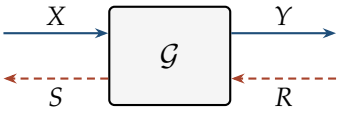
Higher Categorical Semantics

Games and Information in Higher Categories

Composition, descent, and coinduction in noncooperative game theory



Anonymous

2026

About this text. This monograph is a higher-categorical treatment of the territory mapped by Eric Rasmusen's *Games and Information: An Introduction to Game Theory* (4th ed., Blackwell, 2006): noncooperative games, the economics of asymmetric information, and their applications to bargaining, auctions and pricing. It is an independent work. It reproduces none of that book's text, tables, figures, or problems; it is not affiliated with or endorsed by its author or publisher. Appendix B gives a concordance so that the two can be read side by side.

Status of results. Every numbered statement carries a tag: classical (standard game theory, attributed), literature (established in the categorical or type-theoretic literature, attributed), reformulation (a known result restated in a different language, with the restatement proved or sketched here), observation (an elementary remark we have not found stated in this form), or programme (a conjecture or a direction, not a result). No cited theorem is claimed as new.

Authorship and assistance. The text was drafted and typeset with a large language model (Anthropic's Claude) under the direction of the submitting author, who is responsible for it. Generative AI assisted with drafting, literature recall, mathematical consistency checks, computation, and typesetting. The numerical examples and several structural claims are recomputed by the accompanying script `verify.py` (Appendix C). Proofs and bibliographic details have not been independently refereed; corrections are welcome.

Open access. Rights remain with the author.

Contents

Preface viii

Introduction xi

I Games 1

1 The Rules of the Game 2
- 1.1 Decisions . . . 2
- 1.2 Games in strategic form, and their types . . . 2
- 1.3 Dominance as rewriting . . . 5
- 1.4 Nash equilibrium . . . 6
- 1.5 Symmetry and focal points . . . 7

2 Open Games 10
- 2.1 Lenses . . . 10
- 2.2 Open games . . . 11
- 2.3 Atoms and the strategic form . . . 12
- 2.4 Selection functions and backward induction . . . 13
- 2.5 Optics, parameters, and the higher structure . . . 14
- 2.6 What composition does not do . . . 15

3 The Extensive Form 17
- 3.1 Trees are presheaves . . . 17
- 3.2 A taxonomy of information . . . 18
- 3.3 Perfect recall . . . 19
- 3.4 When are two trees the same game? . . . 20

4 Knowledge 25
- 4.1 Information structures . . . 25
- 4.2 Knowledge is base change . . . 25
- 4.3 Common knowledge . . . 27
- 4.4 The homotopy colimit . . . 28
- 4.5 Many agents: simplicial models . . . 29
- 4.6 Knowledge and time . . . 30

5 Belief **32**
5.1 Probability as a monad . . . 32
5.2 Games with incomplete information . . . 33
5.3 Types as a terminal coalgebra . . . 34
5.4 The common prior as a gluing problem . . . 36
5.5 Agreement . . . 38

6 Mixing and Existence **40**
6.1 The mixed extension is free . . . 40
6.2 Existence . . . 41
6.3 The equilibrium correspondence . . . 42
6.4 Correlation . . . 44
6.5 Continuous games . . . 45

7 Order **47**
7.1 Posets as categories . . . 47
7.2 Complementarity . . . 48
7.3 Supermodular games . . . 48
7.4 Quantity competition . . . 49
7.5 Price competition . . . 50
7.6 Ordinal conditions . . . 50
7.7 The metric analogue . . . 51

8 Perfection **53**
8.1 Finite games are an initial algebra . . . 53
8.2 Subgame perfection as a modality . . . 54
8.3 Commitment . . . 55
8.4 The open-set problem . . . 56
8.5 Renegotiation . . . 57

9 Repetition **59**
9.1 Finite repetition . . . 59
9.2 Strategies are coalgebras . . . 60
9.3 Perfection, coinductively . . . 60
9.4 Discounting is a guard . . . 61
9.5 Self-generation . . . 62
9.6 The folk theorem . . . 63
9.7 Markov perfection . . . 64
9.8 Evolution . . . 64
9.9 Computability . . . 65

10 Dynamic Games of Incomplete Information **67**
10.1 Sequential rationality . . . 67
10.2 The geometry of conditioning . . . 68
10.3 Reputation . . . 69
10.4 Almost common knowledge . . . 71

II Asymmetric Information **74**

11 Moral Hazard **75**
11.1 Five models and one diagram . . . 75
11.2 The principal–agent problem . . . 76
11.3 Informativeness . . . 77
11.4 Linear contracts . . . 79

12 Moral Hazard, Continued **81**
12.1 Efficiency wages . . . 81
12.2 Tournaments . . . 82
12.3 Teams . . . 82
12.4 Many tasks . . . 83
12.5 Renegotiation . . . 84
12.6 Extremes: forcing contracts and their failure to exist . . . 84

13 Adverse Selection **87**
13.1 Lemons . . . 87
13.2 Insurance . . . 88
13.3 Securities . . . 89

14 Mechanism Design **92**
14.1 The revelation principle . . . 92
14.2 Transfers: the incentive metric . . . 93
14.3 One-dimensional types . . . 95
14.4 Duality: prices and utilities . . . 96
14.5 Optimal screening . . . 96
14.6 Many agents: Groves mechanisms . . . 98

15 Signalling **100**
15.1 Signalling games . . . 100
15.2 Education as a signal . . . 100
15.3 Refinements . . . 101
15.4 Screening . . . 102
15.5 Cheap talk . . . 102

15.6 Persuasion . . . 103
15.7 Signal jamming . . . 104

III Applications **107**

16 Bargaining **108**
16.1 The axiomatic solution . . . 108
16.2 The strategic solution . . . 109
16.3 Private information and delay . . . 110
16.4 The impossibility of efficient bargaining . . . 111

17 Auctions **114**
17.1 Formats, and two kinds of equivalence . . . 114
17.2 Private values . . . 115
17.3 Revenue equivalence . . . 115
17.4 Optimal auctions . . . 116
17.5 Risk aversion . . . 116
17.6 Common values . . . 117
17.7 Asymmetric information among bidders . . . 118

18 Pricing **120**
18.1 Quantities or prices? . . . 120
18.2 Escaping the Bertrand paradox . . . 120
18.3 Location . . . 121
18.4 Comparative statics without calculus . . . 122
18.5 Durable goods . . . 122
18.6 Entry . . . 123

IV Assessment **126**

19 What the Structure Buys **127**
19.1 What was shown to be the same . . . 127
19.2 What was said that could not be said before . . . 128
19.3 What the language does not do . . . 128
19.4 Problems . . . 129
19.5 A last remark on method . . . 131

A Categorical Toolkit **132**
A.1 Categories, adjunctions, (co)algebras . . . 132
A.2 Monoidal categories, lenses, optics . . . 133

A.3 Enriched categories . . . 133
A.4 Toposes and modalities . . . 134
A.5 Homotopy types and type theory . . . 134
A.6 Probability . . . 135
A.7 Four fixed-point theorems . . . 135

B Concordance with the Source **137**

C Computational Checks **140**

D Solutions to the Exercises **143**

References **161**

Preface

Game theory was written down in the language of sets and functions because, in 1944, that was the language there was. The result is a subject whose central objects—games—are routinely presented by data that the theory itself regards as arbitrary: utility numbers defined only up to affine rescaling, trees defined only up to a list of "inessential transformations", type spaces defined only up to the hierarchies of belief they induce. The working game theorist learns which features of a presentation are real and which are artefacts, and carries the distinction in her head. The premise of this book is that the distinction can be carried by the mathematics instead.

Three developments make this possible. The first is the theory of *open games* and optics, which gives games a compositional algebra: a game is a morphism, and large games are built from small ones by sequential and parallel composition. The second is the categorical understanding of *information*: knowledge as a modality arising from base change, belief as a monad, hierarchies of belief as a terminal coalgebra. The third is the homotopical habit of mind—made systematic in higher category theory and homotopy type theory—of asking of every object what it is *up to equivalence*, of every construction whether it is invariant, and of every existence problem what obstructs it.

What this book is

The book follows the arc of a first graduate course in game theory and information economics, the arc canonically drawn by Rasmusen's *Games and Information*: strategic and extensive forms; information; mixed strategies; dynamic and repeated games; incomplete information; moral hazard, adverse selection, mechanism design and signalling; bargaining, auctions and pricing. On each topic it does two things. It states the classical theory precisely, with proofs or references and at least one worked example. And it re-reads that theory in a language in which the structural content is explicit. Sometimes the re-reading is only a change of notation, and we say so. Sometimes it produces a statement that the classical formulation obscures; the main instances are these.

- *Solution concepts are functions on a type of games, and a type has an identity structure.* Invariance requirements (Kohlberg–Mertens), symmetry requirements (Harsanyi–Selten), the Markov restriction (Maskin–Tirole)

and the insensitivity of interim rationalizability to redundant types are one requirement, applied to four different types (Chapters 1, 3, 5 and 9).

- *Common knowledge is π_0; a common prior is the vanishing of a class in H^1.* The meet of the players' partitions is the set of connected components of the homotopy colimit of their type maps; that homotopy colimit is a graph, and its first cohomology is exactly the obstruction to gluing the players' posteriors into a common prior (Chapters 4 and 5).
- *Perfection is a necessity modality.* Subgame perfection is to subgames what knowledge is to states: the right adjoint of a restriction (Chapter 8).
- *Finite horizons are initial algebras; infinite horizons are terminal coalgebras.* Backward induction is a fold; the one-shot deviation principle and self-generation are coinduction principles; discounting is a guard. The folk theorem lives in the gap between least and greatest fixed points, and Rubinstein's uniqueness theorem is the statement that under a contraction there is no gap (Chapters 9 and 16).
- *Incentive compatibility is an enriched-presheaf condition.* A social choice function induces a generalized metric on the type space; incentive-compatible payment rules are exactly the enriched copresheaves; implementability is the absence of negative cycles; revenue equivalence says the incentive category is a groupoid (Chapter 14).

None of the underlying theorems is ours, and the status tags described on the copyright page are meant to keep that visible. What is offered is a unified development, a number of reformulations with their proofs, a set of worked examples, and—in Chapter 19—an assessment of where the higher-categorical language earns its keep and where it is, so far, only a dialect.

What "higher" means here

A candid remark, since the title invites the question. The structures used in this book are mostly of low categorical dimension: groupoids of games, graphs and their first cohomology, monoidal categories and bicategories of optics, categories enriched in truth values or in extended reals, presheaf toposes. Genuinely infinite-dimensional structure appears at the edges—in the spaces of games and equilibria of Chapter 6, in the simplicial models of many-agent knowledge of Chapter 4, and in the ambient ∞-toposes in which one would like to place smooth and probabilistic games. "Higher categorical semantics" is used here in the sense of the discipline rather than the dimension: objects are considered up to equivalence; equivalences are data; constructions are required to be invariant; quotients are taken homotopically, so that what an

ordinary quotient discards is retained as higher structure. That discipline turns out to organize a surprising amount of game theory. Whether it will prove theorems that the classical language could not is a question the last chapter takes up without sentiment.

Prerequisites and conventions

The reader should know what a category, a functor, an adjunction and a monad are, and should not be alarmed by a string diagram. No game theory is assumed, but the classical material is covered at speed. Appendix A collects the categorical background actually used: monoidal and enriched categories, presheaf toposes and modalities, the vocabulary of homotopy type theory, probability monads and Markov categories, (co)algebras, optics, and the four fixed-point theorems on which the subject runs.

Type-theoretic notation ($\sum$, $\prod$, identity types) is used informally, as in the first chapters of [HoTT13]; nothing depends on a particular proof assistant or a particular model. A normal-form game is written (N, A, u) with $A = \prod_{i \in N} A_i$; a profile is $a = (a_i, a_{-i})$. ΔX is the set of finitely supported probability distributions on X. Each chapter ends with notes on the literature and with exercises; none of the exercises is taken from another text, and Appendix D gives solutions to all of them.

Introduction

The method, and what is missing from it

The method of applied game theory can be stated in a sentence: specify the *rules of the game*—who moves, when, knowing what, with what consequences—and then compute what follows from them under a *solution concept*. The discipline of the method is that everything must be in the rules. An explanation that appeals to something outside the model is not an explanation in the theory.

Logicians will recognize the shape. The rules are *syntax*; the solution concept is a *semantics*; a model is good when the semantics of a simple piece of syntax matches an observed regularity. But the comparison also shows what the classical practice lacks. In logic, syntax is generated by a grammar, semantics is defined by recursion on the grammar, and there is a clear notion of when two pieces of syntax *mean the same thing*. In game theory the syntax is a picture of a tree or a table of numbers; the semantics is defined on the whole object at once; and the question of when two games are the same game has been answered case by case, by people with good taste, over several decades.

Categorical semantics is the systematic study of this shape, and this book applies it. We will ask the three questions a semanticist asks of any language.

1. *What are the types?* That is: what kind of thing is a game, a strategy, a belief, a contract; and when are two of them equal? Here the answer comes from homotopy type theory. A type is not a bare set but a space of things-with-identifications. The type of games is a groupoid (at least), and a solution concept is a function on it. Functions on a type respect its identifications; this single fact reproduces the invariance axioms of refinement theory, and says precisely which popular solution concepts are functions on which types.

2. *How do the terms compose?* Here the answer comes from applied category theory. Games with open boundaries are the morphisms of a symmetric monoidal category; the boundary data—what a game observes, what it does, what utility it receives and what "coutility" it sends back—is an object of a category of *optics*. The same optic structure organizes Bayesian updating and contracts.

3. *Where does the interpretation live?* Classical game theory is interpreted in sets. But a game of incomplete information is a game *varying over* a

space of states, and a dynamic game is a game *varying over* time. The natural homes for variable sets are toposes: presheaves on a space of states, presheaves on the ordinal ω. Their internal logic is modal, and its modalities are the knowledge operator of epistemic game theory and the "later" operator of guarded recursion.

One game, read three times

Take the most familiar object in the subject. Two players each choose C or D; payoffs are

	C	D
C	$2,2$	$0,3$
D	$3,0$	$1,1$

with the row player's payoff listed first.

First reading: an element of a type. The table is a presentation. Rescale either player's payoffs by a positive affine map and nothing a game theorist cares about changes; so the game is at best the equivalence class. But more is true: subtract 2 from the row player's payoffs in the column where her opponent plays C (and symmetrically), and one obtains

	C	D
C	$0,0$	$0,1$
D	$1,0$	$1,1$

in which every player's ranking of her own actions against every opponent action is unchanged—D is still strictly dominant—but mutual defection is now Pareto *efficient*. The two tables have the same best-response structure and different welfare structure. So there are (at least) two types here: games up to *best-response equivalence*, on which Nash equilibrium, dominance and rationalizability are well-defined functions; and games up to *preference equivalence*, a finer type, on which efficiency is also defined. "The Prisoner's Dilemma" is a point of the finer type whose image in the coarser one is unremarkable. The dilemma is not in the strategic structure. Chapter 1 makes this precise.

Second reading: a composite. Each player is a small open system: a box with an output wire carrying her action and an input wire carrying her payoff back. The game is the two boxes side by side, capped by a box that turns the pair of actions into the pair of payoffs (Figure 1). Nothing here is special to two players or to this payoff function. Replace the cap by a box that feeds the actions into a market, or into a second stage of play, and one has a different game built from the same parts. Chapter 2 develops this algebra; the equilibrium condition is defined once, for a single box, and composed.

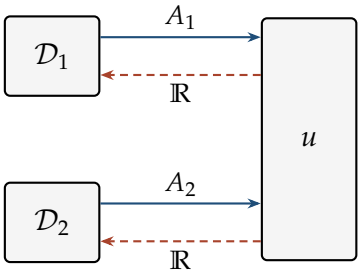


Figure 1: A two-player simultaneous game as a composite of open games: two decisions in parallel, closed off by a payoff function. Solid wires carry moves forward; dashed wires carry utility back.

Third reading: an object varying over time. Repeat the game forever with discount factor δ. A strategy is now a machine: a set of states, an output function, a transition function. Machines are coalgebras; the space of all strategies is a terminal coalgebra; and the equilibrium payoff set is the greatest fixed point of a monotone operator on sets of payoff vectors. That one sentence already contains the automaton representation of strategies, the complexity measure of Kalai and Stanford, the self-generation theorem of Abreu, Pearce and Stacchetti, and the reason the finitely repeated game—an *initial* algebra, solved by recursion from the end—behaves so differently (Chapter 9).

A dictionary

The table below is a map of the book. The left column is the classical vocabulary; the right column is what the corresponding chapter claims that vocabulary was about.

Classical notion	Structure	Ch.
strategic equivalence; invariance	identity type of the type of games	1, 3
focal point	structure that breaks an automorphism	1
simultaneous / sequential play	monoidal product / composition of open games	2
game tree	presheaf on ω	3
perfect recall	the information quotient is a map of forests	3
knowledge operator	comonad $p^*\forall_p$ of base change	4
common knowledge	π_0 of a homotopy colimit	4
universal type space	terminal coalgebra	5
common prior	trivial class in H^1	5
mixed extension	free convex algebra	6
existence of equilibrium	a projection of degree one	6
strategic complements	enrichment in a lattice; Tarski	7
subgame perfection	necessity modality over subgames	8
backward induction	fold over an initial algebra	8
discounting	guard; contraction	9
self-generating set	post-fixed point; coinduction	9
Markov strategy	factors through the bisimulation quotient	9
consistent off-path beliefs	point of a blow-up of the simplex	10
contract	costate of an optic	11
sufficient statistic; garbling	factorization in a Markov category	11
market unraveling	iteration to a fixed point of a monotone map	13
incentive-compatible transfer	enriched copresheaf	14
indirect utility	conjugate under a profunctor	14
Bayes plausibility	fibre of the monad multiplication	15
Rubinstein uniqueness	unique guarded fixed point	16

How to read

Part I (Chapter 1–10) is the theory of games; Part II (Chapter 11–15) is the economics of asymmetric information; Part III (Chapter 16–18) contains applications; Part IV is a single chapter of assessment. Chapters 1 to 5 are the foundation and should be read in order. After them the chapters are largely independent, with two exceptions: Chapter 15 and Part III lean on Chapter 14,

and Chapter 16 on Chapter 9. A reader who wants the classical theory only can read the statements tagged classical and the examples; a reader who knows the classical theory can read the rest.

Part I

Games

Chapter 1

The Rules of the Game

A game is specified by its rules: the players, what each can do, what each knows when doing it, and how each ranks the consequences. This chapter treats the simplest case, in which nobody learns anything before acting, so that the only rules are players, actions and payoffs. Its real subject, however, is the question that the simplest case already raises: when are two specifications *the same game*? The answer organizes the rest of the book.

1.1 Decisions

A game is several decision problems that share variables, so we begin with one.

Definition 1.1. A ***decision problem*** is a set X of alternatives together with a map $k\colon X \to R$ into a set R of outcomes. A ***selection function*** for (X, R) is a map

$$\varepsilon\colon (X \to R) \longrightarrow \mathcal{P}(X)$$

assigning to each problem k the set $\varepsilon(k) \subseteq X$ of alternatives the agent is willing to choose. When R carries a total preorder, the ***maximizing*** selection function is $\arg\max(k) = \{x \mid k(x) \geq k(x') \text{ for all } x'\}$.

The map k is the agent's *context*: everything about the world that matters to her, compressed into the consequence of each alternative. Selection functions were isolated by Escardó and Oliva [EO10, EO11], in single-valued form $(X \to R) \to X$, as the computational content of choice. Separating the agent ε from the context k pays off twice. It lets a context be supplied by other agents, which is what a game is. And it does not commit us to maximization: an agent who wants to be at the median, or to pick whatever a majority picks, is a perfectly good selection function, though no utility function on X represents her [HOSWZ17].

1.2 Games in strategic form, and their types

Definition 1.2. A (finite) ***game in strategic form*** consists of a finite set N of players; for each $i \in N$ a finite nonempty set A_i of actions; and for each i a

payoff function $u_i \colon A \to \mathbb{R}$, where $A = \prod_{j \in N} A_j$. In type-theoretic notation the type of such games is

$$\mathsf{NFG}_u \;:\equiv \sum_{N:\mathsf{FinSet}} \; \sum_{A:N \to \mathsf{FinSet}_+} \; \prod_{i:N} \Big(\prod_{j:N} A_j \to \mathbb{R} \Big).$$

Writing the definition as a type is not decoration. A type comes with its identity types, and the identity type of a type of structures is forced.

Lemma 1.3 (Identity of games). [literature; [HoTT13, §2.14, §9.8], [ANST21]] *For $G = (N, A, u)$ and $G' = (N', A', u')$ in NFG_u, the identity type $G = G'$ is equivalent to the type of triples (ν, α, θ) where $\nu \colon N \simeq N'$ is a bijection of players, $\alpha_i \colon A_i \simeq A'_{\nu(i)}$ is a family of bijections of actions, and θ witnesses $u'_{\nu(i)} \circ \alpha = u_i$ for all i.*

So an identification of games is an isomorphism: a relabelling of players and actions under which payoffs agree on the nose. In particular NFG_u is not a set but a groupoid, and the loop space at a game G is its automorphism group $\mathrm{Aut}(G)$. We return to this in Section 1.5.

The identity type of NFG_u is the right one for the data, and the wrong one for the subject, because payoff numbers are over-specified. Two coarser notions of sameness are in constant informal use.

Definition 1.4. Let u, u' be payoff profiles on the same game form (N, A).

(a) u and u' are ***preference equivalent*** if for each i there are $\alpha_i > 0$ and $\beta_i \in \mathbb{R}$ with $u'_i = \alpha_i u_i + \beta_i$.

(b) u and u' are ***incentive equivalent*** if for each i there are $\alpha_i > 0$ and a function $\beta_i \colon A_{-i} \to \mathbb{R}$ with $u'_i(a) = \alpha_i u_i(a) + \beta_i(a_{-i})$.

We write $\mathsf{NFG}_{\succsim}$ and $\mathsf{NFG}_{\mathrm{inc}}$ for the types of game forms equipped with a payoff profile up to (a), respectively (b), so that there are quotient maps $\mathsf{NFG}_u \twoheadrightarrow \mathsf{NFG}_{\succsim} \twoheadrightarrow \mathsf{NFG}_{\mathrm{inc}}$.

Preference equivalence is the indeterminacy of von Neumann–Morgenstern utility: by the expected-utility theorem a point of $\mathsf{NFG}_{\succsim}$ is a game form with, for each player, a preference relation over lotteries on A satisfying independence and continuity (Chapter 6). Presenting it by numbers is a choice of coordinates. Incentive equivalence is coarser. It remembers only how each player ranks her own (mixed) actions against each conjecture about the others; adding $\beta_i(a_{-i})$ changes how much player i likes what the *others* do, which she cannot affect. Morris and Ui [MU04] study this and neighbouring notions under the names *von Neumann–Morgenstern equivalence* and *best-response equivalence*.

Principle 1.5 (Solution concepts are functions on types)**.** *A solution concept on a type $\mathcal{T}$ of games is a dependent function*

$$S\colon \prod_{G:\mathcal{T}} \mathcal{P}\big(\Sigma(G)\big),$$

where $\Sigma(G)$ is the set of strategy profiles of G. Consequently S respects every identification in $\mathcal{T}$: if $p\colon G = G'$ then $p_*S(G) = S(G')$. A rule defined on presentations defines a solution concept on $\mathcal{T}$ exactly when it is invariant under the equivalence that $\mathcal{T}$ imposes.

This is not an axiom about rational play; it is what the word "function" means once the domain is taken seriously. Its force comes from choosing the domain. Which concepts live on which type?

Proposition 1.6. [observation]

(i) Strict dominance, rationalizability, Nash equilibrium, correlated equilibrium and (for 2×2 games) risk dominance are well defined on $\mathsf{NFG}_{\mathrm{inc}}$.

(ii) Pareto efficiency and payoff dominance are well defined on $\mathsf{NFG}_{\simeq}$ and do not descend to $\mathsf{NFG}_{\mathrm{inc}}$.

Proof. (i) Each of these is defined through differences $u_i(a_i, \mu) - u_i(a_i', \mu)$ of a player's expected payoff between two of her own actions against a fixed distribution μ over the others' actions, and only through the signs of such differences, or—for risk dominance—through comparisons of products over players of such differences. A transformation $u_i \mapsto \alpha_i u_i + \beta_i(a_{-i})$ multiplies every such difference by $\alpha_i > 0$. (ii) Invariance under positive affine maps is clear. For the failure, take the Prisoner's Dilemma of the Introduction and subtract from each player's payoff 2 whenever the opponent plays C. The result is incentive equivalent to the original, but its unique equilibrium (D, D), with payoffs $(1, 1)$, is now Pareto efficient and payoff-dominates (C, C), with payoffs $(0, 0)$. □

The moral is worth stating plainly. The Prisoner's Dilemma is *strategically* trivial—each player has a strictly dominant action—and what makes it a dilemma is extra structure that the strategic analysis never consults. Game theory's most famous example is famous for a feature that is invisible to every solution concept in part (i). Keeping two types apart keeps two kinds of statement apart: predictions (which live on $\mathsf{NFG}_{\mathrm{inc}}$) and evaluations (which need $\mathsf{NFG}_{\simeq}$).

Remark 1.7 (Moduli). In geometric language NFG_u is a moduli stack of games: its points are games, its paths are isomorphisms, and its set of connected

components π_0 is the set of isomorphism classes with which set-level game theory works. A solution concept is a section of the bundle of strategy profiles over the stack. What the stack remembers and the coarse space forgets is the stabilizer $\mathrm{Aut}(G)$ at each point; Section 1.5 shows that this is exactly the information that the theory of focal points is about.

1.3 Dominance as rewriting

Definition 1.8. In a game G, an action a_i is ***strictly dominated*** if some mixed action $\sigma_i \in \Delta A_i$ has $u_i(\sigma_i, a_{-i}) > u_i(a_i, a_{-i})$ for every a_{-i}, and ***weakly dominated*** if the inequality is weak for every a_{-i} and strict for some. A ***restriction*** of G is a product subset $R = \prod_i R_i \subseteq A$ with all R_i nonempty; it is a game $G|_R$ in its own right.

In the Prisoner's Dilemma, D strictly dominates C for both players, and the analysis is over: no assumption about what a player believes of the other is needed. When domination is not immediate one iterates. Write $R \to R'$ if R' is obtained from R by deleting one action that is strictly dominated in $G|_R$. This is an abstract rewriting system on the finite poset of restrictions.

Theorem 1.9 (Order independence). [classical; [GKZ90, Apt04]] *In a finite game the rewriting system $\to$ is terminating and confluent: every maximal sequence of deletions of strictly dominated actions ends at the same restriction R^∞.*

Proof sketch. Termination is finiteness. For confluence it suffices, by Newman's lemma, to check local confluence: if a_i and b_j are each strictly dominated in $G|_R$, then a_i is still strictly dominated after b_j is deleted. If $j \neq i$ the dominating mixture for a_i still dominates on the smaller set of opponent profiles. If $j = i$ and the mixture σ_i dominating a_i puts weight on b_i, replace that weight by the mixture dominating b_i; since the game is finite, repeating the substitution a bounded number of times yields a dominating mixture supported off b_i. □

For weak dominance confluence fails, and with it the claim to be a solution concept in the sense of Principle 1.5: the output depends on the order of deletion, which is not part of the game.

Example 1.10 (Weak dominance is order dependent)**.** In the game

	L	R
T	$1,1$	$0,0$
M	$1,1$	$2,1$
B	$0,0$	$2,1$

both T and B are weakly dominated by M. Deleting T first makes L weakly dominated by R, and the process ends at $\{M, B\} \times \{R\}$ with payoffs $(2, 1)$. Deleting B first makes R weakly dominated by L, ending at $\{T, M\} \times \{L\}$ with payoffs $(1, 1)$.

Strict dominance has a second description that will recur.

Theorem 1.11 (Rationalizability)**.** [classical; [Ber84, Pea84]] *For a restriction R let $\Phi(R) = \prod_i \Phi_i(R)$, where $\Phi_i(R)$ is the set of actions $a_i \in A_i$ that are a best response to some conjecture $\mu \in \Delta(R_{-i})$. Then Φ is monotone; in a finite game its greatest fixed point $\nu\Phi$ exists, equals $\bigcap_n \Phi^n(A)$, and coincides with the set R^∞ of Theorem 1.9.*

The coincidence rests on a separation argument: an action is a best response to no (possibly correlated) conjecture exactly when a mixture strictly dominates it [Pea84]. The first two claims are the Knaster–Tarski theorem on the finite lattice of product subsets. The elements of $\nu\Phi$ are the ***rationalizable*** actions, and the fixed-point description gives a proof principle:

> *Coinduction.* To show that a_i is rationalizable, exhibit a product set $R \ni a_i$ with $R \subseteq \Phi(R)$.

Such an R is a set of actions each justified by conjectures supported inside R—a closed circle of reasons. Rationalizability is therefore not the outcome of a deletion procedure that happens to converge; it is a *coinductive* predicate, the largest self-justifying set, and the procedure is the standard way of computing a greatest fixed point from above. The same pattern—greatest fixed point of a monotone operator, proved by exhibiting a post-fixed point—returns as self-generation in repeated games (Chapter 9) and as common knowledge (Chapter 4).

1.4 Nash equilibrium

Definition 1.12. For a game G let $\mathrm{BR}\colon A \to \mathcal{P}(A)$ be the ***best-response correspondence***, $\mathrm{BR}(a) = \prod_i \arg\max u_i(-, a_{-i})$. A profile a is a (pure) ***Nash equilibrium*** if $a \in \mathrm{BR}(a)$. Thus

$$\mathrm{NE}(G) = \mathrm{Fix}(\mathrm{BR}),$$

the pullback of the graph of BR along the diagonal $A \to A \times A$.

In the language of Definition 1.1: each player is a selection function ε_i; a profile a supplies each player with the context $k_i = u_i(-, a_{-i})$; and a is an equilibrium when each a_i is selected in the context that the others' choices

create, $a_i \in \varepsilon_i(u_i(-, a_{-i}))$. Nothing in this formulation uses maximization, and replacing arg max by other selection functions yields the *selection equilibria* of [HOSWZ17].

Example 1.13 (Three small games)**.**

	C	D
C	$2,2$	$0,3$
D	$3,0$	$1,1$

	O	F
O	$2,1$	$0,0$
F	$0,0$	$1,2$

	a	b
a	$2,2$	$0,0$
b	$0,0$	$1,1$

The first is the Prisoner's Dilemma: $\mathrm{NE} = \{(D, D)\}$, found already by dominance. The second is a *Battle of the Sexes*: no action is dominated, every action is rationalizable, and $\mathrm{NE} = \{(O, O), (F, F)\}$ in pure strategies. The third is a *ranked coordination* game with $\mathrm{NE} = \{(a, a), (b, b)\}$. Both of the latter have a third equilibrium in mixed strategies (Chapter 6).

Rationalizability asks only that each action be justified by *some* conjecture consistent with common knowledge of rationality; equilibrium asks in addition that conjectures be *correct*. The second requirement is not a consequence of rationality, and the examples show its cost: where dominance gives nothing, equilibrium gives several answers. The rest of this chapter is about what, if anything, the game itself says about choosing among them.

1.5 Symmetry and focal points

Let G be a game. By Lemma 1.3 the loops at G in NFG_u are the automorphisms: relabellings (ν, α) of players and actions preserving payoffs. $\mathrm{Aut}(G)$ acts on profiles, and since it preserves best responses it acts on $\mathrm{NE}(G)$ and on the mixed equilibria.

Proposition 1.14 (Equivariance)**.** [observation] *Let $\mathcal{T}$ be a type of games and $s \colon \prod_{G:\mathcal{T}} \Sigma(G)$ a single-valued solution: a function assigning to each game one (mixed) profile. Then for every G the profile $s(G)$ is fixed by every automorphism of G in $\mathcal{T}$.*

Proof. An automorphism is a path $p \colon G = G$. Dependent functions commute with transport: $p_* \, s(G) = s(G)$. □

The classical counterpart is the symmetry-invariance axiom of Harsanyi and Selten [HS88], and the observation that it constrains selection is due to them and to Crawford and Haller [CH90]; see [AK13] for a systematic account. What the type-theoretic phrasing adds is that the axiom is not optional. One cannot write down a selection that violates it without consulting something that is not in the game.

Example 1.15 (No selection without structure)**.** In the pure coordination game with payoffs $(1,1)$ on the diagonal and $(0,0)$ off it, the relabelling that swaps both players' action names is an automorphism, and it exchanges the two pure equilibria. The only equilibrium fixed by $\mathrm{Aut}(G)$ is the mixed one in which each player randomizes evenly, with expected payoff $\frac{1}{2}$. With m actions the invariant equilibrium is uniform and pays $1/m$. In the Battle of the Sexes of Example 1.13, exchanging the players while exchanging O and F is an automorphism swapping the pure equilibria; again only the mixed equilibrium is invariant. In the ranked coordination game no automorphism moves (a,a) to (b,b), so each equilibrium is individually invariant and a rule such as "play the risk-dominant equilibrium" is a legitimate function on $\mathsf{NFG}_{\mathrm{inc}}$.

People asked to meet a stranger somewhere in a city without communicating do far better than $1/m$; this is Schelling's observation [Sch60], and his name for what they use is a ***focal point***. The proposition tells us what kind of thing a focal point must be. It cannot be a property of an equilibrium of G, because $\mathrm{Aut}(G)$ permutes the candidates transitively. It is additional *structure*: a labelling λ of the actions by descriptions that the players share and that are not all alike—"the clock at the main station", "noon"—so that the game actually played is a point (G,λ) of a different type, of ***framed games***, whose automorphism group $\mathrm{Aut}(G,\lambda) \subseteq \mathrm{Aut}(G)$ is smaller. In (G,λ) the salient equilibrium is invariant, and can be selected by a function.

Remark 1.16. This reading agrees with the modern formal treatments of focal points, in which frames or labels are added to the strategic form and symmetries are restricted accordingly [AK13]. It also explains a persistent feature of the literature: every attempt to derive focal points from the unframed game re-imports a frame under another name (a tie-breaking order on actions, a "natural" indexing of players). An order on a set is exactly a structure that kills its automorphisms.

Notes

The vocabulary of players, actions, payoffs and information is standard; the general references used throughout are [OR94, FT91, Mye91, MSZ13], and the source that this book shadows is [Ras06]. Strategic equivalence in the affine sense goes back to [vNM44]. Order independence for strict dominance in finite games is due to Gilboa, Kalai and Zemel [GKZ90]; Apt [Apt04] gives the rewriting-theoretic proof used here and catalogues which elimination procedures are confluent. Rationalizability was introduced by Bernheim [Ber84] and Pearce [Pea84]; its coinductive character is implicit in their "best-response property". Nash equilibrium is from [Nas50a, Nas51]. Higher-order games

and selection equilibria are developed in [HOSWZ17]. The structure identity principle that yields Lemma 1.3 is discussed in [HoTT13] and in full generality in [ANST21].

Exercises

Exercise 1.1. Show that incentive equivalence preserves the set of correlated equilibria (defined in Chapter 6), and give an example showing that it does not preserve the set of Pareto-efficient profiles other than the one in the text.

Exercise 1.2. Let G be a symmetric two-player game with action set $\{1, \dots, m\}$ and payoff 1 if the actions agree and 0 otherwise. Compute $\mathrm{Aut}(G)$. Now add a frame in which action 1 is distinguished and the others are indistinguishable. Compute the automorphism group of the framed game and list its invariant equilibria.

Exercise 1.3. Give a three-action example in which an action is not strictly dominated by any pure action but is strictly dominated by a mixture, and verify directly that it is a best response to no conjecture.

Exercise 1.4. An agent chooses $x \in \{1, \dots, 9\}$ and wishes to be as close as possible to two-thirds of the average of all choices. Write her as a selection function on contexts $k \colon X \to \mathbb{R}$, where k returns the resulting average. Show that her selection function is not $\arg\max$ of any fixed utility on X, and find the selection equilibria of the n-player game.

Exercise 1.5. Prove that the operator Φ of Theorem 1.11 is monotone, and that every Nash equilibrium is contained in $\nu\Phi$, by exhibiting a post-fixed point.

Exercise 1.6. Show that deleting *all* weakly dominated actions of all players simultaneously at each round is a well-defined procedure, invariant under incentive equivalence, and hence a solution concept on $\mathsf{NFG}_{\mathrm{inc}}$ in the sense of Principle 1.5. Compute its output on Example 1.10 and compare it with the two sequential orders. What has been bought, and at what price?

Chapter 2

Open Games

A game in strategic form is a closed system. There is no operation that takes two such games and plugs the outcome of one into the payoffs of the other, or runs one after the other with the second player watching the first; each new situation is modelled from scratch. Yet informally games are built from parts all the time—"a bargaining stage followed by a production stage", "the same market in each of two periods". This chapter presents the calculus, due to Ghani, Hedges, Winschel and Zahn [GHWZ18, Hed16], in which that informal practice is literal: a game with open boundaries is a morphism, and games compose.

2.1 Lenses

The boundary of an open game carries information in two directions. Forwards flow observations and moves; backwards flows value. The algebra of such two-way boundaries is the category of lenses.

Definition 2.1. A ***lens*** $(X, S) \to (Y, R)$ between pairs of sets is a pair $\lambda = (v, c)$ of functions

$$v \colon X \to Y \qquad \text{(view, or play)}, \qquad c \colon X \times R \to S \qquad \text{(update, or coplay)}.$$

The composite of $(v, c) \colon (X, S) \to (Y, R)$ and $(v', c') \colon (Y, R) \to (Z, Q)$ has view $v' \circ v$ and update $(x, q) \mapsto c\big(x, c'(v(x), q)\big)$. Lenses form a symmetric monoidal category Lens with $(X, S) \otimes (X', S') = (X \times X', S \times S')$ and unit $I = (1, 1)$.

Two families of lenses matter immediately. A lens $I \to (X, S)$ is an element $x \in X$: a *state*, or history. A lens $(Y, R) \to I$ is a function $k \colon Y \to R$: a *costate*—which is precisely a decision problem in the sense of Definition 1.1, a context assigning to each move its consequence. Composing a lens $\lambda \colon (X, S) \to (Y, R)$ with a costate k gives a costate $k \circ \lambda \colon x \mapsto c(x, k(v(x)))$ on (X, S): lenses transport contexts backwards. This is the whole idea. A piece of a game is a process that turns what it sees into what it does, and turns the value of what it does into the value of what it saw.

2.2 Open games

Definition 2.2 (Open game). [literature; [GHWZ18]; predicate form as in [BHZ19]] An ***open game*** $\mathcal{G}\colon (X,S) \to (Y,R)$ consists of

- a set $\Sigma_{\mathcal{G}}$ of strategy profiles;
- for each $\sigma \in \Sigma_{\mathcal{G}}$ a lens $\mathcal{G}_\sigma = (P_\sigma, C_\sigma)\colon (X,S) \to (Y,R)$, the *play* and *coplay* under σ;
- an ***equilibrium predicate*** $E_{\mathcal{G}}\colon X \times (Y \to R) \to \mathcal{P}(\Sigma_{\mathcal{G}})$, assigning to each context (x,k)—a state of the source and a costate of the target—the set of profiles that are in equilibrium in that context.

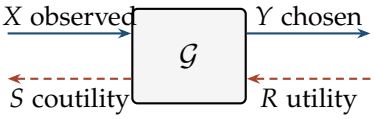


Figure 2.1: An open game $\mathcal{G}\colon (X,S) \to (Y,R)$.

The four boundary sets are read as in Figure 2.1: X is what the game observes of the past, Y what it contributes to the future, R the utility it receives from the future and S the *coutility* it passes back to the past. The original definition records a best-response relation on $\Sigma_{\mathcal{G}}$ in each context rather than its set of fixed points; nothing below depends on the difference.

Definition 2.3 (Composition). Let $\mathcal{G}\colon (X,S) \to (Y,R)$ and $\mathcal{H}\colon (Y,R) \to (Z,Q)$. The ***sequential composite*** $\mathcal{H} \circ \mathcal{G}$ has $\Sigma = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$, lenses $\mathcal{H}_\tau \circ \mathcal{G}_\sigma$, and

$$(\sigma,\tau) \in E_{\mathcal{H}\circ\mathcal{G}}(x,k) \iff \begin{cases} \sigma \in E_{\mathcal{G}}\big(x,\, k \circ \mathcal{H}_\tau\big) \text{ and} \\ \tau \in E_{\mathcal{H}}\big(P_\sigma(x),\, k\big). \end{cases}$$

For $\mathcal{G}\colon (X,S) \to (Y,R)$ and $\mathcal{G}'\colon (X',S') \to (Y',R')$ the ***parallel composite*** $\mathcal{G} \otimes \mathcal{G}'$ has $\Sigma = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}'}$, lenses $\mathcal{G}_\sigma \otimes \mathcal{G}'_{\sigma'}$, and

$$(\sigma,\sigma') \in E_{\mathcal{G}\otimes\mathcal{G}'}((x,x'),k) \iff \sigma \in E_{\mathcal{G}}(x,k_1) \text{ and } \sigma' \in E_{\mathcal{G}'}(x',k_2),$$

where $k_1(y) = \pi_1 k\big(y, P'_{\sigma'}(x')\big)$ and $k_2(y') = \pi_2 k\big(P_\sigma(x), y'\big)$.

Read the two clauses of sequential composition as a sentence. The earlier game $\mathcal{G}$ is in equilibrium when it is so in the context in which its future is *the later game played according to* τ; the later game $\mathcal{H}$ is in equilibrium when it is so at the history that σ actually produces. In parallel composition each component faces a context in which the other's move is held at its equilibrium value. Unilateral deviation, the defining idea of Nash equilibrium, is thus not an axiom about profiles; it is *what composition of contexts means*.

Theorem 2.4. [literature; [GHWZ18]] *Open games, identified up to bijections of their strategy sets compatible with the remaining data, are the morphisms of a symmetric monoidal category* Game *whose objects are pairs of sets. Every lens is an open game with* $\Sigma = 1$ *and trivially true equilibrium predicate, giving a symmetric monoidal functor* Lens $\to$ Game.

2.3 Atoms and the strategic form

Two kinds of atom generate the examples of this book.

Definition 2.5. (a) The ***decision*** $\mathcal{D}_{X,Y} \colon (X,1) \to (Y,\mathbb{R})$ of an agent who observes X, chooses in Y and maximizes has $\Sigma = (X \to Y)$, play $P_\sigma(x) = \sigma(x)$, trivial coplay, and

$$E(x,k) = \{\sigma \mid \sigma(x) \in \arg\max k\}.$$

More generally any selection function ε on (Y,R) gives a decision with $E(x,k) = \{\sigma \mid \sigma(x) \in \varepsilon(k)\}$.

(b) A function $f \colon X \to Y$ gives the ***forward*** game $(X,1) \to (Y,1)$, and a function $u \colon Y \to R$ gives the ***payoff*** game $u \colon (Y,R) \to I$, the lens with coplay $y \mapsto u(y)$.

Proposition 2.6 (Strategic form)**.** [literature; [GHWZ18]] *Let* $G = (N, A, u)$ *be a game in strategic form with* $N = \{1, \dots, n\}$*. The closed game*

$$u \circ \left(\mathcal{D}_{1,A_1} \otimes \cdots \otimes \mathcal{D}_{1,A_n}\right) \colon I \to I$$

has strategy set A *and equilibrium set* $\mathrm{NE}(G)$.

Proof. For $n = 2$: by Definition 2.3, (a_1, a_2) is an equilibrium iff it is one for $\mathcal{D}_1 \otimes \mathcal{D}_2$ in the context $k = u$, iff $a_1 \in \arg\max \pi_1 u(-, a_2)$ and $a_2 \in \arg\max \pi_2 u(a_1, -)$. The general case is the same. □

This is Figure 1. The proposition is not deep; its content is that the classical definition has been *derived* from two generic operations and one atom, so that it is available wherever those operations are.

Example 2.7 (Sequential play, and a warning)**.** Let player 1 choose $a_1 \in A_1$ and player 2, observing a_1, choose $a_2 \in A_2$, with payoffs $u \colon A_1 \times A_2 \to \mathbb{R}^2$. The diagram is Figure 2.2: $\mathcal{D}_{1,A_1}$, then a copy of a_1 alongside $\mathcal{D}_{A_1,A_2}$, then u. The strategy set is $A_1 \times (A_1 \to A_2)$, and unwinding Definition 2.3,

$$(a_1, s) \text{ is an equilibrium } \iff \begin{cases} a_1 \in \arg\max_a u_1(a, s(a)) \text{ and} \\ s(a_1) \in \arg\max u_2(a_1, -). \end{cases}$$

Player 1 optimizes against player 2's whole plan s; player 2 optimizes *at the history that occurs*. At histories $a \neq a_1$ the plan s is unconstrained. These are the Nash equilibria of the sequential game, including those sustained by threats that player 2 would not carry out. Sequential composition as defined here does not by itself deliver subgame perfection. The thesis [Hed16] recovers subgame perfect equilibria of perfect-information games by a further construction, [Hed18] by a notion of state for a double category of open games, and [GKLF18] for iterated games; in Chapter 8 we will see perfection as a modality applied on top of the equilibrium predicate.

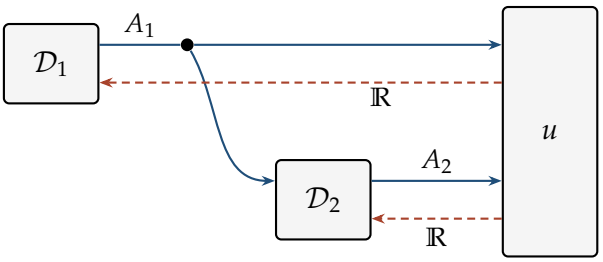


Figure 2.2: A two-stage game: the second decision observes a copy of the first move.

2.4 Selection functions and backward induction

For games of perfect information there is a second compositional structure, older than open games and in one respect sharper. Fix a set R of outcomes and consider single-valued selection functions.

Definition 2.8. [literature; [EO10, EO11]] Let $J_R X = (X \to R) \to X$ and $K_R X = (X \to R) \to R$. A selection function $\varepsilon \in J_R X$ has an associated *quantifier* $\bar{\varepsilon} \in K_R X$, $\bar{\varepsilon}(k) = k(\varepsilon(k))$. The ***product*** of $\varepsilon \in J_R X$ and $\delta \in J_R Y$ is $\varepsilon \otimes \delta \in J_R(X \times Y)$,

$$(\varepsilon \otimes \delta)(q) = (a, b(a)), \qquad b(x) = \delta\big(q(x,-)\big), \quad a = \varepsilon\big(x \mapsto q(x, b(x))\big).$$

The definition is backward induction for two moves: b is the second mover's optimal reply to each first move, and a is the first mover's optimal choice given that reply rule. J_R and K_R are strong monads (the latter is the continuation monad), $\varepsilon \mapsto \bar{\varepsilon}$ is a monad morphism, and $\otimes$ is the monoidal structure of J_R.

Theorem 2.9 (Backward induction is a product). [literature; [EO11]] *Consider a sequential game of perfect information in which player $t = 1, \dots, T$ chooses $x_t \in X_t$ after observing $x_1, \dots, x_{t-1}$, with outcome $q(x_1, \dots, x_T) \in R$, where player t is*

described by a selection function $\varepsilon_t \in J_R X_t$ (for instance $\arg\max$ *with respect to her own coordinate of $R = \mathbb{R}^T$, with a tie-breaking rule). Then*

$$(\varepsilon_1 \otimes \cdots \otimes \varepsilon_T)(q)$$

is the play of a subgame perfect equilibrium, and the equilibrium strategies are obtained from the same formula applied to the subgames.

For history-dependent selection functions the product is dependent, and the theorem is unchanged. What is striking is the infinite case. If each X_t is *searchable*—admits a selection function for every continuous predicate—and q is continuous in the sense of depending on finitely many moves, then the infinite product $\bigotimes_t \varepsilon_t$ exists and computes optimal play in a game with no last move [EO10]. The construction is a form of bar recursion, and it is at the same time a computational proof of the countable Tychonoff theorem: a product of searchable sets is searchable. We will meet the continuity hypothesis again in Chapter 9, where it is called *continuity at infinity*, and the searchability hypothesis in Chapter 8, where its failure is called the *open-set problem*.

2.5 Optics, parameters, and the higher structure

Lenses are the cartesian case of a general construction. Let $(\mathcal{C}, \otimes)$ be a symmetric monoidal category.

Definition 2.10. [literature; [Ril18, CEGLMPR20]] An ***optic*** $(X, S) \to (Y, R)$ in $\mathcal{C}$ is an element of the coend

$$\mathsf{Optic}_{\mathcal{C}}\big((X,S),(Y,R)\big) = \int^{M \in \mathcal{C}} \mathcal{C}(X, M \otimes Y) \times \mathcal{C}(M \otimes R, S).$$

An optic is represented by a forward map that produces a move together with a *residual* M, and a backward map that consumes the residual with the returned utility. The coend identifies representatives that differ by sliding a map of residuals from one side to the other. When $\mathcal{C}$ is cartesian the residual can always be taken to be X itself, and optics reduce to lenses. When $\mathcal{C}$ is a category of stochastic maps (Chapter 5) no such reduction is available—one cannot copy a random variable without correlating the copies—and it is optics, not lenses, that support *Bayesian open games* [BHZ19], in which the forward pass is a Markov kernel and equilibrium is Bayesian Nash equilibrium.

Three features of this more general setting are "higher" in the sense of the Preface.

Equality of optics is a quotient with content. The coend is the set of connected components of a category whose objects are representatives (M, ℓ, r) and

whose morphisms are maps of residuals. Two implementations of a game component with different internal memory are the same optic when they are connected by a zig-zag of such maps. The classical theory has no place for this identification because it has no place for internal memory.

Strategies are parameters, and parameters form a bicategory. For a monoidal category $\mathcal{C}$, $\mathsf{Para}(\mathcal{C})$ has morphisms $X \to Y$ given by a parameter object Π and a map $\Pi \otimes X \to Y$, and 2-cells given by reparametrizations $\Pi' \to \Pi$. An open game is a lens parametrised by its strategy set together with a rule saying which parameters are acceptable. Capucci, Gavranović, Hedges and Rischel [CGHR22] make this the definition: an *arena* is a parametrised optic, the players are attached to the parameter boundary, and the equilibrium predicate is a selection function there. The reformulation separates the *mechanics* of a game (the arena, which composes on the nose) from its *agency* (who controls which parameters, and what they want).

