**

The VCG payment is exactly the **social externality** that player i imposes on the rest of the society.

- $W(\alpha_i, x_{-i})$ is the maximum surplus the *other* players could have achieved if player i simply did not exist (or reported the lowest type α_i).
- $W_{-i}(x)$ is the surplus the *other* players actually end up getting because player i is present and changed the allocation to $Q^*(x)$.

The difference between what others *could have had* and what they *actually have* is the harm (externality) i causes. VCG forces player i to internalize this exact harm by making them pay it out of pocket.

- **Connection to SPA:**

The Second-Price Auction (SPA) is just a special case of VCG. If you win an SPA, your presence took the item away from the second-highest bidder. The value they lost is exactly the second-highest bid. Therefore, your externality on society is the second-highest bid, which is exactly what VCG/SPA makes you pay.

Example (VCG Reduces to SPA: Single Object, Three Bidders).

Suppose there are three bidders for a single object with true valuations $x = (5, 7, 4)$. Let's calculate the VCG payment for the winning bidder (Bidder 2, with value 7):

- If Bidder 2 is absent ($\alpha_2 = 0$), the object goes to Bidder 1 (value 5). The maximum surplus of the others is $W(\alpha_2, x_{-2}) = 5$.
- When Bidder 2 is present, the efficient rule allocates the object to Bidder 2. The surplus realized by the *other* players (Bidder 1 and 3) under this allocation is exactly zero: $W_{-2}(x) = 0$.
- Bidder 2's payment is $M_2^*(x) = W(\alpha_2, x_{-2}) - W_{-2}(x) = 5 - 0 = 5$.

Bidder 2 pays exactly 5, which is the externality they imposed on Bidder 1 by taking the object away.

- **What role does the first term $W(\alpha_i, x_{-i})$ play?**

In the general VCG framework, the first term can be *any* arbitrary function $h_i(x_{-i})$ without violating IC. Because the first term depends **ONLY** on the reports of others (x_{-i}) and a constant (α_i), it acts as a constant shift in player i 's optimization problem. When player i takes the derivative to maximize their payoff, this term vanishes!

However, specifically choosing $h_i(x_{-i}) = W(\alpha_i, x_{-i})$, the maximum social surplus achievable *without* player i , is known as the *Clarke Pivot Rule*. This specific choice analytically guarantees two crucial properties:

- **No Deficit ($M_i \geq 0$):** No one is paid just to participate.

Proof: $M_i^*(x) = W(\alpha_i, x_{-i}) - \sum_{j \neq i} Q_j^*(x)x_j$. The first term is the theoretical maximum surplus the others could achieve on their own. The second term is what they *actually* achieve when i is present. Since the theoretical unconstrained maximum must be greater than or equal to any realized value under constraints, $M_i^* \geq 0$.

- **Individual Rationality ($U_i \geq 0$):** No one gets a negative net utility from participating.

Proof: Player i 's net utility is $U_i(x) = Q_i^*(x)x_i - M_i^*(x)$. Substituting the VCG payment rule, we can rewrite this as:

$$U_i(x) = W(x) - W(\alpha_i, x_{-i})$$

Here, $W(x)$ is the maximum social pie *with* player i , and $W(\alpha_i, x_{-i})$ is the maximum pie *without* player i . Because the social planner could always choose to ignore player i and replicate the “without i ” outcome, adding a player can never shrink the maximum possible social pie. Thus, $W(x) \geq W(\alpha_i, x_{-i})$, which guarantees $U_i(x) \geq 0$.

In short, you only pay if your presence “pivots” the final allocation, but your payment will never exceed the value you personally gain!

5.3.2 Truth-telling in VCG

Proposition 5.11: Truth-telling is a Weakly Dominant Strategy

In the VCG mechanism, reporting the true valuation x_i is a weakly dominant strategy for every player i .

Proof for Proposition.

Let x_i be player i 's true valuation. Suppose player i reports z_i , while the other players report some arbitrary profile x_{-i} .

Notice that we do NOT assume x_{-i} are truthful. They can be any arbitrary lies. We will show $z_i = x_i$ is optimal regardless of what x_{-i} is. This is the definition of a dominant strategy, which is strictly stronger than a Bayesian Nash Equilibrium.

Player i 's true payoff given these reports is:

$$\begin{aligned}\pi_i(z_i|x_i, x_{-i}) &= Q_i^*(z_i, x_{-i})x_i - M_i^*(z_i, x_{-i}) \\ &= Q_i^*(z_i, x_{-i})x_i + W_{-i}(z_i, x_{-i}) - W(\alpha_i, x_{-i}) \\ &= \left[Q_i^*(z_i, x_{-i})x_i + \sum_{j \neq i} Q_j^*(z_i, x_{-i})x_j \right] - W(\alpha_i, x_{-i})\end{aligned}$$

The term in the brackets is exactly the total social surplus evaluated using player i 's true type x_i and the others' reported types x_{-i} , under the allocation rule $Q^*(z_i, x_{-i})$.

By definition, the efficient allocation rule Q^* is the mathematical operator that maximizes the sum of the inputs it is given. Therefore, to make the bracketed term as large as mathematically possible, player i should hand Q^* their true type x_i .

Formally, for any lie z_i :

$$\left[Q_i^*(z_i, x_{-i})x_i + \sum_{j \neq i} Q_j^*(z_i, x_{-i})x_j \right] \leq \max_Q \left[Q_i(x_i, x_{-i})x_i + \sum_{j \neq i} Q_j(x_i, x_{-i})x_j \right]$$

The right-hand side is achieved perfectly by reporting $z_i = x_i$. The final term $W(\alpha_i, x_{-i})$ is completely independent of z_i , so it cannot be manipulated. Thus, truth-telling maximizes i 's payoff for any x_{-i} . ■

Corollary 5.12

(Q^*, M^*) is incentive compatible.

Example (Multi-Unit VCG: Two Identical Items, Three Bidders).

Two identical units of an object are auctioned to three bidders, each of whom wants at most one unit. Their (truthful) valuations are

$$v_1 = 8, \quad v_2 = 5, \quad v_3 = 3.$$

The efficient allocation Q^* assigns one unit to each of the two highest-valued bidders (1 and 2); bidder 3 gets nothing. Total surplus is $W(x) = 8 + 5 = 13$.

To compute payments, calculate the maximum surplus achievable *without* each bidder:

$$W(\alpha_1, x_{-1}) = v_2 + v_3 = 8, \quad W(\alpha_2, x_{-2}) = v_1 + v_3 = 11, \quad W(\alpha_3, x_{-3}) = v_1 + v_2 = 13.$$

Under the Clarke pivot rule, $M_i^* = W(\alpha_i, x_{-i}) - W_{-i}(x)$ where $W_{-i}(x)$ is the surplus the *other* players actually obtain under Q^* :

$$W_{-1}(x) = v_2 = 5, \quad W_{-2}(x) = v_1 = 8, \quad W_{-3}(x) = v_1 + v_2 = 13.$$

Hence

$$M_1^* = 8 - 5 = 3, \quad M_2^* = 11 - 8 = 3, \quad M_3^* = 13 - 13 = 0.$$

Both winners pay 3, the third-highest valuation. This recovers the classic **uniform-price (Vickrey) outcome**: in a k -unit auction, the k winners each pay the $(k + 1)$ -th highest valuation.

Truth-telling check (bidder 1): reporting $\hat{v}_1 = 4 < v_2$ would lose the unit; net utility $0 < 8 - 3 = 5$. Reporting $\hat{v}_1 = 10 > v_1$ keeps the allocation unchanged; payment is still $W(\alpha_1, x_{-1}) - W_{-1}(x) = 3$ (depends only on others' reports), so net utility unchanged at 5. No deviation strictly improves bidder 1's payoff, illustrating that VCG truth-telling extends from the single-object SPA case to multi-unit and combinatorial settings.

5.3.3 VCG Maximizes Revenue Among Efficient Mechanisms

Let the interim expected payoff (indirect utility) of player i in the VCG mechanism be:

$$U_i^*(x_i) = \int [Q_i^*(x_i, x_{-i})x_i - M_i^*(x_i, x_{-i})] f_{-i}(x_{-i}) dx_{-i} \equiv q_i^*(x_i)x_i - m_i^*(x_i).$$

Recall that IC implies $U_i^*(x_i)$ is convex and strictly increasing.

Suppose (Q^*, \bar{M}) is *another* efficient mechanism that is IC. Let \bar{U}_i be its indirect utility function. Because the allocation rule Q^* is identical, the derivative $\bar{U}_i'(x_i) = q_i^*(x_i)$ must also be identical. Therefore, $\bar{U}_i(x_i)$ must have the *exact same shape* as $U_i^*(x_i)$, differing at most by a constant.

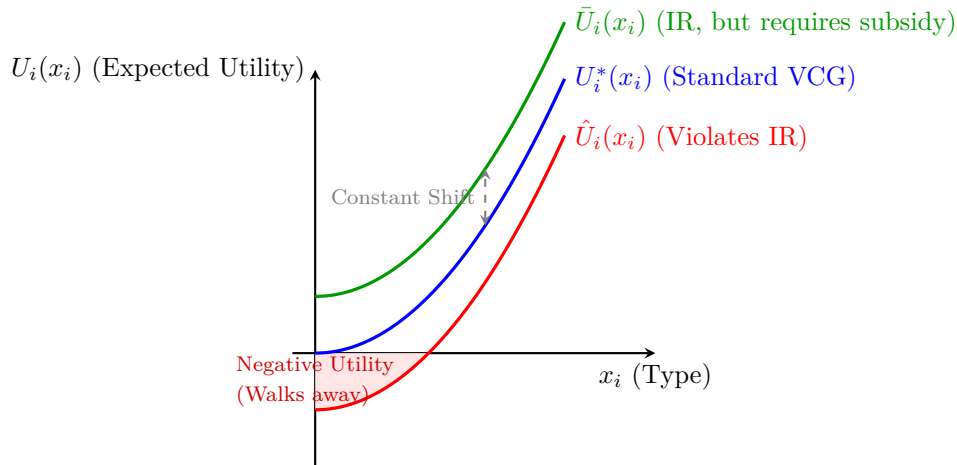
Note that in the definition of the Clarke Pivot Rule payment, we carefully included the constant term $W_{-i}(\alpha_i, x_{-i})$ to guarantee that the *lowest type gets exactly zero surplus*: $U_i^*(\alpha_i) = 0$.

Claim: The Individual Rationality (IR) Boundary

For any IC and efficient mechanism to satisfy the IR constraint for all possible types, its utility curve must lie weakly above the standard VCG utility curve.

We have established that any IC and efficient mechanism yields a utility curve with the exact same shape as the VCG utility curve. By design (the Clarke Pivot Rule), the standard VCG utility curve passes through the origin ($U_i^*(\alpha_i) = 0$).

Any alternative mechanism's utility curve \bar{U}_i that is shifted downwards (lying strictly below U_i^*) would yield negative utility for the lowest types, directly violating the IR constraint. Conversely, any utility curve shifted upwards (lying strictly above U_i^*) perfectly satisfies IR, but implicitly requires the mechanism designer to pay an upfront lump-sum subsidy to the participants.



Remark (Why do we ask for IR?).

Could there be a mechanism that is not IR for some regions, but is still reasonable? Yes! It entirely depends on the socio-economic context. Note that in standard mechanism design, we typically assume “interim” decision-making, where a bidder evaluates IR *after* knowing their own true type x_i , but before knowing others’.

- **Voluntary Participation (Default setting: Auctions, Private Markets):** Here, IR is strictly required for *all* possible types. If $U_i(x_i) < 0$, a bidder of type x_i will simply refuse to participate. Because the designer must ensure the mechanism works regardless of the realized types, the mechanism must universally guarantee $U_i(x_i) \geq 0$.
- **Mandatory Participation (Taxes, Public Goods):** If a government is building a public good or regulating a congestion problem, it has coercive taxing power. You cannot “opt out” of a society just because the mechanism yields a negative payoff for your specific type. In such compulsory environments, IR is *not* a strict constraint.

Proposition 5.13: VCG Maximizes Revenue Among Efficient Mechanisms

Among all efficient, IC, and IR mechanisms, the standard VCG mechanism maximizes the seller's total revenue $\sum_{i=1}^n m_i$.

Proof for Proposition.

From the previous claim, any alternative efficient, IC, and IR mechanism must provide a utility curve that lies weakly above the standard VCG curve. Providing strictly higher utility to the buyers mathematically necessitates extracting strictly less payment from

them. Thus, any deviation from the standard VCG payment rule (such as adding a lump-sum subsidy) strictly decreases the seller's expected revenue. ■

5.3.4 Budget Balance and the AGV Mechanism

We formally define two expected payment concepts:

Definition 5.14: Interim and Ex-ante Expected Payments

Consider a direct mechanism (Q, M) in an environment with n agents, where $M_i(x)$ is the *ex-post* payment exacted from agent i when the reported type profile is $x = (x_i, x_{-i})$. Assume that each agent's type x_j is drawn independently from a type space X_j according to a density function f_j . Let $f_{-i}(x_{-i}) = \prod_{j \neq i} f_j(x_j)$ denote the joint density of the other agents' types.

Define two expected payment concepts corresponding to different information stages:

- **Interim Expected Payment:** The expected payment of agent i evaluated at the interim stage—conditional on knowing their own realized type x_i , but prior to observing the types of other agents. It is defined as:

$$m_i(x_i) = \mathbb{E}_{x_{-i}}[M_i(x_i, x_{-i})] = \int_{X_{-i}} M_i(x_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i}$$

- **Ex-ante Expected Payment:** The expected payment of agent i evaluated at the ex-ante stage—before any agent's type is realized. It is the unconditional expectation of the payment rule over the entire type profile:

$$m_i = \mathbb{E}_{x_i}[m_i(x_i)] = \int_{X_i} m_i(x_i) f_i(x_i) dx_i = \mathbb{E}_x[M_i(x)]$$

While VCG is elegantly efficient and IC, it is notoriously bad at balancing the budget. It often runs at a deficit (e.g., the mechanism has to inject outside money to pay for positive externalities).

Example (VCG Deficit in the Bridge Problem).

Consider the 3-person bridge problem with cost $c = 10$ and values $x = (5, 6, 2)$. The sum of values is $13 \geq 10$, so the bridge is built. Let's calculate the VCG payments:

- Player 1 ($x_1 = 5$): Without P1, others value it at $8 < 10$. Not built. P1's payment is $M_1^* = 0 - (6 + 2 - 10) = 2$.
- Player 2 ($x_2 = 6$): Without P2, others value it at $7 < 10$. Not built. P2's payment is $M_2^* = 0 - (5 + 2 - 10) = 3$.
- Player 3 ($x_3 = 2$): Without P3, others value it at $11 > 10$. Built anyway. P3's payment is $M_3^* = (5 + 6 - 10) - (5 + 6 - 10) = 0$.

The total payment collected from all players is $2 + 3 + 0 = 5$. However, the physical cost to build the bridge is 10. The central planner (or the mechanism) must absorb a massive deficit of 5 to execute this socially optimal decision!

Definition 5.15: Budget Balanced

A mechanism is *Budget Balanced (BB)* if, for all realized profiles x :

$$\sum_{i=1}^n M_i(x) = 0.$$

This is a very strong *ex-post* condition: the agents simply transfer money among themselves under the mechanism.

Naturally, a subsequent question is: **Does there exist a mechanism that is efficient (Q^*), IC, IR, and BB?**

Theorem 5.16: Arrow-d'Aspremont-Gerard-Varet (AGV)

There exists an efficient, IC, IR, and BB mechanism if and only if the standard VCG mechanism runs at an ex-ante surplus. That is:

$$\sum_{i=1}^n m_i^* \geq 0.$$

Proof for Theorem

• \implies :

It is straightforward that if VCG runs at an ex-ante deficit ($\sum m_i^* < 0$), no other mechanism can balance the budget. As discussed, any other IC and IR mechanism must provide a utility curve equal to or above the VCG's curve. Giving everyone more utility means the mechanism must extract even less payment. Thus, the deficit would only widen.

• \impliedby :

We prove the other direction constructively. This is the celebrated AGV mechanism. Define a new ex-post payment rule $\bar{M}_i(x)$ as:

$$\bar{M}_i(x) = \underbrace{\left[m_i^*(x_i) - \frac{1}{n-1} \sum_{j \neq i} m_j^*(x_j) \right]}_{\text{Zero-sum transfers}} + \underbrace{\left[\frac{1}{n-1} \sum_{j \neq i} m_j^* - \frac{1}{n} \sum_{j=1}^n m_j^* \right]}_{\text{Constant adjustment}}$$

Summing the payments across all n players:

$$\begin{aligned} \sum_{i=1}^n \bar{M}_i(x) &= \left\{ \sum_{i=1}^n m_i^*(x_i) - \sum_{i=1}^n \left(\frac{1}{n-1} \sum_{j \neq i} m_j^*(x_j) \right) \right\} \\ &\quad + \left\{ \sum_{i=1}^n \left(\frac{1}{n-1} \sum_{j \neq i} m_j^* \right) - \sum_{i=1}^n \left(\frac{1}{n} \sum_{j=1}^n m_j^* \right) \right\} \end{aligned}$$

Notice that $\sum_{i=1}^n \sum_{j \neq i} m_j^*(x_j)$ is simply summing every player's interim payment ex-

actly $n - 1$ times. Divided by $n - 1$, it precisely cancels out the first term $\sum m_i^*(x_i)$. The first curly brace is strictly 0. The same combinatorial logic applies to the constants in the second curly brace. Thus, $\sum \bar{M}_i(x) = 0$. The budget balances ex-post.

To see if this new payment rule alters bidding incentives, we must find the new interim expected payment $\bar{m}_i(x_i)$. We take the expectation of $\bar{M}_i(x)$ over all other players' types x_{-i} :

$$\bar{m}_i(x_i) = \mathbb{E}_{x_{-i}}[\bar{M}_i(x)]$$

Let's pass the expectation operator through the terms linearly:

1. $\mathbb{E}_{x_{-i}}[m_i^*(x_i)] = m_i^*(x_i)$ (since it is already an interim function depending only on x_i).
2. $\mathbb{E}_{x_{-i}}\left[\frac{1}{n-1} \sum_{j \neq i} m_j^*(x_j)\right] = \frac{1}{n-1} \sum_{j \neq i} \mathbb{E}_{x_j}[m_j^*(x_j)] = \frac{1}{n-1} \sum_{j \neq i} m_j^*$.
3. The terms in the second bracket are already absolute constants, so their expectation is simply themselves.

Substituting these back in:

$$\begin{aligned} \bar{m}_i(x_i) &= m_i^*(x_i) - \frac{1}{n-1} \sum_{j \neq i} m_j^* + \frac{1}{n-1} \sum_{j \neq i} m_j^* - \frac{1}{n} \sum_{j=1}^n m_j^* \\ &= m_i^*(x_i) - \frac{1}{n} \sum_{j=1}^n m_j^* \end{aligned}$$

Note that $\bar{m}_i(x_i)$ only deviates from the standard VCG interim payment $m_i^*(x_i)$ by an absolute constant $(\frac{1}{n} \sum m_j^*)$. Because the slope of the payment function is unchanged, (Q^*, \bar{M}) is perfectly IC.

Furthermore, since the theorem condition explicitly requires the ex-ante surplus to be non-negative ($\sum_{j=1}^n m_j^* \geq 0$), the constant term we are subtracting is positive. This means the overall payment is reduced (or at worst, unchanged) for every player: $\bar{m}_i(x_i) \leq m_i^*(x_i)$. Consequently, players are receiving weakly more utility than in the VCG mechanism, trivially preserving IR. ■

5.3.5 Bilateral Trade Impossibility (Myerson-Satterthwaite)

The AGV construction balances the budget by reshuffling money among $n \geq 2$ buyers, drawing on the fact that the VCG mechanism in such settings typically runs an ex-ante surplus that the designer can simply redistribute. The picture changes dramatically when the two parties to a trade are a buyer *and* a seller, both of whom hold private information. The seminal result below shows that no mechanism can simultaneously be efficient, IC, IR, and budget-balanced—a sharp impossibility theorem with deep implications for market design.

Bilateral Trade Setup

A seller has a privately known cost $C \in [0, 1]$ of producing a single indivisible good; a buyer has a privately known value $V \in [0, 1]$ of consuming it. C and V are independently distributed with full support, common-knowledge priors. A mechanism specifies (i) whether trade occurs, (ii) the amount P the buyer pays, and (iii) the amount R the seller receives. Budget balance requires $P = R$ (no outside funding). A mechanism is **efficient** if trade occurs whenever $V > C$.

Theorem 5.17: Bilateral Trade Impossibility (Myerson-Satterthwaite, 1983)

In the bilateral trade environment, no mechanism is simultaneously efficient, IC, IR, and budget-balanced.

Proof for Theorem

We use the VCG mechanism as a benchmark, paralleling the role it played in the AGV proof.

Step 1: VCG runs a deficit. The VCG mechanism here is the **double Vickrey**: the buyer reports V , the seller reports C ; trade occurs iff $V > C$, in which case the buyer pays C and the seller receives V . Truth-telling is weakly dominant for both sides (standard VCG argument: the buyer's report only affects whether trade happens, and the truthful threshold $V > C$ is exactly her surplus-maximizing condition; symmetric for the seller). The mechanism is IR: a buyer with $V = 0$ never trades and gets 0; any $V > 0$ generates non-negative surplus since the buyer pays $C \leq V$. Symmetrically for the seller at $C = 1$. But on every trading realization $V > C$, the seller receives *more* than the buyer pays:

$$R - P = V - C > 0,$$

so the mechanism runs a strict deficit equal to the realized gains from trade. The ex-ante deficit is the expected gains from trade $\mathbb{E}[(V - C)^+]$, which is strictly positive under the full-support assumption.

Step 2: Every other efficient IC IR mechanism also runs a deficit. Consider any alternative efficient and IC mechanism. By revenue equivalence (applied to the buyer's side as in Section 7.2), the buyer's expected payment under any such mechanism differs from the VCG benchmark by a buyer-side constant K : $E[P_{\text{alt}}(V)] = E[P_{\text{VCG}}(V)] + K$. Symmetrically, the seller's expected receipts differ from VCG by a seller-side constant L : $E[R_{\text{alt}}(C)] = E[R_{\text{VCG}}(C)] + L$.

IR for the buyer at $V = 0$ requires expected payoff ≥ 0 . Since the VCG buyer at $V = 0$ never trades and pays nothing (her IR slack is exactly zero), any IR-preserving alternative must have $K \leq 0$. Symmetrically, the seller's IR at $C = 1$ requires $L \geq 0$.

The expected deficit of the alternative mechanism is

$$E[R_{\text{alt}} - P_{\text{alt}}] = E[R_{\text{VCG}} - P_{\text{VCG}}] + (L - K) > 0,$$

since the VCG deficit is strictly positive and $L - K \geq 0$. So every efficient, IC, IR mechanism inherits VCG's strict deficit, with the deficit possibly even larger. Budget balance ($E[R_{\text{alt}} - P_{\text{alt}}] = 0$) is therefore unattainable. ■

Remark (Why Bilateral Trade Differs from Multi-Buyer Auctions).

The contrast with AGV is illuminating. In a multi-buyer environment, the VCG mechanism's payments come *from* the buyers *to* the seller, accumulating an ex-ante surplus that the designer redistributes. In bilateral trade, both sides have IR constraints binding at opposite ends of the type space (buyer at $V = 0$, seller at $C = 1$), and the mechanism must compensate both: each side's information rent comes out of the social pie, and the two rents jointly exceed the pie itself when types are close to indifferent. Concretely, the VCG payment scheme awards the buyer surplus $V - C$ (she pays only the seller's reported cost, not her own value) and simultaneously awards the seller surplus $V - C$ (he receives the buyer's reported value, not his own cost), so the same pie is paid out twice. No reshuffling can fix this because there are no third parties with budget to draw from.

Remark (Implications for Market Design).

Myerson-Satterthwaite is one of the most consequential negative results in economic theory. It rationalizes a number of empirical observations:

- *Why bargaining breaks down.* When both sides hold private information, even patient and well-intentioned negotiators sometimes fail to consummate mutually beneficial trade. The theorem says this is not a coordination failure—it is a fundamental incentive incompatibility.
- *Why intermediaries exist.* Brokers, market-makers, and exchange platforms are not just middlemen who skim a fee; they are mechanisms for absorbing the unavoidable deficit in efficient bilateral trade. The buyer pays *more* than the seller receives, and the spread funds the inefficiency.
- *Why second-best mechanisms emerge.* Real markets do not aim for full efficiency. Posted-price markets, double auctions, and ascending bilateral negotiations all sacrifice some efficient trades (the ones at small surplus) to break even. The Myerson-Satterthwaite frontier characterizes how much efficiency must be sacrificed.

The same impossibility logic generalizes to any setting where two privately-informed parties must voluntarily transact under budget balance: spectrum reallocation between incumbents and new entrants, water rights between farmers, and labor contracts where both sides have private information about productivity and outside options.

Remark (Chapter Summary).

Mechanism design generalizes auction theory to the question, “what is the best protocol the designer can implement, given the players’ incentives?” The *revelation principle* (Theorem 5.1) reduces the search to direct mechanisms in which agents truthfully report their types. Within direct mechanisms, the *envelope theorem* pins down expected payments as a function of allocation, so the design problem becomes one of choosing an allocation rule alone. Two paradigmatic mechanisms anchor the field. *VCG* (Definition 5.3.1) implements the efficient allocation in dominant strategies but fails budget balance in general. *AGV* restores budget balance via expected externality payments but only achieves Bayesian incentive compatibility, not dominant strategies. The trade-off is sharp and unavoidable: the *Myerson-Satterthwaite theorem* establishes that no mechanism can simultaneously be efficient, individually rational, and budget-balanced when both sides hold private information—an impossibility that explains why bargaining breaks down, why intermediaries exist, and why real-world markets settle for second-best mechanisms like double auctions and posted prices.

Part IV

Matching

Chapter 6

Matching Theory

6.1 The Housing Market: Allocation with Endowments

We now move from settings where monetary transfers are available (auctions, mechanisms with payments) to *allocation problems without money*. No taxes, no subsidies, no side payments—agents simply trade indivisible objects among themselves. This is the classic *housing market* model introduced by Shapley and Scarf (1974).

6.1.1 Housing Market and the Core

Definition 6.1: Housing Market

A *housing market* is a triple (\mathcal{P}, H, \succ) where:

- $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ is a finite set of *agents* (“people”);
- $H = \{h_1, h_2, \dots, h_n\}$ is a finite set of *houses* (“offices”), with $|H| = |\mathcal{P}|$;
- Each agent P_i is initially endowed with house h_i and has a strict preference ordering \succ_i over H .

An *allocation* is a bijection $Q : \mathcal{P} \rightarrow H$. That is, Q assigns exactly one house to each agent.

The key question is: can we find an allocation that is *stable* in the sense that no group of agents can profitably rearrange houses among themselves?

Definition 6.2: Core

An allocation Q^* is in the *core* if there is no coalition S and no allocation Q' of $\{h_i : P_i \in S\}$ to S such that $Q'(P_i) \succ_i Q^*(P_i)$ for all $P_i \in S$ and $Q'(P_j) \succ_j Q^*(P_j)$ for some $P_j \in S$.

Intuitively, Q^* is in the *core* if there is no subset $S \subseteq \mathcal{P}$ (a *blocking coalition*) such that the members of S can reallocate their initial endowments $\{h_i : P_i \in S\}$ among themselves

in a way that makes every member of S weakly better off, and at least one member strictly better off.

Remark (Why “Core” and Not Just Pareto Efficiency?).

Pareto efficiency only rules out improvements where *everyone* can be made better off (the grand coalition $S = \mathcal{P}$). The core is a much stronger requirement: it rules out deviations by *any* subset of agents, including small coalitions. In particular, individual rationality (no agent envies their own initial endowment) is a special case of the core condition with $|S| = 1$.

6.1.2 Top Trading Cycle (TTC)

The *Top Trading Cycle* algorithm, attributed to David Gale by Shapley and Scarf (1974), provides an elegant procedure to find a core allocation.

Top Trading Cycle Algorithm

Given a housing market (\mathcal{P}, H, \succ) with initial endowment $P_i \mapsto h_i$:

1. **Point:** Each remaining agent points to the agent who currently owns their most preferred house (among remaining houses).
2. **Trade:** Since the number of agents is finite and each agent points to exactly one other agent, at least one cycle must exist. Identify all cycles. Each agent in a cycle receives the house they pointed to (i.e., the house owned by the agent they pointed to).
3. **Remove:** Remove all agents and houses involved in cycles from the market.
4. **Repeat:** Return to Step 1 with the reduced market. Terminate when all agents have been assigned a house.

Example (The Office Problem).

Consider 9 people $\mathcal{P} = \{P_1, \dots, P_9\}$ and 9 offices $H = \{1, \dots, 9\}$, arranged around a corridor as follows:

1	2	3
9	4	
8	5	
7		6


Person P_i initially occupies office i . The agents' strict preference orderings (listed from most preferred to least preferred) are:

Rank	P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9
1	7	4	6	7	7	7	9	6	2
2	6	9	2	1	1	1	3	2	5
3	9	6	8	5	8	2	6	4	4
4	4	2	1	8	4	4	2	3	8
5	8	5	7	9	3	5	5	7	3
6	5	3	4	3	5	3	4	5	1
7	1	7	5	2	6	8	8	1	9
8	2	8	9	6	9	9	7	8	6
9	3	1	3	4	2	6	1	9	7

We apply TTC to the office problem defined above. The logic of each round is as follows:

- Every remaining agent *points to* the current owner of their most-preferred available office, forming a directed graph in which each node has out-degree exactly one.
- Following the arrows traces directed **chains**. Because each node has exactly one outgoing edge, every chain must eventually revisit a node—so every chain either *closes into a cycle* or *feeds into* a cycle found by another chain.
- **Cycle members** trade immediately (each receives the office of the agent they pointed to). **Tail agents** (those whose chains lead into the cycle without closing) are held over for the next round.

Round 1. All nine agents with offices $\{1, \dots, 9\}$.



Chains and cycle detection.


- Start from P_1 : $P_1 \rightarrow P_7 \rightarrow P_9 \rightarrow P_2 \rightarrow P_4 \rightarrow P_7$ — the chain revisits P_7 . **Cycle found:** $P_7 \rightarrow P_9 \rightarrow P_2 \rightarrow P_4 \rightarrow P_7$.
- P_1 (arriving at P_7), as well as P_5 (pointing directly to P_7), P_6 (pointing to P_7), P_3 and P_8 (pointing to P_6) are all tails. They do not close any new cycle.

Cycle trade:

$$P_7 \leftarrow \text{office } 9, \quad P_9 \leftarrow \text{office } 2, \quad P_2 \leftarrow \text{office } 4, \quad P_4 \leftarrow \text{office } 7.$$

Remove $\{P_2, P_4, P_7, P_9\}$ and offices $\{2, 4, 7, 9\}$. Tails $\{P_1, P_3, P_5, P_6, P_8\}$ proceed.

Round 2. Remaining agents: $\{P_1, P_3, P_5, P_6, P_8\}$; available offices: $\{1, 3, 5, 6, 8\}$.




Chains and cycle detection.

- Start from P_5 : $P_5 \rightarrow P_1 \rightarrow P_6 \rightarrow P_1$ — revisits P_1 . **Cycle found:** $P_1 \leftrightarrow P_6$.
- P_3 and P_8 both point to P_6 (into the cycle), and P_5 points to P_1 (into the cycle). All three are tails.

Cycle trade: $P_1 \leftarrow$ office 6, $P_6 \leftarrow$ office 1.

Remove $\{P_1, P_6\}$ and offices $\{1, 6\}$. Tails $\{P_3, P_5, P_8\}$ proceed.

Round 3. Remaining: $\{P_3, P_5, P_8\}$; available offices: $\{3, 5, 8\}$.



Chains and cycle detection.

- $P_3 \rightarrow P_8 \rightarrow P_3$ — revisits P_3 . **Cycle found:** $P_3 \leftrightarrow P_8$.
- $P_5 \rightarrow P_8$ feeds into the cycle; P_5 is a tail.

Cycle trade: $P_3 \leftarrow$ office 8, $P_8 \leftarrow$ office 3.

Remove $\{P_3, P_8\}$ and offices $\{3, 8\}$.

Round 4. Only P_5 remains with office 5. With a single agent, the “chain” is a trivial self-loop: $P_5 \rightarrow P_5$. The self-loop is itself a cycle, so P_5 keeps office 5.

The final TTC allocation is:

Agent	P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9
Office	6	4	8	7	5	1	9	3	2

Remark (Cycle Structure in TTC).

In each round, the pointing graph is a *functional graph*: every node has out-degree exactly one. A fundamental property of functional graphs is that each weakly connected component contains *exactly one* directed cycle, with all remaining nodes forming tails that feed into it. Consequently:

- **Existence:** At least one cycle is guaranteed to exist in every round, since the finite graph must eventually revisit a node along any directed path.
- **Multiplicity:** The number of cycles per round equals the number of weakly connected components—it may be one or many, depending on the preference profile. Uniqueness of cycles *per round* does not hold in general.
- **Uniqueness within a chain:** Following the arrows from any fixed starting node traces a single directed path, which must eventually close into *exactly one* cycle. This per-chain uniqueness is what makes TTC well-defined: each agent belongs to at most one cycle and receives a unique assignment in the round that cycle is identified.

6.1.3 Properties of TTC

TTC Implements the Core

Theorem 6.3: TTC Allocation is in the Core

The allocation produced by the TTC algorithm is in the core of the housing market. Moreover, when preferences are strict, the core contains a unique allocation, and TTC finds it.

Proof for Theorem

The intuition is as follows. Agents removed in Round 1 each obtain their overall top choice—no coalition including any of them can make them strictly better off, since they already have their first-best house. Once Round 1 agents are settled, Round 2 agents obtain their best feasible house given Round 1 assignments, and so on. Any blocking coalition would need to offer some agent a house better than what TTC gave them, but that house has already been claimed by someone removed in an earlier round who would never agree to give it up.

We proceed by induction on the round of removal. For each round r , let C_r denote the set of agents removed in round r (i.e., the union of all cycles identified and executed in that round), and let $H_r = \{\mu(i) : i \in C_r\}$ denote the corresponding set of houses allocated in round r .

Base case (Round 1). Every agent removed in Round 1 belongs to a cycle in which each member points to the owner of their globally top-ranked house. Each such agent therefore receives their first-best house. No coalition S can block this allocation: to improve upon the outcome for any Round-1 agent $i \in S$, the coalition would need to reassign i 's own top choice to i —but i already has it.

Inductive step. Suppose that for all rounds $r' < r$, the agents removed in rounds $1, \dots, r' - 1$ receive their best house within the set of houses held by agents in those same rounds, and no coalition can improve upon their assignment using only houses from $\bigcup_{r'' \leq r'} H_{r''}$. Consider agents removed in round r . Each such agent receives their most preferred house among all houses still available—i.e., not yet claimed by any agent removed in rounds $1, \dots, r - 1$.

Suppose for contradiction that some blocking coalition S exists. Let $i \in S$ be an agent who is made strictly better off under the proposed reassignment, and let h' be the house S proposes to give i , where $h' \succ_i \mu(i)$. Since $\mu(i)$ is already i 's best available house in round r , the house h' must have been removed in some earlier round $r' < r$ —it belongs to some agent $j \in C_{r'}$ who exited in round r' with $\mu(j) = h'$. For S to reassign h' to i , it must include j . But by the inductive hypothesis, $\mu(j) = h'$ is j 's best house within $H_{r'}$; giving up h' makes j strictly worse off, so j would not participate in S . Contradiction.

Hence the TTC allocation is in the core. For uniqueness under strict preferences: any core allocation must give Round-1 agents their top choice (otherwise the coalition C_1 itself is a blocking coalition), and by induction any core allocation must coincide with TTC round by round. Thus the core is a singleton, and TTC finds it. ■

Corollary 6.4: TTC is Individually Rational

The allocation produced by TTC is individually rational: every agent weakly prefers their assigned house to their own initial endowment. That is, $\mu(i) \succeq_i h_i$ for all $i \in \mathcal{P}$.

Proof for Corollary.

This follows immediately from the core property. Suppose for contradiction that some agent i strictly prefers their endowment to their TTC allocation: $h_i \succ_i \mu(i)$. Then the singleton coalition $S = \{i\}$, using the trivial allocation $Q'(i) = h_i$ (agent i keeps their own house), satisfies $Q'(i) \succ_i \mu(i)$. This means S blocks μ , contradicting the fact that μ is in the core. Hence no such i can exist. ■

TTC is Order Independent

Theorem 6.5: TTC is Order Independent (Strict Preferences)

Suppose all agents have strict preferences over objects. When there are multiple cycles in a given round, the TTC outcome does not depend on the order in which cycles are identified and removed. That is, under strict preferences, the final allocation is unique regardless of which cycle is processed first.

Proof for Theorem

The proof reduces to one key structural observation about functional graphs.

Claim

Under strict preferences, in any round of TTC, all cycles in the pointing graph are *vertex-disjoint*.

Proof for Claim.

Under strict preferences, each agent has a unique top-ranked available object, so the pointing graph is a functional graph: every node has out-degree exactly one. Suppose for contradiction that two distinct cycles \mathcal{C}_1 and \mathcal{C}_2 share a node v . Then v would need two distinct outgoing edges—one continuing around \mathcal{C}_1 and one around \mathcal{C}_2 —contradicting out-degree one. Hence all cycles are vertex-disjoint. ■

Since any two coexisting cycles \mathcal{C}_1 and \mathcal{C}_2 are vertex-disjoint, the agents and houses in \mathcal{C}_1 are entirely distinct from those in \mathcal{C}_2 . Removing \mathcal{C}_1 first therefore does not affect the membership or the pointing structure of \mathcal{C}_2 , and vice versa. More precisely:

- The agents in \mathcal{C}_2 are still present after \mathcal{C}_1 is removed, and their preferences over the remaining houses are unchanged.
- The agents in \mathcal{C}_2 still point to the same targets they did before (none of their targets were in \mathcal{C}_1), so \mathcal{C}_2 remains a valid cycle in the reduced graph.

By symmetry, removing \mathcal{C}_2 first leaves \mathcal{C}_1 intact. In either order, both cycles are eventually removed and their members receive exactly the same houses. Since the argument applies to every pair of coexisting cycles in every round, the final allocation is independent of the order in which cycles are processed. ■

Remark (Order Dependence Under Weak Preferences).


The strictness assumption embedded in our definition of a housing market is not merely a technical convenience—it is essential for TTC to be well-defined. To see what goes wrong if we relax it, suppose we allow agents to be *indifferent* between objects (i.e., \succ_i is replaced by a weak order \succeq_i). Then an agent with multiple top-ranked objects has out-degree greater than one in the pointing graph, vertex-disjointness of cycles can fail, and different cycle-processing orders may yield different allocations.

Counterexample. Consider three agents with three houses:

agent 1 owns A , agent 2 owns B , agent 3 owns C ,


with weak preferences $1 : B \sim C \succ A$, $2 : A \succ B \succ C$, $3 : A \succ C \succ B$.

Pointing graph. Agent 1 is indifferent between B and C and simultaneously points to both agents 2 and 3. Agents 2 and 3 each point to agent 1. The resulting graph has out-degree 2 at node 1, and contains two cycles sharing that node:



The two cycles $\mathcal{C}_1 = (1 \rightarrow 2 \rightarrow 1)$ and $\mathcal{C}_2 = (1 \rightarrow 3 \rightarrow 1)$ are *not* vertex-disjoint. Processing

them in different orders gives different outcomes:



The two orders produce distinct allocations, confirming that TTC is not well-defined under weak preferences without an additional tie-breaking rule.

The standard remedy is to impose an exogenous **tie-breaking rule** (e.g., a lottery over objects) before running TTC, converting weak preferences into strict ones. This restores out-degree one, vertex-disjoint cycles, and a well-defined outcome. However, the final allocation then depends on the tie-breaking rule chosen, so TTC under weak preferences is best understood as a *family* of mechanisms indexed by tie-breaking rules, rather than a single deterministic procedure.

TTC is Strategy-Proof

Theorem 6.6: TTC is Individual Strategy-Proof

Under the TTC mechanism, truthful reporting of preferences is a dominant strategy for every agent. That is, regardless of what other agents report, no agent can obtain a strictly better outcome by misreporting their preferences.

Proof for Theorem

Fix any agent i and any reports P_{-i} by all other agents. Let μ denote the TTC outcome under truthful reporting, with $\mu(i) = o^*$, and let r be the round in which i exits under truthful TTC. We show that no misreport P'_i yields i a strictly better object.

Key observation. An agent’s report determines only *whom they point to*; it does not affect whom others point to. Since all $j \neq i$ report truthfully, the pointing behavior of every $j \neq i$ depends only on j ’s own preferences and the current available set—both of which are unaffected by i ’s report.

We present two proofs, each illuminating a different aspect of the mechanism.

Proof 1 (by contradiction on the object obtained). Suppose i misreports and obtains $o' \succ_i o^*$. We derive a contradiction in two cases.

- **Case 1: o' was taken in round $r' < r$ under truth-telling.** For i to obtain o' , i must be part of the cycle containing o' in round r' under the misreport. This requires some agent $j \neq i$ to point to i in round r' , closing the cycle through i . But

j 's pointing depends only on j 's preferences and the available set at round r' . As long as i is still present at round r' , the agents and objects removed in rounds $1, \dots, r' - 1$ are exactly those from cycles not involving i —the same as under truth-telling (since those cycles form independently of i 's report). Hence the available set at round r' is identical under both runs, and so is j 's pointing. Since no j points to i in round r' under truth-telling (no cycle involving i forms before round r), no j does so under misreport either. The cycle closes without i . Contradiction.

- **Case 2: o' is available in round r under truth-telling.** Then $o' \in A_r$, the available set when i 's truthful cycle forms. But under truth-telling i points to the owner of their top choice in A_r , which would be o' if $o' \succ_i o^*$ —contradicting $\mu(i) = o^*$.

Proof 2 (by analyzing the exit round). Suppose i misreports and exits in round $r' \neq r$. We show neither case improves i 's outcome.

- **Case 1: $r' > r$ (i exits later).** In TTC, each agent exits in the round their cycle forms, receiving their most preferred object among those still available *at that round*. By construction, $A_{r'} \subseteq A_r$: the longer i stays, the more objects get claimed by earlier cycles. Hence i 's best available object weakly worsens as the round increases, and exiting later cannot improve i 's outcome.
- **Case 2: $r' < r$ (i exits earlier).** For i to exit in round $r' < r$, i must join a cycle in that round. This requires some $j \neq i$ to point toward i (to close the cycle through i). Under truth-telling, j points to i in round r' if and only if h_i is j 's top available object at round r' . Since the available set at round r' is the same under both runs (cycles before round r' do not involve i and form identically), j already points to i in round r' under truth-telling. But then j never exits before round r' , so j is still present in round r under truth-telling. This means $h_j \in A_r$, the available set when i exits truthfully. Since $\mu(i) = o^*$ is i 's top choice from A_r and $h_j \in A_r$, we have $o^* \succsim_i h_j$. Joining j 's cycle early to obtain h_j cannot improve upon o^* .

In all cases, no misreport yields i a strictly better object than o^* . Since i and P_{-i} were arbitrary, truth-telling is a dominant strategy for every agent. ■

The two proofs above capture complementary intuitions about why manipulation fails in TTC.

- **Proof 1** asks: *can i obtain a specific better object o' ?* It shows the answer is no by ruling out each possible location of o' : either o' was already taken (and i cannot intercept its cycle, since no one points back to i), or o' was available all along (and i would have gotten it honestly).
- **Proof 2** asks: *can i benefit by changing which round it exits?* It shows the answer is no in both directions. Exiting later only shrinks the available set. Exiting earlier forces i to join a chain that was already pointing toward i under truth-telling—meaning the object at the end of that chain was already in i 's truthful available set A_r , and therefore no better than o^* .

At the deepest level, both proofs rest on the same asymmetry: *i controls only whom it*

points to, not who points back. The objects reachable by i are ultimately determined by who is willing to trade with i —and that depends on others' preferences, which no misreport can change.

A natural follow-up question is whether TTC is also *group strategy-proof*: can a coalition of agents coordinate their misreports to make every member of the coalition strictly better off? The answer is **NO**—TTC is group strategy-proof. No coalition can jointly misreport preferences and make all coalition members strictly better off. This is a stronger result that follows from the fact that TTC implements the unique strict core allocation, combined with the structure of the trading cycle procedure.

Theorem 6.7: TTC is Strictly Group Strategy-Proof

TTC is group strategy-proof: there is no coalition $S \subseteq \mathcal{P}$ and no coordinated misreport P'_S such that every member of S obtains a strictly better object than under truthful reporting.

Proof for Theorem

The proof has two parts that mirror—and extend—the individual strategy-proofness argument. Individual SP shows that no single agent can infiltrate a better cycle, because they cannot control who points back to them. Group SP faces a richer challenge: even if coalition members coordinate their pointing, they still cannot control the pointing of agents outside S , and—crucially—the core uniqueness of TTC prevents them from profitably rearranging objects among themselves either.

Suppose for contradiction that S misreports P'_S and obtains allocation μ' with $\mu'(i) \succ_i \mu(i)$ for all $i \in S$, where μ is the truthful TTC outcome. Partition S 's gains by the origin of the improved objects: outside S , and inside S (redistribution).

Claim

No S -member can gain an object initially owned by an agent outside S .

Proof for Claim.

Suppose $i \in S$ gains object $o' = \mu'(i) \succ_i \mu(i)$, where o' is initially owned by some $j \notin S$. Fix the reports of all other agents at $(P'_{S \setminus \{i\}}, P_{-S})$ and apply the individual strategy-proofness argument to i alone. Regardless of what others report, i cannot benefit from misreporting; in particular, the same two-case argument applies verbatim:

- If o' was removed in an earlier round under this fixed profile, $j \notin S$ reports truthfully, so j 's pointing behavior depends only on j 's own preferences and the available set—neither of which i 's report can affect. No agent in the cycle containing o' points back to i , so i cannot join that cycle regardless of i 's misreport. Contradiction.
- If o' is still available in i 's exit round, then $o' \succ_i \mu(i)$ implies i would have pointed to j (owner of o') under truthful reporting and obtained o' honestly—contradicting $\mu(i) \neq o'$.

Hence every $i \in S$ can only gain objects initially endowed to other S -members. That is, $\{\mu'(i) : i \in S\} \subseteq \{h_i : i \in S\}$. ■

Claim

S cannot profitably redistribute its own endowments.

Proof for Claim.

By the previous claim, the objects received by S -members under μ' are drawn entirely from $\{h_i : i \in S\}$. Define an allocation ν for S by $\nu(i) = \mu'(i)$ for all $i \in S$. Then ν is a redistribution of $\{h_i : i \in S\}$ among S such that $\nu(i) \succ_i \mu(i)$ for all $i \in S$. This means S constitutes a *blocking coalition* for μ under true preferences, contradicting the fact that μ is in the strict core. ■

Both parts yield contradictions, so no coalition S can jointly misreport to make all its members strictly better off. ■

Theorem 6.8: TTC is Weakly Group Strategy-Proof

The TTC mechanism is weakly group strategy-proof: there is no coalition $S \subseteq \mathcal{P}$ and joint misreport P'_S such that the resulting outcome μ' satisfies

$$\mu'(i) \succsim_i \mu(i) \text{ for all } i \in S, \quad \text{with } \mu'(i^*) \succ_{i^*} \mu(i^*) \text{ for some } i^* \in S.$$

Proof for Theorem

Suppose for contradiction that such S and P'_S exist. Since preferences are strict, every $i \in S$ either keeps exactly the same object ($\mu'(i) = \mu(i)$) or gets a strictly better one. Partition S accordingly:

$$S^+ = \{i \in S : \mu'(i) \succ_i \mu(i)\} \neq \emptyset, \quad S^0 = \{i \in S : \mu'(i) = \mu(i)\}.$$

Run TTC under the misreport profile (P'_S, P_{-S}) . Among all S^+ -members, let r^* be the earliest round in which any S^+ -member exits, and let $i^{**} \in S^+$ be such a member. Let \mathcal{C} be the TTC cycle that i^{**} belongs to in round r^* .

In cycle \mathcal{C} , agents trade their initial endowments cyclically: writing $\mathcal{C} = (a_1 \rightarrow a_2 \rightarrow \dots \rightarrow a_k \rightarrow a_1)$, each a_l gets $h_{a_{l+1}}$ (indices mod k). In particular, $\mu'(a_l) = h_{a_{l+1}}$ for each l .

Case 1: $\mathcal{C} \subseteq S$.

All members of \mathcal{C} belong to S , so by the weak improvement assumption,

$$\mu'(a_l) = h_{a_{l+1}} \succ_{a_l} \mu(a_l) \quad \text{for all } a_l \in \mathcal{C},$$

with $h_{a_{l+1}} \succ_{i^{**}} \mu(i^{**})$ for $i^{**} \in \mathcal{C}$. The coalition \mathcal{C} can execute exactly this trade under their *true* preferences: each a_l contributes endowment h_{a_l} and receives $h_{a_{l+1}}$, a redistribution of $\{h_{a_l} : a_l \in \mathcal{C}\}$. Every member weakly benefits, and i^{**} strictly benefits. This means \mathcal{C} is a blocking coalition for μ under true preferences, contradicting the fact that μ is in the strict core.

Case 2: $\mathcal{C} \not\subseteq S$.

\mathcal{C} contains $i^{**} \in S^+$ and at least one agent $j_0 \notin S$. We consider two sub-cases based on whether the non- S members of \mathcal{C} are harmed by the trade.

Sub-case 2a: $\mu'(j) = h_{n(j)} \succ_j \mu(j)$ for every $j \in \mathcal{C} \setminus S$.

Then every agent in \mathcal{C} weakly benefits from the cycle trade (under true preferences), and i^{**} strictly benefits. Again, \mathcal{C} is a blocking coalition for μ , contradicting the strict core.

Sub-case 2b: Some $j_0 \in \mathcal{C} \setminus S$ is strictly harmed: $\mu'(j_0) \prec_{j_0} \mu(j_0)$.

Let $\mu(j_0) = h_l$ (the object j_0 receives under truthful TTC, which is l 's endowment). Since j_0 reports truthfully in the misreport run and always picks their top available object, the fact that j_0 gets $\mu'(j_0) \prec_{j_0} h_l$ at round r^* means h_l is *not available* at the start of round r^* in the misreport run. So h_l was taken in some round $r' < r^*$ by some agent $m \neq j_0$.

We now derive a contradiction by tracing who took h_l :

- $m \notin S^+$: by choice of r^* , no S^+ -member exits before round r^* . So $m \notin S^+$.
- $m \notin S^0$: if $m \in S^0$, then $\mu'(m) = \mu(m)$. In round r' , m exits and gets h_l , so $\mu'(m) = h_l$. But then $\mu(m) = h_l = \mu(j_0)$, implying $m = j_0$ (since μ is a bijection). This contradicts $m \in S^0 \subseteq S$ and $j_0 \notin S$.
- Therefore $m \notin S$, and m reports truthfully. In the truthful run, m 's exit cycle forms at round r_m (the round m exits under truth-telling). Under truthful TTC, m gets $\mu(m) \neq h_l$ (since $\mu^{-1}(h_l) = j_0 \neq m$ and μ is a bijection). So m is *not* in the cycle containing l in the truthful run: under truthful TTC, m does not point to l at the round where h_l changes hands.

So in the misreport run, the truthful agent $m \notin S$ ends up in a different cycle than in the truthful run, specifically one that includes l and yields h_l to m . Since m

always reports truthfully, m 's pointing changes only if the available set at m 's exit round changes. This change in the available set must itself be caused by some S -member's misreport altering an earlier cycle. Tracing this chain: let \mathcal{C}' be the TTC cycle in round r' that contains m (and l) in the misreport run. By the same case analysis applied to \mathcal{C}' :

- If $\mathcal{C}' \subseteq S$: a sub-coalition of S is trading endowments in \mathcal{C}' . Since $m \notin S$ is in \mathcal{C}' , this is impossible.
- If $\mathcal{C}' \not\subseteq S$: \mathcal{C}' contains some S -member. If all agents in \mathcal{C}' weakly benefit (sub-case 2a for \mathcal{C}'), then \mathcal{C}' is a blocking coalition—contradiction. If some agent in $\mathcal{C}' \setminus S$ is strictly harmed, the same tracing argument applies recursively to a strictly earlier round.

Since the rounds are finite and strictly decreasing at each recursive step, this process must terminate. The terminal case is either sub-case 2a (blocking coalition, contradiction with core) or Case 1 (coalition entirely within S , contradiction with core). Either way, we reach a contradiction.

Since all cases yield contradictions, no coalition S can achieve a weak Pareto improvement via coordinated misreporting. TTC is weakly group strategy-proof. ■

Remark (Why Weak GSP Is Harder Than Strong GSP).

The additional difficulty in the weak GSP proof relative to strong GSP comes entirely from S^0 -members. In strong GSP, every coalition member is assumed strictly better off, so individual strategy-proofness can be applied to each member to show they can only receive objects from within S (Claim 1 of the strong GSP proof). This argument fails for S^0 -members: they are indifferent between their truthful and misreport outcomes, so ISP provides no constraint on where their object originates.

S^0 -members may misreport not to benefit themselves, but to act as “facilitators”—adjusting the cycle structure in early rounds to create a path that allows S^+ -members to enter better cycles. Sub-case 2b captures exactly this scenario and requires the recursive tracing argument above: each step of the chain that S^0 -members create must eventually close into a cycle, and that cycle is forced to be a blocking coalition for μ under true preferences—which is impossible since μ is the unique strict core.

6.2 The Assignment Problem: Allocation without Endowments

The housing market model rests critically on the assumption that each agent holds an initial endowment: this is what gives IR its bite, and what makes the core a meaningful solution concept. Once we drop this assumption—handing all objects to a central planner and asking how to distribute them from scratch—both IR and the core lose their reference point. We enter the *assignment problem*, where no agent has a prior claim to any object and the relevant normative goals reduce to *Pareto efficiency* and *strategy-proofness*.

6.2.1 Serial Dictatorship (SD)

The serial dictatorship (SD) mechanism addresses this setting directly. Agents are given a predetermined priority ordering and pick their favorite available object in turn. Compared to TTC, SD trades the structural richness of cycle-based trading for sheer simplicity: there are no endowments to track, no cycles to identify, and no core to verify—just a queue and a sequence of greedy choices.

Serial Dictatorship

Fix an ordering $\sigma = (\sigma(1), \sigma(2), \dots, \sigma(n))$ of the agents. The mechanism proceeds as follows:

1. Agent $\sigma(1)$ (the “dictator”) picks their most preferred house from H .
2. Agent $\sigma(2)$ picks their most preferred house from the remaining houses.
3. Continue until all agents are assigned.

Example (Serial Dictatorship on the Office Problem).

We run SD on the office problem with the natural ordering $\sigma = (P_1, P_2, \dots, P_9)$. At each step, the active agent scans their preference list top-to-bottom and picks the highest-ranked office that is still available. Taken offices are shown in gray.

Step	Available offices	Agent picks	Gets
1	{1, 2, 3, 4, 5, 6, 7, 8, 9}	P_1 : top choice is 7 ✓	7
2	{1, 2, 3, 4, 5, 6, 8, 9}	P_2 : top choice is 4 ✓	4
3	{1, 2, 3, 5, 6, 8, 9}	P_3 : top choice is 6 ✓	6
4	{1, 2, 3, 5, 8, 9}	P_4 : 7 taken; next available: 1 ✓	1
5	{2, 3, 5, 8, 9}	P_5 : 7, 1 taken; next available: 8 ✓	8
6	{2, 3, 5, 9}	P_6 : 7, 1 taken; next available: 2 ✓	2
7	{3, 5, 9}	P_7 : top choice is 9 ✓	9
8	{3, 5}	P_8 : 6, 2, 4 taken; next available: 3 ✓	3
9	{5}	P_9 : 2 taken; next available: 5 ✓	5

The SD allocation under ordering $\sigma = (P_1, \dots, P_9)$ is:

Agent	P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9
Office	7	4	6	1	8	2	9	3	5

Compare this with the TTC allocation on the same problem:

Agent	P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9
Office	6	4	8	7	5	1	9	3	2

Several agents differ: P_1 gets office 7 under SD (their top choice) versus 6 under TTC;

P_4 gets 1 under SD versus 7 under TTC. This illustrates a key tradeoff: the agent who goes first under SD always receives their top choice, but later agents may fare worse than under TTC, where the cycle structure can give non-first agents access to offices they could not reach by waiting their turn.

6.2.2 Properties of SD

Theorem 6.9: Serial Dictatorship is Pareto Efficient

The allocation produced by SD is Pareto efficient: there is no other allocation that makes every agent weakly better off and some agent strictly better off.

Proof for Theorem

The argument is essentially one of optimality at the margin. Consider the first agent i who would need to be reassigned to achieve the supposed Pareto improvement. Agent i already holds their top choice among everything not taken by higher-priority agents—there is simply nothing better left for them. Any Pareto improvement for i is therefore impossible, and the supposed dominating allocation cannot exist.

Suppose for contradiction that some allocation μ' Pareto dominates the SD allocation μ . Let $i = \sigma(k)$ be the *first* agent in the priority ordering such that $\mu'(i) \neq \mu(i)$. Since μ' Pareto dominates μ , we have $\mu'(i) \succ_i \mu(i)$.

But $\mu(i)$ is i 's most preferred object among all objects available when it was i 's turn—that is, among all objects not already taken by $\sigma(1), \dots, \sigma(k-1)$. Since i is the first agent to differ between μ and μ' , every agent $\sigma(1), \dots, \sigma(k-1)$ receives the same object under both μ and μ' . Therefore the set of objects available to i at step k is identical under both allocations, and $\mu(i)$ is i 's top choice from that set. The existence of $\mu'(i) \succ_i \mu(i)$ with $\mu'(i)$ in the same available set is a contradiction.

