

# Lecture Notes on Game Theory and Information Economics \*

William H. Sandholm<sup>†</sup>

September 9, 2019

## Contents

<b>0</b>	<b>Basic Decision Theory</b>	<b>4</b>
0.1	Ordinal Utility . . . . .	4
0.2	Expected Utility and the von Neumann-Morgenstern Theorem . . . . .	5
0.3	Utility for Money and Risk Attitudes . . . . .	9
0.4	Bayesian Rationality . . . . .	10
<b>1</b>	<b>Normal Form Games</b>	<b>11</b>
1.1	Basic Concepts . . . . .	12
1.1.1	Definition . . . . .	12
1.1.2	Randomized strategies . . . . .	12
1.1.3	Conjectures . . . . .	16
1.1.4	Expected utility . . . . .	18
1.2	Dominance and Iterated Dominance . . . . .	19
1.2.1	Strictly dominant strategies . . . . .	19
1.2.2	Strictly dominated strategies . . . . .	21
1.2.3	Iterated strict dominance . . . . .	23
1.2.4	Weak dominance . . . . .	24
1.3	Rationalizability . . . . .	26
1.3.1	Definition and examples . . . . .	26
1.3.2	Rationalizability and iterated strict dominance . . . . .	31
1.3.3	Rationalizability with three or more players . . . . .	35
1.3.4	A positive characterization of rationalizability . . . . .	37

---

\*Many thanks to Katsuhiko Aiba, Srinivas Arigapudi, Jose Antonio Carrasco Novoa, Emin Dokumacı, Danqing Hu, Rui Li, Allen Long, Ignacio Monzón, Michael Rapp, and Ryoji Sawa for creating the initial draft of this document from my handwritten notes, and to George Loginov and Antonio Penta for their comments.

<sup>†</sup>Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA.  
e-mail: whs@ssc.wisc.edu; website: <http://www.ssc.wisc.edu/~whs>.

1.4	Nash Equilibrium . . . . .	38
1.4.1	Definition . . . . .	38
1.4.2	Computing Nash equilibria . . . . .	40
1.4.3	Existence of Nash equilibrium and structure of the equilibrium set . . . . .	45
1.4.4	Interpretations of Nash equilibrium . . . . .	47
1.4.5	Nash equilibrium in games with continuous strategy sets . . . . .	51
1.5	Correlated Equilibrium . . . . .	54
1.5.1	Definition and examples . . . . .	54
1.5.2	Interpretation . . . . .	58
1.6	The Minmax Theorem . . . . .	60
<b>2</b>	<b>Extensive Form Games</b>	<b>68</b>
2.1	Basic Concepts . . . . .	68
2.1.1	Defining extensive form games . . . . .	68
2.1.2	Pure strategies in extensive form games . . . . .	73
2.1.3	Randomized strategies in extensive form games . . . . .	74
2.1.4	Reduced normal form games . . . . .	78
2.2	The Principle of Sequential Rationality . . . . .	80
2.3	Games of Perfect Information and Backward Induction . . . . .	81
2.3.1	Subgame perfect equilibrium, sequential rationality, one-shot deviations, and backward induction . . . . .	81
2.3.2	Epistemic foundations for backward induction . . . . .	89
2.3.3	Applications . . . . .	93
2.3.4	Subgame perfect equilibrium in more general classes of games . . . . .	96
2.4	Games of Imperfect Information and Sequential Equilibrium . . . . .	98
2.4.1	Subgames and subgame perfection in games of imperfect information . . . . .	98
2.4.2	Beliefs and sequential rationality . . . . .	100
2.4.3	Definition of sequential equilibrium . . . . .	103
2.4.4	Computing sequential equilibria . . . . .	107
2.4.5	Existence of sequential equilibrium and structure of the equilibrium set . . . . .	116
2.5	Forward Induction . . . . .	116
2.5.1	Motivation and discussion . . . . .	117
2.5.2	Forward induction in signaling games . . . . .	119
2.6	Invariance and Proper Equilibrium . . . . .	127
2.7	Full Invariance and Kohlberg-Mertens Stability . . . . .	134
<b>3</b>	<b>Repeated Games</b>	<b>138</b>
3.1	Basic Concepts . . . . .	139
3.2	Stick-and-Carrot Strategies and the Folk Theorem . . . . .	144
3.3	Computing the Set of Subgame Perfect Equilibrium Payoffs . . . . .	153
3.3.1	Dynamic programming . . . . .	153
3.3.2	Dynamic programs vs. repeated games . . . . .	155
3.3.3	Factorization and self-generation . . . . .	155

<b>4</b>	<b>Bayesian Games</b>	<b>164</b>
4.1	Basic Definitions . . . . .	164
4.2	Bayesian Strategies, Dominance, and Bayesian Equilibrium . . . . .	167
4.3	Computing Bayesian Equilibria . . . . .	170
4.4	More on Interpretation . . . . .	173
4.5	Applications . . . . .	175
4.5.1	Auctions . . . . .	175
4.5.2	The revelation principle for normal form games . . . . .	179
4.6	Ex Post Equilibrium . . . . .	181
4.7	Correlated Types and Higher-Order Beliefs . . . . .	183
4.8	Incomplete Information, Hierarchies of Beliefs, and Bayesian Games . . . . .	188
4.9	Rationalizability in Bayesian Games . . . . .	192
	<b>Interlude: Information Economics</b>	<b>195</b>
<b>5</b>	<b>Signalling and Screening with Competing Uninformed Players</b>	<b>197</b>
5.1	The Market for Lemons (Akerlof (1970)) . . . . .	197
5.2	Job Market Signalling (Spence (1973)) . . . . .	199
5.3	Screening in Insurance Markets (Rothschild and Stiglitz (1976)) . . . . .	206
<b>6</b>	<b>Principal-Agent Models</b>	<b>213</b>
6.1	Moral Hazard . . . . .	213
6.1.1	Two effort levels, two output levels . . . . .	214
6.1.2	Finite sets of effort and output levels . . . . .	216
6.2	Adverse Selection (Monopolistic Screening) . . . . .	220
6.2.1	Two types . . . . .	221
6.2.2	A continuum of types . . . . .	225
<b>7</b>	<b>Mechanism Design</b>	<b>233</b>
7.1	Mechanism Design Environments and the Revelation Principle . . . . .	234
7.1.1	Mechanism design environments and social choice functions . . . . .	234
7.1.2	Mechanisms . . . . .	236
7.1.3	Implementation . . . . .	238
7.1.4	Direct mechanisms and the revelation principle . . . . .	239
7.2	Incentive Compatibility, Payoff Equivalence, and Revenue Equivalence . . . . .	243
7.3	Revenue Maximization and Optimal Auctions . . . . .	249
7.4	Allocative Efficiency . . . . .	254
7.4.1	Groves mechanisms and the VCG mechanism . . . . .	255
7.4.2	Budget balance and the AGV mechanism . . . . .	258
7.4.3	The KPW mechanism . . . . .	261
7.4.4	The Krishna-Perry-Williams and Myerson-Satterthwaite theorems . . . . .	269
7.5	Allocative Efficiency with Interdependent Values . . . . .	273
7.6	Correlated Types, Full Surplus Extraction, and the Wilson Doctrine . . . . .	276
7.6.1	Full surplus extraction via the Crémer-McLean mechanism . . . . .	276

7.6.2	Full surplus extraction via menus of lotteries . . . . .	279
7.6.3	The beliefs-determine-preferences property and the Wilson doctrine	280
7.6.4	Robust mechanism design . . . . .	281

<b>References</b>		<b>282</b>
-------------------	--	------------

## 0. Basic Decision Theory

### 0.1 Ordinal Utility

We consider a decision maker (or agent) who chooses among *outcomes* in some set  $Z$ . To begin we assume that  $Z$  is finite.

The primitive description of preferences is in terms of a *preference relation*  $\succeq$ . For any ordered pair of outcomes  $(x, y) \in Z \times Z$ , the agent can tell us whether or not he *weakly prefers*  $x$  to  $y$ . If yes, we write  $x \succeq y$ . If no, we write  $x \not\succeq y$ .

We can use these to define

*strict preference*:  $a \succ b$  means  $[a \succeq b \text{ and } b \not\succeq a]$ .

*indifference*:  $a \sim b$  means  $[a \succeq b \text{ and } b \succeq a]$ .

We say that the preference relation  $\succeq$  is a *weak order* if it satisfies the two *weak order axioms*:

*completeness*: For all  $a, b \in Z$ , either  $a \succeq b$  or  $b \succeq a$  (or both).

*transitivity*: For all  $a, b, c \in Z$ , if  $a \succeq b$  and  $b \succeq c$ , then  $a \succeq c$ .

Completeness says that there are no outcomes that the agent is unwilling or unable to compare. (Consider  $Z = \{\text{do nothing, kill someone to save five others' lives}\}$ .)

Transitivity rules out preference cycles. (Consider  $Z = \{\text{a scoop of ice-cream, an enormous hunk of chocolate cake, a small plain salad}\}$ .)

The function  $u: Z \rightarrow \mathbb{R}$  is an *ordinal utility function* that *represents*  $\succeq$  if

$$u(a) \geq u(b) \text{ if and only if } a \succeq b.$$

**Theorem 0.1.** *Let  $Z$  be finite and let  $\succeq$  be a preference relation. Then there is an ordinal utility function  $u: Z \rightarrow \mathbb{R}$  that represents  $\succeq$  if and only if  $\succeq$  is complete and transitive.*

*Moreover, the function  $u$  is unique up to increasing transformations:  $v: Z \rightarrow \mathbb{R}$  also represents  $\succeq$  if and only if  $v = f \circ u$  for some increasing function  $f: \mathbb{R} \rightarrow \mathbb{R}$ .*

In the first part of the theorem, the “only if” direction follows immediately from the fact that the real numbers are ordered. For the “if” direction, assign the elements of  $Z$  utility values sequentially; the weak order axioms ensure that this can be done without contradiction.

“Ordinal” refers to the fact that only the order of the values of the utility function have meaning. Neither the values nor differences between them convey information about intensity of preferences. This is captured by the second part of the theorem, which says that utility functions are only unique up to increasing transformations.

If  $Z$  is (uncountably) infinite, weak order is not enough to ensure that there is an ordinal utility representation:

*Example 0.2. Lexicographic preferences.* Let  $Z = \mathbb{R}^2$ , and suppose that  $a \geq b \Leftrightarrow a_1 > b_1$  or [ $a_1 = b_1$  and  $a_2 \geq b_2$ ]. In other words, the agent’s first priority is the first component of the outcome; he only uses the second component to break ties. While  $\geq$  satisfies the weak order axioms, it can be shown that there is no ordinal utility function that represents  $\geq$ . In essence, there are too many levels of preference to fit them all into the real line.  $\diamond$

There are various additional assumptions that rule out such examples. One is

*Continuity:*  $Z \subseteq \mathbb{R}^n$ , and for every  $a \in Z$ , the sets  $\{b: b \geq a\}$  and  $\{b: a \geq b\}$  are closed.

Notice that Example 0.2 violates this axiom.

**Theorem 0.3.** *Let  $Z \subseteq \mathbb{R}^n$  and let  $\geq$  be a preference relation. Then there is a continuous ordinal utility function  $u: Z \rightarrow \mathbb{R}$  that represents  $\geq$  if and only if  $\geq$  is complete, transitive, and continuous.*

In the next section we consider preferences over lotteries—probability distributions over a finite set of outcomes. Theorem 0.3 ensures that if preferences satisfy the weak order and continuity axioms, then they can be represented by a continuous ordinal utility function. By introducing an additional axiom, one can obtain a more discriminating representation.

## 0.2 Expected Utility and the von Neumann-Morgenstern Theorem

Now we consider preferences in settings with uncertainty: an agent chooses among “lotteries” in which different outcomes in  $Z$  have different probabilities of being realized.

*Example 0.4.* Suppose you are offered a choice between

lottery 1:      \$1M for sure

$$\text{lottery 2: } \begin{cases} \$2\text{M with probability } \frac{1}{2}, \\ \$0 \text{ with probability } \frac{1}{2}. \end{cases}$$

One tempting possibility is to look at expected values: the weighted averages of the possible values, with weights given by probabilities.

$$\text{lottery 1: } \$1\text{M} \times 1 = \$1\text{M}$$

$$\text{lottery 2: } \$2\text{M} \times \frac{1}{2} + \$0\text{M} \times \frac{1}{2} = \$1\text{M}$$


But most people strictly prefer lottery 1.

The lesson: if outcomes are in dollars, ranking outcomes in terms of expected numbers of dollars may not capture preferences.  $\diamond$


Let  $Z$  be a finite set of outcomes. We denote by  $\Delta Z$  the set of probability distributions over  $Z$ :  $\Delta Z = \{p: Z \rightarrow \mathbb{R}_+ \mid \sum_{z \in Z} p(z) = 1\}$ . Here we interpret  $p \in \Delta Z$  as a *lottery over outcomes* from  $Z$ .

*Example 0.5.*  $Z = \{\$0, \$10, \$100\}$ ,


$$\Delta Z = \{p = (p(\$0), p(\$10), p(\$100))\}$$



$$p = (.2, .8, 0)$$



$$q = (.9, 0, .1)$$



$$r = (0, 0, 1) \quad \diamond$$


An agent has preferences  $\geq$  over  $\Delta Z$ , where “ $p \geq q$ ” means that he likes  $p$  at least as much as  $q$ .

The preference relation  $\geq$  on  $\Delta Z$  admits an *expected utility representation* if there is a function  $u: Z \rightarrow \mathbb{R}$  such that

$$(1) \quad p \geq q \Leftrightarrow \sum_{z \in Z} u(z) p(z) \geq \sum_{z \in Z} u(z) q(z).$$

The function  $u$  is called a *Bernoulli utility function*. The expectations of  $u$  with respect to  $p$  and  $q$  appearing on the right in (1) are called *von Neumann-Morgenstern* (or *NM*) *expected utility functions*.

When can a preference relation on  $\Delta Z$  be represented using expected utilities? To answer this question, we introduce *compound lotteries*, or “lotteries over lotteries”. An example of a compound lottery is shown at left in the figure below. We assume that the agent only cares about the ultimate probabilities of obtaining each outcome in  $Z$ , as shown at right in the figure. The lottery at right is called the *reduced lottery* corresponding to the compound lottery.



We can write a compound lottery with second-round lotteries  $p$  and  $q$  as  $\alpha p + (1 - \alpha)q$ , where  $\alpha \in [0, 1]$ . By the assumption above, we identify this lottery with the reduced lottery in which the probability of outcome  $z$  is  $\alpha p(z) + (1 - \alpha)q(z)$ . That is, we identify the compound lottery with the reduced lottery defined by the appropriate linear combination of probability vectors  $p$  and  $q$ . This captures the reduction from the figure above correctly.

*Example 0.6.* In the figure above, the compound lottery  $.7p + .3q$  is identified with the single-stage lottery with outcome probabilities  $.7(.2, .8, 0) + .3(.9, 0, .1) = (.41, .56, .03)$ .  $\diamond$

Let  $\geq$  be a *preference relation* on  $\Delta Z$ , where  $p \geq q$  means that lottery  $p \in \Delta Z$  is weakly preferred to lottery  $q \in \Delta Z$ . Below are the three *von Neumann-Morgenstern* (or *NM*) *axioms*. Notice that (NM2) and (NM3) are stated in terms of compound lotteries.

(NM1) *Weak order:*  $\geq$  is complete and transitive.

(NM2) *Continuity:* For all  $p, q$ , and  $r$  such that  $p > q > r$ , there exist  $\delta, \varepsilon \in (0, 1)$  such that

$$(1 - \delta)p + \delta r > q > (1 - \varepsilon)r + \varepsilon p.$$

The continuity axiom says that there is no outcome so good or so bad that having any probability of receiving it dominates all other considerations.


*Example 0.7.*  $p$  = win Nobel prize for sure,  $q$  = nothing,  $r$  = get hit by a bus for sure.  $\diamond$

(NM3) *Independence:* For all  $p, q$ , and  $r$  and all  $\alpha \in (0, 1)$ ,

$$p \geq q \Leftrightarrow \alpha p + (1 - \alpha)r \geq \alpha q + (1 - \alpha)r.$$

The independence axiom says that introducing the same probability of a third lottery  $r$  to lotteries  $p$  and  $q$  does not reverse preferences between  $p$  and  $q$ .

*Example 0.8.* Let  $p = (.2, .8, 0)$ ,  $q = (.9, 0, .1)$ ,  $r = (0, 0, 1)$ . Independence implies that if  $p$  is preferred to  $q$ , then  $\hat{p} = .5p + .5r = (.1, .4, .5)$  is preferred to  $\hat{q} = .5q + .5r = (.45, 0, .55)$ .  $\diamond$



**Theorem 0.9** (von Neumann and Morgenstern (1944)).

Let  $Z$  be a finite set, and let  $\succeq$  be a preference relation on  $\Delta Z$ . Then there is a Bernoulli utility function  $u: Z \rightarrow \mathbb{R}$  that provides an expected utility representation for  $\succeq$  if and only if  $\succeq$  satisfies (NM1), (NM2), and (NM3).

Moreover, the function  $u$  is unique up to positive affine transformations. That is,  $v$  also represents  $\succeq$  if and only if  $v \equiv au + b$  for some  $a > 0$  and  $b \in \mathbb{R}$ .

Discussion of Theorem 0.9

- (i) The theorem tells us that as long as (NM1)–(NM3) hold, there is some way of assigning numbers to the outcomes such that taking expected values of these numbers is the right way to evaluate lotteries over outcomes.
- (ii) The values of a Bernoulli utility function are sometimes called *cardinal utilities* (as opposed to ordinal utilities). What more-than-ordinal information do cardinal utilities provide?

The nature of this information can be deduced from the fact that a Bernoulli utility function is unique up to positive affine transformations.

*Example 0.10.* Let  $a, b, c \in Z$ , and suppose that  $u_a > u_c > u_b$ . Let  $\lambda = \frac{u_c - u_b}{u_a - u_b}$ . This quantity is not affected by positive affine transformations. Indeed, if  $v = \alpha u + \beta$ , then

$$\frac{v_c - v_b}{v_a - v_b} = \frac{(\alpha u_c + \beta) - (\alpha u_b + \beta)}{(\alpha u_a + \beta) - (\alpha u_b + \beta)} = \frac{\alpha(u_c - u_b)}{\alpha(u_a - u_b)} = \lambda.$$

To interpret  $\lambda$ , rearrange its definition to obtain

$$u_c = \lambda u_a + (1 - \lambda)u_b.$$



This says that  $\lambda$  is the probability on  $a$  in a lottery over  $a$  and  $b$  that makes this lottery exactly as good as getting  $c$  for sure.  $\diamond$

(iii) While expected utility seems reasonable, it is not innocuous.

*Example 0.11. The Allais (1953) paradox.*

Consider the following lotteries:

$p$	\$1M for sure	$r$	\$1M with probability .11 \$0M with probability .89
$q$	\$5M with probability .10 \$1M with probability .89 \$0M with probability .01	$s$	\$5M with probability .10 \$0M with probability .90

Many people express preferences  $p > q$  and  $s > r$ . These preferences cannot be captured by expected utility!

Why not?  $p > q$  implies that

$$1u(\$1M) > .10 u(\$5M) + .89 u(\$1M) + .01 u(\$0)$$

Now add  $.89(u(\$0) - u(\$1M))$  to each side:

$$.11 u(\$1M) + .89 u(\$0) > .10 u(\$5M) + .90 u(\$0),$$

Thus  $r > s$ .  $\diamond$

The problem in Example 0.11 is that the preferences  $p > q$  and  $s > r$  include a greater sensitivity to low probability events than NM expected utility allows. For alternative models that address this and other problematic examples, see Kahneman and Tversky (1979) and Gilboa (2009).

### 0.3 Utility for Money and Risk Attitudes

Suppose that the outcomes are amounts of money measured in dollars (as in our example of \$1M for sure vs. a 50% chance of \$2M). Then for any lottery  $p \in \Delta Z$ , we can compute the expected dollar return

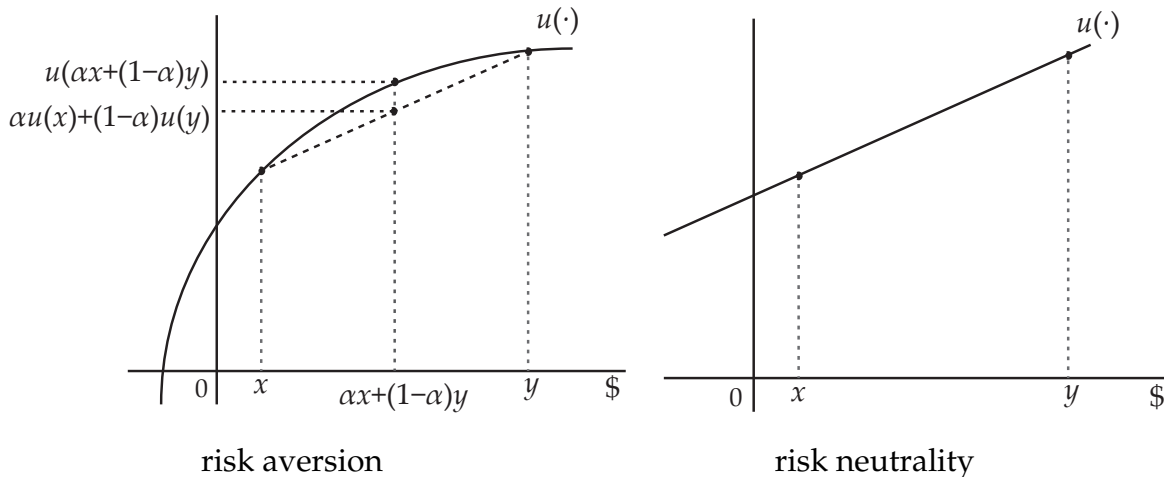
$$\bar{p} = \sum_{z \in Z} z p(z).$$

We say that an agent is *risk averse* if she always weakly prefers getting  $\bar{p}$  for sure to the lottery  $p$ , *risk loving* if she always prefers  $p$  to getting  $\bar{p}$  for sure, and *risk neutral* if she is always indifferent.

Risk preferences can be read from the shape of an agent's utility function  $u: \mathbb{R} \rightarrow \mathbb{R}$  from dollar outcomes to units of utility. We generally assume that  $u$  is *increasing*:  $x < y \Rightarrow u(x) < u(y)$ . If  $u$  is differentiable, this implies that  $u'(x) \geq 0$  for all  $x \in \mathbb{R}$ .

Consider a lottery  $p$  which yields  $\$x$  with probability  $\alpha$  and  $\$y$  with probability  $1 - \alpha$ . The expected value of this lottery is  $\$\bar{p} = \$(\alpha x + (1 - \alpha)y)$ .

$u$ is	if for all $x, y \in \mathbb{R}$ and $\alpha \in [0, 1]$	if $u''$ exists	risk preference
<i>concave</i>	$\alpha u(x) + (1 - \alpha)u(y) \leq u(\alpha x + (1 - \alpha)y)$	$u'' < 0$	risk aversion
<i>affine</i>	$\alpha u(x) + (1 - \alpha)u(y) = u(\alpha x + (1 - \alpha)y)$	$u'' = 0$	risk neutrality
<i>convex</i>	$\alpha u(x) + (1 - \alpha)u(y) \geq u(\alpha x + (1 - \alpha)y)$	$u'' > 0$	risk loving



What about lotteries with more than two outcomes?

**Theorem 0.12 (Jensen's inequality).** Suppose  $u: \mathbb{R} \rightarrow \mathbb{R}$  is concave. If  $x_1, \dots, x_n \in \mathbb{R}$ , and if  $\alpha_1, \dots, \alpha_n \geq 0$  satisfy  $\sum_{i=1}^n \alpha_i = 1$ , then  $\sum_i \alpha_i u(x_i) \leq u(\sum_i \alpha_i x_i)$ .

For  $n = 2$  this is the definition of concavity. For  $n > 2$  this is proved by induction.

## 0.4 Bayesian Rationality

In settings with uncertainty, where all relevant probabilities are objective and known, we call an agent *NM rational* if he acts as if he is maximizing a NM expected utility function.

What if the probabilities are not given? We call an agent *Bayesian rational* (or say that he has *subjective expected utility preferences*) if

- (i) In settings with uncertainty, he forms beliefs about the probabilities of all relevant events.
- (ii) When making decisions, he acts to maximize his expected utility given his beliefs.
- (iii) After receiving new information, he updates his beliefs by taking conditional probabilities whenever possible.

In game theory, it is standard to begin analyses with the assumption that players are Bayesian rational.

Foundations for subjective expected utility preferences are obtained from *state-space models* of uncertainty. These models begin with a set of possible states whose probabilities are not given, and consider preferences over maps from states to outcomes. Savage (1954) provides an axiomatization of subjective expected utility preferences in this framework. Both the utility function and the assignment of probabilities to states are determined as part of the representation. Anscombe and Aumann (1963) consider a state-space model in which preferences are not over state-contingent outcomes, but over maps from states to lotteries à la von Neumann-Morgenstern. This formulation allows for a much simpler derivation of subjective expected utility preferences, and fits very naturally into game-theoretic models. See Gilboa (2009) for a textbook treatment of these and more general models of decision under uncertainty.

## 1. Normal Form Games

*Game theory* models situations in which multiple players make strategically interdependent decisions. *Strategic interdependence* means that your outcomes depend both on what you do and on what others do.

This course focuses on *noncooperative game theory*, which works from the hypothesis that agents act independently, each in his own self interest. *Cooperative game theory* studies situations in which subsets of the agents can make binding agreements.

We study some basic varieties of games and the connections among them:

1. *Normal form games*: moves are simultaneous
2. *Extensive form games*: moves take place over time
3. *Repeated games*: a normal form game is played repeatedly, with all previous moves being observed before each round of play
4. *Bayesian games*: players receive private information before play begins

## 1.1 Basic Concepts

*Example: 1.1. Game show!*

Two contestants on a game show are separated by a partition. In front of them is a table holding \$60,000 in cash. Each contestant chooses between two options,  $D$  (dash) and  $C$  (chill). If both contestants dash, they split the \$60,000. If only one dashes, he gets the entire \$60,000 and his opponent gets nothing. If neither contestant dashes, the host increases the total amount of cash to \$100,000, and the contestants split this.

		2	
		$c$	$d$
1	$C$	5,5	0,6
	$D$	6,0	3,3

- (i) Players:  $\mathcal{P} = \{1, 2\}$
- (ii) Pure strategy sets  $S_1 = \{C, D\}$ ,  $S_2 = \{c, d\}$   
Set of pure strategy profiles:  $S = S_1 \times S_2$ . For example:  $(C, d) \in S$
- (iii) Utility functions  $u_i : S \rightarrow \mathbb{R}$ . For example:  $u_1(C, d) = 0$ ,  $u_2(C, d) = 6$ .

This game is an instance of a *Prisoner's Dilemma*, and the traditional names for the strategies are Cooperate and Defect; see the remarks following Example 1.9.  $\diamond$

### 1.1.1 Definition

A *normal form game*  $G = \{\mathcal{P}, \{S_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$  consists of:

- (i) a finite set of *players*  $\mathcal{P} = \{1, \dots, n\}$ ,
- (ii) a finite set of *pure strategies*  $S_i$  for each player,
- (iii) a *Bernoulli utility function*  $u_i : S \rightarrow \mathbb{R}$  for each player, where  $S = \prod_{i \in \mathcal{P}} S_i$  is the set of *pure strategy profiles* (lists of strategies, one for each player).

If each player chooses some  $s_i \in S_i$ , the strategy profile is  $s = (s_1, \dots, s_n)$  and player  $j$ 's payoff is  $u_j(s)$ .

### 1.1.2 Randomized strategies

In our description of a game above, players each choose a particular pure strategy  $s_i \in S_i$ . But it is often worth considering the possibility that each player makes a randomized choice.

**Mixed strategies and mixed strategy profiles**

If  $A$  is a finite set, then we let  $\Delta A$  represent the set of probability distributions over  $A$ : that is,  $\Delta A = \{p : A \rightarrow \mathbb{R}_+ \mid \sum_{a \in A} p(a) = 1\}$ .

Then  $\sigma_i \in \Delta S_i$  is a *mixed strategy* for  $i$ , while  $\sigma = (\sigma_1, \dots, \sigma_n) \in \prod_{i \in \mathcal{P}} \Delta S_i$  is a *mixed strategy profile*.

Under a mixed strategy profile, players are assumed to randomize *independently*: for instance, learning that 1 played  $C$  provides no information about what 2 did. In other words, the distribution on the set  $S$  of pure strategy profiles created by  $\sigma$  is a *product distribution*.

*Example 1.2. Battle of the Sexes.*

		2	
		$a$	$b$
1	$A$	3, 1	0, 0
	$B$	0, 0	1, 3


Suppose that 1 plays  $A$  with probability  $\frac{3}{4}$ , and 2 plays  $a$  with probability  $\frac{1}{4}$ . Then

$$\sigma = (\sigma_1, \sigma_2) = ((\sigma_1(A), \sigma_1(B)), (\sigma_2(a), \sigma_2(b))) = ((\frac{3}{4}, \frac{1}{4}), (\frac{1}{4}, \frac{3}{4}))$$


The pure strategy profile  $(A, a)$  is played with probability  $\sigma_1(A) \cdot \sigma_2(a) = \frac{3}{4} \cdot \frac{1}{4} = \frac{3}{16}$ . The complete product distribution is presented in the matrix below.

		2		
		$a (\frac{1}{4})$	$b (\frac{3}{4})$	
1	$A (\frac{3}{4})$	$\frac{3}{16}$	$\frac{9}{16}$	$\diamond$
	$B (\frac{1}{4})$	$\frac{1}{16}$	$\frac{3}{16}$	


When player  $i$  has two strategies, his set of mixed strategies  $\Delta S_i$  is the simplex in  $\mathbb{R}^2$ , which is an interval.



When player  $i$  has three strategies, his set of mixed strategies  $\Delta S_i$  is the simplex in  $\mathbb{R}^3$ , which is a triangle.



When player  $i$  has four strategies, his set of mixed strategies  $\Delta S_i$  is the simplex in  $\mathbb{R}^4$ , which is a pyramid.



(use imagination here)

### *Correlated strategies*

It is sometimes useful to introduce the possibility that agents can randomize in a correlated way. This is possible if each player observes a signal that is correlated with the signals observed by others, and conditions his action choice on the signal he observes. Such signals can be introduced artificially, but they may also appear naturally in the environment where the players play the game—see Sections 1.5.2 and 4.5.2.

*Example 1.3. Battle of the Sexes revisited.*

Suppose that the players observe a toss of a fair coin. If the outcome is Heads, they play  $(A, a)$ ; if it is Tails, they play  $(B, b)$ .

A formal description of their behavior specifies the probability of each pure strategy profile:  $\rho = (\rho(A, a), \rho(A, b), \rho(B, a), \rho(B, b)) = (\frac{1}{2}, 0, 0, \frac{1}{2})$ .

		2	
		a	b
1	A	1/2	0
	B	0	1/2

This behavior cannot be achieved using a mixed strategy profile, since it requires correlation: any mixed strategy profile putting weight on  $(A, a)$  and  $(B, b)$  would also put weight on  $(A, b)$  and  $(B, a)$ :

		2	
		$y > 0$	$(1 - y) > 0$
		a	b
1	$x > 0$	A	$xy$
	$(1 - x) > 0$	B	$x(1 - y)$
			$(1 - x)y$
			$(1 - x)(1 - y)$

all marginal probabilities  $> 0 \Rightarrow$  all joint probabilities  $> 0 \quad \diamond$

We call  $\rho \in \Delta\left(\prod_{i \in \mathcal{P}} S_i\right) = \Delta S$  a *correlated strategy*. It is an arbitrary joint distribution on  $\prod_{i \in \mathcal{P}} S_i$ .

*Example 1.4.* Suppose that  $\mathcal{P} = \{1, 2, 3\}$  and  $S_i = \{1, \dots, k_i\}$ . Then a mixed strategy profile  $\sigma = (\sigma_1, \sigma_2, \sigma_3) \in \prod_{i \in \mathcal{P}} \Delta S_i$  consists of three probability vectors of lengths  $k_1, k_2,$  and  $k_3$ , while a correlated strategy  $\rho \in \Delta\left(\prod_{i \in \mathcal{P}} S_i\right)$  is a single probability vector of length  $k_1 \cdot k_2 \cdot k_3$ .  $\diamond$

Because players randomize independently in mixed strategy profiles, mixed strategy profiles generate the product distributions on  $\prod_{i \in \mathcal{P}} S_i$  (rather than all joint distributions on  $\prod_{i \in \mathcal{P}} S_i$ ). Thus:


$$\text{mixed strategy profiles} = \prod_{i \in \mathcal{P}} \Delta S_i \text{ "c" } \Delta \left( \prod_{i \in \mathcal{P}} S_i \right) = \text{correlated strategies.}$$

We write "c" because the sets on each side are subsets of different spaces. A more precise statement is that the set of correlated strategies generated by mixed strategy profiles  $\sigma$  via


$$\rho^\sigma(s) = \prod_{i \in \mathcal{P}} \sigma_i(s_i)$$

is strictly contained in the set of all correlated strategies. These two points are illustrated in the next example.

*Example 1.5.* If  $S_1 = \{A, B\}$  and  $S_2 = \{a, b\}$ , then the set of mixed strategy profiles  $\Delta\{A, B\} \times \Delta\{a, b\}$  is the product of two intervals, and hence a square. The set of correlated strategies  $\Delta(\{A, B\} \times \{a, b\})$  is a pyramid.



The set of correlated strategies that correspond to mixed strategy profiles—in other words, the product distributions on  $\{A, B\} \times \{a, b\}$ —forms a surface in the pyramid. Specifically, this surface consists of the correlated strategies satisfying  $\rho(Aa) = (\rho(Aa) + \rho(Ab)) \cdot (\rho(Aa) + \rho(Ba))$ . (Showing that this equation characterizes the product distributions is a good exercise).



### 1.1.3 Conjectures

One can divide traditional game-theoretic analyses into two classes: equilibrium and non-equilibrium. In equilibrium analyses (e.g., using Nash equilibrium), one assumes that players correctly anticipate how opponents will act. In this case, a Bayesian rational player will maximize his expected utility with respect to a correct prediction about opponents' strategies. In nonequilibrium analyses (e.g., dominance arguments) correct predictions are not assumed. Instead, Bayesian rationality requires players to form probabilistic conjectures about how their opponents will act, and to maximize their expected payoffs given their conjectures.



In a two-player game, a *full conjecture* of player  $i$  is a probability distribution  $v_i$  over player  $j$ 's mixed strategies. We can write this loosely as  $v_i \in \Delta(\Delta S_j)$ . (This is loose because  $\Delta S_j$  is an infinite set.) Player  $i$ 's conjecture is thus a “compound probability distribution” over  $S_j$ : he assigns probabilities to mixed strategies of player  $j$ , which themselves assign probabilities to pure strategies in  $S_j$ .

Since player  $i$  is an expected utility maximizer, he only cares about the “reduced probabilities” over  $S_j$ . These are represented by a *reduced conjecture*  $\mu_i \in \Delta S_j$ . When  $v_i$  has finite support, the reduced conjecture  $\mu_i$  it generates is

$$\mu_i(s_j) = \sum_{\sigma_j} v_i(\sigma_j) \sigma_j(s_j).$$

(An infinite support would require an integral.) Notice that a reduced conjecture is the same kind of object as a mixed strategy of player  $j$ : both are probability measures on  $S_j$ .

We call two full conjectures *equivalent* if they generate the same reduced conjecture.

*Example 1.6.* Let  $S_2 = \{L, R\}$ , let  $\sigma_2 = (\sigma_2(L), \sigma_2(R)) = (\frac{1}{2}, \frac{1}{2})$ , let  $\hat{\sigma}_2 = (\frac{1}{3}, \frac{2}{3})$ , and suppose that player 1's full conjecture is  $v_1(\sigma_2) = \frac{1}{4}$ ,  $v_1(\hat{\sigma}_2) = \frac{3}{4}$ . The corresponding reduced conjecture is

$$\begin{aligned} \mu_1(L) &= v_1(\sigma_2) \sigma_2(L) + v_1(\hat{\sigma}_2) \hat{\sigma}_2(L) = \frac{1}{4} \cdot \frac{1}{2} + \frac{3}{4} \cdot \frac{1}{3} = \frac{3}{8}, \\ \mu_1(R) &= v_1(\sigma_2) \sigma_2(R) + v_1(\hat{\sigma}_2) \hat{\sigma}_2(R) = \frac{1}{4} \cdot \frac{1}{2} + \frac{3}{4} \cdot \frac{2}{3} = \frac{5}{8}. \end{aligned}$$

Two other full conjectures equivalent to  $v_1$  are

$$\begin{aligned} \tilde{v}_1(\sigma_2^L) &= \frac{3}{8}, \tilde{v}_1(\sigma_2^R) = \frac{5}{8}, \text{ where } \sigma_2^L(L) = \sigma_2^R(R) = 1; \\ \tilde{v}_1(\bar{\sigma}_2) &= 1, \text{ where } \bar{\sigma}_2 = (\frac{3}{8}, \frac{5}{8}). \end{aligned}$$

Under  $\tilde{v}_1$ , player 1 isn't sure of which pure strategy player 2 is playing, but he is sure that the strategy is pure. Under  $\bar{v}_1$ , player 1 is certain that player 2 will play mixed strategy  $\bar{\sigma}_2$ . The final probabilities on  $L$  and  $R$  are the same in both cases. (Aside: The specification of  $\tilde{v}_1$  above is rigorous, but we would usually just write  $\tilde{v}_1(L) = \frac{3}{8}$  and  $\tilde{v}_1(R) = \frac{5}{8}$ .)  $\diamond$

In our nonequilibrium analyses in Sections 1.2 and 1.3, it will be easier to work directly with reduced conjectures, which we will call *conjectures* for short. An exception will occur in Example 1.16.

The only novelty that arises in games with three or more players is that player  $i$ 's conjecture may reflect correlation in opponents' choices—for instance, through the observation of

signals that player  $i$  himself does not see. Here a (*reduced*) *conjecture*  $\mu_i \in \Delta(\prod_{j \neq i} S_j)$  is a probability measure on pure strategy profiles of  $i$ 's opponents. It is thus the same sort of object as a correlated strategy among  $i$ 's opponents. See Section 1.3.3 for further discussion.

In all cases, we assume that a player's own randomization is independent of the conjectured behavior of his opponents, in the sense that joint probabilities are obtained by taking products—see (4) below.

#### 1.1.4 Expected utility

To compute a numerical assessment of a correlated strategy or mixed strategy profile, a player takes the weighted average of the utility of each pure strategy profile, with the weights given by the probabilities that each pure strategy profile occurs. This is called the *expected utility* associated with  $\sigma$ . See Section 0.2.

Player  $i$ 's expected utility from correlated strategy  $\rho$  is

$$(2) \quad u_i(\rho) = \sum_{s \in S} u_i(s) \cdot \rho(s).$$

Player  $i$ 's expected utility from mixed strategy profile  $\sigma = (\sigma_1, \dots, \sigma_n)$  is

$$(3) \quad u_i(\sigma) = \sum_{s \in S} u_i(s) \cdot \left( \prod_{j \in \mathcal{P}} \sigma_j(s_j) \right).$$

In (3), the term in parentheses is the probability that  $s = (s_1, \dots, s_n)$  is played.

*Example 1.7. Battle of the Sexes once more.*

		2	
		a	b
1	A	3, 1	0, 0
	B	0, 0	1, 3
		payoffs	

		2	
		$a \left(\frac{1}{4}\right)$	$b \left(\frac{3}{4}\right)$
1	$A \left(\frac{3}{4}\right)$	$\frac{3}{16}$	$\frac{9}{16}$
	$B \left(\frac{1}{4}\right)$	$\frac{1}{16}$	$\frac{3}{16}$
		probabilities	

Suppose  $\sigma = (\sigma_1, \sigma_2) = ((\sigma_1(A), \sigma_1(B)), (\sigma_2(a), \sigma_2(b))) = \left(\left(\frac{3}{4}, \frac{1}{4}\right), \left(\frac{1}{4}, \frac{3}{4}\right)\right)$  is played. Then

$$u_1(\sigma) = 3 \cdot \frac{3}{16} + 0 \cdot \frac{9}{16} + 0 \cdot \frac{1}{16} + 1 \cdot \frac{3}{16} = \frac{3}{4},$$

$$u_2(\sigma) = 1 \cdot \frac{3}{16} + 0 \cdot \frac{9}{16} + 0 \cdot \frac{1}{16} + 3 \cdot \frac{3}{16} = \frac{3}{4}. \quad \diamond$$

In nonequilibrium analyses, player  $i$  only has a conjecture  $\mu_i \in \Delta(\prod_{j \neq i} S_j)$  about his opponents' behavior. If player  $i$  plays mixed strategy  $\sigma_i \in \Delta S_i$ , his expected utility is

$$(4) \quad u_i(\sigma_i, \mu_i) = \sum_{s \in S} u_i(s) \cdot \sigma_i(s_i) \mu_i(s_{-i}).$$

Here we are using the independence between own randomization and conjectured opponents' behavior assumed at the end of the previous section.

There is a standard abuse of notation here. In (2)  $u_i$  acts on correlated strategies (so that  $u_i: \Delta(\prod_{j \in \mathcal{P}} S_j) \rightarrow \mathbb{R}$ ), in (3)  $u_i$  acts on mixed strategy profiles (so that  $u_i: \prod_{j \in \mathcal{P}} \Delta S_j \rightarrow \mathbb{R}$ ), and in (4)  $u_i$  acts on mixed strategy/conjecture pairs (so that  $u_i: \Delta S_i \times \Delta(\prod_{j \neq i} S_j) \rightarrow \mathbb{R}$ ). Sometimes we even combine mixed strategies with pure strategies, as in  $u_i(s_i, \sigma_{-i})$ . In the end we are always taking the expectation of  $u_i(s)$  over the relevant distribution on pure strategy profiles  $s$ , so there is really no room for confusion.

## 1.2 Dominance and Iterated Dominance

Suppose we are given some normal form game  $G$ . How should we expect Bayesian rational players (i.e., players who form conjectures about opponents' strategies and choose optimally given their conjectures) playing  $G$  to behave? We consider a sequence of increasingly restrictive methods for analyzing normal form games. We start by considering the implications of Bayesian rationality and of common knowledge of rationality. After this, we introduce equilibrium assumptions.

We always assume that the structure and payoffs of the game are *common knowledge*: that everyone knows these things, that everyone knows that everyone knows them, and so on.

Notation:

$$\begin{aligned} G = \{\mathcal{P}, \{S_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\} & \quad \text{a normal form game} \\ s_{-i} \in S_{-i} = \prod_{j \neq i} S_j & \quad \text{a profile of pure strategies for } i\text{'s opponents} \\ \mu_i \in \Delta S_{-i} & \quad i\text{'s (reduced) conjecture about his opponents' strategies} \end{aligned}$$

### 1.2.1 Strictly dominant strategies

“Dominance” concerns strategies whose performance is good (or bad) regardless of how opponents behave.

Strategy  $\sigma_i \in \Delta S_i$  is strictly dominant if

$$(5) \quad u_i(\sigma_i, \mu_i) > u_i(\sigma'_i, \mu_i) \text{ for all } \sigma'_i \neq \sigma_i \text{ and } \mu_i \in \Delta S_{-i}.$$

That is, strategy  $\sigma_i$  is strictly dominant if it maximizes  $i$ 's expected payoffs regardless of his conjecture about his opponents' strategies.

While condition (5) directly concerns Bayesian rationality, we now provide a condition that is easier to check.

**Proposition 1.8.** (i) Only a pure strategy can be strictly dominant.

(ii) Strategy  $s_i \in S_i$  is strictly dominant if and only if

$$(6) \quad u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}) \text{ for all } s'_i \neq s_i \text{ and } s_{-i} \in S_{-i}.$$

In words, (ii) says that player  $i$  prefers  $s_i$  to any other pure strategy  $s'_i$  for any pure strategy profile of his opponents.

*Proof.* (i) The payoff to a mixed strategy is the weighted average of the payoffs of the pure strategies it uses. Thus against any fixed conjecture  $\mu_i$ , the mixed strategy cannot do strictly better than all pure strategies. Thus the mixed strategy cannot be strictly dominant.

In detail: for any conjecture  $\mu_i \in \Delta(S_{-i})$ , we have

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

It is thus impossible that  $u_i(\sigma_i, \mu_i) > u_i(s_i, \mu_i)$  for all  $s_i \in S_i$ , so condition (5) cannot hold.

(ii) (5)  $\Rightarrow$  (6) is immediate: consider the conjecture  $\mu_i$  with  $\mu_i(s_{-i}) = 1$ .

(6)  $\Rightarrow$  (5) holds because the inequality in (5) is a weighted average of those in (6). In more detail: By part (i), the dominating strategy must be pure. To check (5) when the alternate strategy is also pure, observe that for any fixed  $\mu_i$ ,

$$u_i(s_i, \mu_i) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \mu_i(s_{-i}) > \sum_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i}) \mu_i(s_{-i}) = u_i(s'_i, \mu_i),$$

where the inequality follows from (6). Checking (5) when  $i$ 's alternate strategy is mixed requires an additional step; this is left as an exercise. ■

*Example 1.9. Game show revisited.*

		2	
		<i>c</i>	<i>d</i>
1	C	5,5	0,6
	D	6,0	3,3

Joint payoffs are maximized if both players chill. But regardless of what player 2 does, player 1 is better off dashing. The same is true for player 2. In other words,  $D$  and  $d$  are strictly dominant strategies.  $\diamond$

### Remarks

- (i) The game above is an instance of a *Prisoner's Dilemma*: a  $2 \times 2$  symmetric normal form game in which the strategy profile that maximizes total payoffs has both players playing a strictly dominated strategy. An asymmetric version of this game first appeared in an experiment run by Flood and Dresher in 1950; see Flood (1952). A symmetric version of the game, and the famous backstory of prisoners being interrogated in separate rooms, is due to Tucker; see Luce and Raiffa (1957).
- (ii) The entries in the payoff bimatrix are the players' Bernoulli utilities. If the game is supposed to represent the game show from Example 1.1, then having these entries correspond to the dollar amounts in the story is tantamount to assuming that (i) each player is risk neutral, and (ii) each player cares only about his own dollar payoffs. If other considerations are important—for instance, if the two players are friends and care about each others' fates—then the payoff matrix would need to be changed to reflect this, and the analysis would differ correspondingly; see Gilboa (2010, Section 7.1) for an illuminating discussion.

To summarize, the analysis above tells us only that *if* each player is rational and cares only about his dollar payoffs, then we should expect to see  $(D, d)$ .

### 1.2.2 Strictly dominated strategies

Most games do not have strictly dominant strategies. How can we get more mileage out of the notion of dominance?

Strategy  $\sigma'_i \in \Delta S_i$  is *strictly dominated* by strategy  $\sigma_i \in \Delta S_i$  if

$$u_i(\sigma_i, \mu_i) > u_i(\sigma'_i, \mu_i) \text{ for all } \mu_i \in \Delta S_{-i}.$$

Thus, Bayesian rational players never choose strictly dominated strategies.

Remarks on strictly dominated strategies:

(i)  $\sigma'_i$  is strictly dominated by  $\sigma_i$  if and only if

(7)  $u_i(\sigma_i, s_{-i}) > u_i(\sigma'_i, s_{-i})$  for all  $s_{-i} \in S_{-i}$ .

As before, this condition is easier to check than the original definition.

(ii) A strategy that is not dominated by any pure strategy may be dominated by a mixed strategy:

		2	
		L	R
1	T	3, -	0, -
	M	0, -	3, -
	B	1, -	1, -

B is not dominated by T or M but it is dominated by  $\frac{1}{2}T + \frac{1}{2}M$ . Note how this conclusion depends on taking expected utility seriously: the payoff of 1.5 generated by playing  $\frac{1}{2}T + \frac{1}{2}M$  against L is just as “real” as the payoffs of 3 and 0 obtained by playing T and M against L.