Open games have morphisms. Hedges [Hed18] constructs a symmetric monoidal double category whose horizontal 1-cells are open games, whose vertical 1-cells are lenses and whose 2-cells are morphisms of open games; the identity type of the type of open games is then the type of invertible 2-cells, and Principle 1.5 applies to it.

2.6 What composition does not do

It is important not to oversell. Definition 2.3 makes the equilibrium *predicate* compositional: whether (σ, τ) is an equilibrium of $\mathcal{H} \circ \mathcal{G}$ is decided by consulting $E_{\mathcal{G}}$ and $E_{\mathcal{H}}$. It does not make the equilibrium *set* of a composite a function of the equilibrium sets of its parts in any fixed contexts, because the context in which $\mathcal{G}$ is consulted depends on τ. To compose, one must know each part's behaviour in *every* context.

Seen from semantics this is as it should be: the meaning of an open term is its behaviour under all environments, and game theory itself insists that a strategy is a complete contingent plan. But it has two consequences that the reader should hold on to.

- Open games give compositional *modelling* and compositional *equilibrium checking*. They do not give compositional *solving*, and cannot: computing a Nash equilibrium is hard for reasons that have nothing to do with notation (Chapter 6). The software built on this theory is, accordingly, an equilibrium checker that is fed candidate profiles [OGE].
- The equilibria of a composite can differ in kind from those of its parts. Multiplicity, coordination failure and the need for refinement are emergent. A compositional language describes emergence precisely; it does

not abolish it.

Notes

Open games were introduced in [GHWZ18] and in Hedges's thesis [Hed16]; the survey-style account closest to the presentation here is [BHZ19], which also introduces contexts for optics. Lenses come from the theory of bidirectional transformations; their systematic generalization to optics is [Ril18, CEGLMPR20]. Alternative axiomatizations that make the proof of Theorem 2.4 itself compositional are given in [AGGKLN21], and the translation from extensive-form games into open games with agency is [CGLN22]. Selection functions and their products are due to Escardó and Oliva [EO10, EO11]; [HOSWZ17] uses them to define equilibria for agents who do not maximize. A metric refinement, in which equilibrium predicates are relaxed by an ϵ budget compatibly with composition, is [Gha25]. Earlier categorical treatments of games, which take games as objects rather than morphisms, are surveyed in [TV23].

Exercises

Exercise 2.1. Verify that composition of lenses is associative, and that a lens $I \to (X, S)$ is an element of X while a lens $(Y, R) \to I$ is a function $Y \to R$.

Exercise 2.2. Using Definition 2.3, compute the equilibrium predicate of the closed game of Example 2.7 for the entry game in which player 1 chooses *Out* (payoffs $(0, 4)$) or *In*, after which player 2 chooses *Accommodate* $(1, 2)$ or *Fight* $(-1, 0)$. List all equilibria and identify the one that is not subgame perfect.

Exercise 2.3. Show that $(\varepsilon \otimes \delta) \otimes \gamma = \varepsilon \otimes (\delta \otimes \gamma)$ for selection functions, and that $\overline{\varepsilon \otimes \delta}(q) = \overline{\varepsilon}\big(x \mapsto \overline{\delta}(q(x, -))\big)$. Interpret the second identity when $\overline{\varepsilon} = \max$ and $\overline{\delta} = \min$.

Exercise 2.4. For $\mathcal{C} = \mathsf{Set}$ with cartesian product, prove that the coend defining $\mathsf{Optic}_{\mathcal{C}}$ reduces to the set of lenses. (Use the bijection $\mathsf{Set}(X, M \times Y) \cong \mathsf{Set}(X, M) \times \mathsf{Set}(X, Y)$ and the co-Yoneda lemma.)

Exercise 2.5. Let $\mathcal{G}, \mathcal{G}' \colon I \to I$ be closed games. Show that the equilibria of $\mathcal{G} \otimes \mathcal{G}'$ are the pairs of equilibria. Then give open games $\mathcal{G} \colon I \to (Y, R)$ and $\mathcal{H} \colon (Y, R) \to I$ such that $\mathcal{H} \circ \mathcal{G}$ has an equilibrium (σ, τ) although σ is not an equilibrium of $\mathcal{G}$ in the context induced by any $\tau' \neq \tau$.

Chapter 3

The Extensive Form

The strategic form hides time and information inside the word "strategy". The extensive form displays them: a tree of histories, a rule saying who moves where, and for each player a partition of her decision nodes into sets she cannot tell apart. This chapter sets up the extensive form in a way that makes its two structural hypotheses—that histories form a tree, and that players remember—into statements about a topos; it then confronts the question of Chapter 1 in its sharpest classical form: when are two trees the same game?

3.1 Trees are presheaves

Let ω be the poset $0 < 1 < 2 < \cdots$ regarded as a category.

Definition 3.1. A ***forest*** is a presheaf $T\colon \omega^{\mathrm{op}} \to \mathsf{Set}$: a sequence of sets $T(n)$ with restriction maps $T(n+1) \to T(n)$. An element of $T(n)$ is a ***node*** of depth n, and its image in $T(n-1)$ is its ***parent***. A forest is a ***tree*** if $T(0) = 1$. The category of forests is the presheaf topos $\widehat{\omega} = \mathsf{Set}^{\omega^{\mathrm{op}}}$.

That forests are exactly presheaves on ω is an old observation; the topos $\widehat{\omega}$ has acquired the name *topos of trees* in the semantics of guarded recursion [BMSS12], and we will use it in that role in Chapter 9. Three pieces of tree vocabulary are topos vocabulary.

- A *play* of infinite length is a compatible sequence of nodes, that is, a global section $1 \to T$; the set of plays is $\Gamma(T) = \lim_n T(n)$. (A tree with terminal nodes is brought under this description by giving each terminal node a single dummy successor forever.)
- The *subtree* at a node $h \in T(m)$ is the forest $T_h(n) = \{t \in T(m+n) \mid t \mapsto h\}$, a fibre of a restriction map.
- A *labelling* of edges by actions is a map of forests from T to the forest of finite words over an alphabet.

Definition 3.2. A finite ***extensive game*** Γ consists of

(i) a finite tree T, with node set H, terminal nodes $Z \subseteq H$ and decision nodes $D = H \setminus Z$;

(ii) a map $\iota\colon D \to N \sqcup \{0\}$ assigning each decision node to a player or to Nature (0), and at each node of Nature a probability distribution on its children;

(iii) for each player i an equivalence relation $\sim_i$ on $D_i = \iota^{-1}(i)$, whose classes are i's ***information sets***, forming a set $\mathcal{I}_i$ with quotient map $q_i\colon D_i \twoheadrightarrow \mathcal{I}_i$;

(iv) for each information set I a set $A(I)$ of actions and, for each $h \in I$, a bijection between $A(I)$ and the children of h;

(v) payoffs $u\colon Z \to \mathbb{R}^N$.

A pure strategy of player i chooses an action at each of her information sets. The actions form a bundle $\coprod_I A(I) \to \mathcal{I}_i$, and a pure strategy is a section of it:

$$S_i = \prod_{I:\mathcal{I}_i} A(I).$$

A ***behavioural strategy*** is a section of the bundle of lotteries, $\beta_i \in \prod_I \Delta A(I)$; a ***mixed strategy*** is a lottery over sections, $\mu_i \in \Delta\big(\prod_I A(I)\big)$. A profile of strategies of any of these kinds, with Nature's distributions, induces a distribution over terminal nodes, and taking expected payoffs gives the ***strategic form*** of Γ, a game in the sense of Definition 1.2. The requirement that a strategy be a function on $\mathcal{I}_i$ rather than on D_i is the entire formal content of "not knowing": a plan must factor through the quotient q_i.

3.2 A taxonomy of information

Informational assumptions are properties of the family of quotient maps (q_i) and of the position of Nature's moves. The source text [Raso6, ch. 2] organizes them into four binary distinctions, which in the present terms read as follows.

Information is…	**Condition on** Γ
perfect	every q_i is a bijection: all information sets are singletons
certain	no node of Nature lies below a node of a player
symmetric	at every node where a player moves, her information set contains every other player's (the mover knows nothing that the others do not), and at terminal nodes all players' information sets coincide
complete	Nature does not move at the root, or all players observe its move there

Perfect information implies symmetric and complete information; certainty is independent of the other three. Simultaneous moves are imperfect but

symmetric information. The interesting distinctions for economics are the failures: *asymmetric* information, where a player moves knowing something that another does not, is the subject of Part II; *incomplete* information, where an unobserved initial move of Nature selects which game is played, is Harsanyi's device for uncertainty about the rules themselves and is taken up in Chapter 5.

3.3 Perfect recall

A player has perfect recall if she never forgets what she knew or did. The standard definition is through her *experience*: for $h \in D_i$ let $X_i(h)$ be the sequence of pairs (information set of i, action taken there) encountered along the path from the root to h. Player i has ***perfect recall*** if $h \sim_i h'$ implies $X_i(h) = X_i(h')$.

There is a more structural way to say this. Each player has her own clock, which ticks when she moves.

Definition 3.3. The ***personal forest*** of player i is the forest P_i whose nodes of depth n are the decision nodes $h \in D_i$ at which i has moved exactly n times before, the parent of h being the last node $g \in D_i$ strictly before h. The edge from g to h is labelled by the action i takes at g on the way to h.

Proposition 3.4 (Perfect recall is internal)**.** [reformulation; cf. [Kuh53, vS96]] *Player i has perfect recall if and only if $\sim_i$ is a congruence on the labelled forest P_i: whenever $h \sim_i h'$, either both are roots of P_i, or their parents are $\sim_i$-equivalent and the connecting edges carry the same action. Equivalently: $\mathcal{I}_i$ carries a (unique) structure of labelled forest making $q_i \colon P_i \twoheadrightarrow \mathcal{I}_i$ a morphism of $\widehat{\omega}$ compatible with labels.*

Proof. If the parents g, g' of h, h' are equivalent with the same action a, then inductively $X_i(h) = X_i(g) \cdot ([g], a) = X_i(g') \cdot ([g'], a) = X_i(h')$; and two roots have empty experience. Conversely if $X_i(h) = X_i(h')$ the sequences have the same length, and if nonempty the same last entry, which says that the parents are equivalent and the actions agree. The second formulation restates the first: a quotient of a presheaf by a levelwise equivalence relation is a presheaf, with the quotient map natural, exactly when the relation is compatible with restriction. □

So perfect recall is the hypothesis that *a player's information is itself a tree*: her information sets are the nodes of a forest of their own, the forest of what she has seen and done, and the quotient from histories to information is a map in the topos of (her) time. Under imperfect recall the partition exists only level by level, or not even that, and there is no object in $\widehat{\omega}$ that deserves to be called her point of view. This is why perfect recall is not a mere technical

convenience: it is the condition for the player to be a *single agent extended in time* rather than a team of agents sharing a payoff.

Theorem 3.5 (Kuhn). [classical; [Kuh53]] *If player i has perfect recall, then for every mixed strategy of i there is a behavioural strategy inducing the same distribution over terminal nodes against every profile of the other players, and conversely.*

Proof sketch, via realization plans [vS96]. By Proposition 3.4 the information sets and actions of i form a labelled forest; call a path from a root a *sequence* of i. A *realization plan* is a function r on sequences with $r(\varnothing) = 1$ and $\sum_{a \in A(I)} r(\varsigma_I a) = r(\varsigma_I)$, where ς_I is the unique sequence leading to I. The probability of a terminal node is the product over players of the realization probabilities of the sequences leading to it, times Nature's probabilities; so outcomes depend on a strategy only through its realization plan. A mixed strategy μ has plan $r_\mu(\varsigma) = \mu\{s \mid s \text{ agrees with } \varsigma\}$; a behavioural strategy β has plan $r_\beta(\varsigma) = \prod_{(I,a)\in\varsigma} \beta_I(a)$; and from any plan one recovers a behavioural strategy by $\beta_I(a) = r(\varsigma_I a)/r(\varsigma_I)$ wherever the denominator is positive. □

Remark 3.6 (The monad-theoretic content). The distribution monad Δ is monoidal: there is a natural map

$$\nabla \colon \prod_I \Delta A(I) \longrightarrow \Delta\Big(\prod_I A(I)\Big)$$

forming product measures. It is not surjective—its image is the independent distributions—and this failure is what correlation *is* (Chapter 6). Kuhn's theorem says that when information is a tree, the outcome map cannot see the difference: it factors through realization plans, on which ∇ becomes surjective. The reason is that a single play visits the information sets of i along one *chain* of her forest, so correlation between choices at incomparable information sets is never observed. Without perfect recall the same play can meet two information sets whose joint behaviour matters, and the theorem fails.

Example 3.7 (Forgetting)**.** A single player chooses L or R, forgets which, and then chooses l or r at one information set containing both resulting nodes. The mixed strategy $\frac{1}{2}(L,l) + \frac{1}{2}(R,r)$ reaches the terminal nodes Ll and Rr with probability $\frac{1}{2}$ each. A behavioural strategy randomizes independently at the two information sets, so it cannot give positive probability to both Ll and Rr without giving positive probability to Lr or Rl.

3.4 When are two trees the same game?

Definition 3.2 defines a type EFG of structures, and as in Lemma 1.3 its identity type is isomorphism of labelled trees. Nobody believes that this is the right

notion of sameness for all purposes. Whether Nature's move is drawn before or after a simultaneous move of the players, whether two successive choices of one player are drawn as one node with four branches or two nodes with two branches each—these are conventions of drawing. Thompson made the point precise.

Theorem 3.8 (Thompson; Elmes–Reny). [classical; [Tho52, ER94]] *Two finite extensive games have the same reduced strategic form, up to isomorphism, if and only if one can be obtained from the other by a finite sequence of four elementary transformations and their inverses:* interchange *of two consecutive moves made in ignorance of each other;* coalescing *of two consecutive moves of the same player into one;* addition of a superfluous move*; and* inflation–deflation *of information sets. Among games with perfect recall, the first three (suitably formulated) suffice and preserve perfect recall.*

Here the *reduced* strategic form identifies pure strategies of a player that yield the same payoffs against everything. In the language of Chapter 1 the theorem computes a quotient. There is a map

$$\mathsf{EFG} \longrightarrow \mathsf{NFG}_u, \qquad \Gamma \mapsto \text{reduced strategic form},$$

and its fibres are the classes of the equivalence relation generated by the Thompson moves. Let $\mathsf{EFG}_{\mathrm{T}}$ be the type obtained from EFG by adjoining a path for each Thompson move—a higher inductive type, or classically the quotient. Then, by Principle 1.5:

> *A solution concept on* $\mathsf{EFG}_{\mathrm{T}}$ *is a solution concept on* EFG *that is constant along Thompson moves; equivalently, one that depends only on the reduced strategic form.*

This is the *invariance* requirement of Kohlberg and Mertens [KM86]. The next example, which is theirs in spirit, shows that it has teeth.

Example 3.9 (Coalescing destroys subgames)**.** In the game Γ on the left of Figure 3.1, player 1 first chooses between an outside option *Out*, worth $(2,2)$, and *In*; after *In* the two players play a simultaneous game in which (T,L) pays $(3,1)$, (B,R) pays $(1,3)$, and miscoordination pays $(0,0)$. In Γ' on the right, player 1's two successive choices have been coalesced into one choice among *Out*, *T*, *B*. The two trees have the same reduced strategic form:

	L	R
Out	$2,2$	$2,2$
T	$3,1$	$0,0$
B	$0,0$	$1,3$

Its Nash equilibria are (T, L) and the profiles in which player 1 plays *Out* and player 2 plays L with probability $q \leq \frac{2}{3}$. In Γ the simultaneous game after *In* is a proper subgame, with three equilibria: (T, L), (B, R), and a mixed one in which player 1 plays T with probability $\frac{3}{4}$ and player 2 plays L with probability $\frac{1}{4}$, worth $\frac{3}{4}$ to player 1. A subgame perfect equilibrium (Chapter 8) must induce one of these, so the subgame perfect equilibria of Γ are $(In, T; L)$ and *Out* followed by $q \in \{0, \frac{1}{4}\}$. In Γ' there is no proper subgame, and every Nash equilibrium is subgame perfect, including *Out* with $q = \frac{1}{2}$. The set of subgame perfect equilibria is not constant along a coalescing move.

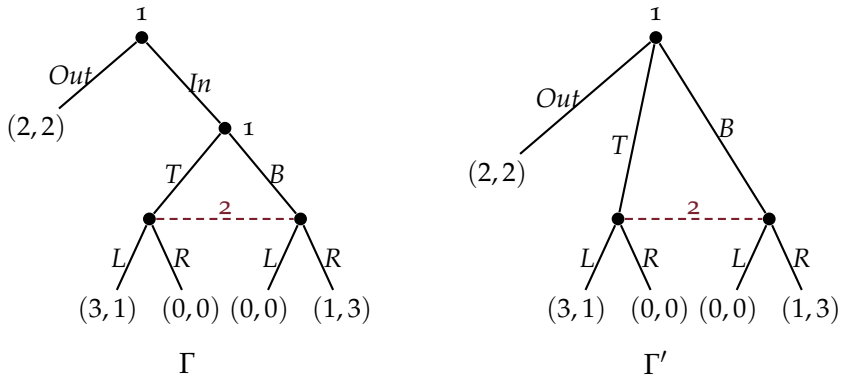


Figure 3.1: Two extensive games related by coalescing player 1's moves. They have the same reduced strategic form and different sets of subgame perfect equilibria.

Corollary 3.10. [classical; [KM86]] *Subgame perfect equilibrium is not a solution concept on* $\mathsf{EFG}_{\mathrm{T}}$. *The same example shows this for sequential equilibrium.*

What should one conclude? Two positions are coherent, and the type-theoretic framing makes clear that they disagree about the domain, not about rationality.

The invariance position [KM86]: the Thompson moves do not change the decision problem any player faces, so the type of games is $\mathsf{EFG}_{\mathrm{T}}$, a solution concept must descend to it, and subgame perfection is a property of a drawing. On this view the correct refinements are those defined on the strategic form whose pullback to every tree is well behaved. Such concepts exist. A *proper* equilibrium of the strategic form [Mye78] induces a sequential equilibrium in every extensive game with that strategic form [vD84, KM86]; *stable sets* are designed to have this and further properties. In Example 3.9 both select (T, L): *Out* strictly dominates B, so player 2, finding herself called on to move, should infer T—the argument known as *forward induction*.

The structural position: the order of moves is part of the game; coalescing changes what player 2 can infer about player 1's state of mind at her information set; the type of games is EFG with its own identity type, and subgame perfection is a perfectly good function on it.

What is not coherent is to use both types at once: to regard the trees Γ and Γ' as "the same game" when drawing them and as different when solving them. The discipline of declaring the type does not settle the debate, but it does make it impossible to ignore, and it recurs in every later chapter in which a solution concept is found to depend on a presentation: redundant types in Chapter 5, history-dependence in Chapter 9.

Remark 3.11 (Subgames). A node h is the root of a ***subgame*** if every information set that meets the subtree T_h is contained in it. Subgames are the subobjects of the tree to which the whole structure of Definition 3.2 restricts; they are the opens of the topology in which Chapter 8 takes place. The example shows that this topology is not a Thompson invariant.

Notes

The extensive form with information sets is due to Kuhn [Kuh53], generalizing von Neumann and Morgenstern. Theorem 3.5 is from the same paper; the realization-plan proof follows von Stengel's sequence form [vS96], which is also the basis of efficient algorithms for extensive games. The transformations of Theorem 3.8 were introduced by Thompson [Tho52]; Elmes and Reny [ER94] showed that inflation–deflation can be dispensed with among games with perfect recall. The invariance debate was opened by Kohlberg and Mertens [KM86]. The relation between proper equilibria of the strategic form and sequential equilibria of the tree is due to van Damme [vD84] and to [KM86]. Presheaves on ω as trees are folklore; the name "topos of trees" is from [BMSS12]. Translations of extensive games into open games are given in [CGLN22].

Exercises

Exercise 3.1. Show that for a tree T in which every node has a child, the set of infinite plays is the set of global sections of T in $\widehat{\omega}$, and that the subtree T_h is a pullback.

Exercise 3.2. Draw an extensive game with two players in which information is symmetric but imperfect, and one in which it is certain but asymmetric, and verify the conditions in the table.

Exercise 3.3. Give an example of a game without perfect recall in which the personal forest P_i and the levelwise quotient both exist but the quotient map is not a map of forests. Identify the failure of naturality.

Exercise 3.4. Verify the claims of Example 3.9: compute all Nash equilibria of the strategic form and all three equilibria of the simultaneous subgame.

Exercise 3.5. Starting from the tree Γ of Example 3.9, apply an interchange of moves to the simultaneous subgame so that player 2 moves before player 1's choice between T and B. Does the set of subgame perfect equilibria change?

Exercise 3.6. Let r be a realization plan for a player with perfect recall. Show that the set of realization plans is a polytope whose dimension is $\sum_I (|A(I)| - 1)$, and compare with the dimension of the mixed-strategy simplex.

Chapter 4

Knowledge

"Player i knows E" is the most heavily used undefined phrase in informal game theory. Aumann's partition model defines it. This chapter derives the partition model from a single categorical construction—base change along the map that sends a state of the world to what a player observes of it—and then computes *common* knowledge as a colimit. Taking that colimit homotopically produces a graph whose components are the classical answer and whose loops are the subject of the next chapter.

4.1 Information structures

Definition 4.1. An ***information structure*** for a set N of agents is a set Ω of ***states*** together with, for each $i \in N$, a surjection $p_i \colon \Omega \twoheadrightarrow T_i$ onto a set of ***types*** (or signals). An ***event*** is a subset $E \subseteq \Omega$.

The fibres of p_i are the cells of agent i's partition: at state ω she observes $p_i(\omega)$ and cannot exclude any state of the fibre $p_i^{-1}p_i(\omega)$. In an extensive game Ω may be taken to be the set of plays and p_i the record of what i observes; in a Bayesian game (Chapter 5) Ω is the set of type profiles and states of nature. We write $\mathrm{Sub}(\Omega)$ for the Boolean algebra of events.

4.2 Knowledge is base change

A function $p \colon \Omega \to T$ induces the inverse-image map $p^* \colon \mathrm{Sub}(T) \to \mathrm{Sub}(\Omega)$, and p^* has adjoints on both sides,

$$\exists_p \;\dashv\; p^* \;\dashv\; \forall_p, \qquad \exists_p E = \{t \mid p^{-1}(t) \cap E \neq \varnothing\}, \quad \forall_p E = \{t \mid p^{-1}(t) \subseteq E\}.$$

This is Lawvere's observation that quantifiers are adjoints to substitution, and it holds in any topos.

Definition 4.2. For an agent with observation map p the ***knowledge*** and ***possibility*** operators on $\mathrm{Sub}(\Omega)$ are

$$K = p^*\forall_p, \qquad P = p^*\exists_p.$$

Explicitly $KE = \{\omega \mid p^{-1}p(\omega) \subseteq E\}$ is the union of the cells contained in E: the set of states at which the agent's observation rules out every state outside E. This is Aumann's operator [Aum76]. The definition as a composite of adjoints makes its properties automatic.

Proposition 4.3 (S5). [classical; the adjoint formulation is standard in categorical logic] *K is an idempotent comonad and P an idempotent monad on the poset* $\mathrm{Sub}(\Omega)$*, with $P \dashv K$. Hence, for all events E, F:*

(i) $KE \subseteq E$ *(truth)*

(ii) $KE = KKE$ *(positive introspection)*

(iii) $K(E \cap F) = KE \cap KF$ *and* $K\Omega = \Omega$*; indeed K preserves all intersections*

(iv) $PE \subseteq F \iff E \subseteq KF$

(v) $\neg KE = K\neg KE$ *(negative introspection)*

The fixed points of K are exactly the events of the form p^*D, those measurable with respect to the agent's partition.

Proof. A composite $p^*\forall_p$ of a left adjoint after its right adjoint is a comonad, and it is idempotent because p is surjective, so that $\forall_p p^* = \mathrm{id}$. Dually for P. (iii) holds because both p^* and $\forall_p$ are right adjoints. (iv): $p^*\exists_p E \subseteq F \iff \exists_p E \subseteq \forall_p F \iff E \subseteq p^*\forall_p F$. (v): KE is in the image of p^*, so is its complement since p^* preserves complements, and $Kp^* = p^*$. □

Properties (i)–(iii) and (v) are the axioms of the modal logic S5. The proof shows where each comes from, and therefore what survives in other settings. In a topos that is not Boolean, (i)–(iv) hold verbatim and (v) may fail: knowledge is an S4 modality whose negative introspection is a classical accident. For an observation *relation* rather than a function one keeps only the adjunction (iv) and monotonicity. The literature on bounded reasoning and unawareness, which gives up partitions, is in this sense the study of which parts of an adjoint triple one is prepared to lose.

Remark 4.4 (Proof-relevant knowledge). In type theory the triple becomes $\sum_p \dashv p^* \dashv \prod_p$ on families of types, and for a family $E \colon \Omega \to \mathsf{Type}$

$$(KE)(\omega) \;\equiv\; \prod_{\omega' \colon \mathrm{fib}_p(p\,\omega)} E(\omega').$$

A witness that the agent knows E at ω is a function producing evidence of E at every state she considers possible. Necessity-as-dependent-product along a map is the standard account of modality in dependent type theory [Cor20]. Note that K is a *comonadic* modality; the modalities of [RSS20] are monadic, like P. The comonadic ones need more type-theoretic care [Shu18], a technical echo of the fact that knowledge is harder to axiomatize than possibility.

4.3 Common knowledge

Now let there be several agents, with operators K_i. Put $KE = \bigcap_i K_i E$ ("everyone knows E").

Definition 4.5. E is ***common knowledge*** at ω if $\omega \in CE :\equiv \bigcap_{n\geq 1} K^n E$: everyone knows E, everyone knows that everyone knows E, and so on. An event F is ***self-evident*** (or *public*) if $KF = F$.

Theorem 4.6 (Aumann). [classical; [Aum76]] *Let $q\colon \Omega \twoheadrightarrow T_\wedge$ be the pushout of the maps p_i in* Set*:*

$$\begin{array}{ccc} \Omega & \xrightarrow{p_1} & T_1 \\ {\scriptstyle p_2}\downarrow & \searrow^{q} & \downarrow \\ T_2 & \longrightarrow & T_\wedge \end{array}$$

(and similarly for more agents). Then $C = q^*\forall_q$. That is, common knowledge is the knowledge of a fictitious agent whose partition is the meet *of the agents' partitions, their finest common coarsening. Moreover CE is the largest self-evident event contained in E:*

$$CE = \nu X.\, K(E \cap X).$$

Proof. Each K_i preserves arbitrary intersections, hence so does K, and $K^{n+1}E \subseteq K^n E$; therefore $K(CE) = \bigcap_{n\geq 2} K^n E = CE$, so CE is self-evident, and it lies in E. If $F \subseteq E$ is self-evident then $F = K^n F \subseteq K^n E$ for all n, so $F \subseteq CE$. This proves the last claim. An event is self-evident iff it is a union of cells of every agent's partition, iff it is saturated for the equivalence relation generated by all the $\ker p_i$, iff it has the form q^*D. The largest such event inside E is $q^*\forall_q E$. □

Two things deserve emphasis. First, common knowledge is defined by an infinite conjunction and computed by a *finite colimit*. The iterative hierarchy "I know that you know that…" is the standard computation of a greatest fixed point from above; the colimit is its closed form. Second, the greatest-fixed-point formula is a coinduction principle, exactly parallel to that for rationalizability in Chapter 1:

> *To show that E is common knowledge at ω, exhibit a self-evident event F with $\omega \in F \subseteq E$.*

A public announcement is the canonical example: the event that the announcement was made is self-evident, so whatever it entails is common knowledge. No infinite regress has to be traversed.

Remark 4.7 (Limits and colimits of information). Dually, let $\langle p_i \rangle \colon \Omega \to \prod_i T_i$ be the pairing and $p_\vee$ its corestriction to its image. The associated operator $p_\vee^* \forall_{p_\vee}$ is ***distributed knowledge***: what the agents would know if they pooled their observations. Its partition is the join, the coarsest common refinement. So the two group-knowledge notions of epistemic logic are the two universal constructions on the diagram of type maps: common knowledge is its *colimit*, distributed knowledge its *limit*.

4.4 The homotopy colimit

A pushout of sets forgets how its identifications were made. The homotopical habit is to compute the colimit in spaces instead, where the identifications are recorded as paths, and then to see what has been gained.

Definition 4.8. The ***knowledge graph*** B of an information structure $(\Omega, (p_i)_{i \in N})$ is the bipartite graph with vertex set $\Omega \sqcup \coprod_i T_i$ and one edge $e_{\omega,i}$ joining ω to $p_i(\omega)$ for each state ω and agent i.

Proposition 4.9. [reformulation; cf. [Tho79]] *B is the homotopy colimit, in the ∞-category of spaces, of the diagram of discrete spaces $\{p_i \colon \Omega \to T_i\}_{i \in N}$. Consequently:*

(i) $\pi_0(B) = T_\wedge$*: the connected components of B are the cells of the meet, the minimal self-evident events;*

(ii) each component is a connected graph, hence homotopy equivalent to a wedge of circles, and

$$b_1(B) = (n-1)\,|\Omega| - \sum_i |T_i| + |T_\wedge|$$

when Ω is finite and $n = |N|$.

Proof. The homotopy colimit of a diagram of sets is the nerve of its category of elements [Tho79]. Here that category has objects $\Omega \sqcup \coprod_i T_i$ and one non-identity arrow $\omega \to p_i(\omega)$ for each pair (ω, i), with no composable pairs; its nerve is the graph B. (i) holds because π_0 is a left adjoint and so carries homotopy colimits of spaces to colimits of sets. (ii) is the Euler characteristic: B has $|\Omega| + \sum_i |T_i|$ vertices, $n|\Omega|$ edges and $|T_\wedge|$ components. □

For two agents every state-vertex has degree two and may be suppressed: B becomes the graph whose vertices are the cells of the two partitions and whose edges are the states, each state joining the two cells that contain it (Figure 4.1). A path in B is a chain of the form "at ω agent 1 cannot exclude ω', at which agent 2 cannot exclude ω'',…"—precisely the chains along which higher-order uncertainty propagates. Aumann's theorem says that only the

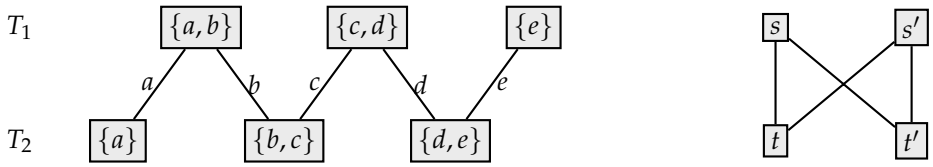


Figure 4.1: Knowledge graphs for two agents (state-vertices suppressed). Left: five states, partitions $\{ab \mid cd \mid e\}$ and $\{a \mid bc \mid de\}$; B is a tree, so no nontrivial event is common knowledge and $b_1 = 0$. Right: the product structure $\Omega = T_1 \times T_2$ with two types each; B is a 4-cycle and $b_1 = 1$.

existence of such a chain matters for common knowledge. The graph also remembers *how many essentially different* chains there are.

Example 4.10. For the product structure $\Omega = T_1 \times T_2$ with $|T_1| = m$, $|T_2| = n$—independent private types with full support, the standard setting of auction theory—B is the complete bipartite graph $K_{m,n}$, connected with $b_1 = (m-1)(n-1)$. If agent 1 is fully informed ($p_1 = \mathrm{id}$), every state is a leaf hanging from a cell of agent 2, and B is a disjoint union of stars: $b_1 = 0$.

So π_0 of the homotopy colimit is common knowledge, and there is a π_1 about which the classical theory of knowledge is silent. That silence is justified: Proposition 4.3 shows that knowledge operators factor through Sub, and subobjects cannot detect loops. The loops become visible as soon as one puts something over B that can have monodromy. Probabilistic beliefs are such a thing, and Chapter 5 shows that the first cohomology of B is exactly the obstruction to a common prior.

4.5 Many agents: simplicial models

For n agents there is a second geometric object in the literature, coming from distributed computing.

Definition 4.11. [literature; [GLR21, HKR13]] The ***simplicial model*** of an information structure is the simplicial complex $\mathcal{C}$ with vertex set $\coprod_i T_i$—a vertex is a pair (agent, local state), "coloured" by its agent—in which each state ω spans the $(n-1)$-simplex $\{p_1(\omega), \dots, p_n(\omega)\}$.

Two states are indistinguishable to agent i exactly when their simplices share the i-coloured vertex, and indistinguishable to every member of a coalition J when they share a J-coloured face. Goubault, Ledent and Rajsbaum [GLR21] show that such complexes correspond to (proper) Kripke models of multi-agent S5, and that standard communication protocols act on them by subdivision. The topology matters: in the theory of distributed computing

[HKR13] a task is solvable only if a certain map of complexes exists, and the obstructions are higher connectivity conditions—the impossibility of wait-free consensus is the statement that a protocol complex is connected and so cannot be mapped compatibly onto a disconnected output complex.

The graph B and the complex $\mathcal{C}$ are related by a map $B \to \mathcal{C}$ sending each state-vertex to the barycentre of its simplex, and they have the same π_0. They do not have the same π_1. Take three agents and two states that agents 1 and 2 cannot distinguish and agent 3 can: $\mathcal{C}$ is two triangles glued along an edge, which is contractible, whereas B has a cycle. Both are legitimate invariants: B is the homotopy colimit of the type maps and controls gluing problems for data attached to individual agents; $\mathcal{C}$ additionally fills in what coalitions distributedly know. Which higher invariants of $\mathcal{C}$ are game-theoretically meaningful is, as far as we know, open (Chapter 19).

4.6 Knowledge and time

In an extensive game with perfect recall, each player's information refines as play proceeds: by Proposition 3.4 her information sets form a forest, and the map from plays to her information at her n-th move is a tower of quotients $\Omega \to \cdots \to T_i^{(n+1)} \to T_i^{(n)}$—classically, a filtration. Each stage has its own knowledge operator $K_i^{(n)}$ and these increase with n. The object underlying the tower is a presheaf on ω, that is, an information structure *in the topos of trees*; the internal knowledge modality of that topos packages all the $K_i^{(n)}$ at once. We shall not need this internal formulation in detail, but it explains why the temporal modality of Chapters 8 and 9 and the epistemic modality of this chapter have the same algebraic form: both are $p^*\forall_p$ for a suitable map p.

Notes

The partition model and the characterization of common knowledge by the meet are due to Aumann [Aum76]; the notion itself is usually credited to Lewis. The fixed-point formulation is standard in epistemic logic. Quantifiers as adjoints are due to Lawvere; the systematic reading of modal operators as base-change (co)monads in dependent type theory is expounded in [Cor20], and modalities in homotopy type theory are studied in [RSS20]. The identification of the homotopy colimit of a diagram of sets with the nerve of its category of elements is a special case of Thomason's theorem [Tho79]. Simplicial models for epistemic logic are due to Goubault, Ledent and Rajsbaum [GLR21], building on the topological theory of distributed computing [HKR13]. The approximation of common knowledge by common p-belief, which matters when

knowledge is replaced by high-probability belief, is taken up in Chapter 10.

Exercises

Exercise 4.1. Let $p\colon \Omega \to T$ be a function that is not surjective. Which clauses of Proposition 4.3 survive?

Exercise 4.2. Show directly from the adjunctions that $K_i K_j$ need not equal $K_j K_i$, and that $CE \subseteq K_{i_1} \cdots K_{i_m} E$ for every finite sequence of agents.

Exercise 4.3. For the information structure on the left of Figure 4.1, compute $K_1 E$, $K_2 E$, $K_1 K_2 E$ and CE for $E = \{a, b, c, d\}$. At which states does everyone know E without E being common knowledge?

Exercise 4.4. Verify the formula for $b_1(B)$ in Proposition 4.9 on both graphs of Figure 4.1, and prove Example 4.10.

Exercise 4.5. Prove the claim of Remark 4.7: distributed knowledge is $p_\vee^* \forall_{p_\vee}$ for the join partition. Show that for every event $CE \subseteq K_i E \subseteq DE \subseteq E$, where D denotes distributed knowledge.

Exercise 4.6. Three agents, two states, with p_1, p_2 constant and p_3 injective. Draw B and $\mathcal{C}$ and compute their fundamental groups.

Chapter 5

Belief

Knowledge excludes; belief weighs. This chapter adds probability to the information structures of Chapter 4 and treats the three classical topics that result: Harsanyi's reduction of games with incomplete information to games with imperfect information; the universal space of hierarchies of belief; and the common prior assumption with its most famous consequence, the impossibility of agreeing to disagree. The categorical tools are a monad, a terminal coalgebra and a cohomology class.

5.1 Probability as a monad

For a set X let ΔX be the set of finitely supported probability distributions on X. Δ is a monad on Set: the unit $\eta\colon X \to \Delta X$ sends a point to its Dirac distribution, and the multiplication $m\colon \Delta\Delta X \to \Delta X$ averages a distribution over distributions. For measurable spaces the same structure is carried by the ***Giry monad*** $\mathcal{G}$ on Meas, which sends a space to the space of probability measures on it [Gir82]. A map $X \to \Delta Y$ is a ***stochastic map*** or Markov kernel; such maps compose by integration, and form the Kleisli category of the monad, which we denote FinStoch on finite sets and Stoch in general.

It is often cleaner to axiomatize the Kleisli category directly. A ***Markov category*** [Fri20, CJ19] is a symmetric monoidal category in which every object has a *copy* map $X \to X \otimes X$ and a *delete* map $X \to I$, compatibly with $\otimes$, such that deletion is natural. A morphism is *deterministic* if it commutes with copying. In FinStoch the deterministic maps are the functions. Three notions will be used repeatedly, and all three are expressible in any Markov category with conditionals:

- the ***marginals*** of a joint state $I \to X \otimes Y$, obtained by deleting one factor;
- the ***Bayesian inverse*** $f^\dagger_\pi\colon Y \to X$ of a kernel $f\colon X \to Y$ with respect to a prior $\pi\colon I \to X$, characterized by the equality of the two ways of forming the joint distribution on $X \otimes Y$, and unique πf-almost surely;
- ***conditional expectation*** along a deterministic $p\colon \Omega \to T$, which for a prior μ is integration against the Bayesian inverse $p^\dagger_\mu$.

Bayesian inversion is contravariantly functorial, $(g \circ f)^\dagger = f^\dagger \circ g^\dagger$ with the

evident priors. A kernel together with its inverse is therefore a lens-like object—forward prediction, backward inference—and the compositional theory of such *Bayesian lenses* [BHS23] is what makes the Bayesian open games of Section 2.5 work.

5.2 Games with incomplete information

Suppose the players do not know the game: payoffs depend on a state of nature $\theta \in \Theta$ about which they have different information. Harsanyi's proposal [Har67] was to model each player's information by a *type* and to let the uncertainty be resolved by a move of Nature.

Definition 5.1. A ***Bayesian game*** consists of players N, a set Θ of states of nature, type sets T_i, action sets A_i, payoff functions $u_i \colon A \times \Theta \to \mathbb{R}$, and *beliefs* $\beta_i \colon T_i \to \Delta(\Theta \times T_{-i})$. A strategy of player i is a function $s_i \colon T_i \to A_i$. The profile s is a ***Bayesian Nash equilibrium*** if for every i and every t_i the action $s_i(t_i)$ maximizes the interim expected payoff

$$\mathbb{E}_{(\theta, t_{-i}) \sim \beta_i(t_i)} \, u_i\big(a_i, s_{-i}(t_{-i}), \theta\big).$$

The game has a ***common prior*** if there is $\pi \in \Delta(\Theta \times T)$ whose conditionals are the β_i.

With a common prior the Bayesian game is the extensive game in which Nature first draws (θ, t) from π and each player observes her own coordinate t_i. Information is *incomplete* in the taxonomy of Chapter 3, and its incompleteness has been converted into *asymmetry* about a move of Nature. In the language of the earlier chapters, a Bayesian game is a family (G_θ) of strategic-form games indexed by the states, together with an information structure $\Omega = \Theta \times T$, $p_i \colon \Omega \to T_i$; the Harsanyi transformation is the dependent sum $\sum_\theta G_\theta$, and the requirement that s_i be a function of t_i is the factorization through p_i that defines not knowing. A strategy is a section over T_i, so $S_i = \prod_{t_i} A_i$, and when every type has positive probability the Bayesian equilibria are the Nash equilibria of the strategic form with these strategy sets—or equally of the *agent form* in which each type is a separate player.

Example 5.2 (A contribution game)**.** Two players simultaneously decide whether to contribute to a public good. The good is provided if at least one contributes and is then worth 1 to each. Player i's cost of contributing, c_i, is her private type, drawn independently and uniformly from $[0, 2]$. Look for an equilibrium in cutoff strategies: contribute iff $c_i \le c^*$. If the other player uses the cutoff c^* she contributes with probability $c^*/2$. Contributing yields $1 - c_i$; not contributing yields $c^*/2$. So player i contributes iff $c_i \le 1 - c^*/2$,

and consistency requires $c^* = 1 - c^*/2$, that is $c^* = \frac{2}{3}$. The good is provided with probability $1 - (\frac{2}{3})^2 = \frac{5}{9}$, although it would be efficient to provide it whenever $\min(c_1, c_2) < 2$, which is almost always. Each player free-rides on the possibility that the other is a low-cost type.

5.3 Types as a terminal coalgebra

The definition of a Bayesian game is circular in a way that troubled its first readers: a type is something that determines a belief about, among other things, the other players' types. Harsanyi cut the circle by fiat, taking the T_i as primitive. The alternative is to unfold it. Player i has a belief about Θ; a belief about Θ and j's belief about Θ; and so on—a *hierarchy*. Is every hierarchy a type? Is there a space containing them all?

Definition 5.3. Fix a measurable space Θ and, for definiteness, two players. A ***type space*** is a coalgebra for the endofunctor

$$F(X_1, X_2) = \big(\mathcal{G}(\Theta \times X_2),\ \mathcal{G}(\Theta \times X_1)\big)$$

of $\mathsf{Meas} \times \mathsf{Meas}$: a pair of measurable spaces with measurable maps $\beta_i \colon T_i \to \mathcal{G}(\Theta \times T_{-i})$. A ***type morphism*** is a coalgebra homomorphism: a pair of measurable maps $h_i \colon T_i \to T_i'$ with $\beta_i' \circ h_i = \mathcal{G}(\mathrm{id}_\Theta \times h_{-i}) \circ \beta_i$.

A type morphism is a map that preserves beliefs, including beliefs about beliefs, to all orders; this is the notion of morphism of Mertens–Zamir and Heifetz–Samet, and the coalgebraic formulation is due to Moss and Viglizzo.

Theorem 5.4 (Universal type space)**.** [classical; [MZ85, BD93, HeSa98]; coalgebraic form [MV06]] *The category of type spaces over Θ has a terminal object T^*. By Lambek's lemma its structure map is an isomorphism:*

$$T_i^* \ \cong\ \mathcal{G}\big(\Theta \times T_{-i}^*\big).$$

When Θ is Polish (or compact Hausdorff, with regular measures), T^* is the limit of the terminal sequence

$$1 \longleftarrow F1 \longleftarrow F^2 1 \longleftarrow \cdots,$$

whose n-th term is the space of belief hierarchies of depth n and whose limit is the space of coherent hierarchies.

The isomorphism is the non-circular form of Harsanyi's idea: a universal type *is* a belief about nature and the other's universal type. Terminality says that every type in every type space has a unique image in T^*, namely its hierarchy; two types, possibly in different type spaces, with the same image

are ***behaviourally equivalent*** (bisimilar). The construction by the terminal sequence works when F preserves the limit of that sequence, and this is exactly the Kolmogorov extension theorem: a compatible family of finite-depth beliefs extends uniquely to a belief on the limit. In the bare measurable category the extension theorem fails, there are coherent hierarchies that are not types [HeSa99], and the terminal coalgebra—which still exists [HeSa98, MV06]—must be built otherwise. The analytic subtleties of the universal type space are thus the subtleties of when a functor is ω^{op}-continuous.

The price of universality

Why should a universal space exist at all? For non-probabilistic belief it does not.

Proposition 5.5. [classical; Cantor, Lambek; cf. [BK06, AZ15]] *There is no terminal coalgebra for the functor $X \mapsto \mathcal{P}(\Theta \times X)$ on* Set*, for nonempty Θ.*

Proof. By Lambek's lemma a terminal coalgebra would be a set X with $X \cong \mathcal{P}(\Theta \times X)$, contradicting Cantor's theorem. □

A "possibility type" is a set of pairs (state, type of the other) that the player does not exclude, and the proposition says that there is no space of all possibility types: some conceivable belief about the other player is always missing. Brandenburger and Keisler [BK06] sharpened this into an impossibility theorem with the shape of a paradox—*Ann believes that Bob assumes that Ann believes that Bob's assumption is wrong*—and Abramsky and Zvesper [AZ15] showed that it is an instance of Lawvere's fixed-point theorem [Law69], the common core of the diagonal arguments of Cantor, Russell, Gödel and Tarski. Probabilistic beliefs escape because $\mathcal{G}$ is a small functor: a probability measure on a Polish space is determined by countably much data. So do possibility models with enough topology to tame the powerset [MMP05]. Universality is bought by restricting what counts as a belief, and the restriction is not innocent: it is where the continuity hypotheses of the theory come from, and Chapter 10 shows that strategic behaviour is *not* continuous in the topology that makes the universal type space work.

Redundant types and invariance

Terminality suggests an invariance principle in the style of Principle 1.5: a solution concept for Bayesian games should depend on a type only through its image in T^*. Bayesian Nash equilibrium does not.

Example 5.6 (A redundant type space)**.** Let $\theta \in \{0,1\}$ be equally likely. Each of two players chooses *bet* or *pass*. If both bet, each receives $+1$ when $\theta = 1$ and -2 when $\theta = 0$; otherwise both receive 0. In the type space $\mathbb{A}$ each player has a single type, with the prior belief. In the type space $\mathbb{B}$ each player has a type $t_i \in \{+,-\}$; the pair (t_1, t_2) is uniform and $\theta = 1$ exactly when $t_1 = t_2$. In $\mathbb{B}$ each type believes that $\theta = 1$ with probability $\frac{1}{2}$, believes that the other believes this, and so on: every type of $\mathbb{B}$ has the same hierarchy as the type of $\mathbb{A}$. In $\mathbb{A}$ there is no equilibrium in which both bet, since joint betting is worth $-\frac{1}{2}$. In $\mathbb{B}$ the profile "bet iff $t_i = +$" is an equilibrium: a $+$ type meets a betting opponent exactly when the types agree, hence when $\theta = 1$, and earns $\frac{1}{2}$.

The two type spaces are bisimilar and have different equilibrium outcomes: the types of $\mathbb{B}$ carry a payoff-irrelevant correlating signal that the hierarchies do not record. As with subgame perfection and the Thompson moves, there are two coherent responses. One may enlarge the invariant—record more than the hierarchy [EP06, Liu09]—so that $\mathbb{A}$ and $\mathbb{B}$ become different points. Or one may change the solution concept: *interim correlated rationalizability* depends only on the hierarchy [DFM07], and is in this sense the solution concept that actually lives on T^*.

5.4 The common prior as a gluing problem

Return to a finite information structure $(\Omega, (p_i))$ as in Definition 4.1. A ***posterior profile*** β assigns to each agent i and type $t \in T_i$ a probability distribution $\beta_i(t)$ with full support on the cell $p_i^{-1}(t)$. This is what the agents have: beliefs at each of their types. A ***common prior*** for β is a full-support $\mu \in \Delta\Omega$ whose conditional on each cell $p_i^{-1}(t)$ is $\beta_i(t)$. The common prior assumption asserts that one exists, and almost every applied model makes it. When is it satisfiable?

The question is a gluing problem. Each agent's posteriors are a family of "vertical" measures along the fibres of her type map; the n families overlap at the states; we ask for a single measure on Ω inducing them all. Gluing problems have cocycles. Recall the knowledge graph B of Definition 4.8, with an edge $e_{\omega,i}$ from the type $p_i(\omega)$ to the state ω.

Theorem 5.7 (Common priors and H^1)**.** [reformulation; cf. [RN09, HeSa12]] *Define a real 1-cochain on B by*

$$\lambda_\beta(e_{\omega,i}) = \log \beta_i\big(p_i\omega\big)(\omega).$$

(a) β admits a common prior if and only if the class $[\lambda_\beta] \in H^1(B;\mathbb{R})$ vanishes.