*Notice that this argument does **not** require every agent to get their globally top-ranked object. It only requires that each agent gets their top choice **given the residual supply**, which is precisely what the greedy sequential structure of SD guarantees. ■*

Theorem 6.10: Serial Dictatorship is Strategy-Proof

Under SD, truthful reporting of preferences is a dominant strategy for every agent. No agent can obtain a strictly better object by misreporting their preferences, regardless of what other agents report.

Proof for Theorem

In SD, your report only affects what you pick—it has no effect whatsoever on what is available to you, since the available set is determined entirely by agents ahead of you in the queue. You always face the same menu A_k , and truth-telling simply ensures you pick your genuine favorite from it. Lying can only cause you to pick something you like less.

Fix any agent $i = \sigma(k)$ and any reports by all other agents. Suppose i misreports, yielding object o' under the misreport, compared to $o^* = \mu(i)$ under truthful reporting. We show $o' \not\sim_i o^*$.

The key observation is that, the set of objects available to i at step k is determined entirely by the choices of $\sigma(1), \dots, \sigma(k-1)$, each of whom picks their own top available object according to their own report. Since $i = \sigma(k)$ has no influence over any earlier agent's report or choice, the available set A_k at step k is **identical** whether i reports truthfully or not.

Under truthful reporting, i picks $o^* = \mu(i)$, their most preferred object in A_k (the available set at step k). Under misreport, i picks according to a false preference list, and receives some $o' \in A_k$. Since o^* is i 's true top choice from A_k and $o' \in A_k$, we have $o^* \succsim_i o'$. In particular, $o' \not\sim_i o^*$. Since the reports of other agents and the choice of $i = \sigma(k)$ were arbitrary, truth-telling is a dominant strategy. ■

6.2.3 Equivalence of Random SD and TTC with Random Endowments

The separation between the housing market (with endowments, TTC) and the assignment problem (without endowments, SD) may seem sharp. Yet there is a striking bridge between the two frameworks: once we introduce randomness into either model in the natural way, the two mechanisms become outcome-equivalent. This result, due to Abdulkadiroğlu and Sönmez (1998), reveals that the endowment structure and the priority structure are, at a deeper level, two sides of the same coin.

The key observation is that both frameworks have one remaining degree of freedom that the model leaves unspecified:

- In the *assignment problem*, objects have no owner, so the mechanism designer must choose a **priority ordering** over agents. If this choice seems arbitrary, a natural response is to randomize: draw a uniformly random ordering.
- In the *housing market*, agents must have initial endowments, but if no agent has a natural prior claim to any particular object, a natural response is again to randomize: draw a uniformly random endowment assignment, then run TTC.

Randomizing the one free parameter in each model yields two randomized mechanisms. The theorem below says they are the same mechanism.

Definition 6.11: Randomized Mechanisms

- **Random Serial Dictatorship (RSD)**: Draw a uniformly random ordering σ of agents (each of the $n!$ orderings equally likely), then run SD under σ .
- **Core from Random Endowments (CFRE)**: Draw a uniformly random assignment of objects to agents as initial endowments (each of the $n!$ bijections equally likely), then run TTC to find the core allocation.

Theorem 6.12: RSD = CFRE

RSD and CFRE induce the *same* probability distribution over allocations: for every deterministic allocation μ ,

$$\Pr_{\sigma}[\mu_{\sigma}^{\text{SD}} = \mu] = \Pr_e[\mu_e^{\text{TTC}} = \mu],$$

where σ is a uniformly random priority ordering and e is a uniformly random endowment assignment.

Proof for Theorem**Lemma 6.13: Within-Cycle Commutativity**

Fix any endowment e and let $\mu = \mu_e^{\text{TTC}}$, with TTC producing cycles C_1, C_2, \dots, C_m in successive rounds. Let $\Sigma(\mu, e)$ denote the set of all priority orderings that:

1. place all agents in C_k before all agents in C_{k+1} (for each k), and
2. within each C_k , use any permutation of its members.

Then SD under every $\sigma \in \Sigma(\mu, e)$ produces μ .

Proof for Lemma

We establish one fact about the TTC cycle structure.

Fact 6.14: Within-Cycle Self-Preference

For any cycle C_k and any two distinct members $a, b \in C_k$: $\mu(a) \succ_a \mu(b)$.

Take any $\sigma \in \Sigma(\mu, e)$ and any agent $a \in C_k$. When a 's turn comes in SD, the available objects are a subset of those available in TTC round k , for the following reason: the objects removed before round k in TTC are exactly $\bigcup_{l < k} \{\mu(x) : x \in C_l\}$, and by definition of $\Sigma(\mu, e)$, all agents in C_1, \dots, C_{k-1} pick before a . By the Fact above and induction, each such agent picks their own μ -assignment, so the objects removed before a 's turn in SD are exactly $\bigcup_{l < k} \{\mu(x) : x \in C_l\}$ plus the μ -assignments of those C_k -members who pick before a .

We verify that $\mu(a)$ is a 's top available object when a picks:

- **Objects from earlier cycles are unavailable.** All $\mu(i)$ for $i \in C_l, l < k$ have been claimed. In TTC, $\mu(a)$ is a 's top choice among all objects available in round k —this already accounts for the absence of earlier objects. So no unavailable object is preferred over $\mu(a)$.
- **$\mu(a)$ has not been taken by earlier C_k members.** By the Fact, every other $b \in C_k$ prefers $\mu(b)$ over $\mu(a)$, so no earlier C_k -member claims $\mu(a)$.
- **Objects from later cycles are weakly worse.** All objects $\mu(x)$ for $x \in C_l, l > k$ were available in TTC round k , and a chose $\mu(a)$ over all of them.

Therefore a claims $\mu(a)$ in SD. Since a was arbitrary, SD under σ produces μ . ■

It remains to show that the randomization in CFRE and RSD induce the same distribution. We do this by mapping both the valid orderings σ and the valid endowments e to a canonical block partition uniquely defined by μ and the agents' preferences.

Fix an allocation μ . Define an ordered partition $\mathcal{B} = (B_1, B_2, \dots, B_m)$ of the agents recursively as follows: Let B_1 be the minimal nonempty set of agents whose top choices among all objects are exactly $\mu(B_1)$. Let B_2 be the minimal nonempty set of agents whose top choices among the remaining objects $N \setminus \mu(B_1)$ are exactly $\mu(B_2)$, and so on. Notice that \mathcal{B} is uniquely determined independent of σ or e .

Counting Orderings for RSD:

An ordering σ yields μ under SD if and only if it respects this dependency structure. Specifically, agents in B_k must appear after all agents in B_1, \dots, B_{k-1} to ensure their preferred objects in those blocks are already taken. Because B_k is defined as the minimal closed set under top available choices, the agents within B_k can be processed in *any* internal order without disrupting the assignment μ . Thus, the total number of orderings yielding μ is exactly the number of ways to permute agents within each block:

$$|\{\sigma : \mu_\sigma^{\text{SD}} = \mu\}| = \prod_{k=1}^m |B_k|!$$

Counting Endowments for CFRE:

An endowment e yields μ under TTC if and only if for every block k , e endows the objects $\mu(B_k)$ strictly to the agents in B_k (i.e., $e(B_k) = \mu(B_k)$). To see why, consider the TTC pointer function $p(i) = e^{-1}(\mu(i))$. If $e(B_k) = \mu(B_k)$, then p forms a bijection from B_k to itself. Consequently, p perfectly decomposes into a set of disjoint cycles covering all of B_k . Since their top available choices are in $\mu(B_k)$, every agent $i \in B_k$ participates in a cycle and receives exactly $\mu(i)$.

The number of such valid endowments is simply the number of possible bijections from B_k to $\mu(B_k)$ for each block k . Therefore:

$$|\{e : \mu_e^{\text{TTC}} = \mu\}| = \prod_{k=1}^m |B_k|!$$

Since there are exactly $n!$ total orderings and $n!$ total endowments, and the sizes of the pre-images yielding μ are strictly equal, we conclude:

$$\Pr_{\sigma}[\mu_{\sigma}^{\text{SD}} = \mu] = \frac{\prod_{k=1}^m |B_k|!}{n!} = \Pr_e[\mu_e^{\text{TTC}} = \mu]$$

This holds for any deterministic allocation μ , completing the proof. ■

Remark (Proof Intuition).

The bijection has a clean interpretation. In SD, the priority ordering determines who gets served first. In TTC, the endowment determines who holds what. The bijection $\sigma \mapsto e(\sigma)$ encodes the SD priority order *into* the endowment structure by a cyclic shift: agent σ_k is endowed with the object that σ_{k-1} ends up with under SD. This creates a pointing chain in TTC that exactly mirrors the SD queue—agent σ_k points to σ_{k+1} because σ_{k+1} holds what σ_k ultimately wants. The TTC cycles that form then execute the same trades as the SD steps, in the same order.

At a higher level, both mechanisms are doing the same thing: translating an underlying permutation (of agents, or of objects) into an allocation by a greedy sequential procedure. Randomizing the permutation uniformly in either case produces the same lottery, because the two “greedy procedures” are isomorphic via the cyclic shift bijection.

The equivalence has several consequences worth noting.

- **Mechanism design flexibility.** A planner who knows only agents’ preferences—but has no natural endowment structure or priority structure—can implement the same outcome-distribution using whichever model is more convenient for the context.
- **Properties transfer.** Since RSD inherits strategy-proofness and Pareto efficiency from SD, and CFRE inherits the core property and IR from TTC, the equivalence implies that both randomized mechanisms simultaneously satisfy all four properties in expectation.
- **The endowment is not deep.** The result suggests that the distinction between “having endowments” and “not having endowments” is less fundamental than it appears—at least

once both are randomized. What matters is the underlying combinatorial structure of sequential greedy assignment, not the narrative framing of endowments versus priorities.

6.3 Two-Sided Matching

We now transition from one-sided allocation problems (housing markets) to *two-sided matching* problems, where agents on both sides of the market have preferences over agents on the other side. The canonical example is the *college admissions* problem introduced by Gale and Shapley (1962).

6.3.1 Two-Sided Matching and Pairwise Stability

Definition 6.15: Two-Sided Matching Market

A **two-sided matching market** (or *marriage model*) consists of:

- Two disjoint, finite sets of agents: $\mathcal{S} = \{s_1, \dots, s_n\}$ (students) and $\mathcal{C} = \{c_1, \dots, c_m\}$ (colleges).
- Each student $s \in \mathcal{S}$ has a strict preference ordering \succ_s over \mathcal{C} .
- Each college $c \in \mathcal{C}$ has a strict preference ordering \succ_c over \mathcal{S} .
- Each college has a **quota** $q_c \geq 1$ specifying the number of seats available. In the basic model, $q_c = 1$ for all c , and $m = n$.

Definition 6.16: Matching

A **matching** μ is a bijection $\mu : \mathcal{C} \rightarrow \mathcal{S}$ (when $|\mathcal{C}| = |\mathcal{S}|$ and all quotas are 1). That is, $\mu(c)$ denotes the student assigned to college c , and $\mu^{-1}(s)$ denotes the college assigned to student s .

In the housing market (Shapley–Scarf), only one side—the agents—had preferences; houses were passive objects. Here, *both* sides have preferences, which fundamentally changes the problem. The notion of “core” from the housing market generalizes to the concept of *stability*, but the structure of stable allocations is much richer.

The central solution concept in two-sided matching is *pairwise stability*. Unlike the core (which considers arbitrary coalitions), pairwise stability focuses on deviations by pairs (c, s) .

Definition 6.17: Blocking Pair and Pairwise Stability

A pair $(c, s) \in \mathcal{C} \times \mathcal{S}$ is a **blocking pair** for matching μ if:

1. College c prefers student s to its current match: $s \succ_c \mu(c)$;
2. Student s prefers college c to their current match: $c \succ_s \mu^{-1}(s)$.

A matching μ is **pairwise stable** if there exists no blocking pair.

Remark.

Pairwise stability captures a minimal notion of “**no justified envy**”: there is no student-college pair who would both prefer to be matched with each other over their current assignments. If such a pair existed, the college would have an incentive to admit the student, and the student would have an incentive to accept—destabilizing the matching.

6.3.2 Deferred Acceptance (DA)

Gale and Shapley (1962) introduced the *Deferred Acceptance* (DA) algorithm, which always produces a pairwise stable matching.

Student-Proposing Deferred Acceptance (DA)

1. **Propose:** Each unmatched student proposes to the highest-ranked college on their list to which they have not yet proposed.
2. **Hold/Reject:** Each college that receives proposals *tentatively* holds the most preferred applicant (according to its own ranking) among the new proposals and its currently held student (if any), and rejects all others.
3. **Repeat:** Rejected students propose to their next choice. Continue until no rejections occur (every student is either tentatively held or has been rejected by all colleges).
4. **Finalize:** All tentative matches become permanent.

Example (Student-Proposing DA on the Matching Problem).

Consider 10 students $S = \{s_1, \dots, s_{10}\}$ and 10 colleges $C = \{c_1, \dots, c_{10}\}$, each with quota $q = 1$.

College preferences over students (Final matches highlighted):

Rank	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
1	s_1	s_1	s_{10}	s_6	s_7	s_4	s_1	s_6	s_8	s_6
2	s_2	s_6	s_2	s_1	s_2	s_8	s_4	s_4	s_{10}	s_5
3	s_3	s_9	s_4	s_5	s_8	s_2	s_3	s_2	s_6	s_1
4	s_4	s_2	s_6	s_2	s_4	s_6	s_8	s_{10}	s_9	s_4
5	s_5	s_5	s_8	s_7	s_{10}	s_{10}	s_5	s_8	s_5	s_{10}
6	s_6	s_8	s_3	s_3	s_5	s_3	s_2	s_3	s_7	s_8
7	s_7	s_3	s_5	s_8	s_9	s_7	s_9	s_1	s_4	s_9
8	s_8	s_4	s_7	s_4	s_6	s_1	s_6	s_9	s_3	s_2
9	s_9	s_7	s_9	s_{10}	s_3	s_5	s_7	s_7	s_1	s_3
10	s_{10}	s_{10}	s_1	s_9	s_1	s_9	s_{10}	s_5	s_2	s_7

Student preferences over colleges (Final matches highlighted):

Rank	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
1	c_5	c_4	c_5	c_6	c_3	c_3	c_8	c_1	c_1	c_2
2	c_{10}	c_9	c_6	c_2	c_7	c_4	c_6	c_4	c_5	c_{10}
3	c_6	c_{10}	c_1	c_5	c_9	c_5	c_4	c_2	c_2	c_9
4	c_4	c_6	c_7	c_8	c_{10}	c_1	c_2	c_6	c_{10}	c_4
5	c_7	c_3	c_4	c_1	c_1	c_9	c_9	c_3	c_3	c_1
6	c_2	c_7	c_8	c_7	c_5	c_7	c_7	c_7	c_8	c_7
7	c_8	c_1	c_9	c_3	c_4	c_{10}	c_5	c_5	c_6	c_5
8	c_9	c_5	c_{10}	c_9	c_2	c_8	c_{10}	c_{10}	c_7	c_6
9	c_1	c_2	c_2	c_{10}	c_6	c_2	c_1	c_8	c_4	c_3
10	c_3	c_8	c_3	c_4	c_8	c_6	c_3	c_9	c_9	c_8

We execute the student-proposing DA algorithm. We track the process using a chalkboard-style simulation: the colleges are listed across the top, and proposing students are written below them. In each round, if a college receives multiple proposals (or a new proposal challenges a held student), the college's less preferred student is simply crossed out (~~struck through~~).

Round 1.

All 10 students propose to their first-choice colleges. Colleges c_1 , c_3 , and c_5 receive multiple proposals and must reject the less preferred students.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_8	s_{10}	s_6	s_2	s_3	s_4		s_7		
s_9		s_5		s_1					

Rejected: s_9 (next c_5), s_5 (next c_7), s_1 (next c_{10}).

Round 2.

The three rejected students propose to their next choices. s_9 challenges s_3 at c_5 . College c_5 prefers s_9 (#7) over s_3 (#9).

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_8	s_{10}	s_6	s_2	s_9	s_4	s_5	s_7		s_1
				s_3					

Rejected: s_3 (next c_6).

Round 3.

s_3 proposes to c_6 . College c_6 evaluates the held s_4 and the new s_3 . s_4 is c_6 's top pick and is retained.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_8	s_{10}	s_6	s_2	s_9	s_4	s_5	s_7		s_1
					s_3				

Rejected: s_3 (next c_1).

Round 4.

s_3 proposes to c_1 . College c_1 prefers s_3 (#3) over the currently held s_8 (#8).

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_{10}	s_6	s_2	s_9	s_4	s_5	s_7		s_1

Rejected: s_8 (next c_4).

Round 5.

s_8 proposes to c_4 . College c_4 prefers the held s_2 (#4) over s_8 (#7).

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_{10}	s_6	s_2	s_9	s_4	s_5	s_7		s_1

Rejected: s_8 (next c_2).

Round 6.

s_8 proposes to c_2 . College c_2 prefers s_8 (#6) over the held s_{10} (#10).

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_8	s_6	s_2	s_9	s_4	s_5	s_7		s_1

Rejected: s_{10} (next c_{10}).

Round 7.

s_{10} proposes to c_{10} . College c_{10} prefers the held s_1 (#3) over s_{10} (#5).

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_8	s_6	s_2	s_9	s_4	s_5	s_7		s_1

Rejected: s_{10} (next c_9).

Round 8.

s_{10} proposes to c_9 . College c_9 is empty and accepts immediately.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_8	s_6	s_2	s_9	s_4	s_5	s_7	s_{10}	s_1

No further rejections. The algorithm terminates.

The **student-optimal stable matching** μ_S is directly read from the final chalkboard state:

College	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
Student	s_3	s_8	s_6	s_2	s_9	s_4	s_5	s_7	s_{10}	s_1

6.3.3 Properties of DA

Theorem 6.18: DA is Pairwise Stable

The matching produced by the student-proposing DA algorithm is pairwise stable.

Proof for Theorem

The argument proceeds as follows. Suppose, for contradiction, that (c, s) is a blocking pair under the DA outcome μ . Then:

1. Student s prefers c to $\mu^{-1}(s)$, meaning $c \succ_s \mu^{-1}(s)$. Since DA has students propose in order of preference, s must have proposed to c at some earlier round before being eventually matched with $\mu^{-1}(s)$.
2. But s was rejected by c at that round, which means c was holding (tentatively matched with) a student it prefers to s . Since colleges only ever improve their tentative match over the course of the algorithm, c 's final match $\mu(c)$ is at least as preferred as the student it held when it rejected s . Hence $\mu(c) \succ_c s$.

This contradicts the requirement that c prefers s to $\mu(c)$. Therefore no blocking pair can exist. ■

Theorem 6.19: Strategy-Proofness for the Proposing Side

In the student-proposing DA, truthful reporting of preferences is a dominant strategy for every **student**. That is, no student can obtain a strictly better outcome by misreporting their preferences, regardless of what other students or colleges report.

Proof for Theorem

We present the proof for the general many-to-one matching market where each college c has a quota $q_c \geq 1$.

We proceed by contradiction. Suppose there exists a true preference profile P , but some student s^* submits a fake preference list P'_{s^*} . Let μ be the truthful DA outcome, and let μ' be the DA outcome under the misreported profile $P' = (P'_{s^*}, P_{-s^*})$.

Suppose the misreport is strictly profitable for s^* , meaning $\mu'(s^*) \succ_{s^*} \mu(s^*)$.

Step 1: Identify the “Improving” Coalition.

Let S' be the set of all students who are strictly better off under μ' than under μ according to their true preferences:

$$S' = \{s \in \mathcal{S} \mid \mu'(s) \succ_s \mu(s)\}$$

By our assumption, $s^* \in S'$. Let $C' = \mu'(S')$ be the set of colleges these students are matched to in μ' .

Step 2: The Rejection Implication.

For every student $s \in S'$, because $\mu'(s) \succ_s \mu(s)$, s must have proposed to $\mu'(s)$ and been rejected during the *truthful* DA. Consequently, every college $c \in C'$ must have rejected at least one student from S' in the truthful DA.

Under the rules of DA with quotas, a college only rejects an acceptable student if

its quota is completely full. This implies that in the truthful outcome μ , every college $c \in C'$ exactly filled its capacity ($|\mu(c)| = q_c$), and it strictly prefers **every single student** currently in $\mu(c)$ to the S' student it rejected.

Step 3: All Displaced Students must belong to S' .

Take any college $c \in C'$. Since $c \in C'$, there is some student $s \in S'$ such that $s \in \mu'(c)$. Let $k \in \mu(c)$ be **any** student assigned to c in the truthful DA. We claim k must also belong to S' .

Suppose for contradiction that $k \notin S'$. This means k does not strictly prefer μ' over μ . Since c must be filled with different students or drop some students to accommodate s in μ' , if k is displaced from c , k strictly prefers $\mu(k) = c$ over $\mu'(k)$. Since $k \notin S'$ but $s^* \in S'$, $k \neq s^*$. Thus, k reported truthfully.

Now consider the matching μ' under the reported profile P' . College c admits $s \in S'$ in μ' . But from Step 2, c strictly prefers k over s ($k \succ_c s$). At the same time, k strictly prefers c over $\mu'(k)$. Because k 's reported preference is their true preference, (c, k) forms a **blocking pair** for μ' under P' ! This contradicts the fact that μ' is the stable DA outcome under P' .

Hence, our claim holds: **every** student $k \in \mu(c)$ sitting in college c during the truthful DA must belong to S' .

Step 4: The Counting Contradiction.

Let us count the seats in C' . The total capacity of all colleges in C' is $\sum_{c \in C'} q_c$.

By Step 2, every college in C' is completely full under μ . By Step 3, every single student occupying these seats in μ belongs to S' . Therefore, S' contains **at least** $\sum_{c \in C'} q_c$ students.

Now look at μ' . By the definition of C' , every student in S' is assigned to a college in C' . Since the maximum number of students C' can hold is $\sum_{c \in C'} q_c$, S' can contain **at most** $\sum_{c \in C'} q_c$ students.

Thus, $|S'| = \sum_{c \in C'} q_c$. The students assigned to C' in the truthful matching μ are *exactly* the students in S' , occupying all available seats.

In μ' , this exact same group of students S' is completely reallocated among the exact same seats in C' . But by the definition of S' , **every single student in S' strictly prefers their new seat in μ' to their old seat in μ .**

This implies μ' **strictly Pareto dominates** μ for this subset of students, without affecting anyone else. However, it is a fundamental property of the student-proposing DA that its outcome (μ) is the **student-optimal stable matching** (which is Pareto efficient for the students). A strict Pareto improvement via a pure internal reallocation is mathematically impossible.

This contradiction proves that no such profitable misreport P'_{s^*} can exist. ■

Remark.

The intuition is clean: under DA, a student who misreports can only cause themselves to be rejected by colleges they would have been admitted to, or cause themselves to propose to less preferred colleges earlier. Misreporting can never lead to admission at a college that would have rejected the student under truthful reporting.

Theorem 6.20: Impossibility of Two-Sided Strategy-Proofness

No stable matching mechanism is strategy-proof for **both** sides simultaneously. In particular, under the student-proposing DA, colleges may have incentives to misreport their preferences.

Remark (How Can the Accepting Side Manipulate?).

The accepting side (colleges in student-proposing DA) can potentially manipulate by *truncating* their preference lists or *reordering* them. For instance, a college might strategically rank a less-preferred student higher to trigger a chain of rejections that ultimately results in a more favorable match. This impossibility result (Roth, 1982) highlights a fundamental tension: **any stable matching mechanism must sacrifice strategy-proofness on at least one side.**

6.3.4 Many-to-One Matching (Quotas)

So far, our two-sided matching model assumed $q_c = 1$ for all colleges, mathematically mirroring a strict one-to-one “marriage” market. In reality, the classic Gale-Shapley model was built for the *college admissions* problem, where institutions admit entire incoming classes. Generalizing from one-to-one to many-to-one matching introduces capacity constraints, but the core theoretical machinery translates beautifully.

Definition 6.21: Many-to-One Matching Market

Each college $c \in \mathcal{C}$ now has a maximum capacity or **quota** $q_c \geq 1$. A **matching** μ is a mapping such that:

- For each student s , $\mu(s) \in \mathcal{C} \cup \{\emptyset\}$ (a student attends at most one college or remains unmatched).
- For each college c , $\mu(c) \subseteq \mathcal{S}$ with $|\mu(c)| \leq q_c$ (a college admits a *set* of students up to its quota).
- $\mu(s) = c \iff s \in \mu(c)$.

In a many-to-one matching market, a college c with quota q_c ultimately admits a *set* of students. While the model specifies the college’s strict ranking over individual students (denoted by \succ_c), we must define how it ranks different incoming classes (subsets of students). To keep the model tractable and rule out complex complementarities (e.g., “I only want student A if I also get student B”), we impose a standard behavioral assumption on how colleges evaluate groups.

Assumption 6.22: Responsive Preferences

Let \succ_c^* denote college c 's preference relation over *sets* of students. We assume \succ_c^* is **responsive** to the college's preferences over individual students \succ_c . Specifically, for any set of students S strictly below capacity ($|S| < q_c$), and any two students $s, s' \notin S$:

1. **Substitution:** Replacing a student with a strictly more preferred individual always makes the college strictly better off:

$$S \cup \{s\} \succ_c^* S \cup \{s'\} \iff s \succ_c s'$$

2. **Addition:** Adding an acceptable student (one who is preferred to leaving a seat empty, $s \succ_c \emptyset$) always makes the college strictly better off:

$$S \cup \{s\} \succ_c^* S \iff s \succ_c \emptyset$$

Remark (Why do we need this assumption?).

Without responsiveness, evaluating stability becomes a combinatorial nightmare. A college might reject a globally top-ranked student simply because they do not “fit” well with the rest of the admitted cohort.

By assuming responsive preferences, we guarantee that a college's evaluation of a group is strictly and additively tied to the independent quality of its individual members. *Crucially, this mathematically justifies why we only need to check for pairwise blocking pairs.* If a college ever wants to deviate and alter its admitted class, responsiveness ensures that this desire will manifest as a simple pairwise blocking condition with a single better student, relieving us from checking complex multi-student coordinated deviations.

Definition 6.23: Pairwise Stability with Quotas

A matching μ is pairwise stable if it is individually rational and contains no **blocking pair**. A pair (c, s) blocks μ if:

1. Student s prefers c to their current match: $c \succ_s \mu(s)$;
2. **Either** c has an empty seat ($|\mu(c)| < q_c$),
Or c is full ($|\mu(c)| = q_c$) but prefers s to at least one of its currently admitted students ($s \succ_c s'$ for some $s' \in \mu(c)$).

Student-Proposing DA (with Quotas)

The DA algorithm adapts to quotas with a single, straightforward modification:

1. **Propose:** Each unassigned student proposes to their most preferred acceptable college that has not yet rejected them.

2. **Hold/Reject:** Each college c considers its new applicants alongside the students it tentatively held from the previous round. It ranks this combined pool, **tentatively holds up to q_c of its most preferred students**, and rejects the rest.
3. **Repeat** until no student is rejected.

Example (DA with Quotas).

Consider 6 students $\{s_1, \dots, s_6\}$ and 2 colleges $\{c_1, c_2\}$. College c_1 has a quota $q_1 = 2$, and college c_2 has a quota $q_2 = 3$. Notice that total demand (6 students) exceeds total supply (5 seats), meaning at least one student will be left unmatched.

College preferences (Quotas in parentheses, Final matches highlighted):

Rank	$c_1 (q = 2)$	$c_2 (q = 3)$
1	s_1	s_6
2	s_2	s_5
3	s_3	s_4
4	s_4	s_3
5	s_5	s_2
6	s_6	s_1

Student preferences (Final matches highlighted):

Rank	s_1	s_2	s_3	s_4	s_5	s_6
1	c_2	c_1	c_1	c_1	c_2	c_1
2	c_1	c_2	c_2	c_2	c_1	c_2

We run the student-proposing DA algorithm. Remember that colleges rank their combined pool of applicants and simply draw a “cutoff line” at their quota.

Round 1.

Students propose to their top choices: $s_1, s_5 \rightarrow c_2$, while $s_2, s_3, s_4, s_6 \rightarrow c_1$.

$c_1 (q = 2)$	$c_2 (q = 3)$
s_2	s_1
s_3	s_5
s_4	
s_6	

Explanation: c_1 receives 4 applicants for 2 seats. Looking at c_1 's preference list, $s_2 \succ s_3 \succ s_4 \succ s_6$. It holds the top two (s_2, s_3) and rejects the rest. c_2 only has 2 applicants for 3 seats, so it holds both.

Rejected: s_4 (next c_2), s_6 (next c_2).

Round 2.

The rejected students s_4 and s_6 propose to c_2 . c_2 's new pool is $\{s_1, s_5\}$ (held) $\cup \{s_4, s_6\}$ (new) = $\{s_1, s_4, s_5, s_6\}$.

$c_1 (q = 2)$	$c_2 (q = 3)$
s_2	s_6
s_3	s_5
	s_4
	s_1

Explanation: c_2 ranks its pool: $s_6 \succ s_5 \succ s_4 \succ s_1$. It keeps the top 3. Notice that s_1 , who was safely held in Round 1, is now bumped out by stronger applicants!

Rejected: s_1 (next c_1).

Round 3.

s_1 proposes to c_1 . c_1 's pool is now $\{s_2, s_3\}$ (held) \cup $\{s_1\}$ (new).

$c_1 (q = 2)$	$c_2 (q = 3)$
s_1	s_6
s_2	s_5
s_3	s_4

Explanation: c_1 ranks its pool: $s_1 \succ s_2 \succ s_3$. It keeps the top 2. Another devastating bump occurs: s_3 is kicked out to make room for s_1 .

Rejected: s_3 (next c_2).

Round 4.

s_3 proposes to c_2 . c_2 's pool is $\{s_4, s_5, s_6\}$ (held) \cup $\{s_3\}$ (new).

$c_1 (q = 2)$	$c_2 (q = 3)$
s_1	s_6
s_2	s_5
	s_4
	s_3

Explanation: c_2 ranks its pool: $s_6 \succ s_5 \succ s_4 \succ s_3$. The cutoff is 3, so s_3 is rejected immediately.

Rejected: s_3 (no choices remaining).

Round 5.

Student s_3 has exhausted their preference list. No unassigned student has an acceptable college left to propose to. The algorithm terminates.

The final student-optimal stable matching μ_S with quotas is:

$$\mu_S(c_1) = \{s_1, s_2\}, \quad \mu_S(c_2) = \{s_4, s_5, s_6\}, \quad \mu_S(s_3) = \emptyset$$

6.3.5 Truncated Preferences (Unacceptable Matches)

In our previous examples, we implicitly assumed that every college was acceptable to every student, and every student was acceptable to every college. In reality, matching markets

are characterized by **outside options**. A student might prefer to stay home or enter the labor market rather than attend a poorly ranked college. Similarly, a college might prefer to leave a seat empty rather than admit a student who does not meet its minimum academic standards.

We model this by introducing **truncated preferences**. An agent ranks a subset of the opposite side and deems the rest *unacceptable*.

Definition 6.24: Unacceptable Matches, Matching with Truncations and Individual Rationality

We use the symbol \emptyset to denote the outside option of remaining unmatched.

- A college c is **acceptable** to student s if $c \succ_s \emptyset$.
- A student s is **acceptable** to college c if $s \succ_c \emptyset$.

Any agent ranked below \emptyset in a preference list is strictly worse than being unmatched.

Definition 6.25: Matching with Truncations

Let $\mathcal{C} = \{c_1, \dots, c_m\}$ (with quotas q_j) and $\mathcal{S} = \{s_1, \dots, s_n\}$. Each student s has a strict preference \succ_s over \mathcal{C} , and each college c has a strict preference \succ_c over \mathcal{S} . Agents may rank only a subset of the other side as acceptable, leaving the ranking below \emptyset as unacceptable. A **matching** is a function $\mu : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$ (with the constraint that $|\{s : \mu(s) = c\}| \leq q_c$ for each c). Here $\mu(s) = \emptyset$ means student s is unmatched.

Definition 6.26: Individual Rationality

A matching μ is **Individually Rational (IR)** if no agent is matched with an unacceptable partner: $\mu(s) \succ_s \emptyset$ and $\mu(c) \succ_c \emptyset$ for all $s \in \mathcal{S}, c \in \mathcal{C}$.

Definition 6.27: Stability with Truncations

A matching μ is **stable** if:

- Individual rationality:** Every matched student prefers their match to being unmatched: if $\mu(s) \in \mathcal{C}$, then $\mu(s) \succ_s \emptyset$. Similarly, every college finds each of its matched students acceptable.
- No blocking pair:** There is no pair (s, c) with $\mu(s) \neq c$ such that $c \succ_s \mu(s)$ and either c has an unfilled seat or $s \succ_c s'$ for some $s' \in \mu(c)$.

The DA algorithm handles truncated preferences effortlessly by embedding the IR constraint directly into its rules:

1. Students never propose to unacceptable colleges (they drop out of the market once they reach \emptyset on their lists).

2. Colleges immediately reject any proposing student who is unacceptable to them, even if they have empty seats.

This slight modification is not just algorithmic; it provides a constructive proof for a fundamental existence theorem in matching theory.

Theorem 6.28: Existence of Stable Matchings with Truncations

For any two-sided matching market with strict, truncated preferences, there always exists at least one matching that is both pairwise stable and individually rational.

Proof for Theorem

The proof is constructive. We show that the Deferred Acceptance (DA) algorithm, modified with the two IR rules above, always terminates and outputs a matching that satisfies both properties.

Part 1: Individual Rationality (IR).

By the modified rules of the DA algorithm, no student ever proposes to an unacceptable college. Furthermore, no college ever tentatively holds (or permanently admits) an unacceptable student. Therefore, in the final matching μ , every student either receives an acceptable college or remains unmatched ($\mu(s) \succsim_s \emptyset$), and every college only admits acceptable students. Thus, μ is individually rational.

Part 2: Pairwise Stability.

Suppose, for contradiction, that there exists a blocking pair (c, s) in the final matching μ . By definition of a blocking pair under quotas and truncations:

1. Student s strictly prefers c over their final match: $c \succ_s \mu(s) \succsim_s \emptyset$.
2. College c strictly prefers s over \emptyset (meaning s is acceptable to c), AND either c has an empty seat ($|\mu(c)| < q_c$) or c strictly prefers s over some admitted student $s' \in \mu(c)$.

Since $c \succ_s \mu(s)$, student s must have proposed to college c at some round prior to proposing to $\mu(s)$ (or before dropping out, if $\mu(s) = \emptyset$).

Since s is not matched with c in the end, college c must have rejected s at some round. Under the DA rules, a college only rejects an *acceptable* student if it has already filled its entire quota q_c with students it strictly prefers to the rejected student.

Moreover, as the DA algorithm progresses, a college's pool of tentatively held students can only improve (a college only drops a held student if a strictly better one proposes). Therefore, in the final matching μ , college c must be exactly at full capacity ($|\mu(c)| = q_c$), and it must strictly prefer every student in $\mu(c)$ to s .

This directly contradicts the second condition of the blocking pair (that c has an empty seat or prefers s to some admitted student s').

Hence, no blocking pair can exist. The modified DA algorithm always terminates yielding an individually rational and pairwise stable matching, proving existence. ■

Example (DA with Truncated Preferences).

Consider our 10-student, 10-college market again, but now with severe truncations. The symbol \emptyset marks the end of the acceptable options. (For columns that do not explicitly show \emptyset , we assume the list naturally ends there or the unlisted options are unacceptable).

College preferences (Final matches highlighted):

Rank	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
1	s_1	s_1	s_{10}	s_6	s_7	s_4	s_1	s_6	s_8	s_6
2	s_2	s_6	s_2	s_1	s_2	s_8	s_4	s_4	s_{10}	s_5
3	s_3	s_9	s_4	s_5	s_8	s_2	s_3	s_2	s_6	s_1
4	s_4	s_2	s_6	s_2	s_4	s_6	s_8	s_{10}	s_9	s_4
5	s_5	\emptyset	s_8	s_7	\emptyset	s_{10}	s_5	s_8	s_5	s_{10}
6	s_6		s_3	s_3		s_3	s_2	s_3	s_7	s_8
7	\emptyset		s_5	s_8		s_7	s_9	s_1	s_4	s_9
8			s_7	s_4		s_1	s_6	\emptyset	s_3	s_2
9			s_9	s_{10}		s_5	s_7		s_1	s_3
10			s_1	s_9		s_9	s_{10}		s_2	s_7
11			\emptyset	\emptyset						

Student preferences (Final matches highlighted):

Rank	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
1	c_5	c_4	c_5	c_6	c_3	c_3	c_8	c_1	c_1	c_2
2	c_{10}	c_9	c_6	c_2	c_7	c_4	c_6	c_4	c_5	c_{10}
3	c_6	c_{10}	c_1	c_5	c_9	c_5	c_4	c_2	c_2	c_9
4	\emptyset	c_6	c_7	c_8	c_{10}	c_1	c_2	c_6	c_{10}	c_4
5		c_3	c_4	c_1	c_1	c_9	c_9	c_3	\emptyset	c_1
6		c_7	c_8	c_7	c_5	c_7	\emptyset	c_7		c_7
7		c_8	c_3	c_3	c_4	c_{10}		c_5		c_5
8		c_1	c_9	c_9	c_2	c_8		c_{10}		c_6
9		c_5	c_{10}	c_{10}	c_6	c_2		c_8		c_3
10		c_2	c_2		c_8			c_9		c_8

We execute the student-proposing DA. Notice how truncations cause severe “immediate rejections” even when colleges are empty.

Round 1.

Students propose to their top choices.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_8	s_{10}	s_6	s_2	s_1	s_4		s_7		
s_9		s_5		s_3					

Explanation: Look closely at the rejections. College c_5 rejects *both* s_1 and s_3 immediately, because neither student is on c_5 's acceptable list ($s_7, s_2, s_8, s_4, \emptyset$). Similarly,

c_1, c_2, c_8 reject their applicants purely because they fail to meet the individual rationality threshold.

Rejected: $s_1, s_3, s_5, s_7, s_8, s_9, s_{10}$.

Round 2.

The mass of rejected students propose to their second choices.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
		s_6	s_2	s_3	s_4	s_5			s_1
			s_8		s_3				s_{10}
					s_7				

Explanation: c_{10} receives s_1 and s_{10} ; it prefers s_1 . c_6 holds s_4 and rejects both s_3 and s_7 . c_4 rejects s_8 . c_5 continues to be empty because s_9 is unacceptable.

Rejected: $s_3, s_7, s_8, s_9, s_{10}$.

Round 3.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_9	s_6	s_2		s_4	s_5		s_{10}	s_1
	s_8		s_7						

Explanation: c_1 accepts s_3 . c_9 accepts s_{10} . c_2 receives s_9 and s_8 ; s_9 is acceptable, but s_8 is unacceptable! Thus c_2 holds s_9 and rejects s_8 .

Rejected: s_7, s_8 .

Rounds 4 & 5.

$s_7 \rightarrow c_2$ (unacceptable, rejected). $s_8 \rightarrow c_6$ (rejected, $s_4 \succ s_8$).

Next: $s_7 \rightarrow c_9$ (rejected, $s_{10} \succ s_7$). $s_8 \rightarrow c_3$ (rejected, $s_6 \succ s_8$).

Rejected: s_7, s_8 .

Round 6 (The Dropouts Begin).

Student s_7 has exhausted their list (reaching \emptyset after c_9). s_7 **drops out of the market and remains permanently unmatched.**

Meanwhile, s_8 proposes to c_7 .

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_3	s_9	s_6	s_2		s_4	s_8		s_{10}	s_1
						s_5			

Explanation: c_7 prefers the new s_8 over the held s_5 .

Rejected: s_5 (next c_9).

Rounds 7, 8, & 9.

$s_5 \rightarrow c_9$ (rejected, $s_{10} \succ s_5$).

$s_5 \rightarrow c_{10}$ (held, $s_5 \succ s_1$). s_1 is bumped!

$s_1 \rightarrow c_6$ (rejected, $s_4 \succ s_1$).

Round 10.

Student s_1 has exhausted their truncated list (reaching \emptyset after c_6). s_1 **drops out of the market**. No further proposals can be made. The algorithm terminates.

The student-optimal stable matching μ_S under truncated preferences is:

College	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
Student	s_3	s_9	s_6	s_2	\emptyset	s_4	s_8	\emptyset	s_{10}	s_5

Notice that students s_1 and s_7 , as well as colleges c_5 and c_8 , remain completely unmatched. This is the profound impact of the Individual Rationality constraint: stability now requires that some participants go home empty-handed rather than accept sub-par matches.

The introduction of truncated preferences is not just a way to model outside options (Individual Rationality); it also reveals a powerful method of strategic manipulation in matching markets.

If the mechanism asks participants to submit a rank-order list, can an agent benefit by *pretending* that certain acceptable partners are actually unacceptable? In other words, can an agent game the system by submitting an artificially shortened list? This question was famously answered by Alvin Roth.

Theorem 6.29: Strategic Truncation (Roth, 1982)

- **The Proposing Side:** In the college-proposing DA, no college can ever strictly improve its match by truncating its true preference list. (Symmetrically, no student can benefit from truncation in the student-proposing DA).
- **The Receiving Side:** In the student-proposing DA, a college **CAN** strictly improve its match by artificially truncating its preference list!

Remark (The Intuition Behind Strategic Truncation).

That the proposing side cannot benefit from truncation is a direct corollary of the Gale-Shapley Strategy-Proofness theorem: since truthful reporting is a dominant strategy for the proposers, no manipulation (including truncation) can help them.

But why does truncation work for the receiving side? Consider a college under the student-proposing DA. Suppose its true preference is $s_1 \succ_c s_2 \succ_c s_3$.

If the college reports truthfully, it might receive an application from s_2 early in the algorithm and tentatively hold them. Holding s_2 acts as a “safety net” for the college, but it also permanently removes s_2 from the broader applicant pool, effectively stabilizing the rest of the market.

If the college instead **truncates** its list to only include its top choice ($P'_c : s_1 \succ \emptyset$), it will ruthlessly reject s_2 . This rejection forces s_2 to apply to their next-choice college, potentially triggering a massive **chain of rejections** across the market. If this chain reaction eventually bumps s_1 from whichever college was holding them, s_1 will fall down their own preference list and apply to c . By burning its safety net, the college forcefully

destabilizes the market to shake loose its top choice!

Of course, truncation is a high-risk, high-reward strategy: if the rejection chain does not circle back to bump s_1 , the college will end up completely empty-handed (\emptyset).

6.3.6 The Set of Stable Matchings

We now investigate the structure of the set of all stable matchings. The key results are: (i) the student-proposing DA produces the best stable matching for students, (ii) the interests of the two sides are diametrically opposed, and (iii) the set of stable matchings forms a lattice.

Example (A 4×4 Matching Problem).

Consider 4 students and 4 colleges with the following preferences:

Rank	c_1	c_2	c_3	c_4	Rank	s_1	s_2	s_3	s_4
1	s_1	s_2	s_3	s_4	1	c_4	c_3	c_2	c_1
2	s_2	s_1	s_4	s_3	2	c_3	c_4	c_1	c_2
3	s_3	s_4	s_1	s_2	3	c_2	c_1	c_4	c_3
4	s_4	s_3	s_2	s_1	4	c_1	c_2	c_3	c_4

Notice the “anti-diagonal” structure: colleges prefer students in order s_1, s_2, s_3, s_4 , while students prefer colleges in the reverse order c_4, c_3, c_2, c_1 .

Student-Proposing DA. Each student proposes to their top choice: $s_1 \rightarrow c_4, s_2 \rightarrow c_3, s_3 \rightarrow c_2, s_4 \rightarrow c_1$. Every college receives exactly one proposal. No rejections—the algorithm terminates immediately.

$$\mu_S : (c_1, s_4), (c_2, s_3), (c_3, s_2), (c_4, s_1).$$

Every student gets their *first* choice. Every college gets their *last* choice.

College-Proposing DA. Each college proposes to their top choice: $c_1 \rightarrow s_1, c_2 \rightarrow s_2, c_3 \rightarrow s_3, c_4 \rightarrow s_4$. Every student receives exactly one proposal. No rejections.

$$\mu_C : (c_1, s_1), (c_2, s_2), (c_3, s_3), (c_4, s_4).$$

Every college gets their *first* choice. Every student gets their *last* choice.

A third stable matching. Consider:

$$\bar{\mu} : (c_1, s_3), (c_2, s_1), (c_3, s_4), (c_4, s_2).$$

One can verify that $\bar{\mu}$ has no blocking pair and is therefore stable, yet $\bar{\mu} \neq \mu_S$ and $\bar{\mu} \neq \mu_C$. In fact, this small market has **10 stable matchings** in total—the two DA solutions plus 8 others.

Student-Optimality of DA

Theorem 6.30: Student-Optimality of DA

The stable matching μ_S produced by the student-proposing DA is **weakly preferred by every student** to any other stable matching. That is, for every student s and every stable matching μ :

$$\mu_S^{-1}(s) \succeq_s \mu^{-1}(s).$$

By symmetry, the college-proposing DA produces the college-optimal stable matching μ_C .

Proof for Theorem

We first introduce a key definition.

Definition 6.31: Possible College

A college c is **possible** for student s if there exists some stable matching μ such that $\mu(s) = c$.

Claim

In the student-proposing DA, no student s is ever rejected by a college c that is possible for s .

Proof for Claim.

By induction on the number of steps in the algorithm.

Base case. At step 1, every student proposes to their top choice. No rejections have occurred yet, so the claim holds vacuously.

Inductive step. Suppose that up to step $t - 1$, no student has been rejected by a college that is possible for them. Suppose at step t , student s is rejected by college c , because c prefers another student s' (whom it holds) to s :

$$s' \succ_c s.$$

We need to show that c is not possible for s .

Since s' is currently proposing to c at step t , student s' must have already been rejected by every college they prefer to c . By the induction hypothesis, all those colleges that rejected s' were not possible for s' . Therefore, c is the best possible college for s' : in any stable matching μ , we have $c \succeq_{s'} \mu(s')$.

Now suppose, for contradiction, that c is possible for s . Then there exists a stable matching μ with $\mu(s) = c$. In this matching μ , student s' is matched to $\mu(s') \neq c$ (since s has c). But we just argued $c \succeq_{s'} \mu(s')$. Since $\mu(s') \neq c$ (because s occupies c in μ) and preferences are strict, this implies $c \succ_{s'} \mu(s')$. Moreover, c strictly prefers s' to $s = \mu(c)$. This means (c, s') blocks μ —contradicting the stability of μ . Therefore c is not possible for s . ■

The theorem follows immediately. Since no student is ever rejected by a possible college during the algorithm, each student's final match under μ_S is the *best* college that is achievable in any stable matching. ■

Opposition of Interests

The following proposition reveals that the two sides of the market have *diametrically opposed* interests over stable matchings.

Proposition 6.32: Opposition of Interests

If μ and μ' are both stable matchings and every student weakly prefers μ to μ' , then every college weakly prefers μ' to μ .

Proof for Proposition.

Suppose all students weakly prefer μ to μ' . Suppose for contradiction that some college c strictly prefers μ to μ' , i.e., $\mu(c) \succ_c \mu'(c)$. Let $s = \mu(c)$. Then:

- College c prefers s to $\mu'(c)$: $s \succ_c \mu'(c)$.
- Student s weakly prefers μ to μ' , meaning $c = \mu^{-1}(s) \succsim_s \mu'^{-1}(s)$.

If s strictly prefers c to $\mu'^{-1}(s)$, then (c, s) blocks μ' , contradicting stability. If s is indifferent, then since preferences are strict, $\mu^{-1}(s) = \mu'^{-1}(s) = c$, which implies $\mu(c) = \mu'(c) = s$, contradicting our assumption that c strictly prefers μ to μ' . ■

Corollary 6.33: Pessimality

The student-optimal stable matching μ_S is simultaneously the **college-pessimal** stable matching: every college gets its worst partner among all stable matchings. Symmetrically, μ_C is the student-pessimal stable matching.

The Lattice of Stable Matchings

Definition 6.34: Lattice of Stable Matchings

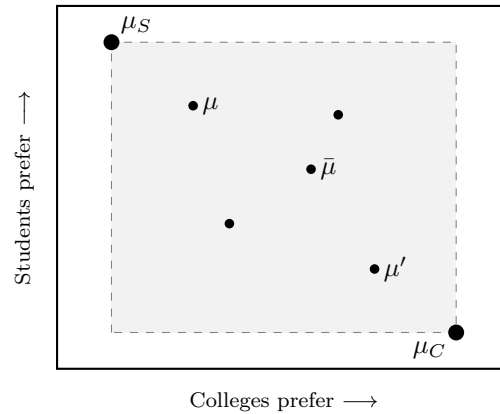
Let \mathcal{M} denote the set of all stable matchings. Define a partial order \geq_S on \mathcal{M} from the students' perspective: $\mu \geq_S \mu'$ if every student weakly prefers μ to μ' . The set (\mathcal{M}, \geq_S) forms a **distributive lattice**. This means:

- For any two stable matchings μ and μ' , there exist stable matchings $\mu \vee \mu'$ (the *join*: student-best of the two) and $\mu \wedge \mu'$ (the *meet*: student-worst of the two).
- The lattice has a unique maximum element μ_S (student-optimal) and a unique minimum element μ_C (college-optimal, equivalently student-pessimal).

By the Opposition of Interests, the college ordering is exactly the reverse: what is the join for students is the meet for colleges, and vice versa.

Remark.

In the 4×4 example above, μ_S sits at the top of the lattice (every student's favorite) and μ_C at the bottom. All 8 remaining stable matchings lie strictly between these two extremes. The following diagram illustrates the structure: the vertical axis represents student welfare (higher is better for students) and the horizontal axis represents college welfare (further right is better for colleges). By the Opposition of Interests, μ_S occupies the upper-left corner and μ_C the lower-right corner. All other stable matchings are confined to the rectangular region spanned by these two extremes.



Moving toward μ_S makes every student weakly better off and every college weakly worse off. Moving toward μ_C does the opposite. Any stable matching must lie weakly within the bounds set by these two extremes.

The Rural Hospital Theorem

When we allow preferences to include truncations—i.e., agents may find some partners *unacceptable*, so that not every agent is necessarily matched—the following striking result holds.

Theorem 6.35: Rural Hospital Theorem (Lone Wolf Theorem)

If a student (or college) is unmatched in **some** stable matching, then that student (or college) is unmatched in **all** stable matchings.

Proof for Theorem

We prove the result for the case of one-to-one matching ($q_c = 1$ for all c).

Let $S_u \subseteq \mathcal{S}$ be the set of students unmatched in μ_S (the student-optimal stable matching), and let $C_u \subseteq \mathcal{C}$ be the set of colleges unmatched in μ_S . Similarly, let S'_u and C'_u be the sets of students and colleges unmatched in μ_C (the college-optimal stable matching).

Step 1: $S_u \subseteq S'_u$. Since μ_S is the student-optimal stable matching, every student gets their best possible partner across all stable matchings. If a student s is unmatched even in μ_S , then s cannot be matched in any stable matching. In particular, s is unmatched in μ_C . Hence $S_u \subseteq S'_u$.

Step 2: $C'_u \subseteq C_u$. By symmetric reasoning, μ_C is the college-optimal stable matching, so every college gets its best possible partner. If a college c is unmatched in μ_C , then c is unmatched in every stable matching, including μ_S . Hence $C'_u \subseteq C_u$.

Step 3: Counting argument. In any one-to-one matching, the number of matched students equals the number of matched colleges. Therefore:

$$|\mathcal{S}| - |S_u| = |\mathcal{C}| - |C_u| \quad \text{and} \quad |\mathcal{S}| - |S'_u| = |\mathcal{C}| - |C'_u|.$$

From Step 1, $|S_u| \leq |S'_u|$, so the number of matched students in μ_S is at least the number in μ_C . From Step 2, $|C'_u| \leq |C_u|$, so the number of matched colleges in μ_C is at least the number in μ_S . But matched students = matched colleges in each matching, so:

$$\underbrace{|\mathcal{S}| - |S_u|}_{\text{matched in } \mu_S} \geq \underbrace{|\mathcal{S}| - |S'_u|}_{\text{matched in } \mu_C} = \underbrace{|\mathcal{C}| - |C'_u|}_{\text{matched in } \mu_C} \geq \underbrace{|\mathcal{C}| - |C_u|}_{\text{matched in } \mu_S}.$$

The leftmost and rightmost terms are equal (both count matched pairs in μ_S), so all inequalities are equalities. This forces $|S_u| = |S'_u|$ and $|C_u| = |C'_u|$. Combined with $S_u \subseteq S'_u$ and $C'_u \subseteq C_u$, we conclude:

$$S_u = S'_u \quad \text{and} \quad C_u = C'_u.$$

Since every stable matching μ satisfies $S_u \subseteq S'_u \subseteq S''_u$ (the first inclusion by student-optimality of μ_S , the second by college-optimality of μ_C), and $S_u = S'_u$, we have $S''_u = S_u$ for all stable matchings μ . ■

Remark.

The name “Rural Hospital” comes from the practical observation in medical residency matching: hospitals in rural areas that fail to fill all their positions under one stable matching will fail to fill them under *every* stable matching. Switching from the hospital-proposing to the resident-proposing DA (or any other stable mechanism) cannot help a rural hospital that is fundamentally undesirable to residents. The set of matched agents is invariant across stable matchings—only the *partners* can change, not whether one is matched at all.

6.4 The Boston Mechanism (Immediate Acceptance)

To truly appreciate the elegance and power of the Deferred Acceptance (DA) algorithm, it is instructive to study a mechanism that attempts to solve the same problem but fails spectacularly in its theoretical properties. The most famous example is the **Boston Mechanism**, historically used by the Boston Public School system (and many other school districts worldwide) before economists pointed out its severe flaws.

The Boston Mechanism is also known as the **Immediate Acceptance** mechanism. The critical difference from DA lies in the word “immediate”: once a college accepts a student, that decision is permanent. Seats are locked in, and future applicants cannot displace already-admitted students, regardless of how highly the college ranks the new applicants.

Definition 6.36: The Boston Mechanism

1. **Round 1:** Each student proposes to their first-choice college. Each college looks at its proposals, accepts its most preferred students up to its quota, and rejects the rest. **These acceptances are final.** The college's quota is reduced accordingly.
2. **Round k ($k \geq 2$):** Each student rejected in Round $k - 1$ proposes to their k -th choice college. Each college looks at its *remaining* unfilled seats (if any). It accepts its most preferred students among the new applicants up to its remaining quota, and rejects the rest. **These acceptances are final.**
3. **Repeat** until all students are matched or all available seats are filled.

Example (Boston Mechanism on the 10x10 Matching Problem).

Let us apply the Boston Mechanism to the exact same 10-student, 10-college problem. All colleges have quota $q = 1$.

College preferences (Boston Mechanism final matches highlighted):

Rank	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
1	s_1	s_1	s_{10}	s_6	s_7	s_4	s_1	s_6	s_8	s_6
2	s_2	s_6	s_2	s_1	s_2	s_8	s_4	s_4	s_{10}	s_5
3	s_3	s_9	s_4	s_5	s_8	s_2	s_3	s_2	s_6	s_1
4	s_4	s_2	s_6	s_2	s_4	s_6	s_8	s_{10}	s_9	s_4
5	s_5	s_5	s_8	s_7	s_{10}	s_{10}	s_5	s_8	s_5	s_{10}
6	s_6	s_8	s_3	s_3	s_5	s_3	s_2	s_3	s_7	s_8
7	s_7	s_3	s_5	s_8	s_9	s_7	s_9	s_1	s_4	s_9
8	s_8	s_4	s_7	s_4	s_6	s_1	s_6	s_9	s_3	s_2
9	s_9	s_7	s_9	s_{10}	s_3	s_5	s_7	s_7	s_1	s_3
10	s_{10}	s_{10}	s_1	s_9	s_1	s_9	s_{10}	s_5	s_2	s_7

Student preferences (Boston Mechanism final matches highlighted):

Rank	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
1	c_5	c_4	c_5	c_6	c_3	c_3	c_8	c_1	c_1	c_2
2	c_{10}	c_9	c_6	c_2	c_7	c_4	c_6	c_4	c_5	c_{10}
3	c_6	c_{10}	c_1	c_5	c_9	c_5	c_4	c_2	c_2	c_9
4	c_4	c_6	c_7	c_8	c_{10}	c_1	c_2	c_6	c_{10}	c_4
5	c_7	c_3	c_4	c_1	c_1	c_9	c_9	c_3	c_3	c_1
6	c_2	c_7	c_8	c_7	c_5	c_7	c_7	c_7	c_8	c_7
7	c_8	c_1	c_9	c_3	c_4	c_{10}	c_5	c_5	c_6	c_5
8	c_9	c_5	c_{10}	c_9	c_2	c_8	c_{10}	c_{10}	c_7	c_6
9	c_1	c_2	c_2	c_{10}	c_6	c_2	c_1	c_8	c_4	c_3
10	c_3	c_8	c_3	c_4	c_8	c_6	c_3	c_9	c_9	c_8

We use the same chalkboard-style simulation. However, green text now represents a **permanent acceptance (admitted)**, rather than a tentative hold. Once a column is filled with a green student, that college is closed forever.

Round 1.

All 10 students propose to their first-choice colleges.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_8	s_{10}	s_6	s_2	s_3	s_4		s_7		
s_9		s_5		s_1					

Explanation: Colleges c_2, c_4, c_6, c_8 receive exactly one applicant and admit them immediately. Colleges c_1, c_3, c_5 face competition. They admit their preferred applicant (s_8, s_6, s_3 respectively) and reject the rest.

Status: Colleges $c_1, c_2, c_3, c_4, c_5, c_6, c_8$ are now **permanently closed**.

Rejected: s_9 (next c_5), s_5 (next c_7), s_1 (next c_{10}).

Round 2.

The three rejected students propose to their second-choice colleges.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
				FULL		s_5			s_1
				s_9					

Explanation: Students s_1 and s_5 apply to c_{10} and c_7 , which are still empty. They are admitted permanently. Student s_9 applies to c_5 . However, c_5 was already filled by s_3 in Round 1! Even though c_5 prefers s_9 (#7) over s_3 (#9), it is too late. The seat is gone. s_9 is automatically rejected.