(iii) If a pure strategy  $s_i$  is strictly dominated, then so is any mixed strategy  $\sigma_i$  with  $s_i$  in its support (i.e., that uses  $s_i$  with positive probability). This is because any weight placed on a strictly dominated strategy can instead be placed on the dominating strategy, which raises player  $i$ 's payoffs regardless of how his opponents act. For instance, in the example from part (ii),  $\frac{2}{3}M + \frac{1}{3}B$  is strictly dominated by

$$\frac{2}{3}M + \frac{1}{3}(\frac{1}{2}T + \frac{1}{2}M) = \frac{1}{6}T + \frac{5}{6}M.$$

Turning this idea into a proof is a good exercise.

(iv) But even if a group of pure strategies are not dominated, mixed strategies that combine them may be:

		2	
		L	R
1	T	3, -	0, -
	M	0, -	3, -
	C	2, -	2, -



T is optimal against L and M is optimal against R, so T and M are not strictly dominated. (This is immediate—cf. Observation 1.17). But  $\frac{1}{2}T + \frac{1}{2}M$  (guarantees  $\frac{3}{2}$ ) is strictly dominated by C (guarantees 2). In fact, any mixed strategy with both T and

$M$  in its support is strictly dominated: if mixed strategy  $\sigma_1$  has  $\sigma_1(T), \sigma_1(M) \geq p > 0$ , then by reducing  $\sigma_1(T)$  and  $\sigma_1(M)$  by  $p$  and increasing  $\sigma_1(C)$  by  $2p$ , we can increase expected payoffs by  $2p \cdot (2 - \frac{3}{2}) = p$  against both of player 2's pure strategies.

### 1.2.3 Iterated strict dominance

Some games without a strictly dominant strategy can still be solved using the idea of dominance.

*Example 1.10.* Consider the two-player normal form game  $G$  with players  $A$  (Al) and  $B$  (Bess), strategy sets  $S_A = S_B = \{1, 2, 3\}$  (representing effort levels), and payoff functions

$$u_A(s_A, s_B) = (\min\{s_A, s_B\})^2 - |s_A - s_B|, \quad u_B(s_A, s_B) = \begin{cases} (s_A)^2 - s_B & \text{if } (s_A, s_B) \neq (3, 1), \\ -s_B & \text{if } (s_A, s_B) = (3, 1). \end{cases}$$

In words, Al receives the square of the lower strategy chosen and pays a cost equal to the distance between the two strategies. Bess pays  $s_B$ , and she receives the square of  $s_A$  unless she chooses the minimum effort and Al chooses the maximum. This game's payoff matrix is presented below.

		B		
		1	2	3
A	1	1, 0	0, -1	-1, -2
	2	0, 3	4, 2	3, 1
	3	-1, -1	3, 7	9, 6

Analysis: Al does not have a dominated pure strategy. Bess has one dominated pure strategy:  $s_B = 3$  is dominated by strategy  $s_B = 2$ . Thus if Bess is rational she will not play  $s_B = 3$ .

If Al knows that Bess is rational, he knows that she will not play  $s_B = 3$ . So if Al is rational, he won't play  $s_A = 3$ , which is strictly dominated by  $s_A = 2$  once  $s_B = 3$  is removed.

Now if Bess knows:

- (i) that Al knows that Bess is rational
- (ii) that Al is rational

then Bess knows that Al will not play  $s_A = 3$ . Hence, since Bess is rational she will not play  $s_B = 2$ .

Continuing in a similar vein: Al will not play  $s_A = 2$ .

Therefore,  $(s_A, s_B) = (1, 1)$  solves the game by iterated strict dominance.  $\diamond$

Iterated strict dominance is driven by *common knowledge of rationality*—by the assumption that all statements of the form “ $i$  knows that  $j$  knows that . . .  $k$  is rational” are true—as well as by common knowledge of the game itself.

To see which strategies survive iterated strict dominance it is enough to

- (i) Iteratively remove all dominated pure strategies.
- (ii) When no further pure strategies can be removed, check all remaining mixed strategies.

The remarks from Section 1.2.2 provide the reasons that we don’t need to consider mixed strategies explicitly until the end. During the iterated elimination, removing a pure strategy implicitly removes the mixed strategies that use it. This is justified by remark (iii), which says that if a pure strategy is dominated, so are all mixed strategies that use it. Remark (i) says that to see which of a player’s strategies are dominated, we only need to consider performance against opponents’ pure strategy profiles. Because of this, removing opponents’ dominated mixed strategies at some stage of the iteration won’t affect which strategies can be removed in the next stage. But remark (iv) says we need to look for dominated mixed strategies at the end, because a mixed strategy can be dominated even if none of the pure strategies it uses are. (And remark (iii) says that we need to consider mixed strategies as possible dominating strategies.)

A basic fact about iterated strict dominance is:

**Proposition 1.11.** *The set of strategies that remains after iteratively removing strictly dominated strategies does not depend on the order in which the dominated strategies are removed.*

See Dufwenberg and Stegeman (2002) or Ritzberger (2002) for a proof.

Often, iterated strict dominance will eliminate a few strategies but not completely solve the game.

#### 1.2.4 Weak dominance

Strategy  $\sigma'_i \in \Delta S_i$  is *weakly dominated* by  $\sigma_i$  if

$$(8) \quad \begin{aligned} u_i(\sigma_i, \mu_i) &\geq u_i(\sigma'_i, \mu_i) && \text{for all } \mu_i \in \Delta S_{-i}, \text{ and} \\ u_i(\sigma_i, \mu'_i) &> u_i(\sigma'_i, \mu'_i) && \text{for some } \mu'_i \in \Delta S_{-i}. \end{aligned}$$



As with strict domination, it is enough to consider opponents' pure strategy profiles: (8) is equivalent to

$$(9) \quad \begin{aligned} u_i(\sigma_i, s_{-i}) &\geq u_i(\sigma'_i, s_{-i}) && \text{for all } s_{-i} \in S_{-i}, \text{ and} \\ u_i(\sigma_i, s'_{-i}) &> u_i(\sigma'_i, s'_{-i}) && \text{for some } s'_{-i} \in S_{-i} \end{aligned}$$

Strategy  $\sigma_i \in \Delta S_i$  is *weakly dominant* if it weakly dominates all other strategies. As the names suggest, all strictly dominant strategies are also weakly dominant. And as with strict dominance, only pure strategies can be weakly dominant.

(We use the term *very weak dominance* when only the first requirement in (8) or (9) is imposed. This notion is more commonly applied in Bayesian games (Section 4.2) and especially mechanism design (Section 7.1.3).)

*Example 1.12.* Weakly dominated strategies are not ruled out by Bayesian rationality alone.

		2		
		L	R	
1	T	1, -	0, -	◇
	B	0, -	0, -	

While the use of weakly dominated strategies is not ruled out by Bayesian rationality alone, the avoidance of such strategies is often taken as a first principle. In decision theory, this principle is referred to as *admissibility*; see Kohlberg and Mertens (1986) for discussion and historical comments. In game theory, admissibility is sometimes deduced from the principle of *cautiousness*, which requires that players not view any opponents' behavior as impossible; see Asheim (2006) for discussion.

It is natural to contemplate iteratively removing weakly dominated strategies. However, iterated removal and cautiousness conflict with one another: removing a strategy means viewing it as impossible, which contradicts cautiousness. See Samuelson (1992) for discussion and analysis. One consequence is that the order of removal of weakly dominated strategies can matter—see Example 1.13 below. (For results on when order of removal does not matter, see Marx and Swinkels (1997) and Østerdal (2005).) But versions of iterated weak dominance can be placed on a secure epistemic footing (see Brandenburger et al. (2008)), and moreover, iterated weak dominance is a powerful tool for analyzing extensive form games (see Section 2.5.1).

*Example 1.13.* Order of removal can matter under IWD.

		2	
		L	R
U		5, 1	4, 0
1 M		6, 0	3, 1
D		6, 4	4, 4

In the game above, removing weakly dominated strategy  $U$ , then weakly dominated strategy  $L$ , and then strictly dominated strategy  $M$  leads to prediction  $(D, R)$ . But removing weakly dominated strategy  $M$ , then weakly dominated strategy  $R$ , and then strictly dominated strategy  $U$  leads to the prediction  $(D, L)$ .  $\diamond$

An intermediate solution concept between ISD and IWD is introduced by Dekel and Fudenberg (1990), who suggest one round of elimination of all weakly dominated strategies, followed by iterated elimination of strictly dominated strategies. Since weak dominance is not applied iteratively, the tensions described above do not arise. Strategies that survive this *Dekel-Fudenberg procedure* are sometimes called *permissible*. See Section 2.6 for further discussion.

### 1.3 Rationalizability

#### 1.3.1 Definition and examples

Iterated strict dominance may not exhaust the set of strategies that won't be played by Bayesian rational players under common knowledge of rationality. Bayesian rational players not only avoid dominated strategies; they also avoid strategies that are not a best response to any conjecture about opponent's behavior. If we apply this idea iteratively, we obtain the sets of rationalizable strategies for each player.

Strategy  $\sigma_i$  is a *best response* to (reduced) conjecture  $\mu_i \in \Delta S_{-i}$  if

$$(10) \quad u_i(\sigma_i, \mu_i) \geq u_i(\sigma'_i, \mu_i) \quad \text{for all } \sigma'_i \in \Delta S_i.$$

In contrast with dominance, there is no "for all  $\mu_i$ " in (10).

When (10) holds we write  $\sigma_i \in B_i(\mu_i)$ . Thus (10) defines a set-valued map  $B_i: \Delta S_{-i} \Rightarrow \Delta S_i$ , which is called *player  $i$ 's best response correspondence*.

The following proposition provides an easy way of checking whether a mixed strategy is a best response.

**Proposition 1.14.** *Strategy  $\sigma_i$  is a best response to  $\mu_i$  if and only if every pure strategy  $s_i$  in the support of  $\sigma_i$  is a best response to  $\mu_i$ .*

*Proof.* Rewrite player  $i$ 's expected utility from mixed strategy  $\sigma_i$  given conjecture  $\mu_i$  as

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

The final expression is a weighted average of the expected utilities  $u_i(s_i, \mu_i)$  of each pure strategy  $s_i$  given conjecture  $\mu_i$ . The probability distributions  $\sigma_i \in \Delta S_i$  that maximize this weighted average are precisely those that place all mass on pure strategies that maximize  $u_i(s_i, \mu_i)$ . ■

The *rationalizable strategies* (Bernheim (1984), Pearce (1984)) are those that remain after we iteratively remove all strategies that are not best responses to any allowable conjectures. Because it only requires common knowledge of rationality, rationalizability is a relatively weak solution concept. Still, when rationalizability leads to many rounds of removal, it can result in stark predictions.

*Example 1.15. Guessing  $\frac{3}{4}$  of the average.*

There are  $n$  players. Each player's strategy set is  $S_i = \{0, 1, \dots, 100\}$ .

The *target integer* is defined to be  $\frac{3}{4}$  of the average strategy chosen, rounding down.

All players choosing the target integer split a prize worth  $V > 0$ . If no one chooses the target integer, the prize is not awarded.

Which pure strategies are rationalizable in this game?

To start, we assert that for any pure strategy profile  $s_{-i}$  of his opponents, player  $i$  has a response  $r_i \in S_i$  such that the target integer generated by  $(r_i, s_{-i})$  is  $r_i$ . (Proving this is a short but tricky exercise.) Thus for any conjecture  $\mu_i$  about his opponents, player  $i$  can obtain a positive expected payoff by playing a best response to some  $s_{-i}$  in the support of  $\mu_i$ .

So: Since  $S_i = \{0, 1, \dots, 100\}$ ,

- ⇒ The highest possible average is 100.
- ⇒ The highest possible target is  $\lfloor \frac{3}{4} \cdot 100 \rfloor = 75$ .
- ⇒ Strategies in  $\{76, 77, \dots, 100\}$  yield a payoff of 0.
- ⇒ Since player  $i$  has a strategy that earns a positive expected payoff given his conjecture, strategies in  $\{76, 77, \dots, 100\}$  are not best responses.

Thus if players are rational, no player chooses a strategy above 75.

- ⇒ The highest possible average is 75.
- ⇒ The highest possible target is  $\lfloor \frac{3}{4} \cdot 75 \rfloor = \lfloor 56\frac{1}{4} \rfloor = 56$ .
- ⇒ Strategies in  $\{57, \dots, 100\}$  yield a payoff of 0.
- ⇒ Since player  $i$  has a strategy that earns a positive expected payoff given his conjecture, strategies in  $\{57, \dots, 100\}$  are not best responses.

Thus if players are rational and know that others are rational, no player chooses a strategy above 56.

Proceeding through the rounds of eliminating strategies that cannot be best responses, we find that no player will choose a strategy higher than

$$75 \dots 56 \dots 42 \dots 31 \dots 23 \dots 17 \dots 12 \dots 9 \dots 6 \dots 4 \dots 3 \dots 2 \dots 1 \dots 0.$$

Thus, after 14 rounds of iteratively removing strategies that cannot be best responses, we conclude that each player's unique rationalizable strategy is 0.  $\diamond$

*Example 1.16. Determining the rationalizable strategies in a normal form game.*

		2		
		L	C	Q
1	T	3,3	0,0	0,2
	M	0,0	3,3	0,2
	D	2,2	2,2	2,0

Write  $\mu_1 = (l, c, q)$  and  $\mu_2 = (t, m, d)$ .

To find  $B_1 : \Delta S_2 \Rightarrow \Delta S_1$


$$\begin{aligned} u_1(T, \mu_1) \geq u_1(M, \mu_1) &\Leftrightarrow 3l \geq 3c \\ &\Leftrightarrow l \geq c \end{aligned}$$

$$\begin{aligned} u_1(T, \mu_1) \geq u_1(D, \mu_1) &\Leftrightarrow 3l \geq 2 \\ &\Leftrightarrow l \geq \frac{2}{3} \end{aligned}$$

To find  $B_2 : \Delta S_1 \Rightarrow \Delta S_2$

$$\begin{aligned} u_2(\mu_2, L) \geq u_2(\mu_2, C) &\Leftrightarrow 3t + 2d \geq 3m + 2d \\ &\Leftrightarrow t \geq m \end{aligned}$$

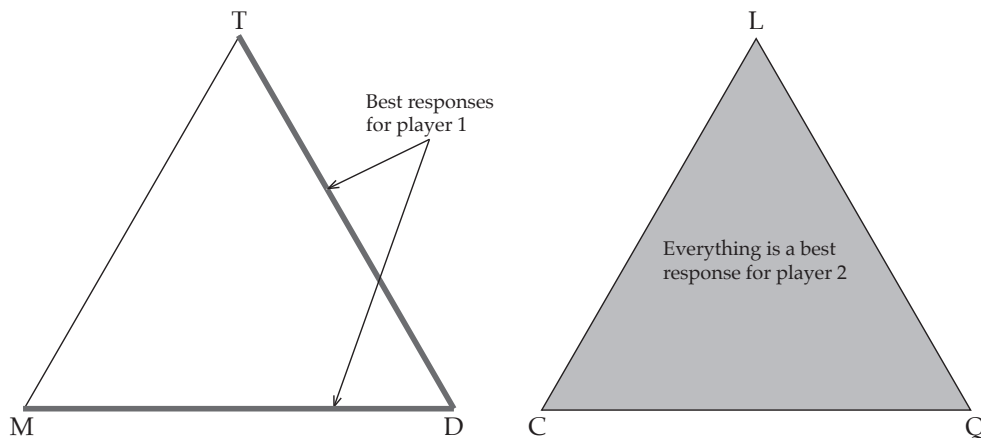
$$\begin{aligned} u_2(\mu_2, L) \geq u_2(\mu_2, Q) &\Leftrightarrow 3t + 2d \geq 2t + 2m \\ &\Leftrightarrow t + 2d \geq 2m \end{aligned}$$



To compute each player's set of rationalizable strategies, we start by computing each player's set of best responses to all conjectures about their opponent's strategies. We use Proposition 1.14: a mixed strategy is optimal if and only if all strategies in its support are optimal.

For player 1, all mixed strategies that do not put positive probability on both  $T$  and  $M$  are best responses:  $B_1(\Delta S_2) = \{\sigma_1 \in \Delta S_1 : \sigma_1(T) = 0 \text{ or } \sigma_1(M) = 0\}$ .

For player 2, all mixed strategies are best responses:  $B_2(\Delta S_1) = \Delta S_2$ .



In the next round of eliminating strategies that are not best responses to any conjectures, player 2 knows that player 1 will never play a mixture that uses both  $T$  and  $M$ . Can it be a best response for her to play  $Q$ ?

Yes. Player 2 knows that  $T$ ,  $M$ , and  $D$  are all possible best responses for 1. Thus player 2 may hold any conjectures that put all of their weight on these pure strategies. One

possible conjecture for player 2 is  $\mu_2(T) = \mu_2(M) = \frac{1}{2}$ , and against this conjecture,  $Q$  is the unique best response.

Likewise, if 2's conjecture is  $\mu_2(T) = \mu_2(M) = \frac{2}{5}$  and  $\mu_2(D) = \frac{1}{5}$ , then all of her strategies are best responses. Thus 2's set of rationalizable strategies is  $\Delta S_2$ . Since none of 2's strategies are removed, 1's set of rationalizable strategies is his set of best responses,  $\{\sigma_1 \in \Delta S_1: \sigma_1(T) = 0 \text{ or } \sigma_1(M) = 0\}$ .

To explain what is happening more carefully we need to bring back full conjectures (Section 1.1.3). The reduced conjecture  $\mu_2(T) = \mu_2(M) = \frac{1}{2}$  is generated by the full conjecture  $v_2(\frac{1}{2}T + \frac{1}{2}M) = 1$ . Since  $\frac{1}{2}T + \frac{1}{2}M$  is not a best response, this full conjecture is not allowed. But  $\mu_2(T) = \mu_2(M) = \frac{1}{2}$  is also generated by the full conjecture  $v_2(T) = v_2(M) = \frac{1}{2}$ . Since  $T$  and  $M$  are both best responses, this full conjecture is allowed, which explains why the reduced conjecture is allowed.  $\diamond$

When we compute the rationalizable strategies, we must account for each player's uncertainty about his opponent's strategies. Thus, during each iteration we must leave in all of his best responses to all conjectures over the surviving pure strategies, even conjectures that correspond to mixed strategies that can be ruled out. We thus have the following procedure to compute the players' rationalizable strategies:

- (i) Iteratively remove pure strategies that are never a best response to any allowable conjecture.
- (ii) When no further pure strategies can be removed, remove mixed strategies that are never a best response to any allowable conjecture.

There are refinements of rationalizability based on assumptions beyond common knowledge of rationality that generate tighter predictions in some games, while still avoiding the use of equilibrium knowledge assumptions—see Section 2.6.

### *Formal definition of rationalizability*

Here is a concise formal definition of rationalizability for two-player games.

For each player  $i$ , let  $\Sigma_i^0 = \Delta S_i$ . For  $k \geq 1$ , define  $\Sigma_i^k$  recursively by

$$(11) \quad \Sigma_i^k = B_i\left(\text{conv}(\Sigma_j^{k-1})\right) \equiv \{\sigma_i \in \Sigma_i^{k-1}: \sigma_i \in B_i(\mu_i) \text{ for some } \mu_i \in \text{conv}(\Sigma_j^{k-1})\}.$$

where  $\text{conv}(A)$  denotes the convex hull of the set  $A$  (and  $j$  is  $i$ 's opponent). Then the set  $R_i$  of player  $i$ 's rationalizable strategies is the intersection of the nested sets  $\Sigma_i^k$ :

$$(12) \quad R_i^* = \bigcap_{k=0}^{\infty} \Sigma_i^k.$$

In a finite game, this procedure terminates in a finite number of steps, so that  $R_i^* = \Sigma_i^k$  for all large enough finite  $k$ .

In this definition,  $\Sigma_j^{k-1} \subseteq \Delta(S_j)$  is the set of mixed strategies of player  $j$  that survive  $k - 1$  rounds of removal. Then  $\text{conv}(\Sigma_j^{k-1}) \subseteq \Delta(S_j)$  is the set of allowable conjectures of player  $i$  used in round  $k$ . Notice how this step takes advantage of the fact that  $j$ 's mixed strategies and  $i$ 's conjectures are both elements of  $\Delta(S_j)$ . Finally  $\Sigma_i^k = B_i(\text{conv}(\Sigma_j^{k-1}))$  is the set of  $i$ 's best responses to these allowable conjectures, and so is the set of  $i$ 's strategies that survive  $k$  rounds of removal.

As in the computation procedure described above, we could replace the set of mixed strategies  $\Sigma_j^k$  with the set of pure strategies it contains until the final round of removal; both sets generate the same set of conjectures, or equivalently, both have the same convex hull.

### 1.3.2 Rationalizability and iterated strict dominance

#### *Rationalizability and iterated strict dominance in two-player games*

It is obvious that:

**Observation 1.17.** *If  $\sigma_i$  is strictly dominated, then  $\sigma_i$  is not a best response to any conjecture.*

In two-player games, the converse of this observation is also true: any strategy that is not a best response to any conjecture is strictly dominated. (Below we explain how to obtain this result as an application of the supporting hyperplane theorem.) We therefore have an equivalence:

**Theorem 1.18.** *In a two-player game,  $\sigma_i$  is strictly dominated if and only if  $\sigma_i$  is not a best response to any conjecture.*

Applying this result iteratively yields:

**Theorem 1.19.** *In a two-player game, a strategy satisfies iterated strict dominance if and only if it is rationalizable.*


Whether Theorems 1.18 and 1.19 hold in games with more players depends on how we define rationalizability in these games—see Section 1.3.3.

#### *Background on hyperplanes*

For each nonzero vector  $p \in \mathbb{R}^n$ , the set  $H_{p,0} = \{z \in \mathbb{R}^n : p \cdot z = 0\}$  is a *subspace* of  $\mathbb{R}^n$  of dimension  $n - 1$ . The vector  $p$  called a *normal vector* to this subspace. If we interpret each


$x \in H_{p,0}$  as a *vector* (namely, the vector from the origin to the *point*  $x$ ), then  $H_{p,0}$  is the set of vectors that are orthogonal to  $p$ . Being able to interpret elements of  $\mathbb{R}^n$  both as points and as vectors is key for understanding what follows.

*Example 1.20.* In  $\mathbb{R}^2$ , a subspace of dimension 1 is a line through the origin. For instance, the line  $x_2 = ax_1$  defines a subspace. Since we can rewrite this equation as  $(-a, 1) \cdot x = 0$ , it is the subspace  $H_{p,0}$  with normal vector  $p = (-a, 1)$ . Below is the case in which  $a = -\frac{1}{2}$ , so that  $p = (\frac{1}{2}, 1)$ .  $\diamond$



A *hyperplane* in  $\mathbb{R}^n$  is a set of the form  $H_{p,c} = \{x \in \mathbb{R}^n : p \cdot x = c\}$  for some *normal vector*  $p$  and *intercept*  $c$ . Viewing points in the various spaces as vectors,  $H_{p,c}$  can be obtained by translating the subspace  $H_{p,0}$  by any element of  $H_{p,c}$  itself. That is, for any  $\hat{x} \in H_{p,c}$ , we have  $H_{p,c} = H_{p,0} + \{\hat{x}\}$ , where the latter set is defined as  $\{z + \hat{x} \in \mathbb{R}^n : z \in H_{p,0}\}$ . Thus  $H_{p,c}$  and  $H_{p,0}$  are parallel to one another. To put the last fact in different words, let  $T(H_{p,c}) = \{x - \hat{x} : x, \hat{x} \in H_{p,c}\}$  be the set of vectors tangent to  $H_{p,c}$  (i.e., the displacements that can be obtained by moving between points in  $H_{p,c}$ ). Then  $T(H_{p,c}) = H_{p,0}$ . These claims are best understood ideas geometrically; again, moving between the interpretations of elements of  $\mathbb{R}^n$  as points and vectors is key. It is worth working through the (simple) proof of one of these facts.

*Example 1.21.* In  $\mathbb{R}^2$ , a hyperplane is a line. For instance, the line  $x_2 = ax_1 + c$  corresponds to the hyperplane  $H_{p,c}$  with  $p = (-a, 1)$ . We again show the case in which  $a = -\frac{1}{2}$ , so that  $p = (\frac{1}{2}, 1)$ .  $\diamond$





When the normal vector  $p$  is drawn with its tail at the origin, it points towards hyperplanes  $H_{p,c}$  with  $c > 0$ . This follows from the law of cosines:  $p \cdot x = \|p\| \|x\| \cos \theta > 0$  where  $\theta$  is the angle between vectors  $p$  and  $x$ ; thus  $p \cdot x > 0$  when  $\theta$  is acute.

A set of the form  $\{x \in \mathbb{R}^n : p \cdot x \geq c\}$  is called a *half space*; it contains the points on the side of the hyperplane  $H_{p,c}$  toward which  $p$  is pointing.

### *The supporting and separating hyperplane theorems*

**Theorem 1.22** (The supporting hyperplane theorem).

Let  $A \subset \mathbb{R}^n$  be a closed convex set and let  $y \in \text{bd}(A)$ . Then there exists a  $p \in \mathbb{R}^n - \{0\}$  such that  $p \cdot x \leq p \cdot y$  for all  $x \in A$ .


$H_{p,c}$  with  $c = p \cdot y$  is called a *supporting hyperplane* of  $A$  at  $y$ .

Theorem 1.22 is a special case of:

**Theorem 1.23** (The separating hyperplane theorem).

Let  $A, B \subset \mathbb{R}^n$  be closed convex sets such that  $A \cap B \subseteq \text{bd}(A) \cap \text{bd}(B)$ . Then there exists a  $p \in \mathbb{R}^n - \{0\}$  such that  $p \cdot x \leq p \cdot \hat{x}$  for all  $x \in A$  and  $\hat{x} \in B$ .

Let  $c$  satisfy  $p \cdot x \leq c \leq p \cdot \hat{x}$  for all  $x \in A$  and  $\hat{x} \in B$ . Then  $H_{p,c}$  is a hyperplane that *separates*  $A$  and  $B$ . Any  $y$  in  $\text{bd}(A) \cap \text{bd}(B)$  must satisfy  $p \cdot y = c$ .



For proofs, discussion, examples, etc. see Hiriart-Urruty and Lemaréchal (2001).

### *Best responses and dominance in two-player games*

We recall this earlier result:

**Theorem 1.18.** In a two-player game,  $\sigma_i$  is strictly dominated if and only if  $\sigma_i$  is not a best response to any conjecture.

The ( $\Rightarrow$ ) direction is immediate. The ( $\Leftarrow$ ) direction follows from the supporting hyperplane theorem. We illustrate the proof using an example.


*Example 1.24.* Our goal is to show that in the two-player game below,  $[\sigma_i \in \Delta S_i$  is not strictly dominated] implies that  $[\sigma_i$  is a best response to some  $\mu_i \in \Delta S_{-i}]$ .

		2	
		L	R
1	A	2, -	5, -
	B	6, -	3, -
	C	7, -	1, -
	D	3, -	2, -

Let  $v_1(\sigma_1) = (u_1(\sigma_1, L), u_1(\sigma_1, R))$  be the vector payoff induced by  $\sigma_1$ . Note that  $u_1(\sigma_1, \mu_1) = \mu_1 \cdot v_1(\sigma_1)$ .

Let  $V_1 = \{v_1(\sigma_1) : \sigma_1 \in \Delta S_1\}$  be the set of such vector payoffs. Equivalently,  $V_1$  is the convex hull of the vector payoffs to player 1's *pure* strategies. It is closed and convex.

Now  $\sigma_1 \in \Delta S_1$  is not strictly dominated if and only if  $v_1(\sigma_1)$  lies on the northeast boundary of  $V_1$ . For example,  $\hat{\sigma}_1 = \frac{1}{2}A + \frac{1}{2}B$  is not strictly dominated, with  $v_1(\hat{\sigma}_1) = (4, 4)$ . We want to show that  $\hat{\sigma}_1$  is a best response to some  $\hat{\mu}_1 \in \Delta S_2$ .



A general principle: when you are given a point on the boundary of a convex set, the normal vector at that point often reveals something interesting.

The point  $v_1(\hat{\sigma}_1)$  lies on the hyperplane  $\hat{\mu}_1 \cdot w_1 = 4$ , where  $\hat{\mu}_1 = (\frac{1}{3}, \frac{2}{3})$ .

This is a supporting hyperplane for the set  $V_1$ , where  $\hat{\mu}_1 \cdot w_1 \leq 4$ .


Put differently,

$$\begin{aligned} & \hat{\mu}_1 \cdot w_1 \leq \hat{\mu}_1 \cdot v_1(\hat{\sigma}_1) \quad \text{for all } w_1 \in V_1 \\ \Rightarrow & \hat{\mu}_1 \cdot v_1(\sigma_1) \leq \hat{\mu}_1 \cdot v_1(\hat{\sigma}_1) \quad \text{for all } \sigma_1 \in \Delta S_1 \text{ (by the definition of } V_1) \\ \Rightarrow & u_1(\sigma_1, \hat{\mu}_1) \leq u_1(\hat{\sigma}_1, \hat{\mu}_1) \quad \text{for all } \sigma_1 \in \Delta S_1 \end{aligned}$$

Therefore,  $\hat{\sigma}_1$  is a best response to  $\hat{\mu}_1$ .

The same argument shows that *every* mixture of  $A$  and  $B$  is a best response to  $\tilde{\mu}_1$ .

We can repeat this argument for all mixed strategies of player 1 corresponding to points on the northeast frontier of  $V_1$ , as in the figure below at left. The figure below at right presents player 1's best response correspondence, drawn beneath graphs of his pure strategy payoff functions. Both figures link player 1's conjectures and best responses: in the left figure, player 1's conjectures are the normal vectors, while in the right figure, player 1's conjectures correspond to different horizontal coordinates.



### 1.3.3 Rationalizability with three or more players

There are two definitions of rationalizability for games with three or more players. The commonly used definition, called *correlated rationalizability*, allows for correlation in player  $i$ 's conjecture about his opponents' choices. This is how we defined conjectures in games with three or more players in Section 1.1.3. Player  $i$  might hold a correlated conjecture if he thought his opponents could coordinate their choices by observing a (not explicitly modeled) signal that  $i$  himself does not observe. (This argument is related to the *revelation principle for normal form games*; see Section 1.5.2 and especially Section 4.5.2.)

With correlated conjectures, the supporting hyperplane argument from the two-player case goes through essentially unchanged, so we recover the earlier equivalences:

**Theorem 1.25.** *In a finite-player game,  $\sigma_i$  is strictly dominated if and only if  $\sigma_i$  is not a best response to any correlated conjectures.*

**Theorem 1.26.** *In a finite-player game, a strategy satisfies iterated strict dominance if and only if it is correlated rationalizable.*

The original definition of rationalizability from Bernheim (1984) and Pearce (1984) is now called *independent rationalizability*. This definition assumes that a player separately forms conjectures about each of his opponents' mixed strategies, and that his conjectures about different opponents' choices are independent of one another, in that the induced correlated conjecture is a product measure. With such independent conjectures, Theorem 1.25 fails: there are three-player games with strategies that are not a best response to any independent conjectures but are not strictly dominated (Example 1.27). The reason the separating hyperplane argument does not work is that with three or more players, the normal vectors to the analogue of the set  $V_1$  at a "northeast" boundary point are correlated conjectures, and none of these may correspond to an independent conjecture.

*Example 1.27.* Consider the following three-player game in which only player 3's payoffs are shown.

		3:A		3:B		3:C			
		2		2		2			
		L	R	L	R	L	R		
1	T	-, -, 5	-, -, 2	T	-, -, 4	-, -, 0	T	-, -, 1	-, -, 2
	B	-, -, 2	-, -, 1	B	-, -, 0	-, -, 4	B	-, -, 2	-, -, 5

Strategy  $B$  is not strictly dominated, since a dominating mixture of  $A$  and  $C$  would need to put at least probability  $\frac{3}{4}$  on both  $A$  (in case 1 and 2 play  $(T, L)$ ) and  $C$  (in case 1 and 2 play  $(B, R)$ ). If player 3's conjectures about player 1's choices and player 2's choices are independent,  $B$  is not a best response: Independence implies that for some  $t, l \in [0, 1]$ , we can write  $\mu_3(T, L) = tl$ ,  $\mu_3(T, R) = t(1 - l)$ ,  $\mu_3(B, L) = (1 - t)l$ , and  $\mu_3(B, R) = (1 - t)(1 - l)$ . Then

$$\begin{aligned}
 u_3(C, \mu_3) &> u_3(B, \mu_3) \\
 \Leftrightarrow \quad &tl + 2t(1 - l) + 2(1 - t)l + 5(1 - t)(1 - l) > 4tl + 4(1 - t)(1 - l) \\
 \Leftrightarrow \quad &1 + t + l > 6tl,
 \end{aligned}$$

which is true whenever  $t + l \leq 1$  (why?); symmetrically,  $u_3(A, \mu_3) > u_3(B, \mu_3)$  whenever  $t + l \geq 1$ . But  $B$  is a best response to the correlated conjecture  $\mu_3(T, L) = \mu_3(B, R) = \frac{1}{2}$ .  $\diamond$

### 1.3.4 A positive characterization of rationalizability

The procedure introduced earlier defines rationalizability in a “negative” fashion, by iteratively removing strategies that are not best responses to any conjecture. It is good to have a “positive” characterization, describing rationalizability in terms of what it requires rather than what it rules out. Theorem 1.28 provides this, focusing on pure strategies.

**Theorem 1.28.** *Let  $R_i \subseteq S_i$  for all  $i \in \mathcal{P}$ , and consider the following condition:*

- (\*) For each  $i \in \mathcal{P}$  and each  $s_i \in R_i$ , there is a  $\mu_i \in \Delta S_{-i}$  such that
- (a) the support of  $\mu_i$  is contained in  $R_{-i} = \prod_{j \neq i} R_j$ , and
  - (b)  $s_i \in B_i(\mu_i)$ .
- (i) If sets  $R_1, \dots, R_n$  satisfy condition (\*), then the strategies in these sets are rationalizable.
- (ii) The sets  $R_1^*, \dots, R_n^*$  of rationalizable pure strategies satisfy condition (\*).

*Proof.* For part (i), note that condition (\*) implies that strategies in the sets  $R_i$  cannot be removed during the first round of elimination, and thus cannot be removed in the second round, and thus the third either. . .

For part (ii), let  $R_1^*, \dots, R_n^*$  be the sets of rationalizable pure strategies. Since these sets are obtained after a finite number of rounds of elimination, each  $s_i \in R_i^*$  must be a best response to a conjecture concentrated on  $R_{-i}^*$ ; otherwise it would have been removed before the iteration terminated. ■

The following is the special case of Theorem 1.28(i) obtained when all conjectures  $\mu_i$  are concentrated on a single pure strategy profile.

**Corollary 1.29.** *Let  $R_i \subseteq S_i$  for all  $i \in \mathcal{P}$ . Suppose that for each  $i \in \mathcal{P}$  and each  $s_i \in R_i$  there is a  $s_{-i} \in R_{-i}$  such that  $s_i \in B_i(s_{-i})$ . Then the strategies in each set  $R_i$  are rationalizable.*

*Example 1.30.* Consider the following game:

		2		
		x	y	z
1	X	3,3	5,2	5,2
	Y	2,5	7,0	0,7
	Z	2,5	0,7	7,0

We show that all pure strategies are rationalizable by applying Corollary 1.29.

First consider  $(R_1, R_2) = (\{X\}, \{x\})$ . Since  $X \in B_1(x)$  and  $x \in B_2(X)$ , these strategies are rationalizable. Since each strategy is a best response to the opponent’s strategy,  $(X, x)$  is a *Nash equilibrium*.

Now consider  $(R_1, R_2) = (\{Y, Z\}, \{y, z\})$ . Since  $Y \in B_1(y)$ ,  $Z \in B_1(z)$ ,  $y \in B_2(Z)$ , and  $z \in B_2(Y)$ , these strategies are also rationalizable. Here no strategy profile consists of mutual best responses; instead, the four strategy profiles form a *best response cycle*. Clearly, the rationalizability of these strategies relies on the fact that correct conjectures are not required.  $\diamond$

With some additional notation, we can restate (\*) from Theorem 1.28 concisely as a (set-valued) fixed point condition. For each set  $R_{-i}$ , let

$$B_i(R_{-i}) = \bigcup_{\mu_i: \mu_i(R_{-i})=1} B_i(\mu_i)$$

be the set of strategies of player  $i$  that are best responses to some conjecture that is concentrated on opponents' strategy profiles in  $R_{-i}$ . Then condition (\*) can be restated as the requirement that

$$(13) \quad R_i \subseteq B_i(R_{-i}) \text{ for all } i \in \mathcal{P}.$$

## 1.4 Nash Equilibrium

Rationalizability only relies on common knowledge of rationality. Unfortunately, it often fails to provide tight predictions of play. To obtain tighter predictions, we need to impose stronger restrictions on players' conjectures about their opponents' behavior. If we add the requirement that conjectures are correct, we obtain Nash equilibrium, the central solution concept of noncooperative game theory.

### 1.4.1 Definition

We first define Nash equilibrium in a way that is indirect (and abuses notation) but connects Nash equilibrium with rationalizability (cf. Theorem 1.28). Strategy profile  $\sigma \in \Sigma$  is a *Nash equilibrium supported by full conjectures*  $\{v_i\}_{i \in \mathcal{P}}$  if for all  $i \in \mathcal{P}$ ,

- (a)  $v_i(\sigma_{-i}) = 1$ , and
- (b)  $\sigma_i \in B_i(v_i)$ .

In words, (a) says that player  $i$  has correct beliefs—he assigns probability 1 to the mixed strategy profile his opponents will play; (b) says that player  $i$ 's strategy is optimal given these beliefs.

Now that we have made our point about correct conjectures, we can define Nash equi-

librium directly, without reference to conjectures. To reduce the amount of notation, let  $\Sigma_i = \Delta S_i$  denote player  $i$ 's set of mixed strategies; also let  $\Sigma = \prod_{j \in \mathcal{P}} \Delta S_j$  and  $\Sigma_{-i} = \prod_{j \neq i} \Delta S_j$ . Finally, recycling notation, let  $B_i : \Sigma_{-i} \Rightarrow \Sigma_i$  be player  $i$ 's best response correspondence defined over opponents' mixed strategy profiles:

$$B_i(\sigma_{-i}) = \operatorname{argmax}_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_{-i})$$

Then strategy profile  $\sigma \in \Sigma$  is a *Nash equilibrium* (Nash (1950)) if for all  $i \in \mathcal{P}$ ,

$$(14) \quad \sigma_i \in B_i(\sigma_{-i}).$$

In words: each player plays a best response to the strategy profile of his opponents. (Compare this to characterization (13) of rationalizability.)

*Example 1.31. Good Restaurant, Bad Restaurant.*

		2	
		$g$	$b$
1	G	2, 2	0, 0
	B	0, 0	1, 1

Everything is rationalizable.

The Nash equilibria are:  $(G, g)$ ,  $(B, b)$ ,  $(\frac{1}{3}G + \frac{2}{3}B, \frac{1}{3}g + \frac{2}{3}b)$ .

Checking the mixed equilibrium:

$$u_2(\frac{1}{3}G + \frac{2}{3}B, g) = \frac{1}{3} \cdot 2 + \frac{2}{3} \cdot 0 = \frac{2}{3}$$

$$u_2(\frac{1}{3}G + \frac{2}{3}B, b) = \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 = \frac{2}{3}$$

$\Rightarrow$  All strategies in  $\Sigma_2$  are best responses.

In the mixed equilibrium each player is indifferent between his mixed strategies.

Each chooses the mixture that makes his opponent indifferent.  $\diamond$

We can refine our prediction applying the notion of strict equilibrium.  $s^*$  is a *strict equilibrium* if for each  $i$ ,  $s_i^*$  is the unique best response to  $s_{-i}^*$ . That is,  $B_i(s_{-i}^*) = \{s_i^*\}$  for all  $i$ .

Strict equilibria seem especially compelling.

But strict equilibria do not exist in all games (unlike Nash equilibria: see Section 1.4.3).

In the previous example, the Nash equilibrium  $(G, g)$  maximizes both players' payoffs. One might be tempted to say that a Nash equilibrium with this property is always the one to focus on. But this criterion is not always compelling:

Example 1.32. *Joint investment.*

Each player can make a safe investment that pays 8 for sure, or a risky investment that pays 9 if the other player joins in the investment and 0 otherwise.

		2	
		<i>r</i>	<i>s</i>
1	R	9,9	0,8
	S	8,0	8,8

The Nash equilibria here are  $(R, r)$ ,  $(S, s)$  and  $(\frac{8}{9}R + \frac{1}{9}S, \frac{8}{9}r + \frac{1}{9}s)$ . Although  $(R, r)$  yields both players the highest payoff, each player might be tempted by the sure payoff of 8 that the safe investment guarantees.  $\diamond$

### 1.4.2 Computing Nash equilibria

Finding the pure strategy Nash equilibria of a normal form game is simple. For each player  $i$  and each pure strategy profile  $s_{-i}$  of  $i$ 's opponents, mark the payoff corresponding to each of  $i$ 's best responses. The boxes whose payoffs are all marked correspond to the pure Nash equilibria.

Example 1.33. *A Cournot duopoly.* Two firms, Acme and Tyrell, compete in quantities in  $\{0, 1, 2, 3, 4\}$ . Neither has fixed costs. Acme has a marginal cost of 2 for each of its first 2 units, and a marginal cost of 4 for subsequent units. Tyrell has a marginal cost of 3 for each of its first 2 units, and a marginal cost of 4 for subsequent units. When a total of  $Q$  units are produced, the market price of the good is  $10 - Q$ .

		Tyrell				
		0	1	2	3	4
Acme	0	0,0	0,6	0,10	0, <b>11</b>	0,10
	1	7,0	6,5	5, <b>8</b>	4, <b>8</b>	3,6
	2	12,0	<b>10</b> ,4	<b>8</b> ,6	<b>6</b> ,5	<b>4</b> ,2
	3	<b>13</b> ,0	<b>10</b> ,3	7, <b>4</b>	4,2	1,-2
	4	12,0	8, <b>2</b>	4, <b>2</b>	0,-1	-4,-6

The unique pure Nash equilibrium is  $(2, 2)$ . One can verify that this is also the solution by iterated strict dominance, and thus the unique rationalizable strategy profile.  $\diamond$

Example 1.34. *Example 1.16 revisited.*



		2		
		L	C	Q
1	T	<b>3, 3</b>	0, 0	0, 2
	M	0, 0	<b>3, 3</b>	0, 2
	D	<b>2, 2</b>	<b>2, 2</b>	<b>2, 0</b>

This game has two pure Nash equilibria,  $(T, L)$  and  $(M, C)$ .  $\diamond$

This method extends directly to games with continuous strategy sets—see Section 1.4.5.

To proceed further, it is helpful to derive explicit links between Nash equilibrium and rationalizability.

**Proposition 1.35.** (i) *Any pure strategy used with positive probability in a Nash equilibrium is rationalizable.*

(ii) *If each player has a unique rationalizable strategy, the profile of these strategies is a Nash equilibrium.*

*Proof.* By Theorem 1.28, the strategies in  $R_i \subseteq S_i$  are rationalizable if for each  $i \in \mathcal{P}$  and each  $s_i \in R_i$ , there is a conjecture  $\mu_i \in \Delta S_{-i}$  such that

- (a) the support of  $\mu_i$  is contained in  $R_{-i}$ , and
- (b)  $s_i$  is a best response to  $\mu_i$ .

To prove part (i) of the proposition, let  $\sigma \in (\sigma_1, \dots, \sigma_n)$  be a Nash equilibrium, and let  $R_i$  be the support of  $\sigma_i$ . Each  $s_i \in R_i$  is a best response to  $\sigma_{-i}$  (by Proposition 1.14), so (a) and (b) hold with  $\mu_i$  determined by  $\sigma_{-i}$  (i.e.  $\mu_i(s_{-i}) = \prod_{j \neq i} \sigma_j(s_j)$ ).

Clearly, if all players' rationalizable strategies are unique, they are pure. To prove part (ii) of the proposition, suppose that  $s = (s_1, \dots, s_n)$  is the unique rationalizable strategy profile. Then (a) and (b) imply that  $s_i$  is a best response to  $s_{-i}$ , and so  $s$  is a Nash equilibrium. ■

With Proposition 1.35 as a starting point, one can specify guidelines for computing all mixed strategy Nash equilibria of a game:

- (i) Eliminate pure strategies that are not rationalizable.
- (ii) For each profile of supports, find all equilibria.

Step (ii) is sometimes called the *support enumeration algorithm*. Once the profile of supports is fixed, one identifies all equilibria with this profile of supports by introducing the optimality conditions implied by the supports. For a player to use a certain support, the pure strategies in the support of a player's equilibrium strategy receive the same payoff, which is at least as high as payoffs for strategies outside the support. In this way,


each player's optimality conditions restrict what *the other players'* strategies may be. In addition, restricting all but one player to small supports, especially to singleton supports, imposes direct restrictions on remaining player's best response—see the first two cases of Example 1.36. To sum up, the support enumeration algorithm divides the analysis into a set of exhaustive cases, each of which places additional structure on the problem.

Inevitably, this approach is computationally intensive: if player  $i$  has  $k_i$  strategies, there are  $\prod_{i \in \mathcal{P}} (2^{k_i} - 1)$  possible profiles of supports, and each can have multiple equilibria. (In practice, one fixes the supports of only  $n - 1$  players' strategies, and restricts the  $n$ th player's strategy using the implied optimality conditions—see the examples below.)

Section 1.4.3 discusses the structure of the set of Nash equilibria and algorithms for computing Nash equilibria in two-player games.

*Example 1.36. (Example 1.16 revisited).*

		2		
		L	C	Q
1	T	3,3	0,0	0,2
	M	0,0	3,3	0,2
	D	2,2	2,2	2,0



We showed in Example 1.16 that the players' sets of rationalizable strategies are  $R_1 = \{\sigma_1 \in \Delta S_1: \sigma_1(T) = 0 \text{ or } \sigma_1(M) = 0\}$  and  $R_2 = \Delta S_2$ .

$Q$  is rationalizable because it is a best response to the conjecture  $\mu_2(T) = \mu_2(M) = \frac{1}{2}$ . But since  $Q$  is not a best response to any  $\sigma_1 \in R_1$ ,  $Q$  is never played in a Nash equilibrium. The key point here is that in Nash equilibrium, player 2's conjecture is correct (i.e., place

probability on player 1's actual strategy). Thus, we need not consider any support for  $\sigma_2$  that includes  $Q$ .

Three possible supports for  $\sigma_2$  remain:

$\{L\} \Rightarrow$  1's BR is  $T \Rightarrow$  2's BR is  $L \quad \therefore (T, L)$  is Nash

$\{C\} \Rightarrow$  1's BR is  $M \Rightarrow$  2's BR is  $C \quad \therefore (M, C)$  is Nash

$\{L, C\} \Rightarrow$  Optimality for player 2 requires that  $u_2(\sigma_1, L) \stackrel{(i)}{=} u_2(\sigma_1, C) \stackrel{(ii)}{\geq} u_2(\sigma_1, Q)$

To determine the set of  $\sigma_1$  for which (i) and (ii) hold, look at the picture of  $B_2$ , or compute:

$$\begin{aligned} (i) \quad 3t + 2d = 3m + 2d & \quad (ii) \quad 3m + 2d \geq 2t + 2m \quad (\text{use } t = m, d = 1 - m - t = 1 - 2t) \\ \Rightarrow t = m & \quad \Rightarrow 3t + 2(1 - 2t) \geq 4t \\ & \quad \therefore t = m \leq \frac{2}{5} \end{aligned}$$

But the only strategy in  $B_1(\Delta S_2) = \{\sigma_1 \in \Delta S_1 : \sigma_1(T) = 0 \text{ or } \sigma_1(M) = 0\}$  that puts equal weight on  $T$  and  $M$  is the pure strategy  $D$ . Player 1 is willing to play  $D$  if

$$\begin{aligned} u_1(D, \sigma_2) \geq u_1(T, \sigma_2) & \Leftrightarrow l \leq \frac{2}{3}, \text{ and} \\ u_1(D, \sigma_2) \geq u_1(M, \sigma_2) & \Leftrightarrow c \leq \frac{2}{3} \end{aligned}$$

Since we are assuming that player 2's strategy has support  $\{L, C\}$ , her strategy must be of the form  $L + (1 - \alpha)C$  with  $\alpha \in [\frac{1}{3}, \frac{2}{3}]$ .