(b) *If it does, the conditional of a common prior on each connected component of B—each cell of the meet—is unique; common priors correspond to full-support distributions on $\pi_0(B)$.*

(c) *Every class in $H^1(B;\mathbb{R})$ is $[\lambda_\beta]$ for some posterior profile. Hence if $b_1(B) = 0$ every posterior profile has a common prior, and if $b_1(B) > 0$ those that do form a closed submanifold of codimension $b_1(B)$ in the manifold of posterior profiles; in particular a nowhere dense set.*

Proof. (a) If μ is a common prior with marginals $m_i(t) = \mu(p_i^{-1}t)$ then $\beta_i(t)(\omega) = \mu(\omega)/m_i(t)$, so $\lambda_\beta = \delta\varphi$ for the o-cochain $\varphi(\omega) = \log\mu(\omega)$, $\varphi(i,t) = \log m_i(t)$. Conversely if $\lambda_\beta = \delta\varphi$, put $\mu(\omega) = e^{\varphi(\omega)}$ and $m_i(t) = e^{\varphi(i,t)}$. Then $\beta_i(t)(\omega) = \mu(\omega)/m_i(t)$ on each cell; summing over the cell gives $m_i(t) = \mu(p_i^{-1}t)$, so the conditionals of μ (normalized) are the β_i. (b) The potential φ is determined up to adding a function constant on each component of B, that is, μ is determined up to a positive factor on each component. (c) On a graph every 1-cochain is a cocycle. Given a cochain c, subtract from the edges at each type vertex (i,t) the constant $\log\sum_{\omega\in p_i^{-1}t} e^{c(e_{\omega,i})}$; this changes c by a coboundary and makes it λ_β for a posterior profile. The same computation shows that, in logarithmic coordinates, the posterior profiles are a slice transverse to the coboundaries supported on type vertices, so that $\beta \mapsto [\lambda_\beta]$ is a submersion onto $H^1(B;\mathbb{R}) \cong \mathbb{R}^{b_1}$. □

Unwinding the vanishing of a class on a graph gives a condition on cycles: *around every cycle of B, the alternating product of posterior probabilities is* 1. These are the cycle equations of Rodrigues-Neto [RNo9]. Part (c) recovers the theorem of Hellman and Samet [HeSa12], who call a partition profile *tight* when every posterior profile has a common prior, and characterize tightness (for a connected structure) by the count $\sum_i |T_i| = (n-1)|\Omega| + 1$. By Proposition 4.9 this count is the equation $b_1(B) = 0$: *tight means that the homotopy colimit of the type maps is contractible on each component.*

Example 5.8. In the product structure with two types each (Figure 4.1, right), write the states as $a = (s,t)$, $b = (s,t')$, $c = (s',t)$, $d = (s',t')$. Agent 1's posteriors give the ratios $\mu(b)/\mu(a)$ and $\mu(d)/\mu(c)$; agent 2's give $\mu(c)/\mu(a)$ and $\mu(d)/\mu(b)$. Four numbers are prescribed and a prior has three degrees of freedom; the single cycle equation is

$$\frac{\beta_1(s)(b)}{\beta_1(s)(a)} \cdot \frac{\beta_2(t')(d)}{\beta_2(t')(b)} \cdot \frac{\beta_1(s')(c)}{\beta_1(s')(d)} \cdot \frac{\beta_2(t)(a)}{\beta_2(t)(c)} = 1.$$

Generic posteriors violate it. In the structure on the left of the figure there is no cycle and every posterior profile comes from a prior.

The picture is that of a connection. The posteriors tell each agent how to transport relative probability between two states she cannot distinguish; composing along a path in B transports it between any two states of a component; a common prior exists when transport is path-independent, that is, when the connection is flat with trivial holonomy. Nontrivial holonomy around a cycle means that the agents' relative odds, chained around the cycle, multiply to something other than 1—which is exactly an arbitrage opportunity for an outsider, or a mutually agreeable bet for the insiders.

Theorem 5.9 (No-betting characterization). [classical; [Mor94, Sam98, Fei00]] *A finite posterior profile on a connected information structure has a common prior if and only if there is no* agreeable bet*: no family of functions $f_i\colon \Omega \to \mathbb{R}$ with $\sum_i f_i = 0$ such that every type of every agent assigns strictly positive expected value to her own f_i.*

The classical proofs are separation arguments. Read through Theorem 5.7, the theorem is the familiar duality between potentials and circulations: either the log-odds are a gradient, or some cycle carries a profitable flow.

5.5 Agreement

Theorem 5.10 (Agreeing to disagree). [classical; [Aum76]] *Let μ be a common prior, E an event, and $f_i(\omega) = \mu\big(E \mid p_i^{-1}p_i(\omega)\big)$ agent i's posterior probability of E. If at some state of positive probability it is common knowledge that $f_i = a_i$ for each i, then all the a_i are equal.*

Proof. For a map $p\colon \Omega \to T$ write $\mathbb{E}_p\colon \mathbb{R}^\Omega \to \mathbb{R}^T$ for conditional expectation with respect to μ. Conditional expectation is functorial for composition of maps (the tower property): if $r\colon T \to T'$ then $\mathbb{E}_{rp} = \mathbb{E}_r\mathbb{E}_p$. Apply this to the two factorizations $q = r_i p_i$ of the map to the pushout in Theorem 4.6. Common knowledge of $f_i = a_i$ at ω means that $\mathbb{E}_{p_i}1_E$ is constant, equal to a_i, on the fibre of r_i over $m = q(\omega)$. Hence $(\mathbb{E}_q 1_E)(m) = (\mathbb{E}_{r_i}\mathbb{E}_{p_i}1_E)(m) = a_i$ for every i. □

The proof uses two things and only two: the pushout square of the type maps, and the fact that a single measure is being pushed around it. If the posteriors descend to the pushout, they are both equal to the pushed-forward measure there. Without a common prior there is nothing to push, and Theorems 5.7 and 5.9 say precisely how disagreement can then be sustained. Two classical corollaries: agents who announce their posteriors back and forth must converge to agreement [GP82]; and rational traders who begin from an efficient allocation cannot come to trade on private information alone, since common knowledge of willingness to trade is common knowledge of disagreement [MS82]. Every theory of speculative trade has to break one of the hypotheses, and most break the common prior.

Notes

The monad of probability measures is due to Giry [Gir82], following Lawvere; Markov categories were introduced under that name by Fritz [Fri20], building on [CJ19]. Games with incomplete information are Harsanyi's [Har67]. The universal type space was constructed by Mertens and Zamir [MZ85] and by Brandenburger and Dekel [BD93]; the measure-theoretic construction, and the notion of type morphism, are due to Heifetz and Samet [HeSa98], and the identification with a terminal coalgebra to Moss and Viglizzo [MV06]. An early domain-theoretic treatment is [Vas92]. The role of the common prior was analysed by Aumann [Aum87]; the characterization of its existence by the absence of agreeable bets is due to Morris [Mor94], Samet [Sam98] and Feinberg [Fei00]; the cycle equations are from [RN09] and the genericity analysis from [HeSa12]. The cohomological packaging of Theorem 5.7 is, as far as we know, not in the literature in this form, but every ingredient is.

Exercises

Exercise 5.1. In Example 5.2, let costs be uniform on $[0, \bar{c}]$ with $\bar{c} > 1$. Find the symmetric cutoff equilibrium and the probability that the good is provided as $\bar{c} \to \infty$ and as $\bar{c} \to 1$.

Exercise 5.2. Show that a type morphism between finite type spaces preserves first- and second-order beliefs, and deduce by induction that it preserves hierarchies.

Exercise 5.3. Verify the claims of Example 5.6, and exhibit explicitly the unique type morphism from $\mathbb{B}$ to $\mathbb{A}$.

Exercise 5.4. For three agents and the information structure of the last exercise of Chapter 4, write down the cycle equation and interpret it.

Exercise 5.5. In the setting of Example 5.8, suppose the cycle equation fails. Construct an agreeable bet explicitly, and check that its existence contradicts the existence of a common prior.

Exercise 5.6. Prove the tower property $\mathbb{E}_{rp} = \mathbb{E}_r \mathbb{E}_p$ on a finite probability space directly, and identify $\mathbb{E}_p$ with composition with the Bayesian inverse $p^\dagger_\mu$.

Chapter 6

Mixing and Existence

Matching pennies has no equilibrium in the sense of Definition 1.12. The classical remedy is to let players randomize, and Nash's theorem then guarantees an equilibrium in every finite game. This chapter presents the mixed extension as a free construction, Nash's theorem as a statement about the degree of a map, and correlated equilibrium as what one obtains by declining to apply a certain non-invertible natural transformation.

6.1 The mixed extension is free

Definition 6.1. A ***convex space*** is an algebra for the monad Δ: a set C with a map $\Delta C \to C$ assigning to each finite formal mixture of points a point, compatibly with units and with mixtures of mixtures. A map of convex spaces is ***affine***. The free convex space on a set X is ΔX itself.

Theorem 6.2 (Expected utility). [classical; [vNM44]] *Let $\succsim$ be a complete, transitive relation on ΔX, with X finite, satisfying independence and continuity. Then there is an affine map $U\colon \Delta X \to \mathbb{R}$ representing $\succsim$, and U is unique up to composition with a map $x \mapsto \alpha x + \beta$, $\alpha > 0$.*

The categorical reading is short. An affine map out of a free algebra is determined by its restriction to generators, so U *is* a function $u\colon X \to \mathbb{R}$, extended by expectation. The uniqueness clause says that the freedom in U is the group of automorphisms of $\mathbb{R}$ as an ordered convex space, which is the group of positive affine maps. That is where the preference equivalence of Definition 1.4 comes from: a point of $\mathsf{NFG}_{\succsim}$ carries, for each player, an order-and-convexity structure on lotteries, and utility functions are its coordinate charts.

Definition 6.3. The ***mixed extension*** of $G = (N, A, u)$ is the game with strategy sets ΔA_i and payoffs

$$\prod_i \Delta A_i \xrightarrow{\nabla} \Delta\Big(\prod_i A_i\Big) \xrightarrow{\overline{u}_i} \mathbb{R},$$

where ∇ forms the product distribution and $\overline{u}_i$ is the affine extension of u_i. A ***mixed equilibrium*** of G is a Nash equilibrium of its mixed extension.

The payoff is affine in each player's own strategy separately. Consequently a mixed strategy is a best response iff every pure strategy in its support is, and all of those earn the same payoff. This *indifference principle* is how mixed equilibria are computed.

Example 6.4 (Three computations)**.** (a) In the Battle of the Sexes of Example 1.13, let the row player choose O with probability p and the column player with probability q. Column is indifferent when $p = 2(1-p)$, row when $2q = 1-q$; so $p = \frac{2}{3}$, $q = \frac{1}{3}$, and each player's expected payoff is $\frac{2}{3}$, less than at either pure equilibrium.

(b) *Inspection.* An agent complies at cost e or cheats; an inspector inspects at cost c or does not. A cheater who is inspected pays a fine $F > e$, and the inspector gains $b > c$ from catching one. There is no pure equilibrium. The inspector is indifferent when the probability of cheating is $\gamma = c/b$; the agent is indifferent when the probability of inspection is $\iota = e/F$. Raising the fine leaves the frequency of cheating unchanged and lowers the frequency of inspection: in a mixed equilibrium each player's behaviour is pinned down by the *other's* payoffs.

(c) *Volunteering* [Die85]. Each of n bystanders may act at cost c; if anyone acts all receive $b > c$. In the symmetric mixed equilibrium each abstains with probability r where $b - c = b(1 - r^{n-1})$, so $r = (c/b)^{1/(n-1)}$, and the probability that no one acts is $(c/b)^{n/(n-1)}$, which *increases* with n towards c/b.

What does a mixed strategy mean? Four answers are in circulation: deliberate randomization; a description of an opponent's uncertainty about what a player will do; the distribution of pure strategies in a population; and Harsanyi's *purification* [Har73], in which each player has a small private payoff shock and plays a pure strategy of her type, the mixture being the shadow of unobserved heterogeneity. The last has a structural gloss. In a Markov category with enough structure, every kernel $1 \to A_i$ factors as a state on a "noise" object followed by a deterministic map [Fri20, §11]: a mixed action is a pure function of something private. Purification is the statement that this factorization can be chosen to be an equilibrium of a nearby game of incomplete information.

6.2 Existence

Theorem 6.5 (Nash). [classical; [Nas50a, Nas51]] *Every finite game has a mixed equilibrium.*

Proof sketch. Let $\Sigma = \prod_i \Delta A_i$, a product of simplices. For $\sigma \in \Sigma$ and $a_i \in A_i$ let $g_{a_i}(\sigma) = \max\{0, u_i(a_i, \sigma_{-i}) - u_i(\sigma)\}$ be the gain from deviating to a_i, and

define $f\colon \Sigma \to \Sigma$ by

$$f(\sigma)_i(a_i) = \frac{\sigma_i(a_i) + g_{a_i}(\sigma)}{1 + \sum_{b_i} g_{b_i}(\sigma)}.$$

f is continuous, so by Brouwer's theorem it has a fixed point, and one checks that the fixed points are exactly the points where all gains vanish. □

Three remarks place the theorem.

It is a topological theorem with no order-theoretic or algebraic proof. Chapter 7 will prove existence by Tarski's theorem for a different class of games; for general finite games one needs Brouwer, or its set-valued form due to Kakutani, or an equivalent.

It is not constructive. Brouwer's theorem is not provable in constructive mathematics; what is provable is the existence of approximate fixed points. In homotopy type theory the situation is instructive: Shulman [Shu18] proves Brouwer's theorem in a *cohesive* extension of the theory, in which the topology of the disc is an intrinsic part of the type rather than a structure on a set of points, and the fixed point is asserted in propositionally truncated form. One learns that a fixed point exists; one is not handed a function that finds it.

It is computationally hard. Finding a mixed equilibrium, even of a two-player game, is complete for the complexity class PPAD [DGP09, CDT09]—the class of total search problems whose totality is guaranteed by the parity argument that underlies Sperner's lemma, hence Brouwer. The complexity result is the quantitative form of the logical one: the only general reason an equilibrium exists is topological, and topology does not come with an algorithm.

6.3 The equilibrium correspondence

Fix the players and action sets and let the payoffs vary. The space of games is then a Euclidean space $\mathcal{U} = \mathbb{R}^{N \times A}$, and the object of interest is the ***graph of the equilibrium correspondence***

$$\mathcal{E} = \{(u, \sigma) \in \mathcal{U} \times \Sigma \mid \sigma \text{ is a mixed equilibrium of } u\},$$

with its projection $\mathrm{pr}\colon \mathcal{E} \to \mathcal{U}$. The fibre of pr over a game is its set of equilibria.

Theorem 6.6 (Structure theorem). [classical; [KM86]] *There is a homeomorphism $h\colon \mathcal{E} \to \mathcal{U}$ such that $\mathrm{pr} \circ h^{-1}$ is homotopic to the identity by a homotopy that extends to the one-point compactification. In particular* pr *is a proper map of degree $+1$.*

The total space of equilibria is as simple as possible—a Euclidean space, lying over the space of games with degree one, like a sheet of paper that has

been folded but not torn. All of the following are consequences of this one statement.

- *Existence.* A map of degree one is surjective. This is Nash's theorem.
- *Oddness.* Over a regular value the fibre is finite, and its points, counted with the sign of the local degree, sum to $+1$. Hence generic games have an odd number of equilibria [Wil71, Har73b].
- *Index.* The local degree at an isolated equilibrium, or more generally on a connected component of the fibre, is its ***index***; indices sum to $+1$ over every game. A component of nonzero index cannot disappear under small perturbations of the payoffs: it is *essential*.

Example 6.7. The ranked coordination game of Example 1.13 has three equilibria: the two pure ones, each of index $+1$, and the mixed one, of index -1; and $1 + 1 - 1 = 1$. The mixed equilibrium is the fold in the paper. Perturb the payoffs until the game has a dominant action, and the mixed equilibrium and one pure equilibrium approach each other, merge, and vanish, leaving a single equilibrium of index $+1$. Indices are computed numerically for this game in Appendix C.

What the structure theorem does not give is as important.

Proposition 6.8 (No continuous selection)**.** [observation] *There is no continuous map $s\colon \mathcal{U} \to \Sigma$ with $s(u)$ an equilibrium of u for all u, as soon as some player has two actions.*

Proof. Restrict to a one-parameter family of one-player games with two actions and payoffs $(t, 0)$. For $t > 0$ the unique equilibrium is the first action, for $t < 0$ the second. A selection has no limit at $t = 0$. □

So pr has degree one and no section. The equilibrium correspondence is upper but not lower hemicontinuous; its fibres change homotopy type from point to point. They can be arbitrarily complicated: by a theorem of Datta [Dat03], every real algebraic variety is isomorphic to the set of totally mixed equilibria of some three-player game. The situation is summarized by saying that pr is as far from a fibration as a degree-one map of Euclidean spaces can be.

Remark 6.9 (Existence is a truncated statement). [programme] In a type theory whose functions are continuous, Theorem 6.5 and Proposition 6.8 together say that the type $\prod_u \left\| \sum_\sigma \mathsf{isNash}(u, \sigma) \right\|$ is inhabited while $\prod_u \sum_\sigma \mathsf{isNash}(u, \sigma)$ is not: equilibria merely exist. This is more than a manner of speaking. It locates the problem of *equilibrium selection* exactly: a selection is an untruncation,

and by Proposition 1.14 and Proposition 6.8 it must be both equivariant and discontinuous. It also suggests what a refinement should be. An index, or the property of lying on an essential component, is a function on $\mathcal{E}$ that is locally constant where it should be and invariant under the homeomorphism h. The programme of characterizing refinements by such invariants was begun by Kohlberg and Mertens and carried furthest by Govindan and Wilson [GW05]; whether it can be phrased as the study of *homotopy-invariant* functions on the map pr, in a precise ∞-categorical sense, is taken up in Chapter 19.

6.4 Correlation

The mixed extension passes through the map $\nabla\colon \prod_i \Delta A_i \to \Delta(\prod_i A_i)$, which is not surjective. Suppose we do not apply it.

Definition 6.10. [classical; [Aum74]] A ***correlated equilibrium*** of G is a distribution $\pi \in \Delta(A)$ such that for every player i and every pair of actions a_i, a_i',

$$\sum_{a_{-i}} \pi(a_i, a_{-i}) \big[u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}) \big] \;\geq\; 0.$$

The interpretation is a mediator who draws a profile from π and privately recommends to each player her coordinate; the inequalities say that obeying is optimal given what the recommendation reveals.

Proposition 6.11. [classical; [Aum74, Aum87]; geometry as in [NGH04]] *The set* $\mathrm{CE}(G)$ *is a nonempty convex polytope in* $\Delta(A)$*, defined by linear inequalities, and well defined on* $\mathsf{NFG}_{\mathrm{inc}}$*. The mixed equilibria of* G *are*

$$\mathrm{NE}(G) = \mathrm{CE}(G) \;\cap\; \nabla\Big(\prod_i \Delta A_i\Big),$$

the intersection of the polytope with the variety of product distributions (the Segre variety of rank-one tensors).

The two halves of the right-hand side divide the difficulty of the subject. The polytope is linear programming: a correlated equilibrium can be computed in polynomial time, and the set is a convex body with a centre, extreme points and a geometry one can reason about. The Segre variety is where the nonlinearity, the multiplicity, the non-convexity of the Nash set and the PPAD-hardness all come from. Categorically, ∇ is the structure map that makes Δ a *monoidal* monad; it would be an isomorphism if Δ preserved products, and it is not. Independence is not a property that a joint distribution has for free. Nash equilibrium imposes it; correlated equilibrium does not.

Example 6.12. In the Battle of the Sexes, a public coin flip choosing between (O,O) and (F,F) is a correlated equilibrium worth $\frac{3}{2}$ to each player, against $\frac{2}{3}$ in the mixed equilibrium. It lies in the convex hull of the Nash equilibria. Aumann's example [Aum74] shows that correlated equilibria can leave that hull: with payoffs

	C	D
C	$6,6$	$2,7$
D	$7,2$	$0,0$

the distribution putting weight $\frac{1}{3}$ on each of (C,C), (C,D), (D,C) is a correlated equilibrium worth 5 to each player, which no mixture of Nash equilibria achieves.

A correlated equilibrium is also a Bayesian Nash equilibrium of a game of incomplete information in which the types are payoff-irrelevant signals with a common prior—exactly the redundant types of Example 5.6. Aumann's theorem [Aum87] that correlated equilibrium *is* the expression of common knowledge of rationality under a common prior is the same observation read in the other direction, and its generalization to genuine incomplete information is the *Bayes correlated equilibrium* of Bergemann and Morris [BM16], which reappears in Chapter 15 as the set of outcomes an information designer can induce.

6.5 Continuous games

Much of economics uses games with continuous action sets. If each A_i is a compact convex subset of a Euclidean space and each u_i is continuous, and quasi-concave in a_i, then a pure equilibrium exists: best responses are convex-valued with closed graph, and Kakutani's theorem applies (Debreu, Glicksberg, Fan). Without quasi-concavity one needs mixed strategies over a continuum, and existence follows from Glicksberg's infinite-dimensional version. Discontinuous payoffs—ubiquitous in price competition, where a slightly lower price captures the market—require a separate theory [DM86]; we will meet both the phenomenon and its cure in Chapters 8 and 18.

A continuous-time example shows mixing over a continuum. In the ***war of attrition*** two players compete for a prize v by waiting; waiting costs 1 per unit time to each, and the first to quit loses. In the symmetric equilibrium each player quits at a constant hazard rate $1/v$: over a short interval dt the cost of waiting, dt, is balanced by the chance dt/v that the opponent quits and yields v. Expected payoffs are zero. The prize is entirely dissipated by the contest for it, a theme that returns with rent-seeking and with signalling.

Notes

Convex spaces as algebras of the distribution monad go back to Stone and Świrszcz; see [Fri20] for references. Nash's two proofs are [Nas50a, Nas51]. The structure theorem is due to Kohlberg and Mertens [KM86]; generic oddness to Wilson [Wil71] and Harsanyi [Har73b]; the theory of index and essential components, and its equivalence with degree, to Govindan and Wilson (see [GW05] and its references). The connection between dynamic stability and index is [DR03]. The complexity of equilibrium is [DGP09, CDT09]. Correlated equilibrium was introduced in [Aum74]; [NGH04] studies the position of the Nash set in the correlated polytope. Purification is from [Har73]. The inspection game's comparative statics are sometimes called the "Robinson Crusoe fallacy"; the volunteer's dilemma is [Die85]; the war of attrition originates in evolutionary biology (Chapter 9).

Exercises

Exercise 6.1. Verify that the fixed points of the map f in the proof of Theorem 6.5 are exactly the mixed equilibria. Why can the gain g_{a_i} not be positive for every a_i in the support of σ_i?

Exercise 6.2. Find all mixed equilibria of the ranked coordination game of Example 1.13, and determine for which payoff perturbations of the form $u_1(b,b) \mapsto 1 + t$ the mixed equilibrium survives.

Exercise 6.3. Show that the set of correlated equilibria of the Prisoner's Dilemma is a single point, and that in any game an action that is strictly dominated receives zero probability in every correlated equilibrium.

Exercise 6.4. For the 2×2 game of Example 6.12, compute the polytope of correlated equilibria that are symmetric under exchange of the players, and find the one maximizing the sum of payoffs.

Exercise 6.5. In the volunteer's dilemma, suppose the n players can use a public correlating device. Describe a correlated equilibrium in which the good is always provided. Is it invariant under $\mathrm{Aut}(G)$? Find the best invariant correlated equilibrium.

Exercise 6.6. Show that $\nabla \colon \Delta X \times \Delta Y \to \Delta(X \times Y)$ is injective, natural in X and Y, and surjective only if X or Y has at most one element.

Chapter 7

Order

Nash's theorem needs topology and gives little information about the equilibria it produces. For a large and economically central class of games—those in which actions are ordered and a player wants to do more when others do more—existence, structure and comparative statics all follow from order alone. This chapter is about that class. Categorically it is the decategorified case of everything else in the book: categories with at most one arrow between two objects, where the theory of fixed points of functors collapses to Tarski's theorem.

7.1 Posets as categories

A preorder is a category with at most one morphism between any two objects; equivalently, a category enriched in the two-element lattice $\mathbf{2} = \{\bot \leq \top\}$. Under this dictionary a monotone map is a functor, a Galois connection is an adjunction, a meet is a limit and a join a colimit, and a complete lattice is a category that is complete and cocomplete. An algebra for a monotone $f\colon L \to L$ is a pre-fixed point $f(x) \leq x$, a coalgebra is a post-fixed point $x \leq f(x)$, and Lambek's lemma says that an initial algebra or a terminal coalgebra is a fixed point.

Theorem 7.1 (Knaster–Tarski). [classical; [Tar55]] *Let L be a complete lattice and $f\colon L \to L$ monotone. Then the set $\mathrm{Fix}(f)$ is a nonempty complete lattice. Its least element is $\mu f = \bigwedge\{x \mid f(x) \leq x\}$, the initial algebra, and its greatest is $\nu f = \bigvee\{x \mid x \leq f(x)\}$, the terminal coalgebra.*

We have already used the theorem twice: rationalizability is $\nu\Phi$ (Theorem 1.11) and common knowledge is $\nu X.K(E \cap X)$ (Theorem 4.6). The associated proof rules—*induction*: to show $\mu f \leq x$ show $f(x) \leq x$; *coinduction*: to show $x \leq \nu f$ show $x \leq f(x)$—are the thin-category versions of the universal properties of initial algebras and terminal coalgebras, which in Chapters 8 and 9 will govern finite and infinite horizons.

7.2 Complementarity

Definition 7.2. Let X be a lattice and T a poset. A function $f\colon X \to \mathbb{R}$ is ***supermodular*** if $f(x \vee y) + f(x \wedge y) \geq f(x) + f(y)$ for all x, y. A function $f\colon X \times T \to \mathbb{R}$ has ***increasing differences*** in (x, t) if $f(x', t) - f(x, t)$ is nondecreasing in t whenever $x' \geq x$.

On a product of chains, supermodularity is equivalent to increasing differences in each pair of coordinates, and for smooth functions on $\mathbb{R}^k$ to nonnegative cross partial derivatives. Increasing differences is a second-order condition: it constrains not how f varies but how the *variation* of f in one argument varies with another. It is the discrete form of complementarity.

Theorem 7.3 (Topkis). [classical; [Top78, Top98]] *Let X be a lattice, T a poset, and $f\colon X \times T \to \mathbb{R}$ supermodular in x with increasing differences in (x, t). Then $t \mapsto \arg\max_x f(x, t)$ is increasing in the strong set order: if $t \leq t'$, $x \in \arg\max f(-, t)$ and $x' \in \arg\max f(-, t')$, then $x \wedge x' \in \arg\max f(-, t)$ and $x \vee x' \in \arg\max f(-, t')$. In particular the greatest and least maximizers, when they exist, are monotone in t.*

The theorem makes no use of concavity, differentiability, convexity of X, or uniqueness of the maximizer. It converts a hypothesis about the objective directly into a conclusion about the selection function arg max: under complementarity, arg max is a *monotone* selection.

7.3 Supermodular games

Definition 7.4. [classical; [Top79, Viv90, MR90]] A ***supermodular game*** is a game in which each A_i is a complete lattice (in applications a compact sublattice of $\mathbb{R}^k$), and each u_i is supermodular and upper semicontinuous in a_i and has increasing differences in (a_i, a_{-i}).

Supermodularity of a game is preserved by the transformations $u_i \mapsto \alpha_i u_i + \beta_i(a_{-i})$ of Definition 1.4, so it is a property of a point of $\mathsf{NFG}_{\mathrm{inc}}$ together with an order on each action set.

Theorem 7.5. [classical; [Top79, Viv90, MR90, Zho94]] *In a supermodular game:*

(i) the set of pure Nash equilibria is a nonempty complete lattice; in particular there are a greatest equilibrium $\overline{a}$ and a least equilibrium $\underline{a}$;

(ii) every rationalizable action, and the support of every mixed or correlated equilibrium, lies in the order interval $[\underline{a}, \overline{a}]$; the game is dominance solvable iff the pure equilibrium is unique;

(iii) *if payoffs depend on a parameter t with increasing differences in (a_i, t), then $\overline{a}(t)$ and $\underline{a}(t)$ are nondecreasing in t.*

Proof sketch. By Theorem 7.3 the greatest best response $\overline{b}_i(a_{-i})$ is well defined and monotone. The map $\overline{b}\colon A \to A$ is monotone on a complete lattice, so by Theorem 7.1 it has a greatest fixed point, which is a Nash equilibrium and is $\geq$ every equilibrium; dually for the least. For (ii), iterate $\overline{b}$ from the top element: the sequence decreases to $\overline{a}$, and at each step the actions above the current point are strictly dominated. (iii) is Theorem 7.3 again, applied to the fixed-point operator. □

Compare with Chapter 6. Existence is in pure strategies, by an order-theoretic fixed-point theorem whose proof is a transfinite iteration rather than a topological degree; the extremal equilibria are found by iterating best responses from the top and the bottom, which is both an algorithm and a model of adaptive learning [MR90]; and comparative statics, which for general games founders on multiplicity and on the sign-indefiniteness of the index, is monotone for the extremal equilibria without further assumptions. The price is the hypothesis. But the hypothesis is satisfied by a remarkable share of the models economists care about: price competition with differentiated products, technology adoption, search, bank runs, arms races, team production.

7.4 Quantity competition

Example 7.6 (Cournot)**.** Two firms choose quantities $q_1, q_2 \geq 0$ of a homogeneous good with inverse demand $P = a - b(q_1 + q_2)$ and constant marginal cost $c < a$. Firm i's profit is $q_i(a - c - b(q_i + q_j))$, and her best response is

$$q_i = \frac{a - c - bq_j}{2b} \quad (\text{or } 0).$$

Best responses slope *downward*: quantities are ***strategic substitutes*** [BGK85]. The unique equilibrium is $q_1 = q_2 = \frac{a-c}{3b}$, with price $\frac{a+2c}{3}$. With n firms each produces $\frac{a-c}{(n+1)b}$ and the price $\frac{a+nc}{n+1}$ tends to marginal cost.

The Cournot duopoly has *decreasing* differences in (q_i, q_j), so it is not supermodular as it stands. But reverse the order on firm 2's action set and it is: with two players, strategic substitutes become strategic complements under the duality $L \mapsto L^{\mathrm{op}}$ applied to one of them. All of Theorem 7.5 then applies, with "greatest" meaning "most favourable to firm 1". The trick is categorical—an adjunction with one side dualized is still an adjunction—and it is special to two players: with three firms there is no consistent reversal, and genuinely different behaviour appears. The best-response map of the

linear duopoly has slope $-\frac{1}{2}$, so it is a contraction and iterated best responses converge; with three firms the slopes sum to -1 and the iteration cycles; with more it diverges [The60].

7.5 Price competition

Example 7.7 (Bertrand)**.** If instead the two firms choose prices and the good is homogeneous, the lower-priced firm takes the whole market. Any price pair above cost invites undercutting, and the unique equilibrium is $p_1 = p_2 = c$: two firms suffice for the competitive outcome. This is the *Bertrand paradox*. The payoff is discontinuous, and the conclusion is fragile: capacity limits, product differentiation, repeated interaction and incomplete information each overturn it, and Chapter 18 examines them.

With differentiated products, let demand be $q_i = a - p_i + s p_j$ with $0 \le s < 2$. Profit $(p_i - c)(a - p_i + s p_j)$ has best response

$$p_i = \frac{a + c + s p_j}{2},$$

which slopes *upward*: prices are ***strategic complements***. The game is supermodular, and its unique equilibrium is $p^* = \frac{a+c}{2-s}$. An increase in a or in c raises both prices, by part (iii) of the theorem, without any computation.

The distinction between substitutes and complements governs how a firm should try to influence a rival's later choice. If actions are strategic substitutes, a commitment to be aggressive makes the rival retreat; if complements, it makes the rival aggressive too. The whole taxonomy of strategic investment [FT84] is the sign of a cross partial derivative composed with the sign of the effect of the investment, and we return to it under the heading of commitment in Chapter 8.

7.6 Ordinal conditions

Supermodularity is a cardinal condition: it is not preserved by monotone transformations of the objective, although the conclusion of Topkis's theorem obviously is. From the standpoint of Principle 1.5 this is a defect of the hypothesis, and Milgrom and Shannon repaired it.

Theorem 7.8 (Monotone comparative statics)**.** [classical; [MS94]] *Let X be a lattice and T a poset. The correspondence $(t, S) \mapsto \arg\max_{x \in S} f(x, t)$ is increasing in t and in the sublattice S (in the strong set order) if and only if f is* ***quasi-supermodular*** *in x and satisfies the* ***single-crossing property*** *in (x, t): for $x' > x$ and $t' > t$, $f(x', t) \ge f(x, t)$ implies $f(x', t') \ge f(x, t')$, and likewise with strict inequalities.*

The conditions are ordinal, and the theorem is an equivalence: single crossing is exactly what monotone comparative statics requires, no more. For games with von Neumann–Morgenstern payoffs the relevant invariance is only affine and the cardinal conditions are appropriate; for decision problems under certainty the ordinal ones are. Single crossing will carry Part II: it is the condition under which types can be sorted by their choices, and so the condition under which screening and signalling work (Chapters 14 and 15).

7.7 The metric analogue

Lawvere observed [Law73] that a category enriched in the ordered monoid $([0,\infty],\geq,+)$ is a generalized metric space: the hom-object $d(x,y)$ is a distance, composition is the triangle inequality $d(x,y)+d(y,z)\geq d(x,z)$, and the identity is $0\geq d(x,x)$. Enriched functors are the maps that do not increase distance. The table compares the two enrichments.

	enriched in 2	**enriched in** $[0,\infty]$
object	preorder	generalized metric space
functor	monotone map	nonexpansive map
completeness	complete lattice	Cauchy completeness
good endomaps	monotone	contractions
fixed-point theorem	Tarski: a lattice of fixed points	Banach: a unique fixed point
proof	transfinite iteration	geometric convergence

Game theory uses both columns, and the difference between them is the difference between multiplicity and uniqueness. A monotone operator has a least and a greatest fixed point and possibly many between; a contraction has one. When a strategic situation is governed by a monotone operator—best responses in a supermodular game, the Abreu–Pearce–Stacchetti operator of Chapter 9, the unraveling operator of Chapter 13—expect a lattice of equilibria and ask which end is selected. When it is governed by a contraction—the Bellman operator of a discounted decision problem, the Shaked–Sutton recursion for alternating-offer bargaining in Chapter 16—expect uniqueness. Discounting is what turns the first kind of problem into the second, when it does.

Notes

Tarski's theorem is [Tar55]. The lattice-theoretic approach to optimization and games is due to Topkis [Top78, Top79]; his book [Top98] is the standard

reference. Vives [Viv90] and Milgrom and Roberts [MR90] developed supermodular games and their applications, and Zhou [Zho94] proved that the equilibrium set is a complete lattice. Strategic substitutes and complements were named by Bulow, Geanakoplos and Klemperer [BGK85]. The ordinal theory is from Milgrom and Shannon [MS94]. The instability of Cournot adjustment with three or more firms is Theocharis's observation [The60]. Lawvere's paper [Law73] is the source of the enriched view of metric spaces; the parallel between order-theoretic and metric fixed points is developed in the domain-theory literature, for instance [Rut96].

Exercises

Exercise 7.1. Prove Theorem 7.1 for finite lattices by iterating f from the top and bottom elements, and give an example of a monotone map on a non-complete lattice with no fixed point.

Exercise 7.2. Show that a function on a product of two chains is supermodular iff it has increasing differences, and that $f(x, y) = g(x + y)$ is supermodular iff g is convex.

Exercise 7.3. Verify the equilibrium quantities, price and profits in Example 7.6 for n firms, and show that total industry profit is decreasing in n.

Exercise 7.4. In the differentiated Bertrand game of Example 7.7, compute the greatest and least rationalizable prices by iterating the best-response map from the bounds of the interval $[c, \bar{p}]$ for large $\bar{p}$. Conclude that the game is dominance solvable.

Exercise 7.5. A coordination game: each of n players chooses effort $e_i \in \{0, 1, \dots, m\}$ and receives $2 \min_j e_j - e_i$. Show that the game is supermodular, describe its lattice of pure equilibria, and identify the greatest and least.

Exercise 7.6. Give an example of a function satisfying the single-crossing property but not increasing differences, and a monotone transformation of a supermodular function that is not supermodular.

Chapter 8

Perfection

A Nash equilibrium of a dynamic game may rest on a threat that would not be carried out. Selten's remedy [Sel65, Sel75] is to require equilibrium not only in the game but in every part of it that is itself a game. This chapter gives two descriptions of the resulting concept. For finite games of perfect information it is the value of a recursive function on trees. In general it is a modal operator applied to the Nash predicate—the temporal sibling of the knowledge operator of Chapter 4.

8.1 Finite games are an initial algebra

Fix a set O of outcomes (say $\mathbb{R}^N$) and a countable set of action names. Consider the polynomial endofunctor of Set

$$P(X) = O + \sum_{i \in N} \sum_{A} X^A,$$

where A ranges over finite nonempty sets of action names. An element of $P(X)$ is either an outcome or a triple (player, action set, assignment of an element of X to each action).

Proposition 8.1. [observation] *The initial algebra μP is the set of finite game trees of perfect information with outcomes in O, without moves of Nature. For every P-algebra $\alpha\colon P(V) \to V$ there is a unique map $[\![-]\!]_\alpha\colon \mu P \to V$ with $[\![\mathit{leaf}\, o]\!] = \alpha(o)$ and $[\![(i, A, (\Gamma_a)_{a \in A})]\!] = \alpha\big(i, A, ([\![\Gamma_a]\!])_a\big)$.*

The unique map is the *fold* determined by α; it is definition by recursion on the tree, and its uniqueness is proof by induction on the tree.

Theorem 8.2 (Zermelo–Kuhn)**.** [classical; [Kuh53]] *Every finite game of perfect information has a subgame perfect equilibrium in pure strategies. If no player is indifferent between two terminal nodes, it is unique.*

Proof. Take $V = \mathbb{R}^N$ and let α be the identity on outcomes and $\alpha(i, A, v) = v(a^*)$ for some $a^* \in \arg\max_{a \in A} v(a)_i$. The fold $[\![-]\!]_\alpha$ assigns to every tree its *backward-induction value*, and recording the maximizers chosen along the way gives a strategy profile that, by induction, is a Nash equilibrium of every subtree. Without ties α is a function with no choices to make. □

Backward induction is therefore not an algorithm that happens to solve games; it is *the* canonical map out of the type of finite games, and each choice of algebra is a different question asked of the tree. The algebra of the proof gives the value. An algebra on sets of values, taking all maximizers at each node, gives the set of all subgame-perfect payoffs. The algebra built from the product of selection functions gives Theorem 2.9. An algebra that at each node applies a Pareto filter before maximizing gives the renegotiation-proof values of Section 8.5. What they share is the hypothesis: the tree is an element of an *initial* algebra, so there is a last move from which to start. Chapter 9 is about what happens when there is not.

8.2 Subgame perfection as a modality

Now let Γ be any finite extensive game. Recall that a node h roots a *subgame* Γ_h when no information set straddles the boundary of its subtree. Subgame roots are preordered by precedence, and a strategy profile σ of Γ restricts to a profile $\sigma|_h$ of each Γ_h.

Definition 8.3. [classical; [Sel65]] σ is a ***subgame perfect equilibrium*** of Γ if $\sigma|_h$ is a Nash equilibrium of Γ_h for every subgame root h.

Let W be the set of pairs (h, σ)—a profile together with a subgame at which to evaluate it—and let $\mathsf{Nash} \subseteq W$ be the set of pairs with $\sigma|_h \in \mathrm{NE}(\Gamma_h)$. For a subset $Q \subseteq W$ define

$$\Box Q = \{(h,\sigma) \mid (h',\sigma) \in Q \text{ for every subgame root } h' \text{ at or after } h\}.$$

Proposition 8.4 (Perfection is necessity)**.** [observation]

(i) $\Box$ *is an idempotent comonad on* $\mathcal{P}(W)$ *preserving all intersections; it satisfies the S4 axioms* $\Box Q \subseteq Q$ *and* $\Box Q = \Box\Box Q$.

(ii) $\Box Q$ *is the largest subset of* Q *that is* hereditary*: closed under passing to later subgames.*

(iii) $\mathrm{SPE}(\Gamma) = \{\sigma \mid (\mathrm{root}, \sigma) \in \Box\,\mathsf{Nash}\}$.

(iv) $\Box = p^* \forall_p$*-type base change: for the projection* $p \colon W^{\leq} \to W$ *from the set of pairs* $\big((h,\sigma),h'\big)$ *with* h' *at or after* h*, and the evaluation* $e \colon W^{\leq} \to W$*,* $e((h,\sigma),h') = (h',\sigma)$*, one has* $\Box = \forall_p e^*$.

Proof. All four are immediate from the definition; (i) is the general fact that the universal modality of a preorder is an S4 necessity, and (iv) exhibits it as a composite of right adjoints. □

Compare with Definition 4.2. Knowledge is $p^*\forall_p$ for the map from states to an agent's observations: a proposition is known when it holds at every state the agent cannot distinguish from the actual one. Perfection is the same construction for the relation "is a later subgame of": a profile is perfect when it is an equilibrium at every subgame that could be reached. The first modality is S5 because indistinguishability is symmetric; the second is only S4 because time is not. *Subgame perfection is to histories what knowledge is to states.* The parallel is exact enough to predict the next refinement: when no subgame begins at an information set, one must combine the two modalities and ask for optimality at every information set given what the player believes there. That is sequential rationality (Chapter 10).

Item (ii) deserves a second look. The Nash predicate is not hereditary—an equilibrium of the whole game may restrict to nonsense off the path—and $\Box$ repairs this in the only possible way, by passing to the hereditary core. Corollary 3.10 can now be read as saying that the preorder of subgames, and with it the modality $\Box$, is not an invariant of the reduced strategic form.

Example 8.5 (Entry)**.** An entrant chooses *Out*, ending the game with payoffs $(0, 4)$, or *In*; the incumbent then chooses *Accommodate*, giving $(1, 2)$, or *Fight*, giving $(-1, 0)$. The strategic form has two pure equilibria, $(\mathit{In}, \mathit{Accommodate})$ and $(\mathit{Out}, \mathit{Fight})$. The second is Nash because the incumbent's choice is never tested; at the subgame after *In* it prescribes an action that is not a best response. Only $(\mathit{In}, \mathit{Accommodate})$ is subgame perfect, and it is the backward-induction solution. This is the equilibrium that the composite of open games in Example 2.7 failed to exclude.

8.3 Commitment

The incumbent of Example 8.5 would like the entrant to believe that entry will be fought. Subgame perfection says the entrant will not believe it, because once entry has occurred fighting is not in the incumbent's interest. The only way to make the threat credible is to change the game: to take an earlier, observable, irreversible action after which fighting *is* optimal, or accommodation unavailable. This is Schelling's insight [Sch60]: *a player can gain by reducing her own options or her own payoffs*. In a decision problem, deleting alternatives never helps. In a game it can, because it changes the context that others face. Formally a commitment device is a morphism of games—a new first stage composed in front—and its value is the difference between two backward-induction values.

Example 8.6 (Stackelberg)**.** In the linear market of Example 7.6, let firm 1 choose its quantity first and firm 2 observe it. Firm 2 plays its best response

$q_2 = \frac{a-c-bq_1}{2b}$, so firm 1 maximizes $q_1(a - c - bq_1 - bq_2) = \frac{1}{2}q_1(a - c - bq_1)$ and chooses $q_1 = \frac{a-c}{2b}$; then $q_2 = \frac{a-c}{4b}$. The leader's profit $\frac{(a-c)^2}{8b}$ exceeds the Cournot profit $\frac{(a-c)^2}{9b}$, and the follower's is lower. Quantities are strategic substitutes: committing to produce more makes the rival produce less. Had the variables been strategic complements, as prices are in Example 7.7, the leader would commit to be soft, and both would gain.

In the calculus of Chapter 2 the difference between the Cournot and Stackelberg games is one wire: the copy of q_1 that runs into firm 2's decision in Figure 2.2. Commitment is observability. Accordingly its value is fragile to noise in the wire: if the follower observes the leader's choice only imperfectly, however slightly, the pure-strategy equilibrium reverts to the simultaneous-move outcome [Bag95]. The algebraic reading is the Escardó–Oliva product of Section 2.4: the leader's selection function is applied to the context that already contains the follower's response.

Sunk costs and threats

Consider a plaintiff with a claim of expected value J at trial, where going to trial costs her c_P and the defendant c_D. If $J < c_P$ her threat to litigate is empty and backward induction gives a settlement of zero: she will drop the suit, so the defendant offers nothing. Now suppose that she can pay part s of her costs before the settlement stage. Sunk costs are sunk: at the settlement stage only the remaining cost $c_P - s$ matters. If $J > c_P - s$ the threat has become credible, and the defendant will pay up to $J + c_D$ to avoid trial. A suit with negative expected value at the outset can extract a positive settlement, because the order in which costs are incurred is part of the tree [Beb96]. The example repeats the lesson of the Thompson moves: what is strategically irrelevant in a decision problem—when a cost is paid—is relevant in a game.

8.4 The open-set problem

Suppose the plaintiff makes a take-it-or-leave-it demand S and the defendant accepts if $S < J + c_D$, rejects if $S > J + c_D$. What does the defendant do at $S = J + c_D$? If she rejects when indifferent, the plaintiff wishes to demand the largest number strictly below $J + c_D$, and there is none: the plaintiff's maximization is over an open set and the supremum is not attained. No best response exists, hence no equilibrium. The only subgame perfect equilibrium has the defendant accept when indifferent.

This is usually presented as a technicality, and resolved by a convention. It is better understood as a theorem about selection functions. arg max is a total

selection function on a space of alternatives only when maxima exist, which for continuous contexts means that the space is compact—*searchable*, in the terminology of Section 2.4. When a player's feasible set depends on another's tie-breaking rule, existence of equilibrium determines the tie-breaking rule: indifference must be resolved in the way that closes the first mover's feasible set. Simon and Zame [SZ90] turned this observation into a general existence theorem for discontinuous games in which the rule for resolving ties is part of the solution rather than part of the data. The economic content is that an equilibrium cannot rely on a party's rejecting an offer to which she is indifferent; the mathematical content is that selection functions exist on closed sets.

8.5 Renegotiation

Subgame perfection requires each subgame to be played by *some* equilibrium. When a subgame has several, the choice among them can depend on history, and this is how punishments are sustained: play the bad equilibrium if someone deviated. But if the players can talk when they reach the subgame, why would they agree to play an equilibrium that is worse for all of them?

Example 8.7. Two players play the Prisoner's Dilemma of Example 1.13 and then, having observed the outcome, a coordination game with equilibrium payoffs $(2,2)$ and $(0,0)$. The profile "cooperate; then play the good equilibrium if both cooperated and the bad one otherwise" is subgame perfect: defecting gains 1 in the first stage and loses 2 in the second. But after a defection the players stand at the start of a coordination game in which both prefer the good equilibrium. If they can renegotiate, the bad equilibrium will not be played, the punishment is void, and cooperation in the first stage unravels.

Pareto perfection [BPW87, BR89] formalizes this for finite-horizon games by a recursion on the tree: a profile is Pareto perfect if in every subgame the continuation is Pareto perfect and not Pareto-dominated by another profile that is Pareto perfect in all later subgames. In the language of the first section it is a fold whose algebra filters by Pareto dominance before choosing. For infinite horizons no last subgame anchors the recursion and several inequivalent definitions compete [FM89]; this is the first sign of the difficulty that occupies the next chapter. Renegotiation has the unusual property of being a refinement that makes the players worse off, because the ability to reopen an agreement destroys the ability to commit to a punishment. We will meet the same phenomenon in contract theory (Chapter 12) and in durable-goods pricing (Chapter 18).

Notes

Subgame perfection was introduced by Selten [Sel65] and extended to trembling-hand perfection in [Sel75]. Backward induction is usually traced to Zermelo's analysis of chess; the general theorem for finite trees is in [Kuh53]. The treatment of game trees as an initial algebra is standard in the semantics of programming languages and is the basis of the coalgebraic approach to infinite games discussed in Chapter 9. Commitment is the central theme of Schelling's book [Sch60]; the fragility of first-mover advantage to imperfect observation is due to Bagwell [Bag95]. The credibility of threats to sue when costs are incurred in stages is analysed in [Beb96]. Endogenous sharing rules are from [SZ90]. Renegotiation-proofness and Pareto perfection are developed in [BPW87, BR89, FM89]; the finite-horizon example is in the spirit of [BK85].

Exercises

Exercise 8.1. Write the entry game of Example 8.5 as an element of μP and compute its fold for (a) the value algebra, (b) the algebra of sets of values, (c) the algebra that computes the number of terminal nodes.

Exercise 8.2. Prove items (i)–(iii) of Proposition 8.4, and show by example that Nash is not hereditary.

Exercise 8.3. In Example 8.6, verify the profits, and compute the Stackelberg outcome for the differentiated price game of Example 7.7. Who gains from moving first, and who gains from the existence of a first mover?

Exercise 8.4. A seller makes a take-it-or-leave-it price offer to a buyer whose value v is common knowledge. Show that in the unique subgame perfect equilibrium the buyer accepts when indifferent. Now restrict prices to a grid of mesh ϵ and describe all subgame perfect equilibria.

Exercise 8.5. In Example 8.7, change the bad second-stage equilibrium to pay $(3,0)$ after a defection by player 2 and $(0,3)$ after a defection by player 1, where the coordination game is modified to have these as equilibria. Is the cooperative profile Pareto perfect?

Exercise 8.6. Show that in a finite game of perfect information without ties, the backward-induction profile is the only profile surviving iterated elimination of weakly dominated strategies in a suitable order, and discuss this in light of Example 1.10.

Chapter 9

Repetition

When the same players meet repeatedly, today's behaviour can be rewarded or punished tomorrow, and cooperation that is impossible in one encounter can be sustained. Whether it *is* sustained in equilibrium turns, notoriously, on whether there is a last encounter. This chapter argues that the dependence is not a quirk. A finitely repeated game is an element of an initial algebra and is solved by induction; an infinitely repeated game lives in a terminal coalgebra and is solved by coinduction; and the two universal properties give different answers.

9.1 Finite repetition

Let $G = (N, A, u)$ be a stage game. In the ***repeated game*** G^T the players play G in periods $t = 1, \dots, T$, observing all past actions; a history is an element of $A^* = \coprod_t A^t$, a pure strategy of i is a function $A^* \to A_i$, and payoffs are discounted sums of stage payoffs.

Theorem 9.1. [classical] *If G has a unique Nash equilibrium a^*, then for every finite T the unique subgame perfect equilibrium of G^T is to play a^* after every history.*

Proof. Fold over the tree (Theorem 8.2, in the version for simultaneous-move stages): in the last period only a^* is an equilibrium whatever the history; given that, the continuation value is independent of play in period $T-1$, so only a^* is an equilibrium there; and so on. □

In the Prisoner's Dilemma repeated a thousand times, the only subgame perfect equilibrium is a thousand defections. Selten's *chain-store paradox* [Sel78] makes the discomfort vivid. An incumbent faces the entry game of Example 8.5 in twenty towns in sequence. Induction says that she accommodates in the last town, hence in the nineteenth, hence in all, and that every entrant enters. Everyone's intuition says that she fights early entrants to deter later ones, and that it works. Selten took the paradox to show a gap between the theory and reasonable behaviour. We will take it to show that the intuition is computing a different fixed point from the induction. (When G has several equilibria the theorem fails and finite repetition can sustain a good deal [BK85]; Example 8.7 was an instance.)