Status: Colleges c_7, c_{10} are now **permanently closed**. Only c_9 remains open.

Rejected: s_9 (next c_2).

Rounds 3 through 9 (The Free-Fall of s_9).

Student s_9 has suffered a “cascading failure.” Because he missed his first choice, he is now applying to colleges that were filled in Round 1.

- **Round 3:** $s_9 \rightarrow c_2$. **FULL** (filled by s_{10}). Rejected.
- **Round 4:** $s_9 \rightarrow c_{10}$. **FULL** (filled by s_1). Rejected.
- **Round 5:** $s_9 \rightarrow c_3$. **FULL** (filled by s_6). Rejected.
- **Round 6:** $s_9 \rightarrow c_8$. **FULL** (filled by s_7). Rejected.
- **Round 7:** $s_9 \rightarrow c_6$. **FULL** (filled by s_4). Rejected.
- **Round 8:** $s_9 \rightarrow c_7$. **FULL** (filled by s_5). Rejected.
- **Round 9:** $s_9 \rightarrow c_4$. **FULL** (filled by s_2). Rejected.

Round 10.

Student s_9 finally reaches his 10th and absolute worst choice, c_9 . Luckily, c_9 is still empty.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
								s_9	

The algorithm terminates.

The **Boston Mechanism matching** μ_B is:

College	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
Student	s_8	s_{10}	s_6	s_2	s_3	s_4	s_5	s_7	s_9	s_1

6.4.1 Properties of the Boston Mechanism

The simulation above provides everything we need to expose the theoretical failures of the Boston Mechanism. We examine the two most important properties: Pairwise Stability and Strategy-Proofness.

Theorem 6.37: The Boston Mechanism is NOT Stable

The allocation produced by the Boston Mechanism may contain blocking pairs, meaning it fails to guarantee a pairwise stable matching.

Proof for Theorem

We can prove this by simply identifying a blocking pair in the outcome μ_B of our example.

Under μ_B , student s_9 is matched with c_9 (his 10th choice), and college c_5 is matched with s_3 (its 9th choice). Now look at their true preferences:

- Student s_9 ranks c_5 as his 2nd choice, meaning $c_5 \succ_{s_9} \mu_B(s_9) = c_9$.
- College c_5 ranks s_9 as its 7th choice, meaning $s_9 \succ_{c_5} \mu_B(c_5) = s_3$.

Both s_9 and c_5 would strictly prefer to be matched with each other rather than their current assignments. Thus, (c_5, s_9) **forms a blocking pair**.

The instability arises precisely because of the “immediate acceptance” rule: c_5 was forced to permanently commit to s_3 in Round 1, preventing it from accepting the much stronger applicant s_9 who arrived in Round 2. ■

Theorem 6.38: The Boston Mechanism is NOT Strategy-Proof

Under the Boston Mechanism, truthful reporting is not a dominant strategy. Students can often obtain a strictly better outcome by misreporting their preferences.

Proof for Theorem

Because missing out on your first choice is heavily penalized (the “free-fall” effect seen with s_9), the Boston Mechanism forces students to play a high-stakes psychological game. Instead of ranking their true favorite first, students are incentivized to rank a “safe” school first to secure a seat in Round 1.

We demonstrate this with student s_9 from our example. Under truthful reporting,

s_9 listed c_1 as his 1st choice and c_5 as his 2nd choice. Because he lost the tie-breaker at c_1 in Round 1, he free-fell all the way to c_9 , his worst possible outcome.

Suppose s_9 anticipates this and decides to deviate. He submits a fake preference list P'_{s_9} where he promotes his 2nd choice (c_5) to his 1st choice, abandoning the highly competitive c_1 . The rest of the students report truthfully.

Modified Student Preferences (with s_9 's misreport highlighted in blue):

Rank	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s'_9	s_{10}
1	c_5	c_4	c_5	c_6	c_3	c_3	c_8	c_1	c_5	c_2
2	c_{10}	c_9	c_6	c_2	c_7	c_4	c_6	c_4	c_1	c_{10}
3	c_6	c_{10}	c_1	c_5	c_9	c_5	c_4	c_2	c_2	c_9
4	c_4	c_6	c_7	c_8	c_{10}	c_1	c_2	c_6	c_{10}	c_4
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

Let us trace **Round 1** of the Boston Mechanism under this new profile. All students apply to their reported 1st choices.

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
s_8	s_{10}	s_6	s_2	s_9	s_4		s_7		
		s_3		s_5					
				s_1					

Explanation of c_5 's decision: College c_5 now receives Round 1 proposals from s_1 , s_3 , and the strategically misreporting s_9 . Looking at c_5 's true preference list: s_9 (#7) \succ_{c_5} s_3 (#9) \succ_{c_5} s_1 (#10). Because c_5 evaluates all three applicants simultaneously in Round 1, it permanently admits its favorite among them: s_9 .

By lying and placing c_5 first, student s_9 completely avoided the free-fall and secured his true 2nd choice (c_5). This outcome is strictly better for him than the 10th choice (c_9) he received when telling the truth. Since a profitable deviation exists, the Boston mechanism is **not strategy-proof**. ■

Remark (The Tragedy of the Boston Mechanism).

The Boston Mechanism systematically rewards sophisticated, strategic families while punishing honest, naive ones. A naive student who truthfully applies to a competitive “reach” school wastes their critical Round 1 proposal. By the time they are rejected and move to their “match” schools in Round 2, those schools have already been filled by strategic students who ranked them first. This profound inequity is why many school districts have replaced it with the Strategy-Proof Deferred Acceptance algorithm.

Remark (Chapter Summary).

Matching theory studies allocation problems in which prices cannot do all the work—either because money is unavailable (school choice, kidney exchange, public housing)

or because preferences are intrinsically two-sided (medical residents and hospitals, men and women). Two foundational algorithms organize the field. *Top trading cycles* (TTC) solves the one-sided housing-market problem: starting from arbitrary endowments, it produces the unique core allocation, is individually rational, Pareto efficient, and group strategy-proof. *Deferred acceptance* (Gale-Shapley DA) solves the two-sided market: it produces a stable matching, is strategy-proof for the proposing side, and the proposing side gets its optimal stable matching. Three trade-offs run through the chapter. *Stability vs. optimality*: stability picks out a lattice of matchings; the proposing side gets the top, the receiving side gets the bottom. *Strategy-proofness vs. efficiency*: in priority-based settings, DA is strategy-proof but not Pareto efficient; TTC is Pareto efficient but ignores priorities. *Truthfulness vs. welfare*: the Boston mechanism rewards strategic sophistication and harms honest participants, motivating the move to DA in real-world school choice systems (NYC 2003, Boston 2005).

Part V


Information and Dynamic
Games

Chapter 7

Common Knowledge

7.1 Motivating Example: The Corrupted Industry

Consider an industry with 10 firms C_1, C_2, \dots, C_{10} . Each firm C_i has a manager M_i . The industry is thoroughly corrupted: every manager M_i is secretly selling their own company's private information to all other companies. In addition, each company is buying private information from the managers of other companies.



Every M_i sells to all C_j ($j \neq i$); only M_1 's arrows shown.

There is one corporate governance rule: if a company discovers that its own manager is selling information to outsiders, it will fire the manager on the **Friday of that week**.

What does each firm know? Consider firm C_1 . Since C_1 is buying information from the managers M_2, M_3, \dots, M_{10} , it directly observes that those 9 managers are corrupt. However, C_1 **cannot observe whether its own manager M_1 is corrupt**—it does not know whether M_1 is selling to others. The same logic applies symmetrically to every firm. In summary:

- Each firm knows that *at least 9* managers are corrupt (the ones it buys from).
- No firm knows for certain whether its *own* manager is corrupt.

Now suppose that at the start of Week 1, a regulator makes a **public announcement**: “*At least one manager in this industry is corrupt.*” This statement is heard by all 10 firms simultaneously, and everyone knows that everyone heard it.

What happens? Nothing happens in Week 1. Nothing happens in Week 2. In fact, nothing happens until **Week 10**, at which point all 10 managers are fired simultaneously.

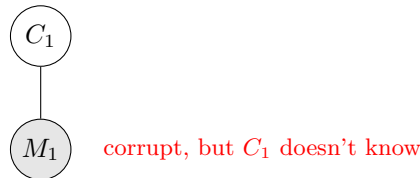
At first glance, this is deeply puzzling. Every firm already knew that at least 9 managers were corrupt—far more than the “at least one” in the announcement. How can such a

seemingly vacuous statement trigger any action at all, let alone a cascade that takes 10 weeks to unfold?

7.1.1 Iterated Reasoning


The answer lies in the chain of *iterated hypothetical reasoning* that the public announcement enables.

The $n = 1$ Case. Suppose there is only 1 firm with 1 corrupt manager.



Before the announcement, C_1 does not know its own manager is corrupt. After the announcement (“at least one is corrupt”), C_1 immediately deduces that M_1 must be the corrupt one, and fires him on Friday of **Week 1**.

The $n = 2$ Case. Suppose there are 2 firms, both with corrupt managers. Each firm observes that the other firm’s manager is corrupt (by buying information from him), but does not know about its own.



Consider the reasoning from C_1 ’s perspective after the announcement. C_1 observes that M_2 is corrupt, but entertains two hypotheses about its own manager:

- **Hypothesis A:** M_1 is not corrupt. Then only M_2 is corrupt. Under this hypothesis, after the public announcement, C_2 would find itself in the $n = 1$ situation: C_2 would see no other corrupt manager, learn from the announcement that at least one exists, and deduce that M_2 is the one. So C_2 would fire M_2 in **Week 1**.
- **Hypothesis B:** M_1 is also corrupt.

C_1 waits and watches. At the end of Week 1, M_2 is *not* fired. This rules out Hypothesis A: if M_1 had been clean, C_2 would have acted already. Therefore C_1 concludes that M_1 must also be corrupt, and fires M_1 in **Week 2**.

By perfectly symmetric reasoning, C_2 also watches whether C_1 fires M_1 in Week 1. When it does not happen, C_2 likewise infers that M_2 is corrupt. Both managers are fired simultaneously in **Week 2**.

The General Case: Induction. By induction on the number of corrupt managers n . Each firm sees $n - 1$ corrupt managers and entertains the hypothesis that its own manager is clean (so that only $n - 1$ managers are corrupt). Under that hypothesis, by the inductive assumption, the other $n - 1$ firms would have fired their managers by Week $n - 1$. When

this does not happen by the end of Week $n - 1$, every firm deduces that the total must be n , meaning its own manager is also corrupt. All n managers are fired simultaneously in **Week n** .

For our example with $n = 10$: nothing happens for 9 weeks, and then all 10 managers are fired in Week 10.

7.1.2 Why Does a “Vacuous” Announcement Matter?

The content of the announcement (“at least one is corrupt”) was already known to every firm—each firm could see at least 9 corrupt managers. What was *not* in place before the announcement was the full tower of *hypothetical reasoning about hypothetical reasoning*.

To see why, trace the chain from C_1 ’s perspective. C_1 reasons: “I see 9 corrupt managers. But suppose M_1 is clean. Then C_2 would see only 8 corrupt managers. And C_2 might suppose M_2 is also clean. Then C_3 would see only 7...” Each level of nesting reduces the number of known-corrupt managers by one. After 9 levels of hypothetical reasoning, we reach a scenario where some firm sees *zero* other corrupt managers and does not know if any exist at all. Without the announcement, the induction has no base case in this deepest hypothetical—the chain of reasoning collapses before reaching its logical conclusion.

The public announcement plugs exactly this gap. By making “at least one is corrupt” publicly and simultaneously known to all, it provides the base case in every hypothetical scenario, no matter how deeply nested. This distinction—between knowledge that is shared but *finitely* nested, and knowledge that is shared *infinitely* deep—is precisely the distinction between *mutual knowledge* and *common knowledge*, which we now formalize.

7.2 Formal Definition

Definition 7.1: Mutual Knowledge and Common Knowledge

Let E be an event and let there be n agents.

- E is **mutual knowledge of order 1** if every agent knows E .
- E is **mutual knowledge of order k** if every agent knows that E is mutual knowledge of order $k - 1$.
- E is **common knowledge** if E is mutual knowledge of order k for *every* $k \in \mathbb{N}$.

In other words, common knowledge means: everyone knows E , everyone knows that everyone knows E , everyone knows that everyone knows that everyone knows E , and so on *ad infinitum*.

In the corrupted industry example, the event $E =$ “at least one manager is corrupt” was, before the announcement, mutual knowledge of high order but **not** common knowledge. Each firm’s chain of hypotheticals (“suppose mine is clean, then you see one fewer, then suppose yours is clean too...”) bottoms out after finitely many levels at a scenario in which one firm sees zero corrupt managers and cannot determine whether E holds. Common knowledge requires that the chain never bottoms out—that at every depth of hypothetical

reasoning, every agent in the hypothetical scenario still knows E .

The public announcement achieves exactly this. Because every firm heard the announcement, and every firm knows every other firm heard it, and so on without bound, E becomes common knowledge. The base case of the induction (“a firm that sees zero corrupt managers and hears the announcement will fire its manager in Week 1”) is now available at every depth, enabling the full chain of iterated reasoning to unfold.

7.3 Aumann’s Agreement Theorem

A celebrated consequence of common knowledge is that two Bayesian agents who share a common prior cannot “agree to disagree” about the posterior of any event, provided the posteriors themselves are common knowledge.

Theorem 7.2: Aumann (1976)

Let two Bayesian agents have a common prior on a finite state space Ω and partitions $\mathcal{P}_1, \mathcal{P}_2$ representing what each agent observes. Let $E \subseteq \Omega$ be an event. If at state ω^* the values of both posteriors $\Pr_1(E | \mathcal{P}_1)$ and $\Pr_2(E | \mathcal{P}_2)$ are common knowledge, then they are equal.

Proof for Theorem

Let $q_1 = \Pr_1(E | \mathcal{P}_1(\omega^*))$ and $q_2 = \Pr_2(E | \mathcal{P}_2(\omega^*))$ be the posteriors. Common knowledge of q_1 at ω^* means that on every state in the meet-cell containing ω^* , agent 1’s posterior is exactly q_1 . Let $M = \mathcal{P}_1 \wedge \mathcal{P}_2$ be the meet (the join of the two partitions in the lattice of partitions, refined by both); the cell of M containing ω^* is the set of states agent 1’s \mathcal{P}_1 -cell could be, intersected with the corresponding \mathcal{P}_2 -cells.

Within this meet-cell $M(\omega^*)$, agent 1’s cells partition $M(\omega^*)$ into sub-cells C_1, C_2, \dots with $\Pr(E | C_k) = q_1$ on each. By the law of total probability,

$$\begin{aligned} \Pr(E | M(\omega^*)) &= \sum_k \Pr(C_k | M(\omega^*)) \cdot \Pr(E | C_k) \\ &= q_1 \cdot \sum_k \Pr(C_k | M(\omega^*)) = q_1. \end{aligned}$$

Symmetrically, $\Pr(E | M(\omega^*)) = q_2$. Hence $q_1 = q_2$. ■

Remark (The Force of the Common-Prior Assumption).

The key non-triviality is that the agents need not have access to the same information; they only need to know the same prior and to commonly know each other’s posterior. The argument iterates: each agent learns nothing about E beyond what is already encoded in the meet of partitions, because anything else would cause the posterior to update, breaking the common-knowledge assumption. Aumann’s theorem has been used to argue against speculative trade between rational agents (the no-trade theorem of Milgrom-Stokey 1982): if both traders share a common prior and rationality is common knowledge, they cannot strictly benefit from trade, because the willingness to trade itself becomes informative.

7.4 The Coordinated Attack Problem

We now present a game-theoretic example that dramatically illustrates the difference between “almost common knowledge” and true common knowledge. This is based on Rubinstein’s (1989) *Electronic Mail Game*.

7.4.1 Setup

Two divisions of an army (Player 1 and Player 2) must decide whether to **attack** (A) or **not attack** (N). The enemy is either **weak** (state G , the “good” state) or **strong** (state B , the “bad” state). The army wins if and only if *both* divisions attack *and* the enemy is weak. The prior probabilities are $\Pr(G) = \pi < \frac{1}{2}$ and $\Pr(B) = 1 - \pi$.

The payoff matrices are:

	State G (prob π)		State B (prob $1 - \pi$)	
	A	N	A	N
A	2, 2	-3, 0	-2, -2	-2, 0
N	0, -3	0, 0	0, -2	0, 0

In state B , not attacking (N) is strictly dominant for both players. In state G , the game is a *coordination game*: (A, A) is a Nash equilibrium with payoff $(2, 2)$, but attacking alone is severely punished (-3) . Crucially, if the state G were **common knowledge**, both (A, A) and (N, N) would be Nash equilibria—the game is a coordination game with two pure-strategy equilibria. The interesting question is whether the players can coordinate on the Pareto-superior equilibrium (A, A) .


7.4.2 Information Structure

The communication protocol is as follows:

1. If the state is G , Nature sends a signal to Player 1 with certainty (no error at this stage). If the state is B , no signal is sent.
2. Upon receiving a signal, Player 1’s system *automatically* forwards a message to Player 2. This transmission fails (the message is lost) with probability $\varepsilon > 0$, independently.
3. If Player 2 receives the message, Player 2’s system automatically sends a confirmation back to Player 1, again lost with probability ε .
4. This back-and-forth continues indefinitely: each successfully received message triggers an automatic reply, each of which is independently lost with probability ε .

The message sending is *not* a strategic choice—it is automatic. The only strategic decision is whether to attack or not.

Let t_i denote the **number of messages received** by Player i . In state B , no messages are sent, so $t_1 = t_2 = 0$. In state G , the message chain proceeds as:



The chain stops at the first lost message. The possible outcomes given state G are:

Chain stops at	t_1	t_2
$P_1 \rightarrow P_2$ lost	1	0
$P_2 \rightarrow P_1$ lost	1	1
$P_1 \rightarrow P_2$ lost	2	1
$P_2 \rightarrow P_1$ lost	2	2
\vdots	\vdots	\vdots

Observe that t_1 and t_2 always differ by at most 1: given $t_1 = T \geq 1$, the value of t_2 is either $T - 1$ or T .

Remark (“Almost” Common Knowledge).

When ϵ is small, the message chain is very likely to be long. With high probability, both players receive many confirmations, so each player is “almost sure” the state is G , “almost sure” the other knows it, “almost sure” the other knows the other knows it, and so on. The state G is “almost” common knowledge. One might expect that (A, A) should be sustainable as an equilibrium. As we will see, this intuition is spectacularly wrong.

7.4.3 The Unique Equilibrium

A strategy for Player i is a function $\sigma_i : \{0, 1, 2, \dots\} \rightarrow \{A, N\}$ mapping the number of messages received to an action.

Proposition 7.3: Unique Equilibrium of the Coordinated Attack Game

There is a unique equilibrium: both players play N regardless of the number of messages received. That is, $\sigma_i(t_i) = N$ for all $t_i \geq 0$ and $i \in \{1, 2\}$.

Proof for Proposition.

By induction on the number of messages received.

Base case: $t_1 = 0$. If Player 1 receives no messages, then state B is certain (since Nature always signals Player 1 in state G). In state B , N is strictly dominant. So $\sigma_1(0) = N$.

Base case: $t_2 = 0$. If Player 2 receives no messages, there are two possibilities:

- (i) State is B and no signal was ever sent (prior probability $1 - \pi$);
- (ii) State is G , Player 1 received the signal, but the forward to Player 2 was lost (prior probability $\pi\epsilon$).

By Bayes' rule:

$$\Pr(G \mid t_2 = 0) = \frac{\pi\varepsilon}{\pi\varepsilon + (1 - \pi)}.$$

If Player 2 plays N , the payoff is 0. If Player 2 plays A , the payoff is at most (assuming Player 1 plays A whenever state is G):

$$\Pr(G \mid t_2 = 0) \times 2 + \Pr(B \mid t_2 = 0) \times (-2) = \frac{2\pi\varepsilon - 2(1 - \pi)}{\pi\varepsilon + (1 - \pi)} < 0,$$

since $\pi\varepsilon \ll 1 - \pi$ (recall $\pi < \frac{1}{2}$ and ε is small). So $\sigma_2(0) = N$.

Inductive step. Suppose that for all $t_1, t_2 < T$, both players play N : $\sigma_1(t_1) = N$ and $\sigma_2(t_2) = N$.

Consider Player 1 with $t_1 = T$ (the argument for Player 2 is symmetric). Given $t_1 = T$, Player 2's message count is either $t_2 = T - 1$ or $t_2 = T$. The conditional probability is:

$$p := \Pr(t_2 = T - 1 \mid t_1 = T) = \frac{\varepsilon}{\varepsilon + (1 - \varepsilon)\varepsilon} = \frac{1}{2 - \varepsilon} > \frac{1}{2}.$$

This is because, given $t_1 = T$: with probability ε the next message ($P_1 \rightarrow P_2$) is lost, yielding $t_2 = T - 1$; with probability $(1 - \varepsilon)\varepsilon$ the message reaches P_2 (so $t_2 \geq T$) but the return is lost, yielding $t_2 = T$.

By the induction hypothesis, $\sigma_2(T - 1) = N$. Even in the best case where Player 2 plays A when $t_2 = T$, if Player 1 plays A :

$$\text{Payoff to } P_1 \leq p \times (-3) + (1 - p) \times 2 = \frac{-3 + 2(1 - \varepsilon)}{2 - \varepsilon} = \frac{-1 - 2\varepsilon}{2 - \varepsilon} < 0.$$

The -3 arises because when $t_2 = T - 1$, Player 2 plays N (by induction), so Player 1 attacks alone. The 2 is the best-case payoff if Player 2 also attacks. Since $p > \frac{1}{2}$, the expected payoff from attacking is strictly negative. Therefore $\sigma_1(T) = N$.

By symmetric reasoning applied to Player 2, $\sigma_2(T) = N$. This completes the induction. ■

7.4.4 Interpretation

The result is striking. The following two scenarios yield completely different equilibrium outcomes:

Information structure	Equilibrium outcome
G is common knowledge	(A, A) is a NE
G is "almost" common knowledge	(N, N) is the <i>unique</i> NE

No matter how small $\varepsilon > 0$ is, and no matter how many thousands of confirmations each player has received, the *unique* equilibrium is for both to play N . The problem is not that the players lack confidence that G is true—after many messages, they are nearly certain. The problem is that each player is never quite sure the *other* player is sure that the other player is sure... that the state is G .

Remark (Common p -Belief).

The key quantitative insight is the conditional probability $p = \frac{1}{2-\varepsilon}$. Even when ε is tiny, p is close to $\frac{1}{2}$ —far from 1. This means that no matter how many messages Player i has received, Player i always assigns probability at least $\frac{1}{2}$ to the event that the other player received one fewer message. The information structure “self-replicates” at every level: the uncertainty never shrinks as we go deeper in the epistemic hierarchy.

This connects to the concept of *common p -belief* (Monderer and Samet, 1989): an event is common p -belief if everyone believes it with probability at least p , everyone believes with probability at least p that everyone believes it with probability at least p , and so on. In the electronic mail game, the state G is common p -belief only for $p \leq \frac{1}{2-\varepsilon} \approx \frac{1}{2}$, which is far from the $p = 1$ required for common knowledge. Coordination requires common knowledge (or at least common p -belief with p sufficiently large); “almost common knowledge” is not enough.

Remark (Chapter Summary).

Common knowledge is the precise epistemic condition under which agents’ inferences about each other’s reasoning can iterate without bound: E is common knowledge if everyone knows E , everyone knows that everyone knows E , and so on *ad infinitum*. The chapter developed three pillars. *The fixed-point characterization*: an event is common knowledge in a state ω if and only if it is true at every state in the meet of the players’ partitions containing ω . *Aumann’s agreement theorem*: agents with a common prior who are commonly known to be Bayesian rational cannot “agree to disagree” about the posterior of any event—if their posteriors are common knowledge, they must coincide. *The fragility of common knowledge*: Rubinstein’s electronic mail game shows that arbitrarily many rounds of mutual confirmation can fail to deliver common knowledge, and that small departures from common knowledge can flip the equilibrium discontinuously. The recurring theme: many central results in game theory (backward induction, Folk Theorem, signaling refinements) implicitly assume common knowledge of rationality, and the assumption matters more than its mathematical innocuousness suggests.

Chapter 8

Dynamic Games

8.1 Repeated Games

8.1.1 Setup

Consider a **stage game** $G = (S_i, u_i)_{i=1}^n$, where S_i is the action set of player i and $u_i : S \rightarrow \mathbb{R}$ is the payoff function, with $S = \prod_{i=1}^n S_i$. The **infinitely repeated game** $G^\delta(\infty)$ consists of playing G in every period $t = 1, 2, \dots$, with players discounting future payoffs by a common discount factor $\delta < 1$. We assume **perfect monitoring**: after each period, all players observe the realized action profile.

Remark (Terminology: “Stage Game” and “Action vs. Strategy”).

Two pieces of terminology deserve flagging up front, because they look like reuse of earlier symbols but mean something subtly different here.

- **Stage game** = the one-shot, simultaneous-move normal-form game G that gets played in every period of the repeat. It is the building block; the repeated game $G^\delta(\infty)$ is what we actually study.
- **Action vs. strategy.** In a one-shot game these collapse: each player picks a single $s_i \in S_i$ and the choice is the strategy. In the repeated game they part ways. S_i remains the *action set*—what is chosen each period—while a *strategy* σ_i is a much richer object, a map from histories to actions (see the next definition). When earlier chapters wrote “ S_i is the strategy set,” they were implicitly working in the one-shot setting where the two coincide.

Definition 8.1: Strategy in $G^\delta(\infty)$

A **strategy** for player i in the infinitely repeated game is a sequence $\sigma_i = (\sigma_i^1, \sigma_i^2, \dots)$, where:

- $\sigma_i^1 \in S_i$ specifies the action in period 1;
- For $t > 1$, $\sigma_i^t : S^{t-1} \rightarrow S_i$ maps the history of play $h^{t-1} = (s^1, s^2, \dots, s^{t-1})$ to an action in period t . The domain $S^{t-1} = \prod_{\tau=1}^{t-1} S$ is the set of all possible such histories.

*This definition restricts attention to **pure strategies**, since σ_i^t maps each history to a single action in S_i rather than to a mixed action in $\Delta(S_i)$. All results in this chapter extend to mixed (behavioral) strategies by replacing S_i with $\Delta(S_i)$; the pure-strategy formulation is adopted for notational simplicity.*

Given a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$, the induced action profile in period t is $s^t(\sigma)$. Player i 's **average discounted payoff** in $G^\delta(\infty)$ is:

$$U_i(\sigma) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(s^t(\sigma)).$$

The $(1 - \delta)$ normalization ensures that perpetual play of a constant action profile s yields payoff $u_i(s)$, making payoffs in the repeated game directly comparable to stage-game payoffs.

8.1.2 Subgame Perfect Equilibrium**Definition 8.2: Continuation Strategy**

Fix a strategy σ_i and a history $h^{t-1} \in S^{t-1}$. The **continuation strategy** $\sigma_i|_{h^{t-1}}$ is the strategy in the infinitely repeated game $G^\delta(\infty)$ defined by: for every $\tau \geq 1$ and every continuation history $\tilde{h}^{\tau-1} \in S^{\tau-1}$,

$$(\sigma_i|_{h^{t-1}})^\tau(\tilde{h}^{\tau-1}) = \sigma_i^{t+\tau-1}(h^{t-1}, \tilde{h}^{\tau-1}),$$

where $(h^{t-1}, \tilde{h}^{\tau-1})$ denotes the concatenated history. The superscript τ on the left is needed because a strategy is a *sequence* $(\sigma_i^1, \sigma_i^2, \dots)$ of period-by-period maps; we have to identify which period's component is being applied. The right-hand side accordingly uses the original strategy's period- $(t + \tau - 1)$ component.

Remark (Intuition for the Continuation Strategy).

Intuitively, fix the past h^{t-1} as a prefix, relabel period t as “period 1” of a new copy of $G^\delta(\infty)$, and read off what σ_i prescribes in that new copy. The profile $\sigma|_{h^{t-1}} = (\sigma_1|_{h^{t-1}}, \dots, \sigma_n|_{h^{t-1}})$ is the continuation profile after h^{t-1} .

Definition 8.3: Subgame Perfect Equilibrium (SPE) of $G^\delta(\infty)$

A strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ is a **subgame perfect equilibrium** if:


- (i) σ^* is a Nash equilibrium of $G^\delta(\infty)$: for all i and all alternative strategies $\sigma_i \in \Sigma_i$ (the set of all such history-to-action maps; we write $\Sigma = \prod_i \Sigma_i$ for the joint set),^a

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(s^t(\sigma^*)) \geq (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(s^t(\sigma_i, \sigma_{-i}^*)).$$

- (ii) For every period t and every history $h^{t-1} = (s^1, s^2, \dots, s^{t-1})$, the **continuation strategy** $\sigma^*|_{h^{t-1}}$ is also a Nash equilibrium of $G^\delta(\infty)$.

^aBy reindexing $t' = t - 1$, this sum is identical to $(1 - \delta) \sum_{t'=0}^{\infty} \delta^{t'} u_i(s^{t'+1})$. We start the index at $t = 1$ so that “period t ” reads as “the t -th period of play”—this is the convention in Osborne-Rubinstein and Mailath-Samuelson.

Condition (i) is ordinary Nash equilibrium: it constrains only on-path play (i.e., the unique history realized when every player follows σ^*). Condition (ii) sharpens this by requiring that σ^* induce a Nash equilibrium in every subgame, including the off-path ones generated by unilateral deviations (i.e., subgames rooted at histories that arise after some player has deviated from σ^* at least once). Because payoffs are time-separable and the stage game repeats unchanged, every subgame is structurally identical to a fresh copy of $G^\delta(\infty)$ —only the calendar index has shifted—so SPE asks for Nash behavior at every node of the history tree



not merely at the root.

Verifying condition (ii) looks forbidding: the tree has infinitely many nodes, and a deviation could in principle span infinitely many periods. Both sources of complexity turn out to be illusory. The **one-deviation principle**, proved next, reduces SPE verification to a local check—no player gains by deviating in *one* period and returning to σ^* —which can often be done history-by-history with a handful of inequalities.

8.1.3 The One-Deviation Principle

Three terms recur throughout the proof and the rest of the chapter; we record them once.

Definition 8.4: Continuation Value, One-Shot Deviation, Profitability

Fix a strategy profile σ , a player i , and a history h^{t-1} .

- The **continuation value** of σ for player i at h^{t-1} is the normalized discounted expected payoff from period t onward,

$$V(\sigma_i, \sigma_{-i} \mid h^{t-1}) = (1 - \delta) \mathbb{E} \left[\sum_{\tau \geq t} \delta^{\tau-t} u_i(s^\tau) \mid h^{t-1} \right].$$

- A strategy σ'_i is a **one-shot deviation** from σ_i at h^{t-1} if σ'_i differs from σ_i only in period t on the realized history—from period $t + 1$ on, it agrees with σ_i .
- A deviation σ'_i is **profitable** (relative to σ_i against σ_{-i} at h^{t-1}) if $V(\sigma'_i, \sigma_{-i} \mid h^{t-1}) > V(\sigma_i, \sigma_{-i} \mid h^{t-1})$.

Theorem 8.5: One-Deviation Principle

To check that σ^* is an SPE, it is sufficient to check that, following any history, no player i can benefit by deviating once and then returning to σ_i^* .

Proof for Theorem

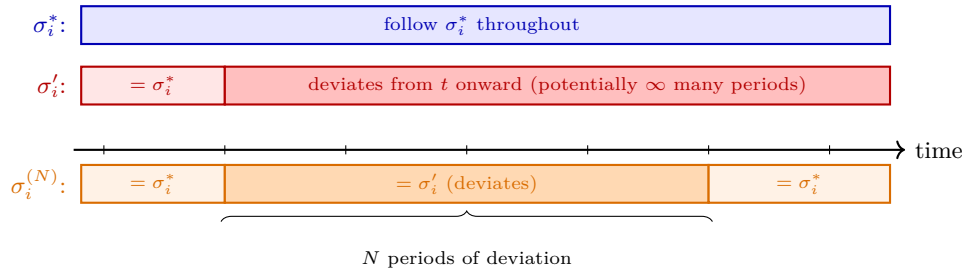
Fix a strategy profile σ^* and assume no player has a profitable one-shot deviation after any history. We show σ^* is an SPE.

Fix a player i , a history h^{t-1} , and a deviation σ'_i with

$$V(\sigma'_i, \sigma_{-i}^* \mid h^{t-1}) > V(\sigma_i^* \mid h^{t-1}, \sigma_{-i}^* \mid h^{t-1}).$$

We derive a contradiction.

Step 1 (Truncation). Since S is finite, stage payoffs are bounded: set $\bar{u} = \max_s |u_i(s)|$. Let $\Delta = V(\sigma'_i, \sigma_{-i}^* \mid h^{t-1}) - V(\sigma_i^* \mid h^{t-1}, \sigma_{-i}^* \mid h^{t-1}) > 0$. Define the truncated strategy $\sigma_i^{(N)}$ to follow σ'_i in the first N periods $t, t + 1, \dots, t + N - 1$ and to revert to σ_i^* from period $t + N$ on. Pictorially:



We claim $\sigma_i^{(N)}$ remains strictly profitable once N is large enough. By construction, $\sigma_i^{(N)}$ and σ'_i are identical for the first N periods, so they induce the same action profile s^τ for $\tau = t, \dots, t + N - 1$. They can differ only from period $t + N$ onward, and the

normalized payoff contribution from that tail—for *any* strategy pair—is at most

$$\left| (1 - \delta) \sum_{\tau=t+N}^{\infty} \delta^{\tau-t} u_i(s^\tau) \right| \leq (1 - \delta) \sum_{\tau=t+N}^{\infty} \delta^{\tau-t} \bar{u} = \delta^N \bar{u}.$$

Hence the difference in continuation values between $\sigma_i^{(N)}$ and σ'_i is at most $2\delta^N \bar{u}$. Choosing N so that $2\delta^N \bar{u} < \Delta$ —possible because $\delta^N \rightarrow 0$ —ensures

$$V(\sigma_i^{(N)}, \sigma_{-i}^* |_{h^{t-1}}) \geq V(\sigma'_i, \sigma_{-i}^* |_{h^{t-1}}) - 2\delta^N \bar{u} > V(\sigma_i^* |_{h^{t-1}}),$$

so $\sigma_i^{(N)}$ is still strictly profitable relative to σ_i^* , while deviating in only finitely many periods.

Step 2 (Backward replacement). Let $T = \{\tau \in \{t, \dots, t + N - 1\} : \sigma_i^{(N), \tau} \neq \sigma_i^{*, \tau} \text{ on the realized history}\}$. Enumerate $T = \{\tau_1 < \dots < \tau_K\}$, so τ_K is the last period where $\sigma_i^{(N)}$ deviates ($K \leq N$, since not every period in the truncation window must actually be a deviation period).

Consider the strategy $\tilde{\sigma}_i$ that agrees with $\sigma_i^{(N)}$ everywhere except at τ_K , where it plays σ_i^{*, τ_K} (conditional on the relevant history). The strategies $\sigma_i^{(N)}$ and $\tilde{\sigma}_i$ are identical in periods $t, \dots, \tau_K - 1$, so they induce identical histories up to period τ_K . At τ_K , $\sigma_i^{(N)}$ takes a one-shot deviation from $\tilde{\sigma}_i$ (after τ_K , both agree with σ_i^*). By the no-profitable-one-shot-deviation assumption, the continuation value at τ_K from $\sigma_i^{(N)}$ is no greater than from $\tilde{\sigma}_i$.

Because earlier periods contribute identically, the total continuation value from h^{t-1} satisfies

$$V(\sigma_i^{(N)}, \sigma_{-i}^* |_{h^{t-1}}) \leq V(\tilde{\sigma}_i, \sigma_{-i}^* |_{h^{t-1}}).$$

Chaining this with Step 1's strict inequality $V(\sigma_i^{(N)}, \sigma_{-i}^* |_{h^{t-1}}) > V(\sigma_i^* |_{h^{t-1}})$ preserves strictness:

$$V(\tilde{\sigma}_i, \sigma_{-i}^* |_{h^{t-1}}) \geq V(\sigma_i^{(N)}, \sigma_{-i}^* |_{h^{t-1}}) > V(\sigma_i^* |_{h^{t-1}}).$$

So $\tilde{\sigma}_i$ is also *strictly* profitable relative to σ_i^* . But $\tilde{\sigma}_i$ has only $K - 1$ deviation periods.

Step 3 (Induction). Iterate Step 2: after K replacements we arrive at σ_i^* itself, whose continuation value equals that of σ_i^* , contradicting that each replacement preserves strict profitability relative to σ_i^* .

Hence no such profitable multi-period deviation σ'_i exists; σ^* is SPE. ■

Remark (Why a Local Check Suffices: The Single-Agent Reduction).

The repeated-game structure turns the verification of SPE into a single-player decision problem at each history. Once σ_{-i}^* is fixed, player i 's problem is no longer a strategic interaction: every other player's behavior is pinned down, so i faces an environment that responds to her actions in a deterministic (or merely stochastic) but non-strategic way. From this single-agent vantage, asking “can i improve on σ_i^* ?” is the same question dynamic programming asks of any sequential decision problem, and that question localizes—improvement requires improvement at *some* period, which the one-shot check rules out.

The assumption doing the work is **discounting**. Because $\delta < 1$, tails of infinite deviations are dominated by their first N periods, which is what lets the truncation step collapse a multi-period deviation to a single-period one. Without discounting—or, relatedly, on a finite horizon with sufficiently rich continuation structure—the principle requires additional assumptions or fails outright.

Example (Cooperative Outcomes in the Prisoner’s Dilemma).

Consider the following stage game (Prisoner’s Dilemma):

	C	D
C	2, 2	−1, 3
D	3, −1	0, 0

In the stage game, D is strictly dominant for both players, so the unique Nash equilibrium is (D, D) with payoff $(0, 0)$. The cooperative outcome (C, C) with payoff $(2, 2)$ is Pareto superior but not a Nash equilibrium. We ask: can cooperation be sustained in the infinitely repeated version $G^\delta(\infty)$?

Consider the *grim trigger* strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$, where each player i plays:

$$\sigma_i^{*,1} = C, \quad \sigma_i^{*,t} = \begin{cases} C & \text{if } s^\tau = (C, C) \text{ for all } \tau < t, \\ D & \text{otherwise.} \end{cases}$$

Under this strategy, both players cooperate as long as no deviation has ever occurred. If any player deviates even once, both players switch to D forever (the “punishment”).

Claim: Grim Trigger Sustains Cooperation

For $\delta \geq \frac{1}{3}$, the grim trigger profile σ^* is an SPE of $G^\delta(\infty)$ with average discounted payoff $(2, 2)$.

Proof for Claim.

By the one-deviation principle, it suffices to check that no player has a profitable one-shot deviation after *any* history. There are two types of subgames to consider.

(a) Cooperative subgame. The history is (C, C) every period. Under σ^* , player i plays C forever and earns 2 each period:

$$U_i^{\text{coop}} = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \cdot 2 = 2.$$

Consider a one-shot deviation to D in the current period. Player i earns 3 today (since the opponent, following σ^* , still plays C), but this triggers the punishment phase from tomorrow on. From tomorrow, the opponent plays D forever, so player i 's best continuation is also D forever, earning 0 per period:

$$U_i^{\text{dev}} = (1 - \delta) \left[3 + \sum_{t=2}^{\infty} \delta^{t-1} \cdot 0 \right] = 3(1 - \delta).$$

The deviation is unprofitable iff $2 \geq 3(1 - \delta) \iff \delta \geq \frac{1}{3}$.

(b) Punishment subgame. Some prior history contains a deviation. Under σ^* , both players play D forever, so player i earns 0 per period, giving continuation payoff 0. A one-shot deviation to C yields -1 today (against opponent's D); from tomorrow, the opponent still plays D forever (the trigger is already on), so i 's best continuation is D , earning 0. Total: $(1 - \delta)(-1) + \delta \cdot 0 = -(1 - \delta) < 0$. Deviation is strictly unprofitable for every $\delta \in (0, 1)$.

Hence σ^* is an SPE whenever $\delta \geq \frac{1}{3}$, with on-path payoff $(2, 2)$. ■

Why can the opponent's strategy be "fixed" at grim trigger when computing the deviation payoff? This is simply the definition of (subgame) Nash equilibrium, not a substantive claim about the opponent's rationality. Checking whether σ^ is an SPE means asking: given that σ_{-i}^* is the opponent's prescribed behavior, does any strategy do better for player i than σ_i^* ? The claim is not that the opponent will stick to grim trigger; rather, that i 's best response to grim-trigger behavior is itself grim trigger. The credibility of the punishment is a separate and crucial point: it must be a NE of the continuation game for the opponent to want to carry it out. That is exactly what condition (ii) of SPE guarantees, and it is what part (b) of the proof above verifies—starting D forever is a NE after a deviation, since (D, D) is already the stage-game NE.*

When players are sufficiently patient ($\delta \geq \frac{1}{3}$), the threat of permanent punishment is credible and severe enough to deter short-run deviations.

8.1.4 Strategies as Automata

Strategies in $G^\delta(\infty)$ are infinite sequences of history-dependent maps. Writing them out period-by-period is unwieldy and obscures the recursive structure that makes them tractable. A more compact representation: encode the strategy as a finite-state automaton whose state summarizes everything about the past that the strategy actually conditions on.

Definition 8.6: Automaton Representation of a Strategy Profile

An **automaton** is a 4-tuple (Q, q_0, f, τ) where


- Q is a finite set of **states**;
- $q_0 \in Q$ is the **initial state**;
- $f : Q \rightarrow S$ is the **output function**, prescribing the action profile played in state q ;
- $\tau : Q \times S \rightarrow Q$ is the **transition function**, mapping (current state, realized action profile) to the next-period state.

An automaton induces a strategy profile σ : in period t , the state q^t is computed from the initial state by iterated application of τ , and the prescribed action profile is $f(q^t)$.

The state q^t is a sufficient statistic for the past: two histories that map to the same state will be played identically going forward. Many natural strategies are finite-state, even though the full history space is infinite—this is the source of compactness.

Example (Grim Trigger as a 2-State Automaton).

In the Prisoner's Dilemma of Section 11.1, the grim-trigger profile from the example is captured by:



$Q = \{q_C, q_D\}$, $q_0 = q_C$, output $f(q_C) = (C, C)$ and $f(q_D) = (D, D)$. The transition function holds in q_C as long as the realized profile is (C, C) and otherwise switches permanently to q_D . SPE verification reduces to a one-shot deviation check at each of the two states.

Example (Tit-for-Tat as a 4-State Automaton).

Tit-for-tat conditions on the previous-period action of each player. The relevant state is the action profile played *last period*, so $Q = \{(C, C), (C, D), (D, C), (D, D)\}$ with transitions $\tau(q, s) = s$ (the current realized profile becomes the next state) and outputs:

$$f(C, C) = (C, C), \quad f(C, D) = (D, C), \quad f(D, C) = (C, D), \quad f(D, D) = (D, D).$$

Player i 's rule is: "in state (s_1, s_2) , play what your opponent played last period." The state graph has every node connected to every other (each new realized profile becomes the next state), making it less compact visually than grim trigger but no less finite-state.

Remark (Why Automata Help).

Three reasons. First, *verification*: SPE under σ^* is equivalent to “no profitable one-shot deviation at every reachable state of the automaton,” which is a check over $|Q|$ states rather than over the infinite history tree. This is exactly the structure used to analyze grim trigger and tit-for-tat in problem sets. Second, *intuition*: many subtle strategies—“one period of punishment then forgive,” “increase punishment severity each violation”—are easy to describe by drawing the state diagram and unwieldy to describe by writing out $\sigma_i^t(h^{t-1})$ in closed form. Third, *forward reference*: Section 11.4 (chain store paradox) implicitly partitions the entrants $1, \dots, K$ into a “reputation phase” state and an “endgame phase” state, with the cutoff k^* as the transition; the same automaton formalism applies, with the wrinkle that beliefs play the role of state.

8.2 The Folk Theorem

We now turn to the central question of repeated games: which payoff profiles can be sustained as SPE outcomes when players are sufficiently patient? The *Folk Theorem* provides a remarkably broad answer.

8.2.1 Feasible and Individually Rational Payoffs

Fix a stage game $G = (S_i, u_i)_{i=1}^n$. Write $u(S) := \{u(s) : s \in S\} \subseteq \mathbb{R}^n$ for the image of the payoff function: the set of all payoff *vectors* $(u_1(s), \dots, u_n(s))$ that arise from some pure action profile $s \in S$. For a finite stage game, $u(S)$ is just the finite collection of cells of the payoff matrix, viewed as points in \mathbb{R}^n .

Definition 8.7: Feasible Set of Payoffs

The **feasible set** of payoffs is the convex hull^a of the pure-action payoff profiles:

$$F = \text{co } u(S) \subseteq \mathbb{R}^n.$$

^aThe **convex hull** of a set $A \subseteq \mathbb{R}^n$, written $\text{co}(A)$, is the set of all *finite convex combinations* of points in A :

$$\text{co}(A) = \left\{ \sum_{k=1}^K \lambda_k a_k : K \in \mathbb{N}, a_k \in A, \lambda_k \geq 0, \sum_{k=1}^K \lambda_k = 1 \right\}.$$

Equivalently, it is the smallest convex set containing A . Geometrically, if A is a finite set of points in \mathbb{R}^2 , $\text{co}(A)$ is the polygon obtained by stretching a rubber band around the points. For a stage game with finite S , $u(S)$ is a finite set of payoff vectors, and $F = \text{co } u(S)$ is the polygon (or polytope, in higher dimensions) of all attainable average payoffs.

Any $u \in F$ can be achieved by appropriate randomization over (or time-averaging across) pure action profiles. This is a purely *mathematical* consequence of the definition of convex hull—any point in $\text{co } u(S)$ is by definition a convex combination $\sum_k \lambda_k u(s_k)$ with $\lambda_k \geq 0$ and $\sum_k \lambda_k = 1$, which is realized either by a public randomization device that draws s_k with probability λ_k each period, or by playing each s_k in a λ_k -fraction of periods and averaging over time. It is *not* the Folk Theorem; the Folk Theorem (proved later) is the much stronger claim that every $u \in F^*$ can be supported as an *SPE* outcome, which involves incentive compatibility, not just feasibility.

To define the lower bound on what players can be forced to accept, we introduce the *minmax payoff*.

Definition 8.8: Minmax Payoff

Player i 's **minmax payoff** is:

$$v_i = \min_{s_{-i}} \max_{s_i} u_i(s_i, s_{-i}).$$

This is the lowest payoff that the other players can force player i to accept, given that player i best-responds. Equivalently, letting m_{-i} denote the minmaxing action profile by $-i$:

$$v_i = \min_{s_{-i}} u_i(\text{BR}_i(s_{-i}), s_{-i}) = u_i(\text{BR}_i(m_{-i}), m_{-i}),$$

and for all s_i , $u_i(s_i, m_{-i}) \leq v_i$.


Example (Computing Minmax Payoffs).

Consider the stage game:

	L	R
U	3, 2	5, 1
D	4, 0	1, -1

To compute v_1 : player 2 chooses s_2 to minimize the maximum payoff player 1 can guarantee. If 2 plays L , player 1's best response gives $\max\{3, 4\} = 4$. If 2 plays R , player 1 gets $\max\{5, 1\} = 5$. So $v_1 = \min\{4, 5\} = 4$, with $m_2 = L$. Analogously, $v_2 = 0$ with $m_1 = D$. In this particular game, the two minmaxing actions happen to form the pure profile (D, L) , which delivers exactly $(v_1, v_2) = (4, 0)$.

The geometry is shown below. The four pure-action payoff profiles are plotted in (u_1, u_2) space; the shaded quadrilateral is their convex hull F . The minmax thresholds $u_1 = v_1 = 4$ and $u_2 = v_2 = 0$ carve out the feasible and individually rational set F^* (the darker triangular region in the upper-right).



Definition 8.9: Feasible and Individually Rational Set

The **feasible and individually rational set** is:

$$F^* = \{u \in F \mid u_i \geq v_i \ \forall i\}.$$

These are the payoff profiles that are both achievable and weakly preferred by every player to their minmax payoff.

8.2.2 Necessity: NE Payoffs Lie in F^* **Proposition 8.10: NE Payoffs Are Individually Rational**

In any Nash equilibrium of $G^\delta(\infty)$, the discounted average payoff profile lies in F^* .

Proof for Proposition.

Two parts:

- (i) **NE payoffs lie in F .** The average discounted payoff in any strategy profile is a convex combination of stage-game payoff vectors $u(s)$ for $s \in S$. Hence the average payoff lies in $F = \text{co } u(S)$.
- (ii) **NE payoffs satisfy $u_i \geq v_i$.** Fix a NE σ^* and any history h^{t-1} . Let $\alpha_{-i}^t \in \Delta(S_{-i})$ denote the (possibly correlated) distribution over s_{-i}^t that σ_{-i}^* induces at h^{t-1} . In a NE, player i best-responds against this distribution, so

$$\mathbb{E}[u_i(s_i^t, s_{-i}^t) \mid h^{t-1}] \geq \max_{s_i \in S_i} \mathbb{E}_{s_{-i}^t \sim \alpha_{-i}^t} [u_i(s_i, s_{-i}^t)] \geq v_i,$$

where the last inequality uses the definition of $v_i = \min_{\alpha_{-i}} \max_{s_i} \mathbb{E}_{\alpha_{-i}} [u_i(s_i, s_{-i})]$ (extending the minmax to mixed s_{-i}). Averaging across periods, player i 's NE payoff is $\geq v_i$.

Combining, NE payoffs $\in F^*$. ■

Remark (Minmax Is Not a Belief—It Is a Guarantee).

The minmax bound is sometimes mis-read as “player i *expects* the opponents to be hostile, so plays the maxmin action.” That is not the content of the proposition. Rather, v_i is the payoff player i can secure *no matter what the opponents do*: against any (possibly correlated, possibly mixed) strategy α_{-i} , player i 's best response delivers at least v_i in expectation. Two caveats clarify the subtlety the reader may have in mind.

Simultaneous moves. Player i does not observe s_{-i}^t before choosing s_i^t . The bound still binds because it is formulated in expectation: against α_{-i} , playing $\text{BR}_i(\alpha_{-i})$ —the best response to the *distribution*, which does not require seeing the realization—already guarantees $\geq v_i$ in expectation.

“Securing” v_i in a specific example. In the 2×2 table above, $v_1 = 4$ is attained when 2 plays L , and $\text{BR}_1(L) = D$. If instead 2 plays R , player 1's best response is U giving

$5 > v_1$. So v_1 is not the payoff i gets by *always* playing D ; it is the payoff i gets by *best-responding to whatever $-i$ actually plays*, and this payoff is $\geq v_1$ regardless of $-i$'s choice. The worst case for i is exactly when $-i$ plays the minmaxing s_{-i} , which is why v_i is a floor.

8.2.3 Sufficiency: The Folk Theorem

The converse—every payoff in F^* can be sustained as an SPE for sufficiently patient players—is the celebrated **Folk Theorem**.

Theorem 8.11: Folk Theorem

Let $u \in F^*$. Then there exists $\bar{\delta} < 1$ such that for all $\delta > \bar{\delta}$, there exists an SPE of $G^\delta(\infty)$ whose average discounted payoff profile is u .

Remark (Pointwise vs. Set-Limit Formulation).

The theorem above is the **pointwise** formulation: each target payoff $u \in F^*$ comes with its own patience threshold $\bar{\delta}(u)$. An equivalent **set-limit** formulation is sometimes more revealing: let $E(\delta) \subseteq \mathbb{R}^n$ denote the set of SPE average payoffs of $G^\delta(\infty)$. Then

$$\lim_{\delta \rightarrow 1} E(\delta) = F^*$$

(in the Hausdorff sense). Equivalently, $E(\delta)$ is monotonically expanding in δ —more patience unlocks more SPE outcomes, never fewer—and fills out all of F^* in the limit. The proposition of Section 11.2.2 gave the easy inclusion $E(\delta) \subseteq F^*$ for every δ ; the Folk Theorem is the non-trivial reverse inclusion in the limit. Neither formulation claims that $E(\delta) = F^*$ for any particular $\delta < 1$: the patience threshold is genuinely needed.

Proof for Theorem

We prove the result for $n = 2$ and assume that there exists a pure action profile s with $u(s) = u$ and $u_i > v_i$ strictly for each i . (The general case uses public randomization over stage profiles or time-averaging; the construction below adapts naturally.)

Let m_i denote the action that minmaxes player $j \neq i$, and write $u_i(m) := u_i(m_1, m_2)$ for i 's payoff when both players execute their minmaxing actions. By the definition of v_i as $\max_{s'_i} u_i(s'_i, m_{-i})$, any action i plays against m_{-i} yields at most v_i ; in particular, $u_i(m) \leq v_i$. Let $\bar{u}_i := \max_{s'_i} u_i(s'_i, s_j)$ denote the one-period temptation payoff on the cooperative path.

Consider the following strategy profile σ^* for each player i , parametrized by an integer $k \geq 1$:

1. **Period 1:** Play s_i .
2. **Cooperative phase:** In period $t > 1$, play s_i if the history is the all- s path (s, s, \dots, s) .
3. **Punishment phase:** If the history ever shows a period in which s was not played,

play m_i for k consecutive periods.

4. **Resume / restart:** After a completed k -period punishment, return to (2). Any deviation from phase (3) restarts the k -period punishment from scratch.

We will choose k^* and $\bar{\delta} < 1$ so that σ^* is an SPE for every $\delta > \bar{\delta}$. By the one-deviation principle, it suffices to rule out profitable one-shot deviations at two kinds of information sets: a round on the cooperative path, and a round within a punishment phase.

Step 1: Deviation from the cooperative path. Compare the two continuation streams over the $(k+1)$ -period window starting at the deviation period; in both, play returns to s from period $k+2$ onward, so the common continuation cancels. Following the prescription yields $(1+\delta+\dots+\delta^k)u_i(s)$ over the window. The best one-shot deviation yields \bar{u}_i today, then k periods of the minmax payoff $u_i(m)$, contributing $\bar{u}_i + (\delta + \delta^2 + \dots + \delta^k)u_i(m)$. The no-deviation inequality is

$$(1 + \delta + \dots + \delta^k) u_i(s) \geq \bar{u}_i + (\delta + \dots + \delta^k) u_i(m),$$

which rearranges to

$$(\delta + \delta^2 + \dots + \delta^k) [u_i(s) - u_i(m)] \geq \bar{u}_i - u_i(s). \quad (\text{S1})$$

The bracket is strictly positive (since $u_i(s) = u_i > v_i \geq u_i(m)$). As $\delta \uparrow 1$ the LHS tends to $k[u_i(s) - u_i(m)]$, while the RHS is a fixed finite number. Hence (S1) binds on k : one must pick k large enough to outrun the temptation.

Step 2: Deviation within a punishment round. Consider any round of an ongoing punishment (the one-deviation principle lets us check only one). Let V^P denote the continuation value from that round onward if the prescription is followed: k more periods of $u_i(m)$, then permanent return to $u_i(s)$:

$$V^P = (1 + \delta + \dots + \delta^{k-1}) u_i(m) + \frac{\delta^k}{1 - \delta} u_i(s).$$

A one-shot deviation to the best response against m_{-i} yields at most v_i today, after which phase (4) restarts a fresh k -period punishment—whose continuation value from tomorrow is again V^P . So the deviation continuation is bounded by $v_i + \delta V^P$. The no-deviation inequality $V^P \geq v_i + \delta V^P$ reduces to $(1 - \delta) V^P \geq v_i$, i.e.

$$(1 - \delta^k) u_i(m) + \delta^k u_i(s) \geq v_i,$$

equivalently

$$\delta^k [u_i(s) - u_i(m)] \geq v_i - u_i(m). \quad (\text{S2})$$

At $\delta = 1$ this becomes $u_i(s) - u_i(m) \geq v_i - u_i(m)$, i.e. $u_i(s) \geq v_i$, which holds strictly since $u_i(s) = u_i > v_i$. By continuity, (S2) holds for all δ sufficiently close to 1 for any fixed k .

Choosing k^ and $\bar{\delta}$.* The binding constraint is (S1), because the RHS $\bar{u}_i - u_i(s) \geq 0$ is a fixed positive quantity that the finite geometric sum on the LHS must overcome. Pick

any integer k^* satisfying

$$k^* [u_i(s) - u_i(m)] > \bar{u}_i - u_i(s) \quad \text{for each } i \in \{1, 2\};$$

such a k^* exists because the bracket is strictly positive. At $\delta = 1$, (S1) holds at $k = k^*$ with strict slack, so there is $\bar{\delta}_1 < 1$ such that (S1) holds for every $\delta \in (\bar{\delta}_1, 1)$. Similarly (S2) gives a threshold $\bar{\delta}_2 < 1$. Setting $\bar{\delta} := \max\{\bar{\delta}_1, \bar{\delta}_2\} < 1$ delivers an SPE of $G^\delta(\infty)$ with average payoff u for every $\delta > \bar{\delta}$. ■

Remark (The Punishment-Phase Intuition: “Delay the Recovery”).

The substantive content of Step 2 is that *minmaxing is itself incentive-compatible*: why does player i obey the prescription to play m_i , which is chosen to hurt j and is not generally i 's best response to m_{-i} ? Inequality (S2) answers cleanly—the only effect of a one-shot deviation during punishment is to *restart the clock on recovery*. Since $u_i(s) > u_i(m)$ strictly, any such delay is costly; with δ close to 1, the δ^k -weighted loss from postponing cooperation dominates the bounded one-period gain of $v_i - u_i(m) \geq 0$ from deviating. The self-referential restart in phase (4)—“a deviation from phase (3) is met with a fresh phase (3)”—is precisely what elevates the construction from a Nash equilibrium of $G^\delta(\infty)$ to a *subgame-perfect* one. Without it, minmaxing during punishment would itself be a profitable one-shot deviation to skip.