Thus each mixed strategy profile  $(D, \alpha L + (1 - \alpha)C)$  with  $\alpha \in [\frac{1}{3}, \frac{2}{3}]$  is a Nash equilibrium.

(It is worth tracking which best response conditions were used and how they were used in each step of finding this component of Nash equilibria. First, the optimality of player 2's strategy  $\sigma_2$  with support  $\{L, C\}$  was used to obtain the restrictions  $t = m \leq \frac{2}{5}$  on player 1's strategy  $\sigma_1$ . Second, the optimality of  $\sigma_1$  against such a  $\sigma_2$  was used to obtain a further restriction on player 1's strategy, namely that he plays  $D$ . Third, the fact that  $D$  is optimal for player 1 against  $\sigma_2$  was used to obtain restrictions on player 2's strategy, namely that  $l, c \leq \frac{2}{3}$ . Combining these last restrictions with the fact that  $\sigma_2$  has support  $\{L, C\}$  let us conclude that  $\sigma_2$  takes the form  $L + (1 - \alpha)C$  with  $\alpha \in [\frac{1}{3}, \frac{2}{3}]$ .)

In conclusion, the set of Nash equilibria consists of three *connected components*:  $\{(T, L)\}$ ,  $\{(M, C)\}$ , and  $\{(D, \alpha L + (1 - \alpha)C) : \alpha \in [\frac{1}{3}, \frac{2}{3}]\}$ .  $\diamond$

*Example 1.37. Zeeman's (1980) game.*


		2		
		A	B	C
1	A	0,0	6,-3	-4,-1
	B	-3,6	0,0	5,3
	C	-1,-4	3,5	0,0

Since the game is symmetric, both players have the same incentives as a function of the opponent's behavior.

$$A \geq B \Leftrightarrow 6b - 4c \geq -3a + 5c \Leftrightarrow a + 2b \geq 3c;$$

$$A \geq C \Leftrightarrow 6b - 4c \geq -a + 3b \Leftrightarrow a + 3b \geq 4c;$$

$$B \geq C \Leftrightarrow -3a + 5c \geq -a + 3b \Leftrightarrow 5c \geq 2a + 3b.$$



Now consider each possible support of player 1's equilibrium strategy.

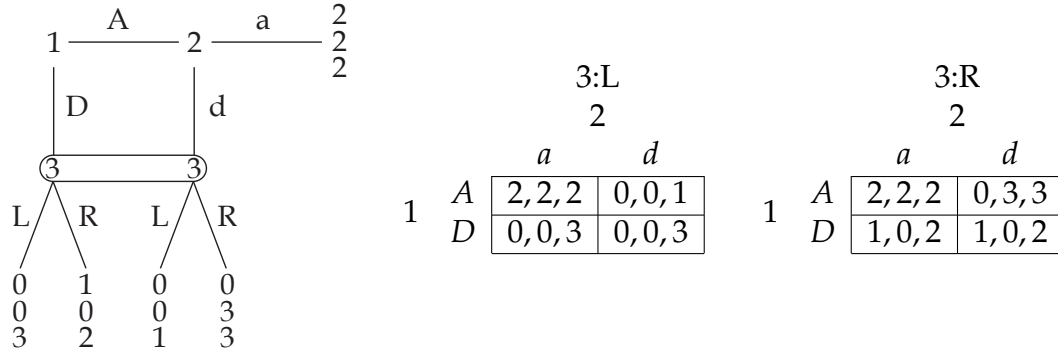
- A     Implies that 2 plays A, and hence that 1 plays A. Equilibrium.
- B     Implies that 2 plays A, and hence that 1 plays A.
- C     Implies that 2 plays B, and hence that 1 plays A.
- A,B   Implies that 2 plays A, and hence that 1 plays A.
- A,C   This allows many best responses for player 2, but the only one that makes both A and C a best response for 1 is  $\frac{4}{5}A + \frac{1}{5}C$ , which is only a best response for 2 if 1 plays  $\frac{4}{5}A + \frac{1}{5}C$  himself. Equilibrium.
- B,C   Implies that 2 plays A, B or a mixture of the two, and hence that 1 plays A.
- all    This is only optimal for 1 if 2 plays  $\frac{1}{3}A + \frac{1}{3}B + \frac{1}{3}C$ , which 2 is only willing to do if 1 plays  $\frac{1}{3}A + \frac{1}{3}B + \frac{1}{3}C$ . Equilibrium.

∴ There are three Nash equilibria: (A, A)

$$\left(\frac{4}{5}A + \frac{1}{5}C, \frac{4}{5}A + \frac{1}{5}C\right)$$

$$\left(\frac{1}{3}A + \frac{1}{3}B + \frac{1}{3}C, \frac{1}{3}A + \frac{1}{3}B + \frac{1}{3}C\right) \quad \diamond$$

Example 1.38. Selten's (1975) horse.



Consider all possible mixed strategy supports for players 1 and 2:

- (D, d) Implies that 3 plays L. Since 1 and 2 are also playing best responses, this is a Nash equilibrium.
- (D, a) Implies that 3 plays L, which implies that 1 prefers to deviate to A.
- (D, mix) Implies that 3 plays L, which with 2 mixing implies that 1 prefers to deviate to A.
- (A, d) Implies that 3 plays R, which implies that 1 prefers to deviate to D.
- (A, a) 1 and 2 are willing to do this if  $\sigma_3(L) \geq \frac{1}{3}$ . Since 3 cannot affect his payoffs given the behavior of 1 and 2, these are Nash equilibria.
- (A, mix) 2 only mixes if  $\sigma_3(L) = \frac{1}{3}$ ; but if 1 plays A and 2 mixes, 3 strictly prefers R – a contradiction.
- (mix, a) Implies that 3 plays L, which implies that 1 strictly prefers A.
- (mix, d) If 2 plays d, then for 1 to be willing to mix, 3 must play L; this leads 2 to deviate to a.
- (mix, mix) Notice that 2 can only affect her own payoffs when 1 plays A. Hence, for 2 to be indifferent,  $\sigma_3(L) = \frac{1}{3}$ . Given this, 1 is willing to mix if  $\sigma_2(d) = \frac{2}{3}$ . Then for 3 to be indifferent,  $\sigma_1(D) = \frac{4}{7}$ . This is a Nash equilibrium.

$\therefore$  There are three components of Nash equilibria:  $(D, d, L)$

$$(A, a, \sigma_3(L) \geq \frac{1}{3})$$

$$(\frac{3}{7}A + \frac{4}{7}D, \frac{1}{3}a + \frac{2}{3}d, \frac{1}{3}L + \frac{2}{3}R) \quad \diamond$$

### 1.4.3 Existence of Nash equilibrium and structure of the equilibrium set

A classic theorem of Nash shows that every finite normal form game admits at least one equilibrium, possibly in mixed strategies. Thus the Nash equilibrium concept always provides at least one prediction of play.

**Theorem 1.39** (Nash (1950)).

*Any finite normal form game has at least one (possibly mixed) Nash equilibrium.*

Equilibrium existence results are usually proved by means of *fixed point theorems*.

**Theorem 1.40** (Brouwer (1912)). *Let  $X \subseteq \mathbb{R}^n$  be nonempty, compact, and convex. Let the function  $f : X \rightarrow X$  be continuous. Then there exists an  $x \in X$  such that  $x = f(x)$ .*

**Theorem 1.41** (Kakutani (1941)). *Let  $X \subseteq \mathbb{R}^n$  be nonempty, compact, and convex. Let the correspondence  $f : X \rightrightarrows X$  be nonempty, upper hemicontinuous, and convex valued. Then there exists an  $x \in X$  such that  $x \in f(x)$ .*

The first published proof of Theorem 1.39 is from Nash (1950):

Let  $B_i : \Sigma_{-i} \rightrightarrows \Sigma_i$  be player  $i$ 's best response correspondence.

Define  $B : \Sigma \rightrightarrows \Sigma$  by  $B(\sigma) = (B_1(\sigma_{-1}), \dots, B_n(\sigma_{-n}))$ .

Then  $\sigma$  is a Nash equilibrium  $\Leftrightarrow \sigma \in B(\sigma)$ .

Show that  $B$  is nonempty, upperhemicontinuous, and convex-valued.

Kakutani's fixed point theorem then implies that a Nash equilibrium exists.

A slicker proof of Theorem 1.39, presented in Nash (1951), builds on Brown and von Neumann (1950):

Define  $C : \Sigma \rightarrow \Sigma$  by

$$C_{i,s_i}(\sigma) = \frac{\sigma_{i,s_i} + [u_i(s_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i})]_+}{1 + \sum_{s'_i \in S_i} [u_i(s'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i})]_+}.$$

Then  $\sigma$  is a Nash equilibrium  $\Leftrightarrow \sigma = C(\sigma)$ . (The interpretation of  $C$  is that each player adds probability weight to pure strategies that do better than his current mixed strategy, and then normalizes so that the probability weights again sum to 1. At a fixed point, no player has a pure strategy that does better than his current mixed strategy; this is another way of saying that all pure strategies that he uses are optimal.)

Since  $C : \Sigma \rightarrow \Sigma$  is continuous, Brouwer's fixed point theorem implies that a Nash equilibrium exists.

We next consider the structure of the set of Nash equilibria. The strongest results are available for two-player games. We call a two-player game *nondegenerate* if for each strategy  $\sigma_i$ , the number of pure best responses to  $\sigma_i$  is at most the cardinality of the support of  $\sigma_i$ . (For instance, if  $\sigma_i$  uses two strategies, no best response to  $\sigma_i$  may use three or more strategies.) Also, we define a *polytope* to be the convex hull of a finite number of points.

**Theorem 1.42.** *Let  $G$  be a finite two-player normal form game.*

- (i) *If  $G$  is nondegenerate, then the number of Nash equilibria of  $G$  is finite and odd, and in each Nash equilibrium, the supports of the two players' strategies have equal cardinality.*

(ii) In general, the set of Nash equilibria is the union of a finite number of polytopes.

Part (i) of the theorem describes the structure of the Nash set in “typical” two-player games. It is due to Lemke and Howson (1964), who also provided an algebraic proof of the existence of Nash equilibrium in two-player games, and the best early algorithm for computing a Nash equilibrium of a two-player game. Part (ii) of the theorem implies that the complete set of Nash equilibria of a two-player game can be described by specifying the extreme points of the equilibrium polytopes. Recent algorithms for finding all Nash equilibria of a two-player game are presented in Avis et al. (2010). A web-based implementation of one of the algorithms from this paper is posted at [cgi.csc.liv.ac.uk/~rahul/bimatrix\\_solver/](http://cgi.csc.liv.ac.uk/~rahul/bimatrix_solver/).

The next result is a structure theorem for games with arbitrary numbers of players. A normal form game is defined by the number of players  $n$ , a set of pure strategies  $S_i$  for each player, and a specification of payoffs  $\{u(s)\}_{s \in S} \in \mathbb{R}^{n \times \#S}$ . A property holds in *generic* finite normal form games if for any choice of the number of players and the pure strategy sets, the set of payoff specifications for which the property fails to hold has measure zero in  $\mathbb{R}^{n \times \#S}$ . Notice that on its own, a result about generic games does not tell us whether the property in question holds in any particular game.

**Theorem 1.43.**

- (i) In generic finite normal form games, the number of Nash equilibria is finite and odd.
- (ii) In any finite normal form game, the set of Nash equilibria is the union of a finite number of connected components (i.e., maximal connected subsets).

Part (i) of the theorem is due to Wilson (1971). Part (ii) is due to Kohlberg and Mertens (1986), who conclude this from a stronger property, namely that the set of Nash equilibria is *semialgebraic*. This property is also satisfied by the equilibrium sets defined by other equilibrium concepts, including sequential equilibrium in extensive form games (Section 2.4). For applications and background, see Blume and Zame (1994) and Coste (2002).

The problem of finding a Nash equilibrium, even in two-player games, is known to be computationally hard, in the sense that the running time of any algorithm that always accomplishes this must grow exponentially in the size of the game (see Daskalakis et al. (2009) and Chen et al. (2009)). This raises doubts about the unrestricted use of Nash equilibrium to predict behavior in games. We discuss this point further in the next section.

1.4.4 Interpretations of Nash equilibrium

*Example 1.44.* (The Good Restaurant, Bad Restaurant game (Example 1.31))

		2		
		$g$	$b$	
1	$G$	2, 2	0, 0	NE: $(G, g)$ $(B, b)$ $(\frac{1}{3}G + \frac{2}{3}B, \frac{1}{3}g + \frac{2}{3}b)$ . $\diamond$
	$B$	0, 0	1, 1	

Example 1.45. Matching Pennies.

		2		
		h	t	
1	H	1, -1	-1, 1	unique NE: $(\frac{1}{2}H + \frac{1}{2}T, \frac{1}{2}h + \frac{1}{2}t)$ . $\diamond$
	T	-1, 1	1, -1	

Nash equilibrium is a minimal condition for self enforcing behavior.

This explains why we *should not* expect players to behave in a way that is *not* Nash, but not why we should expect players to coordinate on a Nash equilibrium.

**Justifications of equilibrium knowledge: why expect correct conjectures?**

There is no general justification for assuming equilibrium knowledge. But justifications can be found in certain specific instances:

- (i) Coordination of play by a mediator.

If a mediator proposes a Nash equilibrium, no player can benefit from deviating. Of course, this only helps if there actually is a mediator.

- (ii) Pre-play agreement.

Even in games with a Pareto dominant Nash equilibrium, communication to coordinate on the equilibrium may not be credible. Consider this variant on the joint investment game from Example 1.32 (Aumann (1990)). Notice that in this game, a player is always better off when his opponent takes the risky action.

		2	
		r	s
1	R	9, 9	0, 8
	S	8, 0	7, 7

In order to coordinate on the Pareto optimum, the players might communicate to one another their intentions to play  $R$  and  $r$ . However, each player wants his opponent to play the risky action regardless of what he himself plans to play. So if player 1 tells player 2 that he will play  $R$ , player 2 could interpret this as player 1 trying to get 8 rather than 7 when he plays  $S$ .

Also, if “pre-play” communication is important, it should be included as part of the game. Then any communication that happens must itself be part of the equilibrium, which brings us back to our starting point. These *cheap talk* games, in which communication has no direct effect on payoffs, have “babbling equilibria”



in which all communication is ignored. For an overview of cheap talk games, see Farrell and Rabin (1996).

(iii) Focal points (Schelling (1960)). Something about the game makes some Nash equilibrium the obvious choice about how to behave.

ex: coordinating on the good restaurant.

ex: meeting in NYC at the information booth at Grand Central Station at noon.

(iv) Learning/Evolution: If players recurrently face the same game, they may find their way from arbitrary initial behavior to Nash equilibrium.

Heuristic learning: Small groups of players, typically employing rules that condition on the empirical distribution of past play (Young (2004), Hart (2005))

Evolutionary game theory: Large populations of agents using myopic updating rules (Sandholm (2009, 2010))

In some classes of games (that include the two examples above), many learning and evolutionary processes do converge to Nash equilibrium.

But there is no general guarantee of convergence:

Many games lead to cycling or chaotic behavior, and in some games any “reasonable” dynamic process fails to converge to equilibrium (Shapley (1964), Hofbauer and Swinkels (1996), Hart and Mas-Colell (2003)).

Indeed, results on the difficulty of computing Nash equilibrium imply that no learning procedure, reasonable or not, can find a Nash equilibrium in a reasonable amount of time in every large game (Daskalakis et al. (2009), Chen et al. (2009)).

Some games introduced in applications are known to have poor convergence properties (Hopkins and Seymour (2002), Lahkar (2011)).

In fact, evolutionary game theory models do not even support the elimination of strictly dominated strategies in all games (Hofbauer and Sandholm (2011)).

***Interpretation of mixed strategy Nash equilibrium: why mix in precisely the way that makes your opponents indifferent?***

In the unique equilibrium of Matching Pennies, player 1 is indifferent among all of his mixed strategies. He chooses  $(\frac{1}{2}, \frac{1}{2})$  because this makes player 2 indifferent. Why should we expect player 1 to behave in this way?

(i) Deliberate randomization

Sometimes it makes sense to expect players to deliberately randomize (ex.: poker).

In *zero-sum games* (Section 1.6), randomization can be used to ensure that you obtain at least the equilibrium payoff regardless of how opponents behave:

In a mixed equilibrium, you randomize to make your opponent indifferent between her strategies. In a zero-sum game, this implies that *you* are indifferent between your opponent's strategies. This implies that you do not care if your opponent finds out your randomization probabilities in advance, as this does not enable her to take advantage of you.

(ii) Mixed equilibrium as equilibrium in conjectures

One can interpret  $\sigma_i^*$  as describing the conjecture that player *i*'s opponents have about player *i*'s behavior. The fact that  $\sigma_i^*$  is a mixed strategy then reflects the opponents' uncertainty about how *i* will behave, even if *i* is not actually planning to randomize.

But as Rubinstein (1991) observes, this interpretation

“... implies that an equilibrium does not lead to a prediction (statistical or otherwise) of the players' behavior. Any player *i*'s action which is a best response given his expectation about the other players' behavior (the other  $n - 1$  strategies) is consistent as a prediction for *i*'s action (this might include actions which are outside the support of the mixed strategy). This renders meaningless any comparative statics or welfare analysis of the mixed strategy equilibrium. . .

(iii) Mixed equilibria as time averages of play: fictitious play (Brown (1951))

Suppose that the game is played repeatedly, and that in each period, each player chooses a best response to the time average of past play.

Then in certain classes of games, the time average of each players' behavior converges to his part in some Nash equilibrium strategy profile.

(iv) Mixed equilibria as population equilibria (Nash (1950))

Suppose that there is one population for the player 1 role and another for the player 2 role, and that players are randomly matched to play the game.

If half of the players in each population play Heads, no one has a reason to deviate. Hence, the mixed equilibrium describes stationary *distributions* of *pure* strategies in each population.

(v) Purification: mixed equilibria as pure equilibria of games with payoff uncertainty (Harsanyi (1973))

*Example 1.46. Purification in Matching Pennies.* Suppose that while the Matching Pennies payoff bimatrix gives player's approximate payoffs, players' actual payoffs also contain small terms  $\varepsilon_H, \varepsilon_h$  representing a bias toward playing heads, and that each player only

knows his own bias. (The formal framework for modeling this situation is called a *Bayesian game*—see Section 4.)

		2	
		$h$	$t$
1	$H$	$1 + \varepsilon_H, -1 + \varepsilon_h$	$-1 + \varepsilon_H, 1$
	$T$	$-1, 1 + \varepsilon_h$	$1, -1$

Specifically, suppose that  $\varepsilon_H$  and  $\varepsilon_h$  are independent random variables with  $\mathbb{P}(\varepsilon_H > 0) = \mathbb{P}(\varepsilon_H < 0) = \frac{1}{2}$  and  $\mathbb{P}(\varepsilon_h > 0) = \mathbb{P}(\varepsilon_h < 0) = \frac{1}{2}$ . Then it is a strict Nash equilibrium for each player to follow his bias. From the ex ante point of view, the distribution over actions that this equilibrium generates in the original normal form game is  $(\frac{1}{2}H + \frac{1}{2}T, \frac{1}{2}h + \frac{1}{2}t)$ .

Harsanyi (1973) shows that any mixed equilibrium can be purified in this way. This includes not only “reasonable” mixed equilibria like that in Matching Pennies, but also “unreasonable” ones like those in coordination games.  $\diamond$

#### 1.4.5 Nash equilibrium in games with continuous strategy sets

Many games appearing in economic applications have continuous strategy sets—for instance, an interval of real numbers. For better or worse, analyses of these games often focus on pure strategy equilibrium.

Let  $\mathcal{G} = \{\mathcal{P}, \{\mathcal{S}_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$  be a game in which the strategy sets  $\mathcal{S}$  may be finite or continuous. Player  $i$ 's *pure best-response correspondence*  $b_i: \mathcal{S}_{-i} \Rightarrow \mathcal{S}_i$  is defined by

$$b_i(s_{-i}) = \operatorname{argmax}_{s_i \in \mathcal{S}_i} u_i(s_i, s_{-i}).$$

Pure strategy profile  $s \in \mathcal{S}$  is a Nash equilibrium if  $s_i \in b_i(s_{-i})$  for all  $i \in \mathcal{P}$ .

In a two-player game in which  $\mathcal{S}_1$  and  $\mathcal{S}_2$  are intervals, one can find all pure Nash equilibria by graphing each player's best response correspondence and finding all points at which the graphs intersect. This is the direct analogue of the procedure for finding all pure Nash equilibria of finite normal form games (Section 1.4.2).

*Example 1.47. Cournot duopoly.* Two firms separately choose what quantity of a homogeneous good to produce. All output is sent to a market maker who sets the price that clears the market.

$q_i \in \mathcal{S}_i = [0, \infty)$	strategies (quantities produced)
$\pi_i(q_i, q_j) = q_i P(q_i + q_j) - C_i(q_i)$	profits
$P(q_T) = \max\{a - q_T, 0\}$	linear inverse demand function
$C_i(q_i) = cq_i, c < a$	identical linear costs


Firm  $i$ 's profit function is

$$\pi_i(q_i, q_j) = \begin{cases} q_i(a - q_i - q_j - c) & \text{if } q_i + q_j \leq a, \\ -cq_i & \text{if } q_i + q_j > a. \end{cases}$$

Thus its best response function is

$$b_i(q_j) = \begin{cases} \frac{a-c-q_j}{2} & \text{if } q_j \leq a-c, \\ 0 & \text{otherwise.} \end{cases}$$

The graphs of the best response correspondences are shown below. The graphs intersect exactly once, at the solution to the system  $q_1 = \frac{a-c-q_2}{2}$ ,  $q_2 = \frac{a-c-q_1}{2}$ . Thus the game's unique Nash equilibrium is  $(q_1, q_2) = (\frac{a-c}{3}, \frac{a-c}{3})$ . ♦



*Example 1.48. Bertrand duopoly.* Each firm chooses a price. The firm choosing the lower price,  $p_{\min}$ , sells  $Q(p_{\min})$  units at this price, producing each unit at marginal cost  $c$ . The firm choosing the higher price produces nothing and sells nothing. If the firms charge the same price, they split the market. (A firm does not produce until it knows how many units it will sell.)

- $p_i \in S_i = [0, \infty)$  strategies (prices)
- $Q(p_{\min}) = \max\{a - p_{\min}, 0\}$  linear demand function
- $c (< a)$  marginal production cost


Firm  $i$ 's profit function is

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

Let  $p^m = \frac{a+c}{2}$  be the monopoly price. Then firm  $i$ 's best response correspondence is

$$b_i(p_j) = \begin{cases} (p_j, \infty) & \text{if } p_j < c, \\ [p_j, \infty) & \text{if } p_j = c, \\ \emptyset & \text{if } c < p_j \leq p^m, \\ \{p^m\} & \text{if } p_j > p^m. \end{cases}$$

The graphs of the best response correspondences are shown below. Their unique intersection point is the game's unique Nash equilibrium,  $(p_1, p_2) = (c, c)$ . ♦



### *Existence of Nash equilibrium in games with continuous strategy sets*

The basic existence theorem for games with continuous strategy sets requires continuous payoff functions and quasiconcavity of payoffs in own strategies. Together these properties ensure that the pure best response correspondences have the properties needed to apply the Kakutani fixed point theorem.

**Theorem 1.49.** *Suppose that (i) each  $S_i$  is a compact, convex subset of  $\mathbb{R}^k$ , (ii) each  $u_i$  is continuous in  $s$ , and (iii) each  $u_i$  is quasiconcave in  $s_i$ . Then the game  $G$  has at least one pure strategy Nash equilibrium.*

Theorem 1.49 applies to the Cournot game (check the quasiconcavity condition!) if one restricts the strategy sets to make them compact intervals (e.g., by not allowing quantities above  $a$ ). The Bertrand game does not have continuous payoffs. However, payoffs are quasiconcave in own strategies (check!), and in addition the game satisfies a property called “better reply security”. Reny (1999) shows that these properties ensure the existence of a pure strategy Nash equilibrium.

Dasgupta and Maskin (1986a) and Reny (1999, 2016) are the basic references on existence of Nash equilibrium in games with discontinuous payoffs.

A closely related topic concerns existence of equilibrium in environments with discontinuous payoffs via *endogenous sharing rules*. In some applications, the map from strategy profiles to payoff profiles is naturally represented by a correspondence. For instance, if firms competing in prices charge the same price, then the consumers may choose between the firms in a variety of ways, leading to a variety of possible payoff profiles for the firms. Then one can ask whether there is a “sharing rule” that generates a game (with single-valued payoff functions) in which a Nash equilibrium exists. Basic references here are Simon and Zame (1990) and Jackson et al. (2002).

One can also consider existence of mixed strategy Nash equilibrium. (Note that a mixed strategy in  $G$  is a probability measure on  $S_i$ .) The basic existence result for games with continuous payoffs drops the requirement of quasiconcavity of payoffs in own strategies. It is proved using an extension of Kakutani’s fixed point theorem to suitable infinite-dimensional spaces (Fan (1952), Glicksberg (1952)).

**Theorem 1.50.** *Suppose that (i) each  $S_i$  is a compact, convex subset of  $\mathbb{R}^k$  and (ii) each  $u_i$  is continuous in  $s$ . Then the game  $G$  has at least one (possibly mixed) Nash equilibrium.*

For results on existence of mixed strategy equilibrium in games with discontinuous payoffs, again see Dasgupta and Maskin (1986a) and Reny (1999, 2016).

## 1.5 Correlated Equilibrium

### 1.5.1 Definition and examples


*Example 1.51.* Consider the following normal form game:

		2	
		L	R
1	T	5, 5	2, 6
	B	6, 2	1, 1

Nash equilibria/payoffs:  $(T, R) \Rightarrow (2, 6)$ ,  $(B, L) \Rightarrow (6, 2)$ ,  $(\frac{1}{2}T + \frac{1}{2}B, \frac{1}{2}L + \frac{1}{2}R) \Rightarrow (3\frac{1}{2}, 3\frac{1}{2})$

If the players can observe a toss of a fair coin before play, they can obtain equilibrium payoffs of  $(4, 4)$  using the correlated strategy  $\frac{1}{2}TR + \frac{1}{2}BL$ . More generally, any point in the convex hull of  $\{(2, 6), (6, 2), (3\frac{1}{2}, 3\frac{1}{2})\}$  can be achieved in equilibrium using some three-outcome randomization device.

We can imagine a mediator using such a device to determine which equilibrium he tells the players to play. If a player expects his opponent to obey the announcement, then the fact that the announced strategy profile is a Nash equilibrium implies that it is optimal for the player to obey the announcement himself.



Can a mediator use a randomizing device to generate expected payoffs greater than 8 in equilibrium?

Yes, by only telling each player what he is supposed to play, not what the opponent is supposed to play.

Suppose the device specifies  $TL$ ,  $TR$ , and  $BL$  each with probability  $\frac{1}{3}$ , so that  $\rho = \frac{1}{3}TL + \frac{1}{3}TR + \frac{1}{3}BL$ .

		2	
		L	R
1	T	$\frac{1}{3}$	$\frac{1}{3}$
	B	$\frac{1}{3}$	0

However, player 1 is only told whether his component is  $T$  or  $B$ , and player 2 is only told if her component is  $L$  or  $R$ :

The correlated strategy  $\rho$  generates payoffs of  $(4\frac{1}{3}, 4\frac{1}{3})$ .

Moreover, we claim that both players obeying constitutes an equilibrium:

Suppose that player 2 plays as prescribed, and consider player 1's incentives.

If player 1 sees  $B$ , he knows that player 2 will play  $L$ , so his best response is  $B$  (since  $6 > 5$ ).

If player 1 sees  $T$ , he believes that player 2 is equally likely to play  $L$  or  $R$ , and so  $T$  is a best response for player 1 (since  $3\frac{1}{2} = 3\frac{1}{2}$ ).

By symmetry, this is an equilibrium.  $\diamond$

Let  $G = \{\mathcal{P}, \{S_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$  be a normal form game.

For any correlated strategy  $\rho \in \Delta(\prod_{j \in \mathcal{P}} S_j)$  with  $\rho(s_i) = \sum_{s_{-i} \in S_{-i}} \rho(s_i, s_{-i}) > 0$ , let  $\rho(s_{-i}|s_i) = \frac{\rho(s_i, s_{-i})}{\rho(s_i)}$ .

We call  $\rho$  a *correlated equilibrium* (Aumann (1974)) if

$$\begin{aligned} \sum_{s_{-i} \in S_{-i}} \rho(s_{-i}|s_i) u_i(s_i, s_{-i}) &\geq \sum_{s_{-i} \in S_{-i}} \rho(s_{-i}|s_i) u_i(s'_i, s_{-i}) \quad \text{for all } s_i \in S_i \text{ with } \rho(s_i) > 0, s'_i \in S_i, \text{ and } i \in \mathcal{P} \\ \Leftrightarrow \sum_{s_{-i} \in S_{-i}} \rho(s_i, s_{-i}) u_i(s_i, s_{-i}) &\geq \sum_{s_{-i} \in S_{-i}} \rho(s_i, s_{-i}) u_i(s'_i, s_{-i}) \quad \text{for all } s_i, s'_i \in S_i, \text{ and } i \in \mathcal{P} \end{aligned}$$

In words, the first condition says that if  $i$  receives signal  $s_i$  and opponents obey their signals,  $i$  cannot benefit from disobeying his signal. The second condition is mathematically simpler—see below.

**Observation 1.52.** *Correlated strategy  $\rho$  is equivalent to a Nash equilibrium if and only if  $\rho$  is a correlated equilibrium and is a product measure (i.e., the players' signals are independent).*

*Example 1.53.* (Describing the set of correlated equilibria)

		2	
		$g$	$b$
1	G	3, 3	0, 5
	B	5, 0	-4, -4

The Nash equilibria are  $(G, b)$ ,  $(B, g)$ , and  $(\frac{2}{3}G + \frac{1}{3}B, \frac{2}{3}g + \frac{1}{3}b)$ .

The set of correlated equilibria of any game is an intersection of a finite number of half spaces. It is therefore a *polytope*: that is, the convex hull of a finite number of points.

The constraint ensuring that player 1 plays  $G$  when told is  $3\rho_{Gg} + 0\rho_{Gb} \geq 5\rho_{Gg} - 4\rho_{Gb}$ , or, equivalently,  $2\rho_{Gb} \geq \rho_{Gg}$ . We compute the constraints for strategies  $B$ ,  $g$ , and  $b$  similarly, and list them along with the nonnegativity constraints:

$$(1) \quad 2\rho_{Gb} \geq \rho_{Gg}$$

$$(5) \quad \rho_{Gg} \geq 0$$




- |                                   |                        |
|-----------------------------------|------------------------|
| (2) $\rho_{Bg} \geq 2\rho_{Bb}$ ; | (6) $\rho_{Gb} \geq 0$ |
| (3) $2\rho_{Bg} \geq \rho_{Gg}$ ; | (7) $\rho_{Bg} \geq 0$ |
| (4) $\rho_{Gb} \geq 2\rho_{Bb}$ ; | (8) $\rho_{Bb} \geq 0$ |


A correlated equilibrium is a correlated strategy which satisfies these eight inequalities along with the equality  $\rho_{Gg} + \rho_{Gb} + \rho_{Bg} + \rho_{Bb} = 1$ .

The equality constraint means that we are essentially working in three dimensions, since the value of  $\rho_{Bb}$  is determined by the value of the other three coordinates. We can therefore sketch the polytope defined by these inequalities in the pyramid-shaped simplex (see the figure on the next page). The polytope has 5 vertices, 9 edges, and 6 faces (and so satisfies Euler's formula:  $F - E + V = 2$ ). Vertices, faces, and edges are defined by (at least) one inequality, two inequalities, and three inequalities, respectively. The 6 faces are defined by inequalities (1)-(5) and (8) (since (1) and (4) (along with the equality) imply (6), and (2) and (3) imply (7)). The 9 edges are defined by the pairs of inequalities (8) and (3), (3) and (2), (2) and (5), (3), and (1), (2) and (4), (5) and (8), (8) and (1), (1) and (4), and (4) and (5). The five vertices of the polytope are at points where the following combinations of inequalities bind:

	$\rho = (\rho_{Gg}, \rho_{Gb}, \rho_{Bg}, \rho_{Bb})$	binding inequalities
$\alpha$	$(0, 1, 0, 0)$	(2),(3),(5),(7), and (8)
$\beta$	$(0, 0, 1, 0)$	(1),(4),(5),(6), and (8)
$\gamma$	$(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}, 0)$	(1),(3), and (8)
$\delta$	$(0, \frac{2}{3}, \frac{2}{3}, \frac{1}{3})$	(2),(4), and (5)
$\varepsilon$	$(\frac{4}{9}, \frac{2}{9}, \frac{2}{9}, \frac{1}{9})$	(1),(2),(3), and (4)

The set of correlated equilibria is drawn below. Notice that vertices  $\alpha$  and  $\beta$  are the pure Nash equilibria, and that vertex  $\varepsilon$  is the mixed Nash equilibrium, at which all four incentive constraints bind.





### 1.5.2 Interpretation

Some game theorists feel that correlated equilibrium is the fundamental equilibrium concept for normal form games:

(i) Mathematical simplicity

The set of correlated equilibria is defined by a finite number of *linear* inequalities. It is therefore a *polytope* (the convex hull of a finite number of points).

Existence of correlated equilibrium can be proved using results from the theory of linear inequalities (Hart and Schmeidler (1989)).

(The basic results from the theory of linear inequalities are much easier to prove than fixed point theorems, and each can be derived from the others without too much difficulty. These results include the Minmax Theorem (Section 1.6), the separating hyperplane theorem for polytopes (Section 1.3.2), Linear Programming Duality, Farkas's Lemma, . . .)

(ii) Learnability

There are procedures which enable players to learn to play correlated equilibria regardless of the game they are playing, at least in the sense of time-averaged play (Foster and Vohra (1997), Hart and Mas-Colell (2000), Young (2004)). (Observe that time-averaging of pure strategy profiles naturally generates correlated strategies rather than mixed strategy profiles.)

There are no such learning procedures for Nash equilibrium.

(iii) Applicability

In some applications players may have the opportunity to design a correlated signal structure. Moreover, the world offers a multitude of correlated signals, and different individuals focus on different ones.

The *revelation principle for normal form games* (Proposition 4.12) says that equilibria

arising from *any* correlated signal structure can be expressed as correlated equilibria. Thus, correlated equilibrium provides a way of implicitly including all signaling structures, possibly provided by a mediator, directly in the solution concept. See Myerson (1991, Chapter 6) for an excellent textbook discussion of this and related ideas.

### *Subjective correlated equilibrium, a posteriori equilibrium, and rationalizability*

We have defined a correlated equilibrium as a probability distribution  $\rho \in \Delta(S)$  over strategy profiles with the property that if player  $i$  believes others will obey their signals, and he himself receives signal  $s_i$ , then it is optimal for player  $i$  to play  $s_i$ . Of course, optimality here depends on player  $i$ 's conditional beliefs  $\rho(\cdot|s_i)$  about others' signals given his own signals.

To obtain a more general solution concept, one can allow each player  $i$  to have his own prior beliefs  $\rho_i \in \Delta(S)$  about the joint distribution of signals. A *subjective correlated equilibrium* is a profile  $(\rho_1, \dots, \rho_n)$  of correlated strategies that satisfies the same optimality condition described in words above; the only difference is that when player  $i$  receives signal  $s_i$ , his optimality condition depends on  $\rho_i(\cdot|s_i)$ , his conditional beliefs based on his own prior  $\rho_i$ .

For what follows we require a somewhat more restrictive solution concept. It may happen that player  $i$ 's beliefs  $\rho_i$  place zero probability on his receiving a particular signal  $s'_i$ . Subjective correlated equilibrium says nothing about what player  $i$  would do if he received such a signal. In contrast, a *a posteriori equilibrium* requires player  $i$  to specify posterior beliefs  $\rho_i(\cdot|s'_i)$  even for signals he assigned prior probability zero, and then to play a best response to such beliefs. In this case the best response may not be  $s'_i$  itself—for instance, it cannot be if  $s'_i$  is a strictly dominated strategy. In an a posteriori equilibrium, player  $i$ 's opponents correctly anticipate what strategy player  $i$  will play in response to signal  $s'_i$ , so their optimality conditions account for any “disobedience”. (Beliefs after probability zero events also play a basic role in analyzing extensive form games—see Section 2.4.)

Brandenburger and Dekel (1987) make the following observation: a pure strategy  $s_i$  is rationalizable if and only if there is an a posteriori equilibrium in which  $s_i$  is played. The proof of this result is a matter of manipulating definitions. If a strategy  $s_i$  is rationalizable, it is a best response to some posterior beliefs  $\rho_i(\cdot|s_i)$ , so these beliefs describe an a posteriori equilibrium. The converse statement is proved by showing that given an a posteriori equilibrium, the sets of strategies each player would play after some signal satisfy the conditions for rationalizability from Theorem 1.28(i). (This result takes some time to digest, but in the end it really is almost trivial.)

One interpretation of this result is that the power of equilibrium concepts to limit predictions further than rationalizability is entirely dependent on the assumption of agreement of beliefs. The setup here separates players' beliefs (about the joint distribution of signals) from players' “strategies” (their maps from signals to normal form game strategies). If players are allowed to hold different beliefs, then the assumption that players best

respond to opponents' "strategies" yields the same conclusions as the assumption of common knowledge of rationality.

For generalizations of these ideas to Bayesian games, see Battigalli and Siniscalchi (2003).

## 1.6 The Minmax Theorem

### *Worst-case analysis*

The analyses so far have assumed that players are Bayesian rational: they form beliefs about opponent's choices of strategies and then maximize expected utility given those beliefs. We start our analysis of zero-sum games by considering an alternative decision criterion based on *worst-case analysis*: each player chooses a mixed strategy that maximizes his minimal expected utility, where the minimum is taken over the possible strategies of the opponents. In general, worst-case analysis can be understood as representing a strong form of pessimism.

Under Bayesian rationality a player optimizes against a fixed belief about the opponents' behavior. In contrast, under worst-case analysis a player's anticipation about the opponent's behavior depends on the agent's own choice. Stoye (2011), building on Milnor (1954), and using an Anscombe-Aumann decision-theoretic framework (see Section 0.4), provides an axiomatization of this *maxmin expected utility* criterion, as well as other classical criteria from statistical decision theory.

### *Zero-sum games*

Let  $G = \{\mathcal{P}, \{S_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$ ,  $\mathcal{P} = \{1, 2\}$  be a two-player game. We call  $G$  a *zero-sum game* if  $u_2(s) = -u_1(s)$  for all  $s \in S$ . In a zero-sum game, the two players' interests are diametrically opposed: whatever is good for one player is bad for the other, and to the same degree.

Since cardinal utilities are only unique up to a positive affine transformation of payoffs ( $v_i(s) = a + bu_i(s)$  for all  $s \in S$ ), games that are not zero-sum may be strategically equivalent to a zero-sum game.

In economics, zero-sum games are somewhat exceptional, since typical economic situations involve a combination of competitive and cooperative elements. Nevertheless, zero-sum games are important in economics because they describe the polar case of completely opposing interests.

When discussing zero-sum games, it is simplest to always speak in terms of player 1's payoffs. This way, player 1 is always the maximizer, and player 2 becomes the minimizer. (This choice is completely arbitrary, and has no bearing on any of the conclusions we

reach. In other words, we could just as easily speak in terms of player 2's payoffs. This would change who we call the maximizer and who the minimizer, but would not alter our conclusions about any given game—see the comment just before Example 1.54 below.)

It is convenient to present worst-case analysis in the context of zero-sum games. In this context, player 2 really does want to minimize player 1's payoffs, since by doing so she maximizes her own payoffs. But the ideas introduced next are useful beyond the context of zero-sum games. For instance, in repeated games, a player (or group of players) may use a minmax strategy to punish an opponent for deviations from equilibrium behavior—see Section 3.2.

### *Maxmin and minmax*

Let  $G$  be a zero-sum game. We first consider a worst-case analysis by player 1. Player 1 supposes that whatever mixed strategy he chooses, player 2 will anticipate it and choose a strategy that punishes player 1 as much as possible. It is *as if* player 1 chooses his strategy first, and player 2 observes this choice and responds with a corresponding punishment.

Let  $\alpha_2(\sigma_1)$  be a *punishment strategy* used by player 2 to punish player 1 when 1 plays  $\sigma_1$ :

$$\alpha_2(\sigma_1) \in \operatorname{argmin}_{\sigma_2 \in \Delta S_2} u_1(\sigma_1, \sigma_2).$$

(Stated in terms of player 2's payoffs,  $\alpha_2(\sigma_1)$  is a *best response* to  $\sigma_1$ , since

$$\alpha_2(\sigma_1) \in \operatorname{argmin}_{\sigma_2 \in \Delta S_2} u_1(\sigma_1, \sigma_2) = \operatorname{argmin}_{\sigma_2 \in \Delta S_2} (-u_2(\sigma_1, \sigma_2)) = \operatorname{argmax}_{\sigma_2 \in \Delta S_2} u_2(\sigma_1, \sigma_2).$$

But we will henceforth call  $\alpha_2(\sigma_1)$  a punishment strategy to keep our nomenclature in terms of player 1's payoffs.)

We say that  $\bar{\sigma}_1 \in \Delta S_1$  is a *maxmin strategy* for player 1 if

$$\bar{\sigma}_1 \in \operatorname{argmax}_{\sigma_1 \in \Delta S_1} \min_{\sigma_2 \in \Delta S_2} u_1(\sigma_1, \sigma_2) = \operatorname{argmax}_{\sigma_1 \in \Delta S_1} u_1(\sigma_1, \alpha_2(\sigma_1)).$$

The resulting payoff  $v_1^{\maxmin}$  is called player 1's *maxmin value*:

$$v_1^{\maxmin} = u_1(\bar{\sigma}_1, \alpha_2(\bar{\sigma}_1)).$$

In words:  $\bar{\sigma}_1$  maximizes player 1's payoffs given that player 2 correctly anticipates 1's choice and punishes him. By definition, strategy  $\bar{\sigma}_1$  is optimal for a player with maxmin expected utility preferences. As we will see, such a player may strictly prefer to play a

mixed strategy; this is not possible for a Bayesian rational player.

Now we consider a worst-case analysis by player 2. Suppose that player 2 anticipates that whatever she does, player 1 will punish her. Since we are considering a zero-sum game, player 1 punishes player 2 by playing a best response.

Let  $\beta_1(\sigma_2)$  be a best response of player 1 against  $\sigma_2$ :

$$\beta_1(\sigma_2) \in \operatorname{argmax}_{\sigma_1 \in \Delta S_1} u_1(\sigma_1, \sigma_2).$$

We say that  $\underline{\sigma}_2 \in \Delta S_2$  is a *minmax strategy* for player 2 if

$$\underline{\sigma}_2 \in \operatorname{argmin}_{\sigma_2 \in \Delta S_2} \max_{\sigma_1 \in \Delta S_1} u_1(\sigma_1, \sigma_2) = \operatorname{argmin}_{\sigma_2 \in \Delta S_2} u_1(\beta_1(\sigma_2), \sigma_2).$$

In words:  $\underline{\sigma}_2$  minimizes player 1's payoffs given that player 1 correctly anticipates 2's choice and best responds. The resulting payoff  $v_1^{\minmax}$  is player 1's *minmax value*:

$$v_1^{\minmax} = u_1(\beta_1(\underline{\sigma}_2), \underline{\sigma}_2).$$

To keep clear what  $v_1^{\maxmin}$  and  $v_1^{\minmax}$  mean, look at whose operation appears first. With  $v_1^{\maxmin}$ , it is player 1 (the maximizer) who chooses a mixed strategy first, and player 2 (the minimizer) who anticipates this choice and punishes. With  $v_1^{\minmax}$ , it is player 2 who chooses a mixed strategy first, and player 1 who anticipates this choice and best responds. (To check your understanding, imagine describing this analysis in terms of player 2's payoffs  $u_2 = -u_1$ . Then player 2 becomes the maximizer and player 1 the minimizer. You can confirm for yourself that player 2's maxmin value  $v_2^{\maxmin}$  is equal to  $-v_1^{\minmax}$ , and player 2's minmax value  $v_2^{\minmax}$  is equal to  $-v_1^{\maxmin}$ . Nothing substantial has changed; only the point of view is different.)

*Example 1.54. Maxmin and minmax with pure strategies.*

Consider the following zero-sum game:

		2		
		L	C	R
1	T	4, -4	2, -2	1, -1
	B	0, 0	1, -1	3, -3

Suppose that the players are only allowed to play pure strategies. In this case


$$\max_{s_1 \in S_1} \min_{s_2 \in S_2} u_1(s_1, s_2) = \max\{1, 0\} = 1;$$

$$\min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2) = \min\{4, 2, 3\} = 2. \quad \diamond$$

It is easy to check that whether players are restricted to pure strategies or are allowed to play mixed strategies,  $\max \min \leq \min \max$ : going last (and getting to react what your opponent did) is at least as good as going first. (The proof is straightforward: see “Proof of the Minmax Theorem” below.) The previous example shows that when players are restricted to pure strategies, we can have  $\max \min < \min \max$ . Can this happen if mixed strategies are allowed?

*Example 1.55. Computing maxmin and minmax strategies and payoffs.*

Consider the payoffs from the previous example. To compute  $\bar{\sigma}_1$  and  $v_1^{\max \min}$ , we first find the graph of  $u_1(\sigma_1, \alpha_2(\sigma_1))$  as a function of  $\sigma_1$ . This is the lower envelope function pictured in the diagram below. Its maximizer is  $\bar{\sigma}_1$ . Evidently,  $u_1(\bar{\sigma}_1, C) = u_1(\bar{\sigma}_1, R)$ , and so  $\bar{\sigma}_1 = \frac{2}{3}T + \frac{1}{3}B$ : if player 1 has maxmin expected utility preferences, he strictly prefers to randomize. 1’s maxmin payoff is  $v_1^{\max \min} = u_1(\bar{\sigma}_1, C) = \frac{5}{3}$ .




Notice that we can divide  $\Delta S_1$  into three punishment regions: writing  $\sigma_1 = \alpha T + (1 - \alpha)B$ , the regions are  $[0, \frac{1}{3}]$  (where the punishment is  $L$ ),  $[\frac{1}{3}, \frac{2}{3}]$  (where it is  $C$ ), and  $[\frac{2}{3}, 1]$  (where it is  $R$ ). Because the lower envelope is a minimum of linear functions, player 1’s maxmin strategy must occur at a vertex of one of the punishment regions: that is, at  $T$ ,  $\frac{2}{3}T + \frac{1}{3}B$ ,  $\frac{1}{3}T + \frac{2}{3}B$ , or  $B$ . An analogous statement is true in the case of player 2’s minmax strategy, as we will see next.

We can find player 2’s minmax strategy in a similar fashion. In this case, we are looking for the strategy of player 2 that *minimizes* an *upper* envelope function. This calculation

uses player 1's best response correspondence. The upper envelope of the payoff functions pictured below is  $u_1(\beta_1(\sigma_2), \sigma_2)$ . Because this upper envelope is the maximum of linear functions, it is minimized at a vertex of one of the best response regions shown at bottom. By computing  $u_1(\beta_1(\sigma_2), \sigma_2)$  at each vertex, we find that  $\underline{\sigma}_2 = \frac{2}{3}C + \frac{1}{3}R$ , where  $u_1(\beta_1(\frac{2}{3}C + \frac{1}{3}R), \frac{2}{3}C + \frac{1}{3}R) = \frac{5}{3} = v_1^{\min\max}$ .