9.2 Strategies are coalgebras

Turn to the infinitely repeated game $G^\infty(\delta)$, with payoffs $U_i = (1-\delta)\sum_{t\geq 0}\delta^t u_i(a^t)$ for a discount factor $\delta \in (0,1)$.

Definition 9.2. An ***automaton*** for player i is a set Q of states with an initial state, an output map $f\colon Q \to A_i$ and a transition map $\tau\colon Q \times A \to Q$. Equivalently it is a pointed coalgebra $Q \to F(Q)$ for the functor $F(X) = A_i \times X^A$ on Set.

Proposition 9.3. [literature; [Rut00]] *The terminal F-coalgebra is the set $A_i^{A^*}$ of functions from histories to actions—the set of pure strategies of the repeated game—with output "evaluate at the empty history" and transition "restrict to histories beginning with a". The unique coalgebra map from an automaton sends a state q to the strategy it implements when started at q.*

So the textbook definition of a strategy (a function on histories) and the working definition (a machine) are related as a terminal coalgebra to an arbitrary one. The image of an automaton in $A_i^{A^*}$ is its set of distinct *continuation strategies*, and the minimal automaton implementing a strategy has exactly one state for each. Kalai and Stanford [KS88] take this number as the *complexity* of a strategy; it is the coalgebraic notion of the size of a behaviour, and the minimal automaton is the quotient by bisimilarity. *Grim trigger*—cooperate until anyone defects, then defect forever—has two states. So does *tit-for-tat*. A strategy of finite complexity is one whose orbit under the transition structure of the terminal coalgebra is finite.

9.3 Perfection, coinductively

Subgame perfection is defined as in Section 8.2: every history roots a subgame, and σ is subgame perfect when $\sigma|_h$ is a Nash equilibrium of the continuation game for every h. But there are now infinitely many subgames and no last one, and checking the definition seems to require comparing σ_i with every one of uncountably many alternative strategies at each of them.

Theorem 9.4 (One-shot deviation principle). [classical; [Bla65, Abr88]; see [FT91]] *In a game that is continuous at infinity—in particular in $G^\infty(\delta)$ with bounded stage payoffs—a profile σ is subgame perfect if and only if no player can gain at any history h by deviating from σ_i at h alone and conforming to it thereafter.*

Here is the coinductive reading. For a set Q of profiles let $\Psi(Q)$ be the set of profiles σ such that (a) no one-shot deviation at the initial history is profitable, payoffs being computed by letting σ run, and (b) each continuation

$\sigma|_a$ after a first-period action profile a lies in Q. Ψ is monotone, and its greatest fixed point $\nu\Psi$ is the set of profiles with no profitable one-shot deviation at any history. The theorem says SPE $= \nu\Psi$. One inclusion is trivial. The other is where the hypothesis enters: a profitable deviation involving infinitely many changes must be approximated by a profitable deviation involving finitely many, which is then reduced to a single one by backward induction over the finitely many. The approximation is *continuity at infinity*: payoffs depend arbitrarily little on what happens after a sufficiently late date.

Example 9.5 (Why the hypothesis is needed)**.** A single agent chooses each period whether to stop. If she ever stops she gets 0; if she never stops she gets 1. The strategy "stop at every history" admits no profitable one-shot deviation: continuing once and stopping next period still yields 0. But it is not optimal. The payoff is discontinuous at infinity: the reward for never stopping is invisible at every finite stage. Coinduction certifies a profile that the semantics rejects.

9.4 Discounting is a guard

Continuity at infinity has a clean formulation. Give the set A^ω of infinite plays the ultrametric $d(x, y) = 2^{-t}$, where t is the first period at which x and y differ. With bounded stage payoffs and $\delta < 1$, the payoff U_i changes by at most $\delta^t \cdot \text{const}$ between plays that agree up to t: it is uniformly continuous. In the terms of Chapter 3, A^ω is the set of global sections of the tree of histories, and a payoff continuous at infinity is one that can be approximated, uniformly, by functions on the finite levels of the tree.

The topos of trees $\widehat{\omega}$ has an endofunctor $\blacktriangleright$, "later", with $(\blacktriangleright X)(n+1) = X(n)$ and $(\blacktriangleright X)(0) = 1$, and a natural map $X \to \blacktriangleright X$. A map is *contractive* if it factors through $\blacktriangleright$: its output at stage $n+1$ depends only on its input at stage n.

Theorem 9.6. [literature; [BMSS12]] *In $\widehat{\omega}$ every contractive map $f\colon \blacktriangleright X \to X$ has a unique fixed point. Complete bounded ultrametric spaces with distances in $\{2^{-n}\}$ and nonexpansive maps embed fully in $\widehat{\omega}$, contractive maps corresponding to maps that factor through $\blacktriangleright$; under the embedding the theorem is Banach's.*

Dynamic programming is the same theorem with real-valued distances. The value of a stationary decision problem satisfies $V = \max_a\{(1-\delta)u(a, s) + \delta V(s')\}$; the right-hand side is an operator on bounded functions that is monotone and discounts, hence is a δ-contraction in the sup norm, hence has a unique fixed point [Bla65]. What δ does is put the continuation value *under a guard*: tomorrow's value enters today's only after being shrunk. In the

enriched language of Section 7.7, δ makes the Bellman operator a contraction of a Lawvere metric space; in the topos-theoretic language it is a quantitative form of $\blacktriangleright$. Either way, a single decision maker's infinite-horizon problem has a unique value, and it is the limit of the finite-horizon values: under a guard, the inductive and the coinductive solutions agree.

For games they do not, and the reason is that equilibrium is not governed by a contraction.

9.5 Self-generation

The fundamental recursive tool for repeated games works not with a value but with the *set* of equilibrium values. For simplicity consider pure strategies and perfect monitoring.

Definition 9.7. [classical; [APS90, Abr88]] For $W \subseteq \mathbb{R}^N$ let $B(W)$ be the set of payoff vectors v for which there are an action profile a and a map $w\colon A \to W$ of *promised continuation values* with

$$v = (1-\delta)u(a) + \delta\, w(a),$$
$$(1-\delta)u_i(a) + \delta\, w_i(a) \geq (1-\delta)u_i(a_i', a_{-i}) + \delta\, w_i(a_i', a_{-i})$$

for all i and a_i': *promise keeping* and *incentive compatibility*. A set W is ***self-generating*** if $W \subseteq B(W)$.

Theorem 9.8 (Abreu–Pearce–Stacchetti). [classical; [APS90]] *Let V^* be the set of subgame-perfect payoff vectors of $G^\infty(\delta)$.*

(i) (Self-generation.) If W is bounded and $W \subseteq B(W)$ then $B(W) \subseteq V^*$.

(ii) (Factorization.) $V^* = B(V^*)$.

(iii) B is monotone and preserves compactness; V^* is compact and equals $\bigcap_n B^n(W_0)$ for any compact $W_0 \supseteq V^*$.

Thus $V^* = \nu B$, the greatest fixed point of B on bounded sets.

Part (i) is the coinduction rule of Theorem 7.1: a post-fixed point lies below the greatest fixed point. But more is true, and it ties this section to the one on automata. A self-generating set W together with a choice, for each $v \in W$, of the pair (a, w) that generates it is a map

$$W \longrightarrow A \times W^A,$$

that is, a *coalgebra* for the strategy functor, with state space the set of promised values. Its unique map to the terminal coalgebra is the equilibrium strategy

profile: start in state v, play $a(v)$, move to state $w(v)$(what was played). Incentive compatibility makes one-shot deviations unprofitable in every state, and Theorem 9.4 does the rest. "Use continuation values as state variables", the slogan of recursive methods in dynamic games and contracts, is the instruction to build a coalgebra.

Example 9.9 (Grim trigger)**.** In the Prisoner's Dilemma with payoffs $T = 3$, $R = 2$, $P = 1$, $S = 0$, the two-point set $W = \{(2,2),(1,1)\}$ is self-generating iff $\delta \geq \frac{1}{2}$. The value $(1,1)$ is generated by (D,D) with $w \equiv (1,1)$. The value $(2,2)$ is generated by (C,C) with $w(C,C) = (2,2)$ and $w = (1,1)$ otherwise; incentive compatibility reads $2 \geq (1-\delta)3 + \delta$, that is, $\delta \geq \frac{T-R}{T-P} = \frac{1}{2}$. The coalgebra on W is the two-state automaton of grim trigger.

9.6 The folk theorem

Let $\underline{v}_i = \min_{\alpha_{-i}} \max_{a_i} u_i(a_i, \alpha_{-i})$ be player i's ***minmax*** value, the worst the others can hold her to, and let $\mathcal{F}^*$ be the set of feasible payoff vectors (convex combinations of stage payoffs) with $v_i > \underline{v}_i$ for all i.

Theorem 9.10 (Folk theorem)**.** [classical; [Fri71, FM86]] *Suppose $\mathcal{F}^*$ has dimension $|N|$ (or $|N| = 2$). Then for every $v \in \mathcal{F}^*$ there is $\underline{\delta} < 1$ such that v is a subgame-perfect payoff of $G^\infty(\delta)$, with public randomization, for all $\delta \geq \underline{\delta}$.*

So for patient players vB is as large as it could possibly be: everything feasible and individually rational. Compare Theorem 9.1. The subgame-perfect payoff sets of the finite repetitions are generated *inductively*—start from the stage equilibria in the last period and apply the finite-horizon analogue of B—and with a unique stage equilibrium they remain a single point for every horizon, so their limit is a point. The infinite-horizon set is defined *coinductively*, as the largest self-consistent set. The same operator, the two extreme fixed points:

> *The folk theorem is the gap between the least and the greatest fixed point of the self-generation operator.*

The chain-store intuition computes the greatest fixed point; backward induction computes the least. Neither is an error. The modelling question is which universal property the situation has—whether the players' interaction is well-founded. A known last period makes it so. A horizon that is uncertain at every date is a discount factor and does not. And the three classical escapes from the paradox each move the problem from the first kind to the second, or blur the line: an infinite horizon; bounded complexity, so that players cannot count down from the end [Ney85]; and a grain of incomplete information about the opponent, which is the subject of Chapter 10.

9.7 Markov perfection

The folk theorem's equilibria condition on history in ways that have no bearing on the future: after any history of the repeated game, the continuation game is the *same game*. Maskin and Tirole proposed to forbid this.

Definition 9.11. [classical; [MT01]] In a dynamic game with observable actions, two histories are ***payoff-equivalent*** if the continuation games they root are strategically equivalent, compatibly at all future histories; the classes of the coarsest such equivalence are the ***Markov states***. A strategy is ***Markov*** if it is measurable with respect to this partition, and a ***Markov perfect equilibrium*** is a subgame perfect equilibrium in Markov strategies.

Proposition 9.12 (Markov means invariant)**.** [reformulation; cf. [MT01]] *The dynamic game is a coalgebra whose states are histories; the Markov states are its quotient by bisimilarity, the image of the histories in the terminal coalgebra of game behaviours. A strategy is Markov iff it factors through this quotient, that is, iff it is a function on the type of continuation games* with its identity type, *in the sense of Principle 1.5.*

Markov perfection is thus the invariance principle of Chapter 1 applied to subgames: equivalent subgames must be played alike. In a repeated game every continuation game is equivalent to every other, a Markov strategy is constant, and the Markov perfect equilibria are the repetitions of stage equilibria. All of the folk theorem's cooperation is therefore sustained by strategies that distinguish indistinguishable subgames. This is Section 1.5 once more. The history is a *frame*: structure that the players attach to a subgame, irrelevant to its payoffs, by which they coordinate on one continuation equilibrium rather than another. Whether such frames are admissible is, exactly as for focal points, a decision about the type of the game. In models where the state is genuinely payoff-relevant—capacity, installed base, reputation, the identity of the firm whose customers are locked in by switching costs—Markov perfection gives sharp predictions, and it is the workhorse of dynamic industrial organization [MT88].

9.8 Evolution

A different justification of equilibrium dispenses with reasoning altogether. A large population is matched at random to play a symmetric game; strategies that do well multiply.

Definition 9.13. [classical; [MSP73]] A mixed strategy x is an ***evolutionarily stable strategy*** if for every $y \neq x$ either $u(x,x) > u(y,x)$, or $u(x,x) = u(y,x)$ and $u(x,y) > u(y,y)$.

Example 9.14 (Hawk–Dove)**.** Two animals contest a resource worth V. Hawks fight, at cost $C > V$ to the loser; doves share or yield. Payoffs to the row player are $\frac{V-C}{2}$ for hawk against hawk, V for hawk against dove, 0 for dove against hawk and $\frac{V}{2}$ for dove against dove. Neither pure strategy is stable; the unique evolutionarily stable strategy plays hawk with probability V/C. The more costly the fight, the rarer it is.

Under the *replicator dynamics* $\dot{x}_a = x_a\big(u(a,x) - u(x,x)\big)$, every evolutionarily stable strategy is asymptotically stable, every stable rest point is a Nash equilibrium, and an asymptotically stable component of equilibria has index equal to its Euler characteristic—$+1$ for a point—in the sense of Section 6.3 [DR03]. The index, which entered as a topological invariant of the equilibrium correspondence, is thus also the shadow of dynamic stability. Axelrod's tournaments [Axe84], in which tit-for-tat prospered among strategies submitted for the repeated Prisoner's Dilemma, belong here: they are an evolutionary argument for a particular point of the folk theorem's large set, and like all such arguments they depend on the population of strategies one starts from.

9.9 Computability

The terminal coalgebra $A_i^{A^*}$ is uncountable; the strategies that any procedure could implement are countable. One might hope that restricting to computable strategies is harmless. It is not.

Theorem 9.15. [classical; [Rab57, Kno94, NZ96]]

(i) There is a game with computable rules in which a player has a winning strategy but no computable winning strategy.

(ii) In the discounted repeated Prisoner's Dilemma there are computable strategies to which a best response exists but no best response is computable.

The proofs are diagonal: a strategy can be built to consult what a given machine would do and to punish it, so that the task of responding optimally embeds an undecidable problem. A player modelled as a Turing machine cannot in general predict a player like herself [Bin87]. This is the third appearance of the diagonal in this book—after the nonexistence of a universal possibility structure (Proposition 5.5) and alongside the non-constructive character of Nash's theorem—and it marks a real boundary: the equilibrium concept quantifies over the whole terminal coalgebra, while agents inhabit its computable part.

Notes

Standard references for repeated games are [MS06, FT91]. The chain-store paradox is [Sel78]. The coalgebraic view of automata is expounded in [Rut00]; automata as strategies were introduced by Rubinstein [Rub86] and their complexity analysed in [KS88, Ney85]. The one-shot deviation principle descends from Blackwell's dynamic programming [Bla65]. Coinductive treatments of infinite games and their equilibria, including cases without discounting, are [LP12, AW17]; the infinite product of selection functions of Section 2.4 is another. The topos of trees and its guarded fixed points are from [BMSS12]. Self-generation is due to Abreu, Pearce and Stacchetti [APS90], building on [Abr88]. The folk theorem with Nash reversion is Friedman's [Fri71]; the general perfect-monitoring result is Fudenberg and Maskin's [FM86]. Markov perfect equilibrium is defined in general in [MT01]. Evolutionary stability is from [MSP73]. Computability limits on strategic reasoning are discussed in [Rab57, Bin87, Kno94, NZ96].

Exercises

Exercise 9.1. Write tit-for-tat and "win-stay, lose-shift" as coalgebras for $F(X) = A_i \times X^A$ and compute the unique maps to the terminal coalgebra on the histories of length at most two.

Exercise 9.2. In the Prisoner's Dilemma of Example 9.9, for which δ is the profile "both play tit-for-tat" a Nash equilibrium? Show that it is not subgame perfect for generic δ, by exhibiting a history at which a one-shot deviation is profitable.

Exercise 9.3. Compute $B(W)$ for $W = \{(1,1)\}$ and for $W = \{(2,2)\}$ in the Prisoner's Dilemma, and conclude that $\{(2,2)\}$ is not self-generating for any δ.

Exercise 9.4. Show that the operator B is monotone, and give an example of a stage game and a δ for which $\bigcup_n B^n(\varnothing) \neq \bigcap_n B^n(W_0)$.

Exercise 9.5. Consider the entry game of Example 8.5 repeated infinitely against a sequence of short-lived entrants, with the incumbent discounting at δ. Find the least δ for which there is a subgame perfect equilibrium with no entry on the path.

Exercise 9.6. Verify that in Hawk–Dove the strategy "hawk with probability V/C" satisfies the definition of evolutionary stability, and compute the mean fitness of the population at that point.

Chapter 10

Dynamic Games of Incomplete Information

When a player moves at an information set that is not the root of a subgame, subgame perfection has nothing to say about her, and yet she may be making an incredible threat. To test her choice one must say what she believes about where she is. This chapter combines the temporal modality of Chapter 8 with the beliefs of Chapter 5; studies the geometry of beliefs at events of probability zero; and then shows, through reputation and the electronic mail game, how violently equilibrium can respond to small changes in what is believed and known.

10.1 Sequential rationality

Definition 10.1. [classical; [KW82a]] An ***assessment*** in an extensive game is a pair (σ, μ) of a behavioural strategy profile and a *system of beliefs*: for each information set I a distribution $\mu(I)$ over the nodes of I. It is ***sequentially rational*** if at every information set I of every player i, the continuation of σ_i maximizes i's expected payoff given $\mu(I)$ and σ_{-i}.

A system of beliefs is a section of the bundle $\coprod_I \Delta(I) \to \mathcal{I}$, just as a strategy is a section of the bundle of actions. Sequential rationality is the necessity modality of Proposition 8.4 with subgames replaced by information sets, each equipped with the belief that turns it into a decision problem: perfection and knowledge, composed.

Example 10.2 (Where subgame perfection is silent)**.** An entrant chooses *Out*, with payoffs $(0, 4)$, or enters at a *small* or a *large* scale. The incumbent observes entry but not its scale, and fights or accommodates. Accommodating small entry gives $(1, 2)$, fighting it $(-1, 3)$; accommodating large entry gives $(2, 1)$, fighting it $(-2, -1)$. There is no proper subgame. Let m be the incumbent's belief that entry is small. Fighting is optimal iff $3m - (1 - m) \geq 2m + (1 - m)$, that is $m \geq \frac{2}{3}$. There are two kinds of sequentially rational assessment: the entrant enters large, the incumbent accommodates, $m = 0$; and the entrant stays out, the incumbent would fight, sustained by any belief $m \geq \frac{2}{3}$ at an information set that is never reached. Had fighting been worse than

accommodating against both scales, no belief would support it, and the second equilibrium would disappear although it remains subgame perfect.

On the equilibrium path beliefs are determined by Bayes' rule. Off the path, the conditioning event has probability zero and Bayes' rule is silent; in the example any m is allowed. The various definitions of ***perfect Bayesian equilibrium*** [FT91b] add requirements of the form "update by Bayes' rule wherever possible". Kreps and Wilson's definition is the cleanest.

Definition 10.3. [classical; [KW82a]] An assessment (σ, μ) is ***consistent*** if it is the limit of assessments (σ^k, μ^k) in which σ^k gives positive probability to every action and μ^k is derived from σ^k by Bayes' rule. A ***sequential equilibrium*** is a consistent, sequentially rational assessment.

Theorem 10.4. [classical; [KW82a, Sel75]] *Every finite extensive game with perfect recall has a sequential equilibrium. Every trembling-hand perfect equilibrium is sequential, and for generic payoffs the two concepts coincide. Every sequential equilibrium is subgame perfect.*

10.2 The geometry of conditioning

Consistency is a closure condition. On the open set of fully mixed strategies, beliefs are a rational function of strategies; the consistent assessments are the closure of the graph of that function. Taking the closure of the graph of a rational map is a classical operation in geometry—it is a *blow-up*—and looking at it this way explains what off-path beliefs are.

Consider the simplest case: a finite set Ω and the map sending a fully supported $p \in \Delta\Omega$ to its family of conditionals $\big(p(\cdot \mid S)\big)_S$, one for each nonempty $S \subseteq \Omega$.

Definition 10.5. [classical; [Mye86, Ren55]] A ***conditional probability system*** on Ω is a family $(p_S)_{\varnothing \neq S \subseteq \Omega}$ with $p_S \in \Delta(S)$ and $p_T(R) = p_S(R)\, p_T(S)$ for $R \subseteq S \subseteq T$.

Proposition 10.6 (Conditional systems are a blow-up)**.** [reformulation; cf. [McL89a, Mye86, BBD91]] *The space of conditional probability systems on a finite set Ω is the closure of the image of the open simplex in $\prod_S \Delta(S)$. It is a compactification of the open simplex in which each face of the closed simplex has been replaced by the product of that face with the space of conditional systems on the complementary set of states. Its strata are indexed by the ordered partitions $(\Omega_1, \dots, \Omega_k)$ of Ω: a point of that stratum is a fully supported distribution on each block, the states of Ω_{j+1} being infinitely less likely than those of Ω_j. Topologically it is a closed ball [McL89a].*

For three states the picture is Figure 10.1. A sequence of priors converging to the vertex δ_1 has, along the way, a well-defined ratio $p_2 : p_3$, and different

sequences converge to different limiting ratios. The simplex forgets the ratio; the blow-up remembers it, replacing the vertex by a segment. Each corner becomes an edge and the triangle becomes a hexagon. In general the face poset is that of the ordered partitions, which is the face poset of the permutohedron.

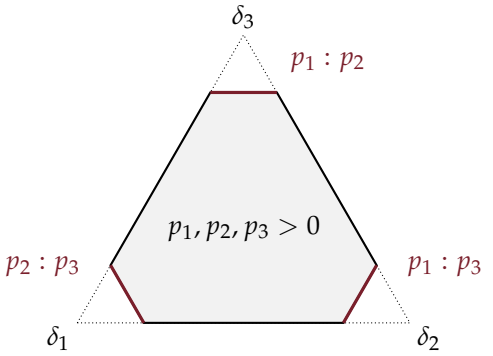


Figure 10.1: Conditional probability systems on three states. Each vertex of the simplex is blown up to a segment recording the limiting ratio of the two vanishing probabilities.

A consistent system of beliefs is a point of such a blow-up, constrained by the product structure that independent trembles of different players impose [McL89b, KR97]. The equivalent descriptions in the literature are different coordinate systems on the same space: *lexicographic probability systems* [BBD91] list the distributions on the blocks in order; *non-Archimedean probabilities* [Ham94, Hal10] are single distributions with values in an ordered field containing infinitesimals, the blocks being the orders of magnitude. The last makes the geometry algebraic: a consistent assessment is the shadow of a fully mixed strategy profile over a larger ordered field, and sequential rationality is ordinary rationality there.

Remark 10.7 (What kind of infinitesimal). The infinitesimals needed here are invertible elements of an ordered field, so that ratios of small probabilities make sense. They are not the nilpotent infinitesimals of synthetic differential geometry, which have no ratios. The appropriate ambient setting is real-closed fields and semialgebraic geometry rather than a smooth topos; conversely, the fact that equilibrium conditions are semialgebraic is what allows the transfer of results between $\mathbb{R}$ and its non-Archimedean extensions.

10.3 Reputation

Return to the chain-store paradox of Chapter 9. Kreps and Wilson [KW82b] and Milgrom and Roberts [MR82a] resolved it by a small change to the type

space: with probability p the incumbent is a *tough* type who fights every entry. Take the entry game with payoffs $(0,4)$ for no entry, $(1,2)$ for accommodated entry and $(-1,1)$ for fought entry, where the second coordinate is the payoff of the ordinary incumbent. An entrant enters iff she thinks the probability of a fight is at most $\frac{1}{2}$.

Example 10.8 (Two towns)**.** In the last town the ordinary incumbent accommodates, so the last entrant enters iff her posterior that the incumbent is tough is at most $\frac{1}{2}$. In the first town, suppose entry occurs. If $p \geq \frac{1}{2}$ the ordinary incumbent fights: this leaves the posterior at p, deters the second entrant, and yields $1+4 > 2+2$. If $p < \frac{1}{2}$ she cannot fight for sure (the posterior would stay below $\frac{1}{2}$ and the fight would be wasted) nor accommodate for sure (a fight would then prove toughness and be well worth it). She fights with probability $\varphi = \frac{p}{1-p}$, which makes the posterior after a fight exactly $\frac{1}{2}$, and the second entrant then enters with probability $\frac{1}{2}$, which makes the incumbent indifferent. The first entrant faces a fight with total probability $p + (1-p)\varphi = 2p$ and so stays out iff $p > \frac{1}{4}$.

One more town halves the threshold again. With n towns remaining entry is deterred whenever $p > 2^{-n}$: a one-in-a-thousand chance of toughness protects the incumbent in all but the last ten towns, however many there are. The chain-store intuition is restored by an arbitrarily small departure from complete information.

Theorem 10.9. [classical; [KMRW82, FL89]]

(i) In the finitely repeated Prisoner's Dilemma in which one player is, with probability $\epsilon > 0$, a type that plays tit-for-tat, the number of periods in which either player defects in a sequential equilibrium is bounded by a constant that depends on ϵ and not on the number of repetitions.

(ii) A patient long-run player facing a sequence of short-run opponents, who assign positive probability to her being committed to her Stackelberg action, obtains in every Nash equilibrium a payoff close to her Stackelberg payoff.

Two structural comments. First, the natural state variable of a reputation equilibrium is the posterior. The equilibrium is Markov in the sense of Proposition 9.12 once the state space is enlarged from histories to beliefs: with incomplete information the continuation game depends on what is believed, so beliefs are payoff-relevant and the bisimulation quotient of the game is a space of posteriors. Second, part (ii) says that reputation *selects*: out of the folk theorem's continuum it picks the outcome that commitment would produce. A possibility of being committed is a substitute for commitment. Both comments are instances of a single observation, to which we now turn: the equilibrium

correspondence is discontinuous in the type space, and small perturbations of beliefs can have effects that small perturbations of payoffs cannot.

10.4 Almost common knowledge

Example 10.10 (Electronic mail game)**.** [classical; [Rub89]] Two players would like to coordinate on a risky action when the state is b and on a safe action when the state is a; the risky action is costly unless both take it in state b. Player 1 learns the state. If it is b, her computer sends a message to player 2's, which sends a confirmation, which is confirmed in turn, and so on; each message is lost with small probability ϵ, and the exchange stops at the first loss. Each player observes only the number of messages her own machine sent. Then in the unique rationalizable outcome both players take the safe action in every state, however many messages have been exchanged.

The information structure is the simplest possible illustration of Chapter 4. The states are the pairs (number of messages sent by 1, number sent by 2), which are $(0,0),(1,0),(1,1),(2,1),\dots$; player 1 cannot distinguish $(n,n-1)$ from (n,n), and player 2 cannot distinguish (n,n) from $(n+1,n)$. The knowledge graph B is a single infinite path. It is connected, so by Proposition 4.9 the meet is trivial and *no* nontrivial event is ever common knowledge, although after $2n$ messages the players have n levels of mutual knowledge of the state. The proof of the example is an induction along the path: safe play at one vertex forces safe play at the next. Truncations K^n of the common-knowledge hierarchy do not approximate C as far as behaviour is concerned.

What does? Monderer and Samet's answer [MS89] is to weaken the modality rather than truncate the iteration. Let $\mathsf{B}_i^p E$ be the event that i assigns probability at least p to E, and define ***common p-belief*** from the operators B_i^p as common knowledge was defined from K_i. Then equilibrium behaviour under common p-belief of the game converges to behaviour under common knowledge as $p \to 1$. The family $(\mathsf{B}^p)_{p\in[0,1]}$ is a *graded* modality, and the theorem says that strategic behaviour is continuous in the grade although not in the depth.

The same discontinuity appears on the universal type space. T^* was built in Theorem 5.4 as an inverse limit and carries the product topology, in which two types are close when their beliefs agree to high finite order.

Theorem 10.11. [classical; [WY07, DFM06]] *For every type t in the universal type space and every action a that is rationalizable for t, there are types arbitrarily close to t in the product topology for which a is the* unique *rationalizable action (under a richness assumption on payoffs). Consequently no refinement of rationalizability is robust*

to perturbations of higher-order beliefs, and the coarsest topology on types in which rationalizable behaviour is continuous is strictly finer than the product topology.

The universal type space exists because of a continuity property of the functor $\mathcal{G}$; strategic behaviour is discontinuous in exactly the topology that this property provides. The semantics of types and the semantics of play live on the same set and disagree about its topology. One may regard this as bad news for refinements, or as an instrument. *Global games* [CvD93, MS98] use it as an instrument: adding small private noise to a coordination game with multiple equilibria produces a game with a unique rationalizable outcome, obtained by an infection argument like that of the mail game, and the outcome selected is the risk-dominant one. Note that this is consistent with Proposition 1.6: the argument uses only best-response structure, and risk dominance is the selection criterion that lives on $\mathsf{NFG}_{\mathrm{inc}}$.

Notes

Sequential equilibrium is due to Kreps and Wilson [KW82a], perfect equilibrium to Selten [Sel75], and perfect Bayesian equilibrium in its usual form to Fudenberg and Tirole [FT91b]. Conditional probability systems go back to Rényi [Ren55] and were brought into game theory by Myerson [Mye86]; McLennan [McL89a, McL89b] studied the topology of the spaces of conditional systems and of consistent assessments, and Kohlberg and Reny [KR97] characterized consistency by an independence condition on relative probabilities. Lexicographic and non-Archimedean representations are in [BBD91, Ham94, Hal10]. The reputation papers are [KW82b, MR82a, KMRW82], and the general bounds are due to Fudenberg and Levine [FL89]; Diamond's model of reputation acquisition in debt markets [Dia89] is an application in which a borrower's incentive to choose safe projects grows with her record. The electronic mail game is [Rub89]; common p-belief is [MS89]; the structure theorem for rationalizability is [WY07] and the strategic topology [DFM06]. Global games were introduced by Carlsson and van Damme [CvD93] and applied to currency crises by Morris and Shin [MS98].

Exercises

Exercise 10.1. In Example 10.2, find all Nash equilibria of the strategic form, including mixed ones, and determine which are sequential.

Exercise 10.2. List the thirteen strata of the space of conditional probability systems on a three-element set, and match each with a face of the hexagon of Figure 10.1.

Exercise 10.3. Represent the conditional probability system on $\{1,2,3\}$ in which state 1 is infinitely more likely than the others and $p_2 : p_3 = 2 : 1$ (a) as a lexicographic probability system and (b) as a distribution with values in $\mathbb{R}(\varepsilon)$.

Exercise 10.4. Extend Example 10.8 to three towns and verify the threshold $p > \frac{1}{8}$ for deterrence in the first.

Exercise 10.5. Draw the knowledge graph of the electronic mail game, and compute K_1E, K_2K_1E and $K_1K_2K_1E$ for the event E that the state is b.

Exercise 10.6. Show that B_i^p is monotone, satisfies $\mathsf{B}_i^p E \cap \mathsf{B}_i^p F \subseteq \mathsf{B}_i^{2p-1}(E \cap F)$, and reduces to an operator containing K_i when $p = 1$. In what sense is the family a graded comonad?

Part II

Asymmetric Information

Chapter 11

Moral Hazard

Part II is about transactions between parties one of whom knows, or will know, or will do, something the other cannot observe. The uninformed party—the *principal*—designs the terms; the informed party—the *agent*—responds to them. Everything turns on the fact that a contract can depend only on what is observed. In the language of Part I, a contract is a function that must factor through an observation map, and contract design is the choice of a context for someone else's decision.

11.1 Five models and one diagram

The classical models of asymmetric information differ in the order of four events: Nature's choice of a state θ observed only by the agent; the signing of a contract; the agent's action or message; and the realization of an outcome. A convenient classification, which the source text [Raso6, ch. 7] makes systematic, is the following.

Model	Order of events	What the principal cannot see
hidden action	contract $\to$ action $\to$ noise $\to$ outcome	the action
hidden knowledge	contract $\to$ θ $\to$ report or action	a state learned after contracting
adverse selection	θ $\to$ contract	the type, known before contracting
signalling	θ $\to$ signal $\to$ contract	the type; the informed party moves first
screening	θ $\to$ menu $\to$ choice	the type; the uninformed party moves first

The first two are called *moral hazard*, because the asymmetry arises after the parties are bound; the last three, in which the agent is informed at the outset, are the province of Chapters 13 to 15. As open games (Chapter 2) the five are five ways of composing the same boxes: a move of Nature, a decision of the

principal, a decision of the agent, and an observation channel that limits which wires may run into which box. What is constant is the shape of the principal's problem. She does not choose the agent's action. She chooses the agent's *context*, and the action is whatever the agent's selection function returns in that context.

11.2 The principal–agent problem

An agent chooses an effort $e \in E$. Effort determines a distribution over observable outcomes $x \in X$: the *technology* is a Markov kernel $f\colon E \to X$, written $f(x \mid e)$. A ***contract*** is a function $w\colon X \to \mathbb{R}$ specifying the agent's pay. The agent has utility $u(w) - c(e)$, with u increasing and concave and c the cost of effort, and an outside option worth $\overline{U}$. The principal is risk neutral and receives $x - w(x)$.

In the vocabulary of Definition 2.1, the contract composed with u is a costate on the outcome boundary: a function assigning value to each observable result. Pulling it back along the technology gives the agent's decision problem,

$$k_w(e) = \mathbb{E}_{x \sim f(e)}\big[u(w(x))\big] - c(e),$$

and the agent selects $e \in \arg\max k_w$. The principal, moving first, chooses w to maximize her own payoff at the selected effort. This is the product of selection functions of Section 2.4, with one addition, the agent's right to refuse:

$$\begin{aligned} \max_{w,\,e} \quad & \mathbb{E}_{f(e)}[x - w(x)] \\ \text{subject to} \quad & \mathbb{E}_{f(e)}[u(w(x))] - c(e) \;\geq\; \overline{U} && \text{(participation)} \\ & e \in \arg\max_{e'} \; \mathbb{E}_{f(e')}[u(w(x))] - c(e') && \text{(incentive compatibility).} \end{aligned}$$

The incentive constraint is not an extra assumption; it is the statement that the second player in a sequential composite best-responds.

The first best

If effort were observable the contract could depend on it, incentive compatibility would be vacuous, and the principal would solve an insurance problem. A risk-neutral party should bear all the risk: the optimal contract pays a constant wage w^* with $u(w^*) = \overline{U} + c(e)$ when the agreed effort is supplied (*Borch's rule*). This is the ***first best***. With hidden effort a constant wage induces the least costly effort. The agent must be exposed to the outcome in order to care about it, and exposure to risk is costly to a risk-averse agent. That conflict between insurance and incentives is the whole subject.

There is one case with no conflict. If the agent is risk neutral, the contract $w(x) = x - K$ makes the agent the residual claimant of the entire outcome. She then chooses the efficient effort, bears all the risk at no cost, and pays the principal the fixed fee K. Moral hazard is a problem only in combination with risk aversion, or with a limit on how much the agent can be made to pay.

The second best

Example 11.1 (Two efforts, two outcomes)**.** Effort is high or low. The project succeeds with probability 0.8 under high effort and 0.4 under low. The agent has $u(w) = \sqrt{w}$, $c(H) = 1$, $c(L) = 0$, $\overline{U} = 1$; a success is worth 20 to the principal and a failure 0.

Observable effort. To obtain H the principal pays the constant wage with $\sqrt{w} = 2$: $w = 4$, for a profit of $16 - 4 = 12$. To obtain L she pays $w = 1$ for a profit of $8 - 1 = 7$.

Hidden effort. Let u_s, u_f be the agent's utility from pay after success and failure. Incentive compatibility for H is $0.8u_s + 0.2u_f - 1 \geq 0.4u_s + 0.6u_f$, that is $u_s - u_f \geq 2.5$; participation is $0.8u_s + 0.2u_f - 1 \geq 1$. Both bind at the optimum: $u_f = 0$, $u_s = 2.5$, so $w_f = 0$ and $w_s = 6.25$. The expected wage is 5, and the profit is 11. The principal still prefers high effort, but pays an *agency cost* of 1: the expected wage exceeds the first-best wage by exactly the risk premium the agent requires for carrying the incentive scheme.

With a continuum of efforts and outcomes, and when the agent's problem can be replaced by its first-order condition, the optimal contract satisfies

$$\frac{1}{u'(w(x))} = \lambda + \mu \, \frac{f_e(x \mid e)}{f(x \mid e)} \tag{11.1}$$

for multipliers $\lambda, \mu > 0$ [Hol79]. Pay depends on the outcome only through the likelihood ratio f_e / f: the contract rewards outcomes not because they are good for the principal but because they are *evidence* of effort. If the likelihood ratio is increasing in x (the monotone likelihood ratio property), pay is increasing in x. The first-order approach is valid under further convexity conditions [Rog85], and can fail without them.

11.3 Informativeness

Formula (11.1) says that the contract should use exactly the information about effort contained in the outcome. This can be made into a theorem that involves no first-order conditions, and it is here that the Markov-category language of Chapter 5 earns its place.

Definition 11.2. [classical; [Bla51]; categorical form [Fri20, FGPR23]] Let $f\colon E \to X$ and $g\colon E \to Y$ be Markov kernels out of the same object (two *experiments* on E). g is a ***garbling*** of f, written $g \preceq f$, if $g = h \circ f$ for some kernel $h\colon X \to Y$. A statistic $s\colon X \to V$ is ***sufficient*** for f if f is in turn a garbling of $s \circ f$.

So the informativeness order on experiments is the factorization preorder on the coslice under E in the category of kernels, and a sufficient statistic is a deterministic map that loses nothing in this order. Blackwell's theorem says that $g \preceq f$ iff every Bayesian decision maker weakly prefers observing f to observing g. In an agency problem the decision maker and the source of the uncertainty are different people, and the corresponding statement needs its own proof.

Theorem 11.3 (Informativeness principle)**.** [classical; [Hol79, GH83]]

(i) If the monitoring technology $g\colon E \to Y$ is a garbling of $f\colon E \to X$, then for every effort e the least expected wage that implements e under f is no greater than under g.

(ii) In particular, if s is a sufficient statistic for the pair (x, y) with respect to e, an optimal contract can be taken to depend on (x, y) only through s; and a signal y has value in addition to x only if x is not sufficient for (x, y).

Proof. (i) Let $w'\colon Y \to \mathbb{R}$ implement e under $g = h \circ f$. Under f, consider first the *random* contract that, on observing x, draws $y \sim h(x)$ and pays $w'(y)$. For every effort it induces the same distribution of pay as w' does under g, so the agent's incentives and participation are unchanged. Now replace the random payment at each x by the sure payment $\tilde{w}(x)$ with $u(\tilde{w}(x)) = \mathbb{E}_{y\sim h(x)} u(w'(y))$. The agent's expected utility from every effort is the same, so $\tilde{w}$ implements e; and by Jensen's inequality $\tilde{w}(x) \le \mathbb{E}_{h(x)} w'(y)$, so it costs no more. (ii) follows by applying (i) in both directions. □

The proof is two moves: pull the contract back along the garbling, then derandomize. The first uses only composition of kernels. The second is the one place where risk aversion enters, and it shows why the agency ordering can be strictly finer than Blackwell's [Kim95]: what matters is not all the information in the signal but the information about the agent's *deviations*, the dispersion of the likelihood ratio.

The principle has a normative reading that is often forgotten. Any variation in pay that is not evidence about the agent's action is pure cost. Executive pay that moves with the stock market as a whole, when the market is observable and not influenced by the executive, violates it; so does punishing an agent for bad luck that can be verified to be luck.

11.4 Linear contracts

The second-best contract of (11.1) is typically a complicated function, sensitive to the exact shape of the likelihood ratio. Real contracts are simple, often linear. Holmström and Milgrom [HM87] gave a model in which linearity is optimal, and its closed form has become the workhorse of applied contract theory.

Example 11.4 (The linear–exponential–normal model)**.** Output is $x = e + \varepsilon$ with $\varepsilon \sim N(0, \sigma^2)$. The agent has constant absolute risk aversion r and effort cost $\frac{1}{2}ce^2$; the contract is $w = \alpha + \beta x$. The agent's certainty equivalent is

$$\alpha + \beta e - \tfrac{1}{2}ce^2 - \tfrac{1}{2}r\beta^2\sigma^2,$$

so she chooses $e = \beta/c$. The fixed part α only transfers surplus, so the optimal slope maximizes total certainty-equivalent surplus $\frac{\beta}{c} - \frac{\beta^2}{2c} - \frac{1}{2}r\beta^2\sigma^2$, giving

$$\beta^* = \frac{1}{1 + rc\sigma^2}.$$

The slope is the agent's share of the marginal product. It is 1—sell the firm to the agent—when she is risk neutral or output is noiseless, and falls towards 0 as risk aversion, noise, or the convexity of the cost of effort increase.

Why should attention be restricted to linear contracts? In [HM87] the agent controls the drift of a Brownian motion continuously in time and observes its path; the stationarity of the problem makes the optimal contract linear in final output. A different and more recent justification is robustness: if the principal knows only some of the actions available to the agent and evaluates contracts by their worst-case performance over the rest, linear contracts are optimal, because they are exactly the contracts that tie the two parties' payoffs together whatever the agent does [Car15]. Both arguments say that linearity is what remains when the fine detail of the technology f cannot be relied on.

Notes

The hidden-action model is due to Mirrlees and Holmström [Hol79]; the rigorous treatment without the first-order approach is Grossman and Hart's [GH83], and conditions validating the first-order approach are given by Rogerson [Rog85]. The informativeness principle is from [Hol79]; its relation to Blackwell's order [Bla51] is examined in [GH83, Kim95]. Sufficiency, the factorization theorem and Blackwell's theorem have been redeveloped in Markov categories by Fritz and coauthors [Fri20, FGPR23]. Linear contracts are from [HM87], and the robustness argument is Carroll's [Car15]. The textbook reference for all of Part II is [BD05]; see also [MWG95]. The source text develops the

same ideas through a sequence of variations on a single production example [Raso6, ch. 7], varying who makes the offer and what is observed; the lesson of that sequence is that bargaining power determines the division of surplus and information determines its size.

Exercises

Exercise 11.1. In Example 11.1, find the range of values of a success for which the principal prefers low effort under hidden action although she would prefer high effort if it were observable.

Exercise 11.2. In Example 11.1 add a limited-liability constraint $w \geq 0$ and let the agent be risk neutral with $\overline{U} = 0$. Find the cheapest contract implementing H and the agent's rent.

Exercise 11.3. Suppose that in addition to the outcome the principal observes a signal $y \in \{0, 1\}$ that equals 1 with probability 0.9 under high effort and 0.5 under low effort, independently of the outcome. Show that the outcome is not sufficient for (outcome, y), and compute the cheapest contract implementing H in Example 11.1 that uses both.

Exercise 11.4. Show that the relation $\preceq$ on experiments is a preorder, that deterministic maps s with $s \circ f \preceq f$ always exist, and that s is sufficient iff equality holds up to equivalence.

Exercise 11.5. Derive β^* in Example 11.4, and compute the agency cost—the loss of total surplus relative to observable effort—as a function of r, c and σ^2.

Exercise 11.6. Draw the principal–agent problem with hidden action as a composite of open games in the style of Figure 2.2, with a box for Nature's noise. Which wire is missing compared with the observable-effort problem?

Chapter 12

Moral Hazard, Continued

The one-period, one-agent, one-task problem of Chapter 11 isolates the trade-off between insurance and incentives. This chapter relaxes each restriction in turn—many periods, many agents, many tasks, and the possibility of rewriting the contract—and finds in each case that a construction from Part I is doing the work.

12.1 Efficiency wages

Suppose performance cannot be contracted on at all, but a worker who shirks is detected with probability q each period and can then be dismissed. Dismissal is a punishment only if the job is worth more than the alternative. Let the wage be w, the cost of effort e, the discount factor δ, and let unemployment be worth V_U. A worker who never shirks has value $V_N = \frac{w-e}{1-\delta}$. By the one-shot deviation principle (Theorem 9.4) it suffices to deter a single period of shirking:

$$V_N \;\geq\; w + \delta\big[(1-q)V_N + qV_U\big] \quad\Longleftrightarrow\quad \delta q\,(V_N - V_U) \;\geq\; e.$$

The left side is the discounted expected loss from being caught: detection probability times patience times the *rent* from employment. With $V_U = \frac{\overline{u}}{1-\delta}$ the condition becomes

$$w \;\geq\; \overline{u} + e + \frac{(1-\delta)\,e}{\delta q}.$$

The wage must exceed the worker's opportunity cost $\overline{u} + e$ by a premium that is larger the harder shirking is to detect. This is an ***efficiency wage*** [SS84]. Two features are worth noting. The incentive is *coinductive*: it is provided not by any payment contingent on performance but by the continuation value of the relationship, which is exactly the promised-value state of Theorem 9.8. And it has a striking aggregate implication. If all firms pay the premium, a dismissed worker must not be able to step into an equally good job at once; equilibrium requires $V_U < V_N$, hence a pool of involuntarily unemployed workers who would strictly prefer to work at the going wage. Unemployment is the punishment state of the automaton.

12.2 Tournaments

With several agents the principal can compare them. Let $y_i = e_i + \eta + \varepsilon_i$, where η is a shock common to all agents and the ε_i are idiosyncratic. A ***tournament*** pays a prize W_1 to the agent with the higher output and $W_2 < W_1$ to the other [LR81]. Agent i wins when $e_i + \varepsilon_i > e_j + \varepsilon_j$; the common shock cancels. If G is the distribution function of $\varepsilon_j - \varepsilon_i$ and g its density, agent i's problem is to maximize $W_2 + (W_1 - W_2)G(e_i - e_j) - c(e_i)$, and in a symmetric equilibrium

$$(W_1 - W_2)\, g(0) = c'(e).$$

Effort is governed by the *spread* of prizes and by how responsive the probability of winning is to effort; the level of prizes serves only to satisfy participation. For risk-neutral agents the first best $c'(e) = 1$ is attained with $W_1 - W_2 = 1/g(0)$: the more luck matters, the larger the prize for winning must be.

Why compare agents rather than pay each on her own output? The informativeness principle answers (Theorem 11.3). When η is present, y_j is informative about the noise in y_i, so y_i alone is not a sufficient statistic for e_i; the difference $y_i - y_j$ is free of η. Relative performance evaluation is valuable exactly to the extent that agents face common shocks [Hol82]. A rank-order tournament is a coarse way to use this information—it throws away the margin of victory—but it has the merit that the principal's total payment is fixed, so that she has no incentive to misreport who won. Its costs are equally structural: agents gain from each other's failure, and so may sabotage or collude.

12.3 Teams

Now let the agents produce a single joint output $x = x(e_1, \dots, e_n)$, increasing and concave, with private costs $c_i(e_i)$. Only x is observed. The agents form a partnership: output is divided among them according to sharing rules $s_i(x)$ with $\sum_i s_i(x) = x$ for all x.

Theorem 12.1 (Holmström). [classical; [Hol82]] *No budget-balanced sharing rule makes the efficient effort profile a Nash equilibrium.*

Proof for differentiable rules. Efficiency requires $\partial x/\partial e_i = c_i'(e_i)$ for each i. Agent i's best response satisfies $s_i'(x)\,\partial x/\partial e_i = c_i'(e_i)$. For the efficient profile to be an equilibrium every agent must therefore face $s_i' = 1$ at the efficient output; but budget balance gives $\sum_i s_i' = 1$. □

Each agent must be the residual claimant of the whole marginal product, and there is only one marginal product to go round. The cure is to break the budget: introduce a principal who is not a member of the team, and pay

$s_i(x) = b_i$ if x reaches the efficient level and nothing otherwise, the principal keeping any shortfall. Then every agent is pivotal, and efficient effort is an equilibrium. On this account the economic function of an owner is neither to monitor nor to supply capital, but to be a sink.

The theorem has a clean reading in the calculus of Chapter 2. A partnership is a *closed* game: all utility generated inside is distributed inside. Holmström's result says that the efficient profile is not an equilibrium of any closed game of this shape, and that it becomes one when a boundary wire is added along which value can leave. The Groves mechanism of Chapter 14 is the same phenomenon for hidden information: making each agent internalize the total surplus is possible, and incompatible with budget balance.

12.4 Many tasks

Let the agent supply a vector of efforts $e \in \mathbb{R}^k$ at cost $\frac{1}{2} e^\top C e$, generating benefits $b^\top e$ for the principal and signals $x = e + \varepsilon$, $\varepsilon \sim N(0, \Sigma)$. With a linear contract $w = \alpha + \beta^\top x$ and constant absolute risk aversion r, the agent chooses $e = C^{-1}\beta$, and maximizing total surplus gives

Proposition 12.2 (Multitask incentives). [classical; [HM91]]

$$\beta^* = (I + rC\Sigma)^{-1} b.$$

If there are two tasks and the second cannot be measured at all ($\sigma_2^2 \to \infty$, so $\beta_2 = 0$), then

$$\beta_1^* = \frac{b_1 - b_2\, C_{12}/C_{22}}{1 + r\sigma_1^2 \left(C_{11} - C_{12}^2/C_{22}\right)}.$$

The first formula generalizes Example 11.4. The second is the interesting one. When the tasks are substitutes in the agent's cost ($C_{12} > 0$), rewarding the measured task draws effort away from the unmeasured one, and the optimal incentive on the measured task is *reduced* by the value of what is crowded out. If the unmeasured task matters enough, β_1^* is zero or negative: the best contract is a flat salary, not because effort cannot be measured but because *some* of it cannot.

This is the formal theory of a familiar complaint. When a measure of performance becomes a target, effort flows to what is measured, and the measure ceases to track what it was introduced to track. The complaint goes under the names of Goodhart and Campbell; the proposition says precisely when it has force (tasks are substitutes; the unmeasured task is valuable) and what to do (weaken the incentive; or redesign jobs so that measurable and unmeasurable tasks are assigned to different agents [HM91]). It also explains

why incentives inside organizations are so much weaker than a single-task model would recommend, and why the strongest incentives are found where the output is one-dimensional.

12.5 Renegotiation

A contract is a commitment by both sides. Suppose the parties can tear it up and write another whenever both wish. After the agent has chosen her effort but before the outcome is known, the incentive scheme has done its work: effort is sunk, and all that remains is risk, which the risk-neutral principal should bear. Both parties gain by replacing the contract with a fixed payment. But if the agent anticipates this, she has no reason to work.