Two further points are worth underlining. First, the binding constraint on patience is (S1), not (S2): (S2) has strict slack at $\delta = 1$ for *any* k , so it is Step 1 that forces k large and in turn forces δ close to 1. Second, the role of k is asymmetric across the two steps—larger k strengthens (S1) by extending the punishment, but *weakens* (S2) slightly because δ^k shrinks. Both tensions resolve in the limit $\delta \rightarrow 1$; the finite- δ threshold $\bar{\delta}$ depends on the specific stage game through $\bar{u}_i - u_i(s)$ and $v_i - u_i(m)$.

Remark (The Folk Theorem in Words).

The Folk Theorem states that *anything that is feasible and individually rational can be sustained as an SPE*, provided players are patient enough. This is both a celebration and a caution: while cooperation, fair sharing, and many other socially desirable outcomes can be supported, so can a vast multitude of other outcomes—including very inefficient or asymmetric ones. The theorem highlights the fundamental **indeterminacy** of repeated games: the equilibrium concept alone does not pin down a unique prediction, and selection among the many SPE requires additional considerations (such as renegotiation-proofness, evolutionary stability, or focal points).

Example (Collusion in Cournot Duopoly).

The Folk Theorem's canonical application is **tacit collusion** in oligopoly. Consider a symmetric Cournot duopoly with inverse demand $P(Q) = a - Q$ and constant marginal cost $c < a$, where $Q = q_1 + q_2$. In the one-shot game, the unique Nash equilibrium

is the *Cournot outcome* with $q_i^N = (a - c)/3$ per firm, yielding per-firm profit $\pi^N = (a - c)^2/9$. The *collusive* outcome would be to split monopoly production: each firm produces $q_i^M = (a - c)/4$, total output equals the monopoly output $Q^M = (a - c)/2$, and per-firm profit is $\pi^M = (a - c)^2/8 > \pi^N$. Collusion is not an NE of the stage game: holding $q_j = (a - c)/4$ fixed, firm i 's best response is $q_i^D = 3(a - c)/8$ (the “cheat” quantity), yielding $\pi^D = 9(a - c)^2/64 > \pi^M$.

Consider the infinitely repeated game $G^\delta(\infty)$ with grim-trigger strategies: each firm plays q_i^M as long as the history is all-collusive, and reverts permanently to the Cournot output q_i^N after any deviation. By the one-deviation principle, this is an SPE iff

$$\underbrace{\frac{\pi^M}{1 - \delta}}_{\text{collusion forever}} \geq \underbrace{\pi^D}_{\text{cheat today}} + \underbrace{\frac{\delta \pi^N}{1 - \delta}}_{\text{Cournot forever after}} .$$

Plugging in and solving,

$$\delta \geq \bar{\delta} = \frac{\pi^D - \pi^M}{\pi^D - \pi^N} = \frac{9/64 - 1/8}{9/64 - 1/9} = \frac{9}{17} \approx 0.529.$$

Whenever the common discount factor exceeds ≈ 0.53 , the collusive monopoly split is sustainable as an SPE—no written contract required, and no coordination beyond the common trigger rule.

Because the Cournot NE is itself an NE of the stage game, reverting to q^N is a credible threat: phase (b) of the proof structure is automatic. This is why Cournot-reversion is the textbook punishment for tacit collusion; stronger punishments (full min-max, which can drive profits to zero) would sustain collusion for an even lower $\bar{\delta}$, at the cost of requiring the “delay-the-recovery” machinery to enforce the min-max itself.

Remark (Why the Folk Theorem Is Substantive).

It is tempting to dismiss the Folk Theorem as a formal triviality—“if you threaten enough, you can sustain anything.” The Cournot example above shows why this reading undersells it. The theorem provides a *rationalization of observed behavior* in settings where explicit contracts are infeasible or illegal: tacit collusion among firms, resource-sharing among non-kin animals, peace-keeping among tribes without a central authority, self-enforcing treaties between nations. In each case, repeated play converts a one-shot game that would dissipate into defection into a long-run game in which cooperation is self-enforcing. The Folk Theorem identifies the precise patience threshold, the structure of credible punishments, and the range of sustainable outcomes—making it one of the most empirically deployed results in microeconomic theory.

8.3 Finitely Repeated Games

The results of the previous two sections all took the horizon to be infinite. A natural question is what survives when G is played only finitely many times. The short answer splits by stage-game multiplicity. When the stage game has a *unique* NE, finite repetition

is too weak to sustain anything but stage-NE play in every period (Section 11.3.1)—weaker than one might hope. When the stage game has *multiple* NE with payoff variation, finite-horizon folk-theorem-style results return (Section 11.3.2)—stronger than one might hope. Multiple stage NE are not exotic: coordination games, battle of the sexes, Cournot games with entry, and many others routinely admit several pure or mixed NE.

8.3.1 Setup

Fix a stage game $G = (S_i, u_i)_{i=1}^n$ and a horizon $T \in \mathbb{N}$. The **finitely repeated game** $G(T)$ consists of playing G in each period $t = 1, 2, \dots, T$. Histories, strategies, and subgames are defined exactly as in the infinite case but truncated at T . For most of what follows, discounting is inessential¹; we use the **arithmetic average** payoff

$$U_i(\sigma) = \frac{1}{T} \sum_{t=1}^T u_i(s^t(\sigma)),$$

which coincides with the long-run stage payoff when a constant profile is played every period. All statements below go through unchanged if one instead uses the discounted average $(1 - \delta)/(1 - \delta^T) \cdot \sum_{t=1}^T \delta^{t-1} u_i(s^t)$; the finite horizon is what is doing the work, not the absence of discounting.

Because the horizon is finite and the game is one of perfect information across periods, SPE of $G(T)$ can be found by **backward induction**: solve the period- T subgames first, plug the continuation values into period $T - 1$, and so on.

8.3.2 Unique Stage-Game NE: Backward Induction Collapse

Proposition 8.12: Unique-NE Backward Induction

Suppose the stage game G has a unique Nash equilibrium $s^* \in S$. Then $G(T)$ has a unique SPE outcome path: $s^t = s^*$ for every $t = 1, \dots, T$.

Proof for Proposition.

By backward induction. In the last period, the continuation is empty, so every subgame starting at $t = T$ is just G itself; its unique NE prescribes s^* . With period- T payoffs fixed at $u(s^*)$ on every possible history, the period- $(T - 1)$ subgame is strategically equivalent to G (the period- T continuation is a constant that enters additively into every player's payoff and therefore does not affect best responses); its unique NE again prescribes s^* . Iterating, s^* is played in every period along the unique SPE path. ■

Remark (Finite-Horizon Prisoner's Dilemma).

In the Prisoner's Dilemma of Section 11.1, reproduced here for convenience,

¹Discounting played two roles in the infinite-horizon case: (i) ensuring the payoff sum converges, and (ii) driving the truncation step in the proof of the one-deviation principle, which collapsed an infinite deviation into a long-but-finite one. Both roles are vacuous in finite horizon—the sum is automatically finite, and there is nothing to truncate—so the SPE set is the same whether one uses arithmetic average, time-discounted average, or any monotone transformation of the two.

	<i>C</i>	<i>D</i>
<i>C</i>	2, 2	-1, 3
<i>D</i>	3, -1	0, 0

(D, D) is the unique stage NE. The proposition says that the only SPE of the T -period game, for any finite T , prescribes (D, D) in every period—cooperation unravels. The contrast with the grim-trigger equilibrium of $G^\delta(\infty)$ is stark: patience sustains cooperation in the infinite-horizon game precisely because there is no last period from which to unravel. In a finite horizon, the last period pins down (D, D) , which pins down (D, D) in the period before, and so on all the way back.

8.3.3 Multiple Stage-Game NE: Using Continuation NE as Rewards and Punishments

When G has multiple pure NE, the backward-induction collapse breaks down: a subgame need not have a unique continuation equilibrium, so earlier periods can support non-NE stage play by using different continuation equilibria as carrots and sticks.

Example (Two-Period Game with Two Stage-Game NE).

Consider the stage game

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	5, 5	0, 1	0, 6
<i>M</i>	1, 0	3, 3*	0, 0
<i>D</i>	6, 0	0, 0	2, 2*

The pure-strategy Nash equilibria of G are (M, M) with payoff $(3, 3)$ and (D, R) with payoff $(2, 2)$. Note that neither is Pareto-dominant, and the cooperative profile (U, L) yields $(5, 5)$, which is not an NE: against L , player 1 strictly prefers D ($6 > 5$); against U , player 2 strictly prefers R ($6 > 5$).

Consider the two-period game $G(2)$. Three pure-strategy SPE outcome paths are worth noting.

- Repeating (M, M) for both periods.** Play (M, M) unconditionally. Average payoff $(3, 3)$. SPE because in the period-2 subgame (M, M) is NE, and in period 1 each player has no profitable one-shot deviation given the opponent plays M and the period-2 continuation is fixed at (M, M) .
- Repeating (D, R) for both periods.** Analogous, with average payoff $(2, 2)$.
- “Cooperate then reward”—the non-trivial SPE.** Play (U, L) in period 1. If (U, L) was realized, play (M, M) in period 2 (reward); if *anything else* occurred, play (D, R) in period 2 (punishment). Average payoff $(5+3)/2 = 4$ for each player.

Verification of (c) via the one-deviation principle. In period 2, both (M, M) and (D, R) are stage-game NE, so no one-shot deviation in the final period is profitable under either continuation rule. In period 1, suppose player 1 contemplates a deviation from U given opponent plays L . The best stage-game deviation is D (payoff 6 vs. on-path 5), so the temptation is

$$\bar{u}_1 - u_1 = 6 - 5 = 1.$$

A deviation, however, triggers the period-2 punishment (D, R) instead of the reward (M, M) : the period-2 payoff drops from 3 to 2, a loss of 1. The sum of period payoffs under deviation is $6 + 2 = 8$; on path it is $5 + 3 = 8$. Player 1 is indifferent, hence has no *strictly* profitable one-shot deviation. The symmetric computation for player 2 (temptation R , gain $6 - 5 = 1$; punishment loss $3 - 2 = 1$) gives the same conclusion.

Because the deviation condition just barely binds, (c) is a *weak* SPE. A small perturbation of the stage-game payoffs—say, raising the reward NE payoff to $(3 + \varepsilon, 3 + \varepsilon)$ or lowering the punishment NE payoff to $(2 - \varepsilon, 2 - \varepsilon)$ —would make the deterrence strict.

Remark (Why $G(2)$ Escapes the Backward-Induction Collapse).

The proof of Proposition 11.2 required the period- T subgame to have a *unique* NE. When multiple NE exist, the continuation equilibrium played in period T can depend on the period- $(T - 1)$ outcome—and this *conditioning* is what gives earlier periods leverage. The gap between the reward NE payoff (3) and the punishment NE payoff (2) acts as an incentive device: deviating today trades a one-period gain against a future-period loss. As long as the gain does not exceed the loss, cooperation is sustainable.

8.3.4 The Benoit-Krishna Folk Theorem (Sketch)

The idea behind Example 11.2 generalizes. If the stage game has enough distinct NE payoffs to reward and punish each player separately, one can chain reward/punishment phases over many periods and recover a finite-horizon analogue of the folk theorem.

Theorem 8.13: Benoit-Krishna (1985)

Suppose G has, for each player i , two pure-strategy NE $s^{(i)}, s'^{(i)}$ with $u_i(s^{(i)}) > u_i(s'^{(i)})$. Then for every $u \in F^*$ (with u in the interior, and for u achievable in pure strategies) and every $\varepsilon > 0$, there exists \bar{T} such that for all $T \geq \bar{T}$, the finitely repeated game $G(T)$ has an SPE whose average payoff is within ε of u .

The construction allocates the T periods into three phases: a long “main” phase of length $\approx T - K$ during which the target profile is played, followed by a “last- K ” phase in which the NE $s^{(i)}$ is played if no deviation occurred and $s'^{(i)}$ is played otherwise. As T grows, the K -period reward/punishment phase has a per-period weight that vanishes in the average, so the average payoff converges to u ; but the reward/punishment gap remains fixed in absolute terms, so deterrence continues to work.

Remark (Contrast with the Infinite-Horizon Folk Theorem).

The infinite-horizon Folk Theorem (Section 11.2) can use *minmax* as the punishment threat. The construction is the one we proved earlier: when player i deviates from the cooperative path, the others minmax i for k periods and then return to cooperation. What makes this strategy profile itself a subgame perfect equilibrium—and not merely a Nash equilibrium—is the **delay-the-recovery** structure of Step 2 in that proof: deviating from the minmaxing phase *restarts* the k -period punishment, postponing the

eventual return to cooperation. With δ close to 1, the discounted loss from postponed cooperation dominates the one-period gain from refusing to minmax, so minmaxing is incentive-compatible at every step. In other words, infinite play has enough “room after the punishment” for the punishment to be sustained, and that room is what lets minmax (the harshest threat) double as the credible threat.

The finite-horizon Folk Theorem cannot replicate this construction: at horizon T , a k -period punishment phase starting near the end leaves no time for the cooperative return that disciplines the punishment phase itself. So in finite horizon one must fall back on *stage-game Nash equilibria* as punishments—these are by definition self-enforcing in any single stage, requiring no follow-up phase. This is why the hypothesis of Benoit-Krishna requires multiple NE with distinct payoffs for each player: only when the stage game has multiple NE with payoff variation does “play one NE on path, switch to another after deviation” deliver any deterrence. A Prisoner’s Dilemma, having a unique stage NE, is therefore outside the theorem’s reach: cooperation in finitely repeated Prisoner’s Dilemma remains impossible at any horizon.

8.4 The Chain Store Paradox


Proposition 11.2—that a finitely repeated game with a unique stage-game NE has a unique SPE in which the stage NE is played every period—raises a follow-up puzzle. Consider a monopolist chain store facing a sequence of K potential entrants, one per market. Intuitively, the store should *fight* early entrants to build a reputation for fighting, deterring later entrants from entering at all. Yet if fighting is dominated in each stage game taken in isolation, Proposition 11.2 rules this out: the unique SPE has every entrant entering and the store accommodating every time. Selten (1978) sharpened this tension into the **chain store paradox**, paradoxical because real-world chain stores manifestly *do* fight entrants and real-world entrants often stay out. Kreps and Wilson (1982), along with Milgrom and Roberts (1982), resolved the paradox by introducing a vanishingly small amount of incomplete information: a tiny probability that the store is “crazy” and enjoys fighting is enough to sustain reputation-building as a sequential equilibrium.

Remark (Reputation Needs Uncertainty).

Before the formal construction, it is worth pausing on why the repaired model needs incomplete information at all. Reputation, as an economic concept, only comes alive when the audience is unsure about the type of the agent they face. If the store’s payoffs were perfectly known, there would be nothing for entrants to “learn” from past fights—the fight would just be an irrational flinch that no forward-looking calculation would read as evidence of future behavior. The central question the resolution asks is: *can a normal store sustain a reputation for being crazy?*—equivalently, can it, by acting crazy, generate enough doubt in later entrants’ minds that they stay out? The answer turns out to be yes, and it works even when the prior probability of the crazy type is arbitrarily small.

8.4.1 The Stage Game

A single long-lived **store** S faces K short-lived **entrants** E_1, E_2, \dots, E_K in sequence. In period k , E_k decides whether to enter (IN) or stay out (OUT). Conditional on entry, the store decides whether to accommodate (A) or fight (F). The stage tree is:



Parameters satisfy $0 < b < 1 < a$. The first coordinate is the entrant's payoff; the second is the store's stage payoff. The store's total payoff is the undiscounted sum across the K periods.

Remark (Interpreting the Payoffs).

Entering a contested market and earning duopoly profit nets the entrant $b \in (0, 1)$, a positive but small gain relative to the outside option 0. Entering and being fought costs the entrant $1 - b > 0$. For the store: maintaining a monopoly earns rent $a > 1$; sharing the market (accommodating) earns 0; fighting costs 1 relative to accommodating, which is why fighting is dominated in isolation.

8.4.2 Selten's Paradox: Backward-Induction Collapse

Proposition 8.14: SPE of the Chain Store Game Under Common Knowledge

Assume the store's payoffs above are common knowledge. Then the unique SPE of the K -period chain store game has every entrant entering and the store accommodating in every period.

Proof for Proposition.

By backward induction. Because markets are additively separable across periods, the store's current-period choice of F vs. A does not affect any later-period payoffs whenever both types are common knowledge. So the store's choice at each IN node is governed by the stage game alone: A (payoff 0) strictly dominates F (payoff -1). Anticipating A , each entrant compares IN ($b > 0$) to OUT (0) and strictly prefers IN. Iterating from period K backward, every period features (IN, A). ■

Remark (What Makes This "Paradoxical").


Common-sense reasoning about reputation says the store should fight early entrants so that later entrants, seeing the track record, stay out. The proposition says this intuition is wrong whenever the store’s payoffs are common knowledge: fighting today does nothing to alter tomorrow’s common-knowledge rationality calculation. Selten’s contribution was to pose the tension sharply—the game-theoretic solution contradicts robust empirical behavior. Kreps and Wilson’s contribution was to show that a *tiny* perturbation of the information structure rescues the intuition.

8.4.3 Resolution: A “Crazy” Store Type


Suppose the store’s type is not common knowledge. Nature picks the type once, before period 1. With probability $1 - \varepsilon$ the store is **normal**, with the same payoffs as above. With probability ε the store is **crazy**: payoffs are reversed on the F - A axis, so that fighting yields 0 and accommodating yields -1 . Entrants observe the store’s past actions but not its type; they share the common prior ε .

The two types’ stage trees are below. The only difference is the store’s payoff in the bottom two leaves.

Normal store (prob. $1 - \varepsilon$)



Crazy store (prob. ε)



For the crazy type, F strictly dominates A at every IN node, so the crazy store *always fights* when entered—regardless of continuation, regardless of k .

This one perturbation is transformative. The normal store, knowing the crazy type always fights, has a strict incentive to *pool* with the crazy type by fighting early entrants, so that later entrants cannot distinguish normal from crazy and stay out. Accommodating, by contrast, immediately reveals the store as normal (the crazy type would never accommodate), after which every subsequent entrant enters.

8.4.4 Deriving the Equilibrium in Three Steps

We now characterize the unique sequential equilibrium of the game with one normal and one crazy type. Rather than guess the answer and verify it, we proceed forward: solve the entrant’s stage problem; back out the normal store’s value function; then close the loop by showing Bayes makes the two consistent.

Notation. Let $\mu_k \in [0, 1]$ denote the belief, held at the start of period k , that the store is crazy; the prior is $\mu_1 = \varepsilon$. Let σ_k denote the normal store’s $\Pr(F \mid \text{Normal}, \text{IN})$ in period k . Since the crazy store fights with probability 1, the unconditional fight probability conditional on IN is

$$\pi_k \equiv \Pr(F \mid \text{IN}) = \mu_k \cdot 1 + (1 - \mu_k) \cdot \sigma_k.$$

Two structural observations close off easy cases. *Once accommodated, forever revealed:* if the store ever plays A , every subsequent entrant infers $\mu = 0$ and the game collapses to Selten's complete-information paradox. *Beliefs move only on entry:* if entrant $k - 1$ chose OUT, the store has no chance to act and $\mu_k = \mu_{k-1}$. Belief updating happens only at IN-then-action events.

We will prove inductively that the equilibrium organizes itself around a geometric *threshold sequence*

$$\bar{\mu}_k \equiv b^{K-k+1}, \quad k = 1, \dots, K.$$

Step 1: Entrant's best response

Consider entrant k holding belief μ_k and conjecturing σ_k (which together fix π_k). Her expected payoff from IN is

$$\Pr(A \mid \text{IN}) \cdot b + \Pr(F \mid \text{IN}) \cdot (b - 1) = (1 - \pi_k)b + \pi_k(b - 1) = b - \pi_k.$$

The OUT payoff is 0. Hence the entrant strictly prefers IN iff $\pi_k < b$, strictly prefers OUT iff $\pi_k > b$, and is indifferent iff $\pi_k = b$. The mechanism by which π_k is generated—through the prior μ_k on the crazy type or through the normal store's mixing σ_k —is irrelevant to the entrant; only the bottom line matters.

Step 2: Normal store's value, by backward induction

Let $V_k(\mu)$ denote the normal store's expected total payoff from period k onward, given belief $\mu_k = \mu$. We compute V_k from $k = K$ backward, proving along the way the invariant

$$V_k(\bar{\mu}_k) = 1 \quad \text{for every } k = 1, \dots, K.$$

Why 1? Because at the threshold, the normal store needs to be indifferent between F (paying -1 now, harvesting some continuation) and A (paying 0 now, then exposed forever, harvesting 0). Indifference requires the harvest from F to equal 1.

Last period $k = K$. Reputation has no future use, so the normal store's stage choice given IN is A (0) over F (-1). Hence $\sigma_K = 0$, and $\pi_K = \mu_K$. By Step 1, entrant K goes OUT iff $\mu_K > b$, IN iff $\mu_K < b$, and mixes at $\mu_K = b = \bar{\mu}_K$. Let entrant K at the threshold play IN with probability q_K and OUT with probability $1 - q_K$. Then

$$V_K(\bar{\mu}_K) = q_K \cdot 0 + (1 - q_K) \cdot a.$$

Setting this to 1 pins down q_K :

$$(1 - q_K)a = 1 \implies q_K = \frac{a - 1}{a}.$$

Here is where the mixing weight $(a - 1)/a$ comes from. It is the entrant's IN-probability that leaves the normal store exactly indifferent at the threshold. (Off the threshold, $V_K(\mu_K) = 0$ for $\mu_K < \bar{\mu}_K$ and $V_K(\mu_K) = a$ for $\mu_K > \bar{\mu}_K$.)

Inductive step at $k < K$. Assume $V_{k+1}(\bar{\mu}_{k+1}) = 1$. We compute $V_k(\bar{\mu}_k)$ and verify the same invariant.

If entrant k goes OUT, no information is revealed and the store collects $a + V_{k+1}(\mu_k)$. If entrant k comes IN, the normal store compares:

$$\begin{aligned} \text{Accommodate (A): } & 0 + V_{k+1}(0) = 0, \\ \text{Fight (F): } & -1 + V_{k+1}(\mu_{k+1}), \end{aligned}$$

where μ_{k+1} is determined by Bayes given the chosen σ_k . For the normal store to mix with $\sigma_k \in (0, 1)$, the two options must be equal:

$$0 = -1 + V_{k+1}(\mu_{k+1}) \iff V_{k+1}(\mu_{k+1}) = 1.$$

By the inductive hypothesis, $V_{k+1}(\mu_{k+1}) = 1$ holds at $\mu_{k+1} = \bar{\mu}_{k+1}$. So the normal store's mixing must drag the next-period posterior to exactly the next threshold:

$$\mu_{k+1} = \bar{\mu}_{k+1} = b^{K-k}.$$

This is the indifference condition that determines σ_k ; we solve for it in Step 3.

Step 3: Bayes closes the loop

The normal store's mixing must engineer $\mu_{k+1} = b^{K-k}$. By Bayes,

$$\mu_{k+1} = \Pr(\text{Crazy} \mid F) = \frac{\mu_k \cdot 1}{\mu_k \cdot 1 + (1 - \mu_k) \cdot \sigma_k}.$$

Setting $\mu_{k+1} = b^{K-k}$ and solving:

$$b^{K-k} = \frac{\mu_k}{\mu_k + (1 - \mu_k)\sigma_k} \implies \sigma_k = \frac{1 - b^{K-k}}{b^{K-k}} \cdot \frac{\mu_k}{1 - \mu_k}.$$

This is the formula for σ_k , derived rather than stated. Substituting back into the total fight probability,

$$\pi_k = \mu_k + (1 - \mu_k)\sigma_k = \mu_k + \mu_k \cdot \frac{1 - b^{K-k}}{b^{K-k}} = \frac{\mu_k}{b^{K-k}}.$$

By Step 1, the entrant is indifferent iff $\pi_k = b$, i.e.,

$$\mu_k = b \cdot b^{K-k} = b^{K-k+1} = \bar{\mu}_k. \quad \checkmark$$

The threshold formula $\bar{\mu}_k = b^{K-k+1}$ falls out as the fixed point: each player's randomization is calibrated so that, given the other's, both are indifferent.

It remains to compute $V_k(\bar{\mu}_k)$ to close the induction. At $\mu_k = \bar{\mu}_k$, entrant k mixes at IN with probability q_k . If entrant goes IN, the normal store mixes (and is indifferent), so its conditional continuation value equals either branch, namely 0 (the value of A). If entrant goes OUT, the store earns $a + V_{k+1}(\bar{\mu}_k)$. Now $\bar{\mu}_k = b \cdot \bar{\mu}_{k+1} < \bar{\mu}_{k+1}$, which puts $\bar{\mu}_k$ below

the next-period threshold; in that region $V_{k+1}(\cdot) = 0$. Hence

$$V_k(\bar{\mu}_k) = q_k \cdot 0 + (1 - q_k) \cdot a = (1 - q_k)a.$$

The invariant $V_k(\bar{\mu}_k) = 1$ pins down q_k :

$$q_k = \frac{a - 1}{a}.$$

The same mixing weight reappears at every period: **entrant k , when at the threshold, always goes IN with probability $(a - 1)/a$.** The induction is complete.

Equilibrium summary

For each $k \in \{1, \dots, K\}$:

Entrant k , given belief μ_k :

OUT	if $\mu_k > \bar{\mu}_k = b^{K-k+1}$
IN	if $\mu_k < \bar{\mu}_k$
IN w.p. $\frac{a-1}{a}$, OUT w.p. $\frac{1}{a}$	if $\mu_k = \bar{\mu}_k$

Normal store, given belief μ_k and that entrant chose IN:

Play A	if $k = K$
Play F w.p. $\sigma_k = \frac{1-b^{K-k}}{b^{K-k}} \cdot \frac{\mu_k}{1-\mu_k}$	if $k < K$ and $\mu_k < b^{K-k}$
Play F for sure	if $k < K$ and $\mu_k \geq b^{K-k}$

Belief update from k to $k + 1$:

$\mu_{k+1} = \mu_k$	if entrant k chose OUT
$\mu_{k+1} = \max\{b^{K-k}, \mu_k\}$	after (IN, F) at period k , $\mu_k > 0$
$\mu_{k+1} = 0$	after (IN, A) at period k , or once $\mu = 0$

Remark (Why the Two Indices Differ by 1).


A common point of confusion: entrant k 's indifference threshold uses b^{K-k+1} , while the normal store's pure-fight cutoff uses b^{K-k} . The shift of one is the period-by-period erosion of reputation: at the start of period k , the entrant is indifferent at $\mu_k = b^{K-k+1}$; if a fight occurs, Bayes pushes the posterior up to b^{K-k} , which is exactly the next entrant's indifference threshold. So the threshold sequence $\{\bar{\mu}_k\}_k$ rises geometrically toward b as k approaches the endgame, and each fight ratchets the belief one step along that ladder.

Remark (The Duality of Mixing Weights).

Two probabilities are calibrated. The entrant's mix $\frac{a-1}{a}$ over IN keeps *the normal store* indifferent between F and A : the “ a factor” (rent from holding the monopoly) controls the store's incentive. The normal store's mix σ_k over F keeps *the next entrant* indifferent between IN and OUT: the “ b factor” (the entrant's per-round profit) controls the entrant's incentive. Each player randomizes not to gain anything for itself, but to fix the opponent's indifference. This is the standard logic of mixed-strategy equilibrium, played out over a sequence of belief thresholds rather than a single stage game.

8.4.5 Dynamics: Reputation and Endgame

Plot the threshold curve $k \mapsto b^{K-k+1}$ against the belief trajectory μ_k in equilibrium:



The dynamics split into two phases. Let k^* be defined implicitly by $\varepsilon = b^{K-k^*+1}$, i.e.

$$k^* = K + 1 - \frac{\log \varepsilon}{\log b}.$$

- **Reputation phase**, $k \leq k^*$. The threshold b^{K-k+1} is below the prior ε , so every entrant stays OUT. Beliefs remain at ε ; nothing is learned; the store collects monopoly rent a every period. The normal store's reputation is untested, and untested reputation is free.
- **Endgame phase**, $k > k^*$. The threshold has risen above ε , so without further updating the entrant would prefer IN to OUT. Entrant $k^* + 1$ enters pure; the normal store mixes F and A . If F is realized, the belief jumps up to b^{K-k^*} —the next entrant's indifference threshold—and from there on each entrant mixes IN/OUT while the normal store mixes F/A , with beliefs tracking the threshold curve exactly. If instead A is realized at any point, the store is revealed as normal and all remaining entrants enter unopposed.

The length of the reputation phase is $k^* = K + 1 - \log \varepsilon / \log b$. As $\varepsilon \rightarrow 0$ with K fixed, $k^* \rightarrow K + 1$: the reputation phase swallows *the entire game* and all entrants stay out. As $K \rightarrow \infty$ with ε fixed, $K - k^* = (\log \varepsilon) / (\log b) - 1$ is a *constant*, so the endgame takes up only a fixed number of periods regardless of horizon—the reputation phase takes up the proportion $(K - k^*) / K \rightarrow 0$ of a long horizon.

Contrast this with the complete-information equilibrium, in which every entrant enters and the store accommodates every period. The reputation equilibrium inverts this: *nobody* enters in the long opening phase, the store harvests monopoly rent without ever being tested, and only in the final handful of periods does any entry occur—gradually, through mixing, and only because the finite horizon makes the rent from further reputation-building smaller than the cost of fighting. The interpretive one-liner is that the normal store is *taking advantage of the fact that the entrant thinks he might be crazy, by slightly mimicking what the crazy store would do*. Randomized fighting is the cheapest way to refuse to reveal one's type.

Remark (What the Chain Store Paradox Teaches).

Three lessons. First, the backward-induction prediction of Proposition 11.2 is *fragile*: a vanishingly small perturbation of the information structure changes the equilibrium qualitatively, not just quantitatively. The prediction discontinuity at $\varepsilon = 0$ is what earned the result the word “paradox”—the limit of the reputation equilibrium as $\varepsilon \rightarrow 0$ is not the complete-information equilibrium.

Second, *reputation requires a commitment device*, and in sequential-move games against short-lived opponents the natural device is uncertainty about one’s own type. The normal store is willing to pay a short-run cost (fighting, -1 instead of 0) precisely because doing so preserves ambiguity about its type, and ambiguity is what deters future entry.

Third, the chain store paradox is less about chain stores than about how small amounts of rationality-relevant uncertainty can rescue behaviors that strict common-knowledge reasoning rules out. The same logic underlies Kreps-Milgrom-Roberts-Wilson models of finitely repeated prisoner’s dilemma (where a tiny probability of a “tit-for-tat” type rescues cooperation), reputation models in industrial organization (FTC enforcement, product quality signaling), and macroeconomic models of monetary policy commitment.

Remark (Chapter Summary).

Dynamic games extend the static equilibrium concepts of Chapter 2 to environments with explicit time. Two settings dominate. The *infinitely repeated game* $G^\delta(\infty)$ is governed by the Folk Theorem (Theorem 8.2.3): for δ sufficiently close to 1, every feasible and individually rational payoff profile is sustainable as an SPE outcome. The flip side is that this multiplicity makes prediction hard—infinately repeated games predict almost everything that is consistent with individual rationality. The *finitely repeated game* $G(T)$ is governed by Benoit-Krishna (Theorem 8.3.4): if the stage game has multiple stage-NE payoffs, almost any feasible IR payoff is sustainable for T large enough; if the stage NE is unique, only the stage-NE outcome is sustainable. The *chain store paradox* bridges to the next chapter: with one-sided private information about the long-lived player’s type, a vanishingly small mass of commitment types is enough to rescue cooperation/deterrence even when standard backward induction rules it out. The general moral: cooperation in dynamic games does not require infinity; it requires either patience plus repetition or a small grain of doubt.

Chapter 9

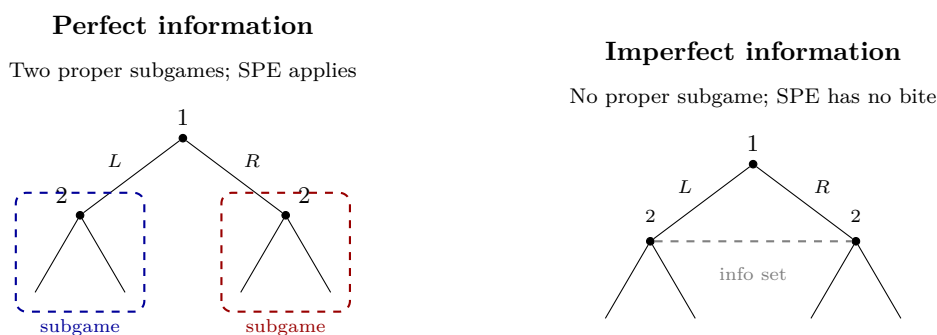
Perfect Bayesian Equilibrium

9.1 From SPE to PBE

Subgame-perfect equilibrium (Chapter 2) refines Nash equilibrium by demanding sequentially rational play in every *subgame*. The definition is powerful in extensive-form games of **perfect information**, where every singleton history initiates a subgame and SPE is verified by backward induction at each node. It is also adequate for extensive-form games of *imperfect* information whose information sets have a tree-like structure that respects subgame boundaries (e.g., simultaneous-move stages embedded in a larger sequential game).

It loses bite as soon as informational nontrivialities cut across the would-be subgames. Consider a game in which player 1 moves first, player 2 then chooses without observing player 1's move, and the resulting branches are linked by a single information set spanning multiple post-move-1 nodes. The candidate “subgame” rooted at any single post-move-1 node is not a subgame in the formal sense, because it severs an information set. Backward induction has nowhere to start. SPE then either coincides with NE (when there is only one trivially-defined subgame, namely the whole game) or rules out very little.

The contrast between the two cases is best seen visually:



In the left tree, each player-2 node initiates its own proper subgame, and SPE pins down player 2's optimal action separately in each—backward induction works. In the right tree, the information set spanning both player-2 nodes makes the candidate “subgame” rooted at either one ill-defined: severing it would cut the information set in half. To evaluate optimality at the information set, player 2 needs to know how likely she is to be at the left versus the right node—a probability distribution over the nodes inside the information set.

SPE has no language for that; PBE does.

The remedy is to add *beliefs* to the equilibrium concept. At every information set the relevant player should hold a probability distribution over the nodes inside that information set, evaluate continuation play under that distribution, and respond optimally. The resulting solution concept is **perfect Bayesian equilibrium (PBE)**: an extension of SPE to extensive-form games *without* well-behaved subgame structure.

Remark (Beliefs as “Why Is This Player Doing This?”).

All of the equilibrium concepts up to this chapter—NE, SPE, even sequential rationality in finite games of perfect information—make no formal use of beliefs. They speak only of strategies and best replies. But in extensive-form games, asking *why* a player chooses a particular action is often more illuminating than verifying that the action is a best reply: “because she thinks the state of the world is X with probability p .” She still plays a best response—but a best response to her belief, not to the actual (unobserved) state. PBE makes that intuitive “thinks” a primitive of the equilibrium. Strategies tell you what each player does; beliefs tell you what each player conjectures about the unobserved nodes of the tree. The first novelty of PBE is to put both objects at the same level of formality.

Remark (Why Beliefs Are the Right Object to Add).

In an SPE, the absence of off-path play is what makes “backward induction at every subgame” coherent: the subgame partition cleanly separates “what happens here” from “what happens elsewhere.” When information sets cross those boundaries, evaluating “what happens here” from the inside requires knowing *which node inside the information set is in fact the current one*—and the only thing capable of summarizing that knowledge is a probability distribution over the nodes. PBE is what you get when you treat that distribution as a primitive of the equilibrium and impose two requirements: (i) actions are best replies given the beliefs, (ii) the beliefs are consistent with the strategies via Bayes’ rule wherever Bayes’ rule applies.

9.2 Definition: Strategies and Beliefs

A PBE is a pair (β, μ) in which $\beta = (\beta_i)_i$ is a profile of behavioral strategies and $\mu = (\mu_i)_i$ is a profile of belief systems, one per player.

Definition 9.1: Belief System

Fix a player i and an information set I_i of player i . A **belief** of player i at I_i is a probability distribution $\mu_i(I_i) \in \Delta(I_i)$ over the decision nodes in I_i . A **belief system** for player i is the collection $\mu_i = \{\mu_i(I_i)\}_{I_i \in \mathcal{I}_i}$ of such distributions, one for every information set of player i .

Definition 9.2: Perfect Bayesian Equilibrium

A pair (β, μ) is a **perfect Bayesian equilibrium** if

1. *Sequential rationality.* For every player i and every information set I_i , the action $\beta_i(I_i)$ that the strategy prescribes at I_i maximizes i 's conditional expected payoff at I_i , with the conditional expectation taken under the belief $\mu_i(I_i)$:

$$\beta_i(I_i) \in \arg \max_{a \in A_i(I_i)} \sum_{n \in I_i} \mu_i(I_i)(n) \cdot U_i(a, \beta_{-i}; n),$$

where $U_i(a, \beta_{-i}; n)$ denotes i 's continuation payoff from node n if i plays a at I_i and the others play β_{-i} thereafter.

2. *Bayesian consistency.* Given β , the belief $\mu_i(I_i)$ is derived from Bayes' rule whenever Bayes' rule applies—that is, whenever the strategy profile reaches I_i with positive probability.

Remark (Conditional, Not Unconditional).

The expectation in (1) is the *conditional* expected payoff at I_i , not the unconditional ex-ante one. The conditioning on “having reached I_i ” is encoded in the support of $\mu_i(I_i)$: by definition $\mu_i(I_i) \in \Delta(I_i)$ is a posterior *over* I_i , so summing $\mu_i(I_i)(n) \cdot U_i(a, \beta_{-i}; n)$ over $n \in I_i$ already restricts attention to the event “ I_i has been reached.” Two practical consequences. First, multiplying by the (positive) probability of reaching I_i would only rescale the objective by a constant and not change the argmax—so writing the conditioning explicitly outside, e.g. $\mathbb{E}[\cdot \mid I_i]$, is redundant if the belief is already a posterior. Second, this is exactly what makes sequential rationality *local*: each information set is evaluated on its own terms, with continuation payoffs $U_i(\cdot; n)$ that depend only on what comes after n and not on how the play got to n .

The first condition is sequential rationality at every information set; it is what replaces “optimality on every subgame” from the SPE definition. The second condition is the new piece: it prevents free choice of beliefs at on-path information sets but—deliberately—leaves beliefs at zero-probability information sets unrestricted by Bayes' rule alone. This loophole is what generates the multiplicity of PBE one observes in signaling games (Section 12.3) and motivates further refinements (intuitive criterion, divinity, sequential equilibrium).


Remark (“Whenever Possible”—the Source of Multiplicity).

The phrase *whenever possible* in (2) is doing real work. Bayes' rule pins down the posterior $P(\text{node} \mid I_i)$ only when the prior probability of reaching I_i under β is positive. At an information set that the strategy profile reaches with probability zero (a so-called *off-path* information set), Bayes' rule is silent—0/0—and the equilibrium concept allows *any* belief there. Equilibria are then often distinguished only by what beliefs are postulated off the equilibrium path, and these off-path beliefs are exactly what determines whether a candidate deviation is profitable. The discipline of refinements like the intuitive criterion

is essentially a proposal for how to tame this freedom.

A worked tree

Consider the extensive-form fragment below. Nature first selects one of two states with equal probability (the initial $[\frac{1}{2}]$ labels). Player 2's information set spans both post-Nature nodes, so 2 does not directly observe the state; the labels $[\frac{2}{3}]$ and $[\frac{1}{3}]$ are 2's belief at that information set. After 2 moves, on one branch a payoff is reached; on the other, player 3 moves at an information set spanned by two nodes with belief $[\frac{1}{4}], [\frac{3}{4}]$.



The numbers in brackets along edges are conditional probabilities induced by β ; the numbers in brackets next to the dashed information sets are the player's beliefs μ . Bayesian consistency forces an arithmetic relationship between the two. For instance, write N_1 and N_2 for the two nodes in player 2's information set (the left- and right-hand children of Nature, as marked on the diagram). Then 2's belief at that information set must satisfy

$$\mu_2(N_1) = \frac{\Pr(\text{Nature picks left}) \cdot 1}{\Pr(\text{Nature picks left}) \cdot 1 + \Pr(\text{Nature picks right}) \cdot 1} = \frac{1}{2}$$

under the prior $(\frac{1}{2}, \frac{1}{2})$. Any other belief there (such as $[\frac{2}{3}, \frac{1}{3}]$) would be inconsistent with the prior unless something earlier in β biased the path—and that something would itself need to be computed and reconciled with Bayes' rule. The board's specific numbers were illustrative of how to read the diagram, not a uniquely-solved equilibrium.

9.3 Spence Signaling

The leading workhorse of PBE is the **Spence (1973) signaling model** of education and labor markets. It crystallizes the idea that a costly action with no direct productive value can nevertheless be informative about hidden type, and—more cautiously—that the resulting equilibria come in continua, none of which is selected by PBE alone.

9.3.1 Setup

A worker has private **type** $\theta \in \{L, H\}$ with $0 < L < H$ (so θ doubles as the worker's productivity: a type- θ worker produces output worth exactly θ , and L, H are two numerical

productivity levels). The prior is

$$\Pr(\theta = H) = \pi, \quad \Pr(\theta = L) = 1 - \pi, \quad \pi \in (0, 1).$$

The worker first chooses an **education level** $e \in [0, \infty)$. Two competitive firms then observe e and simultaneously offer wages; Bertrand competition forces both wage offers up to the firm's expected productivity given e , so the realized wage is

$$w(e) = \mathbb{E}[\theta | e] = \mu(H | e) \cdot H + (1 - \mu(H | e)) \cdot L,$$

where $\mu(H | e)$ is the firms' posterior that the worker is high-type after observing the signal. The worker's payoff is

$$u(\theta, e, w) = w - \frac{e}{\theta}.$$

The cost of education is e/θ : high types find education cheaper. Critically, education *does not affect productivity*: θ enters firm payoffs directly, e only through firm beliefs. Education is a pure signal.

Remark (Why Not Just Say You're High-Type?).

The setup carefully forbids the worker from declaring her type. If declarations were allowed and free, every worker—low or high—would announce “I am high-type,” the announcement would carry no information, and firms would ignore it: this is the standard *cheap talk* non-result. The whole point of the model is that a verbal claim is unverifiable, so the worker must instead choose a costly action whose marginal cost depends on the unobserved type. The single-crossing structure e/θ guarantees that *some* levels of e are too painful for the low type to mimic but bearable for the high type, opening a channel through which type can be communicated.

Single-crossing. The cost function e/θ implements the **single-crossing property** (a.k.a. Spence-Mirrlees): for any two education levels $e' > e$ and any wage gap Δw that makes the low type just willing to switch from e to e' , the high type strictly prefers to switch. This monotone comparative statics in type is what makes signaling possible at all—without it, no signal could separate types in equilibrium.

9.3.2 An Order Constraint in Every Equilibrium

Before classifying equilibria, we record a useful observation that holds at every PBE.

Proposition 9.3: Monotone Education in Type

In any PBE of the Spence signaling game, $e^H \geq e^L$.

Proof for Proposition.

Let $w^\theta = \mathbb{E}[\theta | e^\theta]$ be the equilibrium wage offered to a worker who chooses e^θ . The two

incentive-compatibility constraints are

$$\begin{aligned} \text{(IC for } H) \quad w^H - \frac{e^H}{H} &\geq w^L - \frac{e^L}{H}, \\ \text{(IC for } L) \quad w^L - \frac{e^L}{L} &\geq w^H - \frac{e^H}{L}. \end{aligned}$$

Adding the two inequalities, the wage terms cancel and one obtains

$$\frac{e^L - e^H}{H} + \frac{e^H - e^L}{L} \geq 0 \iff (e^H - e^L)\left(\frac{1}{L} - \frac{1}{H}\right) \geq 0.$$

Since $\frac{1}{L} - \frac{1}{H} > 0$, we conclude $e^H \geq e^L$. ■

The proof is general: it nowhere uses the specific values of w^H, w^L beyond the IC inequalities. It illustrates a recurring trick in the analysis of separating-type games—add the two ICs and watch the wage terms drop out, leaving an inequality among the actions and the type-sensitivity of cost.

9.3.3 Separating Equilibria

A **separating equilibrium** is one in which $e^H \neq e^L$, so that θ is fully revealed by the signal. The proposition forces $e^H > e^L$ in such equilibria.

Two structural features. First, $e^L = 0$ in any separating equilibrium: if the low type is identified by his signal, his wage is L regardless of e , and any $e^L > 0$ is strictly dominated by $e = 0$. Second, the high type's signal $e^H = e^*$ must satisfy both ICs simultaneously; since separation forces $w^H = H, w^L = L$, the ICs become

$$\begin{aligned} \text{(IC for } H) \quad H - \frac{e^*}{H} &\geq L - 0 = L, & \Leftrightarrow \quad e^* &\leq H(H - L), \\ \text{(IC for } L) \quad L - 0 &\geq H - \frac{e^*}{L}, & \Leftrightarrow \quad e^* &\geq L(H - L). \end{aligned}$$

A note on why the low-type IC must be checked at all, given that we just argued $e^L = 0$. The argument $e^L = 0$ said the low type, *conditional on being identified as low* (so that wage is L regardless of e), optimally chooses no education. But the deviation IC for L rules out is a different one: the low type *pretending to be high* by paying e^* and pocketing wage H . Whether this masquerade is profitable depends on the cost-of-education-to-low-type, e^*/L , against the wage gain $H - L$. The lower bound $e^* \geq L(H - L)$ is precisely the condition that makes the masquerade not pay.

Combining, $e^* \in [L(H - L), H(H - L)]$. The lower bound says the signal must be costly enough that the low type would not pay it for the wage gain $H - L$; the upper bound says it must not be *so* costly that even the high type would prefer to mimic the low type.

Proposition 9.4: Separating PBE: A Continuum

For every $e^* \in [L(H - L), H(H - L)]$ there exists a PBE in which $e^H = e^*, e^L = 0$, for some choice of off-path beliefs (constructed in the proof below). Hence the Spence signaling game admits a *continuum* of separating equilibria.

Proof for Proposition.

The strategy profile is $e^L = 0$, $e^H = e^*$. The belief system specifies, on-path, $\mu(H | 0) = 0$ and $\mu(H | e^*) = 1$, generating wages L and H respectively. Off-path beliefs can be chosen to discourage any deviation; the simplest choice is the **pessimistic** specification

$$\mu(H | e) = \begin{cases} 1 & e = e^*, \\ 0 & e \neq e^*, \end{cases}$$

i.e., any out-of-equilibrium signal is read as “low type.”

- *Low type’s IC.* If L stays at 0, payoff is $L - 0 = L$. If L deviates to any $e' \neq 0$ with $e' \neq e^*$, payoff is $L - e'/L < L$, strictly worse. The only deviation worth examining is $e' = e^*$ (which would give wage H); this is unprofitable iff $L \geq H - e^*/L$, i.e. $e^* \geq L(H - L)$.
- *High type’s IC.* If H stays at e^* , payoff is $H - e^*/H$. The best deviation under the pessimistic belief gives wage L , so the most profitable deviation is to $e' = 0$ (cheapest), with payoff L . No-deviation requires $H - e^*/H \geq L$, i.e. $e^* \leq H(H - L)$.

Both ICs hold for $e^* \in [L(H - L), H(H - L)]$. Sequential rationality of the firms is automatic given Bertrand competition and the posited beliefs. ■

Remark (Why a Continuum?).

The interval $[L(H - L), H(H - L)]$ has positive length whenever $L < H$, so there is always more than one separating e^* . The reason is that the on-path beliefs $\mu(H | e^*) = 1, \mu(H | 0) = 0$ pin down only *two* of the firm’s posteriors; the remaining beliefs at every other education level are free, and the equilibrium concept lets us use that freedom to deter deviations. Each e^* in the interval is supported by its own pessimistic specification: the recipe is uniform (“any $e \notin \{0, e^*\}$ is read as low type”), but the location of the on-path spike at e^* shifts across equilibria. So all members of the family share a single rule for how to interpret deviations; they differ only in which e^* the rule targets as the high-type signal.

9.3.4 Pooling Equilibria

A **pooling equilibrium** is one in which both types choose the same signal: $e^L = e^H = \bar{e}$ for some common $\bar{e} \geq 0$. The firm cannot extract any information from \bar{e} alone, so on-path beliefs equal the prior:

$$\mu(H | \bar{e}) = \pi, \quad w(\bar{e}) = \pi H + (1 - \pi)L.$$

Off-path beliefs are not pinned down by Bayes’ rule. We adopt the standard **pessimistic specification**

$$\mu(H | e) = 0 \quad \text{for every } e \neq \bar{e},$$

under which any deviation is read as “low type” and rewarded with wage L . We show below that this is the harshest off-path belief, hence supports the largest set of pooling equilibria;

weaker off-path specifications collapse most of them.

Identifying the binding deviation. A worker contemplating a deviation from \bar{e} to some $e' \neq \bar{e}$ obtains wage L regardless of e' (under pessimistic beliefs), so her deviation payoff is $L - e'/\theta$. Since this is decreasing in e' , the most attractive deviation is the cheapest off-path signal: $e' = 0$. Both types' IC therefore reduces to a single inequality each, comparing the equilibrium payoff at \bar{e} to the deviation payoff at $e' = 0$.

Low-type IC. The low type must (weakly) prefer \bar{e} to deviating to $e' = 0$:

$$\underbrace{\pi H + (1 - \pi)L - \frac{\bar{e}}{L}}_{\text{equilibrium payoff } \bar{U}_L} \geq \underbrace{L - 0}_{\text{deviation payoff}} = L.$$

Rearranging,

$$\pi H + (1 - \pi)L - L \geq \frac{\bar{e}}{L} \iff \pi(H - L) \geq \frac{\bar{e}}{L} \iff \bar{e} \leq \pi L(H - L).$$

High-type IC. The high type must (weakly) prefer \bar{e} to deviating to $e' = 0$:

$$\pi H + (1 - \pi)L - \frac{\bar{e}}{H} \geq L.$$

Rearranging,

$$\bar{e} \leq \pi H(H - L).$$

This bound is strictly weaker than the low-type bound, since $\pi H(H - L) > \pi L(H - L)$. So the high-type IC is automatically satisfied whenever the low-type IC is. Geometrically: education is cheaper per unit for the high type ($1/H < 1/L$), so the same on-path payoff requires a smaller utility cost from \bar{e} for H , giving more slack.

Sequential rationality of the firm. Two competitive firms observe e and bid up to expected productivity. On path, expected productivity given the prior π is exactly $w(\bar{e}) = \pi H + (1 - \pi)L$. Off path under the pessimistic belief, expected productivity is L . Both wages are best responses, so the firm's strategy is sequentially rational.

Combining the three conditions:

Proposition 9.5: Pooling PBE: A Continuum

For every $\bar{e} \in [0, \pi L(H - L)]$ there exists a pure-strategy PBE in which both types choose $e = \bar{e}$, the on-path belief is $\mu(H | \bar{e}) = \pi$, the off-path belief is $\mu(H | e) = 0$ for $e \neq \bar{e}$, and firms offer the wage $\pi H + (1 - \pi)L$ at \bar{e} and the wage L at every other signal.

Proof for Proposition.

Sequential rationality of the firm follows from Bertrand competition together with the on-path posterior π and the off-path pessimistic posterior 0. Both types' IC against the cheapest deviation $e' = 0$ reduce to $\bar{e} \leq \pi L(H - L)$ (low type) and $\bar{e} \leq \pi H(H - L)$ (high type); the first implies the second. ■

The boundary case $\bar{e} = 0$ is the trivial pooling equilibrium in which neither type acquires education and the firm offers the prior wage. The upper boundary $\bar{e} = \pi L(H - L)$ is the most education-intensive pooling equilibrium consistent with the low type's IC; any larger \bar{e} would induce L to drop down to $e' = 0$, breaking the equilibrium.

Why pessimistic off-path beliefs? If we instead specified the off-path belief as $\mu(H | e) = 1$ for some $e \neq \bar{e}$ (say $e' < \bar{e}$), then a low type weighing \bar{e} vs. e' would compare $\pi H + (1 - \pi)L - \bar{e}/L$ to $H - e'/L$. Since $H > \pi H + (1 - \pi)L$ and $e' < \bar{e}$, deviation strictly dominates whenever the off-path optimistic wage is offered cheaply enough—so the pooling equilibrium unravels. The pessimistic specification is the unique off-path belief structure under which no deviation looks more attractive than the equilibrium signal. This is the same observation that motivates the intuitive-criterion analysis in §9.3.6 below: the IC objects precisely to the pessimistic belief whenever the deviation is unattractive for L but attractive for H .


Remark (The Two Roles of Incentive Compatibility).

PBE asks two things: that strategies be optimal given beliefs (sequential rationality), and that beliefs be Bayes-consistent given strategies—wherever Bayes applies. The qualifier is what gives PBE its slack: Bayes pins down on-path beliefs, but off-path beliefs are unconstrained by Bayes alone. As a result, IC ends up doing two distinct jobs.

- *On-path discipline.* IC rules out deviations from one on-path action to another on-path action. At on-path information sets, posteriors are determined by Bayes, and IC must hold against those posteriors. There is no freedom here—this is a necessary condition.
- *Off-path discipline as a design choice.* IC also rules out deviations from on-path actions to off-path actions. But here whether a deviation pays depends on what the firm *thinks* when it sees the off-path signal, and that thought is not pinned down by Bayes. The analyst must select off-path beliefs strategically: pick beliefs harsh enough that every candidate off-path deviation looks unprofitable. The pessimistic specification (“any unexpected signal is a low type”) is the harshest possible choice and therefore sustains the largest set of equilibria.

Under more permissive off-path beliefs (e.g., “any unexpected signal is a high type”), most of these equilibria collapse: a low type would mimic the high type's signal, the firm would interpret it as high, and the equilibrium would unravel. The off-path belief is exactly where the multiplicity lives—and where refinements like the intuitive criterion or divinity intervene to discipline it.

The trade-off underlying separating and pooling equilibria can be visualized in (e, w) -space using each type's indifference curves through the no-education / low-wage anchor $(0, L)$:



Reading the figure. Each type’s indifference curve from $(0, L)$ has slope $1/\theta$: the low type’s curve (slope $1/L$, blue) is steeper than the high type’s (slope $1/H$, red) because education is more costly per unit for the low type. A separating equilibrium with high-type signal e^* requires the firm-offered wage H at $e = e^*$ to lie *below* the low type’s indifference curve through $(0, L)$ (so the low type does not mimic) and *above* the high type’s indifference curve through $(0, L)$ (so the high type prefers signaling to staying at $e = 0$). The wage H touches the low-type curve at $e = L(H - L)$ and the high-type curve at $e = H(H - L)$, giving the interval $[L(H - L), H(H - L)]$ as the range of e^* that supports separation. A pooling equilibrium at $\bar{e} \in (0, \pi L(H - L)]$ analogously sits where the off-path wage L at $e = 0$ and the on-path wage $\pi H + (1 - \pi)L$ at \bar{e} leave the low type weakly content—geometrically the same kind of inequality, with π replacing the 1 that enters the separating bound.

9.3.5 Partial Pooling and the Multiplicity Problem

Beyond pure separation and pure pooling, the model also admits **semi-separating** (partial-pooling) equilibria, in which one type randomizes across two signals. For example, H randomizes between e' and e'' while L plays e'' for sure; firms then assign different posteriors to e' (pure H) and e'' (mixture). Each such construction generates an additional family of PBE indexed by the mixing probability and the off-path beliefs.

The cumulative picture, then, is uncomfortable: the Spence model has *tons* of perfect Bayesian equilibria. Separating and pooling alone yield two continua; partial pooling adds further families. PBE, the natural extension of SPE to incomplete-information games, is too permissive to deliver sharp predictions in this canonical example.

Remark (Why So Many Equilibria, and What to Do About It).

The proliferation comes from two sources. First, off-path beliefs are unrestricted by Bayes’ rule, so different choices generate different equilibria. Second, even on-path the IC constraints in the separating case define a closed interval, not a single point. The literature responds with *refinements*: the intuitive criterion (Cho-Kreps 1987), divinity (Banks-Sobel 1987), undefeated equilibrium (Mailath-Okuno-Fudenberg-Postlewaite 1993), and others. Each restricts off-path beliefs by an additional logical principle (“a deviation that only one type would ever consider should be attributed to that type”) and

selects subsets of the PBE set. In the next subsection we develop the intuitive criterion in full, showing that it selects the unique *Riley outcome*—the least-cost separating equilibrium, $e^L = 0$, $e^H = L(H - L)$. The lesson is general: PBE is the right concept for incomplete-information dynamics, but it usually needs a refinement to deliver predictive bite.

9.3.6 Refinement: The Intuitive Criterion

The intuitive criterion of Cho and Kreps (1987) is a forward-induction restriction on off-path beliefs. Its conclusion in the Spence model is sharp: among the continua of separating, pooling, and partial-pooling PBE, exactly one outcome survives—the least-cost separating equilibrium $(e^L, e^H) = (0, L(H - L))$. This subsection states the criterion and proves the selection result.

Motivation: a forward-induction speech. Fix any PBE and consider an off-path signal e' . Suppose a high-type worker contemplates choosing e' and accompanies the choice with the following speech to the firms:

“I am the high type. Here is why you should believe me. If you grant me the most charitable belief and pay me wage H at e' , my payoff $H - e'/H$ exceeds my equilibrium payoff. So I have a reason to send e' . The same most charitable belief, paying H at e' , gives a low type only $H - e'/L$, which is strictly below his equilibrium payoff. Even under the most generous interpretation, the low type would never want to send e' . So your belief at e' should put zero weight on L .”

The intuitive criterion is precisely this argument formalized. It rules out off-path beliefs that attribute a deviation to a type for whom the deviation is unprofitable under *every* possible belief, when some other type could plausibly benefit.