Notice that  $v_1^{\max\min} = v_1^{\min\max}$ : the payoff that player 1 is able to guarantee himself is equal to the payoff that player 2 can hold him to. This is a consequence of the Minmax Theorem, which we state next.



We can also use a version of the first picture to compute  $\underline{\sigma}_2$  and  $v_1^{\min\max}$ . We do so by finding the convex combination of the graphs of  $u_1(\sigma_1, L)$ ,  $u_1(\sigma_1, C)$ , and  $u_1(\sigma_1, R)$  whose highest point is as low as possible. This is the horizontal line shown in the diagram below at right. (It is clear that no line can have a lower highest point, because no lower payoff for player 1 is feasible when  $\sigma_1 = \frac{2}{3}T + \frac{1}{3}B = \bar{\sigma}_1$ .) Since this horizontal line is the graph of  $\frac{2}{3}u_1(\sigma_1, C) + \frac{1}{3}u_1(\sigma_1, R) = u_1(\sigma_1, \frac{2}{3}C + \frac{1}{3}R)$  (check the endpoints), we conclude that  $\underline{\sigma}_2 = \frac{2}{3}C + \frac{1}{3}R$ . 1's minmax payoff is the constant value of  $u_1(\sigma_1, \frac{2}{3}C + \frac{1}{3}R)$ , which is  $v_1^{\min\max} = \frac{5}{3}$ .

In similar fashion, one can determine  $\bar{\sigma}_1$  and  $v_1^{\max\min}$  using the second picture by finding the convex combination of the two planes whose lowest point is as high as possible. This convex combination corresponds to the mixed strategy  $\bar{\sigma}_1 = \frac{2}{3}T + \frac{1}{3}B$ . (Sketch this plane in to see for yourself.)





### The Minmax Theorem

**Theorem 1.56** (The Minmax Theorem (von Neumann (1928))).

Let  $G$  be a two-player zero-sum game. Then

- (i)  $v_1^{\max\min} = v_1^{\min\max}$ .
- (ii)  $(\sigma_1^*, \sigma_2^*)$  is a Nash equilibrium of  $G$  if and only if  $\sigma_1^*$  is a maxmin strategy for player 1 and  $\sigma_2^*$  is a minmax strategy for player 2. (This implies that a Nash equilibrium of  $G$  exists.)
- (iii) Every Nash equilibrium of  $G$  yields payoff  $v_1^{\max\min} = v_1^{\min\max}$  for player 1.

The common value of  $v_1^{\max\min} = v_1^{\min\max}$  is known as the value of the game, and is denoted  $v_1^*$ .

The Minmax Theorem tells us that in a zero-sum game, player 1 can guarantee that he gets at least  $v_1^*$ , and player 2 can guarantee that player 1 gets no more than  $v_1^*$ ; moreover, in such a game, worst-case analysis and Bayesian-rational equilibrium analysis generate the same predictions of play.

For more on the structure of the equilibrium set in zero-sum games, see Shapley and Snow (1950); see González-Díaz et al. (2010) for a textbook treatment. For minmax theorems for games with infinite strategy sets, see Sion (1958).


*Example 1.57.* In the zero-sum game  $G$  with player 1's payoffs defined in Example 1.55, the unique Nash equilibrium is  $(\frac{2}{3}T + \frac{1}{3}B, \frac{2}{3}C + \frac{1}{3}R)$ , and the game's value is  $v(G) = \frac{5}{3}$ .  $\diamond$

*Example 1.58.* What are the Nash equilibria of this normal form game?

		2		
		X	Y	Z
1	A	10,0	-1,11	-1,11
	B	9,1	1,9	1,9
	C	2,8	8,2	1,9
	D	2,8	2,8	7,3
	E	4,6	5,5	6,4

This game is a constant-sum game, and so creates the same incentives as a zero-sum game (for instance, the game in which each player's payoffs are always 5 units lower). Therefore, by the Minmax Theorem, player 2's Nash equilibrium strategies are her minmax strategies. To find these, we draw player 1's best response correspondence (see below). The numbers in brackets are player 1's best response payoffs against the given strategies of player 2. Evidently, player 1's minmax payoff is  $\underline{v}_1 = \frac{51}{11} = 4\frac{7}{11}$ , which player 2 can enforce using her unique minmax strategy,  $\underline{\sigma}_2 = \frac{5}{11}X + \frac{5}{11}Y + \frac{1}{11}Z$ .

Now let  $\sigma_1^*$  be an equilibrium strategy of player 1. Since player 2 uses her minmax strategy  $\underline{\sigma}_2 = \frac{5}{11}X + \frac{5}{11}Y + \frac{1}{11}Z$  in any equilibrium, it must be a best response to  $\sigma_1^*$ ; in particular, X, Y, and Z must yield her the same payoff against  $\sigma_1^*$ . But  $\sigma_1^*$ , being a best response to  $\underline{\sigma}_2$ , can only put weight on pure strategies B, C, and E. The only mixture of these strategies that makes player 2 indifferent among X, Y, and Z is  $\sigma_1^* = \frac{13}{77}B + \frac{8}{77}C + \frac{56}{77}E$ . We therefore conclude that  $(\frac{13}{77}B + \frac{8}{77}C + \frac{56}{77}E, \frac{5}{11}X + \frac{5}{11}Y + \frac{1}{11}Z)$  is the unique Nash equilibrium of G.  $\diamond$



### Proofs of the Minmax Theorem

The proof of the minmax theorem is straightforward if one uses the Nash equilibrium existence theorem to show that  $v_1^{\max\min} \geq v_1^{\min\max}$ . We do this next. Then Example 1.59 gives a proof that  $v_1^{\max\min} \geq v_1^{\min\max}$  that uses only the separating hyperplane theorem.

To prove that  $v_1^{\maxmin} = v_1^{\minmax}$ , we first show that  $v_1^{\maxmin} \leq v_1^{\minmax}$ :

$$(\dagger) \quad v_1^{\maxmin} = u_1(\bar{\sigma}_1, \alpha_2(\bar{\sigma}_1)) \leq u_1(\bar{\sigma}_1, \underline{\sigma}_2) \leq u_1(\beta_1(\underline{\sigma}_2), \underline{\sigma}_2) = v_1^{\minmax}.$$

(The pure-strategy version of this argument is valid too.)

Proving that  $v_1^{\maxmin} \geq v_1^{\minmax}$  is also straightforward if one takes as given that Nash equilibria exist. Letting  $(\sigma_1^*, \sigma_2^*)$  be a Nash equilibrium of  $G$ , we have that

$$(\ddagger) \quad v_1^{\maxmin} = u_1(\bar{\sigma}_1, \alpha_2(\bar{\sigma}_1)) \geq u_1(\sigma_1^*, \alpha_2(\sigma_1^*)) = u_1(\sigma_1^*, \sigma_2^*) = u_1(\beta_1(\sigma_2^*), \sigma_2^*) \geq u_1(\beta_1(\underline{\sigma}_2), \underline{\sigma}_2) = v_1^{\minmax}.$$

The argument above shows that all of the inequalities in  $(\dagger)$  and  $(\ddagger)$  are equalities. We now use this fact to prove statements (ii) and (iii) of the theorem. First, the two new equalities in  $(\dagger)$  tell us that any pair of maxmin strategies forms a Nash equilibrium. (The second new equality shows that 1's maxmin strategy  $\bar{\sigma}_1$  is a best response to 2's minmax strategy  $\underline{\sigma}_2$ ; the first new equality shows the reverse.) Second, the two new equalities in  $(\ddagger)$  tell us that a Nash equilibrium strategy for a given player is a maxmin strategy for that player. (The first new equality shows that the Nash equilibrium strategy  $\sigma_1^*$  is maxmin strategy for 1; the second new equality shows that  $\sigma_2^*$  is a minmax strategy for 2.) And third, the equalities in  $(\ddagger)$  also show that all Nash equilibria of  $G$  yield payoff  $v_1^{\maxmin} = v_1^{\minmax}$  for player 1. This completes the proof of the theorem. ■

*Example 1.59.* Consider the two player normal form game  $G$  below, in which player 1 is the row player and player 2 is the column player. (Only player 1's payoffs are shown.)

		2			
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
1	<i>T</i>	5, ·	3, ·	1, ·	0, ·
	<i>B</i>	-1, ·	0, ·	4, ·	7, ·

Define the following sets:

$$J = \{v \in \mathbb{R}^2 : v = (u_1(T, s_2), u_1(B, s_2)) \text{ for some } s_2 \in S_2\};$$

$K$  = the convex hull of  $J$ ;

$$L(c) = \{w \in \mathbb{R}^2 : w_1 \leq c \text{ and } w_2 \leq c\}.$$

- (i) Explain in game theoretic terms what it means for a vector to be an element of the set  $K$ .
- (ii) Let  $c^*$  be the smallest number such that the point  $(c^*, c^*)$  is contained in  $K$ . What is the value of  $c^*$ ? Relate this number to player 1's minmax payoff in  $G$ , explaining the reason for the relationship you describe.

- (iii) Specify the normal vector  $p^* \in \mathbb{R}^2$  and the intercept  $d^*$  of the hyperplane  $H = \{v \in \mathbb{R}^2 : p^* \cdot v = d^*\}$  that separates the set  $L(c^*)$  from the set  $K$ , choosing the vector  $p^*$  to have components that are nonnegative and sum to one.
- (iv) Interpret the fact that  $p^* \cdot v \geq d^*$  for all  $v \in K$  in game theoretic terms. What conclusions can we draw about player 1's maxmin payoff in  $G$ ?
- (v) Let  $G_n$  be a two player normal form game in which player 1 has  $n \geq 2$  strategies. Sketch a proof of the fact that player 1's minmax payoff in  $G_n$  and his maxmin payoff in  $G_n$  are equal. (When  $G_n$  is zero-sum, this fact is the Minmax Theorem.)

Solution:

- (i)  $v \in K$  if and only if  $v = (u_1(T, \sigma_2), u_1(B, \sigma_2))$  for some  $\sigma_2 \in \Delta S_2$ .
- (ii)  $c^* = 2$ . This is player 1's minmax value: by playing  $\frac{1}{2}b + \frac{1}{2}c$ , the strategy that generates  $(2, 2)$ , player 2 ensures that player 1 cannot obtain a payoff higher than 2. If you draw a picture of  $K$ , you will see that player 2 cannot restrict 1 to a lower payoff.
- (iii)  $p^* = (\frac{2}{3}, \frac{1}{3})$  and  $d^* = 2$ .
- (iv) Let  $\sigma_1^* = p^*$ . Then  $(p^* \cdot v \geq 2$  for all  $v \in K)$  is equivalent to  $(u_1^*(\sigma_1^*, \sigma_2) \geq 2$  for all  $\sigma_2 \in \Delta S_2)$ . Thus, player 1's maxmin payoff is at least 2; by part (ii), it must be exactly 2.
- (v) Define  $n$ -strategy analogues of the sets  $J$ ,  $K$ , and  $L(c)$ . Let  $c^*$  be the largest value of  $c$  such that  $\text{int}(L(c))$  and  $\text{int}(K)$  do not intersect. (Notice that if, for example, player 1 has a dominated strategy,  $L(c^*) \cap K$  may not include  $(c^*, \dots, c^*)$ ; this is why we need the set  $L(c^*)$ .) If player 2 chooses a  $\sigma_2$  that generates a point in  $L(c^*) \cap K$ , then player 1 cannot obtain a payoff higher than  $c^*$ . Hence, player 1's minmax value is less than or equal to  $c^*$ . (In fact, the way we chose  $c^*$  tells us that it is exactly equal to  $c^*$ .)

Let  $p^*$  and  $d^*$  define a hyperplane that separates  $L(c^*)$  and  $K$ ; we know that such a hyperplane exists by the separating hyperplane theorem. Given the form of  $L(c^*)$ , we can choose  $p^*$  to lie in  $\Delta S_1$ ; therefore, since the hyperplane passes through the point  $(c^*, \dots, c^*)$ , the  $d^*$  corresponding to any  $p^*$  chosen from  $\Delta S_1$  (in particular, to any  $p^*$  whose components sum to one) is in fact  $c^*$ . Thus, as in (iv) above, player 1's maxmin value is at least  $c^*$ . Since player 1's maxmin value cannot exceed his minmax value, we conclude that both of these values equal  $c^*$ .  $\diamond$

## 2. Extensive Form Games

### 2.1 Basic Concepts

#### 2.1.1 Defining extensive form games

Extensive form games describe strategic interactions in which moves may occur sequentially. The two main classes of extensive form games are *games of perfect information* and


games of imperfect information. The key distinction here is that in the former class of games, a player's choices are always immediately observed by his opponents.

Simultaneous moves are modeled by incorporating unobserved moves (see Example 2.2), and so lead to imperfect information games (but see Section 2.3.4).


Extensive form games may also include chance events, modeled as *moves by Nature*. For our categorization, it is most convenient to understand "game of perfect information" to refer only to games without moves by Nature.

The definition of extensive form games here follows Selten (1975). Osborne and Rubinstein (1994) define extensive form games without explicitly introducing game trees. Instead, what we call nodes are identified with the sequences of actions that lead to them. This approach, which is equivalent to Selten's, requires somewhat less notation, but at the cost of being somewhat more abstract.

*Example 2.1. Sequential Battle of the Sexes.* In this game, player 2 observes player 1's choice before making her own choice.




*Example 2.2. (Simultaneous) Battle of the Sexes.* Here player 2 does not observe player 1's choice before she moves herself. We represent the fact that player 2 cannot tell which of her decision nodes has been reached by enclosing them in an oval. The decision nodes are said to be in the same *information set*.




*Example 2.3. A simple card game.* Players 1 and 2 each bet \$1. Player 1 is given a card which is high or low; each is equally likely. Player 1 sees the card, player 2 doesn't. Player 1 can raise the bet to \$2 or fold. If player 1 raises, player 2 can call or fold. If player 2 calls, then player 1 wins if and only if his card is high.

The random assignment of a card is represented as move by Nature, marked in the game tree as "player 0". Since player 2 cannot tell whether player 1 has raised with a high card or a low card, her two decision nodes are in a single information set.



**Games of perfect information**

Let  $X$  be a set of *nodes*, and  $E \subseteq X \times X$  a set of (*directed*) *edges*. The edge leading from node  $x$  to node  $y$  is written as  $e = (x, y)$ . We call the pair  $(X, E)$  a *tree* with *root*  $r$  if  $r$  has no incoming edge, and if for each  $y \neq r$  there is a unique *path* (i.e., sequence of edges  $\{(x_0, x_1), (x_1, x_2), \dots, (x_{n-1}, x_n)\}$ ) from  $r = x_0$  to  $y = x_n$  using only edges in  $E$ . Nodes with outgoing edges are called *decision nodes*; those without are called *terminal nodes*. The sets of these are denoted  $D$  and  $Z$ , respectively.



All edges are labeled with actions:

$A$  set of actions

$\alpha: E \rightarrow A$  assigns each edge an action, assigning different actions to distinct edges leaving the same node ( $e = (x, y), \hat{e} = (x, \hat{y}) \neq e \Rightarrow \alpha(e) \neq \alpha(\hat{e})$ )


We let  $A_x = \{\alpha(e) : e = (x, y) \text{ for some } y \in X\}$  denote the set actions available at decision node  $x$ .

Finally, we introduce a set of players, assign each decision node to a player, and specify each player's utility at each terminal node.

$\mathcal{P} = \{1, \dots, n\}$  set of players

$D_i \subseteq D$  set of player  $i$ 's decision nodes: each  $x \in D$  is in exactly one  $D_i$   
(in other words,  $D_1, \dots, D_n$  form a *partition* of  $D$ )

$u_i: Z \rightarrow \mathbb{R}$  player  $i$ 's Bernoulli utility function



This completes the definition of an extensive form game of perfect information. Extensive form games, perfect information or not, are denoted by  $\Gamma$ .

*Example 2.4. Sequential Battle of the Sexes: notation.* In Example 2.1, we have:

Assignments of decision nodes:  $D_1 = \{x\}, D_2 = \{y, z\}$ .

Action sets:  $A_x = \{F, B\}, A_y = \{f, b\}, A_z = \{\hat{f}, \hat{b}\}$

Example of utility:  $(u_1(\zeta_1), u_2(\zeta_1)) = (2, 1)$ .  $\diamond$

### Moves by Nature

Chance events that occur during the course of play are represented by decision nodes assigned to an artificial "player 0", often called Nature, whose choice probabilities are specified as part of the definition of the game.

To be more precise, the set of decision nodes  $D$  is now partitioned into  $D_0, D_1, \dots, D_n$ , with  $D_0$  being the set of nodes assigned to Nature. Let  $x$  be such a node, so that  $A_x$  is the set of actions available at  $x$ . Then the definition of the game includes a probability vector  $p_x \in \Delta A_x$ , which for each action  $a \in A_x$  specifies the probability  $p_x(a)$  that Nature chooses  $a$ .

Although we sometimes refer to Nature as “player 0”, Nature is really just a device for representing chance events, and should not be regarded as a player. That is, the set of players is still  $\mathcal{P} = \{1, \dots, n\}$ .

### *Games of imperfect information*

In games of imperfect information, certain action choices of one or more players or of Nature may not be observed before other players make their own action choices.

To represent unobserved moves, we form a partition  $\mathbb{I}_i$  of each player  $i$ 's set of decision nodes  $D_i$  into *information sets*  $I \subseteq D_i$ . When play reaches player  $i$ 's information set  $I$ ,  $i$  knows that some node in  $I$  has been reached, but he cannot tell which.

For an information set  $I$  to make sense, its owner must have the same set of actions  $A_I$  available at each node in  $I$ . Were this not true, the owner could figure out where he is in  $I$  by seeing what actions he has available. (Formally, we require that if  $x, y \in I$ , then  $A_x = A_y$ , and we define  $A_I$  to be this common action set.)

When an information set is a singleton (like  $\{x\}$ ), we sometimes abuse notation and refer to the information set itself as  $x$ .

Observe that  $\Gamma$  is a game of perfect information if every information set is a singleton and there are no moves by Nature.

*Example 2.5. A simple card game: notation.* In Example 2.3, we have:

Assignments of decision nodes:  $D_0 = \{v\}$ ,  $D_1 = \{w, x\}$ ,  $D_2 = \{y, z\}$ .

Action sets:  $A_v = \{H, L\}$ ,  $A_w = \{R, F\}$ ,  $A_x = \{r, f\}$ ,  $A_y = A_z = \{C, F\}$ .

The probability assignments at Nature's node are  $p_v(H) = p_v(L) = \frac{1}{2}$ .

Only player 2 has a nontrivial information set:  $\mathbb{I}_2 = \{I\}$ , where  $I = \{y, z\}$ . Notice that  $A_y = A_z$  as required; thus  $A_I = A_y = A_z = \{C, F\}$ .  $\diamond$

### *Perfect recall*

One nearly always restricts attention to games satisfying *perfect recall*: each player remembers everything he once knew, including his own past moves.



To express this requirement formally, write  $e < y$  when edge  $e$  precedes node  $y$  in the game



tree (i.e., when the path from the root node to  $y$  includes  $e$ ). Let  $E_x = \{e = (x, y) \text{ for some } y \in X\}$  be the set of outgoing edges from  $x$ . Then perfect recall is defined as follows:

If  $x \in I \in \mathbb{I}_i$ ,  $e \in E_x$ ,  $y, \hat{y} \in I' \in \mathbb{I}_i$ , and  $e < y$ ,  
then there exist  $\hat{x} \in I$  and  $\hat{e} \in E_{\hat{x}}$  such that  $\alpha(e) = \alpha(\hat{e})$  and  $\hat{e} < \hat{y}$ .

*Example 2.6.* Here are two games in which perfect recall is violated. The game at right is known as the *absent-minded driver's problem*; see Piccione and Rubinstein (1997).



### 2.1.2 Pure strategies in extensive form games

A player's strategy specifies *exactly one action* at each of his information sets  $I$ . The player cannot choose different actions at different nodes in an information set because he is unable to distinguish these nodes during play.

Formally, a *pure strategy*  $s_i \in S_i \equiv \prod_{I \in \mathbb{I}_i} A_I$  specifies an action  $s_i(I)$  for player  $i$  at each of his information sets  $I \in \mathbb{I}_i$ .


*Example 2.7. Sequential Battle of the Sexes: strategies.* In Example 2.1,  $S_1 = A_x = \{F, B\}$  and  $S_2 = A_y \times A_z = \{f\hat{f}, f\hat{b}, b\hat{f}, b\hat{b}\}$ . Note that even if player 1 chooses  $F$ , we still must specify what player 2 would have done at  $z$ , so that 1 can evaluate his choices at  $x$ .  $\diamond$

A pure strategy for an extensive form game must provide a complete description of how a player intends to play the game, regardless of what the other players do. In other words, one role of a strategy is to specify a "plan of action" for playing the game.

However, in games where a player may be called upon to move more than once during the course of play (that is, in games that do not have the *single-move property*), a pure strategy

contains information that a plan of action does not: it specifies how the player would act at information sets that are *unreachable* given his strategy, meaning that the actions specified by his strategy earlier in the game ensure that the information set is not reached.

*Example 2.8. Mini Centipede.*



Assignments of decision nodes:  $D_1 = \{x, z\}$ ,  $D_2 = \{y\}$ .

Action sets:  $A_x = \{A, B\}$ ,  $A_y = \{C, D\}$ ,  $A_z = \{E, F\}$ .

Strategy sets:  $S_1 = A_x \times A_z = \{AE, AF, BE, BF\}$  and  $S_2 = A_y = \{C, D\}$ .

Note that when player 1 chooses  $B$  at node  $x$ , node  $z$  is not reached. Nevertheless, 1's strategy still specifies what he would do at  $z$ .  $\diamond$

One interpretation of a pure strategy is that a player specifies what he would do at unreachable information sets in order to be prepared in the event that he fails to implement his choices at early nodes correctly. The possibility of such failures is the basis for many analyses of extensive form games via their normal form versions—see Section 2.6.

An alternate interpretation is that the specification of a player's strategy at unreachable information sets describes opponents' beliefs about what that player would do there—see Rubinstein (1991).

Either way, this notion of strategy for extensive form games is essential for capturing the *principle of sequential rationality* (Section 2.2), which is the basis for most analyses of extensive form games. The awkwardness in the notion of strategy for games without the single-move property will resurface as difficulties with the notion of sequential rationality in these games—see Section 2.3.2.

We will see that the distinction between strategies and plans of action disappears in the usual normal form representation of an extensive form game (Section 2.1.4), and that this distinction is irrelevant to the notion of Nash equilibrium.


### 2.1.3 Randomized strategies in extensive form games

A *mixed strategy*  $\sigma_i \in \Delta(\prod_{I \in I_i} A_I)$  randomizes over pure strategies.

Example 2.9. For the game tree below,

Pure strategies:  $S_2 = \{L, R\} \times \{l, r\} = \{Ll, Lr, Rl, Rr\}$ ;

Mixed strategies:  $\sigma_2 = (\sigma_2(Ll), \sigma_2(Lr), \sigma_2(Rl), \sigma_2(Rr))$  with  $\sum_{s_2 \in S_2} \sigma_2(s_2) = 1$ .




Another way to specify random behavior is to suppose that each time a player reaches an information set  $I$ , he randomizes over the actions in  $A_I$ . (These randomizations are assumed to be independent of each other and of all other randomizations in the game.)

This way of specifying behavior is called a *behavior strategy*:  $\beta_i \in \prod_{I \in \mathcal{I}_i} \Delta A_I$ .

Example 2.10. For the game from Example 2.9,

$$\beta_2 = (\beta_2^x, \beta_2^y), \text{ where } \beta_2^x(L) + \beta_2^x(R) = 1 \text{ and } \beta_2^y(l) + \beta_2^y(r) = 1.$$



Note that mixed strategies are joint distributions, while behavior strategies are collections of marginal distributions from which draws are assumed to be independent. (Their relationship is thus the same as that between correlated strategies and mixed strategy profiles in normal form games.)

It follows immediately that every distribution over pure strategies generated by a behavior strategy can also be generated by a mixed strategy. However, the converse statement is false in general, since behavior strategies cannot generate correlation in choices at different information sets.

*Example 2.11.* In the game from Example 2.9, consider the behavior strategy  $\beta_2 = ((\beta_2^x(L), \beta_2^x(R)), (\beta_2^y(l), \beta_2^y(r))) = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{3}, \frac{2}{3}))$ . Since the randomizations at the two nodes are independent,  $\beta_2$  generates a mixed strategy that is a product distribution over pure strategies, as shown inside the table below:

		$y$		
		$\frac{1}{3}$	$\frac{2}{3}$	
		$l$	$r$	
$x$	$\frac{1}{2}$	$L$	$\frac{1}{6}$	$\frac{1}{3}$
	$\frac{1}{2}$	$R$	$\frac{1}{6}$	$\frac{1}{3}$

Or, in notation:

$$\begin{aligned} \sigma_2 &= (\sigma_2(LL), \sigma_2(Lr), \sigma_2(Rl), \sigma_2(Rr)) \\ &= (\frac{1}{2} \cdot \frac{1}{3}, \frac{1}{2} \cdot \frac{2}{3}, \frac{1}{2} \cdot \frac{1}{3}, \frac{1}{2} \cdot \frac{2}{3}) \\ &= (\frac{1}{6}, \frac{1}{3}, \frac{1}{6}, \frac{1}{3}). \quad \diamond \end{aligned}$$

*Example 2.12.* The mixed strategy  $\hat{\sigma}_2 = (\frac{1}{3}, \frac{1}{6}, 0, \frac{1}{2})$  entails correlation between the randomizations at player 2's two decision nodes. In behavior strategies, such correlation is forbidden. But in this case the correlation is strategically irrelevant, since during a given play of the game, only one of player 2's decision nodes will be reached. In fact,  $\hat{\sigma}_2$  is also "strategically equivalent" to  $\beta_2$ . We make this idea precise next.  $\diamond$

Correlation in a player's choices at different information sets is strategically irrelevant in any game with perfect recall. Thus in such games, mixed strategies and behavior strategies provide exactly the same strategic possibilities.

To formalize this point, we say that player  $i$ 's behavior strategy  $b_i$  and mixed strategy  $\sigma_i$  are *outcome equivalent* if for any mixed strategy profile  $\sigma_{-i}$  of player  $i$ 's opponents,  $(b_i, \sigma_{-i})$  and  $(\sigma_i, \sigma_{-i})$  generate the same distribution over terminal nodes.

It is immediate from the foregoing that every behavior strategy is outcome equivalent to a mixed strategy. For the converse, we have

**Theorem 2.13** (Kuhn's (1953) Theorem). *Suppose that  $\Gamma$  has perfect recall. Then every mixed strategy  $\sigma_i$  is outcome equivalent to some behavior strategy  $\beta_i$ .*

For a proof, see González-Díaz et al. (2010).

Suppose one is given a mixed strategy  $\sigma_i$ . How can one find an equivalent behavior strategy  $\beta_i$ ? Specifically, how do we define  $\beta_i^I(a)$ , the probability placed on action  $a$  at player  $i$ 's information set  $I$ ?

Intuitively,  $\beta_i^I(a)$  should be the probability that  $a$  is chosen at  $I$ , conditional on  $i$  and his opponents acting in such a way that  $I$  is reached.

Formally,  $\beta_i^I(a)$  can be defined as follows:

- (i) Let  $\sigma_i(I)$  be the probability that  $\sigma_i$  places on pure strategies that do not preclude  $I$ 's being reached.
- (ii.a) If  $\sigma_i(I) > 0$ , let  $\sigma_i(I, a)$  be the probability that  $\sigma_i$  places on pure strategies that do not preclude  $I$ 's being reached and that specify action  $a$  at  $I$ . Then let  $\beta_i^I(a) = \sigma_i(I, a)/\sigma_i(I)$ .
- (ii.b) If  $\sigma_i(I) = 0$ , specify  $\beta_i^I \in \Delta A_I$  arbitrarily.

In games with the single-move property, a player cannot prevent any of his own information sets from being reached, so  $\sigma_i(I) = 1$ . Thus case (ii.a) always applies, so the procedure above specifies a unique  $\beta_i$ , and in fact this  $\beta_i$  is the only behavior strategy that is outcome equivalent to  $\sigma_i$ .

*Example 2.14.* Consider the game from Example 2.9, which has the single-move property. Consider mixed strategy  $\sigma_2 = (\sigma_2(Ll), \sigma_2(Lr), \sigma_2(Rl), \sigma_2(Rr)) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{12}, \frac{1}{4})$ . Writing this inside a joint distribution table, then the marginal distributions are the equivalent behavior strategy:


		$y$		
		$\frac{5}{12}$	$\frac{7}{12}$	
		$l$	$r$	
$x$	$\frac{2}{3}$	$L$	$\frac{1}{3}$	$\frac{1}{3}$
	$\frac{1}{3}$	$R$	$\frac{1}{12}$	$\frac{1}{4}$

In notation:

$$\beta_2^x(L) = \sigma_2(Ll) + \sigma_2(Lr) = \frac{1}{3} + \frac{1}{3} = \frac{2}{3};$$

$$\beta_2^y(l) = \sigma_2(Ll) + \sigma_2(Rl) = \frac{1}{3} + \frac{1}{12} = \frac{5}{12}. \quad \diamond$$

*Example 2.15.* Again let  $\sigma_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{12}, \frac{1}{4})$  as above, but consider the game below:



Then  $\beta_2^x(L) = \sigma_2(Ll) + \sigma_2(Lr) = \frac{2}{3}$  as before. This game does not have the single move property, and player 2's second node can only be reached if  $L$  is played. In a joint distribution table, we can represent this fact by crossing out the  $R$  row of the table (remembering that we do this only when computing choice probabilities at  $y$ , not at  $x$ ), and then scaling up to get the conditional probabilities.

		$y$			
		$\frac{1}{2}$	$\frac{1}{2}$		
		$l$	$r$		
$x$	$\frac{2}{3}$	$L$	$\frac{1}{3}$	$\frac{1}{3}$	$\diamond$
	$\frac{1}{3}$	<del><math>R</math></del>	<del><math>\frac{1}{2}</math></del>	<del><math>\frac{1}{4}</math></del>	

In notation, the computation of  $\beta_2^y(l)$  is

$$\beta_2^y(l) = \frac{\sigma_2(Ll)}{\sigma_2(Ll) + \sigma_2(Lr)} = \frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{3}} = \frac{1}{2}.$$

*Example 2.16.* In the game from the Example 2.15, if  $\sigma_2 = (0, 0, \frac{1}{3}, \frac{2}{3})$ , then  $\beta_2^x(L) = 0$ , but  $\beta_2^y(l)$  is unrestricted, since if 2 plays  $\sigma_2$  her second node is not reached.  $\diamond$


When studying extensive form games directly, it is generally easier to work with behavior strategies than with mixed strategies. Except when the two are being compared, the notation  $\sigma_i$  is used to denote both.

#### 2.1.4 Reduced normal form games

Mixed strategies appear when we consider a normal form game  $G(\Gamma)$  derived from an extensive form game  $\Gamma$ .

To motivate normal form representations, imagine that before play begins, the players specify complete contingent plans for playing  $\Gamma$ . Each reports his plan to a moderator, who then implements the plan. Since in this scenario strategies are chosen beforehand, there is no loss in describing the interaction as a simultaneous move game.


Example 2.17.



$G(\Gamma_1)$

		2				
		Ll	Lr	Rl	Rr	
1	U	3,3	3,3	2,1	2,1	◇
	D	1,2	4,4	1,2	4,4	


Example 2.18.



$G(\Gamma_2)$

		2		
		L	R	
1	U	3,3	2,1	◇
	D	1,2	4,4	

Example 2.19.



$G(\Gamma_3)$

		2				
		Ll	Lr	Rl	Rr	
1	U	4,1	2,2	3,1	3,1	◇
	D	3,0	3,0	3,1	3,1	

We are not done: the (purely) reduced normal form consolidates equivalent pure strategies:

$G(\Gamma_3)$		2			
		Ll	Lr	R	
1	U	4, 1	2, 2	3, 1	
	D	3, 0	3, 0	3, 1	$\diamond$


In the previous example, the two pure strategies that were consolidated correspond to the same “plan of action” (see Section 2.1.2). This is true in generic extensive form games.

Many extensive form games can have the same reduced normal form. For example, switching the order of moves in Example 2.18 does not change the reduced normal form. Thompson (1952) and Elmes and Reny (1994) show that two extensive form games with perfect recall have the same reduced normal form if it is possible to convert one game to the other by applying three basic transformations: interchange of simultaneous moves, addition of superfluous moves, and coalescing of information sets.

## 2.2 The Principle of Sequential Rationality

Equilibrium concepts for normal form games can be applied to a finite extensive form game  $\Gamma$  by way of the reduced normal form. For instance, strategy profile  $\sigma$  is a *Nash equilibrium* of  $\Gamma$  if it is equivalent to a Nash equilibrium of the reduced normal form of  $\Gamma$ . However, the solution concepts we have seen so far do not address questions of credibility of commitment that arise in dynamic contexts.

*Example 2.20. Entry deterrence.* The players are firms, an entrant (1) and an incumbent (2). The entrant moves first, deciding to stay Out or to Enter the market. If the entrant stays Out, he gets a payoff of 0, while the Incumbent gets the monopoly profit of 3. If the entrant Enters, the incumbent must choose between Fighting (so that both players obtain  $-1$ ) or Accommodating (so that both players obtain the duopoly profit of 1).



Nash equilibria:  $(E, A)$  and  $(O, \sigma_2(F) \geq \frac{1}{2})$ .



If player 1 chooses  $E$ , player 2's only reasonable response is  $A$ . Thus  $F$  is an empty threat. But Nash equilibrium does not rule  $F$  out.

Suppose we require player 2 to make credible commitments: if she specifies an action, she must be willing to carry out that action if her decision node is reached. This requirement would force her to play  $A$ , which in turn would lead player 1 to play  $E$ .  $\diamond$

Once we specify a strategy profile in an extensive form game, certain information sets cannot be reached. They are said to be *off-path*.

While all choices at unreached information sets are optimal, these choices can determine which choices are optimal at information sets that are reached.

To make sure that off-path behavior is reasonable—in particular, that only credible threats are made—we introduce the *principle of sequential rationality*: Predictions of play in extensive form games should require optimal behavior starting from *every* information set, not just those on the equilibrium path.

We will study two main solution concepts that formalize this principle: subgame perfect equilibrium for perfect information games (Section 2.3), and sequential equilibrium for imperfect information games (Section 2.4).

In generic games of perfect information, we can capture the principle of sequential rationality without the use of equilibrium assumptions, but with strong assumptions about players' belief in opponents' rationality—see Section 2.3.2. Beyond this setting, equilibrium knowledge assumptions will be necessary.

## 2.3 Games of Perfect Information and Backward Induction

### 2.3.1 Subgame perfect equilibrium, sequential rationality, one-shot deviations, and backward induction

Recall that extensive form game  $\Gamma$  has *perfect information* if every information set of  $\Gamma$  is a singleton, and if  $\Gamma$  contains no moves by Nature.

We now formalize the principle of sequential rationality in the the context of perfect information games by defining three solution concepts. While each definition will seem less demanding than the previous ones, Theorem 2.21 will show that all three are equivalent.

The first two definitions make use of the notion of a subgame. The *subgame* of  $\Gamma$  starting from  $x$ , denoted  $\Gamma_x$ , is the game defined by the portion of  $\Gamma$  that starts from decision node  $x \in D$  and includes all subsequent nodes. Intuitively, a subgame is a portion of the game that one can consider analyzing without reference to the rest of the game. (In games of

imperfect information, not every decision node is the beginning of a subgame: see Section 2.3.4.)

If  $\sigma$  is a strategy profile in  $\Gamma$ , then  $\sigma|_x$  denotes the strategy profile that  $\sigma$  induces in subgame  $\Gamma_x$ .

*Definition (i).* Strategy profile  $\sigma$  is a *subgame perfect equilibrium* of  $\Gamma$  (Selten (1965)) if in each subgame  $\Gamma_x$  of  $\Gamma$ ,  $\sigma|_x$  is a Nash equilibrium.

*Definition (ii).* Strategy  $\sigma_i$  is *sequentially rational* given  $\sigma_{-i}$  if for each decision node  $x$  of player  $i$ ,  $(\sigma|_x)_i$  is a best response to  $(\sigma|_x)_{-i}$  in  $\Gamma_x$ . If this is true for every player  $i$ , we say that strategy profile  $\sigma$  itself is *sequentially rational*.

*Definition (iii).* Strategy  $\sigma_i$  *admits no profitable one-shot deviations* given  $\sigma_{-i}$  if for each decision node  $x$  of player  $i$ , player  $i$  cannot improve his payoff in  $\Gamma_x$  by changing his action at node  $x$  but otherwise following strategy  $(\sigma|_x)_i$ . If this is true for every player  $i$ , we say that strategy profile  $\sigma$  itself *admits no profitable one-shot deviations*.

Each of these definitions imposes requirements on strategies within each subgame  $\Gamma_x$ : subgame perfection (i) requires that *all players' strategies* be optimal; sequential rationality (ii) requires that *the strategy of the owner of node  $x$*  be optimal; no profitable one-shot deviations (iii) requires that *the action chosen at node  $x$*  be optimal. But the fact that the conditions are required to hold at *all subgames* makes the three definitions equivalent, as we discuss next.

### ***Finite games***

A *finite perfect-information game* has a finite number of decision nodes and a finite number of actions at each decision node.

**Theorem 2.21.** *Let  $\Gamma$  be a finite perfect information game. Then the following are equivalent:*

- (i) *Strategy profile  $\sigma$  is a subgame perfect equilibrium.*
- (ii) *Strategy profile  $\sigma$  is sequentially rational.*
- (iii) *Strategy profile  $\sigma$  admits no profitable one-shot deviations.*

The implications (i)  $\Rightarrow$  (ii)  $\Rightarrow$  (iii) follow easily from their definitions, and the implication (ii)  $\Rightarrow$  (i) is mainly notation-juggling. The implication (iii)  $\Rightarrow$  (ii) is a direct consequence of a basic result on single-agent sequential choice called the (*finite-horizon*) *one-shot deviation principle* (Theorem 2.28). Thus the principle of sequential rationality has both game-theoretic content (credibility of commitments) and decision-theoretic content (optimal

sequential decision-making). The proof of Theorem 2.21 and a detailed presentation of the one-shot deviation principle are provided at the end of this section.

We can find all profiles satisfying (iii) using the following *backward induction procedure*: Find a decision node  $x$  that is only followed by terminal nodes. Specify an action at that node that leads to the terminal node  $z$  following  $x$  that yields the highest payoff for the owner of node  $x$ . Then replace decision node  $x$  with terminal node  $z$  and repeat the procedure.

If indifferences occur, the procedure branches, with each branch specifying a different optimal choice at the point of indifference. Different strategy profiles that survive may have different outcomes: while the player making the decision is indifferent between his actions, other players generally are not (see Example 2.35).

The backward induction procedure systematically constructs strategy profiles from which there are no profitable one-shot deviations: it first ensures this at nodes followed only by terminal nodes, and then at the nodes before these, and so on.

**Observation 2.22.** *In a finite perfect information game, a strategy profile admits no profitable one-shot deviations if and only if it survives the backward induction procedure.*


Remarks:

- (i) Since the backward induction procedure always generates at least one strategy profile, existence of subgame perfect equilibrium is guaranteed. In fact, since it is always possible to specify a pure action at every point of indifference, a *pure strategy* subgame perfect equilibrium always exists.
- (ii) If the backward induction procedure never leads to an indifference, it generates a unique subgame perfect equilibrium, sometimes called the *backward induction solution*. This is the case in generic finite perfect information games (specifically, finite perfect information games in which no player is indifferent between any pair of terminal nodes).

In this case “subgame perfect equilibrium” is somewhat of a misnomer. This term suggests that equilibrium knowledge assumptions are needed to justify the prediction. In Section 2.3.2 we explain why this is not true, but also why the assumptions that are needed are still rather strong in many cases.

*Example 2.23. Entry Deterrence: solution.* In the Entry Deterrence game (Example 2.20), the backward induction procedure selects  $A$  for player 2, and hence  $E$  for player 1. Thus the backward induction solution is  $(E, A)$ .  $\diamond$

*Example 2.24. Multiple entrants.* There are two entrants and an incumbent. The entrants decide sequentially whether to stay out of the market or enter the market. Entrants who stay out get 0. If both entrants stay out, the incumbent gets 5. If there is entry, the incumbent can fight or accommodate. If the incumbent accommodates, per firm profits are 2 for duopolists and  $-1$  for triopolists. On top of this, fighting costs the incumbent 1 and the entrants who enter 3.



The backward induction solution of this game is  $(E, (e, o'), (a, a', a''))$ , generating outcome  $(2, 0, 2)$ . This is an instance of *first mover advantage*.

Note that Nash equilibrium does not restrict possible predictions very much in this game: of the 64 pure strategy profiles, 20 are Nash equilibria of the reduced normal form;  $(0, 0, 5)$ ,  $(0, 2, 2)$ , and  $(2, 0, 2)$  are Nash equilibrium outcomes. Thus, requiring credibility of commitments refines the set of predictions substantially.

1:O					1:E						
2					2						
oo'    oe'    eo'    ee'					oo'    oe'    eo'    ee'						
3	$ff'f''$	<b>0 0 5</b>	<b>0 0 5</b>	0 -1 1	0 -1 1	3	$ff'f''$	-1 0 1	-4 -4 -2	-1 0 1	-4 -4 -2
	$ff'a''$	<b>0 0 5</b>	<b>0 0 5</b>	0 -1 1	0 -1 1		$ff'a''$	-1 0 1	-1 -1 -1	-1 0 1	-1 -1 -1
	$fa'f''$	0 0 5	<b>0 0 5</b>	0 -1 1	0 -1 1		$fa'f''$	<b>2 0 2</b>	-4 -4 -2	<b>2 0 2</b>	-4 -4 -2
	$fa'a''$	0 0 5	<b>0 0 5</b>	0 -1 1	0 -1 1		$fa'a''$	<b>2 0 2</b>	-1 -1 -1	<b>2 0 2</b>	-1 -1 -1
	$af'f''$	0 0 5	0 0 5	<b>0 2 2</b>	<b>0 2 2</b>		$af'f''$	-1 0 1	-4 -4 -2	-1 0 1	-4 -4 -2
	$af'a''$	0 0 5	0 0 5	<b>0 2 2</b>	<b>0 2 2</b>		$af'a''$	-1 0 1	-1 -1 -1	-1 0 1	-1 -1 -1
	$aa'f''$	0 0 5	0 0 5	0 2 2	<b>0 2 2</b>		$aa'f''$	<b>2 0 2</b>	-4 -4 -2	<b>2 0 2</b>	-4 -4 -2
	$aa'a''$	0 0 5	0 0 5	0 2 2	<b>0 2 2</b>		$aa'a''$	<b>2 0 2</b>	-1 -1 -1	<b>2 0 2</b>	-1 -1 -1

◇

### Infinite games

Theorem 2.21 continues to hold in games with a finite number of decision nodes but infinite numbers of actions at each node. But subgame perfect equilibria need not exist because maximizing actions need not exist. (For instance, consider a one-player game in which the player chooses  $x \in [0, 1)$  and receives payoff  $x$ .)

Observation 2.22 stated that in a finite perfect information game, if an indifference occurs during the backward induction procedure, then each choice of action at the point of indifference leads to a distinct subgame perfect equilibrium. This conclusion may fail in games with infinite action sets, because some ways of breaking indifferences may be inconsistent with equilibrium play—see Example 2.37.

Theorem 2.21 also holds in games with infinitely many decision nodes provided that a mild condition on payoffs is satisfied. We say that a game's payoffs are *continuous at infinity* if for every  $\varepsilon > 0$  there exists a  $K$  such that choices made after the  $K$ th node of any play path cannot alter any player's payoffs by more than  $\varepsilon$ . This property holds in games in which payoffs are discounted over time and (undiscounted) payoffs in each period are bounded—see Example 2.39 and Section 3.

**Theorem 2.25.** *Let  $\Gamma$  be an arbitrary perfect information game whose payoffs are continuous at infinity. Then the following are equivalent:*

- (i) Strategy profile  $\sigma$  is a subgame perfect equilibrium.
- (ii) Strategy profile  $\sigma$  is sequentially rational.
- (iii) Strategy profile  $\sigma$  admits no profitable one-shot deviations.

Again, implications (i)  $\Rightarrow$  (ii)  $\Rightarrow$  (iii) follow easily from their definitions, and implication (ii)

$\Rightarrow$  (i) is straightforward to check. Implication (iii)  $\Rightarrow$  (ii) follows from the *infinite horizon one-shot deviation principle*, which is an easy extension of its finite-horizon counterpart; see the end of this section for a proof. In these results, the assumption that payoffs are continuous at infinity cannot be dispensed with; see Example 2.31.

In games with infinitely many decision nodes, subgame perfect equilibria cannot be found using backward induction, since there are no last decision nodes from which to launch the backward induction procedure. Instead, the one-shot deviation principle must be used directly—again, see Example 2.39 and Section 3.

### *Proofs of Theorems 2.21 and 2.25*

To begin the proof of Theorem 2.21, we explain why the implications (i)  $\Rightarrow$  (ii) and (ii)  $\Rightarrow$  (iii) are immediate from the definitions. For the first implication, consider what definitions (i) and (ii) require in each subgame  $\Gamma_x$ : in Definition (i), every player's strategy must be optimal given the strategies of the others; in definition (ii), only the strategy of the owner of node  $x$  is required to be optimal.

For the second implication, consider the optimization problems the players are required to solve in definitions (ii) and (iii). In definition (ii), when we evaluate the strategy  $(\sigma|_x)_i$  of the owner  $i$  of node  $x$  in subgame  $\Gamma_x$ , we must compare the performance of  $(\sigma|_x)_i$  to that of each other strategy for the subgame, including ones that specify different actions than  $(\sigma|_x)_i$  at multiple decision nodes. Under definition (iii), the only changes in strategy that are ever considered consist of a change in action at a single decision node.

Thus, the content of Theorem 2.21 is in the reverse implications, (ii)  $\Rightarrow$  (i) and (iii)  $\Rightarrow$  (ii).

Roughly speaking, (ii) implies (i) because in a sequentially rational strategy profile  $\sigma$ , the optimality in subgame  $\Gamma_x$  of a strategy  $(\sigma|_x)_j$  of a player  $j$  who does not own node  $x$  can be deduced from the optimality of  $j$ 's strategies in subgames  $\Gamma_y$  of  $\Gamma_x$  whose initial nodes are owned by  $j$ . The proof to follow makes this argument explicit.


*Proof of Theorem 2.21, (ii)  $\Rightarrow$  (i).* We give the proof for pure strategy profiles  $s$ ; the proof for mixed strategy profiles is similar.

Suppose that  $s$  is a sequentially rational strategy profile, and let  $x$  be a decision node of player  $i$ . We need to show that  $s|_x$  is a Nash equilibrium of  $\Gamma_x$ . First, note that by definition,  $(s|_x)_i$  is a best response to  $(s|_x)_{-i}$ . Second, if a decision node belonging to player  $j \neq i$  is on the play path under  $s|_x$ , then let  $y$  be the first such node. Since  $y$  is the first node of  $j$ 's that is reached,  $j$ 's choices outside of  $\Gamma_y$  do not affect his payoff in  $\Gamma_x$  when his opponents play  $s|_x$ , and sequential rationality implies that  $(s|_y)_j$  is a best response to  $(s|_y)_{-j}$ . Together, these facts imply that  $(s|_x)_j$  is a best response to  $(s|_x)_{-j}$ . Third, if no node of player  $k$  is reached on the play path under  $s|_x$ , then any behavior for  $k$  is a best response to  $(s|_x)_{-k}$ . We therefore conclude that  $s|_x$  is a Nash equilibrium of  $\Gamma_x$ . ■

As noted earlier, the implication (iii)  $\Rightarrow$  (ii) is a consequence of a fundamental result about single-agent decision problems called the *one-shot deviation principle*, which we discuss

next. To understand this decision-theoretic aspect of backward induction, it is easiest to focus on perfect information games with one player, which we call *sequential decision problems*.