Theorem 12.3. [classical; [FT90]] *In the hidden-action model with renegotiation after effort is chosen and before the outcome is realized, no contract induces high effort with probability one. In equilibrium the agent randomizes over effort, and the renegotiated contract is a menu that screens the agent according to the effort she privately knows she has taken.*

Two lessons. First, renegotiation hurts, for the reason given in Section 8.5: the ability to reopen an agreement destroys the ability to commit. Second, it *changes the type of the problem*. At the renegotiation stage the agent's past effort is private information about the distribution of the outcome; the principal faces an agent with a hidden *type*. A model of hidden action has become a model of adverse selection, and must be solved with the tools of Chapter 14. The classification of Chapter 11 is not stable under adding a stage.

Structurally, renegotiation is an operator ρ on the set of contracts: $\rho(w)$ is what w will be replaced by. A contract is *renegotiation-proof* if $\rho(w) = w$. Since renegotiation leads to a contract that is not itself renegotiated, ρ is idempotent, and its fixed points are its image. The *renegotiation-proofness principle* says that nothing is lost by restricting attention to the image: whatever can be achieved with renegotiation on the equilibrium path can be achieved by a contract that is never renegotiated. It is the first of several statements in Part II—the revelation principle is the most famous—to the effect that a class of mechanisms may be replaced by the fixed points of an idempotent operator on it.

12.6 Extremes: forcing contracts and their failure to exist

If the support of the outcome distribution moves with effort, some outcomes prove that the agent shirked, and a sufficiently severe penalty for those outcomes implements the first best at no cost in risk: the penalty is never incurred. Mirrlees observed that something similar is true even when no outcome is

conclusive. With normally distributed noise, very low outcomes are overwhelmingly more likely under low effort than high; a contract that pays the first-best wage except after an extremely low outcome, when it imposes an extremely large penalty, approximates the first best as closely as desired, and no contract attains it. The principal's problem has a supremum and no maximum. This is the open-set problem of Section 8.4 in another guise, and it carries the same warning: the existence of an optimal contract is a compactness property—here, of bounds on punishment—and not a formality. Limited liability, the impossibility of penalties beyond the agent's wealth, is what closes the set in practice, and it changes the character of the solution from risk-sharing to rent-sharing.

Notes

Efficiency wages as a discipline device are due to Shapiro and Stiglitz [SS84]. Tournaments are from Lazear and Rosen [LR81]; relative performance evaluation and the team problem from Holmström [Hol82]. The multitask model is [HM91]. Renegotiation in agency is analysed by Fudenberg and Tirole [FT90]. Mirrlees's approximation result is discussed in [BD05]. The source text [Ras06, ch. 8] also treats insurance under moral hazard by state-space diagrams: full insurance removes the incentive to take care, and the second-best policy leaves the insured with a deductible or a share of the loss—the insurance-market form of $\beta^* < 1$ in reverse.

Exercises

Exercise 12.1. In the efficiency-wage model add an exogenous probability s of separation each period. Derive the no-shirking wage and show that it increases with s.

Exercise 12.2. In the two-agent tournament let ε_i be independent and uniform on $[-a, a]$. Compute $g(0)$, the first-best prize spread for risk-neutral agents with $c(e) = \frac{1}{2}e^2$, and the prizes that make the participation constraint bind when the outside option is zero.

Exercise 12.3. Let $x = e_1 + e_2$ and $c_i(e) = \frac{1}{2}e^2$. Compute the efficient efforts, the equilibrium efforts under equal sharing, and the loss of surplus. Verify that the budget-breaking scheme of the text makes efficient effort an equilibrium, and find another equilibrium of the same scheme.

Exercise 12.4. Derive both formulas of Proposition 12.2. For $b = (1, 2)$, $C = \left(\begin{smallmatrix} 1 & \kappa \\ \kappa & 1 \end{smallmatrix}\right)$, $r\sigma_1^2 = 1$, and task 2 unmeasurable, find the values of κ for which the optimal contract is a flat wage or worse.

Exercise 12.5. Formalize the claim that renegotiation is an idempotent operator in a model with two contracts and one renegotiation stage, and state the renegotiation-proofness principle as a property of the image of ρ.

Chapter 13

Adverse Selection

In moral hazard the agent's information is created after the contract. In adverse selection it is there from the start: the agent knows her type and the principal, or the market, does not. The terms of trade can depend only on what is observed, so different types receive the same terms; and the set of types willing to accept given terms depends on the terms. This chapter studies the resulting fixed-point problem in three markets—goods, insurance and securities.

13.1 Lemons

A seller knows the quality θ of her car and values it at θ. Buyers cannot observe quality; a buyer values a car of quality θ at $v\theta$, with $v > 1$, so every trade is efficient. Buyers are competitive and risk neutral. At a price p the sellers who offer their cars are those with $\theta \leq p$, and a buyer is willing to pay the expected value of what is offered. An equilibrium price is a fixed point of the operator

$$\Phi(p) = v \cdot \mathbb{E}[\theta \mid \theta \leq p].$$

Example 13.1 (Collapse)**.** [classical; [Ake70]] Let θ be uniform on $[0,1]$. For $p \leq 1$ the average quality offered is $p/2$ and $\Phi(p) = vp/2$. If $v < 2$ the only fixed point is $p = 0$: no price at which cars are offered is justified by the cars offered at it, and the market disappears although every trade would create value. If $v > 2$ there is also an equilibrium with $p = v/2 \geq 1$ at which every car is sold.

Example 13.2 (Partial collapse)**.** Let θ be uniform on $[1,2]$ and $v = 1.2$. For $p \in [1,2]$, $\Phi(p) = 0.6(1+p)$, with unique fixed point $p = 1.5$. Cars of quality up to 1.5 trade at price 1.5; their average quality is 1.25, worth exactly 1.5 to buyers. The better half of the market does not trade.

Proposition 13.3. [reformulation; cf. [Ake70, Wil80]] *Φ is monotone on an interval of prices. Hence, by Theorem 7.1, the equilibrium prices form a nonempty complete lattice; the greatest equilibrium is $v\Phi$, reached by iterating Φ downward from any price above it; and since sellers and buyers both weakly prefer higher equilibrium prices, the greatest equilibrium Pareto-dominates the others.*

The familiar story of unraveling—the best cars withdraw, the price falls, the next-best withdraw—is the iteration $p, \Phi(p), \Phi^2(p), \dots$ from above, converging to the greatest fixed point. Akerlof's market is an instance of the left-hand column of the table in Section 7.7: a monotone operator, a lattice of equilibria, and a question about which end is selected. When several equilibria exist they are Pareto-ranked, and the market may be stuck at a low one; this is a coordination failure layered on an information failure.

It is worth isolating what the information failure consists in. The price is a function on what buyers observe, and buyers observe nothing: the price must factor through the map from qualities to a point. Under symmetric information, with quality unobserved by both sides, the same factorization holds and nothing goes wrong: all cars trade at the price $v\,\mathbb{E}[\theta]$. The trouble is the conjunction of a pooled price and a privately informed decision to participate, which makes the pool *depend on* the price. Adverse selection is selection: the conditioning event $\{\theta \le p\}$ moves with the variable being determined. The remedies all work by letting the terms of trade depend on something correlated with type: a costly signal chosen by the informed side (Chapter 15), a menu from which types select (Chapter 14), a warranty, a reputation, a certifier.

13.2 Insurance

Selection is equally adverse when the informed party is the buyer. Individuals have wealth W and suffer a loss d with probability π_H or $\pi_L < \pi_H$ according to their privately known risk type; a fraction λ are high risks. Competitive insurers offer contracts specifying a premium and a coverage. A contract that breaks even on the population average is a bargain for high risks and overpriced for low risks.

Theorem 13.4 (Rothschild–Stiglitz). [classical; [RS76]] *In the competitive insurance market with two risk types, where an equilibrium is a set of contracts such that none makes losses and no further contract would be profitable if offered:*

(i) there is no pooling equilibrium;

(ii) in a separating equilibrium high risks buy full insurance at their fair premium, and low risks buy partial insurance at their fair premium, the coverage being the largest that high risks do not prefer to their own contract;

(iii) this candidate is an equilibrium if and only if the fraction of high risks is at least some $\lambda^* \in (0,1)$; otherwise no equilibrium exists.

The reasoning for (i) is *cream-skimming*. Indifference curves of the two types cross once (Theorem 7.8): low risks are less willing to pay for additional

coverage. Hence near any pooling contract there is another with slightly less coverage and a lower premium that attracts only low risks, and is profitable. For (iii): if high risks are rare, the population-average price is close to the low-risk price, and a pooling contract with nearly full coverage beats the separating pair for *both* types; it would upset the separating candidate, but by (i) cannot itself be an equilibrium.

Example 13.5. Let $u = \ln$, $W = 4$, $d = 3$, $\pi_H = \frac{1}{2}$, $\pi_L = \frac{1}{4}$. High risks insure fully at the fair premium 1.5 and have wealth 2.5 in both states. A low-risk contract covering a fraction α of the loss at the fair premium 0.75α leaves wealth $4 - 0.75\alpha$ without a loss and $1 + 2.25\alpha$ with one. High risks are indifferent between it and their own contract when $(4 - 0.75\alpha)(1 + 2.25\alpha) = 2.5^2$, that is at $\alpha^* \approx 0.29$. Low risks, who under symmetric information would be fully insured, obtain less than a third of full coverage: the informational externality is borne entirely by the *good* risks. The separating pair is an equilibrium iff $\lambda \geq \lambda^* \approx 0.28$ (computed in Appendix C).

Nonexistence is uncomfortable, and it reflects the fact that the equilibrium notion here is not a Nash equilibrium of a well-specified game but a free-entry condition with a deviation set that is not closed under the reactions it provokes. Three repairs are standard. One may specify the game and allow mixed strategies, which restores existence [DM86]. One may change the conjecture that a deviating insurer makes about the response of its rivals: if unprofitable contracts will be withdrawn, pooling can survive [Wil77]. Or one may abandon the strategic story and define a competitive equilibrium in which every contract, traded or not, has a price that reflects the cost of those who would buy it; such equilibria exist generally and select the separating outcome in this model [AG17]. The last is the modern form of Akerlof's fixed point, with the price now a function on the space of contracts rather than a number.

13.3 Securities

In financial markets the uninformed party is a market maker who quotes prices at which she will trade with whoever arrives.

Example 13.6 (The bid–ask spread). [classical; [GM85]] An asset is worth 0 or 1 with equal probability. A fraction μ of traders know the value and buy if it is 1, sell if it is 0; the rest buy or sell with probability $\frac{1}{2}$ for unrelated reasons. A competitive market maker sets the ask at the expected value conditional on a buy order and the bid at the expected value conditional on a sale:

$$\text{ask} = \mathbb{P}(v = 1 \mid \text{buy}) = \frac{1 + \mu}{2}, \qquad \text{bid} = \frac{1 - \mu}{2}.$$

The spread equals μ. It is not a cost of doing business but a transfer: the market maker loses to the informed and recovers the loss from the uninformed, and the spread is the price of the adverse selection in the order flow.

Example 13.7 (Kyle's lambda). [classical; [Kyl85]] An asset has value $v \sim N(p_0, \Sigma_0)$. An insider who knows v submits an order x; noise traders submit $u \sim N(0, \sigma_u^2)$; the market maker sees the total $y = x + u$ and sets the price $p = \mathbb{E}[v \mid y]$. Look for a linear equilibrium $x = \beta(v - p_0)$, $p = p_0 + \lambda y$. The insider maximizes $\mathbb{E}[(v - p)x] = (v - p_0 - \lambda x)x$, so $\beta = \frac{1}{2\lambda}$. The market maker's rule is a regression, $\lambda = \frac{\beta\Sigma_0}{\beta^2\Sigma_0 + \sigma_u^2}$. Solving,

$$\lambda = \frac{\sqrt{\Sigma_0}}{2\sigma_u}, \qquad \beta = \frac{\sigma_u}{\sqrt{\Sigma_0}}, \qquad \operatorname{Var}(v \mid p) = \frac{\Sigma_0}{2}.$$

The insider trades less aggressively the more her trades move the price; the price moves more the more informative the order flow; and in equilibrium exactly half of her information is impounded in the price, whatever the amount of noise.

Both examples have the shape described in Chapter 5: a forward kernel from the hidden state to an observable (the order), and a backward map from the observable to a belief about the state (the price), which is its Bayesian inverse. The informed trader chooses the forward kernel taking the backward map as given; the market maker's backward map must be the inverse of the forward kernel actually chosen. A rational-expectations equilibrium is a fixed point of this pair—a Bayesian lens that is consistent with the best response it induces. The same structure, with a manager's unobserved action in place of the insider's order, is the signal-jamming model of Chapter 15.

Notes

The market for lemons is Akerlof's [Ake70]; the lattice structure of its equilibria and their Pareto ranking are discussed in [Wil80, MWG95]. The insurance model is due to Rothschild and Stiglitz [RS76]; the reactive equilibrium concepts are from [Wil77], mixed-strategy existence from [DM86], and the price-theoretic equilibrium from Azevedo and Gottlieb [AG17]. The sequential trade model of the bid–ask spread is Glosten and Milgrom's [GM85]; the batch-auction model is Kyle's [Kyl85]. The source text [Ras06, ch. 9] develops lemons markets through several variants with different distributions of quality and tastes; Examples 13.1 and 13.2 are in that spirit with different numbers.

Exercises

Exercise 13.1. Let θ be uniform on $[0,1]$, let sellers value a car at θ and buyers at $\theta + k$ with $k > 0$. Find all equilibrium prices as a function of k and the fraction of cars traded.

Exercise 13.2. Construct a distribution of quality for which Φ has three fixed points, and verify that they are Pareto-ranked.

Exercise 13.3. In Example 13.5, verify α^*, compute both types' expected utilities at the separating pair and at full insurance at the pooled fair premium when $\lambda = 0.1$, and confirm that the pooling contract would attract both.

Exercise 13.4. In Example 13.6 let the asset be worth 1 with prior probability π. Compute the bid and the ask, and the market maker's posterior after two consecutive buy orders.

Exercise 13.5. In Example 13.7, compute the insider's expected profit and the expected loss of the noise traders, and show that they are equal. How do they depend on σ_u?

Chapter 14

Mechanism Design

So far the game has been given and the equilibrium sought. Mechanism design reverses the problem: an outcome is wanted, as a function of information the designer does not have, and the question is whether some game has it as an equilibrium. The answer is organized by one principle, which reduces the search over all games to a system of inequalities, and by the structure of that system, which turns out to be a piece of enriched category theory.

14.1 The revelation principle

An agent has a type $\theta \in \Theta$, known only to her, and utility $U(y,\theta)$ over outcomes $y \in Y$. A ***mechanism*** is a set M of messages with an outcome function $g\colon M \to Y$. A strategy is a map $\sigma\colon \Theta \to M$; it is a best response if $U(g\sigma(\theta),\theta) \geq U(g(m),\theta)$ for all θ and m. The mechanism then ***implements*** the social choice function $f = g \circ \sigma\colon \Theta \to Y$.

Fix f and consider the category $\mathsf{Impl}(f)$ whose objects are triples (M,g,σ) with σ a best response in (M,g) and $g\sigma = f$, and whose morphisms $(M,g,\sigma) \to (M',g',\sigma')$ are maps $h\colon M \to M'$ with $g'h = g$ and $h\sigma = \sigma'$.

Theorem 14.1 (Revelation principle). [reformulation; cf. [Gib73, Mye79]]

(i) (Best responses pull back.) Let σ be a best response in (M,g) and let $h\colon M' \to M$ be any map such that $\sigma = h\sigma'$ for some $\sigma'\colon \Theta \to M'$. Then σ' is a best response in (M',gh).

(ii) (Direct mechanisms are initial.) If $\mathsf{Impl}(f)$ is nonempty, then $(\Theta, f, \mathrm{id}_\Theta)$ is an object of it, and is initial.

*Hence f is implementable by some mechanism if and only if it is **incentive compatible**:*

$$U(f(\theta),\theta) \;\geq\; U(f(\theta'),\theta) \qquad \text{for all } \theta,\theta'.$$

Proof. (i) Every deviation in M' maps under h to a deviation in M: $U(gh(m'),\theta) \leq U(g\sigma(\theta),\theta) = U(gh\sigma'(\theta),\theta)$. (ii) Given (M,g,σ), apply (i) with $M' = \Theta$, $h = \sigma$, $\sigma' = \mathrm{id}$: truth-telling is a best response in the direct mechanism $(\Theta, g\sigma) = (\Theta, f)$. The map σ is a morphism $(\Theta,f,\mathrm{id}) \to (M,g,\sigma)$, and any morphism k from the direct mechanism satisfies $k = k \circ \mathrm{id} = \sigma$. $\square$

The two halves have different characters. Initiality is formal: any implementation receives a unique map from the direct one. The content is in (i): the equilibrium property is stable under pulling a mechanism back along a map through which the equilibrium factors, because restricting the message space can only remove deviations. Being an equilibrium is a property that is *contravariant* in the mechanism. With several agents the same argument applies to each agent's interim problem (Bayesian implementation) or state by state (dominant strategies), and the conclusion is the same.

It is as well to be clear about what the principle does not say. It says nothing about *other* equilibria of the direct mechanism; implementation in all equilibria is a harder problem with its own theory. It assumes that the designer can commit to g; without commitment the designer is a player and the principle fails (Chapters 12 and 18). And it treats communication as free.

14.2 Transfers: the incentive metric

Suppose outcomes are pairs of an allocation $x \in X$ and a payment $t \in \mathbb{R}$, and utility is quasi-linear, $U = v(x,\theta) - t$. A direct mechanism is an allocation rule $f\colon \Theta \to X$ with a payment rule $t\colon \Theta \to \mathbb{R}$, and incentive compatibility reads

$$t(\theta) - t(\theta') \;\le\; \ell_f(\theta',\theta) :\equiv v(f(\theta),\theta) - v(f(\theta'),\theta) \qquad \text{for all } \theta,\theta'. \tag{14.1}$$

The number $\ell_f(\theta',\theta)$ is what type θ gains from receiving her own allocation rather than that of θ'. Read (14.1) as a system of difference constraints on a complete directed graph with vertex set Θ and arc lengths ℓ_f: payments are a *potential*, a function whose increase along each arc is at most the arc's length.

Recall from Section 7.7 that a category enriched in $([0,\infty],\ge,+)$ is a generalized metric space. Lengths here may be negative, so enrich instead in the extended reals $\overline{\mathbb{R}} = ([-\infty,\infty],\ge,+)$, as in Lawvere's treatment of the Legendre transform [Law73, Wil15]. An $\overline{\mathbb{R}}$-category is a set with a function d to $\overline{\mathbb{R}}$ satisfying $d(a,b) + d(b,c) \ge d(a,c)$ and $0 \ge d(a,a)$. The free $\overline{\mathbb{R}}$-category on a weighted graph is its shortest-path closure.

Definition 14.2. The ***incentive category*** Θ_f of an allocation rule f is the free $\overline{\mathbb{R}}$-category on the graph (Θ,ℓ_f): its hom-values are

$$d_f(\theta',\theta) = \inf\Big\{ \textstyle\sum_{j=1}^{k} \ell_f(\theta_{j-1},\theta_j) \;\Big|\; \theta' = \theta_0,\theta_1,\dots,\theta_k = \theta \Big\}.$$

Theorem 14.3 (Implementability and revenue equivalence, enriched). [reformulation; cf. [Roc87, HMUV09, Voh11]] *Let $f\colon \Theta \to X$ be an allocation rule.*

(a) There is a payment rule making f incentive compatible if and only if the graph (Θ,ℓ_f) has no cycle of negative length; equivalently $d_f(\theta,\theta) = 0$ for all θ; equivalently d_f never takes the value $-\infty$. (Cyclic monotonicity.)

(b) *The incentive-compatible payment rules for f are exactly the real-valued $\overline{\mathbb{R}}$-functors $\Theta_f \to \overline{\mathbb{R}}$, that is, the functions t with $t(\theta) - t(\theta') \leq d_f(\theta', \theta)$.*

(c) *For each θ_0, the corepresentable $d_f(\theta_0, -)$ is the pointwise greatest incentive-compatible payment rule vanishing at θ_0, and $-d_f(-, \theta_0)$ is the least.*

(d) (Revenue equivalence.) *Any two incentive-compatible payment rules for f differ by a constant if and only if $d_f(\theta, \theta') + d_f(\theta', \theta) = 0$ for all θ, θ': every hom-value of Θ_f is invertible, with inverse the hom-value in the opposite direction.*

Proof. (b) A function satisfies the constraint on every arc iff it satisfies it along every path, iff it satisfies it for the infimum over paths. (a) Summing (14.1) around a cycle shows that a potential forces every cycle to have nonnegative length. Conversely if there are no negative cycles then d_f is finite, $d_f(\theta, \theta) = 0$, and $d_f(\theta_0, -)$ satisfies (b) by the triangle inequality. (c) If t satisfies (b) and $t(\theta_0) = 0$ then $t(\theta) \leq d_f(\theta_0, \theta)$ and $-t(\theta) \leq d_f(\theta, \theta_0)$; both bounds are themselves payment rules by the triangle inequality. (d) By (c) the payment rules vanishing at θ_0 fill the interval between $-d_f(-, \theta_0)$ and $d_f(\theta_0, -)$, which is a point iff the two coincide. □

Parts (a) and (d) are Rochet's theorem [Roc87] and the characterization of revenue equivalence by Heydenreich, Müller, Uetz and Vohra [HMUV09]; the network formulation is developed at length in [Voh11]. What the enriched phrasing contributes is the identification of the objects. An allocation rule *is* a generalized metric on the type space. A payment rule *is* a copresheaf. Part (c) is the enriched Yoneda lemma: the payment rules are bounded by the representables, and the representable at the type whose participation constraint binds is the revenue-maximizing one. Condition (d) says that the incentive category is a groupoid in the enriched sense; it then has no room for two different potentials.

Example 14.4 (Two types)**.** Let $v(x, \theta) = \theta x$ with types $\theta_L = 2$, $\theta_H = 3$ and allocation $f(\theta_L) = 1$, $f(\theta_H) = 3$. Then $\ell_f(L, H) = \theta_H(3 - 1) = 6$ and $\ell_f(H, L) = \theta_L(1 - 3) = -4$. The only cycle has length $2 \geq 0$, so f is implementable; $d_f(L, H) = 6$, $d_f(H, L) = -4$, and the incentive-compatible payments are those with $4 \leq t_H - t_L \leq 6$. Revenue equivalence fails, since $6 + (-4) \neq 0$: with finitely many types there is slack. A seller who wants the largest payments consistent with the low type's participation, $t_L = \theta_L \cdot 1 = 2$, takes the representable: $t_H = t_L + d_f(L, H) = 8$. The "binding downward incentive constraint" of the textbooks is a shortest path from the lowest type.

14.3 One-dimensional types

Let $\Theta = [\underline{\theta}, \overline{\theta}]$ and $v(x, \theta) = \theta x$, with $x \in \mathbb{R}$ a quantity or a probability of receiving a good.

Corollary 14.5 (Myerson's lemma). [classical; [Mye81]] *An allocation rule $x(\cdot)$ is implementable iff it is nondecreasing. The payment rule is then unique up to a constant:*

$$t(\theta) = t(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} s\, dx(s) = \theta x(\theta) - \int_{\underline{\theta}}^{\theta} x(s)\, ds - \text{const},$$

and the agent's indirect utility satisfies the envelope formula $U(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} x(s)\, ds$.

Proof. Here $\ell(\theta', \theta) = \theta\,(x(\theta) - x(\theta'))$. The two-cycle condition $\ell(\theta', \theta) + \ell(\theta, \theta') = (\theta - \theta')(x(\theta) - x(\theta')) \geq 0$ is monotonicity, and for monotone x longer cycles add nothing. For $\theta' < \theta$, a path that climbs from θ' to θ through a fine partition has length equal to an upper Riemann–Stieltjes sum of $\int s\, dx(s)$, and refining decreases it; so $d(\theta', \theta) = \int_{\theta'}^{\theta} s\, dx(s)$. A descending path gives a lower sum with a minus sign, so $d(\theta, \theta') = -\int_{\theta'}^{\theta} s\, dx(s)$. The two are inverse, and Theorem 14.3(d) applies. □

The contrast with Example 14.4 isolates what revenue equivalence needs. With two types there is one path each way and the lengths do not cancel. On an interval, paths can be refined without limit, and upper and lower sums converge to the same integral. Revenue equivalence is a consequence of the *connectedness* of the type space.

In several dimensions, with v linear in $\theta \in \mathbb{R}^n$, the incentive metric is obtained by integrating the vector field $\theta \mapsto f(\theta)$ along paths, and implementability requires two things: monotonicity, and that the integral around every closed loop vanish. The second is a flatness condition: the allocation rule must be a gradient. On a simply connected domain it can be checked locally; this is the content of the local-to-global theorems of mechanism design, in which weak monotonicity together with path-independence around small loops implies incentive compatibility [JMS99, SY05, AK08, Car12]. The reader will recognize Theorem 5.7. Posteriors are local data that glue to a common prior when a cocycle is trivial; an allocation rule is local data that glues to a payment rule when a cocycle is trivial and a sign condition holds. In both, the obstruction lives in loops, and uniqueness of the potential is a statement about connected components.

14.4 Duality: prices and utilities

There is a second description of the same structure, in which the mechanism is a price list rather than a direct revelation game. A *tariff* is a function $p\colon X \to \overline{\mathbb{R}}$; faced with it, type θ obtains

$$p^v(\theta) = \sup_x \big[v(x,\theta) - p(x)\big].$$

Conversely a utility profile $U\colon \Theta \to \overline{\mathbb{R}}$ determines the tariff $U^v(x) = \sup_\theta [v(x,\theta) - U(\theta)]$, the lowest prices at which no type would do better than U.

Proposition 14.6 (Taxation principle as a nucleus). [reformulation; cf. [Roc87, Car01, Wil15]] *The maps $p \mapsto p^v$ and $U \mapsto U^v$ form a Galois connection between tariffs on X and utility profiles on Θ: the adjunction induced by the valuation v regarded as an $\overline{\mathbb{R}}$-profunctor from X to Θ. Its closed elements on the utility side, the profiles with $U = (U^v)^v$, are the **v-convex** functions, and*

(i) an allocation rule f is implementable iff there is a v-convex U with $f(\theta)$ in the v-subdifferential of U at every θ, that is, $U(\theta') \geq U(\theta) + v(f(\theta),\theta') - v(f(\theta),\theta)$;

(ii) every incentive-compatible direct mechanism is equivalent to letting the agent choose from a tariff, and conversely.

When $v(x,\theta) = \langle x,\theta\rangle$ the transform is the Legendre–Fenchel conjugate, v-convex means convex and lower semicontinuous, and (i) is Rockafellar's theorem that the cyclically monotone correspondences are those contained in the subdifferential of a convex function.

Willerton [Wil15] observed that the Legendre–Fenchel transform is the nucleus of the pairing between a vector space and its dual, viewed as a profunctor enriched in $\overline{\mathbb{R}}$; the proposition says that nonlinear pricing is the same construction for the valuation. The dual pair of descriptions—the agent's indirect utility on the type space, the designer's tariff on the allocation space—is an instance of the duality between presheaves and copresheaves that a profunctor always induces. In the theory of optimal transport the same transforms appear under the name of c-conjugates, and the modern treatment of multidimensional screening exploits the identification [FKM11].

14.5 Optimal screening

The designer now has an objective. A monopolist sells a quantity x produced at cost $C(x)$ to a buyer with utility $\theta x - t$, and wants to maximize expected profit.

Example 14.7 (Two types)**.** Let $C(x) = \frac{1}{2}x^2$ and $\theta \in \{2, 3\}$ with equal probability. With types observable the seller would supply $x = \theta$ and extract all the surplus, earning $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 4.5 = 3.25$. With types hidden, the constraints that bind are the low type's participation and the high type's incentive constraint (the shortest path of Example 14.4). The seller's problem reduces to choosing x_L to maximize $\frac{1}{2}(2x_L - \frac{1}{2}x_L^2) - \frac{1}{2}(3-2)x_L$, the low type's surplus less the *information rent* $(\theta_H - \theta_L)x_L$ that must be left to the high type. The solution is

$$x_L = 1, \quad t_L = 2; \qquad x_H = 3, \quad t_H = 8,$$

with profit 2.5. The high type receives the efficient quantity and a rent of 1; the low type's quantity is distorted downward, from 2 to 1, not because serving her is unprofitable but because every unit sold to her raises the rent that the high type can claim by imitating her.

With a continuum of types distributed with density ϕ and distribution function F on $[\underline{\theta}, \overline{\theta}]$, Corollary 14.5 lets payments be eliminated. Expected profit equals

$$\int \Big[\Big(\theta - \frac{1 - F(\theta)}{\phi(\theta)} \Big) x(\theta) - C(x(\theta)) \Big] \phi(\theta)\, d\theta$$

less the utility of the lowest type, to be maximized over nondecreasing $x(\cdot)$. The seller behaves as if each type's willingness to pay were the ***virtual type*** $\theta - \frac{1-F}{\phi}$: the true type less the marginal rent that an increase in her allocation gives to all higher types. If the virtual type is increasing, pointwise maximization gives $C'(x(\theta)) = \theta - \frac{1-F(\theta)}{\phi(\theta)}$: no distortion at the top, downward distortion everywhere else, exclusion of types whose virtual type is negative. If it is not increasing, the monotonicity constraint binds and the solution is obtained by *ironing*: replacing the integrated virtual type by its convex hull [Mye81, MR84]—a closure under the Galois connection of Proposition 14.6 once more.

Example 14.8. With θ uniform on $[0, 1]$ and $C(x) = \frac{1}{2}x^2$, the virtual type is $2\theta - 1$ and $x(\theta) = \max\{0, 2\theta - 1\}$. Types below $\frac{1}{2}$ are excluded. The payment is $t(\theta) = \theta^2 - \frac{1}{4}$, which as a tariff on quantities is $p(x) = \frac{1}{2}x + \frac{1}{4}x^2$, and the profit is $\frac{1}{12}$.

The same analysis, with a regulator in the role of the seller and a firm with private cost in the role of the buyer, gives the theory of regulation under asymmetric information [BM82, LT86]: the optimal policy offers a menu, efficient firms choose high-powered contracts and earn rents, inefficient firms are given low-powered contracts and produce too little, and the regulator's problem is the trade-off between rent extraction and efficiency expressed by the virtual type.

14.6 Many agents: Groves mechanisms

Let agents $i = 1, \dots, n$ have valuations $v_i(x, \theta_i)$ for a collective decision x. An efficient decision rule $x^*(\theta)$ maximizes $\sum_i v_i(x, \theta_i)$.

Theorem 14.9 (Vickrey–Clarke–Groves). [classical; [Vic61, Cla71, Gro73, GL77]] *For any functions h_i of the others' reports, the payments*

$$t_i(\theta) = h_i(\theta_{-i}) - \sum_{j \neq i} v_j\big(x^*(\theta), \theta_j\big)$$

make truthful reporting a dominant strategy for every agent. If the type spaces are smoothly connected, every efficient mechanism with truth-telling in dominant strategies has this form. In general no choice of the h_i balances the budget.

Proof of the first claim. Agent i's payoff from reporting θ_i' is

$$\sum_j v_j\big(x^*(\theta_i', \theta_{-i}), \theta_j\big) - h_i(\theta_{-i}),$$

total surplus at the decision her report induces, minus a term she cannot affect. Truth induces the surplus-maximizing decision. □

Three threads of this book meet in the proof. The payment adds to v_i a function of the others' valuations, so that agent i's payoff becomes incentive equivalent, in the sense of Definition 1.4, to the common objective: the mechanism turns the game into one of common interest without changing anyone's intrinsic preferences. The uniqueness clause is revenue equivalence (Theorem 14.3(d)) applied agent by agent: once the decision rule is fixed, the payments are determined up to a constant in own type, that is, up to h_i. And the failure of budget balance is the team problem of Theorem 12.1: each agent is made residual claimant of the whole surplus, and there is only one surplus.

Example 14.10 (The pivot mechanism)**.** A project costing 10 is to be financed by equal shares if built. Two agents value it at $\theta_1 = 9$ and $\theta_2 = 3$, so their net values are $+4$ and -2; the efficient decision is to build. Choosing h_i to be the others' best attainable net surplus without i gives the *pivot* payments: agent 1, whose presence changes the decision and thereby imposes a loss of 2 on agent 2, pays 2 on top of her share; agent 2 pays nothing extra. The designer collects a surplus of 2 that cannot be returned to the agents without disturbing their incentives.

Notes

The revelation principle is due to Gibbard [Gib73] for dominant strategies and to Myerson [Mye79] and others for Bayesian equilibrium. Rochet [Roc87] characterized implementability in quasi-linear environments by cyclic monotonicity; the graph-theoretic view was developed by Vohra and coauthors [Voh11], and revenue equivalence characterized in [HMUV09]. The envelope formula in generality is due to Milgrom and Segal [MS02]. Lawvere's enriched reading of metric spaces is [Law73] and Willerton's of the Legendre transform [Wil15]; generalized convexity in the principal–agent problem is from [Car01], and the link with optimal transport is exploited in [FKM11]. Optimal nonlinear pricing is due to Mussa and Rosen [MR78] and Maskin and Riley [MR84]; optimal regulation to Baron and Myerson [BM82] and Laffont and Tirole [LT86]. The efficient mechanisms are from [Vic61, Cla71, Gro73], with uniqueness in [GL77]. The statement of the revelation principle as an initial object, and of incentive-compatible payments as enriched copresheaves, are reformulations for which we know no reference but which are immediate from the cited results.

Exercises

Exercise 14.1. State and prove Theorem 14.1 for n agents and Bayesian Nash equilibrium, making explicit the interim utility functions to which part (i) is applied.

Exercise 14.2. For three types $\theta \in \{1,2,3\}$ with $v = \theta x$ and allocation $(x_1, x_2, x_3) = (0, 1, 3)$, compute the incentive category Θ_f (all nine hom-values), the set of incentive-compatible payment rules with $t_1 = 0$, and the greatest one. Check the triangle inequality.

Exercise 14.3. Discretize $[0,1]$ into n equally spaced types and let $x(\theta) = \theta$. Compute $d_f(0,1) + d_f(1,0)$ as a function of n and verify that it tends to zero.

Exercise 14.4. Show that $p \mapsto p^v$ and $U \mapsto U^v$ are order-reversing and that $p^{vvv} = p^v$. For $v = \theta x$ on $X = \Theta = \mathbb{R}$, compute the tariff conjugate to $U(\theta) = \frac{1}{2}\theta^2$.

Exercise 14.5. Verify Example 14.8, and compute the seller's profit if she is restricted to offering a single bundle: one quantity x at one price t.

Exercise 14.6. In the pivot mechanism for a binary public project with n agents, show that the payments are nonnegative, that at most the pivotal agents pay, and that truthful reporting remains optimal even if an agent knows the others' reports.

Chapter 15

Signalling

If the informed party moves first, her action can reveal what she knows. Whether it does, and how much, is determined in equilibrium: the meaning of a signal is the Bayesian inverse of the strategy that produces it. This chapter treats four models of strategic information transmission, ordered by how much the sender can commit: costly signals, costless messages, signals whose meaning the sender designs in advance, and signals that the sender distorts in secret.

15.1 Signalling games

Nature draws a type $\theta \in \Theta$ from a prior π. The sender observes θ and chooses a signal $e \in E$; the receiver observes e and chooses an action a. A sender strategy is a kernel $\sigma\colon \Theta \to E$. The receiver's belief after e is $\mu(e) \in \Delta\Theta$, and in a perfect Bayesian equilibrium μ is the Bayesian inverse $\sigma_\pi^\dagger$ wherever that is defined, the receiver best-responds to μ, and the sender best-responds to the receiver.

The classical trichotomy of equilibria is a trichotomy of kernels. In a ***separating*** equilibrium σ is an injective function, so the inverse is deterministic and the receiver learns the type. In a ***pooling*** equilibrium σ is constant, the inverse is the prior, and the receiver learns nothing. In a ***hybrid*** equilibrium σ is properly stochastic. Equilibria are partially ordered by the Blackwell order of Section 11.3 on the experiments σ: the question "how much information is transmitted?" has a precise meaning and, usually, several answers.

15.2 Education as a signal

Example 15.1 (Spence). [classical; [Spe73]] A worker's productivity is $\theta_L = 1$ or $\theta_H = 2$; the fraction of high types is λ. Before entering the labour market she chooses a level of education $e \geq 0$ at cost e/θ. Education does not affect productivity. Competing employers observe e and pay the expected productivity given e.

Separating equilibria. The low type chooses $e = 0$ and is paid 1; the high type chooses e^* and is paid 2. The low type must not wish to imitate, $2 - e^* \leq 1$, and the high type must not wish to give up, $2 - e^*/2 \geq 1$. So every $e^* \in [1, 2]$

gives a separating equilibrium, supported by the belief that any $e < e^*$ comes from the low type.

Pooling equilibria. Both types choose the same e_p and are paid $1 + \lambda$, with deviations attributed to the low type; this is an equilibrium iff $e_p \le \lambda$.

Welfare. In the least wasteful separating equilibrium, $e^* = 1$, the high type nets 1.5 and the low type 1. In the pooling equilibrium at $e_p = 0$ both net $1 + \lambda$. If $\lambda > \frac{1}{2}$, pooling is better for *both* types: the high types burn half a unit each to distinguish themselves from a small minority.

What makes separation possible is that the cost of the signal falls with type: the *single-crossing* condition of Theorem 7.8. The high type's indifference curves in the plane of education and wage are flatter, so there are contracts she prefers and the low type does not. What makes signalling socially costly is that the signal is otherwise useless. Education in this model produces nothing and reveals everything. The private return to signalling is the wage difference; the social return is zero, or the value of a better match if the information is used. As in the war of attrition of Chapter 6, a contest for a fixed prize dissipates it.

15.3 Refinements

The example has a continuum of separating equilibria and a continuum of pooling ones. They differ only in beliefs after signals that are never sent, which Bayes' rule does not restrict. Sequential equilibrium does not help; every belief at an unreached information set following a single deviation is consistent. Refinements for signalling games restrict off-path beliefs by asking which types could conceivably gain from a deviation.

Definition 15.2 (Intuitive criterion). [classical; [CK87]] Fix an equilibrium with payoffs $u^*(\theta)$. A type θ is *equilibrium-dominated* at an unsent signal e if her payoff from e is below $u^*(\theta)$ for every best response of the receiver to every belief. The equilibrium *fails the intuitive criterion* if for some unsent e some type would gain from sending e whenever the receiver's belief is concentrated on the types not equilibrium-dominated at e.

Proposition 15.3. [classical; [CK87]] *In Example 15.1 the only equilibrium outcome satisfying the intuitive criterion is the separating equilibrium with $e^* = 1$.*

Proof. At a separating equilibrium with $e^* > 1$, consider $e' \in (1, e^*)$. The low type cannot gain from e' even if believed to be high, since $2 - e' < 1$; the high type gains if believed high. So beliefs after e' must be concentrated on the high type, and she deviates. At a pooling equilibrium at e_p, a signal e' with $1 - \lambda < e' - e_p < 2(1 - \lambda)$ is equilibrium-dominated for the low type and attractive to the high type if believed high. □

The surviving outcome is the *Riley outcome*: the separating equilibrium with the least signalling. The argument is a form of forward induction, and in generic signalling games the outcomes that pass it include those of Kohlberg–Mertens stable sets (Section 3.4); stronger criteria of the same family select the Riley outcome with any number of types [BS87]. The selection has an awkward feature. As $\lambda \to 1$ the Riley outcome does not change: when almost everyone is a high type, each of them still spends $e^* = 1$ to avoid being taken for one of the vanishing few, although pooling at zero would be better for everyone. And at $\lambda = 1$ exactly, there is nothing to signal and $e = 0$. The refined equilibrium correspondence is discontinuous in the prior. This is the phenomenon of Theorem 10.11 in miniature, and a reason for caution about any refinement that is insensitive to probabilities.

15.4 Screening

If instead the uninformed party moves first, offering a menu of contracts from which the informed party chooses, the model is one of *screening*. The monopoly version is the nonlinear pricing of Chapter 14; the competitive version is the insurance market of Theorem 13.4. The comparison with signalling is instructive. The two games have the same boxes composed in a different order, and the equilibrium sets differ sharply. When the informed party moves first, the uninformed party's response to unexpected signals is unconstrained, and there are many equilibria. When the uninformed party moves first she commits to a response to *every* choice in advance, there are no off-path beliefs to play with, and the equilibrium is unique or fails to exist. Commitment by the uninformed side replaces multiplicity by non-existence. The Riley outcome and the Rothschild–Stiglitz separating pair are the same allocation, reached from the two sides.

15.5 Cheap talk

Suppose the signal is a costless, unverifiable message. Nothing then distinguishes the types' costs of sending it, and single crossing in the signal is unavailable. Information can still be transmitted, in proportion to the alignment of interests.

Theorem 15.4 (Crawford–Sobel). [classical; [CS82]] *Let the sender's type θ be uniform on $[0,1]$, the receiver choose $a \in \mathbb{R}$ with payoff $-(a-\theta)^2$, and the sender have payoff $-(a-\theta-b)^2$ with bias $b > 0$. Then every equilibrium is equivalent to a* partition equilibrium*: there are cutoffs $0 = c_0 < c_1 < \cdots < c_N = 1$; types in the same interval send the same message; the receiver responds with the midpoint. The cutoffs*

satisfy

$$c_{i+1} - c_i = c_i - c_{i-1} + 4b,$$

and an equilibrium with N intervals exists iff $2N(N-1)b < 1$.

Example 15.5. With $b = \frac{1}{20}$ the largest N with $2N(N-1) < 20$ is 3. The interval lengths are d, $d + \frac{1}{5}$, $d + \frac{2}{5}$ with sum 1, so $d = \frac{2}{15}$ and the partition is $[0, \frac{2}{15}], [\frac{2}{15}, \frac{7}{15}], [\frac{7}{15}, 1]$. The one-interval *babbling* equilibrium, in which messages are ignored, also exists, as it does for every b.

An equilibrium of the cheap-talk game is an *endogenous information structure*. The sender's strategy is a quotient map from the type space onto a set of messages, its fibres are the intervals, and the receiver's knowledge is the operator $p^*\forall_p$ of Definition 4.2 for that quotient. The players choose, in equilibrium, how coarse the receiver's partition will be. The more divergent the interests, the coarser it must be, because a fine partition would give the sender too many profitable ways to shade her report; and the partition is coarser at the top, where the sender's incentive to exaggerate has more room. The language is vague because the speaker is biased, and vaguest where the bias matters most.

15.6 Persuasion

Now give the sender commitment. Before learning the state she chooses an experiment, a kernel $\sigma\colon \Omega \to S$ from states to signals, and its realization is observed by the receiver. She cannot lie about the realization; she can only choose what will be learned.

A signal realization induces a posterior, so an experiment induces a distribution of posteriors $\tau \in \Delta(\Delta\Omega)$. Which τ arise?

Proposition 15.6 (Bayes plausibility). [reformulation; cf. [KG11, AM95]] *Let $m\colon \Delta\Delta\Omega \to \Delta\Omega$ be the multiplication of the distribution monad. A distribution of posteriors τ is induced by some experiment under the prior μ_0 if and only if $m(\tau) = \mu_0$. The experiments, up to relabelling of signals, are the points of the fibre $m^{-1}(\mu_0)$, and the Blackwell order on experiments is the order of mean-preserving spreads on this fibre.*

The multiplication m averages a distribution over distributions; the condition says that the posteriors average to the prior. That this is the only constraint is the "splitting lemma" of Aumann and Maschler. Given the receiver's best response, let $\hat{v}(\mu)$ be the sender's expected payoff when the receiver holds the posterior μ. The sender's problem is to maximize the linear functional $\tau \mapsto \int \hat{v}\, d\tau$ over the fibre $m^{-1}(\mu_0)$.

Theorem 15.7 (Kamenica–Gentzkow). [classical; [KG11]] *The sender's optimal value is* $(\operatorname{cav}\hat{v})(\mu_0)$*, where* $\operatorname{cav}\hat{v}$ *is the smallest concave function above* $\hat{v}$*. She benefits from persuasion iff* $(\operatorname{cav}\hat{v})(\mu_0) > \hat{v}(\mu_0)$*.*

Example 15.8. A regulator approves a drug iff she believes it effective with probability at least 0.6; the prior is 0.3; the firm wants approval regardless. Without information the drug is rejected. Full disclosure gets it approved with probability 0.3. The optimal experiment induces posteriors 0 and 0.6, with probabilities $\frac{1}{2}$ each: an effective drug always passes the test, and an ineffective one passes with probability $\frac{2}{7}$. The drug is approved half the time, although it works only three times in ten. Here $\hat{v} = 1_{[0.6,1]}$ and $\operatorname{cav}\hat{v}(\mu) = \min\{\mu/0.6, 1\}$. The receiver is exactly as well off as with no information: the sender extracts the entire value of the information she designs.

Persuasion is a mechanism-design problem in which the instrument is information rather than money, and its solution has the same form as those of Chapter 14: a hull. There the hull was the convex envelope that irons the virtual type; here it is the concave envelope of the value over the simplex of beliefs. With several receivers who play a game among themselves, the set of outcomes an information designer can induce is the set of Bayes correlated equilibria [BM16] of Section 6.4, and the designer is the mediator of that section with an objective of her own.

15.7 Signal jamming

Finally, let the sender influence the signal by an action that the receiver does not observe.

Example 15.9 (Career concerns). [classical; [Hol99]] A manager's output is $y = \eta + a + \varepsilon$, where $\eta \sim N(m, \sigma_\eta^2)$ is her unknown talent, a her unobserved effort at cost $c(a)$, and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. After seeing y the market pays her next-period wage equal to its expectation of η. If the market conjectures effort a^*, it computes

$$\mathbb{E}[\eta \mid y] = m + \kappa\,(y - m - a^*), \qquad \kappa = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\varepsilon^2}.$$

The manager, taking the conjecture as given, gains κ in future wage per unit of effort, so with discount factor δ she chooses $c'(a) = \delta\kappa$. In equilibrium $a = a^*$: the market correctly subtracts her effort, and her wage is exactly what it would have been had she and everyone else known she would not exert it. She works to bias an inference that is not biased.

The logic is the lens structure of Chapters 2 and 5 made vivid. The market's inference is the backward pass, and it is computed from the *conjectured* forward pass. A deviation by the manager changes the actual forward pass and leaves the backward pass fixed. So although no one is fooled on the equilibrium path, the equilibrium effort is pinned down by the return to fooling off it. The same structure gives *limit pricing* [MR82b]: an incumbent whose cost is private information prices below the monopoly level to signal low cost to a potential entrant; in the separating equilibrium the entrant infers the cost correctly and enters exactly when she would under complete information, and the only effect of the attempt to deter entry is a lower price before it. In signal-jamming models of predation [FT86] the same holds with an unobserved action in place of a private type. In all three, competition for the receiver's belief is wasteful or beneficial according to whether the distorting action is, as it happens, something third parties value: effort and low prices are; education, in Spence's model, is not.

Notes

Signalling is due to Spence [Spe73]. The intuitive criterion is from Cho and Kreps [CK87] and its strengthenings from Banks and Sobel [BS87]; their relation to strategic stability is discussed in both. Strategic information transmission with costless messages is Crawford and Sobel's [CS82]. Bayesian persuasion is due to Kamenica and Gentzkow [KG11], building on the concavification technique of Aumann and Maschler [AM95]; the many-receiver generalization is [BM16]. Career concerns are from Holmström [Hol99], limit pricing as signalling from Milgrom and Roberts [MR82b], and signal jamming from Fudenberg and Tirole [FT86]. The source text [Ras06, ch. 11] treats the education model in several variants and discusses when signalling is productive; the distinction between informed-moves-first and uninformed-moves-first models is its organizing device and has been kept here.

Exercises

Exercise 15.1. In Example 15.1, find all hybrid equilibria in which the low type randomizes between $e = 0$ and the high type's signal. Which of them satisfy the intuitive criterion?

Exercise 15.2. Suppose education is productive: a worker of type θ with education e produces $\theta(1 + \frac{1}{2}\sqrt{e})$. Find the Riley outcome and compare the high type's education with its efficient level.

Exercise 15.3. For the cheap-talk model, derive the cutoff equation from the indifference of the boundary type, and compute the receiver's expected loss in the N-interval equilibrium for $b = \frac{1}{20}$ and $N = 1, 2, 3$.

Exercise 15.4. Verify Example 15.8, including the probability $\frac{2}{7}$, and compute the optimal experiment when the prior is 0.5.

Exercise 15.5. Show that the map $m\colon \Delta\Delta\Omega \to \Delta\Omega$ is affine and that the fibre $m^{-1}(\mu_0)$ is convex. Exhibit two different experiments with two signals each that induce the same τ.

Exercise 15.6. In Example 15.9 let $c(a) = \frac{1}{2}a^2$. Compute equilibrium effort, and compare it with the efficient effort $c'(a) = 1$. For which parameters does the manager work too hard?

Part III

Applications

Chapter 16

Bargaining

Two parties can create a surplus by agreeing and must divide it. There are two classical theories of the division. Nash's axiomatic theory asks what a solution must look like if it respects the symmetries of the problem. Rubinstein's strategic theory writes down a game of offers and counteroffers and finds its equilibrium. This chapter presents each in the terms developed in Part I—the first as an equivariant natural map, the second as the unique fixed point of a contraction—and then shows, with the tools of Chapter 14, why bargaining under two-sided private information cannot be efficient.

16.1 The axiomatic solution

Definition 16.1. A ***bargaining problem*** is a pair (S, d) of a compact convex set $S \subset \mathbb{R}^2$ of feasible utility pairs and a disagreement point $d \in S$, with some $s \in S$ strictly better than d for both. A ***solution*** assigns to each problem a point of S.