Definition 9.6: Equilibrium-Dominated Signal

Fix a PBE of the signaling game and let \bar{U}_θ denote type θ 's equilibrium payoff. A non-equilibrium signal e' is **equilibrium-dominated for type θ** if

$$\max_{w \in [L, H]} [w - e'/\theta] < \bar{U}_\theta,$$

i.e., even granting the most favorable wage $w = H$, type θ 's payoff at e' falls strictly short of her equilibrium payoff. Equivalently, $H - e'/\theta < \bar{U}_\theta$.

Intuitive Criterion (Cho-Kreps, 1987)

A PBE satisfies the **intuitive criterion** if, for every non-equilibrium signal e' and every pair of types θ, θ' ,

$$e' \text{ is equilibrium-dominated for } \theta \text{ but not for } \theta' \implies \mu(\theta | e') = 0.$$

When e' is dominated for all types or for none, the criterion places no restriction.

The criterion bites only when at least one type strictly cannot benefit from the deviation under any belief while at least one other type might. The off-path belief is then required to put zero weight on the unambiguous loser.

We now apply the criterion to the Spence model and obtain the uniqueness result through three claims.

Proposition 9.7: No Pooling Equilibrium Satisfies the Intuitive Criterion

Every pooling equilibrium of the Spence signaling game (i.e., $\bar{e} \in [0, \pi L(H - L)]$) violates the intuitive criterion.

Proof for Proposition.

In a pooling equilibrium at \bar{e} the wage on path is $w(\bar{e}) = \pi H + (1 - \pi)L$, and the equilibrium payoffs are

$$\bar{U}_L = \pi H + (1 - \pi)L - \bar{e}/L, \quad \bar{U}_H = \pi H + (1 - \pi)L - \bar{e}/H.$$

We show that there exists an off-path signal $e' > \bar{e}$ that is equilibrium-dominated for L but not for H .

Such an e' exists iff $H - e'/L < \bar{U}_L$ and $H - e'/H > \bar{U}_H$ both hold, i.e., iff

$$L(H - \bar{U}_L) < e' < H(H - \bar{U}_H).$$

Substituting the equilibrium payoffs and simplifying,

$$\begin{aligned} L(H - \bar{U}_L) &= L[(1 - \pi)(H - L)] + \bar{e}, \\ H(H - \bar{U}_H) &= H[(1 - \pi)(H - L)] + \bar{e}. \end{aligned}$$

The interval is therefore $(L(1 - \pi)(H - L) + \bar{e}, H(1 - \pi)(H - L) + \bar{e})$, which is non-empty since $L < H$. Pick any e' in this interval.

Then e' is equilibrium-dominated for L (the LHS inequality) but not for H (the RHS). The intuitive criterion demands $\mu(H | e') = 1$, which contradicts the off-path belief $\mu(H | e') = 0$ that supports the pooling equilibrium. The pooling equilibrium fails the criterion. ■

Proposition 9.8: No Separating Equilibrium with $e^H > L(H - L)$ Satisfies the Intuitive Criterion

Every separating equilibrium with $e^L = 0$ and $e^H \in (L(H - L), H(H - L)]$ violates the intuitive criterion.

Proof for Proposition.

In such an equilibrium the wage at e^H is H and the equilibrium payoffs are

$$\bar{U}_L = L, \quad \bar{U}_H = H - e^H/H.$$

Pick any $e' \in (L(H-L), e^H)$. The interval is non-empty by hypothesis ($e^H > L(H-L)$).

- For L : $H - e'/L < H - L(H-L)/L = L = \bar{U}_L$. So e' is equilibrium-dominated for L .
- For H : $H - e'/H > H - e^H/H = \bar{U}_H$ (since $e' < e^H$). So e' is *not* equilibrium-dominated for H .

The intuitive criterion demands $\mu(H | e') = 1$. But the equilibrium supporting e^H relied on $\mu(H | e') = 0$ for $e' \neq 0, e^H$ to deter the high type from undercutting her own signal. Contradiction. ■

Proposition 9.9: The Riley Separating Equilibrium Satisfies the Intuitive Criterion

The separating equilibrium with $(e^L, e^H) = (0, L(H-L))$ satisfies the intuitive criterion.

Proof for Proposition.

In the Riley equilibrium $\bar{U}_L = L$ and $\bar{U}_H = H - L(H-L)/H$. Take any non-equilibrium signal $e' > 0$ with $e' \neq L(H-L)$; we check both cases.

Case $e' < L(H-L)$. Then

$$H - e'/L > H - L(H-L)/L = L = \bar{U}_L,$$

so e' is not equilibrium-dominated for L . Similarly, $H - e'/H > H - L(H-L)/H = \bar{U}_H$, so e' is not equilibrium-dominated for H . The criterion places no restriction.

Case $e' > L(H-L)$. Then

$$H - e'/L < L = \bar{U}_L \quad \text{and} \quad H - e'/H < H - L(H-L)/H = \bar{U}_H,$$

so e' is equilibrium-dominated for both types. The criterion places no restriction.

In neither case does the criterion demand a non-trivial belief, so the Riley equilibrium satisfies it (with any off-path belief, including the standard pessimistic $\mu(H | e') = 0$). ■

Theorem 9.10: Uniqueness Under the Intuitive Criterion

In the Spence signaling game, exactly one PBE outcome survives the intuitive criterion: the least-cost separating equilibrium with $e^L = 0$ and $e^H = L(H-L)$.

Remark (Existence and Wider Applicability).

Two practical points worth flagging.

Existence is not automatic. For a general signaling game, the intuitive criterion can in principle eliminate *all* PBE, leaving an empty set of refined equilibria. This does not happen in the Spence model: the proof of the third claim above is itself an existence proof—the Riley equilibrium is constructed and shown to satisfy the criterion. But in

some games one needs to check existence separately, sometimes by passing to weaker refinements (divinity, $D1$, $D2$, ...).

Compatibility with sequential equilibrium. Cho and Kreps designed the intuitive criterion as an economic re-statement of more abstract stability concepts in the Kohlberg-Mertens (1986) “stable equilibrium” tradition. In generic signaling games, the intuitive criterion is consistent with sequential equilibrium: every IC-surviving PBE is a sequential equilibrium. The Mas-Colell-Whinston-Green textbook (Chapter 13) develops these ideas in greater detail.

What about partial pooling? The same argument that knocks out pooling also knocks out partial-pooling (semi-separating) equilibria. If the high type randomizes between e' and e'' while the low type concentrates on e'' , then any signal slightly above the low-type’s equilibrium can be shown to be equilibrium-dominated for L but not for H , producing the same contradiction. The intuitive criterion thus selects a single outcome out of a three-dimensional family of PBE.

Remark (Costly Signals Beyond Education: The Peacock’s Tail).

An evolutionary biology analogue of the Spence model arose around the same time. A peacock’s elaborate tail is metabolically expensive to grow and makes the bird more conspicuous to predators. It does not improve foraging or fighting ability. Why grow it? The Zahavi (1975) “handicap principle” answers: the tail is a costly, unfakeable signal of underlying genetic fitness. Only a peacock with a strong constitution can afford the energy expenditure and the predation risk; weak peacocks cannot mimic the display, so peahens that prefer big tails systematically end up with high-fitness mates. The structural parallel to Spence is exact: a productive but unobservable trait, a costly action whose marginal cost is decreasing in the trait, and a separating equilibrium in which the high type incurs the cost to communicate her type. Costly signaling is one of those rare unifying ideas that crossed independently into economics and biology in the 1970s and now turns up wherever an agent has private information she would like to credibly convey.

9.4 Reputation in the Finitely Repeated PD

The chain store paradox of Section 11.4 showed that a vanishingly small probability of a “crazy” commitment type can sustain reputational behavior in a sequential game between a long-lived player and short-lived opponents. A natural follow-up question, addressed by Kreps, Milgrom, Roberts, and Wilson (1982) in a companion paper, is whether the same resolution rescues cooperation in the **finitely repeated Prisoner’s Dilemma** between two long-lived players—another canonical setting where backward induction predicts unraveling.

9.4.1 Setup

Two players play the Prisoner’s Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	2, 2	-1, 3
<i>D</i>	3, -1	0, 0

for T periods, with periods numbered *backwards*: $t = 1$ is the last period, $t = T$ is the first. Each player is independently *rational* with probability $1 - \varepsilon$ and a *commitment type* (always plays grim trigger: cooperate while the opponent cooperates; switch to permanent defection after any defection) with probability ε . Players observe past actions but not the opponent's type. Let μ_t denote the probability a player ascribes to the *opponent* being a commitment type at the start of period t , conditional on the equilibrium history.

Remark (Why Backward Induction Fails Here).

Recall from Section 11.3 that under common knowledge of payoffs the only SPE of the finitely repeated PD is (D, D) in every period. The reason is structural: the unique stage NE in the last period pins down (D, D) at $t = 1$, which makes the period-2 subgame strategically equivalent to the stage game, and so on. The reputation model breaks this by giving the rational player something at stake in maintaining ambiguity about her type: defecting reveals rationality (a commitment type would never defect against a cooperator), and revelation collapses the subsequent subgame to (D, D) forever. That continuation loss can outweigh the one-period gain from defecting, sustaining cooperation off the unraveling path.

9.4.2 Constructing the Equilibrium

We solve for a symmetric PBE in which both rational players cooperate with positive probability while both retain reputation, and randomize as the endgame approaches. Let

$$q_t = \Pr(\text{player cooperates in period } t \mid \text{both cooperated through } t + 1),$$

let μ_t be the belief about the opponent being commitment, and let p_t be the probability that a *rational* player cooperates in period t . Conditional on histories in which both players have cooperated through period $t + 1$, Bayesian aggregation gives

$$q_t = (1 - \mu_t)p_t + \mu_t,$$

since commitment types always cooperate so far. Inverting,

$$p_t = \frac{q_t - \mu_t}{1 - \mu_t}.$$

After any defection, the commitment type switches to permanent D , so observing D from a player who was on the cooperation path reveals rationality with certainty—beliefs collapse to $\mu = 0$ and the continuation reverts to mutual defection.

The construction proceeds in three phases:

- **Phase I** ($t > K + 1$): full cooperation, $q_t = 1$, $p_t = 1$.
- **Phase II** ($1 < t \leq K + 1$): rational players randomize.

- **Phase III** ($t = 1$): defect outright.

The cutoff K is chosen as a function of ε . Specifically, fix the integer k satisfying

$$\left[\frac{1}{3}\right]^{k+1} \leq \varepsilon < \left[\frac{1}{3}\right]^k, \quad \Leftrightarrow \quad \frac{1}{3} \leq 3^k \varepsilon < 1,$$

and set $K = 2k$, so that $\frac{1}{3} \leq 3^{K/2} \varepsilon < 1$.

9.4.3 Equilibrium Mixing Probabilities and Beliefs

The equilibrium specifies the on-path mixing weights

$$q_1 = q_3 = \dots = q_{K+1} = 3^{K/2} \varepsilon, \quad q_2 = q_4 = \dots = q_K = \left[\frac{1}{3}\right]^{K/2+1} \frac{1}{\varepsilon},$$

and $q_t = 1$ for all $t > K + 1$. Notice that for $t \leq K + 1$,

$$q_t q_{t+1} = \frac{1}{3},$$

which is the key recursion that makes randomization an equilibrium (derived below). On-path beliefs follow the Bayesian update

$$\mu_t = \begin{cases} \varepsilon & \text{if } t > K, \\ \frac{\varepsilon}{q_{K+1} q_K \dots q_{t+1}} & \text{if } t \leq K, \end{cases}$$

which simplifies to $\mu_t = \mu_{t+1}/q_{t+1}$ for any history of mutual cooperation.

9.4.4 Verification at the Last Period

Plug in $t = 1$:

$$\mu_1 = \frac{\varepsilon}{q_{K+1} q_K \dots q_2} = 3^{K/2} \varepsilon = q_1.$$

At $t = 1$ the probability of cooperation by the opponent equals exactly the probability that the opponent is a commitment type. This means rational players defect with probability 1 in the last period (only commitment types cooperate), as one would expect from the unraveling logic.

9.4.5 Indifference in the Mixing Phase

Consider any period $1 < t \leq K + 1$ on the cooperation path. We verify that a rational player is indifferent between cooperating and defecting.

Defecting today. Playing D when the opponent cooperates with probability q_t yields

$$V_t(D) = 3q_t + 0 \cdot (1 - q_t) = 3q_t,$$

since the opponent's commitment type, after observing D , switches to permanent D , so all continuation payoffs are zero.

Cooperating today. Playing C when the opponent cooperates with probability q_t yields a complicated continuation, but the structure of the equilibrium ensures the simple

expression

$$V_t(C) = (2 + 3q_{t-1})q_t - 1 \cdot (1 - q_t).$$

The decomposition: with probability q_t the opponent cooperates today (instantaneous payoff 2), and tomorrow the rational player will defect against a still-cooperating opponent (yielding $3q_{t-1}$ in $t - 1$); with probability $1 - q_t$ the opponent defects today (instantaneous loss -1), revealing rationality, and all subsequent payoffs are zero.

Setting $V_t(C) = V_t(D)$:

$$3q_t = (2 + 3q_{t-1})q_t - (1 - q_t) \iff q_{t-1}q_t = \frac{1}{3},$$

which is exactly the recursion built into the proposed q_t schedule.

Remark (Inductive Verification).

The indifference equation $q_{t-1}q_t = 1/3$ can be verified inductively. *Base case* ($t = 2$): defecting yields $3q_2$; cooperating yields $1 + 3q_1$ if the opponent cooperates (one will defect tomorrow, gaining $3q_1$ when the opponent's commitment type still cooperates) and -1 if the opponent defects. Setting these equal: $3q_2 = q_2(1 + 3q_1) - (1 - q_2)$, which rearranges to $q_1q_2 = 1/3$. *Inductive step*: assume the relation holds for $t - 1$, so $V_{t-1}(C) = V_{t-1}(D) = 3q_{t-1}$. Then in period t , defecting earns $3q_t$ today and 0 thereafter (the opponent's commitment type retaliates); cooperating earns 2 today and continuation $3q_{t-1}$ if the opponent cooperated, or -1 today and 0 thereafter if the opponent defected. The same algebra delivers $q_{t-1}q_t = 1/3$.

9.4.6 What the Equilibrium Achieves

The construction yields a striking result: with arbitrarily small $\varepsilon > 0$, sufficiently long horizons T admit a PBE in which rational players cooperate for almost the entire game, randomize during a final endgame phase of length $\Theta(\log(1/\varepsilon))$, and defect outright in the very last period. The fraction of periods featuring full cooperation tends to 1 as $T \rightarrow \infty$ with ε fixed.

The mechanism is structurally identical to the chain store paradox: each rational player is willing to cooperate today because doing so preserves the opponent's uncertainty about her type (since a defection would reveal rationality and trigger permanent retaliation). The opponent, in turn, willingly cooperates because the probability that the player is a commitment type is high enough to make cooperation marginally profitable. The reputational equilibrium is held together by the small ε , which provides just enough "cover" for rational players to mimic cooperation without immediate detection.

Remark (Two-Sided vs. One-Sided Reputation).

The PD reputation model differs from the chain store in one important respect: *both* players are long-lived and *both* have private types. In the chain store, only the store has a type and the entrants are short-lived; reputation flows in one direction. Here, each rational player both maintains and reads reputation simultaneously, and the equilibrium balances the two sides via the $q_{t-1}q_t = 1/3$ recursion. A consequence: even if only *one* side has commitment types (asymmetric uncertainty), the result still goes through; what

matters is that *some* doubt about rationality exists somewhere in the game. The lesson generalizes well beyond PD: cooperative behavior in any finitely repeated game with a unique stage NE can be rescued by a tiny mass of behavioral types, regardless of whether the perturbation is one-sided or two-sided.

Remark (Why a Vanishingly Small ε Suffices).

A natural worry: shouldn't a tiny ε produce only a tiny effect? The answer is no, and the reason is the same logarithmic blow-up that drove the chain store: the length of the cooperative phase scales as $\log(1/\varepsilon)/\log 3$. Halving ε adds roughly $\log 2/\log 3 \approx 0.63$ extra periods of cooperation. So even at $\varepsilon = 10^{-6}$, the cooperative phase lasts about 13 periods—enough that in any moderately long horizon, almost all play looks cooperative. This logarithmic sensitivity is what makes “a small grain of doubt” a structural feature of repeated games rather than a minor perturbation.

Remark (Chapter Summary).

Perfect Bayesian equilibrium (PBE, Definition 9.2) extends subgame perfection to games of incomplete information. The concept has two ingredients: a strategy profile that is sequentially rational at every information set, and a belief system that is consistent with Bayes' rule wherever Bayes' rule applies. PBE has bite in two canonical settings. *Spence's job-market signaling model*: the rich type set produces a continuum of separating PBE (each indexed by an off-path belief restriction), a continuum of pooling PBE, and various hybrids; the *intuitive criterion* of Cho-Kreps selects the least-cost separating equilibrium, which has clean welfare implications and matches the empirical pattern of overinvestment in education. *Reputation in finitely repeated PD*: a tiny mass of “tit-for-tat” commitment types can sustain near-fully-cooperative play in long but finite horizons, with the cooperative phase scaling as $\log(1/\varepsilon)$. PBE is the right concept for incomplete-information dynamics, but its multiplicity (especially of pooling and partial-pooling equilibria) means it usually requires a refinement to deliver predictive bite.

Part VI

Problem Sets and Solutions

Problem Set 1

Problems

Problem 1.1

Every morning 1 unit of cars (say, one thousand) travel from point A to point B . Each car can take one of two routes. If x cars take the “upper” route, the time taken by each car is $a + bx$ hours, where a and b are nonnegative numbers. Similarly, if y cars take the “lower” route, the time taken by each car is $c + dy$ hours, where again c and d are nonnegative numbers. Further, suppose that $a > c \geq 0$ and $d > b \geq 0$.

- (a) Under the assumption that each driver is aware of the situation described above, fully describe an equilibrium allocation—that is, how many cars travel on each of the two routes. (An equilibrium configuration is such that given the routes followed by all other cars, no driver can decrease the time taken for him to travel from A to B .) Find an expression for the average amount of time taken by the cars in equilibrium, say t^* .
- (b) Show by example (say, by choosing specific values for the parameters a, b, c and d) that the equilibrium may be inefficient, that is, there is another allocation such that the average amount of time taken by the cars is smaller. Find an expression for the socially efficient average time—one that minimizes the average time taken, say \bar{t} .
- (c) Compare the equilibrium average time t^* to the efficient average time \bar{t} when $a = 1$, $b = 0$, $c = 0$ and $d = 1$.
- (d) (Harder) Show that for all values of the parameters a, b, c and d , the ratio of the efficient average time to the equilibrium average time satisfies $\bar{t}/t^* \geq \frac{3}{4}$. (Hint: show that the situation in part (c) yields the lowest ratio of efficient to equilibrium times.)

Problem 1.2

There are two restaurants, named A and B , in a particular city. There is a prior probability of $\rho > \frac{1}{2}$ that restaurant A is better than restaurant B and, of course, a prior probability of $1 - \rho$ that restaurant B is better. There are 5 potential customers and each one receives a privately observed signal, α or β , indicating which of the restaurants is better. The probability of getting signal α when the true state is A (restaurant A is actually better) is $\Pr(\alpha | A) = q > \rho$. The probability of getting signal β when the true state is B is the same, that is, $\Pr(\beta | B) = q > \rho$ also. Furthermore, conditional on the state, the signals of different customers are independent.

- (a) Show that if a customer gets the signal α (and no other information), the posterior probability that A is the better restaurant is greater than $\frac{1}{2}$. (Use Bayes' rule to calculate the posterior probability.)
- (b) What is a customer's posterior assessment that A is the better restaurant if he gets the signal β and learns that one other customer got the signal α ?
- (c) Suppose that the first customer gets the signal α and the other four get the signal β . What is the posterior probability that A is the better restaurant conditional on all five signals?
- (d) Customers always go to the restaurant for which the posterior probability of being better is higher. Suppose, as above, that the first customer gets the signal α and the other four get the signal β . If customers make their choices sequentially and are observed by all others, what will be the outcome?
- (e) (Harder) Find the probability that all customers go to the "wrong" restaurant. For instance, based on all the signals, all customers go to B but it would be socially optimal for all customers to go to restaurant A .

Solutions

Problem 1.1

(a) **Equilibrium allocation and t^* .** In a Wardrop equilibrium with both routes used, drivers cannot reduce their travel time by switching, so the two routes deliver the *same* time:

$$a + bx = c + d(1 - x) \iff x^* = \frac{c + d - a}{b + d}, \quad y^* = \frac{a + b - c}{b + d}.$$

Because $a > c$ and $b \geq 0$, we always have $y^* > 0$. The interior condition $x^* \geq 0$ requires $d \geq a - c$.

- **Interior case** ($d \geq a - c$). Both routes used; substituting x^* back,

$$t^* = a + bx^* = \frac{ad + bc + bd}{b + d}.$$

- **Corner case** ($d < a - c$). Even with everybody on the lower route ($y = 1$), the lower route takes only $c + d < a$, so no one wants to switch up. Then $x^* = 0$, $y^* = 1$, and $t^* = c + d$.

(b) **Inefficiency and \bar{t} .** Total travel time is

$$T(x) = x(a + bx) + (1 - x)(c + d(1 - x)) = (b + d)x^2 + (a - c - 2d)x + (c + d).$$

Setting $T'(x) = 0$,

$$\bar{x} = \frac{c + 2d - a}{2(b + d)}, \quad \bar{y} = \frac{a + 2b - c}{2(b + d)}.$$

The interior case ($\bar{x} \geq 0$) requires $2d \geq a - c$. Plugging back,

$$\bar{t} = T(\bar{x}) = (c + d) - \frac{(a - c - 2d)^2}{4(b + d)}.$$

Concrete example of inefficiency. Take $(a, b, c, d) = (1, 0, 0, 1)$, treated in part (c) below: equilibrium delivers $t^* = 1$ but the social optimum delivers $\bar{t} = \frac{3}{4}$, so equilibrium is 33% slower than optimum.

A clean identity. A direct expansion shows that whenever the equilibrium is interior,

$$\bar{t} = t^* - \frac{(a - c)^2}{4(b + d)}.$$

The wedge $(a - c)^2/[4(b + d)]$ measures the externality each driver imposes on the others by ignoring how their presence increases congestion.

(c) **Comparison at $(a, b, c, d) = (1, 0, 0, 1)$.** Upper route takes time 1 no matter how many cars; lower takes time y .

- Equilibrium: $x^* = 0$, $y^* = 1$, every car on lower, $t^* = 1$.
- Efficient: $\bar{x} = \frac{1}{2}$, $\bar{y} = \frac{1}{2}$, $\bar{t} = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$.

Hence $\bar{t}/t^* = 3/4$. This is the canonical **Pigou example**: the upper road has *constant* time, so its users impose no externality, while every additional user of the lower road slows everyone

else down. Selfish routing pushes every car to the lower road, exactly the externality-maximizing assignment.

(d) The 3/4 bound on \bar{t}/t^* .

Interior equilibrium ($d \geq a - c$). Using the identity from (b),

$$\frac{\bar{t}}{t^*} = 1 - \frac{(a-c)^2}{4(b+d)t^*} = 1 - \frac{(a-c)^2}{4(ad+bc+bd)}.$$

We need $(a-c)^2 \leq ad+bc+bd$. Set $u = a-c > 0$; then $a = c+u$ and the inequality becomes

$$u^2 \leq (c+u)d+bc+bd = cd+ud+bc+bd \iff u(u-d) \leq cd+bc+bd.$$

The interior assumption gives $u \leq d$, so the left side is ≤ 0 while the right side is ≥ 0 . Equality requires $u = d$ and $cd = bc = bd = 0$; combined with $u, d > 0$, the only possibility is $b = c = 0, a = d$. Up to scaling, this is exactly part (c).

Corner equilibrium ($d < a - c$). Two sub-cases.

- If $2d < a - c$, the social optimum is also corner ($\bar{x} = 0$), so $\bar{t} = c + d = t^*$ and the ratio is 1.
- If $(a - c)/2 \leq d < a - c$, the social optimum is interior with

$$\bar{t} = (c+d) - \frac{(a-c-2d)^2}{4(b+d)} = t^* - \frac{(a-c-2d)^2}{4(b+d)}.$$

In this regime $|a-c-2d| = 2d-(a-c) \leq d$, so $(a-c-2d)^2 \leq d^2 \leq (b+d)(c+d) = (b+d)t^*$. Hence $\bar{t}/t^* \geq 1 - 1/4 = 3/4$, with equality demanding $b = c = 0$ and $a - c - 2d = \pm d$, i.e., the same Pigou point.

In every case, $\bar{t}/t^* \geq 3/4$, achieved uniquely (up to scaling) at $b = c = 0, a = d$.

Remark (Why exactly 3/4, not 1/2 or some other constant?).

The factor 3/4 is the celebrated **Roughgarden-Tardos price of anarchy** for affine cost functions. The proof above already exhibits the worst case: a constant-time route paired with a fully congested one. With more general cost functions, the worst-case ratio can be far worse—for $\text{cost}(x) = x^p$, the price of anarchy grows like $p/(\log p)$.

Problem 1.2

(a) Posterior given a single α . Bayes:

$$\Pr(A | \alpha) = \frac{\rho q}{\rho q + (1-\rho)(1-q)}.$$

This exceeds $\frac{1}{2}$ iff $\rho q > (1-\rho)(1-q)$, equivalently $q + \rho > 1$, which holds since $q, \rho > \frac{1}{2}$.

(b) Posterior given own β and observed other- α . Conditional on the state, signals

are independent, so

$$\Pr(A \mid \alpha, \beta) = \frac{\rho \cdot q(1-q)}{\rho \cdot q(1-q) + (1-\rho) \cdot (1-q)q} = \rho.$$

One α and one β exactly cancel: the posterior reverts to the prior.

(c) Posterior given 1 α and 4 β 's. Working with log-odds is cleanest. Let $\ell_t = \log \Pr(A \mid \cdot) / \Pr(B \mid \cdot)$ after t signals; $\ell_0 = \log[\rho/(1-\rho)]$. Each α adds $\log[q/(1-q)] > 0$, each β subtracts the same amount. With 1 α and 4 β ,

$$\ell_5 = \log \frac{\rho}{1-\rho} - 3 \log \frac{q}{1-q}.$$

Re-exponentiating,

$$\Pr(A \mid 1\alpha, 4\beta) = \frac{\rho(1-q)^3}{\rho(1-q)^3 + (1-\rho)q^3}.$$

Since $q > \rho > \frac{1}{2}$ implies $q/(1-q) > \rho/(1-\rho)$, the cube wipes out the prior tilt: $\Pr(A \mid 1\alpha, 4\beta) < \frac{1}{2}$. The aggregate evidence favors B .

(d) Sequential public choice. Customers act in sequence and each predecessor's action is observed. We trace the inferences.

Customer 1 (α , observes nothing). By (a), posterior $> \frac{1}{2}$; goes to A .

Customer 2 (β , observes $1 \rightarrow A$). Customer 1's choice perfectly reveals their signal: a β would have given posterior $< \frac{1}{2}$ (since $q > \rho$), so customer 1 must have observed α . Customer 2's effective evidence is $1\alpha + 1\beta$; by (b), posterior reverts to $\rho > \frac{1}{2}$. Customer 2 goes to A .

Customer 3 (β , observes $1 \rightarrow A, 2 \rightarrow A$). Now customer 2's action is *uninformative*: regardless of their signal, the inferred prior-plus- α from customer 1 outweighed any private β . Customer 3 is in the same evidential position as customer 2: prior + observed α from customer 1, plus own signal. With own β , posterior is $\rho > \frac{1}{2}$. Goes to A .

Customers 4, 5 (β , observe everyone going to A). Same logic: customers 2-3-... reveal nothing, and the only public information is customer 1's α . With own β , posterior = $\rho > \frac{1}{2}$. Both go to A .

Outcome. Every customer goes to restaurant A , even though the aggregate evidence ($1\alpha, 4\beta$) actually favors B (by (c)). This is an **information cascade**: customers 2-5 ignore their private signals because the public evidence dominates, and as a result the public information stops accumulating.

(e) Probability of a wrong cascade. Let X_t be the cumulative net signal balance ($\#\alpha - \#\beta$) inferable from informative actions through customer t , with $X_0 = 0$. Working through the action thresholds (a customer with own α goes to A iff $X \geq -1$; with own β goes to A iff $X \geq 1$) shows:

- In the inconclusive range $X_t \in \{-1, 0\}$, every customer's action reveals their signal, so $X_{t+1} = X_t \pm 1$.
- As soon as $X_t = 1$, every subsequent customer goes to A regardless of their signal: **A-cascade**.
- As soon as $X_t = -2$, every subsequent customer goes to B : **B-cascade**.

The walk on $\{-2, -1, 0, 1\}$ has absorbing barriers at -2 and $+1$. With at most 5 informative steps, by parity A -cascades can start only at odd steps (1, 3, 5) and B -cascades only at even steps (2, 4). Enumerating the possible prefixes:

Cascade direction	Prefix	Probability under state A ($\Pr(\alpha) = q$)
A	α	q
A	$\beta\alpha\alpha$	$(1 - q)q^2$
A	$\beta\alpha\beta\alpha\alpha$	$(1 - q)^2q^3$
B	$\beta\beta$	$(1 - q)^2$
B	$\beta\alpha\beta\beta$	$q(1 - q)^3$
none	$\beta\alpha\beta\alpha\beta$	$q^2(1 - q)^3$

Writing $\omega = q(1 - q)$, totals under state A are

$$\Pr(A\text{-cascade} \mid A) = q(1 + \omega + \omega^2), \quad \Pr(B\text{-cascade} \mid A) = (1 - q)^2(1 + \omega).$$

By the symmetry $q \leftrightarrow 1 - q$, totals under state B are

$$\Pr(B\text{-cascade} \mid B) = q^2(1 + \omega), \quad \Pr(A\text{-cascade} \mid B) = (1 - q)(1 + \omega + \omega^2).$$

Combining via the prior, the probability that the cascade direction is wrong (everyone ends up at the worse restaurant relative to the true state) is

$$\Pr(\text{wrong cascade}) = \rho \cdot (1 - q)^2(1 + \omega) + (1 - \rho) \cdot (1 - q)(1 + \omega + \omega^2).$$

A numerical sanity check. With $\rho = 0.6$ and $q = 0.7$, $\omega = 0.21$, this gives $0.6 \cdot 0.09 \cdot 1.21 + 0.4 \cdot 0.3 \cdot 1.2541 \approx 0.215$, i.e. a $\sim 22\%$ chance the cascade is wrong despite each individual signal being 70% accurate. As the number of customers grows, the probability of a wrong cascade does *not* vanish—the herding model’s central message.

Remark (Why cascades are robust).


Two features make a cascade self-sustaining: (i) once $|X|$ reaches a threshold, an additional signal of either sign is too weak to flip the decision, so individuals optimally ignore their private information; (ii) as a result, X stops moving, and no later observation can ever break the cascade. Information aggregation fails not because individuals are irrational—they are perfectly Bayesian—but because rational behavior *interacts* in a way that throws away private information.

Problem Set 2


Problems

Problem 2.1

- (a) Consider the following extensive-form game. Carefully write down the corresponding strategic (normal) form of this game.



- (b) Now consider the following extensive-form game. Again, carefully write down the corresponding strategic form of this game.




- (c) How do the two strategic forms differ?

Problem 2.2

Two players jointly own a firm which is initially worth \$2. The players move alternately—first player 1 moves, then player 2 moves, then player 1 moves again and this pattern repeats

until time T (an even number). Whenever it is player 1's turn to move, she can either invest effort (I) in the firm or unilaterally liquidate (L) the firm. If she invests, this adds \$2 to the value of the firm. If player 1 liquidates the firm at some time, the current value of the firm is split so that player 1 gets \$2 more than player 2 does. The same is true whenever it is player 2's turn to move with the roles reversed. The extensive form below depicts the situation for $T = 6$:



- (a) For the extensive form above, what is the backward induction solution? Does your answer depend on the value of T ?
- (b) For $T = 6$, write the corresponding strategic (or normal) form and show that iterated elimination of weakly dominated strategies results in a unique outcome. How does this compare with the outcome in part (a)?

Problem 2.3

There are two firms operating in a market with (inverse) demand function $p = a - q$, where p is the price, q is the total quantity sold by the two firms and $a > 0$ is a constant. Each firm can produce at a constant unit cost of c where $0 < c < a$. The firms choose sequentially: first firm 1 chooses a quantity q_1 to produce, and then firm 2, having seen q_1 , chooses a quantity q_2 to produce. Find the backward induction solution to this game.

Solutions

Problem 2.1

(a) **Strategic form of the first game.** Player 1 has the single decision node at the root with three actions $\{l, m, r\}$. Player 2 has a single information set spanning the two nodes reached by l and m , with actions $\{a, b\}$. After r , the game ends at $(2, 4)$ without Player 2 moving.

Player 1's strategies: $\{l, m, r\}$. Player 2's strategies: $\{a, b\}$.

	a	b
l	4, 2	1, 1
m	3, 2	0, 3
r	2, 4	2, 4

(b) **Strategic form of the second game.** Player 1 now has *two* decision nodes: the root, where she chooses between k (continue) and r (terminate at $(2, 4)$); and a second node reached after k , where she chooses between l and m . Player 2's information set still spans the two post- $\{l, m\}$ nodes. A strategy for Player 1 must specify an action at *every* information set, even those that may be unreachable on the equilibrium path. Hence Player 1 has four strategies: $\{kl, km, rl, rm\}$. Player 2 still has $\{a, b\}$.

	a	b
kl	4, 2	1, 1
km	3, 2	0, 3
rl	2, 4	2, 4
rm	2, 4	2, 4

(c) **Comparison.** The two strategic forms describe the same physical outcomes but differ in the size of Player 1's strategy set: 3 in (a) versus 4 in (b). The redundancy arises because under (b), Player 1 must commit to a hypothetical action at the second decision node even when the chosen root action (r) makes that node unreachable, so rl and rm become payoff-equivalent. The two extensive forms are sometimes called **realization-equivalent**: they generate the same set of probability distributions over outcomes for any behavioral strategy profile, but their normal forms are formally different.

Remark (Why this matters).

Realization-equivalent extensive forms can have *different* sets of mixed-strategy Nash equilibria when one allows correlations across information sets, because doubling Player 1's strategy count expands the simplex over which mixing is defined. For most equilibrium analysis (pure NE, behavioral mixed NE), this distinction is purely cosmetic. But it matters for *trembling-hand perfect* equilibrium, where the choice of representation affects which strategy profiles survive a given ε -trembling.

Problem 2.2

(a) **Backward induction.** Reading the game tree (for $T = 6$), the liquidation payoffs along the bottom are

$$(2, 0), (1, 3), (4, 2), (3, 5), (6, 4), (5, 7),$$

and if everyone invests through period 6, the terminal payoff is (8, 6).

- $t = 6$ (Player 2): L gives Player 2 the payoff 7, I gives 6. Choose L , outcome (5, 7).
- $t = 5$ (Player 1): comparing L (6, 4) to I leading to (5, 7), Player 1 prefers L . Outcome (6, 4).
- $t = 4$ (Player 2): L (3, 5) vs. $I \rightarrow (6, 4)$. Prefer L . Outcome (3, 5).
- $t = 3$ (Player 1): L (4, 2) vs. $I \rightarrow (3, 5)$. Prefer L . Outcome (4, 2).
- $t = 2$ (Player 2): L (1, 3) vs. $I \rightarrow (4, 2)$. Prefer L . Outcome (1, 3).
- $t = 1$ (Player 1): L (2, 0) vs. $I \rightarrow (1, 3)$. Prefer L . Outcome $\boxed{(2, 0)}$.

The unraveling is robust to T . At any period t , let $V_t = 2t$ be the firm’s value just before the move. Liquidating gives the mover $V_t/2 + 1 = t + 1$ and the opponent $V_t/2 - 1 = t - 1$. Investing then having the opponent liquidate next period gives the mover $V_{t+1}/2 - 1 = t$. So the mover always prefers liquidating now ($t+1$) to passing and being squeezed (t). Backward induction collapses to immediate liquidation at $t = 1$ with payoff (2, 0) regardless of T .

(b) **Iterated elimination of weakly dominated strategies ($T = 6$).** Each player has $2^3 = 8$ pure strategies (an action at each of their three decision nodes). Writing out the 8×8 matrix is unwieldy, but IEWDS unfolds period by period and reproduces the backward-induction outcome.

- *Round 1.* At Player 2’s last node ($t = 6$), L delivers 7 and I delivers 6 regardless of any earlier play. Hence *any* Player 2 strategy specifying I at $t = 6$ is weakly dominated by the same strategy with L at $t = 6$. After elimination, Player 2 strategies all have $s_6 = L$.
- *Round 2.* Given $s_6 = L$, at $t = 5$ the comparison for Player 1 is L (6, 4) vs. $I \rightarrow (5, 7)$, so $s_5 = I$ is weakly dominated.
- *Round 3.* Given $s_5 = L$, at $t = 4$ Player 2 compares L (3, 5) vs. $I \rightarrow (6, 4)$; eliminate $s_4 = I$.
- *Round 4.* At $t = 3$: L (4, 2) vs. $I \rightarrow (3, 5)$; eliminate $s_3 = I$.
- *Round 5.* At $t = 2$: L (1, 3) vs. $I \rightarrow (4, 2)$; eliminate $s_2 = I$.
- *Round 6.* At $t = 1$: L (2, 0) vs. $I \rightarrow (1, 3)$; eliminate $s_1 = I$.

After six rounds, each player retains the unique strategy “always L ,” so the only IEWDS outcome is liquidation at $t = 1$ with payoff (2, 0). This matches part (a) exactly: backward induction in a finite extensive-form game of perfect information is precisely IEWDS applied at the agent normal form, period by period.

Problem 2.3

Step 1: Firm 2's best response. Given q_1 , firm 2 solves

$$\max_{q_2 \geq 0} (a - q_1 - q_2 - c)q_2 \implies q_2^*(q_1) = \frac{a - c - q_1}{2}$$

(assuming $q_1 \leq a - c$, so firm 2 produces a positive quantity).

Step 2: Firm 1 anticipates $q_2^*(q_1)$ and maximizes. The induced inverse residual demand is $p = a - q_1 - q_2^*(q_1) = (a + c - q_1)/2$, so firm 1's profit becomes

$$\pi_1(q_1) = (p - c)q_1 = \frac{(a - c - q_1)q_1}{2}.$$

First-order condition $a - c - 2q_1 = 0$ gives

$$q_1^* = \frac{a - c}{2}, \quad q_2^* = \frac{a - c}{4}.$$

Total quantity is $q^* = \frac{3(a-c)}{4}$ and the equilibrium price is $p^* = \frac{a+3c}{4}$.

Equilibrium profits.

$$\pi_1^* = (p^* - c)q_1^* = \frac{(a - c)^2}{8}, \quad \pi_2^* = (p^* - c)q_2^* = \frac{(a - c)^2}{16}.$$

Remark (Stackelberg vs. Cournot).

Recall the simultaneous-move Cournot equilibrium of the same duopoly produces $q_i^C = (a - c)/3$ and $\pi_i^C = (a - c)^2/9$. The Stackelberg leader produces *more* ($\frac{1}{2}$ vs. $\frac{1}{3}$ of $a - c$) and earns *more* ($\frac{1}{8} > \frac{1}{9}$) than the Cournot firm; the follower produces *less* and earns *less*. Total industry output rises from $\frac{2(a-c)}{3}$ to $\frac{3(a-c)}{4}$, so consumers also benefit from sequential play. The first-mover advantage rests on commitment: firm 1's choice of q_1 is observed before firm 2 moves, and firm 2 must accept it as fixed when optimizing.

Problem Set 3

Problems

Problem 3.1

Consider a finite n -player game in strategic form $G = (S_i, u_i)_{i=1}^n$. Define player i 's *maxmin payoff* as

$$w_i = \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}),$$

and his *minmax payoff* as

$$v_i = \min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}).$$

- Show that for all i , $v_i \geq w_i$.
- Suppose σ is a Nash equilibrium of the game G (possibly in mixed strategies). Show that for all i , $u_i(\sigma) \geq v_i$.

Problem 3.2

Consider the two-player game:

	L	R
U	1, 0	0, 1
D	$\frac{1}{2}, \frac{1}{3}$	1, 0

- For this game, find a Nash equilibrium σ (possibly in mixed strategies). Show that it is unique.
- Find player 1's maxmin strategy, that is, find the strategy σ_1 that solves

$$\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2).$$

Similarly, find player 2's maxmin strategy.

- Compare the maxmin strategies to the Nash equilibrium strategies. Are they the same?

Problem 3.3

Consider a finite two-player game $G = (S_i, u_i)_{i=1}^2$. Consider another game $H = (S_i, v_i)_{i=1}^2$ with payoff functions:

$$v_1(s_1, s_2) = u_1(s_1, s_2) + a(s_2),$$

where $a : S_2 \rightarrow \mathbb{R}$ is an arbitrary function, and

$$v_2(s_1, s_2) = u_2(s_1, s_2) + b(s_1),$$

where $b : S_1 \rightarrow \mathbb{R}$ is also arbitrary.

- (a) Show that the set of Nash equilibria (pure or mixed) of G is the same as the set of Nash equilibria of H .

Problem 3.4

Suppose $G = (S_i, u_i)_{i=1}^2$ is a finite two-player zero-sum game, that is, $u_1(s_1, s_2) = -u_2(s_1, s_2)$.

- (a) The Minmax Theorem says that in any finite two-player zero-sum game,

$$\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2).$$

Use the fact that every finite game has a Nash equilibrium (possibly in mixed strategies) to establish the Minmax Theorem.

- (b) Show that if (σ'_1, σ'_2) and (σ''_1, σ''_2) are two Nash equilibria of G , then these are interchangeable, that is, (σ'_1, σ''_2) and (σ''_1, σ'_2) are also Nash equilibria of G .
- (c) Prove that all Nash equilibria of G are payoff equivalent.
- (d) Do the properties derived in parts (b) and (c) hold in two-person non-zero-sum games? Provide examples in support of your answers.

Solutions

Problem 3.1

(a) $w_i \leq v_i$. For every σ_i, σ_{-i} ,

$$u_i(\sigma_i, \sigma_{-i}) \leq \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i}).$$

Taking the minimum over σ_{-i} on both sides preserves the inequality, so

$$\min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) \leq \min_{\sigma_{-i}} \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i}) = v_i.$$

The left-hand side depends on σ_i , and the right-hand side does not, so the inequality holds for the maximum over σ_i :

$$w_i = \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) \leq v_i.$$

(b) **Nash payoff** $\geq v_i$. In a NE σ^* , σ_i^* is a best response to σ_{-i}^* , so

$$u_i(\sigma^*) = \max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^*) \geq \min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}) = v_i.$$

Remark (Why v_i and not w_i is the relevant lower bound).

The Folk Theorem of Chapter 11 uses v_i (minmax), not w_i (maxmin), as the lower bound on individually rational payoffs. Part (a) shows the two coincide when max-min equals min-max, which holds in zero-sum games (Problem 3.4) but generally fails in non-zero-sum games. Part (b) explains why: opponents have flexibility to coordinate against player i , and the worst they can credibly do is the minmax—not the harsher maxmin that a paranoid player might fear.

Problem 3.2

(a) **Unique Nash equilibrium.**

No pure NE. A quick scan: at (U, L) player 2 prefers R (1 vs 0); at (U, R) player 1 prefers D (1 vs 0); at (D, L) player 1 prefers U (1 vs $\frac{1}{2}$); at (D, R) player 2 prefers L ($\frac{1}{3}$ vs 0). No cell is mutually best-responding.

Mixed NE. Let player 1 play U with probability p , player 2 play L with probability q . Player 2's indifference between L and R requires

$$p \cdot 0 + (1 - p) \cdot \frac{1}{3} = p \cdot 1 + (1 - p) \cdot 0 \implies p = \frac{1}{4}.$$

Player 1's indifference between U and D requires

$$q \cdot 1 + (1 - q) \cdot 0 = q \cdot \frac{1}{2} + (1 - q) \cdot 1 \implies q = \frac{2}{3}.$$

The unique NE is $(\sigma_1, \sigma_2) = ((\frac{1}{4}, \frac{3}{4}), (\frac{2}{3}, \frac{1}{3}))$ with payoffs $u_1(\sigma^*) = \frac{2}{3}$, $u_2(\sigma^*) = \frac{1}{4}$.

(b) Maxmin strategies.

Player 1. With $\sigma_1 = (p, 1 - p)$ and $\sigma_2 = (q, 1 - q)$,

$$u_1(p, q) = q \cdot \frac{1+p}{2} + (1-q)(1-p).$$

For fixed p , this is linear in q with coefficient $\frac{1+p}{2} - (1-p) = \frac{3p-1}{2}$.

- If $p < \frac{1}{3}$: minimum at $q = 1$, giving $\frac{1+p}{2}$.
- If $p > \frac{1}{3}$: minimum at $q = 0$, giving $1 - p$.
- If $p = \frac{1}{3}$: u_1 does not depend on q and equals $\frac{2}{3}$.

The maximum over p is reached at $p = \frac{1}{3}$, with value $w_1 = \frac{2}{3}$. Player 1's maxmin strategy plays U with probability $\frac{1}{3}$.

Player 2. Similarly,

$$u_2(p, q) = p(1 - q) + (1 - p)\frac{q}{3}.$$

Linear in p with coefficient $(1 - q) - \frac{q}{3} = 1 - \frac{4q}{3}$.

- If $q < \frac{3}{4}$: minimum at $p = 0$, giving $\frac{q}{3}$.
- If $q > \frac{3}{4}$: minimum at $p = 1$, giving $1 - q$.
- If $q = \frac{3}{4}$: u_2 is constant at $\frac{1}{4}$.

Maximum over q at $q = \frac{3}{4}$ with value $w_2 = \frac{1}{4}$. Player 2's maxmin strategy plays L with probability $\frac{3}{4}$.

(c) Comparison.

	Player 1 plays U	Player 2 plays L
NE strategy	$\frac{1}{4}$	$\frac{2}{3}$
Maxmin strategy	$\frac{1}{3}$	$\frac{3}{4}$

The two strategies coincide neither for player 1 nor for player 2; the game is non-zero-sum. The *values* happen to match—NE payoffs $(\frac{2}{3}, \frac{1}{4})$ equal the maxmin values (w_1, w_2) —a consequence of part 3.1(b), since here $w_i = v_i$ holds. In general, however, NE strategies differ from maxmin strategies, because in a NE each player optimizes against the *equilibrium* opponent strategy, not against an adversary trying to minimize their payoff.

Problem 3.3

The transformation $v_i(s_i, s_{-i}) = u_i(s_i, s_{-i}) + a_i(s_{-i})$ adds to player i 's payoff a term that depends only on the opponents' strategies, not on i 's own strategy. The key observation is that the difference $v_i(s_i, s_{-i}) - v_i(s'_i, s_{-i})$ is invariant:

$$v_i(s_i, s_{-i}) - v_i(s'_i, s_{-i}) = u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i}).$$

For mixed strategies the same identity holds because u_i is linear in σ_i :

$$v_i(\sigma_i, \sigma_{-i}) - v_i(\sigma'_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i}) - u_i(\sigma'_i, \sigma_{-i}).$$

Hence player i 's preference ranking over their own strategies, against any fixed opponent profile σ_{-i} , is identical in G and H . Best-response correspondences coincide, and so do the sets of Nash equilibria.

Remark (The transformation preserves more than just NE).

The same argument shows G and H have identical sets of rationalizable strategies, dominant-strategy equilibria, correlated equilibria, etc.—any solution concept defined through preference comparisons over own strategies. This is the formal underpinning of the “potential games” framework in Chapter 3: a potential game is one equivalent (in this sense) to a common-interest game in which all players share a single payoff function.

Problem 3.4

(a) Minmax Theorem from Nash existence. By Nash's existence theorem, the game has a NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$. We show $\max_{\sigma_1} \min_{\sigma_2} u_1 = \min_{\sigma_2} \max_{\sigma_1} u_1$.

\leq *direction.* Always holds (Problem 3.1(a)).

\geq *direction.* In the NE, σ_1^* is a best response to σ_2^* :

$$u_1(\sigma^*) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \geq \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2). \quad (\star)$$

Similarly, σ_2^* is a best response to σ_1^* in the zero-sum game, which means σ_2^* *minimizes* $u_1(\sigma_1^*, \cdot)$:

$$u_1(\sigma^*) = \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \leq \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2). \quad (\star\star)$$

Combining (\star) and $(\star\star)$,

$$\min_{\sigma_2} \max_{\sigma_1} u_1 \leq u_1(\sigma^*) \leq \max_{\sigma_1} \min_{\sigma_2} u_1,$$

which is the inequality opposite to Problem 3.1(a). The two together give equality.

(b) Interchangeability. Let V denote the common value $\max \min = \min \max$. For any NE σ , the proof of part (a) gives

$$u_1(\sigma) = V, \quad \min_{\sigma_2} u_1(\sigma_1, \sigma_2) = V, \quad \max_{\sigma_1} u_1(\sigma_1, \sigma_2) = V.$$

That is: any NE strategy σ_1 is a maxmin strategy for player 1, and any NE strategy σ_2 is a minmax strategy for player 2.

Now take two NE (σ'_1, σ'_2) and (σ''_1, σ''_2) . We show (σ'_1, σ''_2) is also a NE.

Player 1's best response. Since σ''_2 is a minmax strategy, $\max_{\sigma_1} u_1(\sigma_1, \sigma''_2) = V$. Since σ'_1 is a maxmin strategy, $\min_{\sigma_2} u_1(\sigma'_1, \sigma_2) = V$, so $u_1(\sigma'_1, \sigma''_2) \geq V$. Combining, $u_1(\sigma'_1, \sigma''_2) = V = \max_{\sigma_1} u_1(\sigma_1, \sigma''_2)$, so σ'_1 is a best response to σ''_2 .

Player 2's best response. In the zero-sum game, “best response for player 2” means minimizing u_1 . By the symmetric argument, σ''_2 minimizes $u_1(\sigma'_1, \cdot)$ at value V . So σ''_2 is a best response to σ'_1 .

The pair (σ'_1, σ''_2) satisfies both NE conditions. The same argument works for (σ''_1, σ'_2) .

(c) All NE are payoff-equivalent. The argument in (b) showed $u_1(\sigma) = V$ for every NE σ ; since $u_2 = -u_1$, also $u_2(\sigma) = -V$ for every NE. Payoffs are uniquely determined,

despite the multiplicity of equilibrium strategies.

(d) **Both fail in non-zero-sum games.** Battle of the Sexes:

	<i>O</i>	<i>F</i>
<i>O</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

The two pure NE are (O, O) and (F, F) .

Interchangeability fails. Mixing the strategies gives (O, F) with payoff $(0, 0)$. Player 1 prefers F over O given the opponent plays F (1 vs 0), so (O, F) is not a NE.

Payoff equivalence fails. The two pure NE deliver $(2, 1)$ and $(1, 2)$, opposite asymmetric outcomes.

Remark (The deeper reason).

In a zero-sum game, every pair of NE strategies generates the same value V , so the players' incentives to deviate from a mixed pair are determined entirely by V ; this forces interchangeability and payoff equivalence. In non-zero-sum games, player 1 has a strict preference between the two equilibria (here, (O, O) over (F, F)), and so does player 2 (in the opposite direction). The interplay creates a coordination problem with multiple non-equivalent solutions—the central feature distinguishing non-zero-sum from zero-sum strategic interaction.

Problem Set 4

Problems

Problem 4.1

Consider the following game:

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	2, 1	0, 3	0, 3
<i>D</i>	0, 3	2, 1	2, 0

- For this game, find all Nash equilibria of the game (possibly in mixed strategies).
- What are the rationalizable (pure) strategies for each player?
- Show that there is a rationalizable strategy for player 2 which is not played with positive probability in any Nash equilibrium.

Problem 4.2

Find all Nash equilibria (mixed or pure) of the following game:

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	2, 2	0, 3	1, 3
<i>D</i>	3, 2	1, 1	0, 2

Problem 4.3

Consider the following two-player game $G = (S_i, u_i)_{i=1}^2$ where for $i = 1, 2$, the set of pure strategies $S_i = [0, 1]$ (each player has a continuum of pure strategies) and the payoff function ($j \neq i$):

$$u_i(s_1, s_2) = \begin{cases} 1 - s_i & \text{if } s_i > s_j \\ -s_i & \text{if } s_i < s_j \\ \frac{1}{2} - s_i & \text{if } s_i = s_j \end{cases}$$

- Argue that this game has no pure-strategy Nash equilibria.
- Find a symmetric mixed-strategy equilibrium of this game.

Problem 4.4

Consider a Cournot duopoly in a market with inverse demand function $F(q)$ so that if the two firms produce quantities q_1 and q_2 , then the price is $p = F(q_1 + q_2)$. Suppose that F is continuous, non-increasing, and there exists \bar{q} such that for all $q \geq \bar{q}$, $F(q) = 0$. Moreover, $F(0)$ is finite. Each firm has constant per-unit cost of production c . The profit of firm i as a function of its own production q_i and its rival's production q_j is then

$$\pi_i(q_i, q_j) = F(q_i + q_j)q_i - cq_i.$$

(a) Consider the function

$$P(q_1, q_2) = q_1q_2(F(q_1 + q_2) - c).$$

Show that for all i , q_i, q_j and q'_i ,

$$\pi_i(q_i, q_j) - \pi_i(q'_i, q_j) > 0 \iff P(q_i, q_j) - P(q'_i, q_j) > 0.$$

(b) Using the answer to part (a), show that the Cournot duopoly specified above has a pure-strategy equilibrium (q_1^*, q_2^*) .

Solutions

Problem 4.1

(a) All Nash equilibria.

No pure NE. Cell-by-cell: in each of the six cells, at least one player has a strictly profitable deviation.

Mixed NE. Let $\sigma_1 = (p, 1 - p)$ and $\sigma_2 = (a, b, c)$ with $a + b + c = 1$. Player 2's expected payoffs against σ_1 are

$$u_2(L) = 3 - 2p, \quad u_2(M) = 1 + 2p, \quad u_2(R) = 3p.$$

Comparing M and R : $u_2(M) - u_2(R) = 1 - p > 0$ for all $p < 1$, so R is strictly worse than M whenever player 1 mixes nontrivially. Hence $c = 0$ in any mixed NE.

The reduced game restricted to L, M has player 2 indifferent when $3 - 2p = 1 + 2p$, i.e., $p = \frac{1}{2}$. Player 1's indifference between U and D given $\sigma_2 = (a, 1 - a, 0)$ requires

$$u_1(U) = 2a, \quad u_1(D) = 2(1 - a) \quad \implies \quad a = \frac{1}{2}.$$

So the unique NE is

$$\sigma_1^* = \left(\frac{1}{2}, \frac{1}{2}\right), \quad \sigma_2^* = \left(\frac{1}{2}, \frac{1}{2}, 0\right),$$

with payoffs $(u_1, u_2) = (1, 2)$.

(b) Rationalizable strategies. We claim every pure strategy is rationalizable. The argument uses the IUD = RAT equivalence (Chapter 3): a strategy is rationalizable iff it survives iterated elimination of strictly dominated strategies.

Player 1. U is a best response to $\sigma_2 = L$ (since $u_1(U | L) = 2 > 0 = u_1(D | L)$). D is a best response to $\sigma_2 = M$ (since $u_1(D | M) = 2 > 0 = u_1(U | M)$). Both U and D are rationalizable.

Player 2. L is a best response to $\sigma_1 = D$ (since $u_2(L | D) = 3, u_2(M | D) = 1, u_2(R | D) = 0$). M is a best response to $\sigma_1 = U$ (since $u_2(M | U) = u_2(R | U) = 3 > 1 = u_2(L | U)$, ties at the max). For R , the same belief $\sigma_1 = U$ gives $u_2(R | U) = 3$, tied for the max with M , so R is a (weak) best response. All three pure strategies are rationalizable.

(c) R is rationalizable but absent from every NE. The previous parts give the example. R is a best response only at the boundary belief $\sigma_1 = U$ pure: when $p < 1$, M strictly dominates R for player 2. Equilibrium consistency requires σ_1 to mix (as we showed in (a), no pure NE exists), so $\sigma_1 = U$ pure cannot be sustained as the equilibrium belief, and R is squeezed out.

Remark (Why this distinction matters).

Rationalizability requires *some* consistent belief; Nash equilibrium requires *mutually* consistent beliefs. In a NE, player 2's mixing probabilities must be supported by player 1's actual mixing, which here is $\frac{1}{2}/\frac{1}{2}$ —a belief that excludes R . The example shows that the gap between rationalizability and NE can be strict even in finite games of perfect information.

Problem 4.2

Pure NE. Direct check yields two pure NE: (U, R) with payoff $(1, 3)$ and (D, L) with payoff $(3, 2)$.

Mixed NE. For player 2's payoffs as functions of $\sigma_1 = (p, 1 - p)$:

$$u_2(L) = 2, \quad u_2(M) = 1 + 2p, \quad u_2(R) = 2 + p.$$

Family A: $p = 1$ (player 1 plays U pure). Then $u_2(M) = u_2(R) = 3 > 2 = u_2(L)$, so player 2 mixes between M and R . Write $\sigma_2 = (0, \alpha, 1 - \alpha)$. Player 1's payoffs: $u_1(U) = 1 - \alpha$, $u_1(D) = \alpha$; U remains a best response iff $\alpha \leq \frac{1}{2}$. So

$$\sigma_1 = U \text{ pure}, \quad \sigma_2 = (0, \alpha, 1 - \alpha), \quad \alpha \in [0, \frac{1}{2}].$$

At $\alpha = 0$ this is the pure NE (U, R) .