*Example 2.26. A sequential decision problem.*



Let  $\sigma$  and  $\hat{\sigma}$  be pure strategies of the lone player in sequential decision problem  $\Gamma$ .

We say that  $\hat{\sigma}$  is a *profitable deviation from  $\sigma$  in subgame  $\Gamma_x$*  if  $\hat{\sigma}|_x$  generates a higher payoff than  $\sigma|_x$  in subgame  $\Gamma_x$ . In this definition, it is essential that payoffs be evaluated from the vantage of node  $x$ , not from the vantage of the initial node of  $\Gamma$  (see Example 2.29).

By definition,  $\sigma$  is sequentially rational in  $\Gamma$  if it does not admit a profitable deviation in any subgame. Put differently,  $\sigma$  is sequentially rational if for any decision node  $x$ ,  $\sigma|_x$  yields the lone player the highest payoff available in  $\Gamma_x$ .

*Example 2.27. Sequential decision problem revisited.* In the sequential decision problem from Example 2.26, the sequentially rational strategy is  $(B, C, F, G, I, K, N)$ .  $\diamond$

We call strategy  $\hat{\sigma}$  a *one-shot deviation* from strategy  $\sigma$  if it only differs from  $\sigma$  at a single decision node, say  $\hat{x}$ . This one-shot deviation is *profitable* if it generates a higher payoff than  $\sigma$  in subgame  $\Gamma_{\hat{x}}$ .

**Theorem 2.28** (The one-shot deviation principle (finite horizon)).

Let  $\Gamma$  be a finite sequential decision problem. Then the following are equivalent:

- (i) Strategy  $\sigma$  does not admit a profitable deviation in any subgame. That is,  $\sigma$  is sequentially rational.
- (ii) Strategy  $\sigma$  does not admit a profitable one-shot deviation at any decision node.


Thus to construct a strategy that does not admit profitable deviations of any kind, even ones requiring changes in action at multiple nodes at once, it is enough to apply the backward induction procedure, which never considers such deviations explicitly.

In the context of finite sequential decision problems, the one-shot deviation principle is simple, but it is not trivial—it requires a proof. (One way to see that the theorem has some content is to note that it does not extend to infinite sequential decision problems without additional assumptions are imposed—see Theorem 2.30 and Example 2.31 below.)

*Proof of Theorem 2.28.* It is immediate that (i) implies (ii). To prove the converse, suppose that  $\sigma$  does not admit a profitable one-shot deviation, but that it does admit a profitable deviation that requires changes in action over  $T$  periods of play. By hypothesis, the last stage of this deviation is not profitable when it is undertaken. Therefore, only following the first  $T - 1$  periods of the deviation must yield at least as good a payoff from the vantage of the initial stage of the deviation. By the same logic, it is at least as good to only follow the first  $T - 2$  periods of the deviation, and hence the first  $T - 3$  periods of the deviation . . . and hence the first period of the deviation. But since a one-shot deviation cannot be profitable, we have reached a contradiction. ■

In applying the one-shot deviation principle, it is essential that the profitability of a deviation be evaluated from the vantage of the node where the deviation occurs.

*Example 2.29. Illustration of the one-shot deviation principle.*



If deviations are always evaluated from the point of view of the initial node  $x$ , then there is no way to improve upon strategy  $(O, L)$  by only changing the action played at a single node. But  $(O, L)$  clearly is not sequentially rational. Does this contradict the one-shot deviation principle?

No. When we consider changing the action at node  $y$  from  $L$  to  $R$ , we should view the effect of this deviation from the vantage of node  $y$ , where it is indeed profitable. The choice of  $R$  at  $y$  in turn creates a profitable deviation at  $x$  from  $O$  to  $I$ . ◇

We are now in a position to complete the proof of Theorem 2.21.

*Proof of Theorem 2.21, (iii)  $\Rightarrow$  (ii).* Suppose that strategy profile  $\sigma$  admits no profitable one-shot deviation in the finite perfect information game  $\Gamma$ . We wish to show that  $\sigma$  is sequentially rational in  $\Gamma$ .

Suppose that we view the strategy profile  $\sigma_{-i}$  of player  $i$ 's opponents as exogenous. Then player  $i$  is left with a collection sequential decision problems. In fact, player  $i$ 's strategy  $\sigma_i$  was obtained by applying the backward induction procedure in these decision problems,



so Theorem 2.28 implies that  $\sigma_i$  is sequentially rational in these decision problems. As this is true for all players,  $\sigma$  is sequentially rational in  $\Gamma$ . ■

To prove Theorem 2.25, it is enough to establish

**Theorem 2.30** (The one-shot deviation principle (infinite horizon)).

*Let  $\Gamma$  be an infinite horizon sequential decision problem whose payoffs are continuous at infinity. Then the following are equivalent:*

- (i) *Strategy  $\sigma$  does not admit a profitable deviation in any subgame. That is,  $\sigma$  is sequentially rational.*
- (ii) *Strategy  $\sigma$  does not admit a profitable one-shot deviation at any decision node.*

*Proof:* Clearly (i) implies (ii). To prove the inverse, suppose that (i) is false. Then there is a decision node  $x$  whose owner  $i$  has an infinite-period deviation from  $\sigma$  that starts at node  $x$  and increases his payoff in subgame  $\Gamma_x$  by some  $\delta > 0$ . Since payoffs are continuous at infinity, this player  $i$  must have a finite-period deviation in  $\Gamma_x$  that increases his payoff by  $\frac{\delta}{2}$  in  $\Gamma_x$ . But then Theorem 2.28 implies player  $i$  has a profitable one-shot deviation in  $\Gamma_x$ , and thus (ii) must also be false. ■

*Example 2.31.* Without continuity of payoffs at infinity, the infinite horizon one-shot deviation principle does not hold. For example, suppose that an agent must choose an infinite sequence of  $L$ s and  $R$ s; his payoff is 1 if he always chooses  $R$  and is 0 otherwise. Consider the strategy “always choose  $L$ ”. While there is no profitable finite-period deviation from this strategy, there is obviously a profitable infinite-period deviation. ◇

### 2.3.2 Epistemic foundations for backward induction

#### *Games with and without the single move property*

To begin our discussion of foundations, we introduce the game-theoretic setting in which backward induction is easiest to justify.

A *play path* of  $\Gamma$  is a sequence of nodes reached in  $\Gamma$  under some pure strategy profile.

Perfect information game  $\Gamma$  has the *single move property* if there is no play path on which a single player has multiple decision nodes. Thus, the Entry Deterrence game (Example 2.20) and the Multiple Entrants game (Example 2.24) have the single move property.


Under the single-move property, there is no way that a player’s behavior “early” in the game can provide information about his behavior “later” in the game, since a player who has already moved will not be called upon to move again. For this reason, in generic games with the single-move property the backward induction solution can be justified using mild assumptions. When there are just two players, common certainty of rationality

is sufficient. (Why common *certainty* of rationality? “Knowledge” refers to beliefs that are both certain and correct. In extensive form contexts, using “certainty” rather than “knowledge” allows for the possibility that a player discovers that his beliefs are incorrect during the course of play.)

In games with more than two players, one also needs an assumption of *epistemic independence*: obtaining information about one opponent (by observing his move) does not lead a player to update his beliefs about *other* opponents (see Stalnaker (1998) and Battigalli and Siniscalchi (1999)).

In games where the single-move property is not satisfied, so that some player may have to move more than once during the course of play, common certainty of rationality and epistemic independence are not enough to justify backward induction. In addition, one needs to impose an assumption that regardless of what has happened in the past, there continues to be common certainty that behavior will be rational in the future. See Halpern (2001) and Asheim (2002) for formalizations of this idea. This assumption of “undying common belief in future rational play” is sufficient to justify backward induction so long as indifferences never arise during the backward induction procedure. This assumption is strong, and sometimes leads to counterintuitive conclusions—see Example 2.34. But the assumption is of a different nature than the equilibrium knowledge assumption needed to justify Nash equilibrium play, as *coordination* of different players’ beliefs is not assumed.

*Example 2.32. Mini Centipede.*




The backward induction solution of this game is  $((S, S'), s)$ . For player 1 to prefer to play  $S$ , he must expect player 2 to play  $s$  if her node is reached. Is it possible for a rational player 2 to do this?

To answer this question, we have to consider what player 2 might think if her decision node is reached: specifically, what she would conclude about player 1’s rationality, and about what player 1 will do at the final decision node. One possibility is that player 2 thinks that a rational player 1 would play  $S$  at his initial node, and that a player 1 who plays  $C$  there is irrational, and in particular would play  $C'$  at the final node if given the

opportunity. If this is what player 2 would think if her node were reached, then she would play  $c$  there.

If player 1 is rational and anticipates such a reaction by player 2, he will play  $C$ , resulting in the play of  $((C, S'), c)$ . Indeed, this prediction is consistent with the assumption of common certainty of rationality at the start of play—see Ben-Porath (1997). Therefore, stronger assumptions of the sort described before this example are needed to ensure the backward induction solution.  $\diamond$


*Example 2.33. Three-player Mini Centipede.*



The game above is a three-player version of the one from Example 2.32, with the third node being assigned to player 3, and with player 3's payoffs being identical to player 1's. The game's backward induction solution is  $(S, s, S)$ . In this game, if player 2's node is reached, she may become doubtful of player 1's rationality, but there is no obvious reason for her to doubt player 3's rationality. Thus unlike in Example 2.32, the backward induction prediction can be justified by common certainty of rationality and epistemic independence.  $\diamond$

*Example 2.34. Centipede (Rosenthal (1981)).*

Two players alternate moves, starting with player 1. Each player starts with \$0. When moving, a player can Stop the game or Continue. If a player Continues, he gives up \$1, while his opponent gains \$3. The game ends when a player Stops or when both players have \$100.



In the backward induction solution, everyone always stops, yielding payoffs of  $(0, 0)$ .  $\diamond$


Remarks on the Centipede Game:

- (i) This example raises a general critique of backward induction logic: Suppose you are player 2, and you get to go. Since 1 deviated once, why not expect him to deviate again?
- (ii) The outcome  $(0, 0)$  is not only the backward induction solution; it is also the unique Nash equilibrium outcome.
- (iii) In experiments, most people do better (McKelvey and Palfrey (1992)).
- (iv) There are augmented versions of the game in which some Continuing occurs in equilibrium: see Kreps et al. (1982) (cf. Example 2.53 below) and Jehiel (2005).
- (v) Stable cooperative behavior arises in populations of optimizing agents if one makes realistic assumptions about the information agents possess (Sandholm et al. (2016)).

***Backward induction with indifferences***

When indifferences arise during the backward induction procedure, the procedure branches, with each branch leading to a distinct subgame perfect equilibrium. To justify subgame perfect equilibrium predictions in such cases, one must resort to equilibrium knowledge assumptions to coordinate beliefs about what indifferent players will do.

*Example 2.35. Discipline by an indifferent parent.* Player 1 is a child and player 2 the parent. The child chooses to Behave or to Misbehave. If the child Misbehaves, the parent chooses to Punish or Not to punish; she is indifferent between these options, but the child is not.



If node  $y$  is reached, player 2 is indifferent, so any choice is allowed in this subgame. We therefore must split the computation of subgame perfect equilibrium into cases. Every specification of 2's behavior leads to a subgame perfect equilibrium:

$$\begin{aligned}
 &(B, \beta P + (1 - \beta)N) \text{ with } \beta > \frac{1}{2} \\
 &(\alpha B + (1 - \alpha)M, \frac{1}{2}P + \frac{1}{2}N) \text{ with } \alpha \in [0, 1], \\
 &(M, \beta P + (1 - \beta)N) \text{ with } \beta < \frac{1}{2}.
 \end{aligned}$$

One can argue that the parent should be able to credibly commit to punishing—see Tranæs (1998).  $\diamond$

### 2.3.3 Applications

*Example 2.36. Chess and checkers.*

In chess, moves and payoffs are well-defined. The game has perfect information and is zero-sum. Termination rules ensure that it is finite (at least in some variations—see Ewerhart (2002)).

Since chess has perfect information, a pure subgame perfect equilibrium exists. Since it is zero-sum, this pure equilibrium is in maxmin strategies. And since these maxmin strategies are pure, they must guarantee each player his value for the game not only in expectation, but deterministically: he gets exactly his value if his opponent plays a pure best response, and gets more than this if she does not.

It follows that one of these three statements must be true: (i) White can guarantee a win, (ii) Black can guarantee a win, (iii) both players can guarantee a draw.

In fact, it can be shown that after two rounds of removing weakly dominated strategies, all of the strategy profiles that remain yield the equilibrium outcome (Ewerhart (2000)).

But the number of positions in chess is on the order of  $10^{40}$ , so even writing down the game tree for chess is impossible! In big enough games, backward induction is not computationally feasible.

On the other hand, checkers (aka draughts), which has roughly  $5 \times 10^{20}$  positions, has been solved: perfect play leads to a draw! (Schaeffer et al. (2007))  $\diamond$

*Example 2.37. Finite-horizon alternating-offer bargaining (Ståhl (1972)).*

Two players bargain over how to split \$1. The bargaining lasts for at most  $T + 1$  periods, starting with period 0.

In period 0, player 1 offers  $x \in [0, 1]$  dollars to player 2, leaving  $1 - x$  dollars for himself. If player 2 accepts, the dollar is split and the game ends. If player 2 rejects, then the dollar is not split in period 0, and period 1 begins, with the roles reversed.

Play continues, with player 1 making offers in even periods and player 2 making offers in odd periods, either until an offer is accepted, or until period  $T$  ends with a rejected offer, in which case both players get nothing.

Dollars received in period  $t$  are discounted by  $\delta^t$ , where  $\delta \in (0, 1)$ .

This game has perfect information, but also infinite action sets and many tied outcomes.

**Proposition 2.38.** *In the finite-horizon alternating-offer bargaining game,*

- (i) *There is a unique subgame perfect equilibrium. In it, the initial offer is accepted.*
- (ii) *As  $T \rightarrow \infty$ , equilibrium payoffs converge to  $(\frac{1}{1+\delta}, \frac{\delta}{1+\delta})$ .*

Part (ii) tells us that in a long game, there is a first mover advantage.

*Proof:* Suppose  $T$  is even (i.e., the number of periods is odd).

In period  $T$ , player 2 should accept any positive offer. If she receives an offer of 0, she is indifferent, so we must consider the possibility that she accepts such an offer, rejects such an offer, or randomizes after such an offer.

But it is convenient to consider player 1's period  $T$  offer first. It cannot be a best response for player 1 to offer a positive amount to player 2. If he were going to offer her  $x$ , he improve his payoffs by offering her  $\frac{x}{2}$  instead, since that offer will also be accepted. Thus in any subgame perfect equilibrium, player 1 must offer 0 to player 2.

Now suppose that player 2 rejects an offer of 0 with probability  $y > 0$ . Then player 1's expected dollar return from this offer is  $1 - y$ , and so he could do better by offering  $\frac{y}{2}$ , which ensures a dollar return of  $1 - \frac{y}{2}$ .

Thus in any subgame perfect equilibrium, player 2 must accept an offer of 0. If player 2 does so, then offering 0 is optimal for player 1. This is the unique subgame perfect equilibrium of the period  $T$  subgame.

We repeatedly (but implicitly) use this argument in the remainder of the analysis. (See the remark below for further discussion.)

Using backward induction, and the notation (dollars or payoffs for player 1, dollars or payoffs for player 2):

In period  $T$  (even): 2 accepts any offer.

$\Rightarrow$  1 offers  $(1, 0) \xrightarrow{\times \delta^T} (\delta^T, 0)$  payoffs if this offer is accepted.

In period  $T - 1$  (odd): 1 accepts any offer which yields him  $\delta^T$ .

$\Rightarrow$  2 offers  $(\delta, 1 - \delta) \xrightarrow{\times \delta^{T-1}} (\delta^T, \delta^{T-1} - \delta^T)$  payoffs if this offer is accepted.

In period  $T - 2$  (even): 2 accepts any offer which yields her  $\delta^{T-1} - \delta^T$ .

$\Rightarrow$  1 offers  $(1 - \delta + \delta^2, \delta - \delta^2) \xrightarrow{\times \delta^{T-2}} (\delta^{T-2} - \delta^{T-1} + \delta^T, \delta^{T-1} - \delta^T)$  payoffs if accepted.

... In period 0: 1's offer = payoffs =

$$(u_1^*, u_2^*) = (1 - \delta + \delta^2 - \dots + \delta^T, 1 - u_1^*) = \left( \frac{1 - \delta^{T+1}}{1 - \delta}, \frac{\delta - \delta^{T+1}}{1 - \delta} \right) \rightarrow \left( \frac{1}{1 - \delta}, \frac{\delta}{1 - \delta} \right) \text{ as } T \rightarrow \infty.$$

The same limit result holds for odd  $\hat{T}$ :  $(u_1^*, u_2^*) = \left( \frac{1 - \delta^{\hat{T}+1}}{1 - \delta}, \frac{\delta - \delta^{\hat{T}+1}}{1 - \delta} \right) \rightarrow \left( \frac{1}{1 - \delta}, \frac{\delta}{1 - \delta} \right)$  as  $T \rightarrow \infty$ .

(Why? We saw that if  $T$  is even, 1 offers last, and his period 0 offer is  $(\frac{1 - \delta^{T+1}}{1 - \delta}, \frac{\delta - \delta^{T+1}}{1 - \delta})$  (\*). If  $\hat{T} = T + 1$  is odd, then 2 offers last; since when she makes her period 1 offer there are

$T$  periods to go, this offer is obtained by reversing (\*):  $(\frac{\delta - \delta^{T+1}}{1+\delta}, \frac{1+\delta^{T+1}}{1+\delta}) = (\frac{\delta - \delta^T}{1+\delta}, \frac{1+\delta^T}{1+\delta})$ . Thus, in period 0, 1's offer to 2 is  $u_2^* = \delta \cdot \frac{1+\delta^T}{1+\delta} = \frac{\delta + \delta^{T+1}}{1+\delta}$ , and he keeps  $u_1^* = 1 - \frac{\delta + \delta^{T+1}}{1+\delta} = \frac{1 - \delta^{T+1}}{1+\delta}$  for himself.)  $\diamond$

Remark: The argument above shows that the unique subgame perfect equilibrium has each player making offers that cause the opponent to be indifferent, and has the opponent accepting such offers. The fact that the opponent always accepts when indifferent may seem strange, but with a continuum of possible offers, it is necessary for equilibrium to exist. It is comforting to know that this equilibrium can be viewed as the limit of equilibria of games with finer and finer discrete strategy sets, equilibria in which the opponent rejects offers that leave her indifferent.

*Example 2.39. Infinite-horizon alternating offer bargaining (Rubinstein (1982)).*

The game is the same as above, except that if no offer accepted the game can go on indefinitely. As there is no last period, backward induction cannot be applied. Nevertheless:

**Proposition 2.40.** *In the infinite-horizon alternating-offer bargaining game, there is a unique subgame perfect equilibrium. The equilibrium specifies that in any period, the proposer proposes keeping  $\frac{1}{1+\delta}$  and giving  $\frac{\delta}{1+\delta}$  to his opponent, and the receiver accepts exactly those offers that give her at least  $\frac{\delta}{1+\delta}$ . Thus the subgame perfect equilibrium payoffs are  $(u_1^*, u_2^*) = (\frac{1}{1+\delta}, \frac{\delta}{1+\delta})$ .*

Propositions 2.38 and 2.40 show that *equilibrium payoffs* in alternating-offer bargaining games are “continuous at infinity”. (Continuity of equilibrium payoffs at infinity does not hold in all games—see Section 3.1.)

Verifying that the specified strategy profile is a subgame perfect equilibrium is a simple application of the one-shot deviation principle, and showing that it is uniquely so requires surprisingly little calculation. This is because of the recursive structure of the game: each period  $t$  subgame “looks the same” as the full game, and each period  $t$  subgame “looks the same” as a period 1 subgame (cf. Section 3.1). It is also helpful that subgames in even and odd periods “look the same” apart from the identity of the proposer.

*Proof of Proposition 2.40:* Consider any period  $t \in \{0, 1, 2, \dots\}$  in which player  $i$  is the proposer and player  $j$  the receiver. We argue in terms of *continuation payoffs*: that is, we proceed as though the game started in period  $t$ , so that we need not bother discounting everything by  $\delta^t$ . The one-shot deviation principle tells us that to verify that a strategy profile is a subgame perfect equilibrium, we need only consider deviations occurring at one node at a time. To show that the strategy profile specified in the proposition is a subgame perfect equilibrium, note first that player  $j$  could obtain a payoff of  $\delta \cdot \frac{1}{1+\delta}$  by

declining the offer and becoming the proposer next period. Thus it is not profitable for player  $j$  to deviate from accepting an offer of  $\frac{\delta}{1+\delta}$  today. Knowing this, it is optimal for player  $i$  to offer  $\frac{\delta}{1+\delta}$  to player  $j$ , since this is the smallest amount player  $j$  will accept, and since offering player  $j$  less than this will result in player  $i$  obtaining a payoff of  $\delta \cdot \frac{\delta}{1+\delta} < \frac{1}{1+\delta}$ . The proof of uniqueness follows Shaked and Sutton (1984). Let  $\bar{u}_i$  and  $\underline{u}_i$  be the supremum and infimum of the set of player  $i$ 's subgame perfect equilibrium payoffs.

In any period  $t$  of any subgame perfect equilibrium, the receiver  $j$  will accept any offer worth at least  $\delta \bar{u}_j$  to her. Thus it cannot be optimal for the proposer  $i$  to offer  $j$  more than this, or, equivalently, to keep less than  $1 - \delta \bar{u}_j$  for himself. This provides a lower bound on player  $i$ 's possible subgame perfect equilibrium payoffs:

$$(15) \quad \underline{u}_i \geq 1 - \delta \bar{u}_j.$$

Likewise, in any period  $t$  of any subgame perfect equilibrium, the receiver  $j$  will reject any offer worth less than  $\delta \underline{u}_j$  to her. Therefore, if the proposer  $i$  makes an offer that is accepted in the equilibrium, the offer gives  $j$  at least  $\delta \underline{u}_j$ , and so gives  $i$  at most  $1 - \delta \underline{u}_j$ . Alternatively, if  $i$  makes an offer that is rejected, then  $j$  will obtain a payoff of at least  $\delta \cdot \underline{u}_j$ , and so  $i$  will receive a payoff of at most  $\delta \cdot (1 - \underline{u}_j) < 1 - \delta \underline{u}_j$ . These facts yield an upper bound on player  $i$ 's equilibrium payoff:

$$(16) \quad \bar{u}_i \leq 1 - \delta \underline{u}_j.$$

To avoid confusion, we now refer to the players as  $k$  and  $\ell$ . Using inequalities (15) and (16), we obtain the following bounds on  $\underline{u}_k$  and  $\bar{u}_k$ :

$$\begin{aligned} \underline{u}_k &\geq 1 - \delta \bar{u}_\ell \geq 1 - \delta(1 - \delta \underline{u}_k) & \bar{u}_k &\leq 1 - \delta \underline{u}_\ell \leq 1 - \delta(1 - \delta \bar{u}_k) \\ \Leftrightarrow (1 - \delta^2) \underline{u}_k &\geq 1 - \delta & \Leftrightarrow (1 - \delta^2) \bar{u}_k &\leq 1 - \delta \\ \Leftrightarrow \underline{u}_k &\geq \frac{1}{1+\delta}; & \Leftrightarrow \bar{u}_k &\leq \frac{1}{1+\delta}. \end{aligned}$$

Thus  $\underline{u}_k \geq \frac{1}{1+\delta} \geq \bar{u}_k \geq \underline{u}_k$ , so all of the inequalities bind. We conclude that equilibrium must be of the form specified in the proposition. Since player 1 is the proposer in period 0, it is he who obtains the equilibrium payoff of  $\frac{1}{1+\delta}$ .  $\diamond$

### 2.3.4 Subgame perfect equilibrium in more general classes of games


To discuss subgame perfect equilibrium more generally, we must specify what is meant by a subgame in general extensive form games. A *subgame*  $\Gamma'$  of an general extensive form




game  $\Gamma$  is a subset of  $\Gamma$  which

- (i) Contains a decision node, all succeeding edges and nodes, and no others.
- (ii) Does not tear information sets: if  $x \in \Gamma'$  and  $x, y \in I$ , then  $y \in \Gamma'$ .

*Example 2.41.* Below,  $a, b$ , and  $i$  are the initial nodes of subgames; the other nodes are not.



*Example 2.42.* The game tree below has three subgames: the whole game, and the portions of the game starting after  $(H, t)$  and  $(T, h)$ .



As in the perfect information case, strategy profile  $\sigma$  is a *subgame perfect equilibrium* of  $\Gamma$  if it induces a Nash equilibrium in every subgame of  $\Gamma$ .

Subgame perfect equilibrium is completely adequate for analyzing games of perfect information. It is also completely adequate in *games of stagewise perfect information* (also called *multistage games with observable actions*), which generalize perfect information games by allowing multiple players to move at once at any point in play, with all actions being observed immediately afterward. Thus nontrivial information sets are only used to capture

simultaneous choices. Example 2.42 has stagewise perfect information, but Example 2.41 does not.

The key point is that in stagewise perfect information games, once one fixes the behavior of a player's opponents, the player himself faces a collection of sequential decision problems, implying that we can determine whether the player's behavior is sequentially rational by looking for profitable one-shot deviations; in particular, Theorem 2.21 extends to the present class of games; see also Example 2.48 below.

Once simultaneous moves are possible, performing "backward induction" requires us to find all equilibria at each stage of the game, allowing the analysis to branch whenever multiple equilibria are found. Equilibrium knowledge assumptions are needed to coordinate beliefs about how play will proceed.

*(Infinitely) repeated games* (Section 3) are the simplest infinite-horizon version of games of stagewise perfect information, so subgame perfect equilibrium is also the appropriate solution concept for these games. Since players discount future payoffs, payoffs are continuous at infinity in the sense defined in Section 2.3.1. Because of this, the one-shot deviation principle applies, and the conclusions of Theorem 2.25 hold for repeated games. As previously noted, backward induction cannot be used to solve games with infinite numbers of decision nodes; instead, the one-shot deviation principle must be applied directly.

We will see next that in general games of imperfect information, subgame perfect equilibrium is no longer sufficient to capture sequential rationality.


## 2.4 Games of Imperfect Information and Sequential Equilibrium

### 2.4.1 Subgames and subgame perfection in games of imperfect information

Recall that a subgame  $\Gamma'$  of an (imperfect information) extensive form game  $\Gamma$  is a subset of  $\Gamma$  which contains a decision node, all succeeding edges and nodes, and that does not tear information sets; and that a strategy profile  $\sigma$  is a subgame perfect equilibrium of  $\Gamma$  if it induces a Nash equilibrium in every subgame of  $\Gamma$ .

In this context, subgame perfection is generally inadequate to capture the principle of sequential rationality.

*Example 2.43. Entry deterrence II.*



		2	
		F	A
1	O	0, 2	0, 2
	M	-1, -1	2, 1
	B	-3, -1	4, 0

The entrant can choose stay out ( $O$ ), enter meekly ( $M$ ), or enter boldly ( $B$ ). The incumbent cannot observe which sort of entry has occurred, but is better off accommodating either way.

This game has two components of Nash equilibria,  $(B, A)$  and  $(O, \sigma_2(F) \geq \frac{2}{3})$ .

Since the whole game is the only subgame, all of these Nash equilibria are subgame perfect equilibria.  $\diamond$

To heed the principle of sequential rationality in this example, we need to ensure that player 2 behaves rationally if her information set is reached.

To accomplish this, we will require players to form beliefs about where they are in an information set, regardless of whether the information set is reached in equilibrium. We then will require optimal behavior at each information set given these beliefs.

By introducing appropriate restrictions on allowable beliefs, we ultimately will define the notion of sequential equilibrium, the fundamental equilibrium concept for general extensive form games.

Remark on information sets and evaluation of deviations:

In Example 2.43, in any Nash equilibrium in which player 2's information set is reached with positive probability, equilibrium knowledge ensures that player 2 has correct beliefs about where she is in the information set. (This sort of equilibrium knowledge is precisely the sort occurring in normal form games.)

Nevertheless, the information set plays a role when we evaluate the consequences of a deviation by player 1.

For instance, consider evaluating the strategy profile  $(B, A)$ . When we consider a deviation by player 1 from  $B$  to  $M$ , player 2's strategy remains fixed at  $A$ .

If instead each of player 2's nodes were in its own information set, player 2's equilibrium strategy could specify different behavior at each. In this case, if we considered a deviation by player 1 from  $B$  to  $M$ , player 2's behavior *on the path of play* would change if her choices

at her two nodes differed.

## 2.4.2 Beliefs and sequential rationality

### *Beliefs*

Fix an extensive form game  $\Gamma$ .

Player  $i$ 's beliefs are a map  $\mu_i : D_i \rightarrow [0, 1]$  satisfying  $\sum_{x \in I} \mu_i(x) = 1$  for all  $I \in \mathbb{I}_i$ . (The term "beliefs" refers to something different here than it did in Section 1.)

If  $x \in I$ ,  $\mu_i(x)$  represents the probability that player  $i$  assigns to being at node  $x$  given that his information set  $I$  has been reached.


Note: the player subscript on  $\mu_i$  is often omitted when no confusion will result.

Let  $\mu = (\mu_1, \dots, \mu_n)$  denote the profile of all players' beliefs.

Given a strategy profile  $\sigma$ , we can compute the probability  $P_\sigma(x)$  that each node  $x \in X$  is reached. Let  $P_\sigma(I) = \sum_{x \in I} P_\sigma(x)$ .

The belief profile  $\mu$  is *Bayesian given profile  $\sigma$*  if  $\mu_i(x) = \frac{P_\sigma(x)}{P_\sigma(I)}$  whenever  $P_\sigma(I) > 0$ . In words: beliefs are determined by conditional probabilities whenever possible.

*Example 2.44. Entry deterrence II revisited: Bayesian beliefs.*



Suppose 1 plays  $\frac{1}{4}O + \frac{1}{4}M + \frac{1}{2}B$ . Then  $P_\sigma(x) = \frac{1}{4}$ ,  $P_\sigma(y) = \frac{1}{2}$ , and  $P_\sigma(I) = \frac{3}{4}$ , so  $\mu_2(x) = \frac{1/4}{3/4} = \frac{1}{3}$  and  $\mu_2(y) = \frac{2}{3}$ .  $\diamond$

At information sets  $I$  on the path of play, Bayesian beliefs describe the conditional probabilities of nodes being reached.

But beliefs are most important at unreached information sets. In this case, they represent "conditional probabilities" after a probability zero event has occurred.

In Section 2.4.3 we impose a key restriction on allowable beliefs in just this circumstance.

### *Sequential rationality*

Given a node  $x$ , let  $u_i(\sigma|x)$  denote player  $i$ 's expected utility under strategy profile  $\sigma$  conditional on node  $x$  being reached.

We call strategy  $\sigma_i$  *rational starting from information set*  $I \in \mathbb{I}_i$  given  $\sigma_{-i}$  and  $\mu_i$  if

$$(17) \quad \sum_{x \in I} \mu_i(x) u_i(\sigma_i, \sigma_{-i}|x) \geq \sum_{x \in I} \mu_i(x) u_i(\hat{\sigma}_i, \sigma_{-i}|x) \quad \text{for all } \hat{\sigma}_i.$$

In words: starting at information set  $I$ , strategy  $\sigma_i$  generates the best possible payoff for player  $i$  given his beliefs at  $I$  and the other players' strategies. (Remark: The definition only depends on choices under  $\sigma_i$  and  $\sigma_{-i}$  from information set  $I$  onward, and on beliefs  $\mu_i$  at information set  $I$ .)

If (17) holds for every information set  $I \in \mathbb{I}_i$ , we call strategy  $\sigma_i$  *sequentially rational given*  $\sigma_{-i}$  and  $\mu_i$ . If for a given  $\sigma$  and  $\mu$  this is true for all players, we call strategy profile  $\sigma$  *sequentially rational given*  $\mu$ .

If the information set  $I = \{x\}$  is a singleton (as many are), then beliefs are trivial, so (17) reduces to

$$u_i(\sigma_i, \sigma_{-i}|x) \geq u_i(\hat{\sigma}_i, \sigma_{-i}|x) \quad \text{for all } \hat{\sigma}_i.$$

Thus the full glory of (17) is only needed at nontrivial information sets.

A pair  $(\sigma, \mu)$  consisting of a strategy profile  $\sigma$  and a belief profile  $\mu$  is called an *assessment*.

The assessment  $(\sigma, \mu)$  is a *weak sequential equilibrium* if

- (i)  $\mu$  is Bayesian given  $\sigma$ .
- (ii)  $\sigma$  is sequentially rational given  $\mu$ .

Remarks:


1. One can verify that  $\sigma$  is a Nash equilibrium if and only if (i) there are beliefs  $\mu$  that are Bayesian given  $\sigma$  and (ii) for each player  $i$  and each information set  $I \in \mathbb{I}_i$  on the path of play (i.e., with  $P_\sigma(I) > 0$ ),  $\sigma_i$  is rational at information set  $I \in \mathbb{I}_i$  given  $\sigma_{-i}$  and  $\mu_i$ . (For instance, in Example 2.43, if player 1 plays  $O$ , player 2 need not play a best response given his beliefs, so  $(O, F)$  is Nash.) It follows that weak sequential equilibrium is a refinement of Nash equilibrium.
2. The concept we call "weak sequential equilibrium" is called different names by different authors: "perfect Bayesian equilibrium", "weak perfect Bayesian equilibrium"... Moreover, these names are also used by different authors to mean

slightly different things. But fortunately, everyone agrees about the meaning of “sequential equilibrium”, and this is the important concept.

*Example 2.45. Entry deterrence II once more.* Now player 2 must play a best response to some beliefs  $\mu_2$ . Since  $A$  is dominant at this information set, 2 must play it. Thus 1 must play  $B$ , and so  $((B, A), \mu_2(y) = 1)$  is the unique weak sequential equilibrium.  $\diamond$

Here weak sequential equilibrium was sufficiently restrictive. In general it is not.

*Example 2.46. Entry deterrence III.*



In this game the entrant moves in two stages. First, he chooses between staying out ( $O$ ) and entering ( $E$ ); if he enters, he then chooses between entering foolishly ( $F$ ) and entering cleverly ( $C$ ). Entering foolishly (i.e., choosing  $(E, F)$ ) is strictly dominated.

If the entrant enters, the incumbent cannot observe which kind of entry has occurred. Fighting ( $f$ ) is optimal against foolish entry, but accommodating ( $a$ ) is optimal against clever entry.

This is a game of stagewise perfect information, so it can be analyzed fully using subgame perfect equilibrium. The unique subgame perfect equilibrium of the game is  $((E, C), a)$ : player 1 plays the dominant strategy  $C$  in the subgame, so 2 plays  $a$ , and so 1 plays  $E$ . This corresponds to a weak sequential equilibrium (with beliefs  $\mu_2(y) = 1$  determined by taking conditional probabilities).

In addition,  $((O, \cdot), \sigma_2(f) \geq \frac{1}{2})$  are Nash equilibria.

These Nash equilibria correspond to a component of weak sequential equilibria (that are not subgame perfect!):

$$\begin{aligned} &((O, C), f) \text{ with } \mu_2(x) \geq \frac{2}{3}; \\ &((O, C), \sigma_2(f) \geq \frac{1}{2}) \text{ with } \mu_2(x) = \frac{2}{3}. \end{aligned}$$

Why are these weak sequential equilibria? If 1 plays  $O$ , condition (ii) places no restriction on beliefs at 2's information set. Therefore, player 2 can put all weight on  $x$ , despite the fact  $F$  is a dominated strategy.  $\diamond$

Thus, to obtain an appealing solution concept for games of imperfect information, we need a stronger restriction than Bayesian beliefs.

### 2.4.3 Definition of sequential equilibrium

#### *Definition*

An assessment  $(\sigma, \mu)$  is a *sequential equilibrium* (Kreps and Wilson (1982)) if

- (i)  $\mu$  is *(KW)-consistent* given  $\sigma$ : There exists a sequence of completely mixed strategy profiles  $\{\sigma^k\}_{k=1}^\infty$  such that
  - (1)  $\lim_{k \rightarrow \infty} \sigma^k = \sigma$ , and
  - (2) if  $\mu^k$  are the unique Bayesian beliefs for  $\sigma^k$ , then  $\lim_{k \rightarrow \infty} \mu^k = \mu$ .
- (ii)  $\sigma$  is sequentially rational given  $\mu$ .

Consistency requires that  $\mu$  be close to beliefs that could arise if players had small probabilities of making mistakes at each information set.

#### Remarks

- (i) The “there exists” in the definition of consistency is important. In general, one can support more than one consistent belief profile  $\mu$  for a given strategy profile  $\sigma$  by considering different perturbed strategy profiles  $\sigma^k$ .  
Later we consider perfect and proper equilibrium, which also use “there exists” definitions (see the discussion before Example 2.62). We also consider KM stability, which instead uses a “for all” definition, and so imposes a more demanding sort of robustness.
- (ii) Battigalli (1996) and Kohlberg and Reny (1997) offer characterizations of consistency that clarify what this requirement entails.

*Example 2.46 revisited.* Profile  $((E, C), a)$  is subgame perfect equilibrium and a weak sequential equilibrium (with  $\mu_2(y) = 1$ ). To see whether it is a sequential equilibrium, we

must check that  $\mu_2(y) = 1$  satisfies (ii). To do this, we must construct an appropriate sequence of perturbed strategy profiles, compute the implied beliefs by taking conditional probabilities, and show that we get  $\mu_2(y) = 1$  in the limit.

$$\begin{array}{l} \sigma_1(O) = \varepsilon_O \quad \sigma_1(E) = 1 - \varepsilon_O \\ \sigma_1(F) = \varepsilon_F \quad \sigma_1(C) = 1 - \varepsilon_F \end{array} \implies \mu_2(y) = \frac{P_\sigma(y)}{P_\sigma(y) + P_\sigma(x)} = \frac{(1 - \varepsilon_O)(1 - \varepsilon_F)}{1 - \varepsilon_O} \rightarrow 1 \checkmark$$

The assessments  $(((O, C), f), \mu(x) \geq \frac{2}{3})$  are weak sequential equilibria (but not subgame perfect equilibria).

$$\begin{array}{l} \sigma_1(O) = 1 - \varepsilon_E \quad \sigma_1(E) = \varepsilon_E \\ \sigma_1(F) = \varepsilon_F \quad \sigma_1(C) = 1 - \varepsilon_F \end{array} \implies \mu_2(y) = \frac{\varepsilon_E(1 - \varepsilon_F)}{\varepsilon_E} = 1 - \varepsilon_F \rightarrow 1 \Downarrow$$

Therefore,  $(((E, C), a), \mu(y) = 1)$  is the unique sequential equilibrium.  $\diamond$

### Checking consistency of beliefs


It can be shown that for any strategy profile  $\sigma$ , there exists a corresponding profile of consistent beliefs  $\mu$ . So if one can rule out all but one profile of beliefs as candidates for consistency, this surviving profile must be consistent. This ruling out can be accomplished using the following *implications of consistency*: if  $\mu$  is consistent given  $\sigma$ , then

- (i)  $\mu$  is Bayesian given  $\sigma$ . In fact, if all information sets are reached under  $\sigma$ , then  $\mu$  is consistent if and only if it is Bayesian (cf the previous example).
- (ii) *Preconsistency*: If by changing his strategy a player can force one of his own unreached information sets to be reached, his beliefs there should be as if he did so.

More precisely: If player  $i$ 's information set  $I'$  follows his information set  $I$ , and if  $i$  has a deviation  $\hat{\sigma}_i$  such that play under  $(\hat{\sigma}_i, \sigma_{-i})$  starting from  $I$  reaches  $I'$  with positive probability, then  $i$ 's beliefs at  $I'$  are determined by conditional probabilities under  $(\hat{\sigma}_i, \sigma_{-i})$  starting from  $I$ . (One can show that the beliefs at  $I'$  so defined are independent of the deviation  $\hat{\sigma}_i$  specified.)

*Example 2.47.* Below, preconsistency requires  $\mu_1(x) = \frac{1}{2} = \mu_1(y)$ . The same would be true if the tree shown here were an unreached portion of a larger game.





(iii) *Parsimony*: Let  $D_x$  be the set of deviations from  $\sigma$  required to reach  $x$ . If  $x, y \in I \in \mathbb{I}_i$  and  $D_y$  is a strict subset of  $D_x$ , then  $\mu_i(x) = 0$ .

(This rules out the beliefs in the bad weak sequential equilibria from Example 2.46. Parsimony actually implies preconsistency, but preconsistency is interesting in its own right—see Proposition 2.50.)

(iv) *Stagewise consistency*: Let  $\Gamma_x$  be a subgame of  $\Gamma$  that begins with a “simultaneous move game” among some subset  $Q \subseteq \mathcal{P}$  of the players. ( $\Gamma_x$  thus begins with one information set for each player in  $Q$ . These information sets must occur in some order, but each player’s information set will contain a distinct node for each combination of actions that could be played at the earlier information sets in  $\Gamma_x$ .) Then at information sets in this simultaneous move game, each player’s beliefs are determined by the other players’ strategies in the simultaneous move game.

(This also rules out the beliefs in the bad weak sequential equilibria from Example 2.46.)


*Example 2.48.* In *games of stagewise perfect information* (Section 2.3.4), the players play a sequence of simultaneous move games during stages  $t = 0, 1, \dots, T$ , with the choices made in each stage being observed before the next stage begins, and where the game that is played in stage  $t$  may depend on the choices made in previous stages. (Moves by Nature can also be introduced at the start of each stage.)

Let  $\Gamma$  be such a game. Each history describing all actions played through stage  $t - 1 < T$  in  $\Gamma$  determines a stage  $t$  subgame of  $\Gamma$ . In every stage  $T$  subgame, stagewise consistency ensures that each player’s beliefs about the other players’ choices in the subgame are given by their strategies. Therefore, each player’s behavior in the subgame is sequentially rational if and only if the players’ joint behavior in the subgame is a Nash equilibrium. Applying backward induction then shows that in any game with stagewise perfect infor-

mation, the set of subgame perfect equilibria and the set of sequential equilibrium strategy profiles are identical.  $\diamond$

- (v) *Cross-player consistency*: Players with the same information must have the same beliefs about opponents' deviations.

*Example 2.49.* In the game tree below, if player 1 chooses  $A$ , consistency requires that  $\mu_2(y) = \mu_3(\hat{y})$ . The reason is that both players' beliefs are generated by the same perturbation of player 1's strategy.



One can regard cross-player consistency as an extension of the usual equilibrium assumption: In a Nash equilibrium, players agree about opponents' behavior off the equilibrium path. Cross-player consistency requires them to agree about the relative likelihoods of opponents' deviations from equilibrium play.


***A one-shot deviation principle for games of imperfect information***

When determining whether a strategy profile  $\sigma$  is sequentially rational given beliefs  $\mu$ , do we need to explicitly consider multiperiod deviations? The following theorem says that we need not, so long as beliefs are preconsistent.

**Theorem 2.50** (Hendon et al. (1996)). *Let  $\Gamma$  be a finite extensive form game with perfect recall, let  $\sigma$  be a strategy profile for  $\Gamma$ , and suppose that belief profile  $\mu$  is preconsistent given  $\sigma$ . Then  $\sigma$  is sequentially rational given  $\mu$  if and only if no player  $i$  has a profitable one-shot deviation from  $\sigma_i$  given  $\sigma_{-i}$  and  $\mu_i$ .*

### *Relationships among extensive form refinements*

Kreps and Wilson (1982) show that all sequential equilibrium strategy profiles are subgame perfect equilibria. Combining this with earlier facts gives us the following diagram of the relationships among refinements.



#### 2.4.4 Computing sequential equilibria

To compute all sequential equilibria of a game:

- (i) First take care of easy parts of  $\sigma$  and  $\mu$ .
- (ii) Then try all combinations of remaining pieces, working backward through each player's information sets.

By the one-shot deviation principle, we may evaluate optimality of choices at each information set separately.

Remark: Although sequential equilibrium does not allow the play of strictly dominated strategies, the set of sequential equilibria of a game  $\Gamma$  may differ from that of a game  $\Gamma'$  that is obtained from  $\Gamma$  by removing a strictly dominated strategy: see Section 2.5, especially Examples 2.57 and 2.59.

*Example 2.51.* Ace-King-Queen Poker is a two-player card game that is played using a deck consisting of three cards: an Ace (the high card), a King (the middle card), and a Queen (the low card). Play proceeds as follows:

- Each player puts \$1 in a pot in the center of the table.
- The deck is shuffled, and each player is dealt one card. Each player only sees the card he is dealt.
- Player 1 chooses to Raise ( $R$ ) or Fold ( $F$ ). A choice of  $R$  means that player 1 puts an additional \$1 in the pot. Choosing  $F$  means that player 1 ends the game, allowing player 2 to have the money already in the pot.
- If player 1 raises, then player 2 chooses to Call ( $c$ ) or Fold ( $f$ ). A choice of  $f$  means that player 2 ends the game, allowing player 1 to have the money already in the pot. A choice of  $c$  means that player 2 also puts an additional \$1 in the pot; in this case, the players reveal their cards and the player with the higher card wins the money in the pot.

- (i) Draw the extensive form of this game.
- (ii) Find all sequential equilibria of this game.
- (iii) If you could choose whether to be player 1 or player 2 in this game, which player would you choose to be?

(i) See below. The probabilities of Nature's moves (which are  $\frac{1}{6}$  each) are not shown.

(This is an example of an (extensive form) Bayesian game. The natural interpretation is an ex ante one, in which the players do not know their types until play begins. See Section 4.4 for further discussion.)

(ii) The unique sequential equilibrium is

$$\sigma_1(R|A) = 1, \sigma_1(R|K) = 1, \sigma_1(R|Q) = \frac{1}{3}; \sigma_2(c|a) = 1, \sigma_2(c|k) = \frac{1}{3}, \sigma_2(c|q) = 0,$$

with the corresponding Bayesian beliefs. (Here,  $\sigma_2(c|a)$  is the probability that a player 2 of type  $a$  chooses  $c$  in the event that player 1 raises. Below,  $\mu_2(A|k)$  will represent the probability that player 2 assigns to player 1 being of type  $A$  when player 2 is of type  $k$  and player 1 raises.)

To begin the analysis, notice that  $R$  is dominant for (player 1 of) type  $A$ ,  $c$  is dominant for type  $a$ , and  $f$  is dominant for type  $q$ . In addition,  $R$  is dominant for type  $K$ : If type  $K$  raises, the worst scenario for him has player 2 call with an Ace and fold with a Queen (as we just said she would). Thus type  $K$ 's lowest possible expected payoff from raising is  $\frac{1}{2}(-2) + \frac{1}{2}(1) = -\frac{1}{2}$ , which exceeds the payoff of  $-1$  that type  $K$  would get from folding.

All that remains is to determine  $\sigma_1(R|Q)$  and  $\sigma_2(c|k)$ .


Suppose  $\sigma_2(c|k) = 1$ . Then  $\sigma_1(F|Q) = 1$ , which implies that  $\mu_2(A|k) = 1$ , in which case 2 should choose  $f$  given  $k$ , a contradiction.

Suppose  $\sigma_2(c|k) = 0$ . Then  $\sigma_1(F|Q) = 0$ , which implies that  $\mu_2(A|k) = \frac{1}{2}$ , in which case 2 should choose  $c$  given  $k$ , a contradiction.

It follows that player 2 must be mixing when she is of type  $k$ . For this to be true, it must be that when she is of type  $k$ , her expected payoffs to calling and folding are equal; this is true when  $\mu_2(A|k) \cdot (-2) + (1 - \mu_2(A|k)) \cdot 2 = -1$ , and hence when  $\mu_2(A|k) = \frac{3}{4}$ . For this to be her belief, it must be that  $\sigma_1(R|Q) = \frac{1}{3}$ .

For player 1 to mix when he is of type  $Q$ , his expected payoffs to  $F$  and to  $R$  must be equal. This is true when

$$-1 = \frac{1}{2} \cdot (-2) + \frac{1}{2} (\sigma_2(c|k) \cdot (-2) + (1 - \sigma_2(c|k)) \cdot 1).$$



and hence when  $\sigma_2(c|k) = \frac{1}{3}$ .