Theorem 16.2 (Nash). [classical; [Nas50b]] *There is exactly one solution that is (i) Pareto efficient, (ii) symmetric, (iii) invariant under positive affine transformations of either player's utility, and (iv) independent of irrelevant alternatives: if $S' \subseteq S$ contains the solution of (S, d), then it is the solution of (S', d). It is*

$$\phi(S, d) = \operatorname*{arg\,max}_{s \in S,\ s \geq d} (s_1 - d_1)(s_2 - d_2).$$

Read the axioms as statements about a category. Let the objects be bargaining problems. Axioms (ii) and (iii) say that the solution is *equivariant* under the group generated by exchanging the players and rescaling utilities: it is a function on the type of bargaining problems *up to* the equivalences that von Neumann–Morgenstern utility imposes, in the sense of Principle 1.5 and Theorem 6.2. They are not substantive assumptions about bargaining; they are the requirement that the question be well posed. Axiom (iv) is of a different kind. It says that the solution is *natural* with respect to a class of non-invertible morphisms, the inclusions $S' \hookrightarrow S$ that keep the solution feasible. This is the substantive axiom, and different choices of morphisms give different theories: requiring instead monotonicity with respect to expansions

of S that preserve each player's best feasible payoff characterizes the solution of Kalai and Smorodinsky [KS75]. The axiomatic method in bargaining is the choice of a category on which the solution is to be a natural transformation.

Example 16.3. Two players divide one unit. Player 1 has utility x and player 2 has utility $\sqrt{1-x}$; disagreement gives both zero. The Nash solution maximizes $x\sqrt{1-x}$, so $x = \frac{2}{3}$. The more risk-averse player receives the smaller share: concavity makes her less willing to hold out for more.

16.2 The strategic solution

Definition 16.4. In the ***alternating-offers game***, player 1 proposes a division of a unit surplus; player 2 accepts, ending the game, or rejects, in which case a period passes and player 2 proposes; and so on. Player i discounts by $\delta_i \in (0,1)$ per period.

With a finite number T of rounds, after which disagreement gives zero, the game is an element of the initial algebra of Chapter 8 and is solved by a fold. In the last round the proposer takes everything. In the round before, the proposer must offer the other her discounted continuation value. With a common discount factor the first proposer's share is

$$x_T = 1 - \delta + \delta^2 - \cdots \pm \delta^{T-1} = \frac{1 - (-\delta)^T}{1+\delta},$$

which oscillates and converges to $\frac{1}{1+\delta}$.

Theorem 16.5 (Rubinstein). [classical; [Rub82, SS84b]] *The infinite-horizon alternating-offers game has a unique subgame perfect equilibrium. Player 1 proposes to keep*

$$x^* = \frac{1-\delta_2}{1-\delta_1\delta_2},$$

player 2 accepts at once, and each player always proposes the same division and accepts exactly the offers at least as good as her discounted proposer share.

Proof (Shaked–Sutton). The game is stationary: the subgame beginning at each of player i's proposals is the same game. Let M_i and m_i be the supremum and infimum of i's subgame-perfect payoffs in a subgame where she proposes. A responder gets at most $\delta_j M_j$ by rejecting, so she accepts anything better, and the proposer can secure $m_i \geq 1 - \delta_j M_j$. A responder can secure $\delta_j m_j$ by rejecting, so $M_i \leq 1 - \delta_j m_j$. Combining the four inequalities,

$$M_1 \leq \varphi(M_1), \qquad m_1 \geq \varphi(m_1), \qquad \text{where } \varphi(x) = 1 - \delta_2\,(1 - \delta_1 x).$$

The affine map φ is increasing with slope $\delta_1\delta_2 < 1$, so $x \leq \varphi(x)$ iff $x \leq x^*$ and $x \geq \varphi(x)$ iff $x \geq x^*$, where x^* is its fixed point. Hence $M_1 \leq x^* \leq m_1$. $\square$

The proof is worth comparing with the analysis of repeated games in Chapter 9. There too one considered the set of subgame-perfect payoffs and an operator under which it is a fixed point, and the operator B was monotone, with a greatest fixed point much larger than the least. Here the operator acts on the interval $[m_1, M_1]$ and is a *contraction*, with modulus $\delta_1\delta_2$; a contraction has one fixed point, and the interval collapses. In the terms of Theorem 9.6, bargaining is a *guarded* recursion: the continuation value enters each player's decision only as a single number, scaled by a discount factor. In a repeated game the continuation enters as a function from histories to promised values, whose *spread* supplies incentives, and discounting does not shrink the spread relative to what it has to deter. That is the structural reason why one infinite-horizon game has a unique equilibrium and the other has a folk theorem.

A corollary: the infinite-horizon solution is the limit of the finite-horizon ones, $x_T \to \frac{1}{1+\delta}$. Under a guard the initial algebra and the terminal coalgebra give the same answer, and the question whether there is a last period—so vexed for repeated games—does not matter.

The solution has the comparative statics one would hope for. The proposer has an advantage, which vanishes as offers become frequent. Patience is power: x^* increases in δ_1 and decreases in δ_2. If the time between offers is Δ and $\delta_i = e^{-r_i\Delta}$, then as $\Delta \to 0$,

$$x^* \longrightarrow \frac{r_2}{r_1 + r_2},$$

which is the asymmetric Nash bargaining solution with weights inversely proportional to the rates of impatience [BRW86]. The strategic theory thus recovers the axiomatic one, and tells us what the disagreement point and the bargaining weights are in terms of the primitives of a negotiation: the weights are relative patience, and the disagreement point is the payoff *during* disagreement, not the payoff from walking away. An outside option that a player could take up instead of continuing affects the outcome only if it is worth more to her than her equilibrium share; otherwise the threat to take it is not credible, and by the logic of Chapter 8 it is ignored.

16.3 Private information and delay

In the equilibrium of Theorem 16.5 agreement is immediate, so nothing is wasted. Real negotiations take time, and strikes happen. The standard explanation is private information. If the seller does not know how much the buyer values agreement, a low offer that would be accepted at once by every type leaves money on the table with the high types; a high offer that only high types accept, followed by lower offers later, discriminates among them. Delay is a screening device: only a buyer who values the good little is willing to wait.

The seller is then in the position of the durable-goods monopolist of Chapter 18, competing with her own future offers, and the analysis is postponed to that chapter. With private information on both sides there is a sharper result, which does not depend on any particular bargaining procedure.

16.4 The impossibility of efficient bargaining

A seller owns a good that she values at c; a buyer values it at v. The values are private and independently distributed with positive densities on intervals that overlap, so that trade is sometimes efficient and sometimes not, and nobody knows which. A bargaining procedure, together with an equilibrium of it, is by the revelation principle (Theorem 14.1) a direct mechanism: a probability of trade and a payment for each pair of reports.

Theorem 16.6 (Myerson–Satterthwaite). [classical; [MS83]] *There is no mechanism that is Bayesian incentive compatible, individually rational at the interim stage, budget balanced, and efficient (trade occurs exactly when $v > c$).*

Proof sketch, for a common interval of values and costs. Fix the efficient trading rule. By revenue equivalence (Theorem 14.3(d), applied to each trader's interim problem on her interval of types), every incentive-compatible mechanism with this rule gives each type of each trader the same expected utility as the Groves mechanism of Theorem 14.9, up to an additive constant for each trader. In that mechanism, when trade occurs the buyer pays c and the seller receives v: each captures the entire gain $v - c$. The lowest-value buyer and the highest-cost seller never trade and get zero, so participation forbids lowering either constant. The two traders together must therefore receive at least twice the expected gains from trade, while only the gains themselves are available. The shortfall is $\mathbb{E}[(v-c)^+] > 0$, and someone outside would have to pay it. (When the intervals merely overlap, the worst types may earn something in the Groves mechanism, the constants can be lowered by that much, and the computation of [MS83] shows that a positive shortfall remains.) □

The theorem is an *obstruction*, and its proof has the form that obstructions have had throughout this book. Incentive compatibility and the allocation rule determine payments up to constants, because the type spaces are connected. Participation bounds the constants on one side; budget balance fixes their sum. The three requirements are a system of equations and inequalities with no solution. Nothing in the argument refers to offers, counteroffers or deadlines: no bargaining institution whatever can do better.

Example 16.7. [classical; [CS83, MS83]] Let v and c be independent and uniform on $[0,1]$. In the sealed-bid double auction, where trade occurs at the average of

the bids if the buyer's bid exceeds the seller's, there is an equilibrium in linear strategies

$$b(v) = \tfrac{2}{3}v + \tfrac{1}{12}, \qquad s(c) = \tfrac{2}{3}c + \tfrac{1}{4},$$

in which trade occurs iff $v \geq c + \tfrac{1}{4}$. The expected gains from trade are $\tfrac{9}{64}$, against $\tfrac{1}{6}$ under efficiency: a fraction $\tfrac{27}{32}$ is realized. Myerson and Satterthwaite showed that this is the best that any mechanism satisfying the other three conditions can do. The parties fail to trade when the gains are positive but small, each having shaded her bid to improve the price.

The Coase theorem says that, absent transaction costs, parties bargain to an efficient outcome whatever the initial assignment of rights. Theorem 16.6 identifies private information as a transaction cost in this sense, and an irreducible one. It also shows that the assignment of rights matters: if ownership is shared rather than assigned to one party, the participation constraints are relaxed and efficient dissolution can become possible [CGK87].

Notes

The axiomatic theory is Nash's [Nas50b]; the monotonic alternative is Kalai and Smorodinsky's [KS75]. The strategic theory is due to Rubinstein [Rub82], following Ståhl's finite-horizon analysis; the proof given is Shaked and Sutton's [SS84b]. The relation between the two, and the role of outside options, are worked out in [BRW86]. The impossibility theorem is [MS83]; the linear equilibrium of the double auction is from Chatterjee and Samuelson [CS83]; efficient dissolution of partnerships is [CGK87]. A standard text is [OR90].

Exercises

Exercise 16.1. Verify Example 16.3, and compute the Nash solution when player 2's utility is $(1 - x)^\gamma$ with $0 < \gamma \leq 1$. How does her share vary with γ?

Exercise 16.2. Show that the Nash solution is equivariant under $(s_1, s_2) \mapsto (\alpha_1 s_1 + \beta_1, \alpha_2 s_2 + \beta_2)$, and give an example showing that it is not monotone with respect to expansions of S.

Exercise 16.3. Compute the finite-horizon shares x_T for $T = 1, \dots, 5$ and $\delta = 0.9$, and verify geometric convergence to $\tfrac{1}{1+\delta}$ at rate δ.

Exercise 16.4. In the infinite-horizon game suppose that player 2, whenever she rejects an offer, may leave and take an outside option worth w. Find the subgame perfect equilibrium for $w < \delta_2(1 - \delta_1 x^*)$ and for larger w.

Exercise 16.5. Verify that the strategies of Example 16.7 are mutual best responses, and compute the expected gains from trade.

Exercise 16.6. For v and c uniform on $[0,1]$, compute the expected budget deficit of the Groves mechanism with efficient trade in which the worst types get zero, and confirm that it equals $\mathbb{E}[(v-c)^+] = \frac{1}{6}$.

Chapter 17

Auctions

An auction is a mechanism for selling to buyers whose willingness to pay the seller does not know. It is the part of the subject where theory, experiment and practice are closest, and where the apparatus of Chapter 14 pays its way most visibly: the central theorem of the field is one line once payments are seen as potentials.

17.1 Formats, and two kinds of equivalence

Four formats are standard. In the *English* auction the price rises until one bidder remains. In the *Dutch* auction the price falls until someone claims the object. In the *first-price* sealed-bid auction the highest bidder wins and pays her bid; in the *second-price* auction she pays the highest losing bid. Values are *private* if each bidder knows her own value and it does not depend on what others know, and *common* if the object has the same unknown value to all and bidders hold different estimates of it.

Proposition 17.1. [classical; [Vic61]]

(i) The Dutch and first-price auctions have the same reduced strategic form.

(ii) With private values, the English and second-price auctions each have an equilibrium in weakly dominant strategies, with the same outcome.

The two statements are equivalences of different types, in the sense of Section 3.4. In a Dutch auction a strategy is a price at which to stop the clock, nothing is learned before the auction ends, and the bidder who names the highest price wins and pays it: the tree differs from that of the first-price auction by Thompson moves, and the two are the same point of $\mathsf{EFG}_{\mathrm{T}}$. This holds for any values and any information. The second equivalence is not an identification of games. An English auction reveals when others drop out, and has many more strategies than a sealed bid. With private values that information is useless and the outcomes coincide; with common values it is valuable and they do not. One equivalence is of games, the other of solutions under a hypothesis.

17.2 Private values

Let n bidders have values drawn independently from a distribution F with density ϕ on $[0, \bar{v}]$, and let them be risk neutral.

Second-price. Bidding one's value is weakly dominant: the bid determines whether one wins, not what one pays, and one wants to win exactly when the price is below one's value. The seller's revenue is the second-highest value.

First-price. A bidder must shade her bid. Suppose the others use an increasing strategy $b(\cdot)$. A bidder with value v who bids as if her value were z wins with probability $F(z)^{n-1}$ and earns $(v - b(z))F(z)^{n-1}$. Optimality at $z = v$ gives a differential equation whose solution is

$$b(v) = v - \frac{\int_0^v F(s)^{n-1}\,ds}{F(v)^{n-1}} = \mathbb{E}\big[\, Y \mid Y < v \,\big],$$

where Y is the highest of the other $n-1$ values. Each bidder bids what she expects the runner-up's value to be, conditional on winning. For the uniform distribution on $[0,1]$, $b(v) = \frac{n-1}{n} v$.

17.3 Revenue equivalence

In the first-price auction the winner pays her own shaded bid; in the second-price auction she pays the runner-up's unshaded value. For uniform values the expected revenue is $\frac{n-1}{n} \cdot \frac{n}{n+1}$ in the first and $\frac{n-1}{n+1}$ in the second. They are equal. This is not a coincidence of the uniform distribution.

Theorem 17.2 (Revenue equivalence). [classical; [Vic61, Mye81, RS81]] *Consider risk-neutral bidders with independent private values on an interval. Any two auctions that, in equilibrium, (i) allocate the object in the same way as a function of the values and (ii) give the lowest type the same expected utility, yield the same expected payment from every type of every bidder, and the same expected revenue.*

Proof. From the standpoint of bidder i, an auction with its equilibrium is a direct mechanism in which a report v leads to a probability $P(v)$ of winning and an expected payment $m(v)$. Her type space is an interval and her valuation is vP. By Corollary 14.5, m is determined by P up to a constant: $m(v) = m(0) + \int_0^v s\,dP(s)$. □

All of the content is in Theorem 14.3: on a connected type space the incentive category is a groupoid, and a potential on a connected groupoid is unique up to a constant. The theorem is as useful for what it computes as for what it compares. In any efficient auction $P(v) = F(v)^{n-1}$, so $m(v) = \int_0^v s\,dF(s)^{n-1}$. In a first-price auction the bidder pays her bid when she wins, so

$b(v) = m(v)/P(v)$, which is the formula above without solving a differential equation. In an *all-pay* auction she pays her bid regardless, so $b(v) = m(v)$; for uniform values, $b(v) = \frac{n-1}{n}v^n$. The same method gives equilibrium behaviour in wars of attrition, contests and lobbying games, which are auctions in which losers pay.

17.4 Optimal auctions

Revenue equivalence says that the seller cannot gain by changing the payment rule while keeping the allocation. She can gain by changing the allocation—specifically, by sometimes not selling.

Theorem 17.3 (Myerson). [classical; [Mye81, RS81]] *With independent private values and increasing virtual values $\psi(v) = v - \frac{1-F(v)}{\phi(v)}$, expected revenue equals the expected virtual value of the winner, and the revenue-maximizing auction awards the object to the bidder with the highest virtual value provided it exceeds the seller's own value v_0. With symmetric bidders this is any standard auction with the reserve price r solving $\psi(r) = v_0$.*

The reserve price does not depend on the number of bidders. For uniform values and $v_0 = 0$ it is $\frac{1}{2}$: the seller refuses bids below one half although every bidder values the object more than she does. This is the monopolist's distortion of Chapter 14, exclusion at the bottom to reduce the rent at the top. With two bidders the optimal reserve raises revenue from $\frac{1}{3}$ to $\frac{5}{12}$.

How much is this worth? Less than one more bidder. With three bidders and no reserve, expected revenue is $\frac{1}{2} > \frac{5}{12}$, and the comparison is general.

Theorem 17.4 (Bulow–Klemperer). [classical; [BK96]] *Under the assumptions of Theorem 17.3 with symmetric bidders, an English auction with $n+1$ bidders and no reserve yields at least as much expected revenue as the optimal auction with n bidders.*

A seller does better to spend effort on attracting participation than on designing the mechanism. It is one of the few results in the subject that says what theory is worth compared with marketing.

17.5 Risk aversion

Revenue equivalence needs risk neutrality, because it needs payoffs to be quasi-linear in the payment. If bidders are risk averse, bidding one's value remains dominant in the second-price auction, since the argument does not involve probabilities. In the first-price auction a higher bid buys insurance against losing, and risk-averse bidders bid more than risk-neutral ones would.

So the first-price auction yields strictly higher expected revenue [MR84b]. A risk-averse *seller*, by contrast, prefers the first-price auction for a different reason: its revenue, being a function of the highest value alone, is less variable.

17.6 Common values

Now let the object have a common value V and let each bidder observe a signal x_i correlated with it.

Example 17.5 (The wallet game). [classical; [Kle98]] Two bidders each know the amount x_i in their own wallet, independent and uniform on $[0,1]$; the prize is the sum $V = x_1 + x_2$, sold by second-price auction. A bidder's unconditional estimate of the prize is $x_i + \tfrac{1}{2}$. But if she wins against an opponent who bids increasingly in her signal, the opponent's signal is below her own, and the estimate *conditional on winning* is $\tfrac{3}{2}x_i$, which is lower unless $x_i = 1$. The symmetric equilibrium is $b(x) = 2x$: each bids what the prize would be worth if the opponent's signal were equal to her own. Then the winner pays $2x_j$ for a prize worth $x_i + x_j$ and earns $x_i - x_j > 0$, and the loser has no regret. Expected revenue is $\tfrac{2}{3}$, out of an expected value of 1.

The ***winner's curse*** is the failure to condition on the event of winning. The bidder who wins is the one with the most optimistic signal, so winning is bad news about the value; a bidder who bids her unconditional estimate will, on average, regret winning. Rational bidders shade, and shade more the more competitors there are, since beating more rivals is worse news. The curse is a statement about selection of exactly the kind met in Chapter 13: the event on which the payoff is realized is correlated with the unknown, and the relevant expectation is the one conditional on that event.

Theorem 17.6 (Milgrom–Weber). [classical; [MW82]] *Suppose the bidders' signals and the value are* affiliated*: their joint density g satisfies $g(z \vee z')\,g(z \wedge z') \geq g(z)\,g(z')$. Then*

(i) the second-price auction has the symmetric equilibrium $b(x) = \mathbb{E}\big[V \mid x_i = x,\ \max_{j \neq i} x_j = x\big]$;

(ii) expected revenues are ranked: English $\geq$ second-price $\geq$ first-price;

(iii) a seller with private information affiliated with the value raises expected revenue by committing to disclose it.

Affiliation says that $\log g$ is supermodular on the lattice of signal profiles: it is the complementarity condition of Chapter 7 applied to a density, and it means that high values of some variables make high values of the others more

likely. Parts (ii) and (iii) are instances of the *linkage principle*: the more the price paid by the winner is linked to information affiliated with her own signal—the runner-up's bid, the dropout prices of others, the seller's announcement—the less of an informational rent her private signal commands, and the higher the revenue. This is the informativeness principle of Theorem 11.3 in the service of the other side: there, tying pay to informative signals reduced the cost of incentives; here, tying price to them reduces the cost of eliciting values.

Affiliation also breaks revenue equivalence, and more drastically than (ii) suggests. With correlated private values the seller can extract the *entire* surplus by a mechanism that offers each bidder, alongside a second-price auction, a side bet on the others' reports that is fair for her true type and ruinous for any other [CM88]. That such mechanisms are never seen is a standing reminder that the common-prior, risk-neutral, unlimited-liability model is an idealization, and that robustness to its failure is a design criterion in its own right.

17.7 Asymmetric information among bidders

If one bidder knows the common value and another has only public information, the uninformed bidder cannot profit: any bid that wins against an informed rival wins precisely when the object is worth less. In the first-price equilibrium the uninformed bidder randomizes, earns zero in expectation, and by her presence forces the informed bidder to pay a price that averages the value over the worse states [EMW83]. The informed bidder's rent is the value of her information, and it comes out of the seller's revenue. Oil-lease auctions in which one firm owns the neighbouring tract are the canonical application, and one of the cleaner empirical confirmations of auction theory.

Notes

The theory begins with Vickrey [Vic61], who introduced the second-price auction and proved revenue equivalence for uniform values. The general theorem and the optimal auction are due to Myerson [Mye81] and to Riley and Samuelson [RS81]. The comparison of an extra bidder with an optimal reserve is from Bulow and Klemperer [BK96]. Risk aversion is treated in [MR84b]. The general symmetric model with affiliated values, and the linkage principle, are from Milgrom and Weber [MW82]; the wallet game is Klemperer's [Kle98]. Full surplus extraction is due to Crémer and McLean [CM88]. The model with one informed bidder is from [EMW83]. A standard text is [Kri10].

Exercises

Exercise 17.1. Derive the first-price equilibrium bid for n bidders with values uniform on $[0,1]$ both from the first-order condition and from revenue equivalence, and verify that expected revenue is $\frac{n-1}{n+1}$.

Exercise 17.2. For two bidders with uniform values, compute expected revenue of a second-price auction with reserve r, and verify the maximum $\frac{5}{12}$ at $r = \frac{1}{2}$.

Exercise 17.3. Use revenue equivalence to find the symmetric equilibrium of a *third-price* auction with three uniform bidders. Do bidders bid above or below their values?

Exercise 17.4. In the wallet game, verify that $b(x) = 2x$ is an equilibrium and that it is an ex post equilibrium: neither bidder would change her bid after learning the other's signal. Compute the expected revenue of the first-price auction, whose symmetric equilibrium is $b(x) = x$, and compare.

Exercise 17.5. Show that independent signals are affiliated, and that if (x, y) are affiliated and h is increasing then $\mathbb{E}[h(y) \mid x]$ is increasing in x.

Exercise 17.6. Explain, using Section 3.4, why the Dutch and first-price auctions are identified by Thompson moves while the English and second-price auctions are not, by describing the information sets of each.

Chapter 18

Pricing

The oldest application of game theory is to firms that are neither price-takers nor monopolists. This chapter revisits the models of Chapter 7 with the tools accumulated since: commitment, repetition, information and mechanism design each resolve, or relocate, a classical puzzle about how prices are set.

18.1 Quantities or prices?

Cournot's firms choose quantities and earn positive profits; Bertrand's choose prices and earn none (Examples 7.6 and 7.7). The two models make opposite predictions from the same primitives, and for a long time the choice between them was a matter of taste. The reconciliation is a two-stage game.

Theorem 18.1 (Kreps–Scheinkman). [classical; [KS83]] *Suppose two firms first choose capacities, at a cost, and then, knowing both capacities, choose prices, with demand rationed efficiently when the low-priced firm cannot serve it all. Then in the unique subgame-perfect outcome the capacities are the Cournot quantities and both firms charge the Cournot price.*

Quantity competition is price competition preceded by a commitment. A firm that has installed a small capacity has burned the bridge to aggressive pricing: it cannot serve the customers it would win by undercutting. This is the logic of Chapter 8—reducing one's own options changes the rival's context—and in the language of Chapter 7 the capacity stage converts a game with a discontinuous, knife-edge best response into one of strategic substitutes. Which model applies is then an empirical question about what is hard to change in the short run. Where capacity is the slow variable (cement, steel, shipping), Cournot; where it is not (software, generic drugs sold by firms with ample capacity), Bertrand.

18.2 Escaping the Bertrand paradox

The Bertrand outcome rests on a discontinuity: at equal prices firms share the market, and an arbitrarily small price cut captures all of it. Each of the following removes the discontinuity or its consequence.

Capacity constraints. If neither firm can serve the whole market, undercutting has limited appeal and the high-priced firm retains residual demand. For intermediate capacities there is no pure-strategy equilibrium: prices cycle, in Edgeworth's phrase, and the equilibrium is in mixed strategies, which is how one may read sales and price dispersion.

Differentiation. If products are imperfect substitutes, demand is continuous in prices, the price game is supermodular (Example 7.7), and prices exceed cost by a margin that measures the differentiation.

Repetition. If the firms meet repeatedly, the folk theorem applies. With n identical Bertrand competitors, sharing the monopoly profit π^m under the threat of reversion to marginal-cost pricing is an equilibrium iff $\frac{\pi^m/n}{1-\delta} \geq \pi^m$, that is

$$\delta \;\geq\; 1 - \frac{1}{n}.$$

Collusion is easier with fewer firms, with more frequent interaction, and with quicker detection. The structure is the two-state automaton of Example 9.9. In the terms of Proposition 9.12 the collusive price is sustained by strategies that are not Markov: the firms condition on a payoff-irrelevant history, and antitrust law's hostility to "concerted practices" is hostility to exactly that frame.

Search. If consumers must pay a small cost to learn each additional price, the outcome flips to the other extreme. No firm gains by pricing below its rivals, since nobody will find out; each can raise its price by less than the search cost without losing a customer; and the unique equilibrium has every firm at the *monopoly* price, however many firms there are [Dia71]. Bertrand and Diamond are the same discontinuity seen from its two sides: with zero search cost, price equals marginal cost, and with any positive search cost it equals the monopoly price. The realistic middle ground, with dispersed prices, requires heterogeneity among consumers in what they know.

18.3 Location

Differentiation is itself chosen. In Hotelling's model [Hot29] consumers are spread uniformly along a street of unit length and buy one unit from the firm with the lowest price plus transport cost.

If prices are fixed (by regulation, or because the "firms" are political candidates and the "price" is absent), two firms choosing locations both go to the centre: each can gain market share by moving towards the other, and the only equilibrium is back-to-back at the median. This is the *principle of minimum differentiation*, and in politics the median voter theorem. With three firms there is no pure-strategy equilibrium at all.

If the firms also choose prices, the conclusion reverses.

Example 18.2 (Price competition on a line)**.** Let the firms be at the two ends of the street, with marginal cost c, and let transport cost be quadratic, $t \cdot (\text{distance})^2$. The consumer at x is indifferent when $p_1 + tx^2 = p_2 + t(1-x)^2$, so firm 1's demand is $\hat{x} = \frac{p_2 - p_1 + t}{2t}$. Its best response is $p_1 = \frac{1}{2}(p_2 + c + t)$, increasing in p_2—prices are strategic complements—and the equilibrium is

$$p_1 = p_2 = c + t, \qquad \text{with profits } \tfrac{t}{2} \text{ each.}$$

The margin is the transport cost: firms earn exactly what their distance from each other protects.

Anticipating this, firms choosing locations before prices face two forces: moving towards the rival steals customers at given prices, and intensifies price competition. With quadratic transport costs the second dominates and the firms locate at the opposite ends, *maximal differentiation*; with Hotelling's original linear costs the price game has no pure-strategy equilibrium when the firms are close, so his argument for minimum differentiation was incomplete [dGT79]. The location stage is a commitment to soft pricing, and the sign of the strategic effect is the one predicted in Chapter 7 for complements.

18.4 Comparative statics without calculus

Because differentiated-product price games are supermodular, Theorem 7.5 applies wholesale. If any parameter raises each firm's marginal profitability of a higher price—an increase in its own cost, in a rival's cost, in demand, in a tax—then the greatest and least equilibrium price vectors rise. No assumption of concavity, differentiability or uniqueness is used, and the conclusion holds for any number of firms and any pattern of substitution that keeps the complementarity. The classical approach differentiates first-order conditions and signs the result by a stability condition, which amounts to assuming the conclusion. The lattice-theoretic approach gets it from the order structure of the model, and extends it to games with discrete choices, such as adoption of a technology with network effects, in which the greatest and least equilibria may be far apart and the comparative statics of each is still monotone.

18.5 Durable goods

A monopolist sells a durable good to buyers who each want one unit and can wait. Having sold to those who value it most, she would like to cut the price and sell to the rest; and having done that, to cut it again. Buyers who foresee the cuts will not pay the high price today.

Example 18.3 (Two periods)**.** Buyers' values are uniform on $[0,1]$; production is costless; both sides discount the second period by $\delta = \frac{1}{2}$. If the monopolist could commit to prices she would charge $\frac{1}{2}$ in both periods and earn $\frac{1}{4}$. Without commitment, suppose buyers with values above $\hat{v}$ buy in the first period. In the second period she faces values uniform on $[0,\hat{v}]$ and charges $p_2 = \hat{v}/2$. The marginal first-period buyer is indifferent, $\hat{v} - p_1 = \delta(\hat{v} - p_2)$, so $p_1 = \frac{3}{4}\hat{v}$. Maximizing $p_1(1-\hat{v}) + \delta\, p_2 \cdot \frac{\hat{v}}{2}$ gives

$$\hat{v} = 0.6, \qquad p_1 = 0.45, \qquad p_2 = 0.3, \qquad \text{profit } 0.225 < 0.25.$$

Both prices are below the monopoly price and profit is lower: the monopolist is hurt by her own later self.

Theorem 18.4 (Coase conjecture)**.** [classical; [Coa72, GSW86]] *In the infinite-horizon version, as the time between offers shrinks to zero, in every stationary equilibrium the opening price converges to the lowest buyer value (or marginal cost) and the monopolist's profit converges to the competitive level.*

The seller of a durable good competes with her future selves, and when they can move quickly they compete her profits away. The remedies are commitment devices: leasing instead of selling, so that she bears the loss when the price falls; most-favoured-customer clauses that oblige her to refund earlier buyers; limited capacity; a reputation for not cutting prices, in the sense of Chapter 10.

Three earlier themes converge here. First, this is the renegotiation problem of Section 8.5 and Chapter 12: the commitment plan is not subgame perfect, because after the first sale both the seller and the remaining buyers gain from a new deal. Second, it is the screening problem of Chapter 14 without commitment. Declining prices sort buyers by their values, exactly as a menu would; the revelation principle told us that with commitment the seller can do no better than a single posted price, and Theorem 18.4 tells us that without commitment she does much worse. The revelation principle's assumption that the designer can bind herself is not a technicality. Third, it is the bargaining problem of Chapter 16 with one-sided private information: a seller making repeated offers to a single buyer of unknown value is formally the same game, and the Coase conjecture says that the uninformed party's power to make all the offers is worthless when offers can be made quickly.

18.6 Entry

The threat of entry disciplines an incumbent only if she cannot deter it, and the theory of deterrence is a catalogue of the commitment and information effects

already studied. An incumbent may install capacity that makes post-entry aggression credible [Dix80]; this is Example 8.6 with the follower's option to stay out. Whether she should over-invest or under-invest depends on whether the post-entry game is one of strategic substitutes or complements and on whether investment makes her tough or soft: the fourfold taxonomy of [FT84]. She may price low to signal low cost (Chapter 15), or fight early entrants to maintain a reputation (Example 10.8). In each case the conclusion of the equilibrium analysis is more guarded than the informal story. Capacity deters only if it is sunk; a low price deters only if it separates, and then entry occurs exactly when it would have anyway; a reputation deters only if there is some prior chance that it is deserved. The models do not say that predation and limit pricing are myths. They say what must be true of the world for them to work.

Notes

The quantity and price models go back to Cournot and Bertrand; their reconciliation through capacity choice is due to Kreps and Scheinkman [KS83]. Collusion in repeated price competition follows [Fri71]. The search paradox is Diamond's [Dia71]. Spatial competition is from Hotelling [Hot29], corrected by d'Aspremont, Gabszewicz and Thisse [dGT79]. The durable-goods problem was posed by Coase [Coa72] and resolved for stationary equilibria by Gul, Sonnenschein and Wilson [GSW86]; the two-period example is in the manner of [Bul82]. Capacity as commitment is [Dix80], and the taxonomy of strategic investment is [FT84]. The standard text for this chapter is [Tir88], and the source text's treatment is [Ras06, ch. 14].

Exercises

Exercise 18.1. For the linear market of Example 7.6 with n firms, find the least discount factor at which grim trigger sustains the monopoly outcome under quantity competition, and compare with the Bertrand threshold $1 - \frac{1}{n}$.

Exercise 18.2. Two firms each have capacity k and zero cost; demand is $1 - p$. For which k is $p_1 = p_2 = 1 - 2k$ a pure-strategy price equilibrium? What happens to the pure-strategy equilibrium when $k \geq 1$?

Exercise 18.3. Verify Example 18.2. Then let the firms be located at a and $1 - a$ with $0 \leq a < \frac{1}{2}$, compute the price equilibrium, and show that each firm's profit decreases in a.

Exercise 18.4. Three candidates choose positions on $[0,1]$ with voters uniformly distributed, each voter voting for the nearest candidate, and each candidate maximizing her vote share. Show that there is no pure-strategy equilibrium.

Exercise 18.5. Solve Example 18.3 for a general discount factor δ and show that the monopolist's profit is below $\frac{1}{4}$ for all $\delta \in (0,1)$. What happens as $\delta \to 0$ and as $\delta \to 1$?

Exercise 18.6. In Example 18.3, suppose the monopolist leases the good period by period instead of selling it. Find the equilibrium rental prices and her profit.

Part IV

Assessment

Chapter 19

What the Structure Buys

A change of language should be judged by what it lets one say that could not be said before, by what it shows to be the same that looked different, and by the questions it makes it possible to ask. It should not be judged by the dignity of its vocabulary. This chapter takes stock on those terms, and ends with problems.

19.1 What was shown to be the same

The classical theory presents its ideas one model at a time. The principal benefit of the categorical language has been to exhibit a small number of constructions recurring across models.

One invariance principle. A solution concept is a function on a type of games, and so respects the identifications of that type (Principle 1.5). Applied to the automorphisms of a game, this is the symmetry axiom of equilibrium selection, and it shows that a focal point must be structure (Proposition 1.14). Applied to the Thompson moves, it is Kohlberg–Mertens invariance (Section 3.4). Applied to bisimilar types, it separates interim correlated rationalizability from Bayesian equilibrium (Example 5.6). Applied to equivalent continuation games, it is the Markov restriction, and it shows that the folk theorem's cooperation is a focal structure on histories (Proposition 9.12). Four literatures, developed separately over thirty years, are one requirement and four choices of type.

One modality. Knowledge, common knowledge, distributed knowledge and subgame perfection are all of the form $p^*\forall_p$ or its relatives for a suitable map p (Definition 4.2, Theorem 4.6, and Proposition 8.4), and sequential rationality is their composite.

One obstruction theory. The existence of a common prior for given posteriors, and the existence of incentive-compatible payments for a given allocation, are both the triviality of a 1-cocycle; uniqueness in both is connectedness; and the local-to-global theorems of mechanism design are the statement that the obstruction lives in loops (Theorems 5.7 and 14.3). The impossibility theorems of Holmström, of Green and Laffont, and of Myerson and Satterthwaite are computations with the resulting potentials.

Two universal properties. Finite-horizon games are elements of an initial al-

gebra and infinite-horizon games of a terminal coalgebra. Backward induction, the one-shot deviation principle, self-generation, the automaton representation of strategies, recursive contracts, the universal type space and rationalizability are instances of recursion and corecursion (Chapters 5, 8 and 9). Whether the two agree is governed by whether the relevant operator is a contraction, and this one distinction separates the theory's uniqueness results from its multiplicity results (Section 7.7).

One algebra of composition. Normal-form games, sequential games, principal–agent problems, rational-expectations pricing and signal jamming are composites of the same few boxes in different orders, and the forward/backward discipline of optics is what distinguishes a deviation from a change of conjecture (Chapters 2 and 15).

19.2 What was said that could not be said before

Less than one might like, and it is as well to be exact.

Three statements in this book are, as far as we know, not in the literature in the form given: that the homotopy colimit of the type maps is a graph whose π_0 is common knowledge and whose H^1 obstructs common priors, with Hellman–Samet tightness as its contractibility (Proposition 4.9 and Theorem 5.7); that incentive-compatible payment rules are the copresheaves on an enriched category determined by the allocation, with revenue equivalence as a groupoid condition (Theorem 14.3); and that the revelation principle is an initiality statement resting on the contravariance of the equilibrium property (Theorem 14.1). Each is a repackaging of known theorems, with a short proof from them. None enlarges the set of facts about games. Their value, if they have any, is in suggesting generalizations, and the problems below are meant to test that.

The same must be said of the word "higher". The homotopy types that appeared are graphs. The categories that appeared are 1-categories, a bicategory and a double category of open games, and categories enriched in posets. Where genuinely higher structure is visible—the homotopy types of equilibrium sets, which by Datta's theorem are arbitrary; the simplicial models of many-agent knowledge; the ∞-toposes in which smooth and probabilistic games ought to live—this book has pointed and not built. A reader who came for $(\infty, 1)$-categories has been given the 1-truncation of a programme.

19.3 What the language does not do

It does not solve games. Composition of open games composes the equilibrium check, not the search (Chapter 2). Computing equilibria is hard for reasons no

notation removes (Chapter 6).

It does not select equilibria. The invariance principle says what a selection must respect, and thereby that many proposed selections are not well defined; it does not supply one. The structure theorem says there is no continuous one.

It does not adjudicate modelling choices. Whether the order of moves is part of the game, whether histories may serve as frames, whether types are hierarchies or something finer: the type-theoretic discipline forces these choices to be declared. It does not make them.

It does not replace the examples. Every structural statement in this book was extracted from models that economists built for particular purposes, and the worked examples carry content that the general statements do not.

19.4 Problems

The following seem to us well posed, open, and of a kind that the language of this book makes natural.

Open Problem 19.1 (A univalent type of extensive games)**.** [programme] Section 3.4 defined $\mathsf{EFG}_{\mathrm{T}}$ as a quotient. Find a notion of structure, intrinsic rather than presented, whose structure-preserving equivalences are exactly the Thompson equivalences, so that invariance holds by the structure identity principle rather than by decree. The set-theoretic reformulations of the extensive form in terms of plays and choices [AFR16] are a natural starting point; the reduced strategic form is the trivial answer, and the problem is to find one that retains enough dynamic structure to state sequential rationality.

Open Problem 19.2 (Perfection as an operation on open games)**.** [programme] The modality $\Box$ of Proposition 8.4 was defined for closed extensive games. Define it on the category of open games, as an operation on equilibrium predicates, so that $\Box(\mathcal{H} \circ \mathcal{G})$ is computed from $\Box\mathcal{H}$ and $\Box\mathcal{G}$, recovering [Hed16, GKLF18] as special cases, and extend it to information sets with beliefs so as to obtain a compositional sequential equilibrium for the Bayesian open games of [BHZ19].

Open Problem 19.3 (The exit-path category of the space of games)**.** [programme] The equilibrium correspondence is semialgebraic, so the space $\mathcal{U}$ of games with given action sets admits a stratification over each stratum of which the projection $\mathrm{pr}\colon \mathcal{E} \to \mathcal{U}$ of Section 6.3 is a fibre bundle. The fibres then assemble into a functor from the exit-path ∞-category of the stratification to spaces. Describe this functor for the smallest nontrivial cases. Characterize the essential components and the stable sets of [KM86, GW05] as subfunctors with a universal property. Remark 6.9 suggests that a refinement in the proper sense is a homotopy-invariant of this functor; make that a definition and compare it with the axiomatic characterizations.

Open Problem 19.4 (Higher cohomology of information). [programme] For $n \geq 3$ agents the simplicial model $\mathcal{C}$ of Chapter 4 has higher homotopy that the graph B lacks. In distributed computing its connectivity obstructs the solvability of tasks. Find the game-theoretic meaning of $H^k(\mathcal{C})$ for $k \geq 2$. A candidate: the marginal problem [Vor62], of which Theorem 5.7 is the pairwise case, asks when beliefs specified on a family of coalitions are the marginals of a single joint belief, and its obstructions are known to be cohomological in the closely related setting of contextuality [AB11]. Which mediated correlations among coalitions are implementable by local devices is a question of this kind.

Open Problem 19.5 (Graded modalities). [programme] Three approximations in this book are indexed by a real parameter: common p-belief (Chapter 10), ϵ-equilibrium, and discounting by δ. The first is a graded comonad, the second has a compositional metric formulation [Gha25], and the third is a quantitative form of the later modality (Theorem 9.6). Give a single quantale-graded modal type theory in which all three are instances, and in which the continuity theorem of Monderer and Samet and the limit theorems for finite-horizon approximations [FL83] are theorems about grades.

Open Problem 19.6 (Beyond quasi-linearity). [programme] Theorem 14.3 and Proposition 14.6 use transferable utility in an essential way: the base of enrichment is the additive reals. Nöldeke and Samuelson [NS18] extend the duality between tariffs and utility profiles to preferences that are not quasi-linear, by means of Galois connections. Identify the enriching structure for which their duality is the nucleus of a profunctor, and state implementability and payoff equivalence as conditions on the resulting enriched category.

Open Problem 19.7 (A synthetic existence theorem). [programme] Nash's theorem requires Brouwer's, which holds synthetically in a cohesive type theory [Shu18], while mixed strategies require probability, which is axiomatized by Markov categories [Fri20]. Find a setting with both: a category or type theory in which the mixed extension is a free construction, strategy spaces carry intrinsic topology, and the truncated existence statement of Remark 6.9 is a theorem. Determine which classes of games admit an untruncated one: supermodular games, by Theorem 7.5; two-player zero-sum games, by linear programming; potential games, by maximization.

Open Problem 19.8 (Effective game theory). [programme] Theorem 9.15 shows that best responses to computable strategies need not be computable. Develop the theory of repeated games internally to a realizability topos, where every function is computable by construction. Is there a folk theorem there? The internal terminal coalgebra of the strategy functor consists of computable strategies, and the internal equilibrium condition quantifies only over those; the

classical counterexamples suggest that the internal and external equilibrium sets differ, and the difference would be a precise measure of what bounded rationality costs.

19.5 A last remark on method

The source of this book's subject matter describes its own method as modelling with the minimum of structure: find the simplest set of assumptions that generates the phenomenon, and stop. That is a principle of economy about what goes into a model. The discipline followed here adds a second economy, about what one may get out: only what is invariant under the equivalences of the model's type. The two economies are complementary. The first prevents explanation by excess detail. The second prevents explanation by accident of presentation. Game theory has always practised both, the second by the judgment of its practitioners. The claim of this book is modest: that the judgment can be written down.

Appendix A

Categorical Toolkit

This appendix fixes notation and collects, without proofs, the categorical notions used in the text. It is a list of what is needed, in the order in which it is needed, and not an introduction to any of the subjects named; references are given for each.

A.1 Categories, adjunctions, (co)algebras

We assume categories, functors, natural transformations, limits and colimits, and adjunctions $L \dashv R$ [FS19]. A ***monad*** on $\mathcal{C}$ is an endofunctor T with unit $\eta\colon \mathrm{id} \to T$ and multiplication $m\colon TT \to T$ satisfying the monoid laws; its ***Kleisli category*** $\mathrm{Kl}(T)$ has the objects of $\mathcal{C}$ and morphisms $X \to TY$; its ***algebras*** are maps $TA \to A$ compatible with η and m. Every adjunction $L \dashv R$ gives a monad RL and a comonad LR.

For an endofunctor F, an ***F-algebra*** is a map $F(A) \to A$ and an ***F-coalgebra*** a map $C \to F(C)$. An initial algebra μF supports definition by recursion (the unique map to any other algebra is a *fold*) and proof by induction; a terminal coalgebra νF supports corecursion and coinduction [Rutoo]. *Lambek's lemma*: the structure map of an initial algebra or terminal coalgebra is an isomorphism. When $\mathcal{C}$ has a terminal object and F preserves limits of ω^{op}-chains, νF is the limit of $1 \leftarrow F1 \leftarrow F^2 1 \leftarrow \cdots$. Two states of coalgebras are ***bisimilar*** (behaviourally equivalent) when they have the same image in νF.

Functor	μF **or** νF	**Used for**
$X \mapsto O + \sum_i \sum_A X^A$ on Set	μ: finite game trees	backward induction (Ch. 8)
$X \mapsto A_i \times X^A$ on Set	ν: strategies $A^* \to A_i$	automata (Ch. 9)
$X \mapsto \mathcal{G}(\Theta \times X)$ on Meas	ν: universal types	hierarchies (Ch. 5)
monotone f on a complete lattice	μf, νf	Tarski (Ch. 7)

A.2 Monoidal categories, lenses, optics

A ***symmetric monoidal category*** $(\mathcal{C}, \otimes, I)$ has a tensor product that is associative, unital and symmetric up to coherent isomorphism. Its morphisms are drawn as *string diagrams*: boxes with input and output wires, composed by joining wires and tensored by juxtaposition; an equation holds in all symmetric monoidal categories iff the diagrams are isotopic. A *state* is a morphism $I \to X$; a *costate* is a morphism $X \to I$.

The category Lens has pairs of sets as objects and, as morphisms $(X, S) \to (Y, R)$, pairs of functions $X \to Y$ and $X \times R \to S$ (Definition 2.1). For a symmetric monoidal $\mathcal{C}$, the category $\mathsf{Optic}_{\mathcal{C}}$ has hom-sets $\int^M \mathcal{C}(X, M \otimes Y) \times \mathcal{C}(M \otimes R, S)$ [Ril18, CEGLMPR20]; for cartesian $\mathcal{C}$ it is Lens. The bicategory $\mathsf{Para}(\mathcal{C})$ has morphisms $X \to Y$ given by pairs $(\Pi,\ \Pi \otimes X \to Y)$ and reparametrizations as 2-cells [CGHR22]. An open game is a family of lenses indexed by strategies with an equilibrium predicate (Definition 2.2); open games form a symmetric monoidal category (Theorem 2.4).

A.3 Enriched categories

Let $(V, \leq, \otimes, k)$ be a commutative quantale: a complete lattice with a monoid structure that preserves joins in each variable. A ***V-category*** is a set C with $C(a, b) \in V$ such that $k \leq C(a, a)$ and $C(a, b) \otimes C(b, c) \leq C(a, c)$; a ***$V$-functor*** is a map f with $C(a, b) \leq D(fa, fb)$. Three cases are used.

- $V = \mathbf{2}$ with $\wedge$: preorders and monotone maps.
- $V = ([0, \infty], \geq, +)$: Lawvere metric spaces and nonexpansive maps [Law73]. Note the reversal of order: the "composition" inequality is the triangle inequality.
- $V = \overline{\mathbb{R}} = ([-\infty, \infty], \geq, +)$: the same with negative distances allowed, used for the incentive category of Section 14.2. A $\overline{\mathbb{R}}$-category has $d(a, a) \in \{0, -\infty\}$.

V is itself a V-category with $V(a, b)$ the internal hom (truncated subtraction $b - a$ in the metric cases). A ***copresheaf*** is a V-functor $C \to V$; the ***Yoneda lemma*** identifies $C(a, -)$ as the copresheaf freely generated by an element at a. The ***free V-category*** on a V-graph has hom-values given by joins over paths of tensor products of edge values; in the metric cases, shortest-path distances. A ***profunctor*** $\Phi \colon C \otimes D \to V$ induces an adjunction between presheaves on one side and copresheaves on the other, whose fixed points form its ***nucleus***; for the pairing of a vector space with its dual, valued in $\overline{\mathbb{R}}$, this is the Legendre–Fenchel transform [Wil15].

A.4 Toposes and modalities

A ***presheaf topos*** is a category $\mathrm{Set}^{\mathcal{C}^{\mathrm{op}}}$. We use two: presheaves on a discrete set Ω, which is the slice Set/Ω of families of sets indexed by states, and presheaves on the ordinal ω, the ***topos of trees*** $\widehat{\omega}$ [BMSS12]. In any topos $\mathcal{E}$, each object X has a Heyting algebra $\mathrm{Sub}(X)$ of subobjects, and each map $p\colon X \to Y$ induces adjoints $\exists_p \dashv p^* \dashv \forall_p$ between $\mathrm{Sub}(X)$ and $\mathrm{Sub}(Y)$; on families of objects the corresponding adjoints are dependent sum and dependent product, $\sum_p \dashv p^* \dashv \prod_p$ [MM92]. Composites of these are the modal operators of the text:

Operator	Formula	Kind
knowledge of an agent with observation map p	$p^*\forall_p$	comonad, S5
possibility	$p^*\exists_p$	monad, S5
common knowledge	$q^*\forall_q$, q the pushout of the p_i	comonad, S5
distributed knowledge	$p_\vee^*\forall_{p_\vee}$, $p_\vee$ the joint observation	comonad, S5
subgame perfection	$\forall_p e^*$ over later subgames	comonad, S4
later (in $\widehat{\omega}$)	$(\blacktriangleright X)(n+1) = X(n)$	guard

In $\widehat{\omega}$, a map $f\colon X \to X$ is *contractive* if it factors through the canonical $X \to \blacktriangleright X$; contractive maps have unique fixed points, and locally contractive functors have unique fixed points up to isomorphism, so that for them initial algebras and terminal coalgebras coincide [BMSS12].

A.5 Homotopy types and type theory

A ***homotopy type*** (or space, or ∞-groupoid) is considered up to weak equivalence; sets are the discrete ones. The ∞-category $\mathcal{S}$ of spaces has homotopy colimits. The homotopy colimit of a diagram of sets $D\colon J \to \mathrm{Set}$ is the nerve of its category of elements [Tho79]; in particular the homotopy pushout of $A \leftarrow C \to B$ is the graph with vertices $A \sqcup B$ and edges C. The functor $\pi_0\colon \mathcal{S} \to \mathrm{Set}$ is left adjoint to the inclusion and so carries homotopy colimits to colimits. For a graph B and an abelian group G, $H^0(B;G)$ is the group of locally constant G-valued functions on vertices and $H^1(B;G)$ is the group of G-valued functions on edges modulo coboundaries $(\delta\varphi)(e) = \varphi(\mathrm{head}) - \varphi(\mathrm{tail})$; a 1-cochain is a coboundary iff it sums to zero around every cycle, and $H^1(B;\mathbb{R}) \cong \mathbb{R}^{b_1}$ with $b_1 = |E| - |V| + |\pi_0|$.

In ***homotopy type theory*** [HoTT13] types are homotopy types. We use: dependent sums $\sum_{x:A} B(x)$ and products $\prod_{x:A} B(x)$; the identity type $x =_A y$, whose elements are paths; transport of a dependent element along a path;

propositional truncation $\|A\|$, the proposition that A is inhabited; set-quotients and higher inductive types. The ***univalence axiom*** identifies the identity type of the universe with equivalence of types, and has as a consequence the ***structure identity principle***: for a type of structured types, identifications are structure-preserving equivalences [ANST21]. Every function respects identifications of its domain; this is Principle 1.5. A ***modality*** in the sense of [RSS20] is an idempotent monad on the universe with a dependent elimination rule; comonadic modalities such as $\flat$ require a context discipline [Shu18, Cor20].