Family B: $p = 0$ (player 1 plays D pure). Then $u_2(L) = u_2(R) = 2 > 1 = u_2(M)$, so player 2 mixes between L and R . Write $\sigma_2 = (\alpha, 0, 1 - \alpha)$. Player 1's payoffs: $u_1(U) = 2\alpha + (1 - \alpha) = 1 + \alpha$, $u_1(D) = 3\alpha$; D best iff $3\alpha \geq 1 + \alpha$ iff $\alpha \geq \frac{1}{2}$. So

$$\sigma_1 = D \text{ pure}, \quad \sigma_2 = (\alpha, 0, 1 - \alpha), \quad \alpha \in [\frac{1}{2}, 1].$$

At $\alpha = 1$ this is the pure NE (D, L) .

No NE with player 1 strictly mixing. If $p \in (0, 1)$, indifference $u_1(U) = u_1(D)$ requires $2a + c = 3a + b$, i.e., $c = a + b$, hence $c = \frac{1}{2}$. Player 2 indifference among the strategies in support combined with $u_2(L) = 2, u_2(R) = 2 + p$ and $u_2(M) = 1 + 2p$ has no solution for $p \in (0, 1)$.

Summary. The full NE set is the union of the two one-parameter families:

$$\{(U, \alpha M + (1 - \alpha)R) : \alpha \in [0, \frac{1}{2}]\} \cup \{(D, \alpha L + (1 - \alpha)R) : \alpha \in [\frac{1}{2}, 1]\}.$$

Family A delivers payoffs $(1 - \alpha, 3)$, Family B delivers $(3\alpha, 2)$.

Problem 4.3

(a) No pure NE. Suppose (s_1, s_2) is a pure NE.

Asymmetric: $s_1 > s_2$. Player 2's payoff is $-s_2$. Deviating to $s_2 = 0$ yields $0 \geq -s_2$, strict if $s_2 > 0$, so $s_2 = 0$. Then player 1, holding all $s_1 > 0$, gets $1 - s_1$, which strictly increases as $s_1 \downarrow 0$. There is no minimum positive bid, so no equilibrium exists in this configuration.

Symmetric: $s_1 = s_2 = s$. Both get $\frac{1}{2} - s$. Deviating to $s + \varepsilon$ for small $\varepsilon > 0$ wins outright and gives $1 - s - \varepsilon > \frac{1}{2} - s$. Profitable deviation, so not a NE.

Hence no pure NE.

(b) Symmetric mixed-strategy NE. Look for a symmetric NE in which both players randomize according to a continuous CDF F on $[0, 1]$ (no point mass, so ties have probability zero). Conditional on the opponent's draw $X_2 \sim F$, player 1's expected payoff at the bid s is

$$u(s) = (1 - s)F(s) + (-s)(1 - F(s)) = F(s) - s.$$

For indifference across the support of F , $F(s) - s$ must equal a constant V throughout, i.e., $F(s) = V + s$.

For F to be a valid CDF on its support $[a, b] \subseteq [0, 1]$, we need $F(a) = 0$ and $F(b) = 1$, giving $a = -V$ and $b = 1 - V$. Combined with $a \geq 0$ and $b \leq 1$, we conclude $V = 0$, $[a, b] = [0, 1]$, and

$$F(s) = s, \quad s \in [0, 1].$$

The unique symmetric mixed NE has each player drawing S_i uniformly on $[0, 1]$, with equilibrium expected payoff $V = 0$.

Remark (Connection to the all-pay auction).

This game is isomorphic to a normalized two-bidder all-pay auction with private value 1: bidders pay s_i regardless of who wins, and the higher bidder takes the prize. The mixed NE is exactly the dissipation of the prize through pure rent-dissipation: in expectation, each bidder pays half the prize on average, leaving zero net rent.

Problem 4.4

(a) **Ordinal potential.** For any $q_i, q'_i, q_j > 0$,

$$P(q_i, q_j) - P(q'_i, q_j) = q_j [q_i(F(q_i + q_j) - c) - q'_i(F(q'_i + q_j) - c)] = q_j \cdot [\pi_i(q_i, q_j) - \pi_i(q'_i, q_j)].$$

Since $q_j > 0$, the sign of $P(q_i, q_j) - P(q'_i, q_j)$ is exactly the sign of $\pi_i(q_i, q_j) - \pi_i(q'_i, q_j)$. This proves the ordinal-potential equivalence on the strictly positive orthant.

(b) **Existence of a pure-strategy equilibrium.** Restrict to the compact set $[0, \bar{q}]^2$. The profit function is continuous in (q_1, q_2) , so P is continuous, and a maximizer (q_1^*, q_2^*) exists.

Case 1: $q_1^* > 0$ and $q_2^* > 0$. For any $q_1 \in [0, \bar{q}]$, since $q_2^* > 0$ part (a) applies and $P(q_1^*, q_2^*) \geq P(q_1, q_2^*)$ implies $\pi_1(q_1^*, q_2^*) \geq \pi_1(q_1, q_2^*)$. Symmetrically, q_2^* is a best response to q_1^* . Hence (q_1^*, q_2^*) is a NE.

Case 2: at least one $q_i^* = 0$. Then $P(q_1^*, q_2^*) = 0$. We show this can only happen when monopoly itself is unprofitable, in which case $(0, 0)$ is trivially a NE.

If a firm operating alone could earn positive profit—i.e., there exists $q > 0$ with $(F(q) - c)q > 0$, equivalently $F(q) > c$ for some $q > 0$ —then choose any such q and consider the point (q, q) . Provided $F(2q) > c$, $P(q, q) = q^2(F(2q) - c) > 0$, contradicting that the maximum value is 0. By continuity of F and the assumption $F(0) > c$ (otherwise nobody produces and $(0, 0)$ is the trivial NE), there exists q small enough that $F(2q)$ remains close to $F(0) > c$, hence $P > 0$ somewhere. Therefore Case 2 cannot arise unless the monopoly itself is unprofitable, in which case both firms produce zero in equilibrium.

In both cases a pure-strategy NE exists.

Remark (Why “ordinal” potential is enough).

An exact potential function (Chapter 3) requires $\pi_i(q_i, q_j) - \pi_i(q'_i, q_j) = P(q_i, q_j) - P(q'_i, q_j)$. The Cournot game with linear demand admits one (the lecture notes derived $P_{\text{linear}}(q_1, q_2) = a(q_1 + q_2) - b(q_1^2 + q_2^2) - bq_1q_2 - c(q_1 + q_2)$); but with a general non-linear demand F , only the *ordinal* version survives, scaled by q_j . The existence proof

above shows that ordinal potentials are sufficient for guaranteeing pure-strategy NE: what matters is that the maximizer of P tracks players' joint best-response logic, not the exact magnitude of P .

Problem Set 5

Problems

Problem 5.1

Consider the game G :

	y	y'	y''
x	0, 6	0, 10	15, 5
x'	-20, 1	-2, 3	14, 4
x''	10, 0	-10, 2	0, 10

- Show that G is a supermodular game. (*Note:* You may have to reorder the strategies.)
- Does G have any mixed-strategy equilibria?

Problem 5.2

Suppose $G = (S_i, u_i)_{i=1,2}$ is a two-player game where for each i , $S_i = [0, 1]$ and u_i is supermodular. Suppose that G has exactly three pure-strategy equilibria, say s' , s'' and s''' . (There may or may not be other mixed-strategy equilibria.)

- Show that these can be ordered according to the vector ordering, that is, $s' \leq s'' \leq s'''$, perhaps after renaming.
- Now suppose that in addition, for each i , the payoff function u_i is increasing in s_j , $j \neq i$. Show that the payoffs from the equilibria can also be ordered, that is, $u(s') \leq u(s'') \leq u(s''')$, where $u(s') = (u_1(s'), u_2(s'))$, etc.

Problem 5.3

Consider a market with two symmetric firms making similar but differentiated products. The demand for firm i 's product is given by the function $Q_i(p_1, p_2)$ where Q_i satisfies $\frac{\partial Q_i}{\partial p_i} < 0$, $\frac{\partial Q_i}{\partial p_j} > 0$, and $\frac{\partial^2 Q_i}{\partial p_i \partial p_j} > 0$ (where $j \neq i$). Suppose each firm has a constant cost of production of c per unit. Let $\pi_i(p_1, p_2)$ denote the profits of firm i when the two firms choose prices p_1 and p_2 . Since the firms are symmetric, for any pair of prices p and p' , $Q_1(p, p') = Q_2(p', p)$.

- Show that if each firm chooses prices, then this defines a supermodular game.
- Now suppose that when firms choose prices simultaneously, the resulting game has a unique equilibrium (p_1^*, p_2^*) . Consider an alternative specification in which firms

choose prices sequentially: firm 1 chooses p_1 first, then firm 2 chooses p_2 knowing the price set by firm 1. Suppose $(p_1^\#, p_2^\#)$ is the unique subgame perfect equilibrium of the sequential game. (*Note:* You may assume both prices are positive in both equilibria.)

1. Show that $\pi_1(p_1^\#, p_2^\#) > \pi_1(p_1^*, p_2^*)$.
2. (Harder) Show that $\pi_2(p_1^\#, p_2^\#) > \pi_1(p_1^\#, p_2^\#)$.

Problem 5.4

Demand in a declining industry is given by the function

$$P_t = (100 - t) - Q_t,$$

where $t = 0, 1, \dots$ denotes time, P_t is the price at time t , and Q_t is total production at time t . Demand decreases over time and vanishes at $t = 100$. There are two firms in the industry at $t = 0$: firm 1 has production capacity $K_1 = 40$ and constant production cost $c_1 = 30$ per unit; firm 2 has $K_2 = 20$ and $c_2 = 40$. Production is “all-or-nothing”: in any period, each firm produces its capacity or shuts down (true of some chemical industries). Each firm exits the industry (shuts down forever) once it anticipates no future profits. Once a firm exits it cannot reenter.

- (a) Show that there is a Nash equilibrium of the game in which firm 1 exits at time 10 and firm 2 produces for some time after that. Show that there is another Nash equilibrium in which firm 2 exits at time 0 and firm 1 produces for some time.
- (b) Find a subgame perfect equilibrium of the game. Is it unique? Which firm exits first?

Solutions

Problem 5.1

(a) **Strategic complementarities and unique pure NE via IESDS.** The game admits a chain of strict dominations:

1. For player 1, x strictly dominates x' : row x delivers $(0, 0, 15)$ while row x' delivers $(-20, -2, 14)$, so x beats x' at every column.
2. After eliminating x' , for player 2, y' strictly dominates y : y' delivers $(10, 2)$ versus y 's $(6, 0)$ in the reduced game.
3. After eliminating y , for player 1, x now strictly dominates x'' in the $\{y', y''\}$ -restricted game: $(0, 15)$ versus $(-10, 0)$.
4. After eliminating x'' , y' strictly dominates y'' for player 2: $10 > 5$.

The lone surviving profile is (x, y') with payoff $(0, 10)$, the unique pure NE.

Reordering for partial supermodularity. If we reorder rows as $x'' < x < x'$ and columns as $y < y' < y''$, player 1's payoff matrix has increasing differences in (s_1, s_2) across all three row pairs:

$$\Delta_{x'' \rightarrow x} = (-10, 10, 15), \quad \Delta_{x \rightarrow x'} = (-20, -2, -1), \quad \Delta_{x'' \rightarrow x'} = (-30, 8, 14),$$

each non-decreasing in column index. Player 2's increasing-differences condition fails under any single ordering, so G is not strictly supermodular as a two-player game. The hint “you may have to reorder strategies” covers the row-only restoration of monotonicity. Pure-NE existence then follows directly from IESDS rather than from a Tarski-type argument.

(b) **Mixed equilibria.** Because IESDS converges to the unique profile (x, y') , no strictly dominated strategy can appear in the support of any NE (whether pure or mixed). Hence the only NE is the pure (x, y') , and the game has *no* mixed-strategy equilibria.

Problem 5.2

(a) **The three equilibria are totally ordered.** A general theorem (Topkis-Milgrom-Roberts) states that the set of pure NE in a supermodular game is a non-empty complete lattice under the componentwise vector order: if s and \hat{s} are NE, so are $s \vee \hat{s}$ and $s \wedge \hat{s}$.

Suppose, for contradiction, that two of the three NE—say s' and s'' —are incomparable in the vector order. Then $s' \wedge s''$ and $s' \vee s''$ are both distinct from s' and s'' , and by the lattice property both are NE. Hence the equilibrium set contains at least the four points $\{s', s'', s' \wedge s'', s' \vee s''\}$, plus s''' , giving five distinct elements—contradicting the hypothesis of exactly three. Therefore each pair among s', s'', s''' is comparable, and after relabeling

$$s' \leq s'' \leq s'''.$$

(b) **Payoffs are ordered when payoffs are increasing in opponents' actions.** Take any two ordered NE $s' \leq s''$. By the NE property, s'_1 is a best response to s''_2 :

$$u_1(s'_1, s''_2) \geq u_1(s'_1, s'_2). \quad (*)$$

Since u_1 is increasing in s_2 and $s'_2 \leq s''_2$,

$$u_1(s'_1, s''_2) \geq u_1(s'_1, s'_2). \quad (**)$$

Combining (*) and (**), $u_1(s'') \geq u_1(s')$. The symmetric argument yields $u_2(s'') \geq u_2(s')$. Iterating along $s' \leq s'' \leq s'''$ gives $u(s') \leq u(s'') \leq u(s''')$.

Remark (Order vs. optimality).

This proposition is the source of a useful comparative-static heuristic in supermodular games: if exogenous shifts move all best-response curves upward, the highest-payoff NE moves up and stays the highest, while the lowest moves up and stays the lowest. Coordination in the “high” equilibrium is therefore robust to shocks that shift best responses up.

Problem 5.3

(a) **Pricing is supermodular.** Compute

$$\frac{\partial \pi_i}{\partial p_i} = Q_i + (p_i - c) \frac{\partial Q_i}{\partial p_i},$$

and the cross partial:

$$\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} = \frac{\partial Q_i}{\partial p_j} + (p_i - c) \frac{\partial^2 Q_i}{\partial p_i \partial p_j}.$$

Both terms are non-negative under the assumptions (and the second strictly positive whenever $p_i > c$), so $\partial^2 \pi_i / \partial p_i \partial p_j > 0$. Pricing exhibits **strategic complementarities**: each firm’s best-response price is non-decreasing in the rival’s price.

(b.1) **Firm 1 strictly prefers the sequential outcome.** Firm 1, moving first, solves

$$\max_{p_1} \pi_1(p_1, B_2(p_1)),$$

where B_2 is firm 2’s best-response function. By the envelope theorem,

$$0 = \frac{\partial \pi_1}{\partial p_1}(p_1^\#, p_2^\#) + \frac{\partial \pi_1}{\partial p_2}(p_1^\#, p_2^\#) \cdot B_2'(p_1^\#).$$

At the simultaneous NE, $\partial \pi_1 / \partial p_1(p_1^*, p_2^*) = 0$. Plugging $p_1 = p_1^*$ into the sequential first-order expression,

$$\left. \frac{d}{dp_1} \pi_1(p_1, B_2(p_1)) \right|_{p_1=p_1^*} = \underbrace{\frac{\partial \pi_1}{\partial p_2}}_{>0} \cdot \underbrace{B_2'(p_1^*)}_{>0} > 0.$$

So firm 1’s sequential payoff is strictly increasing at $p_1 = p_1^*$; the sequential optimum lies above p_1^* and yields strictly higher profit:

$$\pi_1(p_1^\#, p_2^\#) > \pi_1(p_1^*, p_2^*).$$

(b.2) **Firm 2 does even better than firm 1 in the sequential game.** By symmetry,

$\pi_1(p, q) = \pi_2(q, p)$, so

$$\pi_2(p_1^\#, p_2^\#) > \pi_1(p_1^\#, p_2^\#) \iff \pi_1(p_2^\#, p_1^\#) > \pi_1(p_1^\#, p_2^\#).$$

Define $f(p_1) := \pi_1(p_1, B_2(p_1)) - \pi_1(B_2(p_1), p_1)$. At the simultaneous NE, $B_2(p^*) = p^*$, so $f(p^*) = 0$. Differentiating and evaluating at $p_1 = p^*$:

$$f'(p^*) = -(1 - B_2'(p^*)) \cdot \frac{\partial \pi_1}{\partial p_2}(p^*, p^*).$$

Stability of the simultaneous NE requires $|B_2'(p^*)| < 1$, so $1 - B_2'(p^*) > 0$. Combined with $\partial \pi_1 / \partial p_2 > 0$, we get $f'(p^*) < 0$. Since $p_1^\# > p^*$ and $f(p^*) = 0$, $f(p_1^\#) < 0$, i.e.,

$$\pi_1(p_1^\#, p_2^\#) < \pi_1(p_2^\#, p_1^\#) = \pi_2(p_1^\#, p_2^\#).$$

Remark (Second-mover advantage in price competition).

In Stackelberg *quantity* competition (Cournot, Problem 2.3), the leader profits more than the follower. Here, in Stackelberg *price* competition with strategic complementarities, the *follower* profits more—a sharp reversal. With strategic complementarities, a high price by the leader pulls the follower’s best-response price up but not as far. Both firms charge above the simultaneous NE, but the follower keeps a slightly lower price and captures more demand at higher margin. The leader effectively pays for the joint upward shift; the follower free-rides on it.

Problem 5.4

Per-period profits in each market structure:

Configuration	Firm 1’s profit	Firm 2’s profit
Both active ($Q_t = 60$)	$(40 - t - 30) \cdot 40 = 400 - 40t$	$(40 - t - 40) \cdot 20 = -20t$
Firm 1 alone ($Q_t = 40$)	$(60 - t - 30) \cdot 40 = 1200 - 40t$	0
Firm 2 alone ($Q_t = 20$)	0	$(80 - t - 40) \cdot 20 = 800 - 20t$

Key thresholds: duopoly is profitable for firm 1 only while $t < 10$; firm 2 strictly loses money in any duopoly period $t \geq 1$. Monopoly is profitable for firm 1 while $t < 30$, for firm 2 while $t < 40$.

(a) Two Nash equilibria.

NE I: Firm 1 exits at $t = 10$, firm 2 stays until $t = 40$.

Strategies: firm 1 produces while $t < 10$, exits at $t = 10$; firm 2 produces until $t = 40$.

Firm 1’s BR: per-period profit $400 - 40t$, strictly positive for $t < 10$, zero at $t = 10$, negative thereafter. Total firm 1 profit: $\sum_{t=0}^9 (400 - 40t) = 4000 - 40 \cdot 45 = 2200$. Exiting earlier forfeits positive periods; staying past $t = 10$ generates losses. So $\tau_1 = 10$ is a best response.

Firm 2’s BR: cumulative loss in periods 0–9 (duopoly) is $\sum_{t=0}^9 (-20t) = -900$. Then monopoly from $t = 10$ to $t = 39$: $\sum_{t=10}^{39} (800 - 20t) = 24,000 - 14,700 = 9300$. Net firm 2 profit: $-900 + 9300 = 8400 > 0$. Exit-at- $t = 0$ would give $0 < 8400$, so staying is optimal.

NE II: Firm 2 exits at $t = 0$, firm 1 monopolizes until $t = 30$.

Strategies: firm 2 exits at $t = 0$; firm 1 produces until $t = 30$.

Firm 2's BR: if firm 1 commits to "stay forever," firm 2's continuation under staying is negative, so exiting at $t = 0$ is optimal.

Firm 1's BR: alone, firm 1 makes $1200 - 40t$ per period, profitable until $t = 30$. Total: $\sum_{t=0}^{29} (1200 - 40t) = 36000 - 17400 = 18,600$. Exiting earlier wastes profitable periods.

(b) Subgame perfect equilibrium. NE II relies on firm 1 *threatening* to "stay forever" even if firm 2 stays in. But that threat is not credible: in a duopoly subgame at $t = 10$, firm 1's per-period profit is 0, and at $t = 11, 12, \dots$ it becomes increasingly negative. In any SPE, firm 1 must exit at $t = 10$ as soon as duopoly stops being profitable.

Walking the equilibrium logic forward: anticipating firm 1's SPE-credible exit at $t = 10$, firm 2 endures losses for 0–9 in exchange for 30 periods of monopoly afterward (net +8400, so firm 2 stays). Knowing firm 2 will not exit during the duopoly phase, firm 1 collects $400 - 40t$ during $t < 10$ and exits at $t = 10$.

The unique SPE is NE I: firm 1 exits at $t = 10$, firm 2 monopolizes from $t = 10$ to $t = 40$.

Remark (Why the smaller, less efficient firm wins).

Firm 1 has lower marginal cost (30 vs. 40) and twice the capacity, yet it exits first. The reason: firm 1's duopoly profitability threshold ($t = 10$) arrives *far earlier* than firm 2's monopoly threshold ($t = 40$), giving firm 2 a long enough post-firm-1 monopoly window to absorb the duopoly losses. In declining industries, exit is determined less by static cost advantages than by who can credibly outwait whom—and a firm with lower fixed losses per duopoly period and access to a long monopoly tail can outlast a more efficient rival who hits its zero-profit duopoly point first.

Problem Set 6

Problems

Problem 6.1

Let (U, d) denote a two-person bargaining problem, where $U \subseteq \mathbb{R}^2$ is a compact convex set of feasible utilities and $d \in U$ is the disagreement point. Suppose also that there exists a $v \in U$ such that $v \gg d$. Consider the bargaining solution K defined by: $K(U, d) = v^*$ is an efficient point in U such that

$$\frac{v_2^* - d_2}{v_1^* - d_1} = \frac{\bar{v}_2 - d_2}{\bar{v}_1 - d_1},$$

where $\bar{v}_i = \max\{v_i : (v_1, v_2) \in U\}$ is the highest possible utility for i in U .

- (a) Which of Nash's four axioms does K satisfy and which of these does it not satisfy?

Problem 6.2

Two people are bargaining over a "pie" of size 1. The bargaining lasts for T periods and proceeds as follows. In period 1, player 1 makes an offer $(x, 1 - x)$ to player 2, who then responds with "yes" or "no." If player 2 says "yes," the game is over and a split of x to player 1 and $1 - x$ to player 2 results. If player 2 says "no," then in period 2, player 2 makes a counteroffer $(y, 1 - y)$ to which player 1 responds with "yes" or "no." Again, if player 1 says "yes," the game is over with a split of y to player 1 and $1 - y$ to player 2. If player 1 says "no," then in period 3 player 1 makes an offer, etc. Both players use a common discount factor $\delta < 1$. Thus if a split $(z, 1 - z)$ is agreed to in some period $t \leq T$, then the payoffs to the two players are $\delta^{t-1}z$ and $\delta^{t-1}(1 - z)$, respectively. Denote the T -period game in which 1 makes the first offer by $G_1(T)$ and the T -period game in which 2 makes the first offer by $G_2(T)$.

- (a) Find the unique subgame perfect equilibrium outcome of $G_1(T)$. Let $(x(T), 1 - x(T))$ be the resulting split.
- (b) Find the limit of $x(T)$ as $T \rightarrow \infty$.

Solutions

Problem 6.1

We check the four Nash axioms one by one.

Scale invariance: ✓. A scale transformation $\bar{v}_i = \alpha_i v_i + \beta_i$ (with $\alpha_i > 0$) carries every quantity in the K-S definition through proportionally:

$$\bar{v}_i^* - \bar{d}_i = \alpha_i(v_i^* - d_i), \quad \bar{\bar{v}}_i - \bar{\bar{d}}_i = \alpha_i(\bar{v}_i - d_i),$$

so the ratio $(\bar{v}_i^* - \bar{d}_i)/(\bar{\bar{v}}_i - \bar{\bar{d}}_i)$ is unchanged. Efficiency of v^* is also preserved by affine rescaling.

Efficiency: ✓. By construction, $K(U, d)$ is an efficient (Pareto-optimal) point of U .

Symmetry: ✓. If U is symmetric $((v_1, v_2) \in U \Leftrightarrow (v_2, v_1) \in U)$ and $d_1 = d_2$, then $\bar{v}_1 = \bar{v}_2$ by symmetry. The K-S equation reduces to $v_1^* - d_1 = v_2^* - d_2$, i.e., $v_1^* = v_2^*$. The solution lies on the symmetry axis.

IIA: × (fails). The Kalai-Smorodinsky solution *violates* independence of irrelevant alternatives because \bar{v}_i depends on the entire feasible set U : enlarging U in a direction that improves an ideal point shifts the K-S solution even when the original solution remains feasible.

Counterexample. Let $d = (0, 0)$.

- $U =$ triangle with vertices $(0, 0), (1, 0), (0, 1)$. Ideals $(\bar{v}_1, \bar{v}_2) = (1, 1)$. K-S solution $v^* = (\frac{1}{2}, \frac{1}{2})$.
- $U' =$ unit square (convex hull of $(0, 0), (1, 0), (1, 1), (0, 1)$), which contains U . Ideals are still $(1, 1)$. K-S solution: along the diagonal at the new efficient frontier, $v^* = (1, 1)$.

We have $(\frac{1}{2}, \frac{1}{2}) \in U \subseteq U'$, but $K(U) = (\frac{1}{2}, \frac{1}{2}) \neq (1, 1) = K(U')$. Enlarging the feasible set changes the K-S solution *even when the original solution remains feasible*, violating IIA.

Remark (Why K-S sacrifices IIA).

The Kalai-Smorodinsky and Nash solutions differ in how they respond to “improving the outside option” of one party. K-S anchors the bargaining ratio to the relative ideal points \bar{v}_i , so changes to either ideal change the solution—a feature its proponents view as a virtue (the solution rewards being in a position to demand more), but which formally violates IIA. Nash, by contrast, uses only a local hyperbola through the disagreement point, ignoring the shape of the feasible set away from the chosen point. The two axiomatizations diverge on this question of “should I care about my best alternative?”—and lead to genuinely different bargaining predictions.

Problem 6.2

(a) Unique SPE outcome of $G_1(T)$. Define $x(T) =$ player 1’s share in the unique SPE outcome of $G_1(T)$, and $y(T) =$ player 2’s share in the unique SPE outcome of $G_2(T)$. By symmetry of structure, $x(T) = y(T)$ (the “first-mover share” depends only on the horizon).

Backward induction recursion. In $G_1(T)$, player 1 offers player 2 the smallest amount that just clears player 2’s continuation value. If player 2 rejects, the game enters $G_2(T - 1)$

with one period of delay, where player 2 (as proposer) earns $y(T - 1)$. Player 2 accepts at $1 - x(T) \geq \delta y(T - 1)$, so the proposer takes

$$x(T) = 1 - \delta y(T - 1) = 1 - \delta x(T - 1).$$

With base case $x(1) = 1$ (in a one-period game the proposer takes everything),

$$x(2) = 1 - \delta, \quad x(3) = 1 - \delta + \delta^2, \quad x(4) = 1 - \delta + \delta^2 - \delta^3, \dots$$

The pattern is the alternating geometric sum

$$x(T) = \sum_{k=0}^{T-1} (-\delta)^k = \frac{1 - (-\delta)^T}{1 + \delta}.$$

Splitting by parity:

$$x(T) = \begin{cases} \frac{1 - \delta^T}{1 + \delta}, & T \text{ even,} \\ \frac{1 + \delta^T}{1 + \delta}, & T \text{ odd.} \end{cases}$$

Player 2's equilibrium share is $1 - x(T) = (\delta + (-\delta)^T)/(1 + \delta)$.

Sanity checks. $x(1) = 1$: player 1 takes all. $x(2) = 1 - \delta$: player 1 “bribes” player 2 with δ , the value of becoming the proposer next period. $x(3) = 1 - \delta + \delta^2$: an extra round of indirect proposing benefits player 1.

(b) Limit as $T \rightarrow \infty$. As $T \rightarrow \infty$, $(-\delta)^T \rightarrow 0$ since $|\delta| < 1$, so

$$\lim_{T \rightarrow \infty} x(T) = \frac{1}{1 + \delta}, \quad \lim_{T \rightarrow \infty} (1 - x(T)) = \frac{\delta}{1 + \delta}.$$

This is exactly the Rubinstein infinite-horizon split for linear utility $u(z) = z$. As $\delta \rightarrow 1$, both shares converge to $\frac{1}{2}$, recovering the symmetric Nash bargaining solution.

Remark (Two ways to read the recursion).

Each finite-horizon period adds a layer of forward-looking strategic substitution: the proposer pays the responder's discounted continuation rather than the full pie, but the continuation itself shrinks in proportion to δ . The alternating sum captures this nesting exactly. As T grows, the deepest layers contribute negligibly (they scale as δ^{T-1}), and the limit is determined entirely by the first two terms blown up by the geometric correction $1/(1 + \delta)$.

Problem Set 7

Problems

Problem 7.1

Complete the proof of Rubinstein's Theorem by showing that $(m_1, m_2) = (x^*, y^*)$. (*Note:* see the "Notes on Strategic Bargaining" for a statement of the result, notation, etc.)

Problem 7.2

Consider the following variant of Rubinstein's alternating-offer infinite-horizon model of bargaining. In each period, the proposer is chosen at random. Specifically, at the beginning of each period a coin is tossed. The probability that the coin turns up "heads" is p and the probability that it turns up "tails" is $1 - p$, where $0 < p < 1$. If the coin turns up "heads," player 1 makes an offer $(x, 1 - x)$ to player 2 who responds with "yes" or "no." If the coin turns up "tails," then player 2 makes an offer $(y, 1 - y)$ to player 1 who responds with "yes" or "no." If no agreement is reached in period 1, then the coin is tossed again at the beginning of period 2 and depending on the outcome, player 1 or player 2 makes another offer. Players have a common discount factor $\delta \in (0, 1)$ and linear utilities, that is, $u_i(z) = z$.

- (a) Construct a subgame perfect equilibrium of this bargaining game. How does the solution change with p ?
- (b) Is the equilibrium from part (a) unique? If so, provide a proof.

Problem 7.3

Consider a three-person bargaining game with alternating offers defined as follows. The three players have to divide a "pie" of size 1. In period 1, player 1 makes an offer (x_1, x_2, x_3) where $x_1 + x_2 + x_3 = 1$ and all $x_i \geq 0$. Then players 2 and 3 respond simultaneously with "yeses" and "nos." If both say "yes," the game ends with the division (x_1, x_2, x_3) . If either player 2 or player 3 says "no," then in the next period player 2 makes an offer (y_1, y_2, y_3) to which players 1 and 3 respond simultaneously with "yeses" and "nos." If both say "yes," the game ends with the division (y_1, y_2, y_3) . If either player 1 or player 3 says "no," then in the next period player 3 makes an offer (z_1, z_2, z_3) , and so on. All players have linear utilities $u(w) = w$ and discount future payoffs with a common discount factor $\delta \in (0, 1)$.

- (a) Find a stationary subgame perfect equilibrium of this game.

(b) (Hard) Can you find another subgame perfect equilibrium of this game?

Solutions

Problem 7.1

The lecture notes derive $(M_1, M_2) = (x^*, y^*)$, where M_i is the supremum of player 1's discounted share over all SPE outcomes of G_i . We show by symmetric reasoning that $(m_1, m_2) = (x^*, y^*)$ as well, where m_i is the corresponding infimum.

Recall x^*, y^* solve

$$\begin{aligned} u_1(y^*) &= \delta u_1(x^*), \\ u_2(1 - x^*) &= \delta u_2(1 - y^*). \end{aligned}$$

Claim 1': $u_1(m_2) \geq \delta u_1(m_1)$. In G_2 player 2 proposes; player 1 (the responder) will *reject* any offer y with $u_1(y) < \delta u_1(m_1)$, because rejection gives player 1 a continuation in G_1 delivering at least $u_1(m_1)$ at the next period, hence at least $\delta u_1(m_1)$ in present value. So in any SPE outcome of G_2 with agreement reached, player 1's share y satisfies $u_1(y) \geq \delta u_1(m_1)$. Taking the infimum, $u_1(m_2) \geq \delta u_1(m_1)$.

Claim 2': $u_1(m_2) \leq \delta u_1(m_1)$. Suppose $v_1 \geq m_1$ is the present value to player 1 of an SPE outcome of G_1 . Then there is an SPE of G_2 in which player 2 begins by offering y such that $u_1(y) = \delta u_1(v_1)$, with the punishment: if player 1 rejects, switch to the SPE of G_1 that delivers v_1 to player 1. Player 1 accepts (indifference broken in favor of agreement). This delivers $u_1(y) = \delta u_1(v_1)$ to player 1 in G_2 . Taking $v_1 = m_1$, $u_1(m_2) \leq \delta u_1(m_1)$.

Combining, $u_1(m_2) = \delta u_1(m_1)$.

Claim 3': $u_2(1 - m_1) = \delta u_2(1 - m_2)$. In G_1 , player 2 (responder) accepts iff $u_2(1 - z) \geq \delta u_2(1 - m_2)$ (since $1 - m_2$ is the maximum player 2 can hope for in G_2). For agreement to occur with player 1 getting m_1 (the lowest possible), player 2's share $1 - m_1$ must clear the highest possible continuation value: $u_2(1 - m_1) \leq \delta u_2(1 - m_2)$. The reverse inequality follows by the same constructive argument as Claim 2' applied to G_1 . Hence $u_2(1 - m_1) = \delta u_2(1 - m_2)$.

Conclusion. The pair (m_1, m_2) satisfies

$$\begin{aligned} u_1(m_2) &= \delta u_1(m_1), \\ u_2(1 - m_1) &= \delta u_2(1 - m_2), \end{aligned}$$

which is exactly the system defining (x^*, y^*) . Since this system has a unique solution (by strict concavity of u_i), $(m_1, m_2) = (x^*, y^*) = (M_1, M_2)$. The infimum equals the supremum, so every SPE outcome of G_i delivers exactly (x^*, y^*) .

Problem 7.2

(a) Stationary SPE. Look for a stationary SPE in which each player's continuation value (at the start of any period before the coin flip) is constant: v_1 for player 1, v_2 for player 2.

Offers. When player 1 proposes, the responder's continuation value next period is v_2 , so player 1 offers player 2 the share δv_2 and keeps $1 - \delta v_2$. When player 2 proposes, symmetrically, player 1 receives δv_1 and player 2 keeps $1 - \delta v_1$.

Stationary value equations.

$$v_1 = p(1 - \delta v_2) + (1 - p)\delta v_1, \quad v_2 = p\delta v_2 + (1 - p)(1 - \delta v_1).$$

The first rearranges to $v_1[1 - (1 - p)\delta] = p - p\delta v_2$, the second to $v_2[1 - p\delta] = (1 - p) - (1 - p)\delta v_1$.

Solving. Adding the two yields $v_1 + v_2 = 1$. Substituting $v_2 = 1 - v_1$ into the first:

$$v_1[1 - (1 - p)\delta] = p(1 - \delta) + p\delta v_1 \implies v_1(1 - \delta) = p(1 - \delta) \implies v_1 = p.$$

And $v_2 = 1 - p$.

Equilibrium offers. When player 1 proposes (probability p), she keeps $1 - \delta(1 - p) = 1 - \delta + \delta p$ and gives $\delta(1 - p)$ to player 2. When player 2 proposes (probability $1 - p$), she keeps $1 - \delta p$ and gives δp to player 1.

Comparative statics in p . As p increases, $v_1 = p$ rises linearly and $v_2 = 1 - p$ falls. In the limit $p \rightarrow 1$, the model collapses to “player 1 always proposes,” yielding $v_1 \rightarrow 1$ and $v_2 \rightarrow 0$.

(b) Uniqueness. Apply the supremum-infimum argument from Problem 7.1. Define M and m as the supremum and infimum of player 1’s discounted share over all SPEs. By stationarity, M and m are the same regardless of period. The reasoning gives

$$M = p(1 - \delta(1 - M)) + (1 - p)\delta M,$$

which simplifies to $M(1 - \delta) = p(1 - \delta)$, so $M = p$. The same equation governs m , giving $m = p = M$. The SPE value is unique, hence the SPE itself is unique.

Problem 7.3

(a) A stationary SPE. Let $V_i(j)$ denote player i ’s expected discounted equilibrium payoff at the start of a period in which player j proposes. Stationarity (the game’s structure is invariant under cyclical shifts) implies $V_i(j) = V_{i+1}(j + 1)$ (indices mod 3), so the values depend only on $i - j \pmod{3}$.

Best-response logic. In equilibrium, the proposer offers the minimum acceptable to each responder. If player j proposes and player i rejects, the cycle advances: next period player $j + 1$ proposes, and player i has a continuation value of $V_i(j + 1)$. The responder accepts iff her share is at least $\delta V_i(j + 1)$.

Specializing. Write $a = V_j(j)$ (proposer’s share), $b = V_{j+1}(j)$ (next-proposer’s share), $c = V_{j-1}(j)$ (third player’s share). At a period where j proposes:

- Responder $j + 1$ will propose next period, continuation = a , receives $x_{j+1} = \delta a$.
- Responder $j - 1$ becomes the proposer-after-next, continuation = b , receives $x_{j-1} = \delta b$.

Hence

$$V_j(j) = a = 1 - \delta a - \delta b, \quad V_{j+1}(j) = b = \delta a, \quad V_{j-1}(j) = c = \delta b = \delta^2 a.$$

Substituting $b = \delta a$:

$$a = 1 - \delta a - \delta^2 a \implies a(1 + \delta + \delta^2) = 1 \implies a = \frac{1}{1 + \delta + \delta^2}.$$

Stationary SPE shares. The proposer offers

$$\left(\frac{1}{1+\delta+\delta^2}, \frac{\delta}{1+\delta+\delta^2}, \frac{\delta^2}{1+\delta+\delta^2} \right),$$

where the proposer keeps the largest share, the next-in-line gets the middle share, and the third player gets the smallest. Both responders accept (indifferent at threshold). As $\delta \rightarrow 1$, all three shares converge to $\frac{1}{3}$.

(b) Other SPEs. Three-player alternating-offer bargaining games admit *multiple* SPE outcomes for any fixed $\delta < 1$, not just the stationary one. With two responders, each unilaterally has the power to delay agreement, and the equilibrium concept does not pin down a unique split when we allow non-stationary punishments.

Construction sketch. Fix a target split (x_1, x_2, x_3) with each $x_i > 0$. Construct strategies:

- In period 1, player 1 proposes (x_1, x_2, x_3) . Players 2 and 3 both accept.
- If anyone deviates, transition to a “stationary punishment SPE” in which the deviator gets the smallest share.

The deviator faces a payoff drop from the target x_i to the smallest stationary share $\delta^2/(1+\delta+\delta^2)$. For δ close to 1, the loss from triggering punishment exceeds any one-shot deviation gain. The folk-theorem-style multiplicity of SPE in multi-player bargaining games was established by Shaked (1985).

Remark (Why two players is special).

In two-person Rubinstein, the responder’s threat to reject must be “credible” in the sense of being subgame-perfect, which (by the Mertens-Shaked-Sutton uniqueness) ties the equilibrium to a single point. With three or more players, two responders can collectively threaten to reject in many internally-consistent ways, each generating a different sustainable proposer share. The breakdown of uniqueness from $n = 2$ to $n \geq 3$ is a quintessential illustration of how the equilibrium concept’s bite weakens once strategic interaction expands.

Problem Set 8

Problems

All questions below concern the following set-up. There is a single object for sale and N potential buyers are bidding for the object. Bidder i assigns a value of X_i to the object. Each X_i is iid on $[0, 1]$ according to the CDF F with density $f = F'$, so $\Pr(X_i \leq x) = F(x)$. Bidder i knows the realization x_i of X_i and only that other bidders' values are independently distributed according to F . The object has use value 0 to the seller. Unless otherwise mentioned, all bidders are risk-neutral.

Problem 8.1

Suppose $N = 2$ and $F(x) = x^k$, where $k \geq 1$ is an integer.

- (a) Find the expected payment $m^{II}(x)$ of a bidder with value x in a second-price auction (II). What is the expected selling price?
- (b) Find a symmetric equilibrium β in a first-price auction (I) and the associated expected payment $m^I(x)$ of a bidder with value x . What is the expected selling price?

Problem 8.2

Suppose that $N = 2$ and $F(x) = x$, that is, values are uniformly distributed. The object is sold using an all-pay auction (AP) in which each bidder pays his bid but only the bidder with the higher bid wins the object. Thus, if bidder 1 with value x bids $b_1 > b_2$, then bidder 1's payoff is $x - b_1$ and bidder 2's is $-b_2$.

- (a) Find a symmetric equilibrium in the all-pay auction.
- (b) What is the expected payment $m^{AP}(x)$ of a bidder with value x ? What is the expected selling price?

Problem 8.3

Suppose $F(x) = x$.

- (a) First, suppose the seller decides to sell the object using a second-price auction with $N = 3$. What is the seller's revenue?

- (b) Now suppose $N = 2$. The seller decides to sell using a second-price auction with reserve price $r \geq 0$. This means that if neither bid is above r , the seller keeps the object. Find a symmetric equilibrium bidding strategy for an arbitrary reserve price r . What is the seller's revenue as a function of r ? What is the optimal reserve price?
- (c) Which of the two situations (three bidders with no reserve vs. two bidders with optimal reserve) results in a higher expected revenue for the seller?

Solutions

Problem 8.1

The CDF of the highest competing value is $G(y) = F(y)^{N-1} = y^k$ (since $N = 2$), with density $g(y) = ky^{k-1}$.

(a) **Second-price auction.** Truth-telling is weakly dominant. The expected payment of a bidder with value x is

$$m^{\text{II}}(x) = \int_0^x y g(y) dy = k \int_0^x y^k dy = \frac{k}{k+1} x^{k+1}.$$

The expected selling price is

$$N \cdot \int_0^1 m^{\text{II}}(x) f(x) dx = 2 \int_0^1 \frac{k}{k+1} x^{k+1} \cdot kx^{k-1} dx = \frac{2k^2}{(k+1)(2k+1)}.$$

(b) **First-price auction.** By the revenue equivalence theorem, the FPA must yield the same expected payment for each value x :

$$m^{\text{I}}(x) = m^{\text{II}}(x) = \frac{k}{k+1} x^{k+1}.$$

The symmetric equilibrium bidding strategy follows from $\beta(x) \cdot G(x) = m^{\text{I}}(x)$:

$$\beta(x) = \frac{m^{\text{I}}(x)}{G(x)} = \frac{(k/(k+1))x^{k+1}}{x^k} = \frac{k}{k+1}x.$$

The linear shading factor $\frac{k}{k+1}$ ranges from $\frac{1}{2}$ at $k = 1$ (uniform values) to nearly 1 as k grows large; intuitively, when F concentrates mass near 1, competition tightens and bidders shade less.

The expected selling price is identical to the SPA: $2k^2/[(k+1)(2k+1)]$.

Problem 8.2

(a) **Symmetric equilibrium.** Bidder i 's expected payoff if she bids b when the other plays $\beta(\cdot)$:

$$\Pi_i(b; x_i) = x_i \cdot \Pr(\beta(X_j) < b) - b = x_i \cdot F(\beta^{-1}(b)) - b.$$

With $F(x) = x$, the FOC is

$$x_i \cdot \frac{1}{\beta'(\beta^{-1}(b))} = 1.$$

At a symmetric equilibrium $b = \beta(x_i)$, so $\beta^{-1}(b) = x_i$, yielding $\beta'(x) = x$. Integrating with $\beta(0) = 0$,

$$\beta(x) = \frac{x^2}{2}.$$

(b) **Expected payment and revenue.** A bidder with value x bids $\beta(x) = x^2/2$ and pays it regardless of outcome:

$$m^{\text{AP}}(x) = \beta(x) = \frac{x^2}{2}.$$

This matches $m^{\text{II}}(x) = x^2/2$ for the uniform case in Problem 8.1 (with $k = 1$), confirming revenue equivalence.

The expected selling price is

$$N \int_0^1 \beta(x) f(x) dx = 2 \int_0^1 \frac{x^2}{2} dx = \frac{1}{3},$$

again matching the SPA/FPA revenues.

Remark (Why all-pay matches second-price).

Despite the radically different incentives—losing bidders pay nothing in SPA, but pay their entire bid in AP—the revenue is identical because both auctions implement the same allocation rule (highest bidder wins) and award zero surplus to the lowest type. Revenue equivalence forces them onto the same expected-revenue surface. The behavioral predictions are very different though: bid *distributions* differ across auction formats, even when the expected revenues coincide.

Problem 8.3

(a) **SPA with $N = 3$, no reserve.** Expected revenue equals the expected second-order statistic from three uniform draws. The density of the second-order statistic is

$$f_{(2:3)}(y) = \frac{3!}{1! \cdot 1! \cdot 1!} F(y)^1 (1 - F(y))^1 f(y) = 6y(1 - y).$$

Hence

$$R = \int_0^1 y \cdot 6y(1 - y) dy = 6 \int_0^1 (y^2 - y^3) dy = 6 \cdot \frac{1}{12} = \frac{1}{2}.$$

(b) **SPA with $N = 2$, reserve r .** Truth-telling, with bids below r ignored: bidder i bids x_i if $x_i \geq r$, stays out otherwise.

Revenue by cases:

- Both $x_1, x_2 < r$: no sale. Probability r^2 .
- Exactly one $\geq r$: that bidder wins, pays r . Probability $2r(1 - r)$.
- Both $\geq r$: highest wins, pays second-highest. Probability $(1 - r)^2$. Conditional on both $\geq r$, the expected second-highest is $r + (1 - r)/3$.

$$R(r) = 2r^2(1 - r) + (1 - r)^2 \left(r + \frac{1-r}{3} \right) = \frac{(1 - r)[4r^2 + r + 1]}{3}.$$

Optimal reserve. Differentiating,

$$R'(r) = \frac{6r(1 - 2r)}{3} = 2r(1 - 2r),$$

giving $r^* = \frac{1}{2}$. Plugging in,

$$R(r^*) = \frac{(1/2)(4 \cdot 1/4 + 1/2 + 1)}{3} = \frac{5}{12}.$$

The optimal reserve $r^* = \frac{1}{2}$ satisfies the Myerson condition $r^* - (1 - F(r^*))/f(r^*) = 0$.

(c) **Three bidders without reserve vs. two bidders with optimal reserve.**

Setting	Expected revenue
$N = 3, r = 0$	$\frac{1}{2} = 0.500$
$N = 2, r^* = \frac{1}{2}$	$\frac{5}{12} \approx 0.417$

Three bidders without reserve generates higher expected revenue ($\frac{1}{2} > \frac{5}{12}$).

This illustrates a robust theme due to Bulow and Klemperer (1996): adding one more bidder, even without optimizing the reserve, dominates the gain from optimal reserve setting in the existing pool.

Remark (Bulow-Klemperer in one slogan).

“Recruit one more bidder, then run a basic English auction” beats “optimize the auction format with the bidders you already have.” The result holds far beyond the uniform-iid case considered here—it is one of the most policy-relevant results in market design, motivating procurement reforms that prioritize broadening participation over tweaking auction rules.

Problem Set 9

Problems

All questions below concern the following set-up. There is a single object for sale and two potential buyers are bidding for the object. Bidder i assigns a value of X_i to the object. Each X_i is independently distributed on $[0, 1]$ according to CDF F_i with density $f_i = F_i'$, so $\Pr(X_i \leq x) = F_i(x)$. Bidder i knows x_i and that the other bidder's value is independently distributed according to F_j ($j \neq i$). The object has use value 0 to the seller. All bidders are risk-neutral. The setting is asymmetric: $F_1 \neq F_2$ in general. Let $R^{\text{SPA}}(F_1, F_2)$ and $R^{\text{FPA}}(F_1, F_2)$ denote the equilibrium expected revenues in second- and first-price auctions, respectively.

Problem 9.1

Verify that asymmetry decreases revenues in a second-price auction by showing that

$$R^{\text{SPA}}(F_1, F_2) < R^{\text{SPA}}(\sqrt{F_1 F_2}, \sqrt{F_1 F_2}).$$

Problem 9.2

Suppose that bidder 1 and bidder 2's value distributions are

$$F_1(x) = \left(\frac{x-1}{a}\right)^a \text{ over } [1, a+1], \quad F_2(x) = \exp\left(\frac{-a}{a+1}x - a\right) \text{ over } [0, a+1],$$

respectively, where $a > 1$. (Notice that $F_2(0) > 0$ and so F_2 has a mass point at 0.)

- (a) Verify that the following constitute equilibrium bidding strategies in a first-price auction:

$$\beta_1(x) = x - 1, \quad \beta_2(x) = \frac{a}{a+1}x.$$

- (b) Show that when $a = 1$, $R^{\text{FPA}}(F_1, F_2) < R^{\text{SPA}}(F_1, F_2)$.
(c) Show that when $a = 2$, $R^{\text{FPA}}(F_1, F_2) > R^{\text{SPA}}(F_1, F_2)$.

Problem 9.3

Find the optimal selling mechanism (as in Myerson, 1981) when the value distributions of the two buyers are $F_1(x) = x$ and $F_2(x) = x^2$, both over $[0, 1]$.

Solutions

Problem 9.1

In a second-price auction, truth-telling is dominant, so revenue equals the expected second-highest value:

$$R^{\text{SPA}}(F_1, F_2) = \mathbb{E}[\min(X_1, X_2)] = \int_0^1 (1 - F_1(t))(1 - F_2(t)) dt.$$

For the symmetric environment with both bidders drawing from $G = \sqrt{F_1 F_2}$,

$$R^{\text{SPA}}(G, G) = \int_0^1 (1 - G(t))^2 dt = \int_0^1 (1 - \sqrt{F_1(t)F_2(t)})^2 dt.$$

Pointwise comparison at each t , with $a = F_1(t), b = F_2(t) \in [0, 1]$:

$$(1 - \sqrt{ab})^2 - (1 - a)(1 - b) = 1 - 2\sqrt{ab} + ab - 1 + a + b - ab = a + b - 2\sqrt{ab} = (\sqrt{a} - \sqrt{b})^2 \geq 0.$$

The inequality is strict whenever $\sqrt{F_1(t)} \neq \sqrt{F_2(t)}$. Since $F_1 \neq F_2$ implies $F_1(t) \neq F_2(t)$ on a set of positive measure, integrating gives the strict comparison

$$R^{\text{SPA}}(F_1, F_2) < R^{\text{SPA}}(G, G).$$

Remark (Asymmetry as a tax on SPA revenue).

In a symmetric SPA, the second-highest of two iid draws is the relevant statistic. Asymmetry shrinks this expectation: the weaker bidder is more often below average, but the resulting reduction in the second-highest—which the winner pays—dominates the small increase from the stronger bidder’s higher draws. The geometric-mean substitution $\sqrt{F_1 F_2}$ averages the two distributions in the “log” sense and recovers a higher revenue. The result foreshadows Problem 9.2: in a first-price auction, the comparison is more delicate, and R^{FPA} may even exceed R^{SPA} when asymmetry is sharp.

Problem 9.2

(a) Verifying equilibrium.

Bidder 1’s first-order condition. Bidder 1’s expected payoff from bidding b when value is x (with bidder 2 playing β_2):

$$\pi_1(b; x) = (x - b)F_2(\beta_2^{-1}(b)) = (x - b)F_2\left(\frac{a+1}{a}b\right).$$

Differentiating with respect to b and evaluating at $b = \beta_1(x) = x - 1$:

$$-F_2(\beta_2^{-1}(b)) + (x - b)f_2(\beta_2^{-1}(b)) \cdot \frac{a+1}{a} = 0 \implies \frac{f_2}{F_2} = \frac{a}{a+1}.$$

A constant hazard ratio $f_2/F_2 = a/(a+1)$ pins down F_2 as exponential: $F_2(z) \propto e^{(a/(a+1))z}$. With $F_2(a+1) = 1$, this matches the F_2 given, with mass e^{-a} at zero.

Bidder 2's first-order condition. Bidder 2's expected payoff:

$$\pi_2(b; x) = (x - b)F_1(\beta_1^{-1}(b)) = (x - b)(b/a)^a.$$

The FOC $-F_1 + (x - b)f_1 = 0$ gives

$$-(b/a)^a + (x - b)(b/a)^{a-1} = 0 \iff x = b \cdot \frac{a+1}{a},$$

i.e., $b = \frac{a}{a+1}x = \beta_2(x)$. Both first-order conditions confirm the proposed strategies.

(b) When $a = 1$: SPA dominates FPA. With $a = 1$, $F_1(x) = x - 1$ on $[1, 2]$ and $F_2(x) = e^{(x-2)/2}$ on $[0, 2]$. Equilibrium bids: $B_1 = X_1 - 1 \sim \text{Uniform}[0, 1]$, $B_2 = X_2/2$ with CDF $F_{B_2}(b) = e^{b-1}$ on $(0, 1]$ (mass e^{-1} at zero).

FPA revenue.

$$R^{\text{FPA}} = \mathbb{E}[\max(B_1, B_2)] = \int_0^1 (1 - F_{B_1}(t)F_{B_2}(t)) dt = \int_0^1 (1 - te^{t-1}) dt.$$

Computing $\int_0^1 te^{t-1} dt = e^{-1}$,

$$R^{\text{FPA}}(a = 1) = 1 - e^{-1} \approx 0.632.$$

SPA revenue.

$$R^{\text{SPA}} = \int_0^1 (1 - F_2(t)) dt + \int_1^2 (1 - F_1(t))(1 - F_2(t)) dt = -\frac{5}{2} + 4e^{-1/2} + 2e^{-1} \approx 0.662.$$

Hence $R^{\text{FPA}}(a = 1) < R^{\text{SPA}}(a = 1)$ by ≈ 0.030 .

(c) When $a = 2$: FPA dominates SPA. With $a = 2$, $B_1 \in [0, 2]$ has CDF $b^2/4$, $B_2 \in [0, 2]$ has CDF e^{b-2} on $(0, 2]$.

FPA revenue.

$$R^{\text{FPA}}(a = 2) = \int_0^2 (1 - (t^2/4)e^{t-2}) dt.$$

Computing $\int_0^2 t^2 e^{t-2} dt = 2 - 2e^{-2}$,

$$R^{\text{FPA}}(a = 2) = 2 - \frac{1}{4}(2 - 2e^{-2}) = \frac{3}{2} + \frac{1}{2}e^{-2} \approx 1.568.$$

SPA revenue. A similar (more tedious) computation gives $R^{\text{SPA}}(a = 2) \approx 1.529$. Hence $R^{\text{FPA}}(a = 2) > R^{\text{SPA}}(a = 2)$ by ≈ 0.039 .

Remark (Maskin-Riley: there is no universal ranking).

The two parameter values $a = 1$ and $a = 2$ flip the comparison: SPA dominates at moderate asymmetry but FPA dominates as asymmetry grows. The general principle (Maskin and Riley, 2000) is that in asymmetric environments, neither auction format is uniformly revenue-superior; the answer depends on the shape of the distributions. The intuition: in FPA the strong bidder shades more aggressively against the weaker rival, sometimes letting the weak bidder “win cheap”; in SPA the strong bidder always wins, but pays only the weaker bidder’s value. Whether shading or efficient allocation costs the seller more depends on the specific asymmetry.

Problem 9.3

Virtual values and reserve prices.

$$\psi_1(x) = x - \frac{1-x}{1} = 2x - 1, \quad \psi_2(x) = x - \frac{1-x^2}{2x} = \frac{3x^2 - 1}{2x}.$$

Both are strictly increasing on their domains, so the regularity condition holds. The individual reserve prices solve $\psi_i(r_i) = 0$:

$$r_1^* = \frac{1}{2}, \quad r_2^* = \frac{1}{\sqrt{3}} \approx 0.577.$$

Optimal allocation rule. The object goes to whichever bidder has the highest non-negative virtual value:

- If both $\psi_1(x_1), \psi_2(x_2) < 0$ ($x_1 < 1/2$ and $x_2 < 1/\sqrt{3}$): seller retains the object.
- Otherwise, allocate to bidder i with $\psi_i(x_i) = \max\{\psi_1(x_1), \psi_2(x_2)\}$.

Equating $\psi_1(x_1) = \psi_2(x_2)$ gives the boundary

$$3x_2^2 - 2(2x_1 - 1)x_2 - 1 = 0 \implies x_2 = \frac{(2x_1 - 1) + \sqrt{(2x_1 - 1)^2 + 3}}{3}.$$

Bidder 1 wins when x_2 falls below this curve (and $x_1 \geq 1/2$).

Payment rule. Each winner pays the smallest value at which they would still have won:

- If bidder 1 wins, payment is the threshold $y_1(x_2)$ with $\psi_1(y_1(x_2)) = \max\{0, \psi_2(x_2)\}$:

$$y_1(x_2) = \begin{cases} \frac{1}{2}, & x_2 < \frac{1}{\sqrt{3}}, \\ \frac{3x_2^2 + 2x_2 - 1}{4x_2}, & x_2 \geq \frac{1}{\sqrt{3}}. \end{cases}$$

- If bidder 2 wins, payment is $y_2(x_1)$ such that $\psi_2(y_2(x_1)) = \max\{0, \psi_1(x_1)\}$:

$$y_2(x_1) = \begin{cases} \frac{1}{\sqrt{3}}, & x_1 < \frac{1}{2}, \\ \frac{(2x_1 - 1) + \sqrt{(2x_1 - 1)^2 + 3}}{3}, & x_1 \geq \frac{1}{2}. \end{cases}$$

Economic interpretation. The stronger bidder ($F_2 = x^2$) faces a *higher* individual reserve ($1/\sqrt{3} \approx 0.577$) than the weaker bidder ($F_1 = x$, reserve $1/2$). This is the **hand-capping principle**: by raising the bar for the strong bidder, the seller forces the strong bidder to surrender more rent while keeping the weak bidder competitive.

The optimal mechanism is not always allocatively efficient. When $x_1 = 0.6$ and $x_2 = 0.55 < r_2^*$, virtual values are $\psi_1(0.6) = 0.2 > 0$ and $\psi_2(0.55) < 0$, so bidder 1 wins despite bidder 2 having a positive valuation—revenue-extracting price discrimination at the cost of social surplus.

Remark (Why the asymmetric reserves seem “backwards”).

Naively, one might expect the seller to set a high price for the weak bidder (extracting their entire willingness-to-pay) and a low price for the strong bidder (encouraging participation). The Myerson optimum does exactly the opposite. The reason: the seller's gain from extracting an extra dollar from the strong bidder is offset by the strong bidder demanding more information rent. The net marginal revenue from raising the strong bidder's reserve is the inverse hazard rate $(1 - F_2(x))/f_2(x)$, which equals $(1 - x^2)/(2x)$ here. Setting this equal to x gives $r_2^* = 1/\sqrt{3}$. The strong bidder's higher reserve reflects that their information rent grows faster than the weak bidder's.