(iii) Considering type profiles  $Ak, Aq, Ka, Kq, Qa,$  and  $Qk$  in that order, we find that player 1's expected payoff is

$$\frac{1}{6} \left( \left( \frac{1}{3} \cdot 2 + \frac{2}{3} \cdot 1 \right) + 1 + (-2) + 1 + \left( \frac{1}{3} \cdot (-2) + \frac{2}{3} \cdot (-1) \right) + \left( \frac{2}{3} \cdot (-1) + \frac{2}{9} \cdot 1 + \frac{1}{9} \cdot (-2) \right) \right) = -\frac{1}{9}.$$

Therefore, since this is a zero-sum game, player 2's expected payoff is  $\frac{1}{9}$ , and so it is better to be player 2.  $\diamond$

*Example 2.52.* Suppose we modify Ace-King-Queen Poker (Example 2.51) as follows: Instead of choosing between Raise and Fold, player 1 chooses between Raise and Laydown ( $L$ ). A choice of  $L$  means that the game ends, the players show their cards, and the player with the higher card wins the pot. Find the sequential equilibria of this game. Which player would you choose to be?

The unique sequential equilibrium of this modified game is

$$\sigma_1(R|A) = 1, \sigma_1(R|K) = 0, \sigma_1(R|Q) = \frac{1}{3}; \sigma_2(c|a) = 1, \sigma_2(c|k) = \frac{1}{3}, \sigma_2(c|q) = 0,$$

with the corresponding Bayesian beliefs. That is, the only change from part (ii) (apart


from replacing  $F$  with  $L$ ) is that player 1 now chooses  $L$  when he receives a King.

To see why, note that compared to the original game, the payoffs in this game are only different when 1 chooses  $L$  after  $Ak, Aq$ , and  $Kq$ : in these cases he now gets 1 instead of  $-1$ . As before,  $c$  is dominant for type  $a$ , and  $f$  is dominant for type  $q$ . Given these choices of player 2, the unique best response for type  $K$  is now  $L$ . (The reason is that player 1's expected payoff from choosing  $L$  when he has a King is  $\frac{1}{2}(-1) + \frac{1}{2}(1) = 0$ , which exceeds his expected payoff from  $R$  of  $\frac{1}{2}(-2) + \frac{1}{2}(1) = -\frac{1}{2}$ .)

Unlike before,  $R$  is now only weakly dominant for  $A$ : since  $\sigma_2(f|q) = 1$ ,  $L$  is a best response for type  $A$  when  $\sigma_2(f|k) = 1$ . But if  $\sigma_2(f|k) = 1$ , then  $\sigma_1(R|Q) = 1$ , which implies that  $\mu_2(A|k) \leq \frac{1}{2}$ , which implies that 2 should choose  $c$  given  $k$ , a contradiction. Therefore,  $\sigma_2(f|k) < 1$ , which implies that  $\sigma_1(R|A) = 1$ . Thus, all that is left is to determine  $\sigma_1(R|Q)$  and  $\sigma_2(c|k)$ , and this is done exactly as before.

In computing player 1's equilibrium payoff, the only change is that when the type profile is  $Ka$ , player 1 gets a payoff of  $-1$  rather  $-2$ . This changes the expected payoff in the game from  $-\frac{1}{9}$  to  $-\frac{1}{9} + \frac{1}{6} = \frac{1}{18}$ . Therefore, the change in the rules makes it preferable to be player 1.  $\diamond$

*Example 2.53. Centipede with a possibly altruistic opponent (Myerson (1991), based on Kreps et al. (1982)).*



Start with a four-round Centipede game (Example 2.34) in which continuing costs you 1 but benefits your opponent 5. (If this was it, "always stop" would be the unique subgame perfect equilibrium.)

Suppose that player 1 assigns probability  $\frac{1}{20}$  to player 2 being altruistic, meaning that she cares about the total amount of money the players receive.

(This is an example of an (extensive form) Bayesian game. While we could give this game an ex ante interpretation, in which player 2 does not know her type until play begins, it is more natural to give it an interim interpretation, in which player 2 is normal, and the initial node is there to capture player 1's uncertainty about player 2's type. See Section 4.4 for further discussion.)

Notation for actions: capital letters for player 1; hats for the altruistic type of player 2, primes for a player's second information set.

Computation of equilibria:

The easy parts:  $\mu_1(w) = \frac{19}{20}$  (since beliefs must be Bayesian);

$\sigma_2(\hat{c}) = \sigma_2(\hat{c}') = 1$  (by sequential rationality);

$\sigma_2(s') = 1$  (by sequential rationality).

From Bayesian beliefs: If  $\sigma_1(C) > 0$ , then

$$(†) \quad \mu_1(y) = \frac{P_\sigma(y)}{P_\sigma(y) + P_\sigma(z)} = \frac{\frac{19}{20}\sigma_1(C)\sigma_2(c)}{\frac{19}{20}\sigma_1(C)\sigma_2(c) + \frac{1}{20}\sigma_1(C)} = \frac{19\sigma_2(c)}{19\sigma_2(c) + 1}.$$

And if  $\sigma_1(C) = 0$ , preconsistency implies that (†) still describes the only consistent beliefs.

At his second information set, 1 can choose  $S'$ ,  $C'$ , or a mixed strategy.

Suppose first that 1 plays  $S'$ . For this to be optimal for player 1, his payoff to  $S'$  is at least as big as his payoff to  $C'$ . His beliefs thus must satisfy  $4 \geq 3\mu_1(y) + 8(1 - \mu_1(y))$ , or, equivalently,  $\mu_1(y) \geq \frac{4}{5}$ . But if 1 plays  $S'$ , it is optimal for player 2 to play  $s$ , which by equation (†) implies that  $\mu_1(y) = 0$ , a contradiction.

Now suppose that 1 plays  $C'$ . For this to be optimal for player 1, it must be that  $\mu_1(y) \leq \frac{4}{5}$ . But if 1 plays  $C'$ , it is optimal for player 2 to play  $c$ , so equation (†) implies that  $\mu_1(y) = \frac{19}{20}$ , another contradiction.

So player 1 mixes at his second information set. For this to be optimal, it must be that

$$\mu_1(y) = \frac{4}{5}.$$

Thus equation (†) implies that  $\frac{19\sigma_2(c)}{19\sigma_2(c)+1} = \frac{4}{5}$ , and hence that

$$\sigma_2(c) = \frac{4}{19}.$$

For 2 to be willing to mix, she must be indifferent between  $s$  and  $c$ .

$$\Rightarrow 5 = (1 - \sigma_1(C')) \cdot 4 + \sigma_1(C') \cdot 9$$

$$\Rightarrow \sigma_1(C') = \frac{1}{5}$$

At his first information set, 1 gets 0 for playing  $S$  versus

$$\frac{19}{20} \left( \frac{15}{19} \cdot (-1) + \frac{4}{19} \left( \frac{4}{5} \cdot 4 + \frac{1}{5} \cdot 3 \right) \right) + \frac{1}{20} \left( \frac{4}{5} \cdot 4 + \frac{1}{5} \cdot 8 \right) = \frac{1}{4}$$

for playing  $C$ . Therefore, player 1 continues, and the unique sequential equilibrium is

$$(C, \frac{4}{5}S' + \frac{1}{5}C'), (\frac{15}{19}s + \frac{4}{19}c, s', \hat{c}, \hat{c}') \text{ with } \mu_1(w) = \frac{19}{20}, \mu_1(y) = \frac{4}{5}.$$

The point of the example: By making 1 uncertain about 2's type, we make him willing to continue. But this also makes 2 willing to continue, since by doing so she can keep player 1 uncertain and thereby prolong the game.

Kreps et al. (1982) introduced this idea to show that introducing a small amount of uncertainty about one player's preferences can lead to long initial runs of cooperation in finitely repeated Prisoner's Dilemmas (Example 3.2).  $\diamond$


*Example 2.54.* Consider the following game played between two unions and a firm. First, union 1 decides either to make a concession to the firm or not to make a concession. Union 2 observes union 1's action, and then itself decides whether or not to make a concession to the firm. The firm then chooses between production plan  $a$  and production plan  $b$ . The firm has two information sets: it either observes that both firms have made concessions, or it observes that at least one firm did not make a concession.

Each union obtains \$4 if the firm chooses plan  $a$  and \$0 if the firm chooses plan  $b$ . In addition, each union loses \$2 for making a concession. If the firm chooses plan  $a$ , it obtains \$2 for each union that makes a concession. If the firm chooses plan  $b$ , it obtains \$1 for certain, and receives an additional \$1 if both unions make concessions. Utilities are linear in dollars.

- (i) Draw an extensive form representation of this game, and find all of its sequential equilibria.
- (ii) Now suppose that union 2 cannot observe union 1's decision. Draw an appropriate extensive form for this new game, and compute all of its sequential equilibria.
- (iii) The games in part (i) and (ii) each have a unique sequential equilibrium outcome, but the choices made on these equilibrium paths are quite different. Explain in words why the choices made on the equilibrium path in (i) cannot be made on the equilibrium path in (ii), and vice versa. Evidently, the differences here must hinge on whether or not player 2 can observe a deviation by player 1.



(i) The extensive form game is below.



In any sequential equilibrium, 3 chooses  $A$  and 2 chooses  $\hat{d}$ . To proceed, we split the analysis into cases according to 3's behavior at her right information set, which we call  $I$ . Notice that  $b$  is a best response for player 3 if and only if  $1 \geq 2(1 - \mu_3(z))$ , which is equivalent to  $\mu_3(z) \geq \frac{1}{2}$ .

If 3 plays  $a$ , then 2 plays  $d$ , so 1 plays  $D$ . But then 3 prefers  $b$ . Contradiction.

If 3 plays  $b$ , then 2 plays  $c$  and 1 plays  $C$ , and so  $I$  is unreachable. For  $b$  to be a best response for 3, it must be that  $\mu_3(z) \geq \frac{1}{2}$ . Since 2 chooses  $\hat{d}$ , parsimony implies that  $\mu_3(y) = 0$ . Given our freedom in choosing  $\varepsilon_D^k$  and  $\varepsilon_{\hat{d}}^k$ , it is not difficult to verify that any beliefs satisfying  $\mu_3(z) \geq \frac{1}{2}$  and  $\mu_3(y) = 0$  are consistent. Thus  $(C, (c, \hat{d}), (A, b))$  with  $\mu_3(z) \geq \frac{1}{2}$  and  $\mu_3(y) = 0$  are sequential equilibria.

Now suppose that 3 plays a mixed strategy. Then  $\mu_3(z) = \frac{1}{2}$ . We divide the analysis into subcases:

First, suppose that  $I$  is off the equilibrium path: that is, 2 plays  $c$  and 1 plays  $C$ . We noted above that the belief  $\mu_3(z) = \frac{1}{2}$  and  $\mu_3(y) = 0$  is consistent in this case. Choosing  $c$  is optimal for 2 if  $4\sigma_3(a) \leq 2$ , or equivalently if  $\sigma_3(a) \leq \frac{1}{2}$ ; the same condition describes when choosing  $C$  is optimal for 1. Thus,  $(C, (c, \hat{d}), (A, \sigma_3(b) \geq \frac{1}{2}))$  with  $\mu_3(z) = \frac{1}{2}$  and  $\mu_3(y) = 0$  are sequential equilibria.

Second, suppose that  $I$  is on the equilibrium path. In this case, since beliefs are Bayesian and 2 plays  $\hat{d}$ , we have that

$$\mu_3(z) = \frac{\sigma_1(D)}{(1 - \sigma_1(D))\sigma_2(d) + \sigma_1(D)}$$

Since  $\mu_3(z) = \frac{1}{2}$ , it follows that

$$(\dagger) \quad \sigma_2(d) = \frac{\sigma_1(D)}{1 - \sigma_1(D)}.$$

We split again into subcases:

If 2 plays  $c$ , then (†) implies that 1 plays  $C$ , in which case  $I$  is off the equilibrium path. Contradiction.

If 2 plays  $d$ , then (†) implies that 1 plays  $\frac{1}{2}C + \frac{1}{2}D$ , but 1's best response is  $D$ . Contradiction.

If 2 mixes, then she is indifferent, which implies that  $\sigma_3(a) = \frac{1}{2}$ . Taking this into account, we find that  $C$  is optimal for 1 if and only if  $2\sigma_2(c) \geq 2$ . Since 2 is mixing ( $\sigma_2(c) < 1$ ), it follows that 1 chooses  $D$ . But this contradicts (†).

To sum up, there is a single component of sequential equilibria:

$$\begin{aligned} &(C, (c, \hat{d}), (A, b)) \text{ with } \mu_3(z) \geq \frac{1}{2} \text{ and } \mu_3(y) = 0; \\ &(C, (c, \hat{d}), (A, \sigma_3(b) \geq \frac{1}{2})) \text{ with } \mu_3(z) = \frac{1}{2} \text{ and } \mu_3(y) = 0. \end{aligned}$$

These equilibria generate payoffs  $(2, 2, 4)$ .

(ii) The extensive form game is below.

In any sequential equilibrium, 3 chooses  $A$ . We again split the analysis into cases according to 3's behavior at her right information set  $I$ ; as before,  $b$  is a best response for 3 if and only if  $\mu_3(z) \geq \frac{1}{2}$ .

If 3 plays  $a$ , then 2 plays  $d$ , so 1 plays  $D$ . But then 3 prefers  $b$ . Contradiction.

Suppose that 3 plays  $b$ . We split the analysis into subcases:

Suppose that 2 plays  $c$ . Then 1 plays  $C$ . Then parsimony implies that  $\mu_3(z) = 0$ , and so 3 prefers  $a$ . Contradiction.

Suppose that 2 plays  $d$ . Then 1 plays  $D$ , and 2 and 3's choices are optimal. Thus  $(D, d, (A, b))$  (with  $\mu_2(w) = 1$  and  $\mu_3(z) = 1$ ) is a sequential equilibrium.

Suppose that 2 mixes. Then she must be indifferent, implying that  $2\sigma_1(C) - 2(1 - \sigma_1(C)) = 0$ , or equivalently that  $\sigma_1(C) = \frac{1}{2}$ . Thus 1 must be indifferent, and a similar calculation shows that this implies that  $\sigma_2(c) = \frac{1}{2}$ . But these choices of 1 and 2 imply that  $\mu_3(z) = \frac{1}{3}$ , and so 3 prefers  $a$ . Contradiction.

Now suppose that 3 mixes. Then  $\mu_3(z) = \frac{1}{2}$ . We again split the analysis into subcases.

First suppose that  $I$  is off the equilibrium path, which is the case if 1 plays  $C$  and 2 plays  $c$ . Then parsimony implies that  $\mu_3(z) = 0$ , and so 3 prefers  $a$ . Contradiction.


Now suppose that  $I$  is on the equilibrium path. Then since beliefs are Bayesian, we have that

$$\mu_3(z) = \frac{\sigma_1(D)\sigma_2(d)}{(1 - \sigma_1(D))\sigma_2(d) + \sigma_1(D)}$$

Since  $\mu_3(z) = \frac{1}{2}$ , it follows that

$$(\ddagger) \quad 3\sigma_1(D)\sigma_2(d) = \sigma_1(D) + \sigma_2(d).$$

This equation defines a hyperbola in the plane. One point of intersection with the unit



square (where legitimate mixed strategies live) is the point  $(\sigma_1(D), \sigma_2(d)) = (0, 0)$ , but these choices prevent  $I$  from being unreachable, a contradiction. The remaining points of intersection have positive components, allowing us to rewrite (‡) as

$$(S) \quad 3 = \frac{1}{\sigma_1(D)} + \frac{1}{\sigma_2(d)}.$$

Along this curve  $\sigma_2(d)$  increases as  $\sigma_1(D)$  decreases, and the curve includes points  $(\sigma_1(D), \sigma_2(d)) = (1, \frac{1}{2}), (\frac{2}{3}, \frac{2}{3}),$  and  $(\frac{1}{2}, 1)$ .

A calculation shows that for player 1 to prefer  $D$  and player 2 to prefer  $d$ , we must have

- (a)  $\sigma_3(a)(1 - \sigma_2(d)) + \sigma_2(d) \geq \frac{1}{2}$  and
- (b)  $\sigma_3(a)(1 - \sigma_1(D)) + \sigma_1(D) \geq \frac{1}{2},$

respectively. We now show that (a) and (b) are inconsistent with (S). If  $(\sigma_1(D), \sigma_2(d)) = (1, \frac{1}{2})$ , then since player 2 is mixing, (b) must bind, but since  $\sigma_1(D) = 1$ , it does not—a contradiction. Similar reasoning shows that  $(\sigma_1(D), \sigma_2(d)) = (\frac{1}{2}, 1)$  is impossible. Thus for (S) to hold, players 1 and 2 must both be mixing, and so (a) and (b) must both bind. It follows that  $\sigma_2(d) = \sigma_1(D)$ , and hence, from (S), that  $\sigma_2(d) = \sigma_1(D) = \frac{2}{3}$ . But then the equality in (a) implies that  $\sigma_3(a) = -\frac{1}{2}$ , which is impossible.

Thus, the unique sequential equilibrium is  $(D, d, (A, b))$  (with  $\mu_2(w) = 1$  and  $\mu_3(z) = 1$ ), which generates payoffs  $(0, 0, 1)$ . Evidently, if player 2 cannot observe player 1's choice, all three players are worse off.

(iii) Why can't  $(D, \hat{d}, b)$  be chosen on the equilibrium path in part (i)? If player 3 plays  $(A, b)$ , then player 2 will play  $c$  at her left node, while  $\hat{d}$  is always a best response at her right node. If player 1 is planning to play  $D$ , he knows that when he switches to  $C$ , player 2 will observe this and play  $c$  rather than  $\hat{d}$ , which makes this deviation profitable.

Why can't  $(C, c, A)$  be chosen on the equilibrium path in part (ii)? If 1 and 2 are playing  $(C, c)$ , player 3's information set  $I$  is unreachable. If a deviation causes  $I$  to be reached, then

since 2 cannot observe 1's choice, it follows from parsimony that 3 may not believe that this is the result of a double deviation leading to  $z$ . Thus 3 must play  $a$  at  $I$ . Since 2 anticipates that 3 will play  $a$  at  $I$ , 2 is better off deviating to  $d$ . (The same logic shows that 1 is also better off deviating to  $D$ .)  $\diamond$

#### 2.4.5 Existence of sequential equilibrium and structure of the equilibrium set

Let  $\Gamma$  be a finite extensive form game with perfect recall.

**Theorem 2.55** (Kreps and Wilson (1982)).  $\Gamma$  has at least one sequential equilibrium.

In the next result, a *sequential equilibrium outcome* is the distribution over terminal nodes generated by some sequential equilibrium.

**Theorem 2.56** (Kreps and Wilson (1982), Kohlberg and Mertens (1986)).

- (i) *The set of sequential equilibrium strategy profiles of  $\Gamma$  consists of a finite number of connected components.*
- (ii) *In games with generic choices of payoffs, there are a finite number of sequential equilibrium outcomes; in particular, the outcome is constant on each component of sequential equilibria.*

For more on the (semialgebraic) structure of the set of sequential equilibria and solution sets for other equilibrium concepts, see Blume and Zame (1994).

Theorem 2.56(ii) is also true if we replace “sequential equilibrium” with “Nash equilibrium”. The restriction to generic choices of payoffs in  $\Gamma$  rather than in  $G$  is important: reduced normal forms of most extensive form games have payoff ties, and thus are non-generic in the space of normal form games. Notice also that the result is only about the set of equilibrium outcomes; often, components of equilibria contain an infinite number of strategy profiles which differ in their specifications of off-path behavior.

There are fundamental difficulties with defining sequential equilibrium in games with infinite action spaces. One basic problem is that with continuous action sets, one cannot perturb players' strategies in a way that makes every play path have positive probability. A definition of sequential equilibrium for this setting has recently been proposed by Myerson and Reny (2015), who “consider limits of strategy profiles that are approximately optimal (among all strategies in the game) on finite sets of events that can be observed by players in the game”.

## 2.5 Forward Induction

So far we have:

- defined Nash equilibrium for normal form games.
- used this definition (and the notion of the reduced normal form game) to define Nash equilibrium for extensive form games.


- considered refinements of Nash equilibrium for extensive form games to capture the principle of sequential rationality.

In the coming sections, we supplement equilibrium and sequential rationality by introducing additional principles for analyzing behavior in games.

### 2.5.1 Motivation and discussion

Does sequential equilibrium rule out all “unreasonable” predictions?

*Example 2.57. Battle of the Sexes with an outside option.* Consider the following game  $\Gamma$  and its reduced normal form  $G(\Gamma)$ :



Nash equilibria of the subgame of  $\Gamma$ :  $(T, L)$ ,  $(B, R)$ ,  $(\frac{3}{4}T + \frac{1}{4}B, \frac{1}{4}L + \frac{3}{4}R)$ .

$\Rightarrow$  subgame perfect equilibria of  $\Gamma$ :  $((I, T), L)$ ,  $((O, B), R)$ , and  $((O, \frac{3}{4}T + \frac{1}{4}B), \frac{1}{4}L + \frac{3}{4}R)$ .

Since  $\Gamma$  has stagewise perfect information, all three subgame perfect equilibria of  $\Gamma$  are sequential equilibria when combined with appropriate beliefs (cf. Example 2.48).

(Why?  $((I, T), L)$  has  $\mu(x) = 1$  (since  $x$  is reached),  $((O, B), R)$  has  $\mu(y) = 1$  (by parsimony), and  $((O, \frac{3}{4}T + \frac{1}{4}B), \frac{1}{4}L + \frac{3}{4}R)$  has  $\mu(x) = \frac{3}{4}$  (which is easy to compute directly).)

Nevertheless, only one of the three equilibria seems reasonable: If player 1 enters the subgame, he is giving up a certain payoff of 2. Realizing this, player 2 should expect him to play  $T$ , and then play  $L$  herself. We therefore should expect  $((I, T), L)$  to be played.  $\diamond$

Kohlberg and Mertens (1986) use this example to introduce the idea of *forward induction*:

“Essentially what is involved here is an argument of ‘forward induction’: a subgame should not be treated as a separate game, because it was preceded by a very specific form of preplay communication—the play leading to the subgame. In the above example, it is common knowledge that, when player 2 has to play in the subgame, preplay communication (for the subgame) has effectively ended with the following message from player 1

to player 2: ‘Look, I had the opportunity to get 2 for sure, and nevertheless I decided to play in this subgame, and my move is already made. And we both know that you can no longer talk to me, because we are in the game, and my move is made. So think now well, and make your decision.’

“Speeches” of this sort are often used to motivate forward induction arguments.

In the example above, forward induction can be captured by requiring that an equilibrium persist after a strictly dominated strategy is removed. Notice that strategy  $(I, B)$  is strictly dominated for player 1. If we remove this strategy, the unique subgame perfect equilibrium (and hence sequential equilibrium) is  $((I, T), L)$ . This example shows that none of these solution concepts is robust to the removal of strictly dominated strategies, and hence to a weak form of forward induction.

In general, capturing forward induction requires more than persistence after the removal of dominated strategies.

A stronger form of forward induction is captured by *equilibrium dominance*: an equilibrium should persist after a strategy that is suboptimal given the equilibrium outcome is removed (see Sections 2.5.2 and 2.7).

A general definition of forward induction for all extensive form games has been provided by Govindan and Wilson (2009).

For intuition, GW say:

“Forward induction should ensure that a player’s belief assigns positive probability only to a restricted set of strategies of other players. In each case, the restricted set comprises strategies that satisfy minimal criteria for rational play.”

GW’s formal definition is along these lines:

“A player’s pure strategy is called relevant for an outcome of a game in extensive form with perfect recall if there exists a weakly sequential equilibrium with that outcome for which the strategy is an optimal reply at every information set it does not exclude. The outcome satisfies forward induction if it results from a weakly sequential equilibrium in which players’ beliefs assign positive probability only to relevant strategies at each information set reached by a profile of relevant strategies.”

### *Iterated weak dominance and extensive form games*

One can capture forward induction in Example 2.57 by applying iterated removal of weakly dominated strategies to the normal form  $G(\Gamma)$ :  $B$  is strictly dominated for player 1; once  $B$  is removed,  $R$  is weakly dominated for player 2; once this is removed,  $O$  is strictly dominated for player 1, yielding the prediction  $(T, L)$ . Furthermore, iterated weak dominance is also powerful in the context of generic perfect information games, where

applying it to the reduced normal form yields the backward induction outcome, though not necessarily the game's backward induction solution: see Osborne and Rubinstein (1994), Marx and Swinkels (1997), and Østerdal (2005). But as we noted in Section 1.2.4, iterated removal of strategies and cautiousness conflict with one another. A resolution to this conflict is provided by Brandenburger et al. (2008), who provide epistemic foundations for the iterated removal of weakly dominated strategies.

Closely related to iterated weak dominance is *extensive form rationalizability* (Pearce (1984); Battigalli (1997)), which is based on there being common knowledge that players hold a hierarchy of hypotheses about how opponents will act, and that observing behavior inconsistent with the current hypothesis leads a player to proceed to the next unfalsified hypothesis in his hierarchy. Extensive form rationalizability generates the backward induction outcome (though not necessarily the backward induction solution) in generic perfect information games, and leads to the forward induction outcome in Example 2.57. Epistemic foundations for extensive form rationalizability are provided by Battigalli and Siniscalchi (2002).


### 2.5.2 Forward induction in signaling games

In a signaling game,

- (i) Player 1 (the *sender*) receives a private signal (his *type*) and then chooses an action (a *message*).
- (ii) Player 2 (the *receiver*), observing only the message, chooses an action herself (a *response*).

These games have many applications (to labor, IO, bargaining problems, etc.)

In signaling games, sequential equilibrium fails to adequately restrict predictions of play. We therefore introduce new refinements that capture forward induction, and that take the form of additional restrictions on out-of-equilibrium beliefs.



#### *Notation for signaling games*

$\mathcal{P} = \{1, 2\}$	the players (1 = the <i>sender</i> , 2 = the <i>receiver</i> )
$T$	finite set of player 1's types
$\pi$	prior distributions; $\pi(t) > 0$ for all $t \in T$
$A_1 = M = \{\dots, m, \dots\}$	player 1's finite set of actions ( <i>messages</i> )
$S_1 = \{s_1 : T \rightarrow M\}$	set of pure strategies for the sender
$\sigma_{1a}(m)$	probability that a type $t_a$ sender chooses message $m$
$A_2 = R = \{\dots, r, \dots\}$	player 2's finite set of actions ( <i>responses</i> )
$R^m \subseteq R$	responses available after message $m$
$S_2 = \{s_2 : M \rightarrow R \mid s_2(m) \in R^m\}$	player 2's set of pure strategies
$\sigma_2^m(r)$	probability that a receiver observing $m$ responds with $r$
$\mu_2^m \in \Delta(T)$	the receiver's beliefs after observing message $m$ .
$u_{1a}(m, r)$	a type $t_a$ sender's utility function
$u_2(t, m, r)$	the receiver's utility function

A type  $t_a$  sender's expected utility from message  $m$  given receiver strategy  $\sigma_2$  is

$$\sum_{r \in R^m} u_{1a}(m, r) \sigma_2^m(r).$$

If message  $m$  is sent, the receiver's expected utility from response  $r$  given beliefs  $\mu^m$  is

$$\sum_{t \in T} u_2(t, m, r) \mu_2^m(t).$$

$(\sigma, \mu)$  is a weak sequential equilibrium of  $\Gamma$  if

- (i) For each  $t_a \in T$ ,  $\sigma_{1a}(m) > 0 \Rightarrow m$  is optimal for player 1 of type  $t_a$  given  $\sigma_2$ .
- (ii) For each  $m \in M$ ,  $\sigma_2^m(r) > 0 \Rightarrow r$  is optimal for player 2 after  $m$  given  $\mu_2^m$ .
- (iii)  $\mu_2$  is Bayesian given  $\sigma$ .

**Proposition 2.58.** *In a signaling game, any Bayesian beliefs are consistent, and so every weak sequential equilibrium is a sequential equilibrium.*

This proposition says that any beliefs after an unsent message can be justified by introducing an appropriate sequence of perturbed strategies for player 1. Verifying this is a good exercise.




Recall that a sequential equilibrium outcome is the distribution over terminal nodes generated by some sequential equilibrium. We now focus on signaling games with a finite number of equilibrium outcomes. By Theorem 2.56, this is true for generic choices of payoffs, and it implies that the equilibrium outcome is constant on each of the finite number of components of sequential equilibrium strategy profiles.

Below we attempt to eliminate entire components of equilibria with unused messages as inconsistent with forward induction. (Surviving components may become smaller too.)

### Refinements for signaling games based on belief restrictions

A weak form of forward induction rules out equilibria that vanish after a dominated strategy is removed.

Example 2.59.



There are two components of sequential equilibria (computed below):

Good sequential equilibrium:  $(s_{1a} = O, s_{1b} = I, s_2 = D)$

Bad sequential equilibria:  $(s_{1a} = O, s_{1b} = O, s_2 = U, \mu_2(a) \geq \frac{1}{2})$   
 $(s_{1a} = O, s_{1b} = O, \sigma_2(U) \geq \frac{1}{2}, \mu_2(a) = \frac{1}{2})$

Suppose  $I$  is played. This message is strictly dominated for  $t_a$ , so the receiver really ought to believe she is facing  $t_b$  (i.e.,  $\mu_2(b) = 1$ ), breaking the bad equilibria.

Story: Suppose you are  $t_b$ . If you deviate, you tell the receiver: " $t_a$  would never want to deviate. If you see a deviation it must be me, so you should play  $D$ ".

If we introduce the requirement that equilibria be robust to the removal of dominated strategies, the bad equilibria are eliminated: If we eliminate action  $I$  for type  $t_a$ , then player 2 must play  $D$ , and so type  $t_b$  plays  $I$ .

Computation of equilibria:

We can treat each type of player 1 as a separate player.

Strategy  $O$  is strictly dominant for type  $t_a$ , so he plays this in any sequential equilibrium.


Now consider type  $t_b$ . If  $\sigma_{1b}(I) > 0$ , then  $\mu_2(t_b) = 1$ , so player 2 must play  $D$ , implying that type  $t_b$  plays  $I$ . Equilibrium.

If type  $t_b$  plays  $O$ , then player 2's information set is unreached, so her beliefs are unrestricted. Also, for  $O$  to be type  $t_b$ 's best response, it must be that  $0 \geq -\sigma_2(U) + (1 - \sigma_2(U))$ , or equivalently that  $\sigma_2(U) \geq \frac{1}{2}$ .  $\sigma_2(U) = 1$  is justified for player 2 whenever  $\mu_2(a) \geq \frac{1}{2}$ , while  $\sigma_2(U) \in [\frac{1}{2}, 1)$  is justified whenever  $\mu_2(a) = \frac{1}{2}$ . These combinations are sequential equilibria.  $\diamond$

Stronger forms of forward induction are based on *equilibrium dominance*: they rule out components of equilibria that vanish after a strategy that is not a best response at any equilibrium in the component is removed.

*Example 2.60. Beer-Quiche.*

The signalling game below is a version of the Beer-Quiche game of Cho and Kreps (1987). We present this game in different terms to avoid ridiculousness.



A broker (player 1) is offering an investment to an investor (player 2). Whether the investment is high quality ( $t_h$ ) or low quality ( $t_l$ ) is the broker's private information; the prior probability on the former is  $\frac{2}{3}$ . The broker tells the investor that the investment is high quality ( $H$ ) or that it is low quality ( $L$ ). The investor, who receives the broker's message but cannot observe the quality of the investment, chooses either to invest ( $I$ ) or not to invest ( $D$ ).

Payoffs are as follows: The investor receives a payoff of 2 for investing when the investment is of high quality, a payoff of  $-3$  for investing when the investment is of low quality, and

a payoff of 0 for not investing. The broker's payoff is the sum of up to two terms: he gets a payoff of 3 if the investor invests or 0 if she does not, and an additional payoff of 1 if his announcement about the quality of the investment is truthful.

Components of sequential equilibria (the computation is below):

$$\begin{array}{ll}
 \sigma_{1h}(H) = 1 = \sigma_{1\ell}(H) & \sigma_{1h}(L) = 1 = \sigma_{1\ell}(L) \\
 \sigma_2^H(I) = 1 & \sigma_2^L(I) = 1 \\
 (1) \quad \sigma_2^L(D) = 1 \quad (\geq \frac{1}{3}) & (2) \quad \sigma_2^H(D) = 1 \quad (\geq \frac{1}{3}) \\
 \mu_2^L(t_\ell) \geq \frac{2}{5} \quad (= \frac{2}{5}) & \mu_2^H(t_\ell) \geq \frac{2}{5} \quad (= \frac{2}{5})
 \end{array}$$

Are the equilibria in component (2) reasonable? Imagine that players expect such an equilibrium to be played. Then a salesman with a low-quality property is getting his highest possible payoff by reporting honestly; dishonestly sending message  $H$  can only reduce his payoff. Therefore, if investor receives message  $H$ , he should conclude that it was sent by a salesman with a high-quality property, and so should invest. Expecting this, a salesman with a high-quality property should deviate to his honest message  $H$ . To sum up, the fact that the salesman with a high-quality property wants to reveal this to the investor should lead his preferred equilibrium to be played.

In Example 2.59, certain beliefs were deemed unreasonable because they were based on expecting a particular sender type to play a dominated strategy. This is not the case here:  $H$  is not dominated by  $L$  for  $t_\ell$ . Instead, we fixed the component of equilibria under consideration, and then concluded that certain beliefs are unreasonable given the anticipation of equilibrium payoffs: the possible payoffs to  $H$  for  $t_\ell$  (which are 3 and 0) are all smaller than this type's equilibrium payoff to  $L$  (which is 4), so a receiver who sees  $H$  should not think he is facing  $t_\ell$ .

Computation of equilibria:

We divide the analysis into cases according to the choices of types  $t_h$  and  $t_\ell$ .

$(H, H)$ . In this case  $\mu_2^H(t_h) = \frac{2}{3}$  and hence  $s_2^H = I$ , since  $\frac{2}{3} \cdot 2 + \frac{1}{3} \cdot (-3) = \frac{1}{3} > 0$ .  $\mu_2^L$  is unrestricted. Playing  $H$  gives  $t_h$  his best payoff. Type  $t_\ell$  weakly prefers  $H$  iff  $3 \geq 4(1 - \sigma_2^L(D)) + \sigma_2^L(D)$ , and hence iff  $\sigma_2^L(D) \geq \frac{1}{3}$ .  $\sigma_2^L(D) = 1$  is justified if  $2(1 - \mu_2^L(t_\ell)) - 3\mu_2^L(t_\ell) \leq 0$ , or equivalently if  $\mu_2^L(t_\ell) \geq \frac{2}{5}$ . And  $\sigma_2^L(D) \in [\frac{1}{3}, 1)$  is justified if  $\mu_2^L(t_\ell) = \frac{2}{5}$ . These combinations form a component of sequential equilibria.

$(L, L)$ . In this case  $\mu_2^L(t_h) = \frac{2}{3}$  and hence  $s_2^L = I$ .  $\mu_2^H$  is unrestricted. Playing  $L$  gives  $t_\ell$  his best payoff. Type  $t_h$  weakly prefers  $L$  iff  $3 \geq 4(1 - \sigma_2^H(D)) + \sigma_2^H(D)$ , and hence iff  $\sigma_2^H(D) \geq \frac{1}{3}$ .  $\sigma_2^H(D) = 1$  is justified if  $\mu_2^H(t_\ell) \geq \frac{2}{5}$ , while  $\sigma_2^H(D) \in [\frac{1}{3}, 1)$  is justified if  $\mu_2^H(t_\ell) = \frac{2}{5}$ . These combinations form a component of sequential equilibria.

(H, L). In this case  $s_2^H = I$  and  $s_2^L = D$ . Thus type  $t_\ell$  obtains 1 for playing L and 3 for playing H, and so deviates to H.  $\uparrow$

(L, H). In this case  $s_2^L = I$  and  $s_2^H = D$ . Thus type  $t_\ell$  obtains 0 for playing H and 4 for playing L, and so deviates to L.  $\uparrow$

(mix, L). In this case  $\mu_2^H(t_h) = 1$ , so  $s_2^H = I$ , in which case  $t_h$  strictly prefers H.  $\uparrow$

(mix, H). In this case  $\mu_2^L(t_h) = 1$ , so  $s_2^L = I$ , in which case  $t_\ell$  strictly prefers L.  $\uparrow$

(H, mix). In this case  $\mu_2^H(t_h) > \frac{2}{3}$  and  $\mu_2^L(t_\ell) = 1$ , so  $s_2^H = I$  and  $s_2^L = D$ . But then  $t_\ell$  strictly prefers H.  $\uparrow$

(L, mix). In this case  $\mu_2^L(t_h) > \frac{2}{3}$  and  $\mu_2^H(t_\ell) = 1$ , so  $s_2^L = I$  and  $s_2^H = D$ . But then  $t_\ell$  strictly prefers L.  $\uparrow$

(mix, mix). For  $t_h$  to be indifferent, it must be that  $4\sigma_2^H(I) + 1(1 - \sigma_2^H(I)) = 3\sigma_2^L(I)$ , and hence that  $\sigma_2^L(I) = \sigma_2^H(I) + \frac{1}{3}$ . But for  $t_\ell$  to be indifferent, it must be that  $4\sigma_2^L(I) + 1(1 - \sigma_2^L(I)) = 3\sigma_2^H(I)$ , and hence that  $\sigma_2^H(I) = \sigma_2^L(I) + \frac{1}{3}$ .  $\uparrow \diamond$

We now consider refinements that formalize the notion of equilibrium dominance in signaling games. Cho and Kreps (1987), using results of Kohlberg and Mertens (1986), prove that at least one equilibrium outcome survives after any one of these refinements is applied.

For set of types  $I \subseteq T$ , let  $BR_2^m(I) \subseteq R^m$  be the set of responses to message  $m$  that are optimal for the receiver under some beliefs that put probability 1 on the sender's type being in  $I$ . Formally:

$$BR_2^m(\mu_2^m) = \operatorname{argmax}_{r \in R^m} \sum_{t \in T} u_2(t, m, r) \mu_2^m(t),$$

$$BR_2^m(I) = \bigcup_{\mu_2^m: \mu_2^m(I)=1} BR_2^m(\mu_2^m).$$

Fix a component of sequential equilibria of signaling game  $\Gamma$ , and let  $u_{1a}^*$  be the payoff received by type  $t_a$  on this component. (Recall that we are restricting attention to games in which payoffs are constant on every equilibrium component.)

(I) For each unused message  $m$ , let

$$D^m = \left\{ t_a \in T: u_{1a}^* > \max_{r \in BR_2^m(T)} u_{1a}(m, r) \right\}$$

$D^m$  is the set of types for whom message  $m$  is dominated by the equilibrium, given that the receiver behaves reasonably.

(II) If for some unused message  $m$  with  $D^m \neq T$  and some type  $t_b$ , we have

$$u_{1b}^* < \min_{r \in BR_2^m(T-D^m)} u_{1b}(m, r)$$

then component of equilibria fails the *Cho-Kreps criterion* (a.k.a. the *intuitive criterion*).

Type  $t_b$  would exceed his equilibrium payoffs by playing message  $m$  if the receiver played a best response to some beliefs that exclude types in  $D^m$ .

*Example 2.60 revisited.* Applying the Cho-Kreps criterion in the Beer-Quiche game.

Component (2):  $H$  is unused.

$$BR_2^H(\{t_h, t_\ell\}) = \{I, D\}$$

$$u_{1\ell}^* = 4 > 3 = u_{1\ell}(H, I)$$

$$> 0 = u_{1\ell}(H, D)$$

$t_\ell$  cannot benefit from deviating to  $H$

$$\text{But } u_{1h}^* = 3 < 4 = u_{1h}(H, I) \quad (*)$$

$t_h$  could benefit from deviating to  $H$

$$\Rightarrow D^H = \{t_\ell\}, T - D^H = \{t_h\}$$

only  $t_h$  could benefit from deviating to  $H$

$$BR_2^H(\{t_h\}) = \{I\}$$

$\Rightarrow$  by (\*), these equilibria fail the Cho-Kreps criterion.

(In words: Type  $t_\ell$  is getting his highest payoff in the equilibrium. Thus if a deviation to  $H$  occurs, the receiver should believe that it is type  $t_h$ , and so should play  $I$ . Given this, type  $t_h$  prefers to deviate, since he will get a payoff of 4 rather than 0. This breaks the equilibrium.)

Component (1):  $L$  is unused.

$$D^L = \{t_h\}, T - D^L = \{t_\ell\}$$

only  $t_\ell$  could benefit from deviating to  $L$

$$BR_2^L(\{t_\ell\}) = \{D\}$$

$$u_{1\ell}^* = 3 > 1 = u_{1\ell}(L, D)$$

$\Rightarrow$  these equilibria satisfy the Cho-Kreps criterion

(In words: Type  $t_h$  is getting his highest payoff in the equilibrium. Thus if a deviation to  $L$  occurs, the receiver should believe that it is type  $t_\ell$ , and so should play  $D$ . This is consistent with the equilibrium.)  $\diamond$

### *Further refinements*

We can eliminate more equilibria by replacing (II) with a weaker requirement.

For set of types  $I \subseteq T$ , let  $MBR_2^m(I) \subseteq R^m$  be the set of responses to message  $m$  that are optimal for the receiver under some beliefs that put probability 1 on the sender's type being in  $I$ . Formally:

$$MBR_2^m(I) = \bigcup_{\mu_2^m: \mu_2^m(I)=1} \{\hat{\sigma}_2^m \in \Delta R^m: \text{support}(\hat{\sigma}_2^m) \subseteq BR_2^m(\mu_2^m)\}.$$

(II') If for some unused message  $m$  and each  $\hat{\sigma}_2^m \in MBR_2^m(T - D^m)$ , there is a type  $t_b$  such that  $u_{1b}^* < \sum_{r \in R^m} u_{1b}(m, r) \hat{\sigma}_2^m(r)$ , then  $(\sigma, \mu)$  fails the *strong Cho-Kreps criterion* (a.k.a. the *equilibrium domination test*).

What's the difference between (II) and (II')?

- (i) Under (II), there is a single type who wants to deviate regardless of the BR the receiver chooses.
- (ii) Under (II'), the type can vary with the BR the receiver chooses. This sometimes allows us to rule out more equilibria.

Iterated versions of these concepts can be obtained by applying (I) repeatedly...

$$(D^m)' = \{t_a \in T: u_{1a}^* > \max_{r \in B_2^m(T-D^m)} u_{1a}(m, r)\}$$

...before applying (II) or (II').

The idea: Once some types have been ruled out as senders of message  $m$ , the set of possible best responses to  $m$  can become smaller, which can in turn cause  $m$  to become dominated for more types, allowing us to rule them out as well.

There are additional refinements that are stronger than equilibrium dominance. Banks and Sobel (1987) introduce refinements that are based on the following idea (called *D2* in Cho and Kreps (1987)): instead of following step (II) above, we exclude type  $t_a$  from having deviated to unused message  $m$  if for any mixed best response  $\hat{\sigma}_2^m$  for which  $t_a$  weakly prefers  $m$  to getting the equilibrium payoff, there is another type  $t_b$  that strictly prefers  $m$ . Iterated versions of this sort of requirement are called (*universal*) *divinity* by Banks and Sobel (1987). Under the *never a weak best response* criterion, we exclude type  $t_a$  from having deviated to unused message  $m$  if for any mixed best response  $\hat{\sigma}_2^m$  for which  $t_a$  is *indifferent* between playing  $m$  and getting the equilibrium payoff, there is another type  $t_b$  that strictly prefers  $m$ . All of these refinements are implied by *KM stability* (see Section 2.7); guaranteeing that components satisfying the refinements exist. For further discussion, see Cho and Kreps (1987) and Banks and Sobel (1987).

It may not be surprising that this profusion of different solution concepts has led to substantial criticism of the signaling game refinements literature.


## 2.6 Invariance and Proper Equilibrium

### *Backward induction, invariance, and normal form refinements*

The principle of *invariance* requires predictions of play in extensive form games with the same reduced normal form to be the same.

The idea is that games with the same normal form differ only in terms of how they are presented, so a theory of rational play should not treat them differently.

*Example 2.61. Entry deterrence III rerevisited.*



We saw earlier that in game  $\Gamma$ ,  $((E, C), a)$  is the unique subgame perfect equilibrium and the unique sequential equilibrium (with  $\mu_2(y) = 1$ ). There are additional weak sequential equilibria in which player 1 plays  $O$ : namely,  $((O, C), f)$  with  $\mu_2(x) \geq \frac{2}{3}$ , and  $((O, C), \sigma_2(f) \geq \frac{1}{2})$  with  $\mu_2(x) = \frac{2}{3}$ .

Notice that  $\Gamma$  and  $\Gamma'$  only differ in the how player 1's choices are presented. In particular, both of these games have the same reduced normal form:  $G(\Gamma) = G(\Gamma')$ .

In  $\Gamma'$ , consistency places no restrictions on beliefs. Therefore, all weak sequential equilibria of  $\Gamma$  above correspond to sequential equilibria in  $\Gamma'$ !

What is going on? When  $F$  and  $C$  are by themselves at a decision node, consistency forces player 2 to discriminate between them. But in  $\Gamma'$ ,  $F$  and  $C$  appear as choices at the same decision node as  $O$ , so when player 1 chooses  $O$ , consistency does not discriminate between  $F$  and  $C$ .  $\diamond$

One way to respect invariance is to perform analyses directly on reduced normal forms, so that invariance holds by default. This also has the advantage of mathematical simplicity, since normal form games are simpler objects than extensive form games.

Shifting the analysis to the reduced normal form may seem illegitimate. First, normal form and extensive form games differ in a fundamental way, since only in the latter is it possible to learn something about one's opponent during the course of play (see Example 2.32). Second, working directly with the normal form appears to conflict with the use of backward induction, whose logic seems tied to the extensive form.

Both of these criticisms can be addressed. First, the differences between extensive and normal form games are much smaller if we only consider equilibrium play: when players adhere to equilibrium strategies, nothing important is learned during the course of play. Second, the fact that a strategy for an extensive form game specifies a player's complete plan for playing a game suggests that the temporal structure provided by the extensive form may not be essential as it might seem. In fact, since an extensive form game creates a telltale pattern of ties in its reduced normal form, one can "reverse engineer" a reduced normal form to determine the canonical extensive form that generates it—see Mailath et al. (1993).

To implement the normal form approach, we require robustness of equilibrium to low probability mistakes, sometimes called *trembles*. Trembles ensure that all information sets of the corresponding extensive form game are reached.

### Perfect equilibrium

Throughout this section, we let  $G$  be a finite normal form game and  $\Gamma$  a finite extensive form game (with perfect recall).


Strategy profile  $\sigma$  is an  $\varepsilon$ -perfect equilibrium of  $G$  if it is completely mixed and if  $s_i \notin B_i(\sigma_{-i})$  implies that  $\sigma_i(s_i) \leq \varepsilon$ .

Strategy profile  $\sigma^*$  is a perfect equilibrium of  $G$  if and only if it is the limit of a sequence of  $\varepsilon$ -perfect equilibria with  $\varepsilon \rightarrow 0$ .

Remarks:

- (i) For  $\sigma^*$  to be a perfect equilibrium, it is only necessary that *there exist* a sequence of  $\varepsilon$ -perfect equilibria converging to  $\sigma^*$ . It need not be the case that *every* sequence of strategy profiles converging to  $\sigma^*$  consists of  $\varepsilon$ -perfect equilibria. An analogous point arose in the definition of consistency for sequential equilibrium (Section 2.4.3). The analogous point holds for proper equilibrium, but not for KM stable sets—see below.
- (ii) The formulation of perfect equilibrium above is due to Myerson (1978). The original definition, due to Selten (1975), is stated in terms of Nash equilibria of perturbed games  $G_p$ ,  $p: \bigcup_{i \in \mathcal{P}} S_i \rightarrow (0, 1)$ , in which player  $i$ 's mixed strategy must put at least probability  $p_{s_i}$  on strategy  $s_i$ .