A.6 Probability

Δ is the monad of finitely supported distributions on Set; $\mathcal{G}$ is the Giry monad on measurable spaces [Gir82]. Both are *commutative* (monoidal) monads: there is a natural $\nabla\colon TX \times TY \to T(X \times Y)$ forming product measures, compatible with the monad structure, and it is not an isomorphism. A ***Markov category*** [Fri20] is a symmetric monoidal category with commutative comonoid structures (copy, delete) on all objects, deletion being natural. The Kleisli categories FinStoch and Stoch are examples. A morphism is *deterministic* if it commutes with copy. Given a state $\pi\colon I \to X$ and $f\colon X \to Y$, a ***Bayesian inverse*** is $f^\dagger_\pi\colon Y \to X$ such that the joint states on $X \otimes Y$ formed from (π, f) and from $(f\pi, f^\dagger_\pi)$ coincide [CJ19]. For f, g out of the same object, g is a ***garbling*** of f if $g = hf$; this preorder is the Blackwell order [Bla51, FGPR23]. Algebras of Δ are convex spaces, and ΔX is the free one.

A.7 Four fixed-point theorems

Theorem	Statement	Uses in this book
Tarski	a monotone map on a complete lattice has a complete lattice of fixed points	rationalizability, common knowledge, supermodular games, self-generation, lemons
Banach	a contraction of a complete metric space has a unique fixed point	dynamic programming, alternating-offers bargaining
Brouwer, Kakutani	a continuous map (closed convex-valued correspondence) on a compact convex set has a fixed point	Nash's theorem, existence of sequential equilibrium
Lawvere	if $A \to B^A$ is point-surjective then every $B \to B$ has a fixed point	nonexistence of universal possibility structures; limits of self-prediction

The first two are the enriched-categorical pair of Section 7.7 and are constructive in content: both come with an iteration. The third is topological and is not constructive. The fourth, used in contrapositive form, is the diagonal argument [Law69].

Appendix B

Concordance with the Source

This book covers the territory of Rasmusen's *Games and Information* (4th edition) [Raso6] in a different order and a different language. The table maps the chapters of that book to the places where their subject matter is treated here, with the structural reading proposed. Topics are described in our own words; the source's examples have been replaced by others throughout, and none of its problems is reproduced.

Src.	Subject matter	Here	Reading
1	players, actions, payoffs; dominance and iterated dominance; Nash equilibrium; focal points	1, 2	types of games and their identity; dominance as confluent rewriting; rationalizability as a greatest fixed point; focal points as symmetry-breaking structure; composition
2	strategic and extensive forms; information sets; perfect, certain, symmetric and complete information; the Harsanyi transformation; Bayesian games	3, 4, 5	trees as presheaves; information as a quotient; perfect recall as internality; knowledge as base change; incomplete information as a dependent sum; types as a terminal coalgebra
3	mixed strategies; chicken and the war of attrition; correlated strategies; many-player mixing; auditing; Cournot and Bertrand; complements and substitutes; existence	6, 7	free convex algebras; the non-invertible map ∇; degree of the equilibrium correspondence; lattices and Tarski
4	subgame perfection; entry deterrence; credible threats and sunk costs; the open-set problem; Pareto perfection	8	folds over an initial algebra; necessity modality; commitment as a morphism; searchability

Src.	Subject matter	Here	Reading
5	finitely and infinitely repeated games; the chain-store paradox; minimax and the folk theorem; reputation and product quality; Markov equilibria; evolutionary equilibrium	9; 12	induction versus coinduction; automata as coalgebras; self-generation; Markov as bisimulation invariance. (The quality-assuring price premium is the efficiency wage of Chapter 12 with the roles relabelled.)
6	perfect Bayesian equilibrium and its refinement; the role of common knowledge; incomplete information in the repeated Prisoner's Dilemma; tournaments of strategies; reputation in credit	10; 9	sections of belief bundles; conditional systems as a blow-up; discontinuity in the type space; graded modalities
7	classification of asymmetric-information models; the principal–agent problem; incentive, participation and competition constraints; optimal contracts	11	five composites of the same boxes; contracts as costates; informativeness as factorization of kernels
8	efficiency wages; tournaments; institutions; renegotiation; insurance under moral hazard; teams	12	coinductive incentives; sufficient statistics; open versus closed systems; idempotent operators on contracts
9	adverse selection in employment, used cars and insurance; market microstructure	13	monotone operators and their lattices of fixed points; Bayesian lenses in equilibrium
10	the revelation principle; hidden knowledge after contracting; price discrimination; regulation and procurement; the Groves mechanism	14	initiality and pullback-stability; the incentive metric; payments as copresheaves; tariffs and utilities as a nucleus
11	signalling and screening in education; multiple signals; signal jamming and limit pricing	15	kernels and their Bayesian inverses; order of composition; forward versus backward passes
12	the bargaining problem; the Nash solution; alternating offers with finite and infinite horizons; incomplete information; efficient mechanisms for bargaining	16	equivariance and naturality; unique guarded fixed points; obstruction by potentials

Src.	Subject matter	Here	Reading
13	auction formats; private values; comparison of rules; risk; common values and the winner's curse; information	17	two kinds of equivalence; revenue equivalence as a groupoid condition; affiliation as log-supermodularity
14	quantity and price competition; location; supermodular comparative statics; durable monopoly	18, 7	commitment; discontinuity and its repairs; mechanism design without commitment
App.	notation; supermodularity; fixed-point theorems; discounting; risk	A, 7	the four fixed-point theorems

Additions and omissions

Added. Material with no counterpart in the source includes the calculus of open games (Chapter 2); the Thompson transformations and the invariance debate (Chapter 3); the universal type space and its limits (Chapter 5); the topology of the equilibrium correspondence and the complexity of equilibrium (Chapter 6); self-generation and computability (Chapter 9); conditional probability systems, higher-order beliefs and global games (Chapter 10); cheap talk and Bayesian persuasion (Chapter 15); and the graph-theoretic treatment of incentive compatibility (Chapter 14). The linear-contract and multitask models of Chapters 11 and 12 are two of the additions that the source's own preface says a further edition would make.

Omitted or compressed. The source's teaching examples and their numbered variants; its treatment of particular applications in law, finance and regulation at the level of institutional detail; its model of two simultaneous signals in the pricing of new share issues; and all end-of-chapter problems. A reader who wants a large stock of worked models built with a consistent recipe should read the source; that is what it is for.

Appendix C

Computational Checks

The script `verify.py` that accompanies this book recomputes every numerical example in the text and tests a number of structural claims on explicit or randomly generated finite instances. It uses only the Python standard library, fixes its random seed, and runs in about a second; `python3 verify.py` prints one line per check and exits with status zero if and only if all pass. Version 1.1 performs 113 checks, all of which pass under Python 3.8 or later; the last 29 recompute the numerical answers of Appendix D.

A check of this kind is evidence about the examples and about the absence of slips in the finite statements. It is not a proof of any theorem, and the theorems tagged classical or literature rest on the references cited for them. What the script does establish is the following.

Ch.	What is checked
1	The transformed Prisoner's Dilemma has the same equilibrium and reversed efficiency (Proposition 1.6). Both elimination orders of Example 1.10 and the simultaneous procedure. Order independence of strict dominance on 200 random games. The invariant profile in pure coordination.
2	A direct implementation of Definitions 2.2 and 2.3: the closed composite of decisions and a payoff has exactly the pure Nash equilibria (Proposition 2.6); the sequential composite of Example 2.7 admits the non-credible equilibrium. The product of selection functions agrees with brute-force backward induction on 100 random three-stage trees (Theorem 2.9).
3	The equilibria of Example 3.9 in both trees; the imperfect-recall example.
4	On 300 random information structures: the S5 laws and the adjunction of Proposition 4.3; equality of iterated mutual knowledge with knowledge relative to the pushout (Theorem 4.6); the formula for $b_1(B)$ against a direct computation of the cycle rank (Proposition 4.9). The graphs of Figure 4.1, the product structure, and the three-agent example.

Ch.	What is checked
5	On 400 random structures: when $b_1(B) = 0$ a common prior is constructed for random posteriors and its conditionals verified; when $b_1(B) > 0$ random posteriors fail to glue, and posteriors derived from a prior always glue (Theorem 5.7); the Hellman–Samet count. The cycle equation of Example 5.8. The tower property used in Theorem 5.10. Examples 5.2 and 5.6.
6	Examples 6.4 and 6.12; the fixed-point indices $(+1, +1, -1)$ of Example 6.7, computed as winding numbers of the displacement field of a Nash map around each equilibrium.
7–8	Cournot, Bertrand, Stackelberg and the convergence of best-response iteration from the top and bottom of the lattice (Examples 7.6, 7.7 and 8.6).
9	The grim-trigger threshold; the operator B of Theorem 9.8 on the sets of Example 9.9; Hawk–Dove.
10	Examples 10.2 and 10.8; the count of strata in Proposition 10.6; the knowledge graph of Example 10.10 is a connected tree.
11–12	Examples 11.1 and 11.4; the multitask formula of Proposition 12.2 against numerical maximization; the efficiency-wage condition.
13	Examples 13.1 and 13.2; $\alpha^* \approx 0.2899$ and $\lambda^* \approx 0.2817$ in Example 13.5; Examples 13.6 and 13.7.
14	The incentive category by an all-pairs shortest-path computation: Example 14.4; a three-type example; detection of negative cycles for a non-monotone rule; incentive compatibility of the representable payment rule on 50 random monotone rules; and the failure of revenue equivalence on a grid of $n+1$ types by exactly $d_f(0,1) + d_f(1,0) = 1/n$, vanishing in the limit (Theorem 14.3 and Corollary 14.5). Examples 14.7 and 14.8; the pivot example.
15	Example 15.1; the cutoffs and boundary indifference of the cheap-talk example; Examples 15.8 and 15.9.
16	Example 16.3; finite-horizon shares and the Rubinstein fixed point; the gains from trade $\frac{1}{6}$ and $\frac{9}{64}$ and the buyer's best response in Example 16.7.
17	The first-price equilibrium as a best response; revenue equivalence for uniform values; the optimal reserve; the all-pay and third-price bids; the wallet game.
18	The collusion threshold; Examples 18.2 and 18.3; the capacity exercise.

Ch.	What is checked
D	The numerical answers to Exercises 1.1, 1.3, 1.4, 5.1, 5.5, 6.2, 6.4, 8.3, 9.2, 10.1, 10.4, 11.3, 12.2, 12.3, 13.1–13.3, 14.5, 15.1–15.4, 16.1–16.4, 18.1, 18.3, 18.5 and 18.6. Where a solution asserts an optimum (the best correlated equilibrium of 6.4, the contract of 11.3, the bundle of 14.5), the script checks the stated point against a grid or a random search over the feasible set as well as the first-order conditions.

The implementation of open games in the script is deliberately naive—strategy sets are enumerated and equilibrium predicates evaluated by brute force—and is intended as an executable form of the definitions rather than as software. For serious use see the engine cited in Chapter 2.

Appendix D

Solutions to the Exercises

Solutions are given for every exercise, numbered as in the chapters. Where an exercise asks for a computation, the numbers below are recomputed by the script of Appendix C.

Chapter 1

1.1. The obedience constraint for player i involves payoffs only through the differences $u_i(a_i, a_{-i}) - u_i(a_i', a_{-i})$. Under $u_i \mapsto \alpha_i u_i + \beta_i(a_{-i})$ each difference is multiplied by $\alpha_i > 0$, so the system of inequalities is unchanged. For efficiency, take the Battle of the Sexes of Example 1.13 and add 5 to the row player's payoff whenever the column player chooses F. The profile (F, F) now pays $(6, 2)$ and Pareto-dominates (O, O), which pays $(2, 1)$ and was efficient before; both remain Nash equilibria.

1.2. An automorphism is a permutation of the players together with bijections of action sets preserving payoffs. Payoffs depend only on whether the actions agree, so the two bijections must be the same permutation $\pi \in S_m$, with or without exchanging the players: $\mathrm{Aut}(G) \cong S_m \times \mathbb{Z}/2$. In the framed game π must fix action 1, and the group is $S_{m-1} \times \mathbb{Z}/2$. An invariant profile is symmetric, with each player putting weight x on action 1 and $(1-x)/(m-1)$ on each other action. A player's payoff from an action is the probability the other puts on it, so the support must consist of actions of maximal probability. This leaves three invariant equilibria: $x = 1$ (both play the distinguished action, payoff 1); $x = 0$ (uniform on the undistinguished actions, payoff $\frac{1}{m-1}$); and $x = \frac{1}{m}$ (uniform on all, payoff $\frac{1}{m}$). The first is now selectable.

1.3. Let the row player's payoffs be $T = (3, 0)$, $M = (1, 1)$, $B = (0, 3)$ against columns L, R. M is not dominated by T (against R) nor by B (against L), but $\frac{1}{2}T + \frac{1}{2}B$ pays $(1.5, 1.5)$. If the column player chooses L with probability q, then T pays $3q$ and B pays $3(1-q)$, and $\max\{3q, 3(1-q)\} \geq 1.5 > 1$: M is never a best response.

1.4. Her selection function is $\varepsilon(k) = \arg\min_x |x - \frac{2}{3}k(x)|$. It is not of the form $\arg\max(v \circ k)$ for any $v \colon \mathbb{R} \to \mathbb{R}$: on a constant context k, $\arg\max(v \circ k)$ is all of X, while $\varepsilon(k)$ is a single point depending on the constant. Her objective

depends on her action and on the outcome jointly. For the game, write the signed distance as $x_i(1-\frac{2}{3n})-\frac{2}{3n}\sum_{j\neq i}x_j$, increasing in x_i. If the largest choice in a profile is $M \geq 2$, a player choosing M is above her target by at least $M/3 \geq \frac{2}{3}$, and lowering her choice by one reduces the signed distance by less than 1 while keeping its absolute value smaller; so she deviates. Hence everyone chooses 1, and one checks directly that this is an equilibrium: at 1 the distance is $\frac{1}{3}$, at 2 it is $\frac{4}{3}-\frac{2}{3n} \geq 1$, and it grows with the choice.

1.5. If $R \subseteq R'$ then $\Delta(R_{-i}) \subseteq \Delta(R'_{-i})$, so every action justified by a conjecture on R_{-i} is justified by one on R'_{-i}: $\Phi(R) \subseteq \Phi(R')$. If σ is a (mixed) Nash equilibrium let R be the product of the supports. Each action in the support of σ_i is a best response to $\sigma_{-i} \in \Delta(R_{-i})$, so $R \subseteq \Phi(R)$, and by coinduction $R \subseteq \nu\Phi$.

1.6. At each round the set of weakly dominated actions of each player in the current restriction is determined, and it is never all of a player's actions (an action maximizing the sum of payoffs over the opponents' profiles is undominated). So the procedure involves no choices and terminates. Weak dominance depends only on the signs of own-payoff differences against each opponent profile, which incentive equivalence preserves. On Example 1.10: round one deletes T and B and neither column; round two deletes nothing. The output $\{M\}\times\{L,R\}$ differs from both sequential outputs, $\{M,B\}\times\{R\}$ and $\{T,M\}\times\{L\}$. What has been bought is a function. The price is that the function is a convention: it is not stable under first removing a dominated action by hand, so the dependence on order has been hidden in the definition rather than removed.

Chapter 2

2.1. For lenses (v_j, c_j), $j = 1,2,3$, both bracketings have view $v_3v_2v_1$ and update $(x,p) \mapsto c_1\big(x, c_2(v_1x, c_3(v_2v_1x, p))\big)$. A lens $I \to (X,S)$ is a view $1 \to X$, an element, with a forced update $1 \times S \to 1$. A lens $(Y,R) \to I$ has a forced view and an update $Y \times 1 \to R$.

2.2. Strategies are pairs (a_1, s) with $s\colon \{Out, In\} \to \{A,F\}$. The predicate requires $a_1 \in \arg\max_a u_1(a, s(a))$ and $s(a_1) \in \arg\max u_2(a_1, -)$. If $s(In) = A$ then $a_1 = In$ and the second condition holds; $s(Out)$ is free. If $s(In) = F$ then $a_1 = Out$, and after Out player 2 is indifferent, so again $s(Out)$ is free. There are four equilibria and two outcomes. The two with $s(In) = F$ are not subgame perfect.

2.3. Write $c(x,y) = \gamma(q(x,y,-))$ and $b(x) = \delta\big(y \mapsto q(x,y,c(x,y))\big)$ and $a = \varepsilon\big(x \mapsto q(x,b(x),c(x,b(x)))\big)$. Unwinding either bracketing gives

$(a, b(a), c(a, b(a)))$. For the second identity, with $b(x) = \delta(q(x, -))$ we have $\overline{\delta}(q(x, -)) = q(x, b(x))$, and $\overline{\varepsilon}$ applied to $x \mapsto q(x, b(x))$ is $q(a, b(a)) = \overline{\varepsilon \otimes \delta}(q)$. For max and min this reads $\max_x \min_y q(x, y)$: the product of the selection functions arg max and arg min plays a two-move zero-sum game optimally, and the product of quantifiers computes its value.

2.4. $\int^M \mathsf{Set}(X, M \times Y) \times \mathsf{Set}(M \times R, S) \cong \mathsf{Set}(X, Y) \times \int^M \mathsf{Set}(X, M) \times \mathsf{Set}(M \times R, S) \cong \mathsf{Set}(X, Y) \times \mathsf{Set}(X \times R, S)$, the last step by the co-Yoneda lemma applied to the contravariant functor $M \mapsto \mathsf{Set}(M \times R, S)$. Concretely, a representative $(M, (\ell_1, \ell_2), r)$ is equivalent, by sliding $\ell_1 \colon X \to M$, to $(X, (\mathrm{id}, \ell_2), r \circ (\ell_1 \times \mathrm{id}))$, which is the lens $(\ell_2, r \circ (\ell_1 \times \mathrm{id}))$.

2.5. The only context for a game $I \to I$ is the trivial one, so by Definition 2.3 the predicate of $\mathcal{G} \otimes \mathcal{G}'$ is $E_{\mathcal{G}} \times E_{\mathcal{G}'}$. For the second part let $Y = \{a, b\}$, $\mathcal{G} = \mathcal{D}_{1,Y}$, and $\mathcal{H} = u \circ (\mathrm{id}_Y \otimes \mathcal{D}_{1,Y})$ with u paying $(1, 1)$ if the two choices agree and $(0, 0)$ otherwise. Then (a, a) is an equilibrium of $\mathcal{H} \circ \mathcal{G}$. The context that τ induces on $\mathcal{G}$ is $y \mapsto 1_{y=\tau}$, and a is an equilibrium of $\mathcal{G}$ only in the context induced by $\tau = a$.

Chapter 3

3.1. A map $1 \to T$ in $\widehat{\omega}$ is a family $t_n \in T(n)$ compatible with restriction, that is, an infinite branch; so $\Gamma(T) = \lim_n T(n)$. When every node has a child every node lies on a branch. For $h \in T(m)$ let $T^{+m}(n) = T(m + n)$; restriction gives a map from T^{+m} to the constant presheaf $T(m)$, and T_h is its pullback along $h \colon 1 \to T(m)$, since limits of presheaves are computed levelwise.

3.2. *Symmetric and imperfect:* player 1 chooses L or R; player 2, not observing this, chooses l or r. Player 2's two nodes form one information set, so information is imperfect. At that set the mover knows less than player 1, and at the root nobody knows anything; at the end all know the play. No mover has an informational advantage. *Certain and asymmetric:* Nature chooses one of two types at the root, observed by player 1 only; player 1 moves, then player 2. Nature does not move after a player, so information is certain. At player 1's nodes her information set is a singleton and does not contain player 2's, so information is asymmetric (and incomplete).

3.3. Nature chooses a or b; player 1 observes it and moves at g_a or g_b, which lie in different information sets; she moves again at h_a (after g_a) or h_b (after g_b), and these lie in one information set. Her personal forest has g_a, g_b at depth 0 and h_a, h_b at depth 1, and the levelwise quotient has two classes at depth 0 and one at depth 1. The parent of $[h_a] = [h_b]$ would have to be both $[g_a]$ and $[g_b]$: the square expressing naturality of the quotient with respect to restriction

does not commute. In the "forgetting" example of the text the parents agree and it is the edge labelling that fails to descend.

3.4. Let player 2 choose L with probability q. *Out* pays 2, T pays $3q$, B pays $1-q$. So *Out* is a best response iff $q \le \frac{2}{3}$, and (T, L) is an equilibrium. Player 1 cannot mix *Out* with T: indifference needs $q = \frac{2}{3}$, but any weight on T and none on B makes L strictly better for player 2. The mixed equilibrium of the subgame pays player 1 only $\frac{3}{4}$, so it is not an equilibrium of the whole game. Hence the Nash equilibria are (T, L) and (Out, q), $q \le \frac{2}{3}$. In the subgame: (T, L), (B, R), and the mixed profile with $p = \frac{3}{4}$ (from $p = 3(1-p)$) and $q = \frac{1}{4}$ (from $3q = 1 - q$).

3.5. No. After the interchange the subgame following *In* is rooted at a node of player 2, but it contains the same terminal nodes and has the same strategic form. The subgame perfect equilibria are as before: $(In, T; L)$ and *Out* followed by $q \in \{0, \frac{1}{4}\}$. Interchange preserved the collection of subgames; coalescing did not.

3.6. The variables are $r(\varsigma)$ for the empty sequence and for each pair (I, a): $1 + \sum_I |A(I)|$ of them. The equations $r(\emptyset) = 1$ and $\sum_a r(\varsigma_I a) = r(\varsigma_I)$, one for each I, are independent, each introducing variables not constrained by the earlier ones, and with $r \ge 0$ they force $r \le 1$. A fully mixed behavioural strategy gives a point with all coordinates positive, so the set is a polytope of dimension $\sum_I (|A(I)| - 1)$. The mixed-strategy simplex has dimension $\prod_I |A(I)| - 1$: with ten binary information sets, 10 against 1023.

Chapter 4

4.1. All of them. $K = p^* \forall_p$ is a comonad on a poset, hence idempotent, whatever p is; (iii) and (iv) use only the adjunctions; (v) uses only that p^* preserves complements. K depends on p only through its kernel partition of Ω. What fails is a step of the proof given, $\forall_p p^* = \mathrm{id}$: without surjectivity $\forall_p p^* D = D \cup (T \setminus \operatorname{im} p)$, statements about unrealized types being vacuously known.

4.2. In the structure on the left of Figure 4.1 with $E = \{a, b, c, d\}$: $K_1 K_2 E = \{a, b\}$ but $K_2 K_1 E = \{a, b, c\}$ (see the next solution). For the inclusion: CE is self-evident, so $CE = K(CE) \subseteq K_i(CE) \subseteq CE$ and $K_i CE = CE$ for each i. By monotonicity $CE = K_{i_1} \cdots K_{i_m} CE \subseteq K_{i_1} \cdots K_{i_m} E$.

4.3. $K_1 E = \{a, b, c, d\}$, $K_2 E = \{a, b, c\}$, $K_1 K_2 E = \{a, b\}$, and $CE = \emptyset$ since the meet is trivial and $E \neq \Omega$. Everyone knows E on $K_1 E \cap K_2 E = \{a, b, c\}$, and E is common knowledge nowhere. The iterates of "everyone knows" shrink one state at a time: $\{a, b, c\}$, $\{a, b\}$, $\{a\}$, $\emptyset$.

4.4. Left: $(2-1)\cdot 5-(3+3)+1=0$. Right: $4-(2+2)+1=1$. In the product structure the state (s,t) is an edge joining s and t, so B is $K_{m,n}$, connected, with $b_1 = mn-(m+n)+1=(m-1)(n-1)$. If p_1 is the identity, each state is an edge from its own singleton cell, a vertex of degree one, to a cell of agent 2: a disjoint union of stars, and $b_1 = |\Omega|-(|\Omega|+|T_2|)+|T_2|=0$.

4.5. Distributed knowledge of E at ω means that the intersection of the agents' cells at ω lies in E. That intersection is the fibre of the pairing $\langle p_i\rangle$ through ω, a cell of the join; so $D = p_\vee^*\forall_{p_\vee}$. $CE\subseteq K_iE$ is the case $m=1$ of Exercise 4.2; $K_iE\subseteq DE$ because the join's cell is contained in agent i's; $DE\subseteq E$ is the truth axiom.

4.6. B has vertices ω_1,ω_2 and t_1,t_2,t_3,t_3', with each state joined to t_1 and t_2 and to its own type of agent 3: six vertices, six edges, one component, so $b_1=1$ and $\pi_1(B)\cong\mathbb{Z}$, generated by the square $\omega_1t_1\omega_2t_2$. $\mathcal{C}$ is the two triangles $\{t_1,t_2,t_3\}$ and $\{t_1,t_2,t_3'\}$ glued along the edge $\{t_1,t_2\}$; it is contractible.

Chapter 5

5.1. If the other player uses the cutoff c^* she contributes with probability $c^*/\bar{c}$, so $c^*=1-c^*/\bar{c}$ and $c^*=\frac{\bar{c}}{\bar{c}+1}$. The good is provided with probability $1-\left(\frac{\bar{c}}{\bar{c}+1}\right)^2$: this is $\frac{5}{9}$ at $\bar{c}=2$, tends to 0 as $\bar{c}\to\infty$, and tends to $\frac{3}{4}$ as $\bar{c}\to 1$. Even when every cost is below the benefit, the good fails to be provided a quarter of the time.

5.2. The first-order belief of t_i is the image of $\beta_i(t_i)$ under the projection to Θ, and $\mathrm{pr}_\Theta\circ(\mathrm{id}\times h_{-i})=\mathrm{pr}_\Theta$; so h preserves first-order beliefs. Let b_i^k send a type to its k-th order belief, so that $b_i^{k+1}=\mathcal{G}(\mathrm{id}\times b_{-i}^k)\circ\beta_i$. If $b'^k\circ h=b^k$ then $b_i'^{k+1}\circ h_i=\mathcal{G}(\mathrm{id}\times b_{-i}'^k)\mathcal{G}(\mathrm{id}\times h_{-i})\beta_i=\mathcal{G}(\mathrm{id}\times b_{-i}^k)\beta_i=b_i^{k+1}$. This is the statement that the maps to the terminal sequence commute with coalgebra homomorphisms.

5.3. A $+$ type meets a betting opponent with probability $\frac{1}{2}$, and exactly when $\theta=1$: betting pays $\frac{1}{2}>0$. A $-$ type would meet a betting opponent exactly when $\theta=0$: betting pays -1. So the profile is an equilibrium of $\mathbb{B}$. In $\mathbb{A}$ joint betting pays $\frac{1}{2}-1<0$. The type morphism $\mathbb{B}\to\mathbb{A}$ is the pair of constant maps: the image of $\beta_i(+)=\frac{1}{2}\delta_{(1,+)}+\frac{1}{2}\delta_{(0,-)}$ is $\frac{1}{2}\delta_{(1,*)}+\frac{1}{2}\delta_{(0,*)}$, the belief of the single type of $\mathbb{A}$, and similarly for $-$.

5.4. The only cycle is $\omega_1t_1\omega_2t_2$, and the equation is $\frac{\beta_1(\omega_2)}{\beta_1(\omega_1)}=\frac{\beta_2(\omega_2)}{\beta_2(\omega_1)}$: agents 1 and 2, who have the same information, must hold the same beliefs. Agent 3's edges are pendant and impose nothing.

5.5. Write O_1, O_1' for agent 1's odds $\mu(b){:}\mu(a)$ and $\mu(d){:}\mu(c)$, and O_2, O_2' for agent 2's odds $\mu(c){:}\mu(a)$ and $\mu(d){:}\mu(b)$; the cycle equation is $O_1O_2' = O_2O_1'$. Suppose $O_1O_2' > O_2O_1'$. Let $f = (-A, B, C, -D)$ on (a, b, c, d) with positive entries, agent 1 receiving f and agent 2 receiving $-f$. All four types gain iff $A/O_1 < B < DO_2'$ and $DO_1' < C < A/O_2$, which is solvable iff $O_2O_1' < A/D < O_1O_2'$. For instance if agent 1's posteriors and agent 2's at t are even and agent 2's at t' is $(\frac{1}{3}, \frac{2}{3})$, then $f = (-3, 3.5, 2.5, -2)$ gives the four types expected gains $\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{6}$. If μ were a common prior, $\mathbb{E}_\mu f$ would be an average of agent 1's conditional expectations, hence positive, and minus an average of agent 2's, hence negative.

5.6. $(\mathbb{E}_p f)(t) = \sum_{\omega \in p^{-1}t} \mu(\omega) f(\omega)/\mu(p^{-1}t)$. Then $(\mathbb{E}_r \mathbb{E}_p f)(t') = \sum_{t \in r^{-1}t'} \mu(p^{-1}t)(\mathbb{E}_p f)(t)/\mu((rp)^{-1}t') = (\mathbb{E}_{rp} f)(t')$. The Bayesian inverse is $p^\dagger_\mu(\omega \mid t) = \mu(\omega)1_{p(\omega)=t}/\mu(p^{-1}t)$, and $\mathbb{E}_p f$ is the composite of $p^\dagger_\mu$ with f in the Kleisli category. The tower property is $(rp)^\dagger = p^\dagger r^\dagger$.

Chapter 6

6.1. If all gains vanish, $f(\sigma) = \sigma$. Conversely $f(\sigma) = \sigma$ gives $\sigma_i(a_i) \sum_b g_b(\sigma) = g_{a_i}(\sigma)$ for all a_i. If $\sum_b g_b > 0$, every action in the support of σ_i would have positive gain, that is, would do strictly better than σ_i against σ_{-i}. But the payoff of σ_i is the average of the payoffs of the actions in its support, and not all of them can exceed their average. So all gains vanish.

6.2. (a, a), (b, b), and the mixed profile in which each plays a with probability $\frac{1}{3}$ (from $2p = 1 - p$), worth $\frac{2}{3}$. With $u_1(b, b) = 1 + t$ the column player's mixture must satisfy $2q = (1 + t)(1 - q)$, so $q = \frac{1+t}{3+t}$, while $p = \frac{1}{3}$ is unchanged. The mixed equilibrium exists iff $t > -1$. At $t = -1$ it merges with (b, b) into a segment of equilibria, and for $t < -1$ action a is dominant for the row player and only (a, a) remains: indices $+1$ and -1 have cancelled.

6.3. Obedience for a row player told C reads $-\pi(C, C) - \pi(C, D) \geq 0$, so $\pi(C, \cdot) = 0$; likewise $\pi(\cdot, C) = 0$. In general let σ_i strictly dominate a_i and average the obedience constraints for a_i over $a_i' \sim \sigma_i$: $\sum_{a_{-i}} \pi(a_i, a_{-i})[u_i(a_i, a_{-i}) - u_i(\sigma_i, a_{-i})] \geq 0$ with every bracket negative, so $\pi(a_i, \cdot) = 0$.

6.4. Put $\pi(C, C) = x$, $\pi(C, D) = \pi(D, C) = y$, $\pi(D, D) = z$, $x + 2y + z = 1$. Obedience after C is $-x + 2y \geq 0$ and after D is $y - 2z \geq 0$. In the (x, z) plane the polytope is $2x + z \leq 1$, $x + 5z \leq 1$, $x, z \geq 0$, with vertices $(0, 0)$, $(\frac{1}{2}, 0)$, $(\frac{4}{9}, \frac{1}{9})$ and $(0, \frac{1}{5})$; the third is the mixed Nash equilibrium. The sum of payoffs is $12x + 18y = 9 + 3x - 9z$, maximal at $(x, y, z) = (\frac{1}{2}, \frac{1}{4}, 0)$, worth 5.25 to each player. This improves on the distribution of Example 6.12, which is interior to an edge of the polytope.

6.5. Let a public device name one player uniformly at random, who then acts alone. The named player gets $b - c > 0$ rather than 0; the others get b rather than $b - c$. The good is always provided, and each player expects $b - c/n$. The distribution over profiles is invariant under all permutations of the players. It maximizes total surplus, so it is the best invariant correlated equilibrium. The symmetric mixed equilibrium, also invariant, gives each player only $b - c$. Correlation can respect a symmetry *ex ante* while breaking it *ex post*, which independent randomization cannot.

6.6. The marginals of $\mu \otimes \nu$ are μ and ν, so ∇ is injective; and $\Delta(f \times g)(\mu \otimes \nu) = \Delta f(\mu) \otimes \Delta g(\nu)$. If $x \neq x'$ and $y \neq y'$, the distribution $\frac{1}{2}\delta_{(x,y)} + \frac{1}{2}\delta_{(x',y')}$ is not a product, since a product with its marginals gives (x, y') probability $\frac{1}{4}$. If X or Y is a point the map is a bijection, and if one is empty both sides are empty.

Chapter 7

7.1. $\bot \leq f(\bot)$, so $f^k(\bot)$ is increasing and stabilizes in a finite lattice at a fixed point, which lies below every fixed point y because $f^k(\bot) \leq f^k(y) = y$. Dually from $\top$. For fixed points x, y, $z = x \vee y$ satisfies $z \leq f(z)$, and iterating f from z reaches the least fixed point above both; dually for meets. So $\mathrm{Fix}(f)$ is a lattice. On $\mathbb{N}$, a lattice with no top, $n \mapsto n + 1$ is monotone without fixed points.

7.2. For $x' \geq x$, $y' \geq y$, supermodularity applied to (x', y) and (x, y') is increasing differences. Conversely two incomparable points (x_1, y_1), (x_2, y_2) with $x_1 > x_2$, $y_1 < y_2$ have join (x_1, y_2) and meet (x_2, y_1), and increasing differences between these four points is the supermodular inequality. For $f = g(x + y)$ the condition is $g(s + h + k) - g(s + k) \geq g(s + h) - g(s)$ for $h, k \geq 0$: increments of g are nondecreasing, which for continuous g is convexity.

7.3. Each firm's first-order condition is $a - c - bQ - bq_i = 0$; summing, $Q = \frac{n(a-c)}{(n+1)b}$, $q = \frac{a-c}{(n+1)b}$, $P - c = \frac{a-c}{n+1}$, and profit per firm $\frac{(a-c)^2}{(n+1)^2 b}$. Industry profit is proportional to $\frac{n}{(n+1)^2}$, whose derivative has the sign of $1 - n$.

7.4. The best-response map $p \mapsto \frac{1}{2}(a + c + sp)$ is increasing with slope $s/2 < 1$. Iterating from $\bar{p}$ gives a decreasing sequence and from c an increasing one, both converging to $p^* = \frac{a+c}{2-s}$. By Theorem 7.5(ii) these limits bound the rationalizable prices, so only p^* is rationalizable.

7.5. Let $m_{-i} = \min_{j \neq i} e_j$. For $e_i' \geq e_i$ the difference $2[\min(e_i', m_{-i}) - \min(e_i, m_{-i})] - (e_i' - e_i)$ is nondecreasing in m_{-i}, which is nondecreasing in e_{-i}; action sets are chains. So the game is supermodular. The best response is

$e_i = m_{-i}$: effort up to the minimum returns 2 per unit at cost 1, and beyond it returns nothing. The pure equilibria are the $m+1$ profiles $(k,\dots,k)$, a chain ordered by payoff k. The greatest is $(m,\dots,m)$ and the least $(0,\dots,0)$, and every action is rationalizable.

7.6. On $\{0,1\}\times\{0,1\}$ let $f(0,t)=0$, $f(1,0)=2$, $f(1,1)=1$. The difference $f(1,t)-f(0,t)$ is positive for both t, so single crossing holds, but it decreases in t. The function $x+y$ on $\{0,1\}^2$ is supermodular; $\sqrt{x+y}$ is not, since $\sqrt{2}+0<1+1$.

Chapter 8

8.1. $\Gamma = \big(1, \{Out, In\},\ Out \mapsto (0,4),\ In \mapsto (2, \{A,F\}, A \mapsto (1,2), F \mapsto (-1,0))\big)$. (a) The inner node evaluates to $(1,2)$ and the root to $(1,2)$. (b) There are no ties, so the fold is $\{(1,2)\}$. (c) With $\alpha(o)=1$ and $\alpha(i,A,v)=\sum_a v(a)$ the fold is 3.

8.2. (i) Take $h'=h$ for $\Box Q \subseteq Q$; transitivity of "at or after" gives $\Box Q \subseteq \Box\Box Q$; a universal quantifier commutes with intersections. (ii) If $(h,\sigma)\in\Box Q$ and h'' is after h, everything after h'' is after h, so $(h'',\sigma)\in\Box Q$. If $H\subseteq Q$ is hereditary and $(h,\sigma)\in H$ then $(h',\sigma)\in H\subseteq Q$ for all later h'. (iii) is the definition. In the entry game the profile (*Out*, *Fight*) is in Nash at the root and not at the node after *In*.

8.3. With $q_1=\frac{a-c}{2b}$ and $q_2=\frac{a-c}{4b}$ the price is $c+\frac{a-c}{4}$, so profits are $\frac{(a-c)^2}{8b}$ and $\frac{(a-c)^2}{16b}$, against $\frac{(a-c)^2}{9b}$ each in Cournot. For prices, the follower plays $p_2=\frac{1}{2}(a+c+sp_1)$ and the leader maximizes $(p_1-c)(a-p_1+sp_2)$, giving

$$p_1=\frac{a(2+s)+c(2+s-s^2)}{2(2-s^2)}.$$

With $a=10$, $c=1$, $s=\frac{1}{2}$: $p_1\approx 7.79$ and $p_2\approx 7.45$, both above the simultaneous price 7.33; profits are about 40.29 for the leader and 41.56 for the follower, against 40.11 each. With strategic substitutes the leader gains and the follower loses. With complements both gain and the follower gains more, so each firm would like the other to lead.

8.4. The buyer accepts $p<v$ and rejects $p>v$. If she accepted $p=v$ with probability $\alpha<1$ the seller would earn αv at v and $v-\eta$ at $v-\eta$, and would have no best response. So $\alpha=1$ and $p=v$. On a grid, if v is not a grid point the seller offers the largest grid price below v. If it is, the buyer may accept $p=v$ with any probability α: for $\alpha\geq 1-\epsilon/v$ the seller offers v, and for $\alpha\leq 1-\epsilon/v$ she offers $v-\epsilon$. All of these are subgame perfect, their outcomes lie within ϵ of v, and the limit as $\epsilon\to 0$ selects acceptance when indifferent.

8.5. Yes. A deviation still gains 1 and loses 2. After a defection by player 1 the continuation pays $(0,3)$, which is not Pareto-dominated by $(2,2)$: the innocent party is better off under the punishment and will not agree to renegotiate. The three second-stage equilibria are mutually undominated, and the cooperative path pays $(4,4)$, which no other profile built from them dominates. Punishments that move along the Pareto frontier survive renegotiation; punishments that move inside it do not.

8.6. Take a decision node all of whose children are terminal. Its owner has a unique best action there. Any strategy choosing otherwise is weakly dominated by the strategy that differs only at that node: equal if the node is not reached, better if it is. Delete those strategies, replace the node by its backward-induction value, and repeat. The surviving strategies agree with backward induction wherever their own earlier choices allow the node to be reached, which is one point of the reduced strategic form. The order of deletion was chosen, and Example 1.10 shows that this matters in general. For these games a theorem of Marx and Swinkels [MaSw97] shows that the payoff outcome, though not the set of surviving strategies, is the same for every order.

Chapter 9

9.1. Tit-for-tat: states $\{\mathsf{c},\mathsf{d}\}$, outputs C and D, initial state c, and transition to c or d according as the opponent's last action was C or D. Win-stay, lose-shift: the same states and outputs, with transition to c iff the two last actions agreed. The induced strategies send the empty history to C; a history of length one or two to the opponent's last action (tit-for-tat) or to C iff the last action profile was (C,C) or (D,D) (win-stay, lose-shift). Both depend only on the last profile.

9.2. Against tit-for-tat the candidates are: conform, worth 2; defect forever, worth $3-2\delta$; alternate D,C, worth $\frac{3}{1+\delta}$. Both deviations are unprofitable iff $\delta \geq \frac{1}{2}$. Now consider the history after player 1 alone has defected: tit-for-tat prescribes C for player 1 and D for player 2, and the profile then alternates. Player 1 earns $\frac{3\delta}{1+\delta}$ by conforming and 1 by defecting once more (after which both defect forever): the deviation pays iff $\delta < \frac{1}{2}$. Player 2 earns $\frac{3}{1+\delta}$ by conforming and 2 by forgiving: the deviation pays iff $\delta > \frac{1}{2}$. So the profile is subgame perfect only at $\delta = \frac{1}{2}$.

9.3. If W is a single point the continuation is constant and supplies no incentives, so the action profile must be the stage equilibrium (D,D). Thus $B(\{(1,1)\}) = \{(1,1)\}$ and $B(\{(2,2)\}) = \{(1+\delta,1+\delta)\}$, which does not contain $(2,2)$.

9.4. If $W \subseteq W'$, any map $w\colon A \to W$ is a map into W', so $B(W) \subseteq B(W')$. In any game $B(\emptyset) = \emptyset$, because there are no continuation maps; so $\bigcup_n B^n(\emptyset) = \emptyset$, while $\bigcap_n B^n(W_0) = V^*$ contains the stage-equilibrium payoff.

9.5. Let entrants stay out as long as the incumbent has never accommodated, and enter forever otherwise; let the incumbent fight iff she has never accommodated. After an entry, fighting yields $(1-\delta)\cdot 0 + \delta \cdot 4$ and accommodating yields 2, so she fights iff $\delta \geq \frac{1}{2}$. This is the least such δ: fighting costs 2 today, her best continuation is 4, and she can guarantee 2 in every period by accommodating, so no punishment is harsher.

9.6. With $p = V/C$, $u(H,x) - u(D,x) = \frac{1}{2}(V - pC) = 0$, so every y is a best response to x. For $y = (q, 1-q)$, $u(x,y) - u(y,y) = (p-q)\big[u(H,y) - u(D,y)\big] = \frac{1}{2}C(p-q)^2 > 0$ for $q \neq p$. Mean fitness is $u(D,x) = \frac{V}{2}\big(1 - \frac{V}{C}\big)$, less than the $\frac{V}{2}$ of a population of doves.

Chapter 10

10.1. Let the incumbent fight with probability φ. *Out* pays the entrant 0, small entry $1 - 2\varphi$, large entry $2 - 4\varphi$. For $\varphi < \frac{1}{2}$ she enters large; for $\varphi > \frac{1}{2}$ she stays out; at $\varphi = \frac{1}{2}$ she is indifferent among all three. The incumbent's gain from fighting is $s - 2l$, where s, l are the probabilities of small and large entry. The equilibria are: (large, accommodate); (*Out*, φ) for every $\varphi \geq \frac{1}{2}$; and the mixed profiles in which the entrant plays $(1-3l,\ 2l,\ l)$ on (*Out*, small, large) with $0 < l \leq \frac{1}{3}$ and $\varphi = \frac{1}{2}$. All are sequential. In the first, Bayes' rule gives $m = 0$. In the second the belief is free: take $m \geq \frac{2}{3}$ if $\varphi = 1$ and $m = \frac{2}{3}$ if $\varphi < 1$. In the third Bayes' rule gives $m = \frac{2l}{3l} = \frac{2}{3}$, at which the incumbent is indeed indifferent.

10.2. One stratum with a single block, the interior of the hexagon. Six with two blocks: for each state k, the stratum $(\{i,j\},\{k\})$ is the open side of the triangle on which $p_k = 0$, and the stratum $(\{k\},\{i,j\})$ is the short side replacing the vertex δ_k, parametrized by $p_i{:}p_j$. Six with three blocks, the orderings (i,j,k): the vertices of the hexagon, each the meeting point of the sides $(\{i,j\},\{k\})$ and $(\{i\},\{j,k\})$. In all $1 + 6 + 6 = 13$.

10.3. (a) The pair $(\delta_1,\ (0, \frac{2}{3}, \frac{1}{3}))$. (b) $(1-\varepsilon,\ \frac{2}{3}\varepsilon,\ \frac{1}{3}\varepsilon)$ for a positive infinitesimal ε. The representation in (b) is not unique: perturbing the second and third coordinates by terms of order ε^2 gives the same conditional system.

10.4. With two towns left and reputation q, the ordinary incumbent's continuation value is 4 if $q < \frac{1}{4}$ (entry, and she is indifferent), 6 if $\frac{1}{4} < q < \frac{1}{2}$ (no entry now, entry in the last town), and 8 if $q \geq \frac{1}{2}$. In the first of three towns,

accommodating reveals her type and is worth $2+4$; fighting is worth 1 plus the value at the posterior. If $p > \frac{1}{4}$ she fights for sure and entry is deterred. If $p < \frac{1}{4}$ she must mix so that the posterior after a fight is exactly $\frac{1}{4}$: $\varphi = \frac{3p}{1-p}$, the second entrant then entering with probability $\frac{1}{2}$ to keep her indifferent. The first entrant faces a fight with probability $p + (1-p)\varphi = 4p$ and stays out iff $p > \frac{1}{8}$.

10.5. States: $(0,0)$ in state a, and $(n, n-1)$, (n,n) for $n \geq 1$ in state b. Player 1's cells are $\{(0,0)\}$ and $\{(n,n-1),(n,n)\}$; player 2's are $\{(n,n),(n+1,n)\}$ for $n \geq 0$. The graph is the path $(0,0) - (1,0) - (1,1) - (2,1) - \cdots$, with edges alternately cells of 2 and of 1. $K_1 E = E$, everything but $(0,0)$; $K_2 K_1 E$ removes $(1,0)$; $K_1 K_2 K_1 E$ removes $(1,1)$ as well. Each further operator removes one more state, and the intersection is empty.

10.6. Monotonicity is that of probability. If $\mathbb{P}(E) \geq p$ and $\mathbb{P}(F) \geq p$ then $\mathbb{P}(E \cap F) \geq 2p - 1$. If a cell is contained in E its probability of E is 1, so $K_i E \subseteq \mathsf{B}_i^1 E$, with equality when posteriors have full support. As for the last question: only loosely. The family is antitone in p; it has a counit only at grade 1 and only up to null events; each B_i^p is idempotent, since $\mathsf{B}_i^p E$ is a union of cells; and conjunction degrades grades by $p \odot q = \max\{0, p+q-1\}$. It is a lax monoidal family graded by the unit interval with that product, comonadic at the top grade.

Chapter 11

11.1. With effort observable she prefers H iff $0.8v - 4 \geq 0.4v - 1$, that is $v \geq 7.5$. With hidden effort she prefers L iff $0.8v - 5 < 0.4v - 1$, that is $v < 10$. The range is $7.5 \leq v < 10$.

11.2. Incentive compatibility is $0.4(w_s - w_f) \geq 1$. With $w \geq 0$ the cheapest contract is $w_f = 0$, $w_s = 2.5$, expected wage 2. The agent's expected utility is $2 - 1 = 1 > 0$: limited liability leaves her a rent of 1, which a negative w_f would otherwise extract.

11.3. The likelihood ratio of low to high effort for the pair is the product of the ratios for the outcome and for the signal, and the latter is $\frac{5}{9}$ at $y = 1$ and 5 at $y = 0$; so the pair is more informative than the outcome. Under H the four results (success/failure, y) have probabilities $0.72, 0.08, 0.18, 0.02$, and under L $0.2, 0.2, 0.3, 0.3$. Minimizing $\sum p_H u_z^2$ subject to the two binding constraints and $u_z \geq 0$ gives $u_z = \max\{0,\ a + m(1-\ell_z)\}$ with ℓ_z the likelihood ratio, and the solution has $u = 0$ at (failure, $y = 0$) and

$$u \approx (2.378,\ 0.663,\ 1.306,\ 0), \qquad w \approx (5.65,\ 0.44,\ 1.71,\ 0),$$

with expected wage ≈ 4.41. The agency cost falls from 1 to about 0.41.

11.4. Reflexivity: $h = \mathrm{id}$. Transitivity: garblings compose. For any deterministic s, $s \circ f$ is by definition a garbling of f. Sufficiency adds $f \preceq s \circ f$, so s is sufficient iff f and $s \circ f$ are equivalent in the preorder.

11.5. The surplus $\frac{\beta}{c} - \frac{\beta^2}{2c} - \frac{1}{2} r\sigma^2 \beta^2$ has first-order condition $1 - \beta(1 + rc\sigma^2) = 0$. At the optimum it equals $\frac{\beta^*}{2c}$. With observable effort $e = \frac{1}{c}$ and full insurance the surplus is $\frac{1}{2c}$. The agency cost is $\frac{1}{2c}\left(1 - \beta^*\right) = \frac{r\sigma^2}{2(1+rc\sigma^2)}$.

11.6. See Figure D.1. The payment box receives the contract w and the outcome x. With observable effort it would also receive a copy of e; that wire, dotted in the figure, is what is missing.

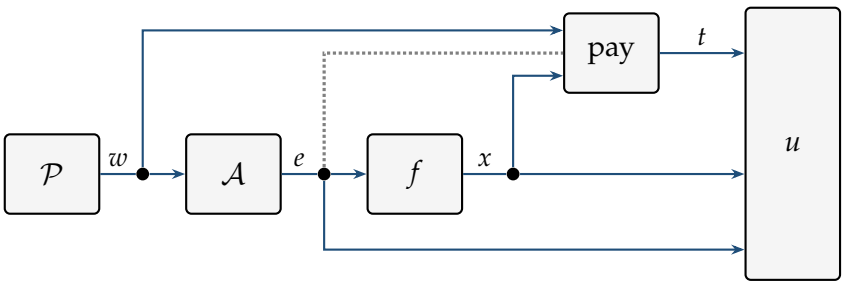


Figure D.1: The principal–agent problem with hidden action: the principal chooses a contract, the agent an effort, Nature the outcome. Backward wires are omitted. The dotted wire, a copy of e into the payment rule, is present only when effort is observable.