Problem Set 10

Problems

Problem 10.1

There is a single object for sale and two risk-neutral bidders are bidding for the object. Bidder i assigns a value X_i to the object where each X_i is independently and uniformly distributed on $[0, 1]$. Bidder i knows the realization x_i and only that the other bidder's value is independently and uniformly distributed.

- (a) Suppose the object is being sold by means of a first-price auction. Find the symmetric equilibrium strategy β .
- (b) Now suppose that after the auction is over, both the losing and winning bids are publicly announced. In addition, there is the possibility of post-auction resale: the winner of the auction may, if he so wishes, offer the object to the other bidder at a fixed take-it-or-leave-it price p . If the other bidder agrees, the object changes hands and the loser pays the winner p . Otherwise, the object stays with the winner and no money changes hands. Show that β remains an equilibrium even if resale is allowed. In particular, show that a bidder with value x cannot gain by bidding an amount $b > \beta(x)$ even with the option to resell.

Problem 10.2

Two agents jointly own an indivisible object (say, a house) and each has an equal share. The value of the whole object to agent i is a random variable X_i which is independently and uniformly distributed on $[0, 1]$. Thus agent 1 derives a value $\frac{1}{2}X_1$ from his share and agent 2 derives $\frac{1}{2}X_2$ from hers. The two agents wish to dissolve their partnership; since the object cannot be subdivided, ownership goes to one of the two agents.

- (a) Consider the following “auction” for reallocating the object: both agents bid amounts b_1, b_2 , and if $b_i > b_j$, then i gets ownership and pays the other agent j the amount b_i . Find a symmetric equilibrium bidding strategy.
- (b) Is the procedure outlined above efficient? Is it individually rational?
- (c) Calculate each agent's payments in the VCG mechanism associated with this problem.

Problem 10.3

A community of N agents is considering whether or not to build a bridge which costs a known amount C to build. If the bridge is built, each agent derives a value X_i , where each X_i is drawn from a distribution F over $[0, 1]$. If the bridge is not built, each agent derives value 0. It is efficient to build the bridge iff $\sum X_i \geq C$.

- (a) Calculate the payments in the VCG mechanism associated with this problem.
- (b) Argue that the VCG mechanism always runs a deficit—that is, for every realization of values, the sum of payments from the agents is less than C .
- (c) Does there exist a mechanism that is efficient, incentive compatible, individually rational, and balances the budget? (By balancing the budget is meant that for every realization, the sum of payments equals C .)

Solutions

Problem 10.1

(a) **Symmetric equilibrium without resale.** Standard derivation: in the symmetric equilibrium, bidder i with value x bids the expected highest competing value conditional on winning,

$$\beta(x) = \mathbb{E}[Y_1 \mid Y_1 < x] = \frac{x}{2}.$$

(b) **The resale option does not break equilibrium.** We show that bidding $b > \beta(x) = x/2$ is not profitable even with the option to resell.

Suppose bidder 1 (with value x) bids $b \leq 1/2$ while bidder 2 plays β . Conditional on winning, $X_2 < 2b$.

Optimal resale price. Conditional on $X_2 \in [0, 2b]$, X_2 is uniform. Bidder 2 accepts a TIOLI offer of p iff $X_2 \geq p$. The deviator's payoff from resale at price p , accounting for keeping the object when rejected, is

$$V(p; b) = p \cdot \frac{2b - p}{2b} + x \cdot \frac{p}{2b} = x + \frac{(2b - p)(p - x)}{2b}.$$

FOC $2b - 2p + x = 0$ gives $p^* = b + x/2$. Substituting,

$$V(p^*; b) = x + \frac{(b - x/2)^2}{2b}.$$

Total expected payoff from bidding b . Win probability is $\Pr(X_2 < 2b) = 2b$, so

$$\pi(b; x) = 2b[-b + V(p^*; b)] = -2b^2 + 2bx + (b - x/2)^2 = -b^2 + bx + \frac{x^2}{4}.$$

Optimum. π is strictly concave in b (second derivative -2), with FOC

$$-2b + x = 0 \implies b^* = \frac{x}{2} = \beta(x).$$

The unique optimum—even with resale—is $\beta(x) = x/2$. Bidding $b > \beta(x)$ strictly decreases expected payoff.

Remark (Why resale fails to disrupt the equilibrium).

Resale gives the winner a way to extract additional surplus from a high-value loser. But the auction has already priced this option in: the loser's value is, on average, half the winner's, so any extra dollar paid in the auction to "win" a high-value loser is exactly cancelled by the resale terms. This is a cousin of the revenue equivalence principle. Resale does, however, change which auction format the seller prefers when bidders have asymmetric distributions.

Problem 10.2

(a) **Symmetric equilibrium of the bid auction.** Looking for symmetric strictly-increasing β , agent 1 with value x bidding b has expected payoff

$$E[\pi_1(b; x)] = \int_0^{\beta^{-1}(b)} \left(\frac{x}{2} - b\right) dt + \int_{\beta^{-1}(b)}^1 \left(\beta(t) - \frac{x}{2}\right) dt.$$

(First integral: winning, paying b , taking object value x ; second: losing, receiving $\beta(t)$, surrendering own half-share value $x/2$.)

Differentiating and using $b = \beta(x)$:

$$-x + \frac{x - 2\beta(x)}{\beta'(x)} = 0 \iff x\beta'(x) + 2\beta(x) = x.$$

Multiplying by x , the LHS becomes $\frac{d}{dx}[x^2\beta(x)]$:

$$\frac{d}{dx}[x^2\beta(x)] = x^2 \implies x^2\beta(x) = \frac{x^3}{3} + C.$$

With $\beta(0) = 0$, $C = 0$ and

$$\boxed{\beta(x) = \frac{x}{3}.}$$

(b) Efficiency and IR.

Efficiency. Since β is strictly increasing, the agent with the higher value submits the higher bid and wins. The object is allocated efficiently.

Individual rationality. Expected payoff for an agent with value x :

$$\begin{aligned} E[\text{payoff}] &= \int_0^x \left(\frac{x}{2} - \frac{x}{3}\right) dt + \int_x^1 \left(\frac{t}{3} - \frac{x}{2}\right) dt \\ &= \frac{x^2}{6} + \frac{1}{6} - \frac{x^2}{6} - \frac{x}{2} + \frac{x^2}{2} = \frac{1}{6} - \frac{x}{2} + \frac{x^2}{2}. \end{aligned}$$

The minimum over $x \in [0, 1]$ is at $x = \frac{1}{2}$, giving $\frac{1}{6} - \frac{1}{4} + \frac{1}{8} = \frac{1}{24} > 0$. So every type strictly improves over the $\frac{1}{2}X_i$ status quo.

(c) **VCG payments.** Apply the standard VCG (Clarke pivot) formulation:

$$M_i^* = W(\alpha_i, v_{-i}) - W_{-i}(v).$$

With $\alpha_i = 0$, $W(\alpha_i, v_{-i}) = v_j$ (the other agent monopolizes if i doesn't participate).

If i wins ($v_i > v_j$): $M_i^* = v_j - 0 = v_j$. Winner pays the loser's value to the planner.

If i loses ($v_i < v_j$): $M_i^* = v_j - v_j = 0$. Loser pays nothing.

Comparison to part (a). The bid auction has the winner paying $\beta(x) = x/3$ to the loser; payments stay within the partnership (budget-balanced). In Clarke-VCG, the winner's payment v_j goes to a third party, and total payments don't balance. Restoring balance via redistribution (sending the winner's payment to the loser) makes the mechanism IR but breaks the dominant-strategy property: the bid auction trades dominant-strategy IC for budget balance and Bayesian IC.

Problem 10.3

(a) **VCG payments.** The efficient allocation is $Q^*(x) = 1\{\sum_j x_j \geq C\}$. Define

$$W(x) = \left(\sum_j x_j - C\right)Q^*(x), \quad W_{-i}(x) = \left(\sum_{j \neq i} x_j - C\right)Q^*(x).$$

If the bridge is not built ($\sum x_j < C$). $Q^* = 0$; $M_i^* = 0$.

If the bridge is built ($\sum x_j \geq C$). Two sub-cases:

- *i non-pivotal* ($\sum_{j \neq i} x_j \geq C$): bridge would be built without i . $M_i^* = 0$.
- *i pivotal* ($\sum_{j \neq i} x_j < C$): without i , $W(0, x_{-i}) = 0$ and $W_{-i}(x) = \sum_{j \neq i} x_j - C < 0$. So

$$M_i^* = 0 - \left(\sum_{j \neq i} x_j - C\right) = C - \sum_{j \neq i} x_j.$$

In words: a pivotal agent pays the *shortfall* $C - \sum_{j \neq i} x_j$ that their participation eliminates; non-pivotal agents pay nothing.

(b) **VCG runs a deficit.** Suppose the bridge is built and let \mathcal{P} be the set of pivotal agents. Total payments are

$$\sum_{i \in \mathcal{P}} M_i^* = |\mathcal{P}| \cdot C - \sum_{i \in \mathcal{P}} \sum_{j \neq i} x_j.$$

Every agent pivotal: $|\mathcal{P}| = N$ and $\sum_{i=1}^N \sum_{j \neq i} x_j = (N-1) \sum x_j$. So

$$\sum_i M_i^* = NC - (N-1) \sum x_j = C - (N-1) \left(\sum x_j - C\right) \leq C,$$

with strict inequality when $\sum x_j > C$.

At least one non-pivotal: non-pivotal agents pay 0, so total \leq all-pivotal sum $\leq C$.

In both cases, $\sum_i M_i^* < C$ (strictly, generically), while the planner pays C . The mechanism runs an **ex-post deficit** whenever the bridge is built.

(c) **Existence of efficient, IC, IR, BB mechanism. No.** By the AGV theorem, such a mechanism exists iff VCG runs an ex-ante non-negative surplus, i.e., $\sum_i E[M_i^*] \geq C \cdot \Pr(\text{build})$. From part (b), $\sum_i M_i^* < C$ whenever the bridge is built. Taking expectations, $\sum_i E[M_i^*] < C \cdot \Pr(\text{build})$. By the AGV impossibility, no mechanism can be simultaneously efficient, IC, IR, and budget-balanced.

Remark (Why public goods are different from private auctions).

The bilateral-trade impossibility (Myerson-Satterthwaite) and the public-goods deficit are two faces of the same theme: when both sides of a transaction have private information, the gains from trade may be unattainable through any voluntary, balanced mechanism. In auctions, the seller subsumes the cost (or has zero cost), so VCG runs a surplus. In public goods, the planner subsumes the cost (C), so VCG runs a deficit. The remedy is typically to weaken IR (mandatory tax-funded provision), weaken efficiency (under-provision via voluntary contribution), or weaken IC (subsidy-financed provision with externally enforced participation).

Problem Set 11

Problems

Problem 11.1

Consider a two-sided matching problem between students in the set $S = \{s_1, s_2, s_3\}$ and colleges in the set $C = \{c_1, c_2, c_3\}$. The (strict) preferences of the individuals are:

$P(s_1)$	$P(s_2)$	$P(s_3)$	$P(c_1)$	$P(c_2)$	$P(c_3)$
c_1	c_1	c_1	s_1	s_1	s_1
c_2	c_2	c_3	s_2	s_3	s_2
c_3	c_3	c_2	s_3	s_2	s_3

- Find the stable matching from the Gale-Shapley algorithm when the students make the offers. Find the corresponding matching when the colleges make offers.
- Find all stable matchings.

Problem 11.2

Consider the problem of matching 4 medical students (labeled s_i) to 3 hospitals (labeled h_j). Hospital 1 has two openings; hospitals 2 and 3 have one opening each. Preferences:

$P(s_1)$	$P(s_2)$	$P(s_3)$	$P(s_4)$	$P(h_1)$	$P(h_2)$	$P(h_3)$
h_3	h_2	h_1	h_1	s_1	s_1	s_3
h_1	h_1	h_3	h_2	s_2	s_2	s_1
h_2	h_3	h_2	h_3	s_3	s_3	s_2
				s_4	s_4	s_4

- Find a stable match μ_S using the deferred acceptance algorithm in which the students make the offers.
- Find a stable match μ_H using the deferred acceptance algorithm in which the hospitals make the offers.

Solutions

Problem 11.1

(a) Student- and college-proposing Gale-Shapley outcomes.

Student-proposing DA.

- *Round 1.* Each student applies to the top: $s_1, s_2, s_3 \rightarrow c_1$. College c_1 holds the highest-ranked applicant (s_1 , since $P(c_1)$ ranks s_1 above s_2 above s_3) and rejects s_2, s_3 .
- *Round 2.* Rejected students apply to their next-best: $s_2 \rightarrow c_2$ and $s_3 \rightarrow c_3$ (since $P(s_3) = c_1, c_3, c_2$). Both c_2 and c_3 hold their unique applicants. No rejections; algorithm terminates.

$$\mu_S = \{(s_1, c_1), (s_2, c_2), (s_3, c_3)\}.$$

College-proposing DA.

- *Round 1.* Each college proposes to its top: $c_1, c_2, c_3 \rightarrow s_1$. Student s_1 keeps c_1 , rejects c_2, c_3 .
- *Round 2.* Rejected colleges propose next-best: $c_2 \rightarrow s_3$ ($P(c_2) = s_1, s_3, s_2$), $c_3 \rightarrow s_2$ ($P(c_3) = s_1, s_2, s_3$). s_3 holds c_2 ; s_2 holds c_3 . No further rejections; algorithm terminates.

$$\mu_C = \{(s_1, c_1), (s_2, c_3), (s_3, c_2)\}.$$

(b) The full set of stable matchings. There are $3! = 6$ candidate matchings. Checking each for blocking pairs:

1. μ_S — stable (verified above).
2. μ_C — stable (verified above).
3. $\{(s_1, c_2), (s_2, c_1), (s_3, c_3)\}$: (s_1, c_1) blocks (s_1 prefers c_1 to c_2 ; c_1 prefers s_1 to s_2).
4. $\{(s_1, c_2), (s_2, c_3), (s_3, c_1)\}$: (s_1, c_1) again blocks.
5. $\{(s_1, c_3), (s_2, c_1), (s_3, c_2)\}$: (s_1, c_1) blocks.
6. $\{(s_1, c_3), (s_2, c_2), (s_3, c_1)\}$: (s_1, c_1) blocks.

The set of stable matchings is exactly $\{\mu_S, \mu_C\}$. The lattice has two elements, with μ_S student-best and μ_C college-best.

Remark (A canonical example of student-college conflict).

In μ_S , student s_2 gets c_2 (her second choice), while in μ_C , s_2 gets c_3 (her third). The two matchings differ exactly on $\{s_2, s_3\}$, who swap colleges. The lone non-trivial degree of freedom is whether c_2 takes its second-favorite student s_3 (and c_3 takes s_2 , as in μ_C) or whether each goes to its own “natural” partner, with s_2 blocking c_2 ’s preferred match (as in μ_S).

Problem 11.2

(a) **Student-proposing DA.** Top-choice applications:

$$s_1 \rightarrow h_3, \quad s_2 \rightarrow h_2, \quad s_3 \rightarrow h_1, \quad s_4 \rightarrow h_1.$$

h_1 has two slots and receives applications $\{s_3, s_4\}$, both held. h_2 holds s_2 . h_3 holds s_1 . No rejections; algorithm terminates after round 1.

$$\mu_S = \{(s_1, h_3), (s_2, h_2), (s_3, h_1), (s_4, h_1)\}.$$

Every student receives their top choice.

(b) **Hospital-proposing DA.** Each hospital proposes for each open slot to its top-choice student.

Round 1. h_1 proposes to $\{s_1, s_2\}$ (top two), $h_2 \rightarrow s_1$, $h_3 \rightarrow s_3$.

- s_1 receives h_1 and h_2 ; with $P(s_1) = h_3, h_1, h_2$, holds h_1 , rejects h_2 .
- s_2 holds h_1 (only offer).
- s_3 holds h_3 .
- s_4 receives no offers.

Round 2. h_2 (rejected) proposes next on $P(h_2) = s_1, s_2, s_3, s_4$, namely s_2 .

- s_2 has h_1 and h_2 ; with $P(s_2) = h_2, h_1, h_3$, holds h_2 , rejects h_1 .

h_1 now has one slot open.

Round 3. h_1 proposes to s_3 (next on $P(h_1)$).

- s_3 has h_3 and h_1 ; with $P(s_3) = h_1, h_3, h_2$, holds h_1 , rejects h_3 .

Round 4. h_3 (rejected) proposes to s_1 (next on $P(h_3) = s_3, s_1, s_2, s_4$).

- s_1 has h_1 and h_3 ; with $P(s_1) = h_3, h_1, h_2$, holds h_3 , rejects h_1 .

h_1 has only s_3 .

Round 5. h_1 proposes to s_4 . s_4 holds h_1 . All slots filled; algorithm terminates.

$$\mu_H = \{(s_1, h_3), (s_2, h_2), (s_3, h_1), (s_4, h_1)\}.$$

Comparison. $\mu_H = \mu_S$. The two extremes of the stable-matching lattice coincide, so the stable matching is *unique*.

Remark (When does $\mu_S = \mu_H$?).

$\mu_S = \mu_H$ iff the set of stable matchings has a single element. A handy sufficient condition is that the student-optimal matching gives every student her first choice and is feasible—which happens here because each student’s top choice differs (modulo h_1 ’s extra capacity). When markets “fit nicely” like this, no student has any strategic motive to misreport, and even the (Roth, 1982) impossibility of two-sided strategy-proofness becomes irrelevant: there is no second matching to manipulate towards.

Problem Set 12

Problems

Problem 12.1

Consider the prisoners' dilemma game G :

	<i>C</i>	<i>D</i>
<i>C</i>	2, 2	0, 3
<i>D</i>	3, 0	1, 1

Suppose G is played infinitely often and both players use the same discount factor $\delta \in (0, 1)$. Consider the following “tit-for-tat” strategy: (i) in period $t = 1$, play C ; (ii) in all periods $t > 1$, play whatever the opponent played in period $t - 1$.

- (a) If both players play tit-for-tat, find all discount factors δ for which there is a subgame perfect equilibrium such that both players choose C in every period.

Problem 12.2

Consider the game G below:

	<i>H</i>	<i>M</i>	<i>L</i>
<i>H</i>	3, 3	1, 1	0, 5
<i>M</i>	1, 1	2, 2	0, 1
<i>L</i>	5, 0	1, 0	-1, -1

and let $G^\delta(\infty)$ denote the game where G is infinitely repeated and both players use a common discount factor $\delta \in (0, 1)$. There is perfect monitoring—all past choices are observed by both players.

- (a) Consider the following outcome paths:
- (i) $(H, H), (H, H), \dots$
 - (ii) $(M, M), (M, M), \dots$

Suppose the players play the “trigger” strategy: start with (i) and continue if there are no deviations. If there is any deviation from (i), play (ii) forever. For what values of δ do the trigger strategies constitute a subgame perfect equilibrium of $G^\delta(\infty)$?

(b) Now consider the following outcome paths:

(i) $(H, H), (H, H), \dots$

(ii) $(L, L), (H, H), (H, H), \dots$

Suppose the players play the “forgiving” strategy: start with (i) and continue if there are no deviations. If there is any deviation from (i), start (ii). If there is any deviation from (ii), restart (ii). For what values of δ do the forgiving strategies constitute a subgame perfect equilibrium?

Problem 12.3

Two firms, labeled 1 and 2, produce differentiated products and face the following demand functions

$$q_1 = a - p_1 + bp_2, \quad q_2 = a + bp_1 - p_2,$$

where $a > 0$ and $b \in (0, 1)$ are given. A decrease in a firm’s own price or an increase in its rival’s price causes its demand to increase. Suppose the cost of production is zero. Each firm chooses its price to maximize its profit.

- (a) Suppose the market lasts for only one period. What are the Nash equilibrium prices (p_1^N, p_2^N) and profits?
- (b) Suppose the two firms merged and became a monopoly. What prices (p_1^M, p_2^M) would the merged firm choose in each market?
- (c) Now suppose the market lasts indefinitely and both firms use a discount factor δ to evaluate future payoffs. Consider the following “trigger” strategy: (i) in $t = 1$ choose p_i^M ; (ii) in $t > 1$, choose p_i^M if in all previous periods both firms chose p_j^M ; otherwise choose the one-shot Nash equilibrium price p_i^N . For what values of δ can the firms collude perfectly and sustain monopoly prices using trigger strategies?

Solutions

Problem 12.1

The on-path outcome under both players following tit-for-tat is (C, C) in every period.

On-path one-shot deviation. Consider deviating to D in some period t , then reverting to tit-for-tat. The deviator earns 3 at period t and induces an alternating pattern thereafter: in $t + 1$ the opponent retaliates with D while the deviator (now back on tit-for-tat) plays C , yielding 0; in $t + 2$ the opponent plays C (matching deviator's C from $t + 1$) while the deviator plays D , yielding 3; and so on. The continuation stream is $3, 0, 3, 0, \dots$, so

$$V_{\text{dev}} = 3 + \delta \cdot 0 + \delta^2 \cdot 3 + \dots = \frac{3}{1 - \delta^2}.$$

Comparison to no-deviation $V_{\text{coop}} = 2/(1 - \delta)$ gives

$$\frac{2}{1 - \delta} \geq \frac{3}{1 - \delta^2} = \frac{3}{(1 - \delta)(1 + \delta)} \iff 2(1 + \delta) \geq 3 \iff \delta \geq \frac{1}{2}.$$

Off-path consideration (caveat). The above only verifies on-path Nash. For a complete SPE check, one must verify that tit-for-tat itself is optimal off-path. Tit-for-tat actually fails this stricter test for $\delta > 1/2$: in the subgame after a single (C, D) history, switching back to permanent C would restore mutual cooperation immediately and dominate the alternating pattern, giving $V_{\text{switch-back}} = 2/(1 - \delta) > 3/(1 - \delta^2)$ exactly when $\delta > 1/2$. Strict subgame perfection of tit-for-tat thus holds only at the knife-edge $\delta = 1/2$.

Answer. For sustaining (C, C) as a Nash equilibrium outcome via tit-for-tat, $\delta \geq 1/2$ is necessary and sufficient. Tit-for-tat is the canonical example of a strategy that achieves cooperation but is not strictly subgame perfect, motivating refinements like grim trigger.

Problem 12.2

Note that (M, M) is the only *symmetric* pure NE of the stage game. Each player's minmax payoff is 0, attained at L . The cooperative outcome (H, H) delivers $(3, 3)$, beating both (M, M) and the minmax.

(a) Trigger strategy. Best on-path deviation: against opponent playing H , the deviator's best is L , yielding 5.

No-deviation condition:

$$V_{\text{coop}} = \frac{3}{1 - \delta} \geq 5 + \delta \cdot \frac{2}{1 - \delta} = V_{\text{dev}}.$$

Multiplying by $(1 - \delta)$,

$$3 \geq 5 - 3\delta \iff \delta \geq \frac{2}{3}.$$

Off-path: the punishment phase plays (M, M) forever, a stage NE; no further deviation incentive. Trigger is SPE iff $\delta \geq 2/3$.

(b) Forgiving strategy. Best on-path deviation: deviator earns 5 at t , then suffers one round of (L, L) giving -1 , then back to (H, H) forever:

$$V_{\text{dev}} = 5 + \delta(-1) + \delta^2 \cdot \frac{3}{1 - \delta}.$$

Comparison:

$$\frac{3}{1-\delta} \geq 5 - \delta + \frac{3\delta^2}{1-\delta} \iff 3(1+\delta) \geq 5 - \delta \iff \delta \geq \frac{1}{2}.$$

Punishment-phase deviation: during (L, L) , deviating to either H or M (best response to L , both yielding 0) gives

$$V_{\text{punish-dev}} = 0 + \delta[-1 + \frac{3\delta}{1-\delta}]$$

since the punishment restarts. The no-deviation continuation is $V_{\text{punish}} = -1 + \delta \cdot \frac{3}{1-\delta}$. The condition $V_{\text{punish}} \geq V_{\text{punish-dev}}$ reduces to $\delta \geq 1/4$, implied by the on-path constraint.

Answer. The forgiving strategy supports (H, H) as a SPE iff $\delta \geq \frac{1}{2}$, more permissive than the trigger threshold.

Remark (Trigger vs. forgiving in one slogan).

Trigger: “forever (M, M) punishment, mild loss per period.” Forgiven: “one period of (L, L) then resume cooperation.” The forgiving strategy hits the deviator harder in a single period (loss 4 relative to coop) but releases the punishment, allowing the cooperative future to discipline today’s behavior more efficiently. The general lesson: **shorter, harsher punishments outperform longer, milder ones** when the goal is sustaining cooperation at low discount factors.

Problem 12.3

(a) **Simultaneous-move Nash equilibrium.** Each firm maximizes $\pi_i = p_i q_i = p_i(a - p_i + bp_j)$. FOC $a - 2p_i + bp_j = 0$ gives the symmetric NE

$$p_1^N = p_2^N = \frac{a}{2-b}, \quad q_i^N = \frac{a}{2-b}, \quad \pi_i^N = \frac{a^2}{(2-b)^2}.$$

(b) **Monopoly (merged firm) prices.** Total profit $\Pi = p_1 q_1 + p_2 q_2$. FOC w.r.t. p_1 :

$$\frac{\partial \Pi}{\partial p_1} = a - 2p_1 + 2bp_2 = 0,$$

so $p_1 = a/2 + bp_2$. By symmetry,

$$p_1^M = p_2^M = \frac{a}{2(1-b)}, \quad q_i^M = \frac{a}{2}, \quad \pi_i^M = \frac{a^2}{4(1-b)}.$$

Monopoly prices exceed Nash prices $(1/(2(1-b))) > 1/(2-b)$ for $b > 0$.

(c) **Trigger strategy collusion.** Best deviation from p_i^M : firm 1 maximizes $p_1(a - p_1 + bp_2^M)$ over p_1 . Optimum:

$$p_1^d = \frac{a(2-b)}{4(1-b)}, \quad q_1^d = p_1^d, \quad \pi_1^d = \frac{a^2(2-b)^2}{16(1-b)^2}.$$

No-deviation condition:

$$\frac{\pi^M}{1-\delta} \geq \pi^d + \frac{\delta\pi^N}{1-\delta} \iff \delta \geq \frac{\pi^d - \pi^M}{\pi^d - \pi^N}.$$

Computing:

$$\pi^d - \pi^M = \frac{a^2b^2}{16(1-b)^2}, \quad \pi^d - \pi^N = \frac{a^2b^2(b^2 - 8b + 8)}{16(1-b)^2(2-b)^2}.$$

Hence

$$\delta_{\min} = \frac{(2-b)^2}{b^2 - 8b + 8}.$$

Limit cases. At $b = 0$: $\delta_{\min} = 4/8 = 1/2$, matching the standard collusion threshold. At $b \rightarrow 1$ (perfect substitutes): $\delta_{\min} \rightarrow 1/1 = 1$.

Remark (Why differentiation eases collusion).

More differentiation (smaller b) reduces both the gain from undercutting (stealing the rival's market is harder) and the harshness of competitive punishment (Nash profits stay healthier). The first effect dominates: collusion becomes easier to sustain. Industries with strong product differentiation—branded consumer goods, pharmaceuticals—are empirically more conducive to tacit collusion than commodity industries, exactly because δ_{\min} is lower there.

Problem Set 13

Problems

Problem 13.1

Consider the game G :

	H	M	L
H	5, 5	0, 1	0, 6
M	1, 0	0, 0	1, 4
L	6, 0	4, 1	0, 0

Let $G(T)$ denote the game in which G is played T times and all past outcomes are observed. Players evaluate repeated game outcome paths by their average payoff over the T periods.

- (a) What are the (pure-strategy) minmax payoffs in G ?
- (b) Find all pure-strategy subgame perfect equilibria of $G(2)$.
- (c) Find all pure-strategy subgame perfect equilibria of $G(3)$.

Problem 13.2 (Tirole)

Two firms, labeled 1 and 2, produce identical products at a constant marginal cost $c > 0$. They face an uncertain market demand function

$$\tilde{D}(p) = \begin{cases} D(p) & \text{with probability } 1 - \eta, \\ 0 & \text{with probability } \eta, \end{cases}$$

where $D(p)$ is such that $D'(p) < 0$ (downward-sloping demand) and $(p - c)D(p)$ has a unique maximizer p^M (the monopoly price). The firms choose prices p_1 and p_2 . When market demand is $D(p)$ and (i) $p_i < p_j$, then firm i sells $D(p_i)$ and firm j sells nothing; (ii) when $p_i = p_j = p$, then each firm sells $\frac{1}{2}D(p)$. If realized market demand is 0, neither firm sells anything.

Suppose the firms compete over an infinite horizon and discount future payoffs using a common discount factor $\delta \in (0, 1)$. Firms do *not* observe their rivals' past or current prices; they only know their own past sales.

- (a) Consider the following strategies for the firms:
 1. Choose price p^M in the first period.

2. In all periods $t > 1$, choose $p_i^t = p^M$ if i 's sales have been $\frac{1}{2}D(p^M)$ in all past periods; otherwise choose $p_i^t = c$ (marginal cost) for k periods and then resume step 1.

(b) Under what circumstances do the strategies constitute a subgame perfect equilibrium?

Solutions

Problem 13.1

(a) Pure-strategy minmax payoffs.

For player 1 the minmax is $v_1 = \min_{s_2} \max_{s_1} u_1(s_1, s_2)$. Computing the column maxima:

$$\max_{s_1} u_1(s_1, H) = 6 \text{ at } L, \quad \max_{s_1} u_1(s_1, M) = 4 \text{ at } L, \quad \max_{s_1} u_1(s_1, L) = 1 \text{ at } M.$$

Taking the smallest, $v_1 = 1$, attained when player 2 plays L . The payoff structure is symmetric across the two players, so $v_2 = 1$ as well, attained when player 1 plays L .

The pure-strategy stage NE of G are

$$(M, L) \text{ with payoff } (1, 4) \quad \text{and} \quad (L, M) \text{ with payoff } (4, 1).$$

Each gives one player exactly the minmax and the other a strictly higher payoff. This payoff asymmetry is the structural fact that constrains the rest of the analysis: any reward to one player using a stage NE simultaneously punishes the other.

(b) Pure SPE of $G(2)$.

Period 2 has no continuation, so its prescription must be a stage NE. The only choices are (M, L) and (L, M) .

For period 1, write Δ_i for player i 's gain from a one-period deviation against the prescribed period-1 profile (s_1, s_2) . Since period 2 must already be a stage NE, the largest continuation differential available to deter player i is the gap between her best and worst stage-NE payoff: $4 - 1 = 3$. So sustainability of any non-NE period-1 prescription requires $\Delta_i \leq 3$ for both players.

The deviation gains (Δ_1, Δ_2) at every cell:

(Δ_1, Δ_2)	H	M	L
H	(1, 1)	(4, 5)	(1, 0)
M	(5, 4)	(4, 4)	(0, 0)
L	(0, 1)	(0, 0)	(1, 1)

Among the non-NE candidates, $\Delta_i \leq 3$ holds at $(H, H), (H, L), (L, H), (L, L)$. But the differential constraint is necessary, not sufficient: with only one period of stage-NE continuation and only two NE to choose from, we can deter at most one player at a time, because the NE that rewards player 1 is the same NE that punishes player 2.

- (H, L) : $\Delta_1 = 1, \Delta_2 = 0$. Only player 1 needs deterring. Use (L, M) on path (rewarding player 1) and (M, L) if player 1 deviates. The $4 - 1 = 3$ differential easily deters $\Delta_1 = 1$. *Sustainable.*
- (L, H) : symmetric. *Sustainable* with (M, L) on path.
- (H, H) : $\Delta_1 = \Delta_2 = 1$. To deter player 1 the period-2 prescription must move from (L, M) on path to (M, L) off path; to deter player 2 it must move in the opposite direction. No single on-path / off-path pair can do both. *Not sustainable.*
- (L, L) : same impossibility.

The two stage NE serve as period-1 prescriptions trivially. Combined with the two period-2 NE choices, this gives four NE-on-NE paths plus the two non-NE starts.

Six pure SPE outcome paths in $G(2)$:

$(M, L)(M, L)$	average (1, 4)
$(M, L)(L, M)$	average (2.5, 2.5)
$(L, M)(M, L)$	average (2.5, 2.5)
$(L, M)(L, M)$	average (4, 1)
$(H, L)(L, M)$	average (2, 3.5)
$(L, H)(M, L)$	average (3.5, 2)

(c) Pure SPE of $G(3)$.

Period 3 must be a stage NE. The structural change relative to (b) is that the period-1 prescription is now followed by a *two-period* continuation, which can be any pure SPE of $G(2)$. The cumulative two-period payoff pairs (R_1, R_2) available across the four NE-on-NE continuations are

$$(2, 8), (5, 5), (5, 5), (8, 2),$$

all summing to 10. The two balanced sequences $(L, M)(M, L)$ and $(M, L)(L, M)$ each give every player exactly 5.

For period-1 prescription (s_1, s_2) with deviation gains (Δ_1, Δ_2) , sustainability now requires $R_i^c - P_i^d \geq \Delta_i$ for both players, where R_i^c is on-path continuation and P_i^d is the continuation following i 's deviation. The minimum P_i^d achievable across the four continuations is 2.

What is newly sustainable as period-1 prescription:

- (H, H) : $\Delta_1 = \Delta_2 = 1$. Take the balanced continuation $(R_1, R_2) = (5, 5)$ on path; if player 1 deviates, switch to $(M, L)(M, L)$ giving $P_1^d = 2$; if player 2 deviates, switch to $(L, M)(L, M)$ giving $P_2^d = 2$. Each differential is $5 - 2 = 3 \geq 1$. *Sustainable.*
- (L, L) : same calculation. *Sustainable.*
- (M, M) : $\Delta_1 = \Delta_2 = 4$. Need $R_1 \geq P_1 + 4 = 6$ and $R_2 \geq P_2 + 4 = 6$, so $R_1 + R_2 \geq 12$. But $R_1 + R_2 = 10$. *Not sustainable.*
- Any cell with $\max(\Delta_1, \Delta_2) > 6$ violates the individual reward bound (since $R_i \leq 8$ and $P_i \geq 2$). This rules out $(M, H), (M, M), (H, M)$ and similar.

What about period 2? The period-2 subgame is itself a $G(2)$ subgame, so its sustainable prescriptions are exactly those from part (b). In particular, the period-2 prescription *cannot* be (H, H) or (L, L) , since those need a two-period continuation that the single period 3 cannot supply. So an attempt at, say, $(H, L), (H, H), (NE)$ fails at the period-2 step.

Summary of new outcomes in $G(3)$. Embedding the six $G(2)$ paths from part (b) into periods $\{2, 3\}$ and varying the period-1 prescription, the new period-1 outcomes are (H, H) and (L, L) , each paired with a balanced two-period continuation on path. The most natural new path is

$$(H, H), (L, M), (M, L) \quad \text{average payoff } (10/3, 10/3) \approx (3.33, 3.33),$$

with the symmetric variant $(H, H), (M, L), (L, M)$ giving the same average. These are strictly closer to the cooperative $(5, 5)$ than any path from $G(2)$. The mechanism is the Benoit-Krishna observation: when the stage game has multiple stage NE with payoff variation, finite-horizon SPE can sustain non-NE play in early periods by using *which* NE is played later as the reward/punishment instrument. A horizon of three is already enough to bring this mechanism into operation.

Problem 13.2

Setup. Let $\pi^M = \frac{1}{2}(p^M - c)D(p^M)$ denote each firm's per-period collusive profit when the market is alive. Let V be the discounted continuation value of being in the cooperative phase (about to price p^M this period and observe demand at the end of it), and let V_W be the value at the start of a k -period war.

Cooperative-phase value. A firm in the cooperative phase prices p^M . With probability $1 - \eta$ the market is alive, the firm sells $\frac{1}{2}D(p^M)$, earns π^M , and the cooperative phase continues into next period. With probability η the market is dead, sales are zero, the rival also sees zero sales and triggers the war. So

$$V = (1 - \eta)\pi^M + \delta[(1 - \eta)V + \eta V_W].$$

A war yields zero profit for k consecutive periods, after which cooperation resumes:

$$V_W = \delta^k V.$$

Substituting and solving for V ,

$$V = \frac{(1 - \eta)\pi^M}{1 - \delta(1 - \eta) - \eta\delta^{k+1}}.$$

Deviation incentive. A firm contemplating deviation chooses its price before observing the demand realization. Undercutting p^M slightly captures all market demand $D(p^M)$ when the market is alive (probability $1 - \eta$) and yields zero when dead (probability η). The expected current-period profit from undercutting is therefore

$$(1 - \eta)(p^M - c)D(p^M) = 2(1 - \eta)\pi^M.$$

After the deviation, the rival's sales are zero regardless of which state was realized, so the rival begins a war next period. The deviation continuation value is

$$V_{\text{dev}} = 2(1 - \eta)\pi^M + \delta V_W = 2(1 - \eta)\pi^M + \delta^{k+1}V.$$

No-deviation condition. Sustainability requires $V \geq V_{\text{dev}}$, equivalently

$$V(1 - \delta^{k+1}) \geq 2(1 - \eta)\pi^M.$$

Substituting the expression for V and canceling the common factor $(1 - \eta)\pi^M > 0$,

$$\frac{1 - \delta^{k+1}}{1 - \delta(1 - \eta) - \eta\delta^{k+1}} \geq 2.$$

Cross-multiplying and rearranging delivers the central condition:

$$2\delta(1 - \eta) - \delta^{k+1}(1 - 2\eta) \geq 1. \quad (\star)$$

Three sanity checks.

Perfect monitoring ($\eta = 0$). Condition (\star) reduces to $2\delta - \delta^{k+1} \geq 1$. As $k \rightarrow \infty$, $\delta^{k+1} \rightarrow 0$, so the threshold becomes $\delta \geq \frac{1}{2}$ —the standard threshold for sustaining Bertrand collusion under grim trigger.

Vanishing patience ($\delta \rightarrow 0$). Both sides go to zero, and the inequality fails. Without patience, the one-period deviation gain $2(1 - \eta)\pi^M$ overwhelms any future loss. Some patience is necessary.

Critical η . As $\eta \rightarrow \frac{1}{2}$ the coefficient $1 - 2\eta$ on δ^{k+1} vanishes, and (\star) reduces to $2\delta(1 - \eta) \geq 1$, i.e. $\delta \geq 1/[2(1 - \eta)]$, which approaches 1 as $\eta \rightarrow \frac{1}{2}$. For $\eta > \frac{1}{2}$ no patience is enough: the expected demand state is too often dead for cooperation to be sustainable at all.

Trade-off in k . For $\eta < \frac{1}{2}$, the LHS of (\star) is increasing in k (since δ^{k+1} shrinks and $1 - 2\eta > 0$), so longer punishments help deter deviation. But longer punishments also reduce the cooperative value V , because innocent demand collapses—which occur with probability η in every cooperative period—trigger correspondingly longer dead-weight punishment phases. The net result is an interior optimal $k^*(\eta, \delta)$ that balances deterrence against on-path waste. For η close to $\frac{1}{2}$, the optimal k becomes finite and the feasibility region in (δ, η, k) shrinks rapidly.

Interpretation. This is the Green-Porter (1984) model of collusion under imperfect monitoring. The defining feature: the equilibrium has price wars erupting *on the equilibrium path*, not because anyone has deviated, but because firms cannot distinguish a rival's deviation from a random demand collapse. A fraction η of cooperative phases inevitably degenerates into wars. Empirically, this matches Porter's case study of the U.S. railroad freight cartel of the 1880s, in which intermittent price wars erupted episodically without any documentary evidence of cheating. Compared to the perfect-monitoring case studied earlier, where the only inefficiency is the patience boundary, imperfect monitoring imposes an additional structural inefficiency: collusion must pay for noisy signals via real punishment rounds in equilibrium.

Part VII

Exams and Solutions

First Mid-Term Exam (2025)

Date: February 18, 2025. Four questions; answer all four.

Problems

Problem 1

Find a Nash equilibrium (pure or mixed) of the following incompletely specified game (where “?” denotes an unknown payoff):

	L	C	R
U	?, 1	3, 2	2, 0
M	?, 0	?, 3	?, ?
D	?, 3	5, 0	1, 6

Problem 2

Consider the following game G :

	A	B
A	3, 2	1, 1
B	4, 3	2, 4

- Find all pure-strategy Nash equilibria of G .
- Now consider a game Γ with the same payoffs as above but where players move sequentially: player 1 (rows) moves first, and player 2 (columns) moves after observing player 1's choice. Find all subgame perfect equilibria of Γ and compare with part (a).
- Now consider a game Γ_ε in which player 1 moves first, but her moves are only imperfectly observed by player 2. When player 1 chooses A , player 2 receives a “signal” a with probability $1 - \varepsilon$ and signal b with probability ε , where $0 < \varepsilon < 1/4$. When player 1 chooses B , player 2 receives signal a with probability ε and b with probability $1 - \varepsilon$.
 - Draw the extensive form associated with this game.
 - Show that there is a Nash equilibrium of the game with imperfectly observed actions in which both players choose B .

Problem 3

Let $G = (S_i, u_i)_{i=1}^2$ be a finite two-player game in strategic form. Define Γ to be an extensive-form game of perfect information where first player 1 chooses $s_1 \in S_1$, player 2 is informed of player 1's choice, and then chooses $s_2 \in S_2$. For any pair (s_1, s_2) , payoffs in Γ equal payoffs in G .

- (a) Show: for every pure-strategy Nash equilibrium of G there is a pure-strategy Nash equilibrium of Γ in which the players make the same choices as in the equilibrium of G .
- (b) Construct an example to show that the statement in (a) does not hold for mixed-strategy equilibria.
- (c) Construct an example to show that the statement in (a) does not hold for pure-strategy subgame perfect equilibria.

Problem 4

Consider the following model of an arms race between two countries. Each country i selects a level of “military capability” $x_i \in [0, 1]$. The payoff to country i when the two countries choose capabilities (x_1, x_2) is

$$u_i(x_1, x_2) = g(x_i - x_j) - c(x_i), \quad j \neq i,$$

where $g : [-1, 1] \rightarrow \mathbb{R}$ is strictly increasing ($g' > 0$) and strictly concave ($g'' < 0$), and $c : [0, 1] \rightarrow \mathbb{R}$ is strictly increasing ($c' > 0$) and strictly convex ($c'' > 0$). The function g represents the “gain” to a country from having a military advantage over the other; c represents the cost of acquiring military capability.

- (a) Show that the game described is supermodular.
- (b) Show that the game has a symmetric pure-strategy equilibrium (x^*, x^*) .
- (c) Can this game have multiple symmetric pure-strategy equilibria?
- (d) Can this game have asymmetric pure-strategy equilibria (x_1^*, x_2^*) where $x_1^* \neq x_2^*$?

Solutions

Problem 1

Six entries in the matrix are unknown (five for player 1, one for player 2 at (M, R)). We seek a NE that does not depend on these unknowns.

Restrict to the $\{U, D\} \times \{C, R\}$ subgame. The known sub-matrix is

	C	R
U	3, 2	2, 0
D	5, 0	1, 6

Looking for a mixed NE: let $p = \Pr(U)$ and $q = \Pr(C)$. Player 2's indifference between C and R :

$$u_2(C | p) = 2p, \quad u_2(R | p) = 6(1 - p) \implies p = \frac{3}{4}.$$

Player 1's indifference between U and D :

$$u_1(U | q) = 3q + 2(1 - q) = 2 + q, \quad u_1(D | q) = 5q + (1 - q) = 1 + 4q \implies q = \frac{1}{3}.$$

The mixed sub-game NE is $\sigma_1 = (\frac{3}{4}, \frac{1}{4})$ over (U, D) and $\sigma_2 = (0, \frac{1}{3}, \frac{2}{3})$ over (L, C, R) , with payoff $(\frac{7}{3}, \frac{3}{2})$.

Why this works in the full game. At $p = \frac{3}{4}$, all three of player 2's pure-strategy expected payoffs coincide:

$$u_2(L) = \frac{3}{4}(1) + \frac{1}{4}(3) = \frac{3}{2}, \quad u_2(C) = 2 \cdot \frac{3}{4} = \frac{3}{2}, \quad u_2(R) = 6 \cdot \frac{1}{4} = \frac{3}{2}.$$

So putting weight 0 on L is optimal for player 2. The unknown payoffs in column L for player 1 do not enter because $q_L = 0$.

Remaining condition. Strategy M must not be a profitable deviation for player 1:

$$u_1(M | \sigma_2) = \frac{1}{3}u_1(M, C) + \frac{2}{3}u_1(M, R) \leq \frac{7}{3}.$$

A constraint on the unknowns, not a contradiction.

$$\text{NE: } \sigma_1 = (\frac{3}{4}, 0, \frac{1}{4}), \sigma_2 = (0, \frac{1}{3}, \frac{2}{3}), \text{ with expected payoffs } (\frac{7}{3}, \frac{3}{2}).$$

Remark (Why this NE is robust to most unknowns).

By choosing $q_L = 0$, the equilibrium “ignores” column L entirely, so player 1's payoffs in that column play no role. The single remaining vulnerability is row M , whose payoffs against the equilibrium mix must not exceed $7/3$. The exam likely intends for the student to recognize this structure—an “unknown-robust” NE in a strict subset of the strategy space.

Problem 2

(a) **Pure NE of G .** Cell-by-cell:

- (A, A) : player 1 prefers B (4 vs 3). Not NE.

- (A, B) : player 1 prefers B (2 vs 1). Not NE.
- (B, A) : player 2 prefers B (4 vs 3). Not NE.
- (B, B) : player 1 BR is B (2 vs 1) ✓; player 2 BR is B (4 vs 3) ✓. **NE**.

Unique pure NE: (B, B) with payoff (2, 4).

(b) **Sequential game Γ (player 1 commits, player 2 observes)**. Backward induction:

- After A : player 2 chooses A ($u_2 = 2 > 1$). Outcome $(A, A) \rightarrow (3, 2)$.
- After B : player 2 chooses B ($u_2 = 4 > 3$). Outcome $(B, B) \rightarrow (2, 4)$.


Player 1 chooses A ($u_1 = 3 > 2$). The unique SPE is

$$\sigma_1^\Gamma = A, \quad \sigma_2^\Gamma : A \text{ if } A, B \text{ if } B,$$

with on-path outcome (A, A) and payoff (3, 2).

Comparison. The unique NE of G delivers (2, 4); the unique SPE of Γ delivers (3, 2). The sequential structure benefits player 1 (commitment value: $3 > 2$) and hurts player 2 ($2 < 4$). Subgame perfection in Γ rules out the threat “always play B ” that sustains (B, B) in the simultaneous game.

(c.1) **Extensive form of Γ_ϵ .**



The dashed lines mark player 2's two information sets.

(c.2) **NE in which both players choose B .**

- Player 1's strategy: B .
- Player 2's strategy: B at every information set.

Verification.

- Player 1: anticipating B at both info sets, A yields $u_1(A, B) = 1$ and B yields $u_1(B, B) = 2$. Prefers B .
- Player 2: given player 1 plays B , the on-path posterior at every info set is $\Pr(B \mid \text{signal}) = 1$, and BR is $u_2(B, B) = 4 > 3 = u_2(B, A)$, so B .

Remark (Why the imperfection helps).

In the perfect-information Γ , (B, B) is not subgame-perfect because if player 1 plays A , player 2 prefers A . With noisy signals, however, player 2's best-response strategy can no longer condition on the action directly. Even after observing signal a , player 2 cannot rule out that player 1 played B , so a strategy of “always B ” can be self-confirming. The $\varepsilon < 1/4$ assumption ensures the noise is not too large.

Problem 3

(a) **Every pure NE of G corresponds to a pure NE of Γ .** Let (s_1^*, s_2^*) be a pure NE of G . Construct strategies in Γ :

$$\sigma_1^\Gamma = s_1^*, \quad \sigma_2^\Gamma(\cdot) = s_2^* \text{ at every information set.}$$

Verification.

- Player 1: anticipating s_2^* regardless, on-path payoff is $u_1(s_1^*, s_2^*)$. Deviation $s_1 \neq s_1^*$ gives $u_1(s_1, s_2^*) \leq u_1(s_1^*, s_2^*)$ by the NE property of G .
- Player 2: on-path the only info set reached is $\{s_1^*\}$; at that info set, BR given $s_1 = s_1^*$ is s_2^* . Off-path optimality is not required for NE.

(b) **Counterexample for mixed equilibria.** Matching pennies-style game:

	L	R
U	1, 0	0, 1
D	0, 1	1, 0

No pure NE; unique mixed NE has both players randomizing $(\frac{1}{2}, \frac{1}{2})$ with expected payoff $(\frac{1}{2}, \frac{1}{2})$.

In Γ , backward induction:

- After U : player 2 plays R . Outcome $(U, R) \rightarrow (0, 1)$.
- After D : player 2 plays L . Outcome $(D, L) \rightarrow (0, 1)$.

Player 1 is indifferent (both give 0). Any pure NE outcome of Γ is (U, R) or (D, L) with payoff $(0, 1)$. The mixed NE expected payoff $(\frac{1}{2}, \frac{1}{2})$ is not realized in any pure NE of Γ .

(c) **Counterexample for SPE.** Use the game from Problem 2: G has unique pure NE (B, B) with payoff $(2, 4)$. The unique pure SPE of Γ has on-path outcome (A, A) with payoff $(3, 2)$. So (B, B) does not arise as the on-path outcome of any pure SPE of Γ .

(In words: (B, B) is a pure NE of Γ via the construction in (a), but it fails subgame perfection because at the off-path info set after A , B is not a best response.)

Problem 4

(a) **Supermodularity.** Cross-partial:

$$\frac{\partial u_i}{\partial x_i} = g'(x_i - x_j) - c'(x_i), \quad \frac{\partial^2 u_i}{\partial x_i \partial x_j} = -g''(x_i - x_j).$$

Since $g'' < 0$, the cross-partial is strictly positive: **strategic complementarities**. Combined with the lattice strategy space $[0, 1]^2$, the game is supermodular.

(b) Symmetric pure-strategy equilibrium. At a symmetric (x, x) , the FOC reduces to

$$g'(0) = c'(x).$$

Since c' is strictly increasing and continuous, this equation has at most one interior solution. Provided $c'(0) \leq g'(0) \leq c'(1)$, an interior $x^* \in (0, 1)$ exists; otherwise the equilibrium is at a corner ($x^* = 0$ if $g'(0) < c'(0)$, $x^* = 1$ if $g'(0) > c'(1)$). Existence follows from Tarski's fixed-point theorem.

(c) Multiple symmetric equilibria? No. The interior FOC $g'(0) = c'(x)$ has a unique solution by strict monotonicity of c' . With at most two corners, generically only one of the three candidates is consistent with the best-response conditions.

(d) Asymmetric pure-strategy equilibria? No. The best-response function $\text{BR}(x_j)$ is implicitly defined by

$$g'(\text{BR}(x_j) - x_j) = c'(\text{BR}(x_j)).$$

Differentiating implicitly,

$$\text{BR}'(x_j) = \frac{g''}{g'' - c''} \in (0, 1),$$

since both numerator and denominator are negative; their ratio is positive but less than 1 because $|c''| > 0$.

The composition $\text{BR} \circ \text{BR}$ is therefore a strict contraction with derivative bounded by $(\text{BR}')^2 < 1$. The unique fixed point lies on the diagonal $x_1 = x_2$. Hence no asymmetric pure-strategy equilibria.

Remark (The role of strict concavity/convexity).

Supermodularity alone allows multiple equilibria in general supermodular games. What kills multiplicity here is the contraction property: $|\text{BR}'| < 1$ strictly, which follows from *strict* concavity of g and *strict* convexity of c . If g were merely weakly concave (e.g., piecewise linear), best-response curves could parallel, generating a continuum of NE. Strict curvature pins down the symmetric NE uniquely in this arms-race model.

Second Mid-Term Exam (2025)

Date: April 3, 2025. Three questions; answer all three.

Problems

Problem 1

Carefully define the following terms:

- (a) Nash's axiom of independence of irrelevant alternatives (IIA) for bargaining problems.
- (b) The revenue equivalence principle of auctions.
- (c) A stable match in a situation in which there is one-to-one matching between students and colleges.

Problem 2

Consider the Rubinstein infinite-horizon alternating-offer bargaining model but with the following preferences. If an agreement (z, t) is reached (a split of $(z, 1 - z)$ at time t), then the payoff to player 1 is $z - (t - 1)c_1$ and the payoff to player 2 is $(1 - z) - (t - 1)c_2$, where c_i is player i 's constant per-period cost of delay and satisfies $0 < c_i < 1$. Note: there is no discounting; rather there is a constant per-period cost of delay. Let G_i denote the game in which player i makes the first offer.

- (a) Suppose $c_1 = c_2 = c$. Show that if $x, y \in [0, 1]$ satisfy $x = y + c$, then there is a subgame perfect equilibrium of G_1 in which a split of $(x, 1 - x)$ is achieved with no delay, and a subgame perfect equilibrium of G_2 in which a split of $(y, 1 - y)$ is achieved with no delay. Write down the equilibrium strategies carefully.
- (b) Suppose $c_1 < c_2$. Show that there is a subgame perfect equilibrium of G_1 in which a split of $(1, 0)$ is achieved without delay, and a subgame perfect equilibrium of G_2 in which a split of $(1 - c_1, c_1)$ is achieved without delay. Write down the equilibrium strategies carefully.

Problem 3

There is a single object for sale and two potential buyers are bidding for the object. Bidder i assigns a value of X_i to the object. Each X_i is iid uniform on $[0, 1]$. Bidder i knows x_i

and only that the other bidder's value is iid uniform on $[0, 1]$. The object has use value 0 to the seller.

- (a) Suppose the seller uses the following “hybrid” auction—a mixture of first-price and second-price. Each bidder i submits a sealed bid b_i . If $b_i > b_j$, bidder i wins and pays $\alpha b_i + (1 - \alpha)b_j$, where α is fixed and known with $0 \leq \alpha \leq 1$. Show that it is an equilibrium for each bidder to follow the strategy

$$\beta(x) = \frac{x}{1 + \alpha}.$$

- (b) Does this hybrid auction satisfy the assumptions of the revenue equivalence principle?
- (c) Can the seller benefit by setting a non-zero reserve price in the hybrid auction? What is the optimal reserve price?

Solutions

Problem 1

(a) Independence of Irrelevant Alternatives (IIA). A bargaining solution F satisfies IIA if, for any two bargaining problems (U, d) and (\bar{U}, d) with $U \subseteq \bar{U}$,

$$F(\bar{U}, d) \in U \implies F(U, d) = F(\bar{U}, d).$$

If the solution chosen on the larger feasible set \bar{U} happens to lie in the smaller set U , then it is also the solution on U . Removing alternatives that were not chosen does not change the chosen point.

(b) Revenue equivalence principle. For any two incentive-compatible auction mechanisms with independent private values, if (i) both mechanisms induce the same interim allocation rule $q_i(x)$, and (ii) both assign the same expected payment to the lowest type ($m_i(0)$, typically zero), then they yield the same expected payment $m_i(x)$ for every type, and the same expected revenue. The result follows from the envelope theorem: the interim utility of a bidder of type x is uniquely determined by the allocation rule plus the boundary condition $u_i(0) = -m_i(0)$.

(c) Stable match. In a one-to-one matching market between students \mathcal{S} and colleges \mathcal{C} with strict preferences, a matching $\mu : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$ is **stable** if there is no *blocking pair*: no $(s, c) \in \mathcal{S} \times \mathcal{C}$ such that

$$c \succ_s \mu(s) \quad \text{and} \quad s \succ_c \mu^{-1}(c).$$

No student-college pair would each prefer to abandon their current match in favor of one another.

Problem 2

(a) Common cost $c_1 = c_2 = c$.

Equilibrium strategies in G_1 .

- Player 1: when proposing, demand x . When responding to player 2's offer z' (player 1's share), accept iff $z' \geq y$.
- Player 2: when proposing, offer $(y, 1 - y)$ to player 1. When responding to player 1's offer (1's share z'), accept iff $z' \leq x$.

Symmetric structure for G_2 .

Verification. At any period of G_1 :

- Player 1 demanding x : player 2's continuation upon rejection is G_2 with one-period delay, getting $1 - y$ at delay cost c , payoff $(1 - y) - c$. Player 2 accepts $1 - x$ iff

$$1 - x \geq (1 - y) - c \iff x \leq y + c.$$

At equality $x = y + c$, player 2 is indifferent and accepts (standard tie-breaking).

- Player 2's offer $(y, 1 - y)$ in G_2 : player 1's continuation is G_1 with delay, getting $x - c = y$. Player 1 accepts y .

Remark (Multiple SPEs under linear cost).

The system $x = y + c$ does not uniquely pin down (x, y) : any $x \in [c, 1]$ paired with $y = x - c \in [0, 1 - c]$ works. Linear cost of delay produces a *continuum* of SPE outcomes, in stark contrast to geometric discounting (Rubinstein), which gives a unique split. With linear cost, each player's continuation value is a constant translate of the proposer's value, admitting any pair (x, y) with $x - y = c$. Geometric discounting introduces an exogenous shrinkage that pins the equilibrium down via two coupled equations.

(b) Asymmetric costs $c_1 < c_2$.

Equilibrium strategies in G_1 .

- Player 1: when proposing, demand 1. When responding to player 2's offer z' , accept iff $z' \geq 1 - c_1$.
- Player 2: when proposing, offer $(1 - c_1, c_1)$. When responding, accept any demand by player 1.

Why it works.

- Player 2 accepts 0 in G_1 . Rejecting yields G_2 continuation: player 2 proposes next period and gets c_1 at one-period delay, payoff $c_1 - c_2 < 0$. Accepting 0 today gives $0 > c_1 - c_2$. Accept.
- Player 1 demands 1. Player 2 accepts any demand ≤ 1 , so player 1 demands the entire pie.
- In G_2 , player 1 accepts $1 - c_1$. Rejecting and entering G_1 next period gives $1 - c_1$. Indifference threshold; accept.
- In G_2 , player 2 offers $1 - c_1$. Player 1's minimum acceptable; player 2 keeps c_1 .