Example 2.62. Entry deterrence revisited.





Nash equilibria:  $(E, A)$  and  $(O, \sigma_2(F) \geq \frac{1}{2})$ . Only  $(E, A)$  is subgame perfect.

What are the perfect equilibria of the normal form  $G(\Gamma)$ ? Since  $F$  is weakly dominated, 2's best response to any completely mixed strategy of 1 is  $A$ , so in any  $\varepsilon$ -perfect equilibrium,  $\sigma_2(F) \leq \varepsilon$ . It follows that if  $\varepsilon$  is small, 1's best response is  $E$ , so in any  $\varepsilon$ -perfect equilibrium,  $\sigma_1(O) \leq \varepsilon$ . Therefore, any sequence of  $\varepsilon$ -perfect equilibria with  $\varepsilon \rightarrow 0$  converges to  $(E, A)$ , which is thus the unique perfect equilibrium of  $G(\Gamma)$ .  $\diamond$

Selten (1975) establishes the following properties of perfect equilibrium:

**Theorem 2.63.**  *$G$  has at least one perfect equilibrium.*

**Theorem 2.64.** *Every perfect equilibrium of  $G$  is a Nash equilibrium which does not use weakly dominated strategies. In two-player games, the converse statement is also true.*

These results imply


**Corollary 2.65.**  *$G$  has at least one Nash equilibrium in which no player uses a weakly dominated strategy.*

Let  $\Gamma$  be a generic extensive form game of perfect information, so that  $\Gamma$  has a unique subgame perfect equilibrium. Will applying perfection to  $G(\Gamma)$  rule out Nash equilibria of  $\Gamma$  that are not subgame perfect?

If  $\Gamma$  has the single move property (i.e., if no player has more than one decision node on any play path), then the perfect equilibrium of  $G(\Gamma)$  is unique, and it is outcome equivalent to the unique subgame perfect equilibrium of  $\Gamma$ .

But beyond games with the single move property, perfect equilibrium is not adequate to capture backward induction.

*Example 2.66.* In the game  $\Gamma$  below,  $((B, D), R)$  is the unique subgame perfect equilibrium.  $((A, \cdot), L)$  are Nash equilibria. (Actually  $A$  with  $\sigma_2(L) \geq \frac{1}{2}$  are Nash too.)



		2		
		L	R	
1	A	2, 4	2, 4	$\approx 1$
	BC	1, 1	0, 0	$\varepsilon$
	BD	1, 1	3, 3	$\leq \frac{\varepsilon}{2}$
		$1 - \varepsilon$	$\varepsilon$	

Both  $(BD, R)$  and  $(A, L)$  are perfect.

Why  $(A, L)$ ? If in the  $\varepsilon$ -perfect equilibria, the weight on  $BC$  is at least double that on  $BD$ , then  $L$  is player 2's best response.

Implicitly, we are assuming that if 1's second node is reached, he is more likely to be dumb than smart. Sequential rationality forbids this, but perfect equilibrium does not.  $\diamond$

To handle this example, we need to force players to behave rationally at all information sets, even those which occur after he himself deviates.

### **Proper equilibrium**

To capture subgame perfection directly in the reduced normal form, we introduce a refinement that requires that more costly trembles be less likely to occur.

Strategy profile  $\sigma$  is  $\varepsilon$ -proper if it is completely mixed and if  $u_i(s_i, \sigma_{-i}) < u_i(s'_i, \sigma_{-i})$  implies that  $\sigma_i(s_i) \leq \varepsilon \sigma_i(s'_i)$ .

Strategy profile  $\sigma^*$  is a *proper equilibrium* (Myerson (1978)) if it is the limit of a sequence of  $\varepsilon$ -proper equilibria with  $\varepsilon \rightarrow 0$ .

*Example 2.66 rerevisited.* Recall that in  $\Gamma$ ,  $((A, \cdot), L)$  is Nash but not subgame perfect. Now we show that  $(A, L)$  is not proper.

$G(\Gamma)$		$2$			
		$L$	$R$		
	$1$	$A$	$2, 4$	$2, 4$	$\approx 1$
		$BC$	$1, 1$	$0, 0$	$\approx \varepsilon^2$
		$BD$	$1, 1$	$3, 3$	$\approx \varepsilon$

Why? If  $\sigma_2(R)$  is small but positive, then  $u_1(A) > u_1(BD) > u_1(BC)$ , so in any  $\varepsilon$ -proper equilibrium we have  $\sigma_1(BD) \leq \varepsilon \sigma_1(A)$  and  $\sigma_1(BC) \leq \varepsilon \sigma_1(BD)$ .

Therefore, player 2 puts most of her weight on  $R$  in any  $\varepsilon$ -proper equilibrium, and so  $L$  is not played in any proper equilibrium.  $\diamond$

Properties of proper equilibrium.

**Theorem 2.67** (Myerson (1978)).

- (i)  $G$  has at least one proper equilibrium.
- (ii) Every proper equilibrium of  $G$  is a perfect equilibrium of  $G$ .

It can be shown that if  $\Gamma$  is a game of perfect information, then every proper equilibrium of  $G(\Gamma)$  is *outcome equivalent* (i.e., induces the same distribution over terminal nodes) to some subgame perfect equilibrium of  $\Gamma$ .

Remarkably, proper equilibrium also captures sequential rationality in games of imperfect information:

**Theorem 2.68** (van Damme (1984), Kohlberg and Mertens (1986)).

- (i) Suppose that  $\sigma^*$  is a proper equilibrium of  $G(\Gamma)$ . Then there is an outcome equivalent behavior strategy profile  $\beta^*$  of  $\Gamma$  that is a sequential equilibrium strategy profile of  $\Gamma$ .
- (ii) Let  $\{\sigma^\varepsilon\}$  be a sequence of  $\varepsilon$ -proper equilibria of  $G(\Gamma)$  that converge to proper equilibrium  $\sigma^*$ . Let behavior strategy  $\beta^\varepsilon$  be outcome equivalent to  $\sigma^\varepsilon$ , and let behavior strategy  $\beta^*$  be a limit point of the sequence  $\{\beta^\varepsilon\}$ . Then  $\beta^*$  is a sequential equilibrium strategy profile of  $\Gamma$ .

What is the difference between parts (i) and (ii) of the theorem? In part (i),  $\sigma^*$  is outcome equivalent to some sequential equilibrium strategy profile  $\beta^*$ . But outcome equivalent strategy profiles may specify different behavior off the equilibrium path; moreover, the strategy  $\sigma_i^*$  for the reduced normal form does not specify how player  $i$  would behave at unreachable information sets (i.e., at information sets that  $\sigma_i^*$  itself prevents from being reached). In part (ii), the  $\varepsilon$ -proper equilibria are used to explicitly construct the behavior strategy profile  $\beta^*$ . Thus, part (ii) shows that the construction of proper equilibrium does not only lead to outcomes that agree with sequential equilibrium; by identifying choices off the equilibrium path, it captures the full force of the principle of sequential rationality.

Theorem 2.68 shows that proper equilibrium achieves our goals of respecting the principle of sequential rationality while ensuring invariance of predictions across games with the same purely reduced normal form.

**Extensive form perfect equilibrium and quasi-perfect equilibrium**

The agent normal form  $A(\Gamma)$  of an extensive form  $\Gamma$  is the reduced normal form we obtain if we assume that each information set is controlled by a distinct player. In particular, whenever the original game has players with more than one information set, we create new players to inhabit the information sets.

Profile  $\sigma$  is an extensive form perfect equilibrium of  $\Gamma$  (Selten (1975)) if it corresponds to a (normal form) perfect equilibrium of  $A(\Gamma)$ .

*Example 2.66 revisited: Direct computation of extensive form perfect equilibrium.*

The agent normal form of the game  $\Gamma$  above is

$A(\Gamma)$	3:C		3:D
	2		2
	L	R	
1	A	2, 4, 2	2, 4, 2
	B	1, 1, 1	0, 0, 0
		L	R
	1	A	2, 4, 2
		B	1, 1, 1
			3, 3, 3

Only  $(B, R, D)$  is perfect in  $A(\Gamma)$ , and so only  $((B, D), R)$  is extensive form perfect in  $\Gamma$ .

Why isn't  $(A, L, \cdot)$  perfect in  $A(\Gamma)$ ?  $C$  is weakly dominated for player 3, so he plays  $D$  in any perfect equilibrium. Facing  $\varepsilon C + (1 - \varepsilon)D$ , 2 prefers  $R$ . Therefore,  $L$  is not played in any perfect equilibrium.  $\diamond$

Extensive form perfect equilibrium is the original equilibrium refinement used to capture backward induction in extensive form games with imperfect information, but it is not an

easy concept to use. Kreps and Wilson (1982) introduced sequential equilibrium to retain most of the force of extensive form perfect equilibrium, but in a simpler and more intuitive way, using beliefs and sequential rationality.

The following result shows that extensive form perfect equilibrium and sequential equilibrium are nearly equivalent, with the former being a just slightly stronger refinement.

**Theorem 2.69** (Kreps and Wilson (1982), Blume and Zame (1994), Hendon et al. (1996)).

- (i) Every extensive form perfect equilibrium is a sequential equilibrium strategy profile.
- (ii) In generic extensive form games, every sequential equilibrium strategy profile is an extensive form perfect equilibrium.


In rough terms, the distinction between the concepts is as follows: Extensive form perfect equilibrium and sequential equilibrium require reasonable behavior at all information sets. But extensive form perfect equilibrium requires best responses to the perturbed strategies themselves, while sequential equilibrium only requires best responses in the limit.

To make further connections, Kreps and Wilson (1982) define *weak extensive form perfect equilibrium*, which generalizes Selten's (1975) definition by allowing slight perturbations to the game's payoffs. They show that this concept is equivalent to sequential equilibrium.

Notice that extensive form perfect equilibrium retains the problem of making different predictions in games with the same reduced normal form, since such games can have different agent normal forms (cf Example 2.61.)

Surprisingly, extensive form perfect equilibria can use weakly dominated strategies:

*Example 2.70.* In the game  $\Gamma$  below,  $A$  is weakly dominated by  $BD$ .



But  $((A, D), R)$  is extensive form perfect: there are  $\epsilon$ -perfect equilibria of the agent normal form in which agent 1b is more likely to tremble to  $C$  than player 2 is to tremble  $L$ , leading agent 1a to play  $A$ .  $\diamond$

In fact, Mertens (1995) (see also Hillas and Kohlberg (2002)) provides an example of a game in which *all* extensive form perfect equilibria use weakly dominated strategies! Again, the difficulty is that some player believes that he is more likely to tremble than his opponents.

In extensive form game  $\Gamma$ , one defines *quasi-perfect equilibrium* (van Damme (1984)) in essentially the same way as extensive-form perfect equilibrium, except that when considering player  $i$ 's best response at a given information set against perturbed strategy profiles, one only perturbs the strategies of  $i$ 's opponents; one does not perturb player  $i$ 's own choices at his other information sets. Put differently, we do not have player  $i$  consider the possibility that he himself may tremble later in the game.

Neither of extensive form perfection or quasi-perfection implies the other. But unlike extensive form perfect equilibria, quasi-perfect equilibria never employ weakly dominated strategies.

van Damme (1984) proves Theorem 2.68 by showing that proper equilibria must correspond to quasi-perfect equilibria, which in turn correspond to sequential equilibria. Mailath et al. (1997) show that proper equilibrium in a given normal form game  $G$  is equivalent to what one might call "uniform quasi-perfection" across all extensive forms with reduced normal form  $G$ .

### *Sequential rationality without equilibrium in imperfect information games*

In this section and the last, our analyses of extensive form games and their reduced normal forms have used equilibrium concepts designed to capture the principle of sequential rationality. But one can also aim to capture sequential rationality through non-equilibrium solution concepts. This requires combining the logic of rationalizability with the "undying common belief in future rational play" used to justify the backward induction solution in generic perfect information games (see Section 2.3.2.) The resulting solution concepts yield the backward induction solution in generic perfect information games, but can be applied to imperfect information games as well.

The basic solution concept obtained in this manner, *sequential rationalizability*, was suggested by Bernheim (1984) (under the name "subgame rationalizability"), defined formally by Dekel et al. (1999, 2002), and provided with epistemic foundations in the two-player case by Asheim and Perea (2005). By adding a requirement of common certainty of "cautiousness", the last paper also defines and provides epistemic foundations for *quasi-perfect rationalizability* for two-player games, which differs from sequential rationalizability only in nongeneric extensive form games.

Similar motivations lead to the notion of *proper rationalizability* for normal form games (Schuhmacher (1999), Asheim (2001), Perea (2011)). The epistemic foundations for this solution concept require common knowledge of (i) cautiousness and (ii) opponents being "infinitely more likely" to play strategies with higher payoffs. (We note that "cautiousness" allows the ruling out of weakly dominated strategies, but *not* of iteratively weakly dominated strategies, because the weakly dominated strategies are never viewed as completely impossible—see Asheim (2006, Sec. 5.3).)

As discussed in Section 2.3.2, the assumption of "undying common belief in future rational play" may be viewed as too strong in games without the single-move property. One weakening of this assumption requires that a player need only expect an opponent to choose rationally at *reachable information sets*, meaning those that are not precluded by the

opponent's own choice of strategy. (For instance, in the Mini Centipede game (Example 2.32), player 1's second decision node is not reachable if he chooses  $B$  at his first decision node.) Foundations for the resulting rationalizability concept, sometimes called *weakly sequential rationalizability*, are provided by Ben-Porath (1997); for the equilibrium analogue of this concept, sometimes called *weakly sequential equilibrium*, see Reny (1992). Adding common certainty of cautiousness yields the *permissible* strategies, which are the strategies that survive the Dekel-Fudenberg procedure; see Dekel and Fudenberg (1990), Brandenburger (1992), and Börgers (1994), as well as Section 1.2.4. The equilibrium analogue of permissibility is normal form perfect equilibrium. See Asheim (2006) for a complete treatment of these solution concepts and their epistemic foundations.


Finally, one can look at solution concepts that require correct expectations about opponents' behavior on the equilibrium path, but allow for differences in beliefs about choices off the equilibrium path. Such solution concepts, which are weaker than Nash equilibrium, include *self-confirming equilibrium* (Fudenberg and Levine (1993), Battigalli and Guaitoli (1997)) and *rationalizable self-confirming equilibrium* (Dekel et al. (1999, 2002)); the latter concept uses rationalizability requirements to restrict beliefs about opponents' play at reachable nodes off the equilibrium path.

## 2.7 Full Invariance and Kohlberg-Mertens Stability


### *Fully reduced normal forms and full invariance*

The *fully reduced normal form*  $G^*(\Gamma)$  of the extensive form game  $\Gamma$  eliminates pure strategies of  $G(\Gamma)$  that are equivalent to a mixed strategy.

*Example 2.71. Battle of the Sexes with an outside option revisited.* Here is the game  $\Gamma$  from Example 2.57, along with its reduced normal form  $G(\Gamma)$ :



Now consider  $\Gamma'$  and its reduced normal form  $G(\Gamma')$ :



But  $M$  is equivalent to  $\frac{3}{4}O + \frac{1}{4}T$ , so the fully reduced normal form is  $G^*(\Gamma') = G(\Gamma)$ .  $\diamond$

Kohlberg and Mertens (1986) propose the principle of *full invariance* (which they just call “invariance”): one should make identical predictions in games with the same fully reduced normal form.

*Example 2.72. Battle of the Sexes with an outside option rerevisited.* We saw in Example 2.57 that  $G(\Gamma)$  has three subgame perfect equilibria,  $((I, T), L)$ ,  $((O, B), R)$ , and  $((O, \frac{3}{4}T + \frac{1}{4}B), \frac{1}{4}L + \frac{3}{4}R)$ . These correspond to three proper equilibria,  $(T, L)$ ,  $(B, R)$ , and  $(\frac{3}{4}T + \frac{1}{4}B, \frac{1}{4}L + \frac{3}{4}R)$ .

$\Gamma'$  has a unique subgame perfect equilibrium,  $((I, D, T), L)$ . (In the subgame, player 1's strategy  $DB$ , which yields him at most 1, is strictly dominated by his strategy  $M$ , which yields him at least  $1\frac{1}{2}$ . Knowing this, player 2 will play  $L$ , so player 1 will play  $(D, T)$  in the subgame and  $I$  initially.) It then follows from Theorem 2.68 that  $(T, L)$  is the unique proper equilibrium of  $G(\Gamma')$ .

However, while  $\Gamma$  and  $\Gamma'$  have different purely reduced normal forms, they share the same fully reduced normal form:  $G^*(\Gamma') = G^*(\Gamma) = G(\Gamma)$ .

If one accepts full invariance as a desirable property, then this example displays a number of difficulties with proper equilibrium:

- (i) Adding or deleting a pure strategy that is equivalent to a mixture of other pure strategies can alter the set of proper equilibria. (One can show that adding or deleting duplicates of *pure* strategies does not affect proper equilibrium.)
- (ii) A proper equilibrium of a fully reduced normal form  $G^*(\Gamma')$  need not even correspond to a subgame perfect equilibrium of  $\Gamma'$ . (Contrast this with Theorem 2.68, a positive result for the *purely* reduced normal form.)
- (iii) If we find solutions for games by applying proper equilibrium to their purely reduced normal forms, then we may obtain different solutions to games with the same fully reduced normal form.  $\diamond$

Hillas (1998) suggests a different interpretation of Example 2.71: by requiring solutions to respect backward induction (sequential equilibrium) and full invariance, one can obtain forward induction for free!

*Example 2.71 rerevisited.* As we have seen,  $\Gamma$  has three subgame perfect (and sequential) equilibria. But  $\Gamma'$  has the same fully reduced normal form as  $\Gamma$ , but its only subgame perfect equilibrium is  $((I, D, T), L)$ . Therefore, our unique prediction of play in  $\Gamma$  should be the corresponding subgame perfect equilibrium  $((I, T), L)$ . As we saw earlier, this is the only equilibrium of  $\Gamma$  that respects forward induction.  $\diamond$

Building on this insight, Govindan and Wilson (2009) argue that together, backward induction and full invariance imply forward induction, at least in generic two-player games.

### ***KM stability and set-valued solution concepts***

With the foregoing examples as motivation, Kohlberg and Mertens (1986) list desirable properties (or “desiderata”) for refinements of Nash equilibrium.

- (D1) *Full invariance:* Solutions to games with the same fully reduced normal form are identical.
- (D2) *Backward induction:* The solution contains a sequential equilibrium.
- (D3) *Iterated dominance:* The solution to  $G$  contains a solution to  $G'$ , where  $G'$  is obtained from  $G$  by removing a weakly dominated strategy.
- (D4) *Admissibility:* Solutions do not include weakly dominated strategies.

Iterated dominance (D3) embodies a limited form of forward induction—see Section 2.5.1. KM argue that admissibility (D4) is a basic decision-theoretic postulate that should be respected, and appeal to various authorities (Wald, Arrow, . . .) in support of this point of view.

In addition, KM require *existence*: a solution concept should offer at least one solution for every game.

For a solution concept to satisfy invariance (D1), backward induction (D2), and existence in all games, the solutions must be *set-valued*: see Example 2.73 below.

Similarly, set-valued solutions are required for the solution concept to satisfy (D1), (D3), and existence (see KM, Section 2.7.B).

Set-valued solutions are natural: Extensive form games possess connected components of Nash equilibria, elements of which differ only in terms of behavior at unreached information sets. Each such component should be considered as a unit.

Once one moves to set-valued solutions, one must consider restrictions on the structure of solution sets. KM argue that solution sets should be *connected sets*.

As build-up, KM introduce two set-valued solution concepts that satisfy (D1)–(D3) and existence, but that fail admissibility (D4) and connectedness.

They then introduce their preferred solution concept: A closed set  $E$  of Nash equilibria (of game  $G = G^*(\Gamma)$ ) is *KM stable* if it is minimal with respect to the following property:

“for any  $\varepsilon > 0$  there exists some  $\delta_0 > 0$  such that for any completely mixed strategy vector  $(\sigma_1, \dots, \sigma_n)$  and for any  $\delta_1, \dots, \delta_n$  ( $0 < \delta_i < \delta_0$ ), the perturbed game where every strategy  $s$



of player  $i$  is replaced by  $(1 - \delta_i)s + \delta_i\sigma_i$  has an equilibrium  $\varepsilon$ -close to  $E$ ."


Remark: If in the above one replaces "for any  $(\sigma_1, \dots, \sigma_n)$  and  $\delta_1, \dots, \delta_n$ " with "for some  $(\sigma_1, \dots, \sigma_n)$  and  $\delta_1, \dots, \delta_n$ ", the resulting requirement is equivalent to perfect equilibrium. Thus, a key novelty in the definition of KM stability is the requirement that equilibria be robust to *all* sequences of perturbations.

KM stability satisfies (D1), (D3), (D4), and existence. In fact, it even satisfies a stronger forward induction requirement than (D3) called *equilibrium dominance*: A KM stable set  $E$  contains a KM stable set of any game obtained by deletion of a strategy that is not a best response to any equilibrium in  $E$  (see Section 2.5.2 for further discussion).

However, KM stability fails connectedness and backward induction (D2): KM provide examples in which a KM stable set (of  $G^*(\Gamma)$ ) contains no strategy profile corresponding to a sequential equilibrium (of  $\Gamma$ ).

A variety of other definitions of strategic stability have been proposed since Kohlberg and Mertens (1986). Mertens (1989, 1991) proposes a definition of strategic stability that satisfies (D1)–(D4), existence, connectedness, and much besides, but that is couched in terms of ideas from algebraic topology. Govindan and Wilson (2006) obtain (D1)–(D4) and existence (but not connectedness) using a relatively basic definition of strategic stability.

*Example 2.73.* Why backward induction and full invariance require a set-valued solution concept.



Analysis of  $G$ : By drawing the payoffs to each of player 1's pure strategies as a function of player 2's mixed strategy, one can see that player 1's unique maxmin strategy is  $O$ , and that player 2's maxmin strategies are  $\alpha L + (1 - \alpha)R$  with  $\alpha \in [\frac{1}{4}, \frac{3}{4}]$ . Since  $G$  is zero-sum, its Nash equilibria are thus the profiles  $(O, \alpha L + (1 - \alpha)R)$  with  $\alpha \in [\frac{1}{4}, \frac{3}{4}]$ .

Analysis of  $\Gamma(p)$ : For each  $p \in (0, 1)$ , this game has a unique sequential equilibrium, namely  $((O, \frac{1}{2-p}M + \frac{1-p}{2-p}B), \frac{4-3p}{8-4p}L + \frac{4-p}{8-4p}R)$ . To see this, first notice that there cannot be an equilibrium

in which player 2 plays a pure strategy. (For instance, if 2 plays  $L$ , then 1's best response in the subgame would be  $T$ , in which case 2 would switch to  $R$ .) Thus, player 2 is mixing. For this to be optimal, her beliefs must satisfy  $\mu(z) = \frac{1}{2}$ . To have an equilibrium in the subgame in which player 2 has these beliefs, player 1's strategy must satisfy

$$(*) \quad \sigma_1(B) = (1 - p)\sigma_1(M) + \sigma_1(T).$$

In particular, player 1 must place positive probability on  $B$  and on at least one of  $M$  and  $T$ . For player 1 to place positive probability on  $T$ , he would have to be indifferent between  $B$  and  $T$ , implying that 2 plays  $\frac{1}{2}L + \frac{1}{2}R$ . But in this case player 1 would be strictly better off playing  $M$  in the subgame, a contradiction. Thus  $\sigma_1(T) = 0$ , and so (\*) implies that player 1 plays  $\frac{1}{2-p}M + \frac{1-p}{2-p}B$  in the subgame.

For player 1 to be willing to randomize between  $M$  and  $B$ , it must be that

$$p + (1 - p)(2\sigma_2(L) - 2(1 - \sigma_2(L))) = -2\sigma_2(L) + 2(1 - \sigma_2(L)),$$

implying that  $\sigma_2(L) = \frac{4-3p}{8-4p}$  and that  $\sigma_2(R) = \frac{4-p}{8-4p}$ .

Finally, with these strategies chosen in the subgame, player 1's expected payoff from choosing  $I$  at his initial node is

$$-2\sigma_2(L) + 2\sigma_2(R) = \frac{4p}{8-4p} = \frac{p}{2-p}.$$

Since  $p < 1$ , this payoff is less than 1, and so player 1 strictly prefers  $O$  at his initial node.

Each choice of  $p \in (0, 1)$  leads to a unique and distinct sequential equilibrium of  $\Gamma(p)$ . These equilibria correspond to distinct Nash equilibria of  $G$ , which itself is the reduced normal form of each  $\Gamma(p)$ . Therefore, if we accept backward induction and invariance, no one Nash equilibrium of  $G$  constitutes an acceptable prediction of play. Thus, requiring invariance and backward induction leads us to set-valued solution concepts.

(If in  $\Gamma(p)$  we had made the strategy  $M$  a randomization between  $O$  and  $B$ , the weight on  $\sigma_2(L)$  would have gone from  $\frac{3}{4}$  to  $\frac{1}{2}$ , giving us the other half of the component of equilibria. This does not give us  $\sigma_2(L) = \frac{1}{2}$ , but this is the unique subgame perfect equilibrium of the game where 1 does not have strategy  $M$ .)  $\diamond$

### 3. Repeated Games

In many applications, players face the same interaction repeatedly. How does this affect our predictions of play?

Repeated games provide a general framework for studying long run relationships.

While we will focus on the basic theory, this subject becomes even more interesting when

one introduces hidden information ( $\Rightarrow$  reputation models (Kreps et al. (1982), Schmidt (1993))), hidden actions ( $\Rightarrow$  imperfect monitoring models (Green and Porter (1984), Abreu et al. (1990), Fudenberg et al. (1994), Sekiguchi (1997), Sannikov (2007))), state variables (Dutta (1995), Maskin and Tirole (2001)), or combinations of these elements. An excellent general reference for this material is Mailath and Samuelson (2006).

### 3.1 Basic Concepts

#### *Definitions*

In an (*infinitely*) *repeated game*, an  $n$  player normal form game  $G$ , known as the *stage game*, is played in periods  $t = 0, 1, 2, \dots$ , with the period  $t$  action choices being commonly observed at the end of the period. Payoffs in the repeated game are the discounted sum of stage game payoffs for some fixed discount rate  $\delta \in (0, 1)$ .

The stage game  $G = \{\mathcal{P}, \{A_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$

$a_i \in A_i$	a pure action
$\alpha_i \in \Delta A_i$	a mixed action
$\alpha \in \prod_{i \in \mathcal{P}} \Delta A_i$	a mixed action profile

The repeated game  $G^\infty(\delta) = \{\mathcal{P}, \{S_i\}_{i \in \mathcal{P}}, \{\pi_i\}_{i \in \mathcal{P}}, \delta\}$

$H^0 = \{h^0\}$	the null history
$H^t = \{(a^0, a^1, \dots, a^{t-1}) : a^s \in A\}$	histories as of periods $t \geq 1$
$H = \bigcup_{t=0}^{\infty} H^t$	finite histories (countable)
$S_i = \{s_i : H \rightarrow A_i\}$	pure strategies
$\Sigma_i = \{\sigma_i : H \rightarrow \Delta A_i\}$	behavior strategies
$\delta \in (0, 1)$	discount rate
$H^\infty = \{(a^0, a^1, \dots) : a^t \in A\}$	infinite histories (uncountable)
$\pi_i : H^\infty \rightarrow \mathbb{R}$	payoff function
$\pi_i(h^\infty) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t)$	

The discount rate  $\delta$  can be interpreted as the probability that the game does not end in any given period. But we don't need to be too literal about infinite repetition: what is important is that the players view the interaction as one with no clear end—see Rubinstein (1991) for a discussion.

The rescaling by  $(1 - \delta)$  is done for convenience: it makes payoffs in the repeated game commensurate with payoffs in the stage game. To see why, recall that  $\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}$ ; thus  $(1 - \delta) \sum_{t=0}^{\infty} \delta^t c = c$ , as desired.

- $\Rightarrow$  In  $G^\infty$ ,  $(c, c, c, \dots)$  is worth  $c$  (as suggested above);  
 $(0, c, c, \dots)$  is worth  $\delta c$ ;  
 $(c, 0, 0, 0, \dots)$  is worth  $(1 - \delta)c$ .

### Subgame perfect equilibrium

Strategy profile  $\sigma$  is a *Nash equilibrium* of  $G^\infty(\delta)$  if no player has a profitable unilateral deviation from  $\sigma$ .

Let  $\sigma|_{h^t}$  denote the *continuation strategy profile* generated by  $\sigma$  after history  $h^t$ . Strategy profile  $\sigma$  is a *subgame perfect equilibrium* of  $G^\infty(\delta)$  if for every history  $h^t \in H$ ,  $\sigma|_{h^t}$  is a Nash equilibrium of  $G^\infty(\delta)$ .

As we noted in Section 2.3.4, the one-shot deviation principle applies in repeated games, so that subgame perfection is equivalent to the absence of profitable one-shot deviations. Formally, we say that strategy  $\sigma_i$  *admits no profitable one-shot deviations* given  $\sigma_{-i}$  if after each history  $h^t$ , player  $i$  cannot improve his payoff in the continuation game that follows  $h^t$  by changing his action immediately after  $h^t$  but otherwise following strategy  $(\sigma|_{h^t})_i$ . If this is true for every player  $i$ , we say that strategy profile  $\sigma$  itself *admits no profitable one-shot deviations*.

**Theorem 3.1.** *Let  $G^\infty(\delta)$  be a repeated game. Strategy profile  $\sigma$  is a subgame perfect equilibrium of  $G^\infty(\delta)$  if and only if  $\sigma$  admits no profitable one-shot deviations.*

This result is our basic tool for analyzing repeated games.

### The repeated Prisoner's Dilemma

Consider the following one-shot Prisoner's Dilemma  $G$ :

		2	
		C	D
1	C	1, 1	-1, 2
	D	2, -1	0, 0

If  $G$  is played once, then each player has a dominant strategy of defecting.

*Example 3.2. The finitely repeated Prisoner's Dilemma  $G^T(\delta)$ ,  $T < \infty$ ,  $\delta \in (0, 1]$ .*

Players play  $T + 1$  times, starting with period 0 (so that the last period is  $T$ .) Before playing period  $t$ , the results of all previous periods are observed. Payoffs are the discounted sum of stage game payoffs

**Proposition 3.3.** *In the unique subgame perfect equilibrium of  $G^T(\delta)$ , both players always Defect.*

*Proof.* By backward induction:

Once the final period  $T$  is reached,  $D$  is dominant for both players, and so is played regardless of the previous history.

Therefore, choices in period  $T - 1$  cannot influence payoffs in period  $T$ .

Hence, backward induction implies that both players defect in period  $T - 1$ .

Repeat through period 0. ■

**Proposition 3.4.** *In any Nash equilibrium of  $G^T(\delta)$ , players always defect on the equilibrium path.*

*Proof.* Fix a Nash equilibrium  $\sigma^*$ .

Clearly, both players play  $D$  in period  $T$  at any node which is reached with positive probability under  $\sigma^*$ .

Therefore, since it cannot affect behavior in period  $T$ , playing  $D$  is the unique best response at any positive probability period  $T - 1$  history, so since  $\sigma^*$  is a Nash equilibrium the players must do so.

Repeat through period 0. ■

*Example 3.5. The infinitely repeated Prisoner's Dilemma  $G^\infty(\delta)$ ,  $\delta \in (0, 1)$ .*

What are the consequences of dropping the assumption that the game has a commonly known final period?

**Proposition 3.6.**

- (i) *"Always defect" is a subgame perfect equilibrium of  $G^\infty(\delta)$  for all  $\delta \in (0, 1)$ .*
- (ii) *If  $\delta \geq 1/2$ , the following defines a subgame perfect equilibrium of  $G^\infty(\delta)$ :  $\sigma_i =$  "Cooperate so long as no one has ever defected; otherwise defect," (the grim trigger strategy).*

There are many other equilibria—see Sections 3.2 and 3.3.

We determine whether a strategy profile is a subgame perfect equilibrium of  $G^\infty(\delta)$  using the one-shot deviation principle (Theorem 3.1). Specifically, it is enough to check that no

player cannot benefit from deviating only in the period immediately following each finite history  $h^t \in H$ . We do so by partitioning the finite histories into a small number of cases according to the nature of continuation play.

*Proof.* We begin by considering part (ii). Since the equilibrium is symmetric, we need only check player 1's behavior. There are two sorts of histories to consider:

Case 1: No one has defected.

We argue in terms of *continuation payoffs*: without loss of generality, we can assume that the deviation occurs in period zero.

$$\text{Equilibrium: } (C, C)_{t=0}, (C, C)_{t=1}, (C, C)_{t=2}, \dots \Rightarrow \pi_1 = 1$$

$$\text{Deviation: } (D, C)_{t=0}, (D, D)_{t=1}, (D, D)_{t=2}, \dots \Rightarrow \pi_1 = (1 - \delta) \left( 2 + \sum_{t=1}^{\infty} \delta^t \cdot 0 \right) = (1 - \delta)2$$

Therefore, the equilibrium behavior is optimal if  $1 \geq (1 - \delta)2$ , or  $\delta \geq 1/2$ .

Case 2: Someone has defected.

$$\text{Equilibrium: } (D, D)_{t=0}, (D, D)_{t=1}, \dots \Rightarrow 0$$


$$\text{Deviation: } (C, D)_{t=0}, (D, D)_{t=1}, (D, D)_{t=2}, \dots \Rightarrow (1 - \delta) \left( -1 + \sum_{t=1}^{\infty} \delta^t \cdot 0 \right) = \delta - 1 < 0$$

Therefore, the equilibrium behavior is strictly better in this case.

Note: The analysis of Case 2 proves part (i) of the proposition. ■

### Remarks

1. The analysis of infinitely repeated games is greatly simplified by their *recursive* structure: the "continuation game" starting from any history  $h^t$  is formally identical to the game starting at the null history  $h^0$ .  
(Dynamic programs have a related (but distinct) recursive structure.)
2. Since the strategies are *contingent* rules for behavior, a change in player  $i$ 's *strategy* in one period can change the *actions* he and his opponents choose in future periods. For instance, if both players are supposed to play the grim trigger strategy, then changing player 1's strategy in the initial period changes the play path from  $((C, C), (C, C), (C, C), \dots)$  to  $((D, C), (D, D), (D, D), \dots)$ .



The (partial) game tree above is that of the repeated Prisoner's Dilemma. The bold edges represent the grim trigger strategy profile. If this strategy profile is played, then play proceeds down the leftmost branch of the tree, and both players cooperate in every period. If we modify one player's strategy so that after a single cooperative history he plays *D* rather than *C*, then play enters a subgame in which both players defect in every period. Thus a deviation from the strategy's prescription after a single history can alter what actions are played in all subsequent periods.

This figure also shows the partition of histories into the two cases above. The decision nodes of player 1 are the initial nodes of subgames. The leftmost subgames follow histories in which no one has defected (including the null history). All other subgames follow histories in which someone has defected.

3. The difference between the equilibrium outcomes of the finitely repeated and infinitely repeated Prisoner's Dilemmas is quite stark. With many other stage games, the difference is not so stark. If the stage game *G* has multiple Nash equilibrium outcomes, one can often sustain the play of non-Nash action profiles of *G* in early periods—for instance, by rewarding cooperative play in early periods with the play of good Nash outcomes of *G* in later periods, and by punishing deviations in early periods with bad Nash outcomes in later periods. For general analyses of the finitely repeated games, see Benoît and Krishna (1985) and Friedman (1985).

## 3.2 Stick-and-Carrot Strategies and the Folk Theorem

Repeated games typically have many subgame perfect equilibria. Therefore, rather than looking for all subgame perfect equilibrium strategy profiles, we instead ask which payoff vectors can be achieved in a subgame perfect equilibrium.

The most basic question is to characterize the set of payoff vectors that can be achieved in a repeated game  $G^\infty(\delta)$  for some fixed discount rate  $\delta \in (0, 1)$ . The solution to this question, due to Abreu et al. (1990), is studied in Section 3.3.

Before addressing this question, we consider one that turns out to be simpler, at least in its basic formulation: What can we say about the set of subgame perfect equilibrium payoff vectors when players are very patient—that is, when the discount rate  $\delta$  approaches 1? The *folk theorem* tells us that all feasible, strictly individually rational payoff vectors can be obtained in subgame perfect equilibrium. Thus in repeated games with very patient players, subgame perfection imposes no restrictions on our predictions of payoffs beyond those evident from the stage game. In proving the folk theorem, we will introduce *stick-and-carrot strategies*, a fundamental device for the provision of intertemporal incentives.

We first define the set of feasible, individually rational payoff vectors.

Let  $V = \{u(a) \in \mathbb{R}^n : a \in A\}$  be the payoff vectors generated by pure action profiles in the stage game.

Then define the set of feasible repeated game payoffs as  $F = \text{conv}(V)$ , the convex hull of  $V$ . By alternating among stage game action profiles over time (and for the moment ignoring questions of incentives), players can obtain any payoff in  $F$  as their repeated game payoff if the discount rate is sufficiently high—see discussion point 2 at the end of the section.

To introduce individual rationality, define player  $i$ 's *minmax value* by

$$\underline{v}_i = \min_{\alpha_{-i} \in \prod_{j \neq i} \Delta A_j} \max_{\alpha_i \in \Delta A_i} u_i(\alpha_i, \alpha_{-i}).$$

Thus  $\underline{v}_i$  is the payoff obtained when his opponents minmax him and he, anticipating what they will do, plays a best response (see Section 1.6). This leads to a lower bound on what player  $i$  obtains in any equilibrium of the repeated game.

**Observation 3.7.** *Player  $i$  obtains at least  $\underline{v}_i$  in any Nash equilibrium of  $G^\infty$ .*

*Proof.* Given any  $\sigma_{-i}$ ,  $i$  can always play his myopic best response in each period, and this will yield him a payoff of at least  $\underline{v}_i$ . ■



The set of feasible, strictly individually rational payoff vectors is

$$F^* = \{v \in F: v_i > \underline{v}_i \text{ for all } i \in \mathcal{P}\}.$$

Note the requirement of strict individual rationality:  $v_i = \underline{v}_i$  sometimes is attainable as an equilibrium payoff (for instance, in the repeated Prisoner's Dilemma), but often it is not. See Mailath and Samuelson (2006, Remark 3.3.1) for details.

**Theorem 3.8** (The folk theorem (Fudenberg and Maskin (1986, 1991), Abreu et al. (1994))).

Let  $v \in F^*$ , and suppose that:

- (i) there are exactly two players; or
- (ii) no two players have identical preferences.  
(i.e., there do not exist  $i, j \in \mathcal{P}$ ,  $a, b > 0$  such that  $u_i(\cdot) \equiv a + bu_j(\cdot)$ ).

Then for all  $\delta$  close enough to 1, there is a subgame perfect equilibrium of  $G^\infty(\delta)$  with payoffs  $v$ .

Condition (ii) is known as *nonequivalent utilities*, or *NEU*.

Remark: Why should we care about being able to sustain low subgame perfect equilibrium payoffs? Consider a stage game  $G$  with an action profile  $a$  that gives all players high payoffs, but is not a Nash equilibrium of  $G$ : there is a player  $i$  with an alternate action  $\tilde{a}_i$  such that action profile  $(\tilde{a}_i, a_{-i})$  gives player  $i$  a very high stage game payoff at the expense of the others. To support repeated play of  $a$  as a subgame perfect equilibrium play path of  $G^\infty(\delta)$ , we must be able to threaten to punish player  $i$  with low payoffs starting tomorrow if he plays  $\tilde{a}_i$  today. But to ensure that such a punishment is credible, the punishment strategy profile itself must be a subgame perfect equilibrium of  $G^\infty(\delta)$ . Thus the ability to enforce low subgame perfect equilibrium payoffs can be crucial to achieving high subgame perfect equilibrium payoffs.

This idea is applied in constructing the stick and carrot strategies used to prove the folk theorem (see Example 3.11). It becomes more prominent still when we consider supporting high equilibrium payoffs in  $G^\infty(\delta)$  for some fixed discount rate  $\delta \in (0, 1)$  (see Example 3.12 and Section 3.3).

The next two examples prove special cases of the folk theorem in order to illustrate two of the basic constructions of subgame perfect equilibria of repeated games. Later we discuss what more must be done to prove the full result.

*Example 3.9. Nash reversion* (Friedman (1971)).

We first prove the folk theorem under the assumptions that

- (1) The payoff vector  $v = u(\hat{a})$  can be obtained from some pure action profile  $\hat{a} \in A$ .
- (2) For each player  $i$ , there is a Nash equilibrium  $\alpha^i \in A$  of the stage game such that  $v_i > u_i(\alpha^i)$ .

Consider this *Nash reversion strategy*, a generalization of the grim trigger strategy:

$\sigma_i$ : "Play  $\hat{a}_i$  if there has never been a period in which exactly one player deviated. Otherwise, if  $j$  was the first to unilaterally deviate, play  $\alpha_i^j$ ."

Clearly,  $\sigma$  generates payoff vector  $v$  in  $G^\infty(\delta)$  for any  $\delta \in (0, 1)$ . To verify that  $\sigma$  is a subgame perfect equilibrium of for  $\delta$  large enough, we check that no player has a profitable one-shot deviation.

Let  $\bar{v}_i = \max_{a \in A} u_i(a)$  denote player  $i$ 's maximal stage game payoff.

There are  $1 + n$  cases to consider. After histories with no unilateral deviations (e.g., on the equilibrium path),  $i$  does not benefit from deviating if

$$v_i \geq (1 - \delta)\bar{v}_i + \delta u_i(\alpha^i).$$


Since  $v_i > u_i(\alpha^i)$ , this is true if  $\delta$  is large enough.

After a history in which player  $i$  was the first to unilaterally deviate, the continuation strategy profile is "always play the Nash equilibrium  $\alpha^i$ ". Clearly, no player has a profitable one-shot deviation here.  $\diamond$

*Example 3.10.* To better understand what can and cannot be achieved using Nash reversion, consider this symmetric normal form game  $G$ :

		2		
		$a$	$b$	$c$
1	$A$	0,0	4,2	1,0
	$B$	2,4	0,0	0,0
	$C$	0,1	0,0	0,0

The Nash equilibria of  $G$  are  $(A, b)$ ,  $(B, a)$ , and  $(\frac{2}{3}A + \frac{1}{3}B, \frac{2}{3}a + \frac{1}{3}b)$ , which yield payoffs  $(4, 2)$ ,  $(2, 4)$ , and  $(\frac{4}{3}, \frac{4}{3})$ . The strategies  $C$  and  $c$  are strictly dominated, but are the pure minmax strategies; they generate the pure minmax payoffs  $(\underline{v}_1^p, \underline{v}_2^p) = (1, 1)$ . The mixed minmax strategies are  $\frac{1}{3}A + \frac{2}{3}C$  and  $\frac{1}{3}a + \frac{2}{3}c$ ; they generate the mixed minmax payoffs  $(\underline{v}_1, \underline{v}_2) = (\frac{2}{3}, \frac{2}{3})$ . (All of these claims can be verified using the figure at left.)



The five black dots in the figure at right correspond to the feasible payoffs from pure strategy profiles of  $G$ . The convex hull of these points is the set  $F$  of feasible repeated game payoffs.

The light gray region consists of feasible payoff vectors in which each player receives more than his payoff in the mixed Nash equilibrium of  $G$ . In principle, these payoff vectors can be obtained by patient players in subgame perfect equilibria of the repeated game using Nash reversion. (Getting points in the interior of this region requires either alternation or randomization on the equilibrium path; see discussion point 1 below.)

The union of the light and medium gray regions consists of payoff vectors that give each player at least his pure minmax payoff. Again modulo alternation or randomization on the equilibrium path, these points can be achieved in subgame perfect equilibrium using the approach introduced in the next example.

The complete shaded region is the set  $F^*$  considered in the folk theorem. To obtain points in the dark gray region, one needs to introduce randomization not only on the equilibrium path, but also on the punishment path; see discussion point 2 below.  $\diamond$

*Example 3.11. Stick and carrot strategies* (Fudenberg and Maskin (1986)).

We now prove the folk theorem under the assumptions that

- (1) The payoff vector  $v = u(\hat{a})$  can be obtained from some pure action profile  $\hat{a} \in A$ .

- (2) There are exactly two players.  
(3) For each player  $i$ ,  $v_i$  is greater than player  $i$ 's pure strategy minmax value,

$$\underline{v}_i^p = \min_{a_j \in A_j} \max_{a_i \in A_i} u_i(a_i, a_j).$$

Let  $a_i^m$  be player  $i$ 's minmaxing pure action, and consider the following *stick and carrot strategy*:

- $\sigma_i$ : "(I) Play  $\hat{a}_i$  initially, or if  $\hat{a}$  was played last period.  
(II) If there is a deviation from (I), play  $a_i^m$   $L$  times and then restart (I).  
(III) If there is a deviation from (II), begin (II) again."

The value of the *punishment length*  $L \geq 1$  will be determined below; often  $L = 1$  is enough.

Again,  $\sigma$  generates payoff vector  $v$  in  $G^\infty(\delta)$  for any  $\delta \in (0, 1)$ . To verify that  $\sigma$  is a subgame perfect equilibrium of for  $\delta$  large enough, we check that no player has a profitable one-shot deviation.

Let  $\bar{v}_i = \max_a u_i(a)$ , and let  $v_i^m = u_i(a^m)$ , where  $a^m = (a_1^m, a_2^m)$  is the joint minmaxing pure action profile. Then

$$(18) \quad v_i^m \leq \underline{v}_i^p < v_i \leq \bar{v}_i.$$

We can therefore choose a positive integer  $L$  such that for  $i \in \{1, 2\}$ ,

$$(19) \quad L(v_i - v_i^m) > \bar{v}_i - v_i, \quad \text{or equivalently,} \\
(L + 1)v_i > \bar{v}_i + Lv_i^m.$$

(In words: if player  $i$  were perfectly patient, he would prefer getting  $v_i$  for  $L + 1$  periods to getting his maximum payoff  $\bar{v}_i$  once followed by his joint minmax payoff  $v_i^m$   $L$  times.)

There is no profitable one-shot deviation from the equilibrium phase (I) if

$$v_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t v_i \geq (1 - \delta) \left( \bar{v}_i + \sum_{t=1}^L \delta^t v_i^m + \sum_{t=L+1}^{\infty} \delta^t v_i \right) \\
\Leftrightarrow \sum_{t=0}^L \delta^t v_i \geq \bar{v}_i + \sum_{t=1}^L \delta^t v_i^m.$$

Equation (19) implies that this inequality holds when  $\delta$  is close enough to 1.

In the punishment phase (II), deviating is most tempting in the initial period, when all  $L$

rounds of punishment still remain. Deviating in this period is not profitable if

$$\begin{aligned}
 (1 - \delta) \left( \sum_{t=0}^{L-1} \delta^t v_i^m + \sum_{t=L}^{\infty} \delta^t v_i \right) &\geq (1 - \delta) \left( \underline{v}_i^p + \sum_{t=1}^L \delta^t v_i^m + \sum_{t=L+1}^{\infty} \delta^t v_i \right) \\
 \Leftrightarrow v_i^m + \delta^L v_i &\geq \underline{v}_i^p + \delta^L v_i^m \\
 \Leftrightarrow \delta^L v_i + (1 - \delta^L) v_i^m &\geq \underline{v}_i^p.
 \end{aligned}$$

Equation (18) implies that this inequality holds when  $\delta$  is close enough to 1.

The intuition for why the punishment phase is an equilibrium is easiest to see in the penultimate inequality. If  $\delta = 1$ , the left-hand side is the sum of the joint minmax payoff  $v_i^m$  and the equilibrium payoff  $v_i$ ; the right-hand side is the sum of the minmax payoff  $\underline{v}_i^p$  and the joint minmax payoff. The joint minmax payoffs cancel, and we are left with the comparison  $v_i \geq \underline{v}_i^p$ , which holds strictly by assumption. Put differently: When players are very patient, the salient difference between obeying the punishment and deviating from it is one round of the equilibrium payoff in the former case vs. one round of the minmax payoff in the latter. The player strictly prefers the former.