Chapter 12

12.1. A non-shirker has $V_N = w - e + \delta[(1-s)V_N + sV_U]$, so $V_N - V_U = \frac{w - e - \overline{u}}{1 - \delta(1-s)}$. Shirking once raises the probability of losing the job from s to $s + (1-s)q$, so the condition is $\delta(1-s)q\,(V_N - V_U) \geq e$, that is

$$w \;\geq\; \overline{u} + e + \frac{e\,[1 - \delta(1-s)]}{\delta(1-s)\,q}.$$

The last term is $\frac{e}{q}\left(\frac{1}{\delta(1-s)} - 1\right)$, increasing in s: a job that may end anyway is a weaker hostage.

12.2. The difference of two independent uniforms on $[-a, a]$ has a triangular density with $g(0) = \frac{1}{2a}$. First-best effort is $e = 1$, so the spread is $W_1 - W_2 = 2a$. Participation binds when $\frac{1}{2}(W_1 + W_2) - \frac{1}{2} = 0$: $W_1 = \frac{1}{2} + a$, $W_2 = \frac{1}{2} - a$. One must check that $e = 1$ is a global best response. For $e < 1$ the marginal net

benefit is $(1-e)\big(1-\frac{1}{2a}\big)$, which is positive only if $a > \frac{1}{2}$. With less noise than that, the first-order condition picks out a minimum and no symmetric pure-strategy equilibrium implements the first best.

12.3. Efficient efforts are $e_i = 1$ with surplus 1. Under equal sharing each maximizes $\frac{1}{2}(e_1+e_2) - \frac{1}{2}e_i^2$, so $e_i = \frac{1}{2}$ and the surplus is $\frac{3}{4}$, a loss of $\frac{1}{4}$. Under the scheme paying $b_i = 1$ iff $x \geq 2$: if the other works 1, working 1 yields $\frac{1}{2}$, and any lower effort yields at most 0. So $(1,1)$ is an equilibrium. So is $(0,0)$: reaching the target alone costs 2 and pays 1. So is every $(e_1, 2-e_1)$ with $2-\sqrt{2} \leq e_1 \leq \sqrt{2}$. The scheme makes efficiency an equilibrium, not the equilibrium.

12.4. The agent's certainty equivalent is $\alpha + \beta^\top e - \frac{1}{2}e^\top Ce - \frac{r}{2}\beta^\top\Sigma\beta$, so $e = C^{-1}\beta$, and total surplus is $b^\top C^{-1}\beta - \frac{1}{2}\beta^\top C^{-1}\beta - \frac{r}{2}\beta^\top\Sigma\beta$. The first-order condition $C^{-1}(b-\beta) = r\Sigma\beta$ gives $\beta = (I + rC\Sigma)^{-1}b$. With $\beta_2 = 0$, $e = \beta_1 C^{-1}(1,0)^\top$ and the surplus in β_1 is quadratic with maximizer $\frac{b_1C_{22} - b_2C_{12}}{C_{22} + r\sigma_1^2 \det C}$, which is the stated formula. In the example $\beta_1^* = \frac{1-2\kappa}{2-\kappa^2}$: a flat wage at $\kappa = \frac{1}{2}$, and a penalty on measured output for $\frac{1}{2} < \kappa < 1$.

12.5. Let the contracts be an incentive scheme w_0 and full insurance w_1. At the renegotiation stage effort is sunk and both parties prefer w_1 to w_0, while nothing is preferred by both to w_1. So $\rho(w_0) = \rho(w_1) = w_1$ and $\rho\rho = \rho$. The principle: the outcome of signing w and then renegotiating equals the outcome of signing $\rho(w)$, which is not renegotiated. Hence the set of achievable outcomes is that of the contracts in the image of ρ, here $\{w_1\}$, and design under renegotiation is design under commitment restricted to $\mathrm{Fix}(\rho)$.

Chapter 13

13.1. For $p \leq 1$ the average quality offered is $p/2$, so $\Phi(p) = \frac{p}{2} + k$; for $p \geq 1$, $\Phi(p) = \frac{1}{2} + k$. If $k < \frac{1}{2}$ the unique equilibrium is $p = 2k$ and a fraction $2k$ of cars trades, although every trade is efficient. If $k \geq \frac{1}{2}$ the price is $k + \frac{1}{2}$ and all cars trade.

13.2. Let quality be uniform on $[1,2]$ with probability $\frac{1}{2}$ and uniform on $[3,4]$ with probability $\frac{1}{2}$, and let $v = 1.7$. For $2 \leq p \leq 3$ only the lower group sells and $\Phi = 1.7 \times 1.5 = 2.55$, a fixed point. For $3 \leq p \leq 4$, $\mathbb{E}[\theta \mid \theta \leq p] = \frac{p^2-6}{2p-4}$, and $\Phi(p) = p$ has the root $p \approx 3.435$. For $p \geq 4$, $\Phi = 1.7 \times 2.5 = 4.25$, a fixed point. So the equilibrium prices are 2.55, 3.435 and 4.25; the middle one is unstable. Buyers earn zero in each. A seller who sells at a lower equilibrium price also sells at a higher one and receives more, and a seller who newly sells prefers the price to her car; so the equilibria are Pareto-ranked by price.

13.3. $(4-0.75\alpha)(1+2.25\alpha)=6.25$ is $1.6875\alpha^2-8.25\alpha+2.25=0$, with smaller root $\alpha^*\approx 0.2899$. At the separating pair the high risks have utility $\ln 2.5\approx 0.916$ and the low risks $\frac{3}{4}\ln(4-0.75\alpha^*)+\frac{1}{4}\ln(1+2.25\alpha^*)\approx 1.123$. With $\lambda=0.1$ the pooled fair premium for full cover is $0.275\times 3=0.825$, leaving wealth 3.175 and utility $\ln 3.175\approx 1.155$ for both types. Both prefer it, so the separating pair is not an equilibrium, consistent with $\lambda<\lambda^*$.

13.4. A buy order has probability $\frac{1}{2}(1+\mu)$ if $v=1$ and $\frac{1}{2}(1-\mu)$ if $v=0$. So

$$\text{ask}=\frac{\pi(1+\mu)}{\pi(1+\mu)+(1-\pi)(1-\mu)},\qquad \text{bid}=\frac{\pi(1-\mu)}{\pi(1-\mu)+(1-\pi)(1+\mu)},$$

and after two buys the posterior is $\frac{\pi(1+\mu)^2}{\pi(1+\mu)^2+(1-\pi)(1-\mu)^2}$: each buy multiplies the odds by $\frac{1+\mu}{1-\mu}$.

13.5. Given v the insider earns $(v-p_0-\lambda x)x=(v-p_0)^2\beta(1-\lambda\beta)=\frac{1}{2}\beta(v-p_0)^2$, with expectation $\frac{1}{2}\beta\Sigma_0=\frac{1}{2}\sigma_u\sqrt{\Sigma_0}$. Noise traders lose $\mathbb{E}[(p-v)u]=\lambda\sigma_u^2=\frac{1}{2}\sigma_u\sqrt{\Sigma_0}$. The two are equal because the market maker breaks even. Both are proportional to σ_u: noise is the insider's camouflage, and more of it makes the market deeper (λ falls) and her profit larger.

Chapter 14

14.1. Let σ be a Bayesian equilibrium of (M,g), $M=\prod_i M_i$. For agent i define the interim utility $V_i(m_i,\theta_i)=\mathbb{E}\big[U_i(g(m_i,\sigma_{-i}(\theta_{-i})),\theta)\mid\theta_i\big]$. Then σ_i is a best response in the single-agent mechanism with messages M_i and utility V_i. Apply part (i) with $h=\sigma_i\colon\Theta_i\to M_i$: truth-telling is a best response in the pulled-back mechanism, in which a report θ_i' yields $g(\sigma_i(\theta_i'),\sigma_{-i}(\theta_{-i}))$. This is agent i's problem in the direct mechanism $(\Theta,g\circ\sigma)$ when the others are truthful. Since it holds for every i, truth-telling is a Bayesian equilibrium there. Initiality is as before, the unique morphism being $\sigma=(\sigma_i)_i$.

14.2. $\ell(i,j)=\theta_j(x_j-x_i)$ gives $\ell(1,2)=2$, $\ell(1,3)=9$, $\ell(2,3)=6$, $\ell(2,1)=-1$, $\ell(3,2)=-4$, $\ell(3,1)=-3$. Shortest paths improve two entries: $d(1,3)=2+6=8$ and $d(3,1)=-4-1=-5$. So d has rows $(0,2,8)$, $(-1,0,6)$, $(-5,-4,0)$. With $t_1=0$ the payment rules are $1\le t_2\le 2$, $4\le t_3-t_2\le 6$, and the greatest is $(0,2,8)$, the first row.

14.3. With $\theta_k=k/n$, $\ell(\theta_i,\theta_j)=\theta_j(\theta_j-\theta_i)$. Splitting an upward step at an intermediate type shortens it, so the shortest upward path takes every step: $d_f(0,1)=\sum_{k=1}^{n}\frac{k}{n}\cdot\frac{1}{n}=\frac{n+1}{2n}$. Likewise downward: $d_f(1,0)=-\sum_{k=0}^{n-1}\frac{k}{n}\cdot\frac{1}{n}=-\frac{n-1}{2n}$. The sum is $\frac{1}{n}$.

14.4. If $p \le p'$ then $v - p \ge v - p'$ pointwise, so $p^v \ge p'^v$; likewise for U. From $p^v(\theta) \ge v(x,\theta) - p(x)$ we get $p^{vv} \le p$, and similarly $U^{vv} \le U$. Hence $p^{vvv} = (p^{vv})^v \ge p^v$ and $p^{vvv} = (p^v)^{vv} \le p^v$. For $U = \frac{1}{2}\theta^2$, $U^v(x) = \sup_\theta[\theta x - \frac{1}{2}\theta^2] = \frac{1}{2}x^2$: facing this tariff type θ buys $x = \theta$.

14.5. The virtual type is $2\theta - 1$, so $x(\theta) = (2\theta - 1)^+$; by Corollary 14.5 $t(\theta) = \theta x - \int_{1/2}^{\theta} x = \theta^2 - \frac{1}{4}$; substituting $\theta = \frac{1}{2}(x+1)$ gives the tariff; and the profit is $\int_{1/2}^{1} \frac{1}{2}(2\theta-1)^2\,d\theta = \frac{1}{12}$. For a single bundle (x,t) the buyers are the types above $\hat\theta = t/x$, and the profit $(\hat\theta x - \frac{1}{2}x^2)(1-\hat\theta)$ is maximized at $x = \hat\theta = \frac{2}{3}$, $t = \frac{4}{9}$, giving $\frac{2}{27}$: eight-ninths of the optimum.

14.6. Let S_{-i} be the sum of the others' reported net values. The pivot payment is $\max\{0, S_{-i}\} - S_{-i}$ if the project is built and $\max\{0, S_{-i}\}$ otherwise. It is nonnegative, and positive only when $S_{-i} < 0$ and the project is built, or $S_{-i} > 0$ and it is not: exactly when agent i changes the decision. Her payoff is $(v_i + S_{-i})1_{\text{built}} - \max\{0, S_{-i}\}$, so she wants the project built iff $v_i + S_{-i} \ge 0$, which a truthful report achieves for every S_{-i}.

Chapter 15

15.1. Let the high type choose e_h and the low type choose e_h with probability $\rho \in (0,1)$ and 0 otherwise. The posterior at e_h is $\mu = \frac{\lambda}{\lambda + (1-\lambda)\rho}$ and the wage $1 + \mu$; at 0 the wage is 1. The low type's indifference gives $e_h = \mu$, and the high type then strictly prefers e_h. With deviations attributed to the low type this is an equilibrium for every ρ, with e_h ranging over $(\lambda, 1)$. None survives the intuitive criterion: a signal $e' \in (1, 2-\mu)$ is equilibrium-dominated for the low type and profitable for the high type if believed.

15.2. Under full information each type equates marginal product and marginal cost: $\frac{\theta}{4\sqrt{e}} = \frac{1}{\theta}$, so $e_L = \frac{1}{16}$ and $e_H = 1$, with payoffs 1.0625 and 2.5. In the Riley outcome the low type keeps $e_L = \frac{1}{16}$. The high type's education must deter imitation: $2 + \sqrt{e} - e \le 1.0625$, that is $\sqrt{e} \ge \frac{1}{2}(1 + \sqrt{4.75})$, so $e_H \approx 2.53$, well above the efficient 1. Her payoff is about 2.33. Productive education does not remove the distortion; it changes the benchmark from which it is measured.

15.3. The boundary type c_i is indifferent between the receiver's actions on the two sides, $\frac{1}{2}(c_{i-1} + c_i)$ and $\frac{1}{2}(c_i + c_{i+1})$, so her ideal point $c_i + b$ is their midpoint; this is the difference equation. An interval of length l contributes $l \cdot \frac{l^2}{12}$ to the receiver's loss. For $b = \frac{1}{20}$: $N = 1$ gives $\frac{1}{12} \approx 0.0833$; $N = 2$ has lengths $\frac{2}{5}, \frac{3}{5}$ and loss ≈ 0.0233; $N = 3$ has lengths $\frac{2}{15}, \frac{5}{15}, \frac{8}{15}$ and loss ≈ 0.0159. In general the loss is $\frac{1}{12N^2} + \frac{b^2(N^2-1)}{3}$.

15.4. Posteriors 0 and 0.6 with mean 0.3 require weight $\frac{1}{2}$ on each. If effective drugs always pass and ineffective ones pass with probability x, then $0.3 + 0.7x = \frac{1}{2}$, so $x = \frac{2}{7}$, and the posterior after a pass is $0.3/0.5 = 0.6$. With prior 0.5 the weight on 0.6 is $\frac{5}{6}$: ineffective drugs pass with probability $\frac{2}{3}$, and approval has probability $\frac{5}{6}$.

15.5. $m(\tau) = \int \mu \, d\tau(\mu)$ is linear in τ, so it is affine and its fibres are convex. Take any experiment with two signals inducing two different posteriors and exchange the names of the signals: the kernel changes and τ does not.

15.6. Effort solves $a = \delta\kappa$, and the efficient level is 1. Since $\delta \leq 1$ and $\kappa < 1$, she works too little for all parameters. Working too hard requires the reputational return to accumulate over several future periods, as in the multi-period version of [Hol99].

Chapter 16

16.1. Maximizing $x(1-x)^\gamma$ gives $1 - x = \gamma x$, so $x = \frac{1}{1+\gamma}$ and player 2 receives $\frac{\gamma}{1+\gamma}$: one half at $\gamma = 1$, one third at $\gamma = \frac{1}{2}$, and nothing in the limit of extreme risk aversion.

16.2. Under the transformation the Nash product is multiplied by $\alpha_1\alpha_2$, so maximizers correspond. For monotonicity let S_1 be the convex hull of $(0,0), (1,0), (0,1), (0.75, 0.75)$ and S_2 that of $(0,0), (1,0), (0,1), (1,0.7)$, with $d = 0$. Then $S_1 \subseteq S_2$, the solutions are $(0.75, 0.75)$ and $(1, 0.7)$, and player 2 loses from the expansion.

16.3. $x_T = 1,\ 0.1,\ 0.91,\ 0.181,\ 0.8371$, and $\frac{1}{1+\delta} \approx 0.5263$. The errors are $-\frac{(-\delta)^T}{1+\delta}$: they alternate in sign and shrink by the factor 0.9 each round.

16.4. By rejecting and continuing, player 2 obtains $\delta_2(1 - \delta_1 x^*) = 1 - x^*$. If w is below this, opting out is never chosen and the equilibrium is unchanged. If w is larger, she opts out after a rejection, so player 1 offers exactly w and keeps $1 - w$, agreement being immediate. (The condition for her to prefer opting out to counter-offering, $w \geq \delta_2(1 - \delta_1(1 - w))$, reduces to the same threshold.) The outside option acts as a constraint on the split and not as a disagreement point.

16.5. Against $s(c) = \frac{2}{3}c + \frac{1}{4}$ a bid b trades with costs below $\bar{c} = \frac{3}{2}(b - \frac{1}{4})$. The derivative of the buyer's payoff is $\frac{3}{2}(v - b) - \frac{1}{2}\bar{c}$, which vanishes at $b = \frac{2}{3}v + \frac{1}{12}$. Symmetrically the seller's first-order condition against $b(v)$ gives $s = \frac{2}{3}c + \frac{1}{4}$. Trade occurs iff $v - c \geq \frac{1}{4}$. The difference $z = v - c$ has density $1 - z$ on $[0,1]$, so the gains are $\int_{1/4}^{1} z(1-z)\,dz = \frac{9}{64}$.

16.6. With efficient trade and each trader paid the other's report, the buyer pays c and the seller receives v whenever $v > c$; a buyer of value 0 and a seller of cost 1 never trade and get zero. The deficit is $v - c$ per trade, with expectation $\int_0^1 z(1-z)\,dz = \frac{1}{6}$.

Chapter 17

17.1. A bidder of value v bidding $b(z)$ earns $(v - b(z))z^{n-1}$; optimality at $z = v$ gives $(b\,v^{n-1})' = (n-1)v^{n-1}$, so $b = \frac{n-1}{n}v$. By revenue equivalence $m(v) = \int_0^v s\,d(s^{n-1}) = \frac{n-1}{n}v^n$ and $b = m/v^{n-1}$. Revenue is $n\int_0^1 m = \frac{n-1}{n+1}$.

17.2. Revenue is r when exactly one value exceeds r, and the lower value when both do: $2r^2(1-r) + \int_r^1 2s(1-s)\,ds = \frac{1}{3} + r^2 - \frac{4}{3}r^3$, maximal at $r = \frac{1}{2}$ with value $\frac{5}{12}$.

17.3. $m(v) = \int_0^v s\,d(s^2) = \frac{2}{3}v^3$. In a third-price auction the winner pays the lowest of the other two bids, whose expected value given that both are below v is $b(\frac{v}{3})$ for linear $b = kv$. So $v^2 \cdot \frac{kv}{3} = \frac{2}{3}v^3$ and $k = 2$: bidders bid twice their values.

17.4. If the opponent bids $2x_j$, winning yields $x_i + x_j - 2x_j = x_i - x_j$, positive exactly when $x_i > x_j$; bidding $2x_i$ wins in exactly those cases, whatever x_j is. So it is optimal even with knowledge of x_j. Revenue is $\mathbb{E}[2\min] = \frac{2}{3}$. In the first-price auction with $b(x) = x$ revenue is $\mathbb{E}[\max] = \frac{2}{3}$. They agree: with independent signals revenue equivalence holds even though values are interdependent.

17.5. For a product density both sides of the affiliation inequality are $\prod_i g_i(z_i)g_i(z_i')$. If (x, y) are affiliated then for $x' > x$ the ratio $g(y \mid x')/g(y \mid x)$ is increasing in y, so the conditional distribution at x' dominates that at x in the first-order stochastic sense, and expectations of increasing functions are ordered.

17.6. In a Dutch auction nothing is observed before the auction ends, so a bidder's successive decisions "not yet" can be coalesced into one choice of a stopping price, and the bidders' choices interchanged: the reduced strategic form is that of the first-price auction, for any structure of values. In an English auction with three or more bidders, exits are observed, information sets distinguish histories with different exit prices, and a strategy is a function of them; the reduced strategic form is strictly larger than that of the sealed bid. With two bidders the only observable event ends the auction, and the two formats are again the same game.

Chapter 18

18.1. In units of $(a-c)^2/b$: the collusive profit is $\frac{1}{4n}$; the best deviation, a best response to the others' $\frac{a-c}{2nb}$, earns $\frac{(n+1)^2}{16n^2}$; the Cournot profit is $\frac{1}{(n+1)^2}$. Grim trigger works iff $\delta \geq \frac{\pi^d-\pi^c}{\pi^d-\pi^N} = \frac{(n+1)^2}{n^2+6n+1}$. This is $\frac{9}{17}$ at $n=2$, above the Bertrand threshold $\frac{1}{2}$, and $\frac{4}{7}$ at $n=3$, below $\frac{2}{3}$; for all $n \geq 3$ it is below $1-\frac{1}{n}$. Price competition makes deviation more tempting and punishment harsher, and for more than two firms the first effect dominates.

18.2. At $p = 1-2k$ both sell their capacity. Undercutting gains nothing. A firm raising its price faces residual demand $1-p-k$, and its best such price is $\frac{1}{2}(1-k)$; it has no incentive to raise iff $1-2k \geq \frac{1}{2}(1-k)$, that is $k \leq \frac{1}{3}$. For $k \geq 1$ either firm can serve the market and the equilibrium is $p=0$. In between there is no pure-strategy equilibrium.

18.3. With the firms at a and $1-a$ the indifferent consumer is at $\hat{x} = \frac{1}{2} + \frac{p_2-p_1}{2t(1-2a)}$, the best response is $p_1 = \frac{1}{2}(p_2 + c + t(1-2a))$, and $p^* = c + t(1-2a)$ with profits $\frac{1}{2}t(1-2a)$, decreasing in a and equal to $\frac{t}{2}$ at $a=0$.

18.4. If all three share a position, each has $\frac{1}{3}$ and moving slightly to the larger side wins nearly half. If exactly two share a position z and the third is at $y > z$, the pair is content only if $y = 3z$, and then the third gains by moving towards z. If all are distinct, the leftmost gains votes by moving right for as long as she stays left of her neighbour, so she is not at a best response (and has none).

18.5. $\hat{v} = \frac{2-\delta}{4-3\delta}$, $p_1 = \hat{v}(1-\frac{\delta}{2})$, $p_2 = \frac{1}{2}\hat{v}$, and the profit simplifies to

$$\Pi(\delta) = \frac{(2-\delta)^2}{4\,(4-3\delta)},$$

which is below $\frac{1}{4}$ iff $\delta^2 < \delta$. It equals $\frac{1}{4}$ at both ends: at $\delta = 0$ the second period is worthless, and at $\delta = 1$ waiting is free, nobody buys early, and she sells once at $\frac{1}{2}$. The loss is greatest at $\delta = \frac{2}{3}$, where $\Pi = \frac{2}{9}$.

18.6. Owning from the first period is worth v and from the second δv, so the first period's services are worth $(1-\delta)v$. A lessor faces the whole population each period. She charges $\frac{1}{2}(1-\delta)$ in the first period and $\frac{1}{2}$ (in current terms) in the second, serving types above $\frac{1}{2}$ both times, and earns $\frac{1}{4}(1-\delta) + \frac{1}{4}\delta = \frac{1}{4}$, the commitment profit. Since the good returns to her each period, a later price cut would cost her exactly what it costs her customers.

References

[AB11] S. Abramsky and A. Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New J. Phys.* 13 (2011) 113036.

[AFR16] C. Alós-Ferrer and K. Ritzberger. *The Theory of Extensive Form Games*. Springer, 2016.

[AG17] E. Azevedo and D. Gottlieb. Perfect competition in markets with adverse selection. *Econometrica* 85 (2017).

[AGGKLN21] R. Atkey, B. Gavranović, N. Ghani, C. Kupke, J. Ledent and F. Nordvall Forsberg. Compositional game theory, compositionally. *Proc. ACT 2020*, EPTCS 333 (2021).

[AK08] A. Archer and R. Kleinberg. Truthful germs are contagious: a local-to-global characterization of truthfulness. *Proc. EC 2008*; *Games Econ. Behav.* 86 (2014).

[AK13] C. Alós-Ferrer and C. Kuzmics. Hidden symmetries and focal points. *J. Econ. Theory* 148 (2013).

[AM95] R. Aumann and M. Maschler. *Repeated Games with Incomplete Information*. MIT Press, 1995.

[ANST21] B. Ahrens, P. North, M. Shulman and D. Tsementzis. The univalence principle. arXiv:2102.06275 (2021).

[APS90] D. Abreu, D. Pearce and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58 (1990).

[AW17] S. Abramsky and V. Winschel. Coalgebraic analysis of subgame-perfect equilibria in infinite games without discounting. *Math. Structures Comput. Sci.* 27 (2017).

[AZ15] S. Abramsky and J. Zvesper. From Lawvere to Brandenburger–Keisler: interactive forms of diagonalization and self-reference. *J. Comput. System Sci.* 81 (2015).

[Abr88] D. Abreu. On the theory of infinitely repeated games with discounting. *Econometrica* 56 (1988).

[Ake70] G. Akerlof. The market for "lemons": quality uncertainty and the market mechanism. *Quart. J. Econ.* 84 (1970).

[Apt04] K. Apt. Uniform proofs of order independence for various strategy elimination procedures. *Contrib. Theor. Econ.* 4 (2004).

[Aum74] R. Aumann. Subjectivity and correlation in randomized strategies. *J. Math. Econ.* 1 (1974).

[Aum76] R. Aumann. Agreeing to disagree. *Ann. Statist.* 4 (1976).

[Aum87] R. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55 (1987).

[Axe84] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[BBD91] L. Blume, A. Brandenburger and E. Dekel. Lexicographic probabilities and equilibrium refinements. *Econometrica* 59 (1991).

[BD05] P. Bolton and M. Dewatripont. *Contract Theory*. MIT Press, 2005.

[BD93] A. Brandenburger and E. Dekel. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* 59 (1993).

[BGK85] J. Bulow, J. Geanakoplos and P. Klemperer. Multimarket oligopoly: strategic substitutes and complements. *J. Polit. Econ.* 93 (1985).

[BHS23] D. Braithwaite, J. Hedges and T. St Clere Smithe. The compositional structure of Bayesian inference. *Proc. MFCS 2023*, LIPIcs 272.

[BHZ19] J. Bolt, J. Hedges and P. Zahn. Bayesian open games. arXiv:1910.03656 (2019); *Compositionality* 5 (2023).

[BK06] A. Brandenburger and H. J. Keisler. An impossibility theorem on beliefs in games. *Studia Logica* 84 (2006).

[BK85] J.-P. Benoit and V. Krishna. Finitely repeated games. *Econometrica* 53 (1985).

[BK96] J. Bulow and P. Klemperer. Auctions versus negotiations. *Amer. Econ. Rev.* 86 (1996).

[BM16] D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theor. Econ.* 11 (2016).

[BM82] D. Baron and R. Myerson. Regulating a monopolist with unknown costs. *Econometrica* 50 (1982).

[BMSS12] L. Birkedal, R. Møgelberg, J. Schwinghammer and K. Støvring. First steps in synthetic guarded domain theory: step-indexing in the topos of trees. *Log. Methods Comput. Sci.* 8(4) (2012).

[BPW87] B. D. Bernheim, B. Peleg and M. Whinston. Coalition-proof Nash equilibria I: concepts. *J. Econ. Theory* 42 (1987).

[BR89] B. D. Bernheim and D. Ray. Collective dynamic consistency in repeated games. *Games Econ. Behav.* 1 (1989).

[BRW86] K. Binmore, A. Rubinstein and A. Wolinsky. The Nash bargaining solution in economic modelling. *RAND J. Econ.* 17 (1986).

[BS87] J. Banks and J. Sobel. Equilibrium selection in signaling games. *Econometrica* 55 (1987).

[Bag95] K. Bagwell. Commitment and observability in games. *Games Econ. Behav.* 8 (1995).

[Beb96] L. Bebchuk. A new theory concerning the credibility and success of threats to sue. *J. Legal Stud.* 25 (1996).

[Ber84] B. D. Bernheim. Rationalizable strategic behavior. *Econometrica* 52 (1984).

[Bin87] K. Binmore. Modeling rational players: part I. *Econ. Philos.* 3 (1987).

[Bla51] D. Blackwell. Comparison of experiments. *Proc. Second Berkeley Symp. Math. Statist. Probab.* (1951).

[Bla65] D. Blackwell. Discounted dynamic programming. *Ann. Math. Statist.* 36 (1965).

[Bul82] J. Bulow. Durable-goods monopolists. *J. Polit. Econ.* 90 (1982).

[CDT09] X. Chen, X. Deng and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56 (2009).

[CEGLMPR20] B. Clarke, D. Elkins, J. Gibbons, F. Loregian, B. Milewski, E. Pillmore and M. Román. Profunctor optics, a categorical update. arXiv:2001.07488 (2020); *Compositionality* 6 (2024).

[CGHR22] M. Capucci, B. Gavranović, J. Hedges and E. F. Rischel. Towards foundations of categorical cybernetics. *Proc. ACT 2021*, EPTCS 372 (2022).

[CGK87] P. Cramton, R. Gibbons and P. Klemperer. Dissolving a partnership efficiently. *Econometrica* 55 (1987).

[CGLN22] M. Capucci, N. Ghani, J. Ledent and F. Nordvall Forsberg. Translating extensive form games to open games with agency. *Proc. ACT 2021*, EPTCS 372 (2022).

[CH90] V. Crawford and H. Haller. Learning how to cooperate: optimal play in repeated coordination games. *Econometrica* 58 (1990).

[CJ19] K. Cho and B. Jacobs. Disintegration and Bayesian inversion via string diagrams. *Math. Structures Comput. Sci.* 29 (2019).

[CK87] I.-K. Cho and D. Kreps. Signaling games and stable equilibria. *Quart. J. Econ.* 102 (1987).

[CM88] J. Crémer and R. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56 (1988).

[CS82] V. Crawford and J. Sobel. Strategic information transmission. *Econometrica* 50 (1982).

[CS83] K. Chatterjee and W. Samuelson. Bargaining under incomplete information. *Oper. Res.* 31 (1983).

[Car01] G. Carlier. A general existence result for the principal-agent problem with adverse selection. *J. Math. Econ.* 35 (2001).

[Car12] G. Carroll. When are local incentive constraints sufficient? *Econometrica* 80 (2012).

[Car15] G. Carroll. Robustness and linear contracts. *Amer. Econ. Rev.* 105 (2015).

[Cla71] E. Clarke. Multipart pricing of public goods. *Public Choice* 11 (1971).

[Coa72] R. Coase. Durability and monopoly. *J. Law Econ.* 15 (1972).

[Cor20] D. Corfield. *Modal Homotopy Type Theory*. Oxford University Press, 2020.

[CvD93] H. Carlsson and E. van Damme. Global games and equilibrium selection. *Econometrica* 61 (1993).

[DFM06] E. Dekel, D. Fudenberg and S. Morris. Topologies on types. *Theor. Econ.* 1 (2006).

[DFM07] E. Dekel, D. Fudenberg and S. Morris. Interim correlated rationalizability. *Theor. Econ.* 2 (2007).

[DGP09] C. Daskalakis, P. Goldberg and C. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39 (2009).

[DM86] P. Dasgupta and E. Maskin. The existence of equilibrium in discontinuous economic games, I and II. *Rev. Econ. Stud.* 53 (1986).

[DR03] S. Demichelis and K. Ritzberger. From evolutionary to strategic stability. *J. Econ. Theory* 113 (2003).

[Dat03] R. Datta. Universality of Nash equilibria. *Math. Oper. Res.* 28 (2003).

[Dia71] P. Diamond. A model of price adjustment. *J. Econ. Theory* 3 (1971).

[Dia89] D. W. Diamond. Reputation acquisition in debt markets. *J. Polit. Econ.* 97 (1989).

[Die85] A. Diekmann. Volunteer's dilemma. *J. Conflict Resolution* 29 (1985).

[Dix80] A. Dixit. The role of investment in entry-deterrence. *Econ. J.* 90 (1980).

[EMW83] R. Engelbrecht-Wiggans, P. Milgrom and R. Weber. Competitive bidding and proprietary information. *J. Math. Econ.* 11 (1983).

[EO10] M. Escardó and P. Oliva. Selection functions, bar recursion and backward induction. *Math. Structures Comput. Sci.* 20 (2010).

[EO11] M. Escardó and P. Oliva. Sequential games and optimal strategies. *Proc. R. Soc. A* 467 (2011).

[EP06] J. Ely and M. Pęski. Hierarchies of belief and interim rationalizability. *Theor. Econ.* 1 (2006).

[ER94] S. Elmes and P. Reny. On the strategic equivalence of extensive form games. *J. Econ. Theory* 62 (1994).

[FGPR23] T. Fritz, T. Gonda, P. Perrone and E. F. Rischel. Representable Markov categories and comparison of statistical experiments in categorical probability. *Theoret. Comput. Sci.* 961 (2023).

[FKM11] A. Figalli, Y.-H. Kim and R. McCann. When is multidimensional screening a convex program? *J. Econ. Theory* 146 (2011).

[FL83] D. Fudenberg and D. Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *J. Econ. Theory* 31 (1983).

[FL89] D. Fudenberg and D. Levine. Reputation and equilibrium selection in games with a patient player. *Econometrica* 57 (1989).

[FM86] D. Fudenberg and E. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54 (1986).

[FM89] J. Farrell and E. Maskin. Renegotiation in repeated games. *Games Econ. Behav.* 1 (1989).

[FS19] B. Fong and D. Spivak. *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge University Press, 2019.

[FT84] D. Fudenberg and J. Tirole. The fat-cat effect, the puppy-dog ploy, and the lean and hungry look. *Amer. Econ. Rev.* 74 (1984).

[FT86] D. Fudenberg and J. Tirole. A "signal-jamming" theory of predation. *RAND J. Econ.* 17 (1986).

[FT90] D. Fudenberg and J. Tirole. Moral hazard and renegotiation in agency contracts. *Econometrica* 58 (1990).

[FT91] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[FT91b] D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *J. Econ. Theory* 53 (1991).

[Fei00] Y. Feinberg. Characterizing common priors in the form of posteriors. *J. Econ. Theory* 91 (2000).

[Fri20] T. Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Adv. Math.* 370 (2020).

[Fri71] J. Friedman. A non-cooperative equilibrium for supergames. *Rev. Econ. Stud.* 38 (1971).

[GH83] S. Grossman and O. Hart. An analysis of the principal-agent problem. *Econometrica* 51 (1983).

[GHWZ18] N. Ghani, J. Hedges, V. Winschel and P. Zahn. Compositional game theory. *Proc. LICS 2018*; arXiv:1603.04641.

[GKLF18] N. Ghani, C. Kupke, A. Lambert and F. Nordvall Forsberg. A compositional treatment of iterated open games. *Theoret. Comput. Sci.* 741 (2018).

[GKZ90] I. Gilboa, E. Kalai and E. Zemel. On the order of eliminating dominated strategies. *Oper. Res. Lett.* 9 (1990).

[GL77] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45 (1977).

[GLR21] É. Goubault, J. Ledent and S. Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task solvability. *Inform. and Comput.* 278 (2021).

[GM85] L. Glosten and P. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financ. Econ.* 14 (1985).

[GP82] J. Geanakoplos and H. Polemarchakis. We can't disagree forever. *J. Econ. Theory* 28 (1982).

[GSW86] F. Gul, H. Sonnenschein and R. Wilson. Foundations of dynamic monopoly and the Coase conjecture. *J. Econ. Theory* 39 (1986).

[GW05] S. Govindan and R. Wilson. Essential equilibria. *Proc. Natl. Acad. Sci. USA* 102 (2005).

[Gha25] N. Ghani. A category theoretic approach to approximate game theory. EPTCS 429 (2025).

[Gib73] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica* 41 (1973).

[Gir82] M. Giry. A categorical approach to probability theory. In *Categorical Aspects of Topology and Analysis*, Lecture Notes in Math. 915, Springer, 1982.

[Gro73] T. Groves. Incentives in teams. *Econometrica* 41 (1973).

[HKR13] M. Herlihy, D. Kozlov and S. Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, 2013.

[HM87] B. Holmström and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica* 55 (1987).

[HM91] B. Holmström and P. Milgrom. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design. *J. Law Econ. Organ.* 7 (1991).

[HMUV09] B. Heydenreich, R. Müller, M. Uetz and R. Vohra. Characterization of revenue equivalence. *Econometrica* 77 (2009).

[HOSWZ17] J. Hedges, P. Oliva, E. Shprits, V. Winschel and P. Zahn. Selection equilibria of higher-order games. *Proc. PADL 2017*, Lecture Notes in Comput. Sci. 10137.

[HS88] J. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.

[Hal10] J. Halpern. Lexicographic probability, conditional probability, and nonstandard probability. *Games Econ. Behav.* 68 (2010).

[Ham94] P. Hammond. Elementary non-Archimedean representations of probability for decision theory and games. In *Patrick Suppes: Scientific Philosopher*, vol. 1, Kluwer, 1994.

[Har67] J. Harsanyi. Games with incomplete information played by "Bayesian" players, I–III. *Management Sci.* 14 (1967–68).

[Har73] J. Harsanyi. Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *Internat. J. Game Theory* 2 (1973).

[Har73b] J. Harsanyi. Oddness of the number of equilibrium points: a new proof. *Internat. J. Game Theory* 2 (1973).

[HeSa12] Z. Hellman and D. Samet. How common are common priors? *Games Econ. Behav.* 74 (2012).

[HeSa98] A. Heifetz and D. Samet. Topology-free typology of beliefs. *J. Econ. Theory* 82 (1998).

[HeSa99] A. Heifetz and D. Samet. Coherent beliefs are not always types. *J. Math. Econ.* 32 (1999).

[Hed16] J. Hedges. *Towards Compositional Game Theory*. PhD thesis, Queen Mary University of London, 2016.

[Hed18] J. Hedges. Morphisms of open games. *Proc. MFPS 2018*, Electron. Notes Theor. Comput. Sci. 341 (2018).

[OGE] CyberCat Institute. From equilibrium checking to learning with the Open Game Engine. Web post, 26 June 2025, https://cybercat.institute/2025/06/26/equilibrium-checking-learning/.

[HoTT13] The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013.

[Hol79] B. Holmström. Moral hazard and observability. *Bell J. Econ.* 10 (1979).

[Hol82] B. Holmström. Moral hazard in teams. *Bell J. Econ.* 13 (1982).

[Hol99] B. Holmström. Managerial incentive problems: a dynamic perspective. *Rev. Econ. Stud.* 66 (1999).

[Hot29] H. Hotelling. Stability in competition. *Econ. J.* 39 (1929).

[JMS99] P. Jehiel, B. Moldovanu and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *J. Econ. Theory* 85 (1999).

[KG11] E. Kamenica and M. Gentzkow. Bayesian persuasion. *Amer. Econ. Rev.* 101 (2011).

[KM86] E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica* 54 (1986).

[KMRW82] D. Kreps, P. Milgrom, J. Roberts and R. Wilson. Rational cooperation in the finitely repeated prisoners' dilemma. *J. Econ. Theory* 27 (1982).

[KR97] E. Kohlberg and P. Reny. Independence on relative probability spaces and consistent assessments in game trees. *J. Econ. Theory* 75 (1997).

[KS75] E. Kalai and M. Smorodinsky. Other solutions to Nash's bargaining problem. *Econometrica* 43 (1975).

[KS83] D. Kreps and J. Scheinkman. Quantity precommitment and Bertrand competition yield Cournot outcomes. *Bell J. Econ.* 14 (1983).

[KS88] E. Kalai and W. Stanford. Finite rationality and interpersonal complexity in repeated games. *Econometrica* 56 (1988).

[KW82a] D. Kreps and R. Wilson. Sequential equilibria. *Econometrica* 50 (1982).

[KW82b] D. Kreps and R. Wilson. Reputation and imperfect information. *J. Econ. Theory* 27 (1982).

[Kim95] S. K. Kim. Efficiency of an information system in an agency model. *Econometrica* 63 (1995).

[Kle98] P. Klemperer. Auctions with almost common values: the "wallet game" and its applications. *Eur. Econ. Rev.* 42 (1998).

[Kno94] V. Knoblauch. Computable strategies for repeated prisoner's dilemma. *Games Econ. Behav.* 7 (1994).

[Kri10] V. Krishna. *Auction Theory*, 2nd ed. Academic Press, 2010.

[Kuh53] H. Kuhn. Extensive games and the problem of information. In *Contributions to the Theory of Games II*, Ann. of Math. Stud. 28, Princeton University Press, 1953.

[Kyl85] A. Kyle. Continuous auctions and insider trading. *Econometrica* 53 (1985).

[LP12] P. Lescanne and M. Perrinel. "Backward" coinduction, Nash equilibrium and the rationality of escalation. *Acta Inform.* 49 (2012).

[LR81] E. Lazear and S. Rosen. Rank-order tournaments as optimum labor contracts. *J. Polit. Econ.* 89 (1981).

[LT86] J.-J. Laffont and J. Tirole. Using cost observation to regulate firms. *J. Polit. Econ.* 94 (1986).

[Law69] F. W. Lawvere. Diagonal arguments and cartesian closed categories. In *Category Theory, Homology Theory and their Applications II*, Lecture Notes in Math. 92, Springer, 1969.

[Law73] F. W. Lawvere. Metric spaces, generalized logic, and closed categories. *Rend. Sem. Mat. Fis. Milano* 43 (1973).

[Liu09] Q. Liu. On redundant types and Bayesian formulation of incomplete information. *J. Econ. Theory* 144 (2009).

[MM92] S. Mac Lane and I. Moerdijk. *Sheaves in Geometry and Logic*. Springer, 1992.

[MMP05] T. Mariotti, M. Meier and M. Piccione. Hierarchies of beliefs for compact possibility models. *J. Math. Econ.* 41 (2005).

[MR78] M. Mussa and S. Rosen. Monopoly and product quality. *J. Econ. Theory* 18 (1978).

[MR82a] P. Milgrom and J. Roberts. Predation, reputation, and entry deterrence. *J. Econ. Theory* 27 (1982).

[MR82b] P. Milgrom and J. Roberts. Limit pricing and entry under incomplete information: an equilibrium analysis. *Econometrica* 50 (1982).

[MR84] E. Maskin and J. Riley. Monopoly with incomplete information. *RAND J. Econ.* 15 (1984).

[MR84b] E. Maskin and J. Riley. Optimal auctions with risk averse buyers. *Econometrica* 52 (1984).

[MR90] P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* 58 (1990).

[MS02] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica* 70 (2002).

[MS06] G. Mailath and L. Samuelson. *Repeated Games and Reputations*. Oxford University Press, 2006.

[MS82] P. Milgrom and N. Stokey. Information, trade and common knowledge. *J. Econ. Theory* 26 (1982).

[MS83] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *J. Econ. Theory* 29 (1983).

[MS89] D. Monderer and D. Samet. Approximating common knowledge with common beliefs. *Games Econ. Behav.* 1 (1989).

[MS94] P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica* 62 (1994).

[MS98] S. Morris and H. S. Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *Amer. Econ. Rev.* 88 (1998).

[MSP73] J. Maynard Smith and G. Price. The logic of animal conflict. *Nature* 246 (1973).

[MSZ13] M. Maschler, E. Solan and S. Zamir. *Game Theory*. Cambridge University Press, 2013.

[MT01] E. Maskin and J. Tirole. Markov perfect equilibrium I: observable actions. *J. Econ. Theory* 100 (2001).

[MT88] E. Maskin and J. Tirole. A theory of dynamic oligopoly, I and II. *Econometrica* 56 (1988).

[MU04] S. Morris and T. Ui. Best response equivalence. *Games Econ. Behav.* 49 (2004).

[MV06] L. Moss and I. Viglizzo. Final coalgebras for functors on measurable spaces. *Inform. and Comput.* 204 (2006).

[MW82] P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica* 50 (1982).

[MWG95] A. Mas-Colell, M. Whinston and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[MZ85] J.-F. Mertens and S. Zamir. Formulation of Bayesian analysis for games with incomplete information. *Internat. J. Game Theory* 14 (1985).

[MaSw97] L. Marx and J. Swinkels. Order independence for iterated weak dominance. *Games Econ. Behav.* 18 (1997).

[McL89a] A. McLennan. The space of conditional systems is a ball. *Internat. J. Game Theory* 18 (1989).

[McL89b] A. McLennan. Consistent conditional systems in noncooperative game theory. *Internat. J. Game Theory* 18 (1989).

[Mor94] S. Morris. Trade with heterogeneous prior beliefs and asymmetric information. *Econometrica* 62 (1994).

[Mye78] R. Myerson. Refinements of the Nash equilibrium concept. *Internat. J. Game Theory* 7 (1978).

[Mye79] R. Myerson. Incentive compatibility and the bargaining problem. *Econometrica* 47 (1979).

[Mye81] R. Myerson. Optimal auction design. *Math. Oper. Res.* 6 (1981).

[Mye86] R. Myerson. Multistage games with communication. *Econometrica* 54 (1986).

[Mye91] R. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

[NGH04] R. Nau, S. Gomez Canovas and P. Hansen. On the geometry of Nash equilibria and correlated equilibria. *Internat. J. Game Theory* 32 (2004).

[NS18] G. Nöldeke and L. Samuelson. The implementation duality. *Econometrica* 86 (2018).

[NZ96] J. Nachbar and W. Zame. Non-computable strategies and discounted repeated games. *Econ. Theory* 8 (1996).

[Nas50a] J. Nash. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* 36 (1950).

[Nas50b] J. Nash. The bargaining problem. *Econometrica* 18 (1950).

[Nas51] J. Nash. Non-cooperative games. *Ann. of Math.* 54 (1951).

[Ney85] A. Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Econ. Lett.* 19 (1985).

[OR90] M. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, 1990.

[OR94] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[Pea84] D. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52 (1984).

[RN09] J. A. Rodrigues-Neto. From posteriors to priors via cycles. *J. Econ. Theory* 144 (2009).

[RS76] M. Rothschild and J. Stiglitz. Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. *Quart. J. Econ.* 90 (1976).

[RS81] J. Riley and W. Samuelson. Optimal auctions. *Amer. Econ. Rev.* 71 (1981).

[RSS20] E. Rijke, M. Shulman and B. Spitters. Modalities in homotopy type theory. *Log. Methods Comput. Sci.* 16(1) (2020).

[Rab57] M. Rabin. Effective computability of winning strategies. In *Contributions to the Theory of Games III*, Ann. of Math. Stud. 39, Princeton University Press, 1957.

[Ras06] E. Rasmusen. *Games and Information: An Introduction to Game Theory*, 4th ed. Blackwell, 2006.

[Ren55] A. Rényi. On a new axiomatic theory of probability. *Acta Math. Acad. Sci. Hungar.* 6 (1955).

[Ril18] M. Riley. Categories of optics. arXiv:1809.00738 (2018).

[Roc87] J.-C. Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *J. Math. Econ.* 16 (1987).

[Rog85] W. Rogerson. The first-order approach to principal-agent problems. *Econometrica* 53 (1985).

[Rub82] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica* 50 (1982).

[Rub86] A. Rubinstein. Finite automata play the repeated prisoner's dilemma. *J. Econ. Theory* 39 (1986).

[Rub89] A. Rubinstein. The electronic mail game: strategic behavior under "almost common knowledge". *Amer. Econ. Rev.* 79 (1989).

[Rut00] J. Rutten. Universal coalgebra: a theory of systems. *Theoret. Comput. Sci.* 249 (2000).

[Rut96] J. Rutten. Elements of generalized ultrametric domain theory. *Theoret. Comput. Sci.* 170 (1996).

[SS84] C. Shapiro and J. Stiglitz. Equilibrium unemployment as a worker discipline device. *Amer. Econ. Rev.* 74 (1984).

[SS84b] A. Shaked and J. Sutton. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica* 52 (1984).

[SY05] M. Saks and L. Yu. Weak monotonicity suffices for truthfulness on convex domains. *Proc. EC 2005*.

[SZ90] L. Simon and W. Zame. Discontinuous games and endogenous sharing rules. *Econometrica* 58 (1990).

[Sam98] D. Samet. Common priors and separation of convex sets. *Games Econ. Behav.* 24 (1998).

[Sch60] T. Schelling. *The Strategy of Conflict*. Harvard University Press, 1960.

[Sel65] R. Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Z. gesamte Staatswiss.* 121 (1965).

[Sel75] R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *Internat. J. Game Theory* 4 (1975).

[Sel78] R. Selten. The chain store paradox. *Theory and Decision* 9 (1978).

[Shu18] M. Shulman. Brouwer's fixed-point theorem in real-cohesive homotopy type theory. *Math. Structures Comput. Sci.* 28 (2018).

[Spe73] M. Spence. Job market signaling. *Quart. J. Econ.* 87 (1973).

[TV23] F. Tohmé and I. Viglizzo. A categorical representation of games. arXiv:2309.15981 (2023).

[Tar55] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* 5 (1955).

[The60] R. Theocharis. On the stability of the Cournot solution on the oligopoly problem. *Rev. Econ. Stud.* 27 (1960).

[Tho52] F. B. Thompson. Equivalence of games in extensive form. RAND Research Memorandum RM-759, 1952.

[Tho79] R. Thomason. Homotopy colimits in the category of small categories. *Math. Proc. Cambridge Philos. Soc.* 85 (1979).

[Tir88] J. Tirole. *The Theory of Industrial Organization*. MIT Press, 1988.

[Top78] D. Topkis. Minimizing a submodular function on a lattice. *Oper. Res.* 26 (1978).

[Top79] D. Topkis. Equilibrium points in nonzero-sum n-person submodular games. *SIAM J. Control Optim.* 17 (1979).

[Top98] D. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.

[Vas92] S. Vassilakis. Some economic applications of Scott domains. *Math. Social Sci.* 24 (1992).

[Vic61] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16 (1961).

[Viv90] X. Vives. Nash equilibrium with strategic complementarities. *J. Math. Econ.* 19 (1990).

[Voh11] R. Vohra. *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, 2011.

[Vor62] N. N. Vorob'ev. Consistent families of measures and their extensions. *Theory Probab. Appl.* 7 (1962).

[WY07] J. Weinstein and M. Yildiz. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica* 75 (2007).

[Wil15] S. Willerton. The Legendre–Fenchel transform from a category theoretic perspective. arXiv:1501.03791 (2015).

[Wil71] R. Wilson. Computing equilibria of N-person games. *SIAM J. Appl. Math.* 21 (1971).

[Wil77] C. Wilson. A model of insurance markets with incomplete information. *J. Econ. Theory* 16 (1977).

[Wil80] C. Wilson. The nature of equilibrium in markets with adverse selection. *Bell J. Econ.* 11 (1980).

[Zho94] L. Zhou. The set of Nash equilibria of a supermodular game is a complete lattice. *Games Econ. Behav.* 7 (1994).

[dGT79] C. d'Aspremont, J. J. Gabszewicz and J.-F. Thisse. On Hotelling's "Stability in competition". *Econometrica* 47 (1979).

[vD84] E. van Damme. A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. *Internat. J. Game Theory* 13 (1984).

[vNM44] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[vS96] B. von Stengel. Efficient computation of behavior strategies. *Games Econ. Behav.* 14 (1996).