Intuition. Player 2 is more impatient (higher per-period cost). The threat to “walk away” is incredible for player 2 since each delay period costs $c_2 > c_1$. This is a discrete analog of the “patience captures the surplus” principle in geometric Rubinstein.

Problem 3

(a) Equilibrium $\beta(x) = x/(1 + \alpha)$. Suppose the opponent uses β . Bidder i with value x bidding b wins iff $\beta(X_j) < b$, equivalently $X_j < (1 + \alpha)b$. Conditional on winning, $X_j \sim U[0, (1 + \alpha)b]$, so $\beta(X_j) \sim U[0, b]$ with mean $b/2$.

Expected winning payment:

$$\alpha b + (1 - \alpha) \cdot \frac{b}{2} = \frac{(1 + \alpha)b}{2}.$$

Expected payoff:

$$\pi(b; x) = (1 + \alpha)b \cdot \left(x - \frac{(1 + \alpha)b}{2}\right).$$

FOC:

$$(1 + \alpha)\left(x - \frac{(1+\alpha)b}{2}\right) - (1 + \alpha)b \cdot \frac{1+\alpha}{2} = 0 \iff x = (1 + \alpha)b \iff b = \frac{x}{1 + \alpha}.$$

Hence $\beta(x) = x/(1 + \alpha)$ is the unique symmetric equilibrium.

(b) Revenue equivalence. Yes. The two ingredients hold:

1. *Same allocation rule.* In the hybrid, the highest bidder wins, and since β is strictly increasing, the highest *value* also wins.
2. *Same payment at the lowest type.* A bidder with $x = 0$ bids 0 and pays 0 in expectation.

Hence the hybrid yields the same expected payment per type as FPA and SPA: $m^{\text{hybrid}}(x) = x^2/2$.

Verification. The expected payment of a winning type- x bidder is $(1 + \alpha)\beta(x)/2 = x/2$; win probability is x ; expected payment $x \cdot x/2 = x^2/2$. ✓

(c) Optimal reserve. By revenue equivalence, the optimal reserve in the hybrid equals the optimal reserve in the SPA, which is the root of the virtual value:

$$\varphi(x) = 2x - 1, \quad \varphi(r^*) = 0 \implies r^* = \frac{1}{2}.$$

$$\boxed{r^* = \frac{1}{2}}, \quad \text{maximum revenue } 5/12.$$

The optimal reserve is independent of α .

Remark (The reserve is format-free).

A striking corollary of revenue equivalence: the optimal reserve is determined entirely by the bidder's type distribution (via the virtual value), not by the auction format. Whether the seller runs FPA, SPA, all-pay, or hybrid, the same threshold $r^* = (1 - F(r^*))/f(r^*)$ extracts the most revenue. Format choice affects the bidder's effort and risk profile, but the optimal threshold sits at the intersection of demand-side incentives and the seller's marginal revenue curve.

Final Exam (2025)

Spring 2025. Three questions; answer all three.

Problems

Problem 1

Consider a two-sided matching problem between students $S = \{s_1, s_2, s_3\}$ and colleges $C = \{c_1, c_2, c_3\}$. Each college has only one slot. The (strict) preferences are:

$P(s_1)$	$P(s_2)$	$P(s_3)$	$P(c_1)$	$P(c_2)$	$P(c_3)$
c_1	c_1	c_1	s_1	s_1	s_1
c_2	c_2	c_3	s_2	s_3	s_2
c_3	\emptyset	c_2	s_3	s_2	\emptyset

where \emptyset denotes that all other matches are unacceptable. Thus s_2 deems c_3 unacceptable and c_3 deems s_3 unacceptable.

- Find the stable matching from the Gale-Shapley algorithm when students make offers.
- Find the corresponding matching when colleges make admission offers.
- Are there any other stable matchings?

Problem 2

Consider the game G :

	H	M	L
H	3, 3	1, 1	0, 5
M	1, 1	2, 2	0, 1
L	5, 0	1, 0	-1, -1

Let $G^\delta(\infty)$ denote the game where G is infinitely repeated with common discount factor $\delta \in (0, 1)$. There is perfect monitoring.

- Consider outcome paths (i) $(H, H), (H, H), \dots$ and (ii) $(M, M), (M, M), \dots$. Players play the “trigger” strategy: start with (i), continue if no deviations; if any deviation from (i), play (ii) forever. For what values of δ do the trigger strategies constitute a subgame perfect equilibrium of $G^\delta(\infty)$?

- (b) Now consider paths (i) $(H, H), (H, H), \dots$ and (ii) $(L, L), (H, H), (H, H), \dots$. Players play the “forgiving” strategy: start with (i), continue if no deviations; if any deviation from (i), start (ii); if any deviation from (ii), restart (ii). For what values of δ do the forgiving strategies constitute a subgame perfect equilibrium?

Problem 3

Consider the signaling game depicted below. Player 1 can be of two types: strong (S) with probability p and weak (W) with probability $1 - p$. Knowing his own type, player 1 takes one of two actions: he can give a gift to player 2 (G) or not (N). If N , the game ends. If G , player 2 takes one of two actions: accept (A) or reject (R). Player 2 does not know player 1’s type. Payoffs (player 1, player 2):

Type	N	(G, A)	(G, R)
S	0, 0	1, 1	-1, 0
W	0, 0	1, -1	-1, 0

- (a) Show that the game has a pooling perfect Bayesian equilibrium (PBE) in which both types of player 1 choose N . Carefully specify the strategies and beliefs.
- (b) Show that the game has a pooling PBE in which both types of player 1 choose G if and only if $p \geq 1/2$.
- (c) Does the game have a PBE which is separating? Does it have a Nash equilibrium which is separating?

Solutions

Problem 1

(a) Student-proposing DA.

Round 1. $s_1, s_2, s_3 \rightarrow c_1$. College c_1 holds s_1 and rejects s_2, s_3 .

Round 2. $s_2 \rightarrow c_2, s_3 \rightarrow c_3$.

- c_2 holds s_2 .
- c_3 rejects s_3 (unacceptable).

Round 3. $s_3 \rightarrow c_2$ (third on $P(s_3)$). c_2 has both s_2 and s_3 ; $P(c_2) = s_1, s_3, s_2$, so holds s_3 , rejects s_2 .

Round 4. s_2 has no further acceptable colleges. s_2 unmatched.

$$\mu_S = \{(s_1, c_1), (s_3, c_2)\}, \quad s_2 \text{ and } c_3 \text{ unmatched.}$$

(b) College-proposing DA.

Round 1. $c_1, c_2, c_3 \rightarrow s_1$. s_1 holds c_1 , rejects c_2, c_3 .

Round 2. $c_2 \rightarrow s_3, c_3 \rightarrow s_2$.

- s_3 holds c_2 (acceptable).
- s_2 rejects c_3 (unacceptable).

Round 3. c_3 has no further acceptable students. c_3 unmatched.

$$\mu_C = \{(s_1, c_1), (s_3, c_2)\}, \quad s_2 \text{ and } c_3 \text{ unmatched.}$$

The two procedures yield the same matching.

(c) Other stable matchings? No. By the Rural Hospital Theorem, the unmatched set is identical across stable matchings; here that's $\{s_2, c_3\}$. So only $\{s_1, s_3\}$ are matched to $\{c_1, c_2\}$. There are two such matchings:

- $\mu = \{(s_1, c_1), (s_3, c_2)\}$ (the DA outcome).
- $\mu' = \{(s_1, c_2), (s_3, c_1)\}$.

Check μ' : (s_1, c_1) blocks (s_1 prefers c_1 to c_2 ; c_1 prefers s_1 to s_3). Not stable.

The unique stable matching is $\mu_S = \mu_C = \{(s_1, c_1), (s_3, c_2)\}$.

Remark (Coincidence of μ_S and μ_C).

When the student-proposing and college-proposing DAs return the same matching, the stable-matching lattice is a singleton. Sufficient conditions: every agent has a unique top choice, and preferences exhibit “fitness compatibility” (strong agents prefer strong agents). Here, all colleges rank s_1 first, all students rank c_1 first. The Rural Hospital Theorem then forces the unmatched set to coincide, narrowing candidate matchings to two; only one is stable.

Problem 2

This is identical to Problem 12.2 of the problem sets; we recapitulate the key thresholds.

(a) **Trigger strategy.** On-path: (H, H) in every period. Off-path (after deviation): switch permanently to (M, M) . Best one-shot deviation from H against H is L , yielding 5. Punishment value $\frac{2}{1-\delta}$. No-deviation condition:

$$\frac{3}{1-\delta} \geq 5 + \delta \cdot \frac{2}{1-\delta} \iff \delta \geq \frac{2}{3}.$$

(b) **Forgiving strategy.** On-path: (H, H) . Off-path: (L, L) for one period, then return to (H, H) . Restart on deviation from L .

On-path constraint: $\frac{3}{1-\delta} \geq 5 + \delta(-1) + \delta^2 \cdot \frac{3}{1-\delta}$ rearranges to $\delta \geq \frac{1}{2}$.

Punishment-phase constraint: $-1 + \delta \cdot \frac{3}{1-\delta} \geq 0 + \delta(-1) + \delta^2 \cdot \frac{3}{1-\delta}$ rearranges to $\delta \geq \frac{1}{4}$, slack relative to the on-path constraint.

Answer. (H, H) is sustained as SPE under forgiving iff $\delta \geq \frac{1}{2}$, looser than the trigger threshold $\delta \geq \frac{2}{3}$.

Remark (Why forgiveness wins).

The forgiving strategy concentrates punishment in a single high-pain period (loss of 4 relative to cooperation) followed by recovery, while the trigger spreads a milder punishment over an infinite horizon. For deterrence, what matters is the present-value sum of losses; forgiveness compresses the loss into a single period and therefore deters at lower δ . *Shorter, harsher* punishments dominate *longer, milder* ones for sustaining cooperation at moderate patience.

Problem 3

(a) Pooling PBE in which both types choose N .

Strategies. $\sigma_1(S) = \sigma_1(W) = N$. Player 2 (off-path, after G): play R .

Beliefs. On-path, neither type chooses G , so player 2's info set after G is reached with probability 0; assign *any* off-path belief $\mu(S | G) \leq 1/2$.

Verification.

- Player 2 at G : given $\mu \leq 1/2$,

$$\mathbb{E}[u_2 | A] = \mu \cdot 1 + (1 - \mu)(-1) = 2\mu - 1 \leq 0, \quad \mathbb{E}[u_2 | R] = 0.$$

R is weakly best.

- Player 1: both types compare $u_1(N) = 0$ vs $u_1(G) = -1$ (since player 2 plays R). Both prefer N .

Strategies and beliefs are mutually consistent. Pooling-on- N is a PBE.

(b) Pooling PBE in which both types choose G , iff $p \geq 1/2$.

Strategies. $\sigma_1(S) = \sigma_1(W) = G$.

On-path belief. By Bayes' rule, $\mu(S | G) = p$. Player 2's expected utility from A is $2p - 1$, from R is 0; player 2 plays A iff $p \geq 1/2$.

Sequential rationality of player 1. Both types compare:

$$u_1(G) = \begin{cases} 1 & \text{if player 2 plays } A \\ -1 & \text{if player 2 plays } R \end{cases}, \quad u_1(N) = 0.$$

Player 1 prefers G iff player 2 plays A , requiring $p \geq 1/2$.

Conclusion. Pooling-on- G PBE exists iff $p \geq 1/2$.

(c) Separating PBE? No. Two candidates:

$S \rightarrow G, W \rightarrow N$. On-path $\mu(S | G) = 1$, so player 2 plays A . W deviating to G : player 2 still plays A (on-path belief), giving W payoff $1 > 0 = u_1(N)$. W deviates—contradicts separation.

$S \rightarrow N, W \rightarrow G$. On-path $\mu(S | G) = 0$, so player 2 plays R . W gets $u_1(G, R) = -1 < 0 = u_1(N)$. W deviates to N —contradicts separation.

No separating PBE.

Separating Nash equilibrium? Also no. NE doesn't require sequential rationality off-path, but player 2's strategy must still be BR to σ_1 . For Candidate 1, $\sigma_1 \rightarrow G$ implies S , BR for player 2 at G is A , W deviates. For Candidate 2, $\sigma_1 \rightarrow G$ implies W , BR is R , W gets $-1 < 0$, prefers N . Neither NE.

Why no separation works. Both types have *identical* payoffs for all action profiles: the gift's cost and the rejection penalty are type-independent. With no type-specific incentive differential, player 1 cannot credibly signal his type via gift-giving.

Remark (Single-crossing as the key).

In the Spence model, the cost of education e/θ is decreasing in the type, so high types find signaling cheaper. The resulting indifference curves cross exactly once (Spence-Mirrlees). Here, the cost of G (relative to N) is identical for both types, so the indifference curves are parallel—no separation possible. Signaling games *require* the marginal cost of the signal to differ across types. Without that asymmetry, no signaling can credibly distinguish the types.

First Mid-Term Exam (2026)

Date: February 17, 2026. Three questions; answer all three.

Problems

Problem 1

Consider the following two-player game where the payoffs are given in monetary amounts (dollars):

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	1, -1	36, -36	0, 0
<i>D</i>	9, -9	4, -4	16, -16

- (a) First, suppose both players have utility functions $u_1(x) = u_2(x) = x$, so the numbers above are also utilities. Find an equilibrium of the game. What are the expected dollar payoffs in this equilibrium?
- (b) Now suppose player 1 becomes risk-averse with Bernoulli utility $v_1(x) = \sqrt{x}$, while player 2 remains risk-neutral with $u_2(x) = x$. How is the equilibrium from part (a) affected? In particular, how does this affect the expected dollar payoff of player 1? What is player 1's equilibrium expected utility, and what is his dollar certainty equivalent?

Problem 2

The following table gives the per-firm profit π as a function of the total number of firms n in an oligopolistic market:

n	1	2	3	4+
$\pi(n)$	450	350	100	(-)

If there is only one firm, profit is 450; if two firms, each earns 350; with three firms, each earns 100; with four or more firms, each makes a loss. Three potential entrants decide *sequentially*: first firm 1, then firm 2, then firm 3. Absent other considerations, all three enter and earn 100.

- (a) Suppose firms operating in the industry can use advertising as an entry barrier. To enter, a new firm must match incumbent advertising. If firms 1 and 2 are in the market, they must each spend 100 to keep firm 3 out. How many firms enter?

- (b) Now suppose the government regulates advertising: no firm may spend more than 25% of its gross profit. If firms 1 and 2 are in the market, the cap is $0.25 \cdot 350 = 87.5$. What effect does this regulation have on the number of firms?

Problem 3

Consider a market with two firms that produce differentiated products. The demand for firm i 's product as a function of its own price p_i and rival's price p_j is

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

where $a > 1$ and $0 < b < 1$. Constant unit cost is $c \in [0, 1]$ and prices $p_i \in [c, 1]$.

- (a) Show that the resulting game is supermodular. What is the set of rationalizable strategies?
- (b) Is the resulting game also a potential game? If so, what is the potential function?

Solutions

Problem 1

(a) **Both risk-neutral.** The game is zero-sum with player 1's payoff matrix

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	1	36	0
<i>D</i>	9	4	16

No row or column is dominated; the equilibrium is in mixed strategies. Let $p = \Pr(U)$. Player 1's expected payoff against each pure column is

$$\text{vs } L : 9 - 8p, \quad \text{vs } M : 4 + 32p, \quad \text{vs } R : 16 - 16p.$$

Player 2 chooses the minimum; player 1 picks p to maximize the lower envelope. The pairwise crossings:

$$L = M : p = \frac{1}{8} \text{ (value 8)}, \quad M = R : p = \frac{1}{4} \text{ (value 12)}, \quad L = R : p = \frac{7}{8} \text{ (value 2)}.$$

On $[0, \frac{1}{8}]$ binding line is M (increasing); on $[\frac{1}{8}, \frac{7}{8}]$ it is L (decreasing); on $[\frac{7}{8}, 1]$ it is R . Maximum at L - M crossing:

$$\boxed{p^* = \frac{1}{8}, \quad V = 8.}$$

At $p^* = \frac{1}{8}$, player 2 mixes only over L and M (since R gives $14 > 8$). To make player 1 indifferent between U and D :

$$q_L \cdot 1 + q_M \cdot 36 = q_L \cdot 9 + q_M \cdot 4 \implies q_L = 4q_M.$$

With $q_L + q_M = 1$: $q_L = \frac{4}{5}, q_M = \frac{1}{5}$.

Equilibrium. $\sigma_1 = (\frac{1}{8}, \frac{7}{8})$ over (U, D) ; $\sigma_2 = (\frac{4}{5}, \frac{1}{5}, 0)$ over (L, M, R) . Expected dollar payoffs (\$8, -\$8).

(b) **Player 1 risk-averse with** $v_1(x) = \sqrt{x}$. Dollar payoffs unchanged; player 1's utility matrix becomes

	<i>L</i>	<i>M</i>	<i>R</i>
<i>U</i>	1	6	0
<i>D</i>	3	2	4

Player 2 still maximizes the negative of dollar payoff, so the same envelope analysis applies: $p^* = \frac{1}{8}$. Update player 2's mix to keep player 1 indifferent under new utilities:

$$q_L + 6q_M = 3q_L + 2q_M \implies q_L = 2q_M.$$

Hence $q_L = \frac{2}{3}, q_M = \frac{1}{3}$. Verify player 2's indifference (in dollars) at $p = \frac{1}{8}$: vs L , $9 - 1 = 8$; vs M , $4 + 4 = 8$; vs R , $14 > 8$. ✓

Effect on the equilibrium. Player 1's mix unchanged at $(\frac{1}{8}, \frac{7}{8})$; player 2's shifts toward M :

$$\sigma_2 : (\frac{4}{5}, \frac{1}{5}, 0) \rightarrow (\frac{2}{3}, \frac{1}{3}, 0).$$

Intuitively, the high-stakes outcome (U, M) pays \$36, which is only 6 in utility for risk-averse player 1; player 2 must put more weight on M to keep player 1 willing to play U .

Player 1's expected dollar payoff. At $p = \frac{1}{8}$, against any L/M mix, expected dollar:

$$\text{vs } L : \frac{1}{8}(1) + \frac{7}{8}(9) = 8; \quad \text{vs } M : \frac{1}{8}(36) + \frac{7}{8}(4) = 8.$$

Still $\boxed{\$8}$, unchanged.

Player 1's expected utility.

$$\mathbb{E}[v_1] = \frac{1}{8} \left[\frac{2}{3}(1) + \frac{1}{3}(6) \right] + \frac{7}{8} \left[\frac{2}{3}(3) + \frac{1}{3}(2) \right] = \frac{8}{3}.$$

Dollar certainty equivalent. CE solves $\sqrt{\text{CE}} = 8/3$:

$$\boxed{\text{CE} = \frac{64}{9} \approx \$7.11.}$$

Risk premium: $\$8 - \$\frac{64}{9} = \$\frac{8}{9} \approx \0.89 . Expected dollar is the same in (a) and (b), but player 1 would forfeit \$0.89 to avoid the equilibrium gamble.

Problem 2

(a) Unrestricted advertising. Backward induction.

Stage 3 (firm 3 decides, given firms 1, 2 in). If 1, 2 each spend $A = 100$, firm 3's net entry payoff is $\pi(3) - 100 = 0$, stays out. Firms 1, 2 each net $\pi(2) - 100 = 250$.

Stage 2 (firm 2 decides).

- Enter: net 250.
- Stay out: firm 1 alone faces firm 3. To deter firm 3 from entering a 2-firm market, firm 1 needs $A \geq \pi(2) = 350$, leaving $\pi(1) - 350 = 100$. Without deterrence, firm 3 enters and both earn $\pi(2) = 350$. Since $350 > 100$, firm 1 chooses no deterrence; firm 3 enters; firm 2 nets 0.

Firm 2 enters ($250 > 0$).

Stage 1 (firm 1 decides).

- Enter: from stage 2, firm 2 enters, both spend 100 to deter firm 3, firm 1 nets 250.
- Stay out: firm 2 enters alone, firm 3 enters, both net 350, firm 1 nets 0.

Firm 1 enters.

$\boxed{\text{Two firms enter.}}$ Profits: 250, 250, 0.

(b) The 25% regulation. With two firms in the market, gross profit is $\pi(2) = 350$, regulatory cap is 87.5. Firm 3's net entry payoff:

$$\pi(3) - 87.5 = 12.5 > 0.$$

Firm 3 enters under the cap—deterrence is no longer feasible. Working backwards: when firm 1 is alone, the cap is 112.5, far below the 350 needed to deter firm 2; firm 2 enters easily. Firms 2 and 3 then enter, with firm 3 not even spending on ads.

$\boxed{\text{All three firms enter.}}$ Each ends with $\pi(3) = 100$.

Lesson. A modest ad cap can break a sustainable deterrence regime by forcing the matching cost *below* the post-entry profit. The regulation, intended to limit anti-competitive behavior, causes a transition from a duopoly to a triopoly. Welfare comparison is mixed: consumers benefit from more competition, but firms collectively lose rent.

Problem 3

(a) **Supermodularity and rationalizable strategies.** Profit:

$$\pi_i(p_1, p_2) = (p_i - c)(a - p_i + bp_j).$$

Compute partials:

$$\frac{\partial \pi_i}{\partial p_i} = a + c - 2p_i + bp_j, \quad \frac{\partial^2 \pi_i}{\partial p_i \partial p_j} = b > 0.$$

Strategy spaces $[c, 1]$ are lattices; positive cross-partial confirms **supermodularity**.

Best response. Setting $\partial \pi_i / \partial p_i = 0$:

$$\text{BR}_i(p_j) = \frac{a + c + bp_j}{2}.$$

Strictly increasing in p_j with slope $b/2 < 1$, so iterated best response is a contraction. Solving $p^* = \text{BR}(p^*)$:

$$p^* = \frac{a + c}{2 - b}.$$

Rationalizable strategies. In a supermodular game with a contraction-mapping best response, iterated elimination of strictly dominated strategies converges to the unique fixed point:

$$R_i = \left\{ \frac{a + c}{2 - b} \right\}.$$

(Implicit: $a + b + c \leq 2$ so $p^* \leq 1$.)

(b) **Potential function.** A smooth game has a potential function iff $\partial^2 P / \partial p_i \partial p_j = \partial^2 P / \partial p_j \partial p_i$ (Schwarz). Both equal b :

$$\frac{\partial}{\partial p_j} \frac{\partial \pi_i}{\partial p_i} = b = \frac{\partial}{\partial p_i} \frac{\partial \pi_j}{\partial p_j}. \quad \checkmark$$

Try the symmetric quadratic candidate

$$P(p_1, p_2) = (a + c)(p_1 + p_2) - (p_1^2 + p_2^2) + bp_1p_2.$$

Verification:

$$\frac{\partial P}{\partial p_1} = (a + c) - 2p_1 + bp_2 = \frac{\partial \pi_1}{\partial p_1},$$

and symmetrically for p_2 . The game is both supermodular and a (cardinal) potential game; the maximizer of P coincides with the unique NE.

Remark (Two routes to existence).

Supermodularity gives existence via Tarski's fixed-point theorem (best-response monotonicity on a lattice). Potential structure gives existence via maximization of a single function. The Bertrand model with linear demand satisfies both. Whenever both hold, the maximizer of P provides a refinement-friendly NE selection.

Second Mid-Term Exam (2026)

Date: March 31, 2026. Three questions; answer all three.

Problems

Problem 1

Carefully define the following terms.

- (a) A *bargaining solution* (in the axiomatic bargaining framework).
- (b) The *virtual value function* associated with a distribution F .
- (c) The *Vickrey-Clarke-Groves (VCG) mechanism*.

Problem 2

Two players bargain via the Rubinstein alternating-offer procedure to divide a “pie” of size 1.

- (a) Suppose the players have different utility functions: $u_1(z) = z$ and $u_2(z) = \sqrt{z}$, and use a common discount factor $\delta \in (0, 1)$ to discount future utilities. How will the pie be split in the unique subgame perfect equilibrium?
- (b) Now suppose both players have identical linear utility functions $u(z) = z$ but use different discount factors $\delta_1 < \delta_2$ (both in $(0, 1)$).
 1. How will the pie be split?
 2. For what value of δ_2 is the split in part (b) the same as in part (a) when the common discount factor $\delta = \delta_1$?

Problem 3

There is a single object for sale and two potential buyers are bidding for the object. Bidder i assigns a value X_i to the object. Each X_i is iid uniform on $[0, 1]$. Bidder i knows x_i and only that the other bidder’s value is iid uniform on $[0, 1]$. The object has use value 0 to the seller.

- (a) Suppose the seller uses a *second-price* auction (SPA) with reserve price $r > 0$. What is the optimal reserve price r^{SPA} ?

- (b) Now suppose the seller uses a *first-price* auction (FPA) with reserve $r > 0$. What is the equilibrium bidding strategy $\beta_r : [0, 1] \rightarrow \mathbb{R}$? What is the optimal reserve r^{FPA} ? How does it compare to r^{SPA} ?
- (c) Finally, suppose the seller uses an SPA but instead of a reserve price, charges each bidder an entry fee $c > 0$ (non-refundable, paid before bidding). What is the optimal entry fee c^{SPA} ?

Solutions

Problem 1

(a) **Bargaining solution.** Let $U \subseteq \mathbb{R}^2$ denote the set of feasible expected utility profiles, assumed compact and convex; let $d \in U$ be the disagreement point with $\exists u \in U$ such that $u \gg d$. A **bargaining solution** is a function

$$F : (U, d) \mapsto F(U, d) \in U$$

that selects, for each bargaining problem (U, d) , a single utility profile in U . Nash's axiomatic characterization shows that the unique solution satisfying scale invariance, efficiency, symmetry, and IIA is the **Nash bargaining solution**

$$F^N(U, d) = \arg \max_{u \in U, u \geq d} (u_1 - d_1)(u_2 - d_2).$$

(b) **Virtual value function.** Let F be a strictly increasing CDF on $[0, \bar{x}]$ with continuous density $f = F'$. The **virtual value function** is

$$\varphi(x) = x - \frac{1 - F(x)}{f(x)}.$$

It captures the marginal revenue of allocating to a buyer of type x : the buyer's value x minus the information rent $(1 - F(x))/f(x)$ that incentive compatibility forces the seller to leave for higher types. The seller's expected revenue from any IC mechanism with allocation rule q equals $\mathbb{E}[\sum_i \varphi(x_i)q_i(x)]$, so the optimal mechanism allocates to the bidder with the highest non-negative virtual value.

(c) **VCG mechanism.** Let $Q^*(x) \in \arg \max_Q \sum_{j=1}^N Q_j(x)x_j$ be an efficient allocation rule, and write $W(x) = \sum_j Q_j^*(x)x_j$ for total maximized social surplus, $W_{-i}(x) = \sum_{j \neq i} Q_j^*(x)x_j$ for the surplus of agents other than i . The **VCG mechanism** pairs Q^* with the Clarke pivot payment rule

$$M_i^*(x) = W(\alpha_i, x_{-i}) - W_{-i}(x),$$

where α_i is the lowest possible type. The payment is the externality i imposes on others. Truth-telling is weakly dominant; the mechanism is efficient, individually rational, and revenue-maximizing among all efficient and IC mechanisms.

Problem 2

The Rubinstein indifference conditions in SPE are

$$u_1(y^*) = \delta_1 u_1(x^*), \quad u_2(1 - x^*) = \delta_2 u_2(1 - y^*),$$

where x^* is player 1's share when she proposes (G_1) and y^* is player 1's share when player 2 proposes (G_2).

(a) **Different utilities, common δ .** With $u_1(z) = z$, $u_2(z) = \sqrt{z}$, $\delta_1 = \delta_2 = \delta$:

$$y^* = \delta x^*, \quad \sqrt{1-x^*} = \delta \sqrt{1-y^*} \iff 1-x^* = \delta^2(1-y^*).$$

Substituting:

$$1-x^* = \delta^2 - \delta^3 x^* \implies x^*(1-\delta^3) = 1-\delta^2 \implies x^* = \frac{1+\delta}{1+\delta+\delta^2}.$$

Hence $y^* = \delta(1+\delta)/(1+\delta+\delta^2)$.

In the unique SPE, player 1 keeps x^* and offers player 2 $1-x^* = \delta^2/(1+\delta+\delta^2)$. Player 2 accepts immediately. As $\delta \rightarrow 1$, $x^* \rightarrow 2/3$: the more risk-averse (concave) player 2 ends up with the smaller share, because concavity reduces effective patience.

(b.1) **The split (linear, asymmetric δ).** With $u(z) = z$ and $\delta_1 < \delta_2$:

$$y^* = \delta_1 x^*, \quad 1-x^* = \delta_2(1-y^*) = \delta_2 - \delta_1 \delta_2 x^*.$$

Solving:

$$1-\delta_2 = x^*(1-\delta_1\delta_2) \implies x^* = \frac{1-\delta_2}{1-\delta_1\delta_2}.$$

Player 1 keeps x^* and offers $1-x^* = \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$.

Since $\delta_1 < \delta_2$ (player 1 more impatient), as $\delta_1, \delta_2 \rightarrow 1$, $x^* \rightarrow \frac{1}{2}$ from below: more impatience hurts the proposer's bargaining position.

(b.2) **Equating splits.** Setting (b.1)'s share equal to (a)'s when $\delta = \delta_1$:

$$\frac{1-\delta_2}{1-\delta_1\delta_2} = \frac{1+\delta_1}{1+\delta_1+\delta_1^2}.$$

Cross-multiplying and expanding:

$$(1-\delta_2)(1+\delta_1+\delta_1^2) = (1+\delta_1)(1-\delta_1\delta_2).$$

After collecting like terms in δ_2 ,

$$\delta_2 = \delta_1^2.$$

Interpretation. The square-root utility in (a) is mathematically equivalent (for the Rubinstein indifference equations) to a linear utility with effective discount δ^2 . Composing $\sqrt{\cdot}$ with one period of δ -discounting yields a payoff that, in linear-utility terms, has been discounted by δ^2 . Concavity is observationally indistinguishable from extra impatience.

Problem 3

(a) **Optimal reserve in SPA.** Truth-telling is dominant. Decomposing revenue by realized event:

$$\begin{aligned} R^{\text{SPA}}(r) &= \underbrace{(1-r)^2 \cdot (r + (1-r)/3)}_{\text{both bid}} + \underbrace{2r(1-r) \cdot r}_{\text{only one bids, pays reserve}} + 0 \\ &= \frac{(1-r)(1+r+4r^2)}{3}. \end{aligned}$$

Differentiating,

$$\frac{d}{dr} [1 + 3r^2 - 4r^3] = 6r - 12r^2 = 6r(1 - 2r).$$

FOC $r = \frac{1}{2}$:

$$\boxed{r^{\text{SPA}} = \frac{1}{2}}, \quad R^{\text{SPA}}\left(\frac{1}{2}\right) = \frac{5}{12}.$$

Cross-check: $\varphi(x) = 2x - 1$, $\varphi(r) = 0 \iff r = \frac{1}{2}$.

(b) Optimal reserve in FPA. In the symmetric equilibrium with strictly increasing β_r on $[r, 1]$, a type- x bidder ($x \geq r$) chooses b to maximize win-probability times surplus. Under symmetry, bidder i wins iff opponent's value is below $\beta_r^{-1}(b)$ (covers both "opponent doesn't enter" and "opponent has a lower bid"). Pretending to be type $v = \beta_r^{-1}(b)$:

$$U(v; x) = v(x - \beta_r(v)).$$

FOC at $v = x$:

$$x - \beta_r(x) - x\beta_r'(x) = 0 \iff \frac{d}{dx} [x\beta_r(x)] = x.$$

Integrate from r to x with boundary $\beta_r(r) = r$:

$$x\beta_r(x) - r^2 = \frac{x^2 - r^2}{2} \implies \boxed{\beta_r(x) = \frac{x^2 + r^2}{2x}, \quad x \in [r, 1]}.$$

Optimal reserve. By the Revenue Equivalence Theorem, FPA and SPA with the same reserve allocate identically and assign the same expected payment to the marginal type, so $R^{\text{FPA}}(r) = R^{\text{SPA}}(r)$ for every r :

$$\boxed{r^{\text{FPA}} = r^{\text{SPA}} = \frac{1}{2}}.$$

(c) Optimal entry fee in SPA. Symmetric equilibrium has a participation cutoff x^* : bidders enter iff $x \geq x^*$. Among participants, bidding is truthful.

A participating type $x \geq x^*$ has interim utility

$$U(x) = x^* \cdot x + \int_{x^*}^x (x - y) dy - c = \frac{x^2 + (x^*)^2}{2} - c.$$

The cutoff type is indifferent: $U(x^*) = 0$:

$$(x^*)^2 = c \implies x^* = \sqrt{c}.$$

Equivalence with reserve SPA. The interim expected payment of a participating type x is

$$m^{\text{entry}}(x) = c + \int_{x^*}^x y dy = \frac{(x^*)^2 + x^2}{2},$$

which (with $x^* = r$) coincides with the reserve- r SPA payment. The two mechanisms are payment-equivalent type by type, hence revenue-equivalent. Maximizing over c is equivalent to maximizing over $r = \sqrt{c}$:

$$\sqrt{c^{\text{SPA}}} = r^{\text{SPA}} = \frac{1}{2} \implies \boxed{c^{\text{SPA}} = \frac{1}{4}}.$$

Maximum revenue is again $5/12$.

Remark (Reserves and entry fees as duals).

Both instruments—a reserve price r and an entry fee c —implement the same allocation rule: only types above some cutoff transact, and the cutoff is set where virtual value equals zero. Their revenue equivalence at the optimum is no coincidence: both mechanisms generate the same interim payment schedule and extract the same information rent. For practical market design, the choice is administrative (entry fees are simpler to enforce; reserves can be set after observing competition).

Final Exam (2026)

Spring 2026. Three questions; answer all three.

Problems

Problem 1

Consider a two-sided matching problem between students in the set $S = \{s_1, s_2, s_3\}$ and colleges in the set $C = \{c_1, c_2, c_3\}$. Each college has *only one slot*. The (strict) preferences of the agents (top is best etc.) are:

$P(s_1)$	$P(s_2)$	$P(s_3)$	$P(c_1)$	$P(c_2)$	$P(c_3)$
c_1	c_2	c_1	s_2	s_1	s_1
c_2	c_1	c_2	s_1	s_2	s_2
c_3	c_3	c_3	s_3	s_3	s_3

- Suppose that matching is determined by the Gale-Shapley algorithm when students make offers (apply). What is the resulting stable matching μ_S ?
- Is the matching $\mu = \{(c_1, s_3), (c_2, s_1), (c_3, s_2)\}$ stable?
- Show that college c_1 can manipulate the outcome in its favor by submitting the false preference $P'(c_1) = s_2 \succ s_3 \succ s_1$ instead of $P(c_1)$. In other words, if μ'_S denotes the outcome of the Gale-Shapley algorithm with students proposing when c_1 submits the false preference $P'(c_1)$ (and all others submit their true preferences), then c_1 is better-off in the matching μ'_S rather than μ_S .

Problem 2

Consider the prisoners' dilemma game G :

	C	D
C	2, 2	-1, 3
D	3, -1	0, 0

and denote by $G(2)$ its two-fold repetition in which before playing in the second period, both players observe all choices in period 1.


- Show that the unique *Nash equilibrium outcome* of $G(2)$ is $(D, D), (D, D)$. (Note that the question is not about subgame perfect equilibria but rather Nash equilibria.)

(b) Next consider the modified prisoner's dilemma game G' :

	C	D	E
C	2, 2	-1, 3	-2, -2
D	3, -1	0, 0	-1, -2
E	-2, -2	-2, -1	-2, -2

Show that $(C, C), (D, D)$ is a Nash equilibrium outcome of $G'(2)$. Is it a subgame perfect equilibrium outcome of $G'(2)$ as well?

Problem 3



Consider the entry game above in which a chain-store (player C) faces a threat of entry from an entrant (player E). E can either enter (play In) or not (play Out) and C can either fight entry (play F) or accommodate it (play A). The payoffs are given in the tree where the first number is the entrant's payoff and the second is the chain-store's payoff.

Suppose that there are two (2) periods and in each period, the game above is played against a *different* entrant. In period 1, the chain-store faces the first entrant, E_1 , and in period 2, it faces E_2 . Outcomes in period 1 are observed by both C and E_2 .

But there is a probability $\varepsilon < \frac{1}{4}$ that the chain-store's payoffs are not as specified above but rather such that the chain-store prefers to fight (play F) rather than to accommodate (play A). Call such a chain-store "irrational".

Consider the following strategies in the two-period game when the probability that the chain-store is irrational is $\varepsilon < \frac{1}{4}$.

- (i) In period 1, the entrant E_1 enters for sure. The chain-store fights with probability $\varepsilon/(1 - \varepsilon)$.
- (ii) In period 2, the entrant E_2 enters for sure if the outcome in period 1 was (In, A) and enters with probability $\frac{1}{2}$ if the outcome in period 1 was (In, F) .
- (a) Show that the strategies in (i) and (ii) constitute a perfect Bayesian equilibrium (PBE). Make sure you specify players' beliefs.
- (b) **(Optional for extra credit)** Show that the outcome in (a) is the unique PBE outcome of this game.

Solutions

Problem 1

(a) Student-proposing DA.

Round 1. $s_1, s_3 \rightarrow c_1$ and $s_2 \rightarrow c_2$. College c_1 has $\{s_1, s_3\}$ and under $P(c_1) = s_2 \succ s_1 \succ s_3$ holds s_1 , rejects s_3 . College c_2 holds s_2 .

Round 2. $s_3 \rightarrow c_2$ (next on $P(s_3)$). c_2 has $\{s_2, s_3\}$; under $P(c_2) = s_1 \succ s_2 \succ s_3$ holds s_2 , rejects s_3 .

Round 3. $s_3 \rightarrow c_3$. c_3 holds s_3 .

$$\mu_S = \{(s_1, c_1), (s_2, c_2), (s_3, c_3)\}.$$

(b) Stability of $\mu = \{(c_1, s_3), (c_2, s_1), (c_3, s_2)\}$.

Consider the pair (s_1, c_1) . Under μ , s_1 is matched to c_2 , but $P(s_1)$ ranks $c_1 \succ c_2$, so s_1 prefers c_1 . Under μ , c_1 is matched to s_3 , but $P(c_1)$ ranks $s_1 \succ s_3$, so c_1 prefers s_1 . Hence (s_1, c_1) is a blocking pair. The matching μ is *not* stable.

(c) Manipulation by c_1 . Run student-proposing GS with c_1 's reported preference $P'(c_1) = s_2 \succ s_3 \succ s_1$ (others truthful).

Round 1. $s_1, s_3 \rightarrow c_1$ and $s_2 \rightarrow c_2$. Under $P'(c_1)$, c_1 holds s_3 and rejects s_1 . c_2 holds s_2 .

Round 2. $s_1 \rightarrow c_2$. c_2 has $\{s_1, s_2\}$; under $P(c_2)$ holds s_1 , rejects s_2 .

Round 3. $s_2 \rightarrow c_1$. c_1 has $\{s_2, s_3\}$; under $P'(c_1) = s_2 \succ s_3 \succ s_1$ holds s_2 , rejects s_3 .

Round 4. $s_3 \rightarrow c_2$. c_2 has $\{s_1, s_3\}$; under $P(c_2)$ holds s_1 , rejects s_3 .

Round 5. $s_3 \rightarrow c_3$. c_3 holds s_3 .

$$\mu'_S = \{(s_1, c_2), (s_2, c_1), (s_3, c_3)\}.$$

Comparing under c_1 's *true* preference $P(c_1) = s_2 \succ s_1 \succ s_3$: in μ_S , c_1 gets s_1 ; in μ'_S , c_1 gets s_2 . Since $s_2 \succ_{c_1} s_1$, college c_1 is strictly better off under μ'_S . Misreporting pays.

Remark (Why students-proposing DA is not strategy-proof for colleges).

The Gale-Shapley algorithm is strategy-proof on the proposing side but not on the receiving side. By dropping s_1 down its list, c_1 rejects s_1 in round 1, which forces s_1 to chase c_2 and displace s_2 . The displaced s_2 then proposes to c_1 , who—having reserved a slot by rejecting its true top—now lands its true top choice s_2 . This is the canonical example showing that no stable mechanism is strategy-proof for both sides; one side must be willing to manipulate.

Problem 2

(a) Unique NE outcome of $G(2)$ is $(D, D), (D, D)$.

The stage game G has D strictly dominant for both players, hence the unique stage NE is (D, D) with payoff 0.

Consider any NE (σ_1, σ_2) of $G(2)$. In any *on-path* period-2 history, the action profile played must be a stage-NE of G : each player's period-2 action solves a one-shot best-response problem given the opponent's period-2 action (continuation payoffs are zero in the

last period). Since (D, D) is the unique NE of G , every on-path period-2 history must produce (D, D) , with continuation payoff 0 for each player regardless of period-1 play.

Knowing period-2 contributes 0 to total payoff irrespective of what happens in period 1, period-1 play in any NE must be a NE of G itself. Hence period-1 must also be (D, D) .

The unique NE outcome is $((D, D), (D, D))$.

(b) $(C, C), (D, D)$ is a NE outcome of $G'(2)$.

Note first that in G' , action E is strictly dominated by D for both players (compare row by row: D payoffs 3, 0, -1 vs E payoffs $-2, -2, -2$, and similarly by columns). So (D, D) is again the unique stage NE of G' .

Consider the strategy profile:

- Period 1: play C .
- Period 2: play D if period-1 outcome was (C, C) ; play E if anyone deviated in period 1.

Beliefs are trivial (perfect monitoring); we check NE.

On-path payoffs. $2 + 0 = 2$ for each player.

Deviation by player 1. The most attractive period-1 deviation is D , paying 3. Under the prescribed strategies, player 2 plays E in period 2. Player 1's strategy in period 2 (after the deviation) is also E , giving -2 . Total: $3 + (-2) = 1 < 2$.

(Player 1 can additionally deviate in period 2 against player 2's E : best response is D , giving -1 . Total: $3 + (-1) = 2 \leq 2$. Not strictly better.)

By symmetry the same holds for player 2. Hence the profile is a NE; the outcome path is $(C, C), (D, D)$. ✓

Is it SPE? *No.* The threat “play E if there was a period-1 deviation” is not credible: in the period-2 subgame following any history, the only NE of the stage game G' is (D, D) , since E is strictly dominated. Playing E off-path is therefore not sequentially rational—neither player is best-responding at the off-path period-2 information set. Thus the profile fails subgame perfection.

In fact, by the same logic as part (a), every SPE of $G'(2)$ plays the unique stage-NE (D, D) in both periods. The outcome $(C, C), (D, D)$ cannot be sustained as an SPE outcome.

Remark (NE vs SPE in finite repetition).

Cooperation can be sustained as a NE outcome in $G'(2)$ but not in $G(2)$, even though both stage games have (D, D) as the unique NE. The difference is that G' contains a third action E that is *worse* than the stage NE D . This gives players an off-path threat that is harsher than the stage NE—enough to deter the period-1 cooperation deviation. In $G(2)$ no such threat exists, since (D, D) is the worst achievable continuation (given that period-2 must be a stage NE). The trick collapses under SPE because the threat itself is not stage-NE; SPE forces both periods to be (D, D) .

Problem 3

Throughout, write σ for the probability that a *rational* chain-store fights when entered upon in period 1, and q for E_2 's probability of entering after observing (In, F) in period 1. The proposed strategies fix $\sigma = \varepsilon/(1 - \varepsilon)$ and $q = \frac{1}{2}$.

(a) Verification of PBE.*Beliefs.*

- E_1 's prior at his initial decision: $\Pr(\text{() rational}) = 1 - \varepsilon$.
- E_2 after (In, A) : rational store with probability 1. (Irrational store always plays F , so A rules out irrationality.)
- E_2 after (In, F) : by Bayes' rule,

$$\mu \equiv \Pr(\text{() rational} \mid F) = \frac{(1 - \varepsilon)\sigma}{(1 - \varepsilon)\sigma + \varepsilon \cdot 1} = \frac{(1 - \varepsilon) \cdot \frac{\varepsilon}{1 - \varepsilon}}{\varepsilon + \varepsilon} = \frac{1}{2}.$$

- E_2 after Out : off path; assign the prior, $\Pr(\text{() rational}) = 1 - \varepsilon$ (Bayes-consistent since Out reveals no type information).

Sequential rationality, period 2 (last period).

- In period 2, the rational store strictly prefers A (gives 0) to F (gives -1) whenever entered upon—it's a one-shot subgame. The irrational store fights by definition. ✓
- E_2 after (In, A) : believes store is rational, so accommodation is certain; payoff from In is 1, from Out is 0. Enters. ✓
- E_2 after (In, F) : $\mu = \frac{1}{2}$, so expected payoff from In is $\frac{1}{2}(1) + \frac{1}{2}(-1) = 0$, equal to Out . Indifferent; mixing $\frac{1}{2}$ - $\frac{1}{2}$ is sequentially rational. ✓

Sequential rationality, period 1.

- Rational store after E_1 enters. Continuation payoffs:

$$\begin{aligned} u_C(F) &= -1 + [q \cdot 0 + (1 - q) \cdot 2] = -1 + 2(1 - q) = 1 - 2q, \\ u_C(A) &= 0 + [1 \cdot 0] = 0, \end{aligned}$$

where after (In, A) the rational store accommodates entry in period 2 (payoff 0), and after (In, F) entry occurs with probability q (giving 0 since the rational store accommodates) and Out occurs with probability $1 - q$ (giving 2). At $q = \frac{1}{2}$: $u_C(F) = u_C(A) = 0$. The rational store is indifferent, so mixing $\sigma = \varepsilon/(1 - \varepsilon)$ is sequentially rational. ✓

- E_1 's expected payoff from In :

$$\Pr(\text{() } F) \cdot (-1) + \Pr(\text{() } A) \cdot 1 = -2\varepsilon + (1 - 2\varepsilon) = 1 - 4\varepsilon > 0,$$

using $\Pr(\text{() } F) = (1 - \varepsilon)\sigma + \varepsilon = 2\varepsilon$. Since $\varepsilon < \frac{1}{4}$, this exceeds the Out -payoff of 0. ✓

All beliefs are consistent with strategies via Bayes' rule on path, and all actions are sequentially rational given beliefs. The pair (β, μ) is a PBE.

(b) Uniqueness of the PBE outcome (extra credit).

Fix any PBE. We show the equilibrium has the same play as in part (a).

Step 1: Period-2 rational store accommodates whenever entered upon. The period 2 subgame (after any history with entry) is a one-shot game; the rational store strictly prefers A to F . So in any PBE, the rational store accommodates in period 2.

Step 2: E_2 's period-2 strategy is pinned down by beliefs. Given step 1:

- After (In, A) , E_2 's posterior must put probability 1 on rational (irrational types never play A). So E_2 strictly prefers In (payoff 1 vs 0); enters for sure.
- After (In, F) , let μ be E_2 's posterior on rational. Expected payoff from In is $2\mu - 1$; from Out is 0. E_2 enters iff $\mu > \frac{1}{2}$, stays out iff $\mu < \frac{1}{2}$, and may mix (any $q \in [0, 1]$) iff $\mu = \frac{1}{2}$.

Step 3: The period-1 rational store mixes. Suppose E_1 enters (we verify below); let σ be the probability the rational store fights and q the probability E_2 enters after (In, F) .

- If $\sigma = 0$: by Bayes, $\mu = 0$, so E_2 stays out after (In, F) , $q = 0$. Then $u_C(F) = 1 > 0 = u_C(A)$ for the rational store, contradicting $\sigma = 0$.
- If $\sigma = 1$: by Bayes, $\mu = 1 - \varepsilon > \frac{1}{2}$ (since $\varepsilon < \frac{1}{4} < \frac{1}{2}$), so E_2 enters after F , $q = 1$. Then $u_C(F) = -1 < 0 = u_C(A)$ for the rational store, contradicting $\sigma = 1$.
- Mixed $\sigma \in (0, 1)$: rational store indifferent requires $1 - 2q = 0$, i.e., $q = \frac{1}{2}$. For E_2 to mix, $\mu = \frac{1}{2}$, which by Bayes requires $(1 - \varepsilon)\sigma = \varepsilon$, i.e., $\sigma = \varepsilon/(1 - \varepsilon)$.

So in any PBE in which E_1 enters, $\sigma = \varepsilon/(1 - \varepsilon)$ and $q = \frac{1}{2}$.

Step 4: E_1 enters. Given the period-1 strategies in step 3, E_1 's expected payoff from In is $1 - 4\varepsilon > 0$. So In strictly dominates Out ; E_1 enters. Hence the assumption “ E_1 enters” in step 3 is in fact forced.

(If we instead try to construct a PBE in which E_1 stays Out , we still need to specify off-path play for the rational store and E_2 . By step 1, the rational store accommodates in period 2; by step 3 applied to the off-path subgame, the rational store mixes with $\sigma = \varepsilon/(1 - \varepsilon)$ in period 1. But then E_1 's expected payoff from In is $1 - 4\varepsilon > 0$, contradicting the optimality of Out . So no such PBE exists.)

The PBE is therefore essentially unique on the path of play, with outcome distribution: E_1 enters; with probability 2ε the store fights and E_2 randomizes; with probability $1 - 2\varepsilon$ the store accommodates and E_2 enters for sure.

Remark (Reputation as the engine).

Without the small probability ε that the store is irrational, the unique SPE of the two-period entry game is for both entrants to enter and the (rational) store to accommodate—fighting is never sequentially rational in the last period, and unraveling kills deterrence in period 1 too. The Kreps-Wilson reputation model shows that an arbitrarily small $\varepsilon > 0$ overturns this: the rational store mimics the irrational type in period 1 with positive probability, exploiting E_2 's residual uncertainty to deter entry. The probability of fighting $\varepsilon/(1 - \varepsilon)$ is calibrated to keep E_2 's posterior at exactly $\frac{1}{2}$, the indifference threshold; any higher and E_2 would surely stay out (so the rational store would always fight, breaking equilibrium); any lower and E_2 would surely enter (so the rational store would never fight, also breaking equilibrium).

Bibliography

- Abdulkadiroğlu, Atila and Tayfun Sönmez (1998). “Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems”. In: *Econometrica* 66.3, pp. 689–701.
- Arrow, Kenneth J. (1979). “The Property Rights Doctrine and Demand Revelation under Incomplete Information”. In: *Economics and Human Welfare*, pp. 23–39.
- Aumann, Robert J. (1959). “Acceptable Points in General Cooperative n -Person Games”. In: *Contributions to the Theory of Games IV*, pp. 287–324.
- (1976). “Agreeing to Disagree”. In: *Annals of Statistics* 4.6, pp. 1236–1239.
- (1995). “Backward Induction and Common Knowledge of Rationality”. In: *Games and Economic Behavior* 8.1, pp. 6–19.
- Banks, Jeffrey S. and Joel Sobel (1987). “Equilibrium Selection in Signaling Games”. In: *Econometrica* 55.3, pp. 647–661.
- Benoît, Jean-Pierre and Vijay Krishna (1985). “Finitely Repeated Games”. In: *Econometrica* 53.4, pp. 905–922.
- Bulow, Jeremy and Paul Klemperer (1996). “Auctions versus Negotiations”. In: *American Economic Review* 86.1, pp. 180–194.
- Cho, In-Koo and David M. Kreps (1987). “Signaling Games and Stable Equilibria”. In: *Quarterly Journal of Economics* 102.2, pp. 179–221.
- Clarke, Edward H. (1971). “Multipart Pricing of Public Goods”. In: *Public Choice* 11, pp. 17–33.
- d’Aspremont, Claude and Louis-André Gérard-Varet (1979). “Incentives and Incomplete Information”. In: *Journal of Public Economics* 11.1, pp. 25–45.
- Fudenberg, Drew and Jean Tirole (1991). *Game Theory*. MIT Press.
- Gale, David and Lloyd S. Shapley (1962). “College Admissions and the Stability of Marriage”. In: *American Mathematical Monthly* 69.1, pp. 9–15.
- Geanakoplos, John (2003). “Nash and Walras Equilibrium via Brouwer”. In: *Economic Theory* 21.2-3, pp. 585–603.
- Green, Edward J. and Robert H. Porter (1984). “Noncooperative Collusion under Imperfect Price Information”. In: *Econometrica* 52.1, pp. 87–100.
- Groves, Theodore (1973). “Incentives in Teams”. In: *Econometrica* 41.4, pp. 617–631.
- Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson (1982). “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma”. In: *Journal of Economic Theory* 27.2, pp. 245–252.
- Kreps, David M. and Robert Wilson (1982). “Reputation and Imperfect Information”. In: *Journal of Economic Theory* 27.2, pp. 253–279.

- Krishna, Vijay (2009). *Auction Theory*. 2nd ed. Academic Press.
- Mailath, George J., Masahiro Okuno-Fujiwara, and Andrew Postlewaite (1993). “Belief-Based Refinements in Signalling Games”. In: *Journal of Economic Theory* 60.2, pp. 241–276.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory*. Oxford University Press.
- Maskin, Eric and John Riley (2000). “Asymmetric Auctions”. In: *Review of Economic Studies* 67.3, pp. 413–438.
- Milgrom, Paul and John Roberts (1990). “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities”. In: *Econometrica* 58.6, pp. 1255–1277.
- Monderer, Dov and Dov Samet (1989). “Approximating Common Knowledge with Common Beliefs”. In: *Games and Economic Behavior* 1.2, pp. 170–190.
- Monderer, Dov and Lloyd S. Shapley (1996). “Potential Games”. In: *Games and Economic Behavior* 14.1, pp. 124–143.
- Myerson, Roger B. (1981). “Optimal Auction Design”. In: *Mathematics of Operations Research* 6.1, pp. 58–73.
- Myerson, Roger B. and Mark A. Satterthwaite (1983). “Efficient Mechanisms for Bilateral Trading”. In: *Journal of Economic Theory* 29.2, pp. 265–281.
- Nash, John F. (1950a). “Equilibrium Points in n -Person Games”. In: *Proceedings of the National Academy of Sciences* 36.1, pp. 48–49.
- (1950b). “The Bargaining Problem”. In: *Econometrica* 18.2, pp. 155–162.
- (1951). “Non-Cooperative Games”. In: *Annals of Mathematics* 54.2, pp. 286–295.
- Osborne, Martin J. and Ariel Rubinstein (1994). *A Course in Game Theory*. MIT Press.
- Roth, Alvin E. (1982). “The Economics of Matching: Stability and Incentives”. In: *Mathematics of Operations Research* 7.4, pp. 617–628.
- Roth, Alvin E. and Marilda A. O. Sotomayor (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press.
- Roughgarden, Tim and Éva Tardos (2002). “How Bad is Selfish Routing?” In: *Journal of the ACM* 49.2, pp. 236–259.
- Rubinstein, Ariel (1982). “Perfect Equilibrium in a Bargaining Model”. In: *Econometrica* 50.1, pp. 97–109.
- (1989). “The Electronic Mail Game: Strategic Behavior under “Almost Common Knowledge””. In: *American Economic Review* 79.3, pp. 385–391.
- Shapley, Lloyd S. and Herbert Scarf (1974). “On Cores and Indivisibility”. In: *Journal of Mathematical Economics* 1.1, pp. 23–37.
- Spence, A. Michael (1973). “Job Market Signaling”. In: *Quarterly Journal of Economics* 87.3, pp. 355–374.
- Topkis, Donald M. (1979). “Equilibrium Points in Nonzero-Sum n -Person Submodular Games”. In: *SIAM Journal on Control and Optimization* 17.6, pp. 773–787.
- Vickrey, William (1961). “Counterspeculation, Auctions, and Competitive Sealed Tenders”. In: *Journal of Finance* 16.1, pp. 8–37.
- Zahavi, Amotz (1975). “Mate Selection: A Selection for a Handicap”. In: *Journal of Theoretical Biology* 53.1, pp. 205–214.

Index

- AGV mechanism, 75
- alternating-offer bargaining, 33
- auction
 - first-price (FPA), 40
 - private-value, 39
 - second-price (SPA), 39
- Aumann's agreement theorem, 129
- bargaining problem, 27
- bargaining solution, 27
- belief, 5
- belief system, 161
- Benoit-Krishna theorem, 151
- best response, 7
- bilateral trade impossibility, 77
- blocking pair, 100
- core, 81
- dominance
 - strict, 7
 - weak, 8
- ε -equilibrium, 25
- equilibrium
 - Nash, 10
 - perfect Bayesian, 162
 - subgame-perfect, 20
 - trembling-hand, 25
- expected externality mechanism, 75
- folk theorem, 145
- game
 - imperfect information, 2
 - perfect information, 2
 - potential, 14
- housing market, 81
- IESDS, 12
- imperfect information, 2
- iteratively undominated set (IUD), 12
- matching
 - housing market, 81
 - two-sided, 100
- minmax payoff, 143
- mixed strategy, 5
- Myerson
 - virtual value, 63
- Myerson-Satterthwaite, 77
- Nash bargaining solution, 27
- Nash equilibrium, 10
- Nash's theorem, 10
- never a best response (NBR), 8
- one-deviation principle, 137
- perfect Bayesian equilibrium (PBE), 162
- perfect information, 2
- potential game, 14
- rationalizability, 13
- rationalizable set (RAT), 13
- reserve price, 47
- revelation principle, 58
- revenue equivalence theorem, 52
- Rubinstein bargaining, 33
- single-crossing property, 164
- Spence-Mirrlees condition, 164
- stable matching, 100
- strategy
 - mixed, 5
 - pure, 4
- strictly dominated strategy, 7
- subgame, 20
- subgame-perfect equilibrium (SPE), 20
- supermodular

function, 17
supermodular game, 17
trembling-hand perfect equilibrium, 25
VCG mechanism, 69
Vickrey auction, 39
Vickrey-Clarke-Groves, 69
virtual value (Myerson), 63
weakly dominated strategy, 8