Why is  $\sigma_i$  called a stick and carrot strategy? The punishment phase (II) is the stick (i.e., the threat) that keeps players from deviating from the equilibrium path phase (I). The equilibrium path phase (I) is the carrot (i.e., the reward) offered to players for carrying out the punishment phase (II).  $\diamond$

*Example 3.12. Stick-and-carrot strategies.*

Let  $G^\infty(\delta)$  be the infinite repetition of the following normal form game  $G$ :

		2		
		A	B	C
1	A	4, 4	0, 5	1, 0
	B	5, 0	3, 3	0, 0
	C	0, 1	0, 0	$-\frac{1}{2}, -\frac{1}{2}$

- (i) Suppose that only actions  $A$  and  $B$  may be played. For what values of  $\delta$  can the payoff vector  $(4, 4)$  be sustained as a subgame perfect equilibrium payoff? Why?
- (ii) Now suppose that all three actions may be played. For what values of  $\delta$  can the payoff vector  $(4, 4)$  be sustained using a stick-and-carrot strategy with a one-period punishment? Why?
- (iii) Would increasing the length of the punishment allow  $(4, 4)$  to be sustained for lower values of  $\delta$ ? Explain why or why not.

Solution:

- (i) If only strategies  $A$  and  $B$  are available, then  $G^\infty(\delta)$  is a repeated Prisoner's Dilemma. The strongest punishment that can be applied is to use the grim trigger strategy. If no one has deviated, then following this strategy is optimal for both players if

$$4 \geq 5(1 - \delta) + 3\delta$$

$$\Leftrightarrow \delta \geq \frac{1}{2}.$$

If someone has deviated, then following the strategy is clearly optimal for any  $\delta$ . Thus payoff  $(4, 4)$  can be sustained if  $\delta \geq \frac{1}{2}$ .

- (ii) Consider the stick-and-carrot strategy whose equilibrium path is  $((A, A), (A, A), \dots)$  and whose punishment path is  $((C, C), (A, A), (A, A), \dots)$ . It is optimal to follow the strategy on the equilibrium path if

$$4 \geq 5(1 - \delta) - \frac{1}{2}\delta(1 - \delta) + 4\delta^2$$

$$\Leftrightarrow 0 \geq 9\delta^2 - 11\delta + 2$$

$$\Leftrightarrow \delta \geq \frac{2}{9}.$$

It is optimal to follow the strategy on the punishment path when

$$-\frac{1}{2}(1 - \delta) + 4\delta \geq 1(1 - \delta) - \frac{1}{2}\delta(1 - \delta) + 4\delta^2$$

$$\Leftrightarrow 0 \geq 9\delta^2 - 12\delta + 3$$

$$\Leftrightarrow \delta \geq \frac{1}{3}.$$

Since  $\frac{1}{3} > \frac{2}{9}$ , payoff  $(4, 4)$  can be sustained if  $\delta \geq \frac{1}{3}$ .

- (iii) For fixed  $\delta$ , increasing the punishment length strengthens the incentive to stay on the equilibrium path, but weakens the incentive to go through with the punishment path. Thus if we increased the punishment length, then the equilibrium path constraint would bind at a smaller value of  $\delta$  than originally, while the punishment path constraint would bind at a higher value of  $\delta$  than previously. Since the latter was originally the active constraint, the smallest  $\delta$  which sustains equilibrium will become larger. In other words, longer punishments are counterproductive.

### *Proving the folk theorem: discussion*

1. How can one support payoff vectors that do not correspond to pure action profiles?

If we do not modify the repeated game, payoff vectors that do not correspond to pure strategy profiles can only be achieved if players alternate among pure action profiles over time. Sorin (1986) shows that any feasible payoff vector can be achieved through

alternation if players are sufficiently patient. However, the need to alternate complicates the construction of equilibrium, since we must ensure that no player has a profitable deviation at any point during the alternation. Fudenberg and Maskin (1991) show that this can be accomplished so long as players are sufficiently patient.

To avoid alternation and the complications it brings, it is common to augment the repeated game by introducing *public randomization*: at the beginning of each period, all players view the realization of a *uniform*(0, 1) random variable, enabling them to play a correlated action in every period. If on the equilibrium path players always play a correlated action whose expected payoff is  $v$ , then each player's continuation payoff is always exactly  $v$ . Since in addition the benefit obtained in a the current period from a one shot deviation is bounded (by  $\max_a u_i(a) - \min_a u_i(a)$ ), the equilibrium constructions and analyses from Examples 3.9 and 3.11 go through with very minor changes.

It is natural to ask whether public randomization introduces equilibrium outcomes that would otherwise be impossible. Since the folk theorem holds without public randomization, we know that for each payoff vector  $v \in F^*$ , there is a  $\underline{\delta}(v)$  such that  $v$  can be achieved in a subgame perfect equilibrium of  $G^\infty(\delta)$  whenever  $\delta > \underline{\delta}(v)$ . Furthermore, Fudenberg et al. (1994) show that any given convex, compact set in the interior of  $F^*$  contains only subgame perfect equilibrium payoff vectors of  $G^\infty(\delta)$  once  $\delta$  is large enough. However, Yamamoto (2010) constructs an example in which the set of subgame perfect equilibrium payoff vectors of  $G^\infty(\delta)$  is not convex (and in particular excludes certain points just inside the Pareto frontier) for any  $\delta < 1$ ; thus, allowing public randomization is not entirely without loss of generality even for discount factors arbitrarily close to 1.

## 2. How can one obtain payoffs close to the players' mixed minmax values?

The the stick and carrot equilibrium from Example 3.11 relied on pure minmax actions as punishments. As Example 3.10 illustrates, such punishments are not always strong enough to sustain all vectors in  $F^*$  as subgame perfect equilibrium payoffs.

The difficulty with using mixed action punishments is that when a player chooses a mixed action, his opponents cannot observe his randomization probabilities, but only his realized pure action. Suppose we modified the stick and carrot strategy profile from Example 3.11 by specifying that player  $i$  play his mixed minmax action  $\alpha_i$  in the punishment phase. Then during this punishment phase, unless player  $i$  expects to get the same stage game payoff from each action in the support of  $\alpha_i$ , he will have a profitable and undetectable deviation from  $\alpha_i$  to his favorite action in the support of  $\alpha_i$ .

To address this problem, we need to modify the repeated game strategies so that when player  $i$  plays one of his less preferred actions in the support of  $\alpha_i$ , he is rewarded with a higher continuation payoff. More precisely, by carefully balancing player  $i$ 's stage game payoffs and continuation payoffs, one can make player  $i$  indifferent among playing any of the actions in the support of  $\alpha_i$ , and therefore willing to randomize among them in the way his mixed minmax action specifies. See Mailath and Samuelson (2006, Sec. 3.8) for a textbook treatment.

## 3. What new issues arise when there are three or more players?

With two players, we can always find an action profile  $(\alpha_1^m, \alpha_2^m)$  in which players 1 and 2 simultaneously minmax each other. This possibility was the basis for the stick and carrot equilibrium from Example 3.11.

Once there are three players, simultaneous minmaxing may no longer be possible: if  $(\alpha_1, \alpha_2)$  minmaxes player 3, there may be no  $\alpha_3$  such that  $(\alpha_1, \alpha_3)$  minmaxes player 2. In such cases, there is no analogue of the two-player stick and carrot equilibrium, and indeed, it is not always possible to achieve payoffs close to the players' minmax values in a subgame perfect equilibrium.

*Example 3.13.* Fudenberg and Maskin (1986). Consider this three-player game  $G$ :

		2	
		A <sub>2</sub>	B <sub>2</sub>
1	A <sub>1</sub>	1, 1, 1	0, 0, 0
	B <sub>1</sub>	0, 0, 0	0, 0, 0

		2	
		A <sub>2</sub>	B <sub>2</sub>
1	A <sub>1</sub>	0, 0, 0	0, 0, 0
	B <sub>1</sub>	0, 0, 0	1, 1, 1

3: A<sub>3</sub>
3: B<sub>3</sub>

Each player's minmax value is 0, so  $F^* = \{(\lambda, \lambda, \lambda) : \lambda \in (0, 1]\}$

Let  $\underline{\lambda} = \min\{\lambda : (\lambda, \lambda, \lambda) \text{ is an subgame perfect equilibrium payoff of } G^\infty(\delta)\}$ . (It can be shown quite generally that this worst payoff exists—see Section 3.3.)

We claim that for all  $\delta \in (0, 1)$ ,  $\underline{\lambda} \geq \frac{1}{4}$ :

Why? Consider the subgame perfect equilibrium that generates payoffs of  $(\underline{\lambda}, \underline{\lambda}, \underline{\lambda})$ . For any mixed strategy  $(\alpha_1, \alpha_2, \alpha_3)$  of  $G$ , there is a player who can guarantee himself  $\frac{1}{4}$ . Both the equilibrium path and the worst punishment path give each player payoffs of  $\underline{\lambda}$ . Hence, for no player to have an incentive to deviate, we must have that  $\underline{\lambda} \geq (1 - \delta)\frac{1}{4} + \delta\underline{\lambda}$ , or equivalently, that  $\underline{\lambda} \geq \frac{1}{4}$ .  $\diamond$

The problem in the previous example is that the players' payoff functions are identical, so that no one player can be rewarded or punished independently of the others. In order to obtain a folk theorem for games with three or more players, we must restrict attention to stage games in which such independent provision of incentives is possible. Fudenberg and Maskin (1986) accomplish this under the assumption that the set  $F^*$  is *full dimensional*, meaning that it has the same dimension as the number of players. Abreu et al. (1994) do so under the weaker restriction that the stage game satisfy the nonequivalent utility condition stated in Theorem 3.8, and show that this condition is essentially necessary for the conclusion of the folk theorem to obtain.

To obtain a pure minmax folk theorem for games with three or more players satisfying the NEU condition using an analogue of the stick and carrot strategy profile from Example 3.11, one first replaces the single punishment stage with distinct punishments for each player (as in Example 3.9). To provide incentives to player  $i$ 's opponents to carry out a punishment of player  $i$ , one must reward the punishers once the punishment is over. See Mailath and Samuelson (2006, Sec. 3.4.2) for a textbook treatment.



### 3.3 Computing the Set of Subgame Perfect Equilibrium Payoffs

Is there a simple characterization of the set of subgame perfect equilibrium payoffs of  $G^\infty(\delta)$  when the value of  $\delta$  is fixed?

Subgame perfect equilibria of  $G^\infty(\delta)$  have the following *recursive characterization*:

- (i) In period 0, no player has a profitable one-shot deviation.
- (ii) Every continuation strategy profile is a SPE of  $G^\infty(\delta)$ .

This decomposition is reminiscent of that of optimal policies for dynamic programs (Bellman (1957), Howard (1960)). Can this connection be exploited?

#### 3.3.1 Dynamic programming

A good reference on this material is Stokey and Lucas (1989, Sec. 4.1–4.2).

Definitions

$X \subseteq \mathbb{R}^n$	state space
$\Gamma : X \Rightarrow X$	feasibility correspondence
	$\Gamma(x)$ = states feasible tomorrow if today's state is $x$
$\Phi : X \Rightarrow X^\infty$	feasibility correspondence for sequences
	$\Phi(x) = \{\{x_t\}_{t=0}^\infty : x_0 = x, x_{t+1} \in \Gamma(x_t) \forall t \geq 0\}$
$F : X \times X \rightarrow \mathbb{R}$	payoff function. $F(x, y)$ = agent's payoff today if today's state is $x$ and tomorrow's state is $y$
$\delta \in (0, 1)$	discount rate.

*Example 3.14.* If  $x_t$  = capital at time  $t$  and  $c_t = f(x_t) - x_{t+1}$  = consumption at time  $t$ , then the payoff at time  $t$  is  $u(c_t) = u(f(x_t) - x_{t+1}) = F(x_t, x_{t+1})$ .  $\diamond$

Assumptions

- (i)  $X$  is convex or finite.
- (ii)  $\Gamma$  is non-empty, compact valued, and continuous.
- (iii)  $F$  is bounded and continuous.

Define the *dynamic program*

$$(D) \quad v(x) = \max_{\{x_t\} \in \Phi(x)} (1 - \delta) \sum_{t=0}^{\infty} \delta^t F(x_t, x_{t+1}).$$

We call  $v : X \rightarrow \mathbb{R}$  the *value function* for (D).

**Theorem 3.15** (The principle of optimality (Bellman (1957))).

Consider the functional equation

$$(F) \quad v(x) = \max_{y \in \Gamma(x)} (1 - \delta)F(x, y) + \delta v(y).$$

Then: (i) If  $v$  satisfies (D), then it satisfies (F);

(ii) If  $v$  satisfies (F) and is bounded, then it satisfies (D).

Theorem 3.15 is a direct analogue of the one-shot deviation principle (Theorem 2.25). But its statement focuses on the value function (payoffs) rather than the policy function (strategies).

Part (ii) of the theorem tells us that in order to solve the dynamic program (D), it is enough to find a bounded solution to the functional equation (F).

Solving (F) via the method of successive approximations

Let  $B(X) = \{w : X \rightarrow \mathbb{R}, w \text{ bounded}\}$ .

Let  $C(X) = \{w : X \rightarrow \mathbb{R}, w \text{ bounded and continuous}\}$ .

Introduce the supremum norm:  $\|w\| = \sup_{x \in X} |w(x)|$ .

Define the *Bellman operator*  $T: C(X) \rightarrow B(X)$  by

$$(20) \quad Tw(x) = \max_{y \in \Gamma(x)} (1 - \delta)F(x, y) + \delta w(y).$$

If  $w: X \rightarrow \mathbb{R}$  describes the continuation values tomorrow, then  $Tw(x)$  is the (optimal) value of being at  $x$  today. Notice that  $Tw \in B(X)$ .

**Observation 3.16.**  $v$  is a fixed point of  $T$  ( $v(x) = Tv(x)$  for all  $x$ ) if and only if  $v$  solves (F).

**Theorem 3.17.**

- (i)  $T(C(X)) \subseteq C(X)$ . ( $T$  maps  $C(X)$  into itself.)
- (ii)  $\|Tw - T\hat{w}\| \leq \delta \|w - \hat{w}\|$ . ( $T$  is a contraction.)
- (iii)  $T$  admits a unique fixed point  $v$ , and  $\lim_{k \rightarrow \infty} T^k w = v$  for all  $w \in C(X)$ .

Part (ii) of the theorem says that the Bellman operator is a *contraction mapping*. Given parts (i) and (ii), part (iii) follows from the Banach (or contraction mapping) fixed point theorem.

Theorem 3.17(iii) says that the functional equation (F) can be solved by *value function iteration*. One starts with an arbitrary guess  $v_0$  about the value function. One then uses  $v_0$

to define the continuation values in (20). Solving (20) generates a new value function  $v_1$ , which is then used to define the continuation values in (20). Theorem 3.17(iii) says that the sequence of value functions generated by this process converges to the the fixed point of the Bellman operator  $T$ , which is the solution to the functional equation (F), and hence the solution to the dynamic program (D).

### 3.3.2 Dynamic programs vs. repeated games

Repeated games differ from dynamic programs in two basic respects:

- (i) There is no state variable (available choices and their consequences do not depend on past choices).
- (ii) There are multiple agents.

The analogue of the value function is a subgame perfect equilibrium payoff vector.

But unlike the value function for (D), the subgame perfect equilibrium payoff vector for  $G^\infty(\delta)$  is typically *not* unique.

Nevertheless, Abreu et al. (1990) introduce a method analogous to value function iteration that computes the set of subgame perfect equilibrium payoff vectors of  $G^\infty(\delta)$ :

- (i) Introduce an operator  $B$  on {sets of payoff vectors}= $\{W \subseteq \mathbb{R}^n\}$   
(compare to the Bellman operator  $T$ , which acts on {value functions}= $\{w : X \rightarrow \mathbb{R}\}$ .)
- (ii) Observe that the set of subgame perfect equilibrium payoffs  $V$  is a fixed point of  $B$ .
- (iii) Establish properties of  $B$  (self-generation, monotonicity, compactness) which imply that iteration of  $B$  from the set of feasible, weakly individually rational payoff vectors leads to  $V$ .

### 3.3.3 Factorization and self-generation

Introduce the finite stage game  $G = \{\mathcal{P}, \{A_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\};$

its infinite repetition  $G^\infty(\delta) = \{\mathcal{P}, \{S_i\}_{i \in \mathcal{P}}, \{\pi_i\}_{i \in \mathcal{P}}, \delta\}, \delta \in (0, 1)$  fixed.

Write:  $u(a) = (u_1(a), \dots, u_n(a))$  for all  $a \in A = \prod_{i \in \mathcal{P}} A_i;$

$\pi(s) = (\pi_1(s), \dots, \pi_n(s))$  for all  $s \in S = \prod_{i \in \mathcal{P}} S_i.$

Our goal is to characterize the equilibrium value set

$V = \{\pi(s) : s \text{ is a (pure strategy) subgame perfect equilibrium of } G^\infty(\delta)\} \subseteq \mathbb{R}^n.$

We assume that a pure strategy subgame perfect equilibrium of  $G^\infty$  exists, so that the set  $V$  is nonempty. For instance, this is true if the stage game  $G$  has a pure equilibrium.

We want to factor elements of  $V$  using pairs  $(a, c)$ , where  $a \in A$  is the initial action profile and  $c : A \rightarrow \mathbb{R}^n$  is a *continuation value function*.

The *value* of pair  $(a, c)$  is the vector  $v(a, c) = (1 - \delta)u(a) + \delta c(a) \in \mathbb{R}^n$ .

Which pairs which satisfy period 0 incentive constraints?

We say that action  $a$  is *enforceable* by continuation value function  $c$  if

$$\begin{aligned} v_i(a, c) &\geq v_i((a'_i, a_{-i}), c) && \text{for all } a'_i \in A_i \text{ and } i \in \mathcal{P} \\ \Leftrightarrow (1 - \delta)u_i(a) + \delta c_i(a) &\geq (1 - \delta)u_i(a'_i, a_{-i}) + \delta c_i(a'_i, a_{-i}) && \text{for all } a'_i \in A_i \text{ and } i \in \mathcal{P}. \end{aligned}$$

We say that  $w \in \mathbb{R}^n$  is *enforceably factored* by  $(a, c)$  into  $W \subseteq \mathbb{R}^n$  if

- (i)  $v(a, c) = w$  (the pair  $(a, c)$  generates value  $w$ )
- (ii)  $a$  is enforceable by  $c$  (the period 0 choices  $a$  are incentive compatible)
- (iii)  $c(A) \subseteq W$  (continuation payoffs are drawn from  $W$ )

Define the operator  $B : 2^{\mathbb{R}^n} \rightarrow 2^{\mathbb{R}^n}$  by

$$B(W) = \{w \in \mathbb{R}^n : w \text{ can be enforceably factored into } W\}.$$

(mnemonic:  $B$  describes values that can come *before* continuation values in  $W$ ).

We say that  $W$  is *self-generating* if  $W \subseteq B(W)$ . This means that one can enforceably factor any payoff vector in  $W$  into an action profile today and continuation values in  $W$  tomorrow, so that the problem of obtaining a value in  $W$  can be “put off until tomorrow”. But when tomorrow comes, we can put off obtaining the value in  $W$  until the following day... If we repeat this indefinitely, then in the end we will have constructed a subgame perfect equilibrium. ( $W$  must be bounded for this to work for certain: see Example 3.23 below.)

**Theorem 3.18.** *If  $W$  is self-generating and bounded, then  $W \subseteq V$ .*

Theorem 3.19 is the recursive characterization of subgame perfect equilibrium payoff vectors:  $v$  is a subgame perfect equilibrium payoff vector if and only if it can be enforceably factored into an action profile today and subgame perfect continuation values tomorrow.

**Theorem 3.19.**  $V = B(V)$ .

Theorem 3.20 provides an algorithm for computing the set of subgame perfect equilibrium payoff vectors. To state it, we define  $W_0 = \{w \in \text{conv}(u(A)) : w_i \geq \underline{v}_i \text{ for all } i \in \mathcal{P}\}$  to be the set of feasible, weakly individually rational payoff vectors, and define  $\{W_k\}_{k=1}^\infty$  inductively by  $W_{k+1} = B(W_k)$ . Obviously  $W_1 = B(W_0) \subseteq W_0$ , and in fact, the sequence  $\{W_k\}_{k=1}^\infty$  is monotone. Since  $V$  is a fixed point of  $B$ , it is a plausible candidate to be the limit of the sequence. Theorem 3.20 confirms that this is the case.

**Theorem 3.20.**  $\bigcap_{k=0}^{\infty} W_k$  exists and equals  $V$ .

In this theorem,  $W_k$  is the set of equilibrium payoff vectors of games of the following form:  $k$  rounds of the stage game  $G$ , followed by (history-dependent) continuation values drawn from  $W_0$  (mnemonic:  $W_k = B^k(W_0)$  comes  $k$  periods before  $W_0$ ). As  $k$  grows larger, the time before the continuation payoffs from  $W_0$  appear is put off further and further into the future.

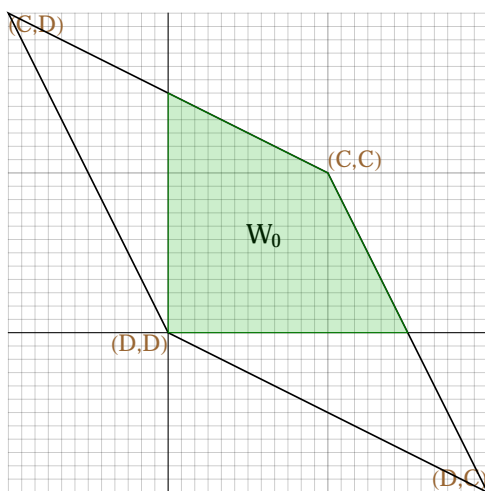
Remarks:

1. All of these results can be generalized to repeated games with imperfect public monitoring—see Abreu et al. (1990).
2. Using these results, one can prove that the set  $V$  of (normalized) subgame perfect equilibrium payoffs is monotone in the discount factor  $\delta \in (0, 1)$ . In other words, increasing players' patience increases the set of equilibrium outcomes.

*Example 3.21.* Consider an infinite repetition of the Prisoner's Dilemma below. What is the set of subgame perfect equilibrium payoffs when  $\delta = \frac{3}{4}$ ?

		2	
		C	D
1	C	1, 1	-1, 2
	D	2, -1	0, 0

Let  $W_0$  be the set of feasible, weakly individually rational payoff vectors:



To begin, we compute the set  $W_1 = B(W_0)$ . For each action profile  $a$ , we determine the set  $W_1^a$  of payoff vectors that can be enforceably factored by some pair  $(a, c)$  into  $W_0$  (meaning that  $c: A \rightarrow W_0$ ). Then  $W_1$  is the union of these sets.

First consider  $a = (D, D)$ . For this action profile to be enforceable, neither player can prefer to deviate to C. Player 1 does not prefer to deviate if

$$\begin{aligned} (1 - \delta)u_1(D, D) + \delta c_1(D, D) &\geq (1 - \delta)u_1(C, D) + \delta c_1(C, D) \\ \Leftrightarrow \frac{1}{4} \cdot 0 + \frac{3}{4}c_1(D, D) &\geq \frac{1}{4} \cdot (-1) + \frac{3}{4}c_1(C, D) \\ \Leftrightarrow c_1(D, D) &\geq -\frac{1}{3} + c_1(C, D). \end{aligned}$$

Similarly, player 2 prefers not to deviate if

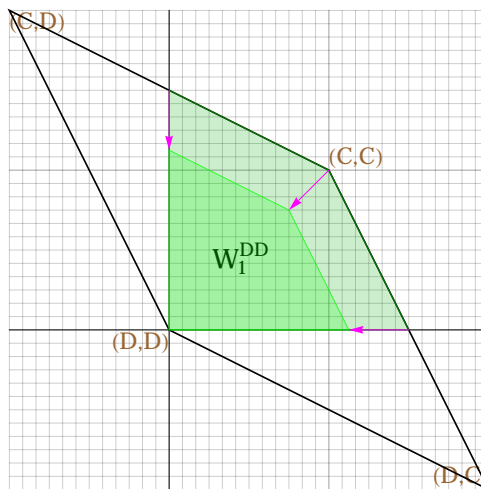
$$c_2(D, D) \geq -\frac{1}{3} + c_2(D, C).$$

These inequalities show that if  $c(D, D) = c(C, D) = c(D, C)$ , the pair  $(a, c)$  will be enforceable. (This makes sense: one does not need to promise future rewards to make players choose a dominant action.) Thus, for any  $w \in W_0$ , we can enforce action profile  $(D, D)$  using a continuation value function  $c$  with  $c(D, D) = w$ .

The value for the pair  $((D, D), c)$  with  $c(D, D) = w$  is

$$(1 - \delta)u(D, D) + \delta c(D, D) = \frac{1}{4} \cdot (0, 0) + \frac{3}{4} \cdot w.$$

In the figure below, the full shaded area is  $W_0$ ; the smaller shaded area is  $W_1^{DD} = \frac{1}{4} \cdot (0, 0) + \frac{3}{4} \cdot W_0$ , the set of payoff vectors that can be enforceably factored by some pair  $((D, D), c)$  into  $W_0$ .



Now consider  $a = (C, C)$ . In this case, the enforceability constraints are

$$(1 - \delta)u_1(C, C) + \delta c_1(C, C) \geq (1 - \delta)u_1(D, C) + \delta c_1(D, C)$$

$$\Leftrightarrow \frac{1}{4} \cdot 1 + \frac{3}{4}c_1(C, C) \geq \frac{1}{4} \cdot 2 + \frac{3}{4}c_1(D, C)$$

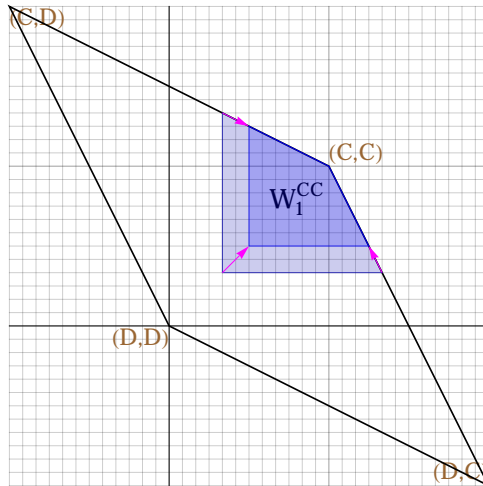
$$\Leftrightarrow c_1(C, C) \geq \frac{1}{3} + c_1(D, C), \text{ and, symmetrically,}$$

$$c_2(C, C) \geq \frac{1}{3} + c_2(C, D).$$

These calculations show that for  $c(C, C) \in W_0$  to be part of an enforceable factoring of  $((C, C), c)$  into  $W_0$ , there must be a point in  $W_0$  that is  $\frac{1}{3}$  units to the left of  $c(C, C)$  (to punish player 1 if he deviates) as well as a point in  $W_0$  that is  $\frac{1}{3}$  units below  $c(C, C)$  (to punish player 2 if she deviates). Since  $(0, 0)$  is both the leftmost and the lowest point in  $W_0$ , it follows that for any  $w \in \{w \in W_0 : w_1, w_2 \geq \frac{1}{3}\}$ , we can enforce  $(C, C)$  using a  $c$  with  $c(C, C) = w$ . The value for the pair  $((C, C), c)$  with  $c(C, C) = w$  is

$$(1 - \delta)u(C, C) + \delta c(C, C) = \frac{1}{4} \cdot (1, 1) + \frac{3}{4} \cdot w.$$

In the figure below, the full shaded area is  $\{w \in W_0 : w_1, w_2 \geq \frac{1}{3}\}$ , the set of allowable values for  $c(C, C)$ ; the smaller shaded area is the set  $W_1^{CC} = \frac{1}{4} \cdot (1, 1) + \frac{3}{4} \cdot \{w \in W_0 : w_1, w_2 \geq \frac{1}{3}\}$ .



Next consider  $a = (C, D)$ . In this case the enforceability constraints are

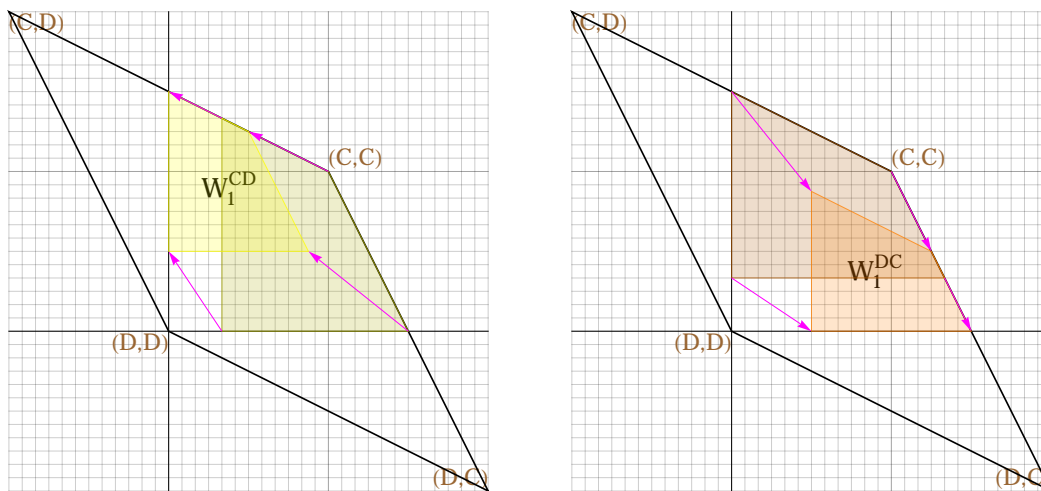
$$c_1(C, D) \geq \frac{1}{3} + c_1(D, D) \text{ and } c_2(C, D) \geq -\frac{1}{3} + c_2(C, C).$$

That is, we only need to provide incentives for player 1. Reasoning as above, we find that for any  $w \in \{w \in W_0 : w_1 \geq \frac{1}{3}\}$ , we can enforce  $(C, D)$  using a  $c$  with  $c(C, D) = w$ . The value for the pair  $((C, D), c)$  with  $c(C, D) = w$  is

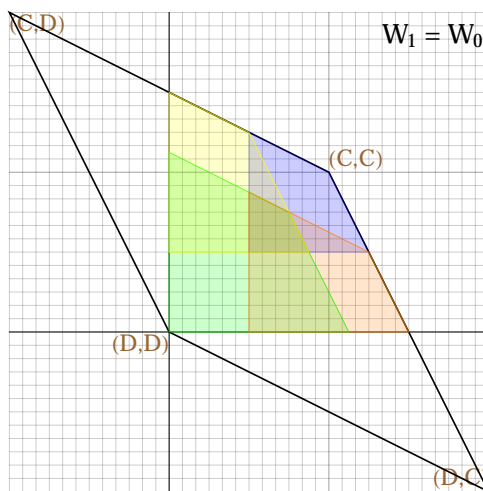
$$(1 - \delta)u(C, D) + \delta c(C, D) = \frac{1}{4} \cdot (-1, 2) + \frac{3}{4} \cdot w.$$

In the figure below at left, the larger shaded area is  $\{w \in W_0 : w_1 \geq \frac{1}{3}\}$ ; the smaller shaded

area is the set  $W_1^{CD} = \frac{1}{4} \cdot (1, 1) + \frac{3}{4} \cdot \{w \in W_0 : w_1 \geq \frac{1}{3}\}$ . The figure below at right shows the construction of  $W_1^{DC}$ , which is entirely symmetric.



Below we draw  $W_1 = W_1^{DD} \cup W_1^{CC} \cup W_1^{CD} \cup W_1^{DC}$ . Evidently,  $W_1 = W_0$ .

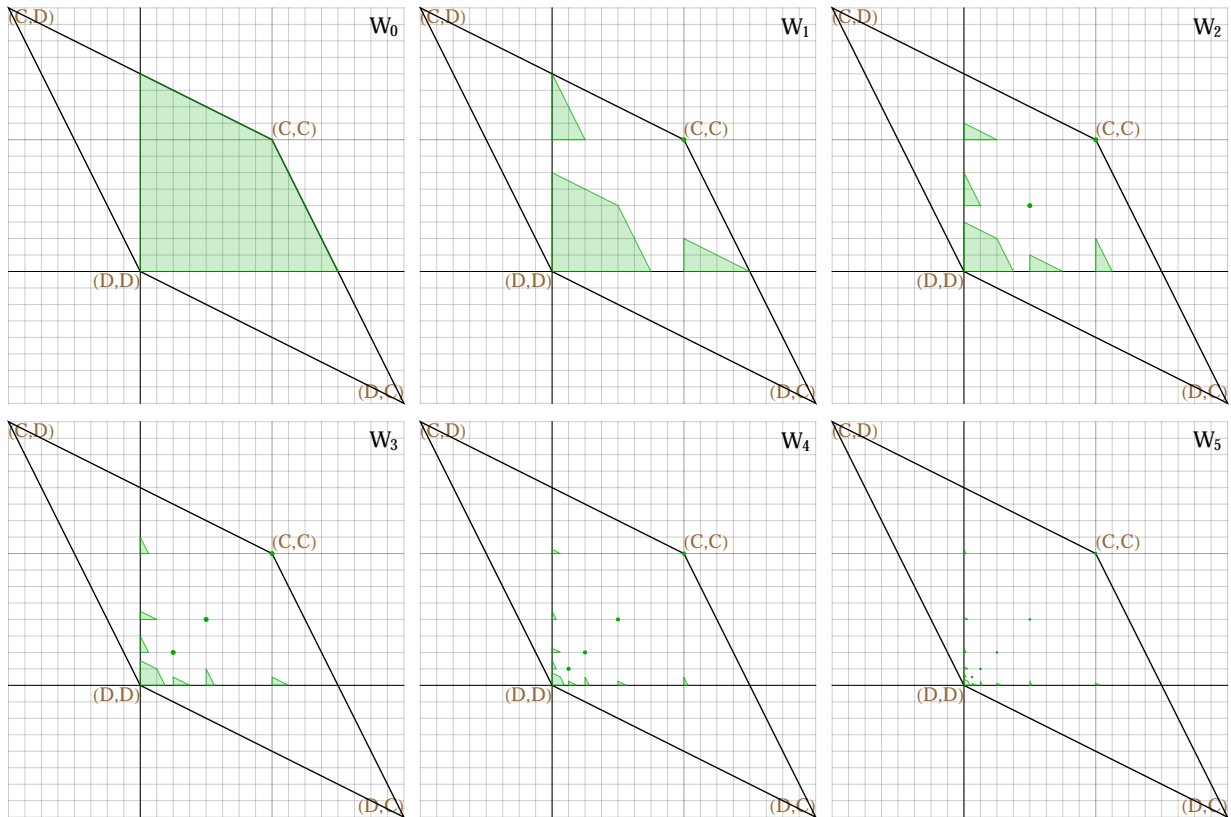


Repeating the argument above shows that  $W_{k+1} = W_k = \dots = W_0$  for all  $k$ , implying that  $V = W_0$ . In other words, all feasible, weakly individually rational payoffs are achievable in subgame perfect equilibrium when  $\delta = \frac{3}{4}$ .  $\diamond$

This example was especially simple. In general, each iteration  $W_0 \mapsto W_1 \mapsto \dots \mapsto W_k \mapsto \dots$  eliminates additional payoff vectors, and  $V$  is only obtained in the limit.

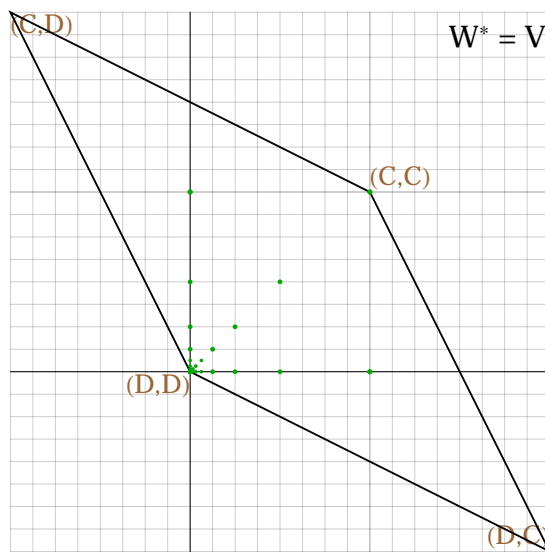
*Example 3.22.* Consider the same Prisoner's Dilemma stage game as in the previous example, but suppose that  $\delta = \frac{1}{2}$ . The sets  $W_0, \dots, W_5$  are shown below.





Taking the limit, we find that the set of subgame perfect equilibrium payoffs is

$$V = \left( \bigcup_{k=0}^{\infty} \left\{ \left( \frac{1}{2^k}, \frac{1}{2^k} \right), \left( \frac{1}{2^k}, 0 \right), \left( 0, \frac{1}{2^k} \right) \right\} \right) \cup \{ (0, 0) \}.$$



The payoff  $(1, 1)$  is obtained in a subgame perfect equilibrium from the grim trigger strategy profile with equilibrium path  $(C, C), (C, C), \dots$ . The payoff  $(1, 0)$  is obtained from the one with equilibrium path  $(D, C), (C, D), (D, C), (C, D), \dots$ ; payoff  $(0, 1)$  is obtained by reversing roles. Payoff  $(\frac{1}{2^k}, \frac{1}{2^k})$  is obtained by playing  $(D, D)$  for  $k$  periods before beginning the cooperative phase; similarly for  $(\frac{1}{2^k}, 0)$  and  $(0, \frac{1}{2^k})$ . Finally the “always defect” strategy profile yields payoff  $(0, 0)$ .  $\diamond$

### *Proofs of the theorems*

Once one understands why Theorems 3.18 and 3.19 are true, the proofs are basically bookkeeping. The proof of Theorem 3.20, which calls on the previous two theorems, requires more work, but it can be explained quickly if the technicalities are omitted. This is what we do below.

*Proof of Theorem 3.18:*  $W$  bounded,  $W \subseteq B(W) \Rightarrow W \subseteq V$ .

*Example 3.23.* That  $W$  is bounded ensures that we can't put off actually receiving our payoffs forever. To see this, suppose that  $W = \mathbb{R}_+$  and  $\delta = 1/2$ . Then we can decompose  $1 \in W$  as  $1 = 0 + \delta \cdot 2$ . And we can decompose  $2 \in W$  as  $2 = 0 + \delta \cdot 4$ . And we can decompose  $4 \in W$  as  $4 = 0 + \delta \cdot 8 \dots$  And the payoff 1 is never obtained.  $\diamond$

The proof: Let  $w^0 \in W$ . We want to construct a SPE with payoffs  $w^0$ .

Since  $w^0 \in W \subseteq B(W)$ , it is enforceably factored by some  $(a^0, c^0)$  into  $W$ .

$\Rightarrow a^0$  is the period 0 action profile in our SPE.

$\Rightarrow$  for each  $a \in A$ ,  $c^0(a) \in W$  is the payoff vector to be obtained from period 1 on after  $h^1 = \{a\}$ .

Now consider the period 1 history  $h^1 = \{a\}$ .

Let  $w^1 = c^0(a)$ . Then  $w^1 \in W \subseteq B(W)$ .

That is,  $w^1$  is enforceably factored by some  $(a_a^1, c_a^1)$  into  $W$ .

$\Rightarrow a_a^1$  is the period 1 action profile occurring after  $h^1 = \{a\}$  in our SPE.

$\Rightarrow$  for each  $\hat{a} \in A$ ,  $c_a^1(\hat{a}) \in W$  is the payoff vector to be obtained from period 2 on after  $h^2 = \{a, \hat{a}\} \dots$

By the one-shot deviation principle, the strategy profile so constructed is a SPE.

Since  $W$  is bounded, this strategy profile achieves the payoff of  $w^0$ .  $\blacksquare$

*Proof of Theorem 3.19:*  $V = B(V)$ .

(I): Proof that  $V \subseteq B(V)$ .

Let  $w \in V$ . Then  $w = \pi(s)$  for some SPE  $s$ .

Let  $a^0 = s(h^0)$  be the initial strategy profile under  $s$ .

For each  $a \in A$ , let  $c(a) = \pi(s|_a)$ , where  $s|_a$  is the strategy profile starting in period 1 if action profile  $a$  is played in period 0. Then:

- (i)  $w = \pi(s) = v(a^0, c)$ . (In words: the pair  $(a^0, c)$  generates the value  $w$ .)
- (ii) Since  $s$  is a SPE, enforceability in period 0 requires that

$$(1 - \delta)u_i(a^0) + \delta\pi_i(s|_{a^0}) \geq (1 - \delta)u_i(a_i, a_{-i}^0) + \delta\pi_i(s|_{(a_i, a_{-i}^0)}) \quad \text{for all } a_i \in A_i, i \in \mathcal{P}.$$

$$\Leftrightarrow v_i(a^0, c) \geq v_i((a_i, a_{-i}^0), c) \quad \text{for all } a_i \in A_i, i \in \mathcal{P}.$$

and so  $a^0$  is enforceable by  $c$ . (In words: period 0 choices are incentive compatible.)

- (iii) Since  $s$  is a SPE, so is  $s|_a$  for every action profile  $a$ , and thus  $c(a) \in V$  for all  $a \in A$ . (In words: continuation values are in  $V$ .)

In conclusion,  $w$  is enforceably factored by  $(a^0, c)$  into  $V$ .

(II): Proof that  $B(V) \subseteq V$ .

Let  $w \in B(V)$ . Then  $w$  is enforceably factored by some  $(a^0, c)$  into  $V$ .

Define strategy profile  $s$  by (i)  $s(h^0) = a^0$ , and (ii) for each  $a \in A$ ,  $s|_a = \hat{s}_a$ , where  $\hat{s}_a$  is an SPE with payoffs  $c(a)$ . (Such an  $\hat{s}_a$  exists because  $c(a) \in V$ .)

Since  $a^0$  is enforceable by  $c$  and each  $s|_a$  is an SPE,  $s$  is also an SPE.

Therefore,  $w = \pi(s) \in V$ . ■

*Proof of Theorem 3.20:*  $W_0$  compact,  $W_0 \supseteq B(W_0) \supseteq V$ ,  $W_{k+1} = B(W_k) \Rightarrow \bigcap_{k=0}^{\infty} W_k = V$ .

(Note: Our choice of  $W_0 = \{w \in \text{conv}(u(A)) : w_i \geq \underline{v}_i \text{ for all } i \in \mathcal{P}\}$  satisfies the requirements above, which are all that are needed for the theorem to hold.)

The map  $B$  is clearly monotone:  $W' \subseteq W \Rightarrow B(W') \subseteq B(W)$ .

One can show that  $B$  preserves compactness.

Thus, the Cantor intersection theorem implies that  $W^* \equiv \bigcap_{k=0}^{\infty} W_k$  is nonempty and compact. A continuity argument allows one to conclude that  $B(W^*) = W^*$ .

We now show that  $W^* = V$ . Since  $W^* \subseteq B(W^*)$ , Theorem 3.18 implies that  $W^* \subseteq V$ . But since  $V \subseteq W_0$  (by assumption) and since  $V = B(V)$  (by Theorem 3.19), monotonicity implies that

$$V = B(V) \subseteq B(W_0) = W_1, \text{ and so that}$$

$$V = B(V) \subseteq B(W_1) = W_2 \dots$$

Continuing, we find that  $V \subseteq W^*$ . We therefore conclude that  $V = W^*$ . ■

## 4. Bayesian Games

A strategic environment is said to have *incomplete information* if when play begins, players possess payoff-relevant information that is not common knowledge. *Bayesian games* (Harsanyi (1967–1968)) provide a tractable way of modeling such environments. While many basic Bayesian games are not difficult to interpret, the Bayesian game framework turns out to be extremely general, and so involves many subtleties and complexities. Sections 4.1–4.6 focus on games whose interpretations are relatively straightforward; more ingenious games, foundational questions, and other challenging material is presented in Sections 4.7–4.9.

### 4.1 Basic Definitions

#### *Defining Bayesian games*

The fundamental new notion in Bayesian games is that of a *type*. A player’s type specifies whatever payoff-relevant information and beliefs he holds that other players do not know. For instance, in an auction environment, a player’s type includes both his private information about the good for sale, as well as his beliefs about other players’ types. More on this shortly.

A *Bayesian game* (Harsanyi (1967–1968)) is a collection  $BG = \{\mathcal{P}, \{A_i\}_{i \in \mathcal{P}}, \{T_i\}_{i \in \mathcal{P}}, \{p_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$ .

$\mathcal{P} = \{1, \dots, n\}$	the set of players
$a_i \in A_i$	a pure action for player $i$
$t_i \in T_i$	a <i>type</i> of player $i$ . Only $i$ knows $t_i$ when play begins.
$p_i : T_{-i} \rightarrow \Delta T_{-i}$	the <i>first-order belief function</i> of player $i$ $p_i(\cdot   t_i)$ describes the first-order beliefs of type $t_i$ $p_i(t_{-i}   t_i)$ is the probability that player $i$ of type $t_i$ assigns to the other players being of types $t_{-i}$
$u_i : A \times T \rightarrow \mathbb{R}$	the <i>utility function</i> of player $i$ (where $A = \prod_{j \in \mathcal{P}} A_j$ and $T = \prod_{j \in \mathcal{P}} T_j$ ) $i$ 's utility can depend on all players' types and actions

As in our previous game models, the definition of the game is common knowledge among the players.

The interaction defined above is a simultaneous move Bayesian game: all players choose actions at the same time. We concentrate on this setting to focus our attention on the novel aspects of Bayesian games.

We usually think of these choices being made when each player  $i$  knows his own type  $t_i$ , but not the other players' types. This moment is known as *interim stage*.

*Example 4.1. A  $k$ -card card game.* Two players play a game using a deck containing  $k$  cards that are numbered 1 through  $k$ . To start, each player antes \$1. The deck is shuffled and each player is dealt one card. Each player simultaneously chooses to Play, which requires putting an additional \$1 into the pot, or to Fold. If both players Play, the player with the higher card takes the pot. If only one player Plays, he takes the pot. If both players Fold, each takes back his ante.

This game is described formally as follows:

$$\begin{aligned}
 A_i &= \{P, F\}, \\
 T_i &= \{1, \dots, k\}, \\
 p_i(t_i | t_j) &= \begin{cases} \frac{1}{k-1} & \text{if } t_i \neq t_j, \\ 0 & \text{if } t_i = t_j, \end{cases} \\
 u_i(a, t) &= \begin{cases} 2 & \text{if } a_i = a_j = P, t_i > t_j, \\ 1 & \text{if } a_i = P, a_j = F, \\ 0 & \text{if } a_i = a_j = F, \\ -1 & \text{if } a_i = F, a_j = P, \\ -2 & \text{if } a_i = a_j = P, t_i < t_j. \quad \diamond \end{cases}
 \end{aligned}$$

### *Common prior distributions*

Beliefs in  $BG$  can be derived from a *common prior distribution*  $p \in \Delta T$  if the first-order beliefs  $p_i(\cdot | t_i)$  are conditional probabilities (or *posterior probabilities*) generated from  $p$ :

$$(21) \quad p_i(t_{-i} | t_i) = \frac{p(t_i, t_{-i})}{\sum_{\hat{t}_{-i} \in T_{-i}} p(t_i, \hat{t}_{-i})}.$$

*Example 4.2.* The common prior distribution of the  $k$ -card card game describes the distribution of possible deals:

$$p(t_1, t_2) = \begin{cases} \frac{1}{k(k-1)} & \text{if } t_1 \neq t_2, \\ 0 & \text{if } t_1 = t_2. \quad \diamond \end{cases}$$

In the card game, there common prior describes both players' beliefs at the *ex ante stage*, meaning the time before the cards are dealt.

Most economic applications of Bayesian games employ the *common prior assumption*—the assumption that the players' first order beliefs  $p_i$  can be derived from some common prior  $p$  using (21). This assumption means that it is *as if* the game has an ex ante stage at which no player has learned his type. See Section 4.4 for a discussion.

The next observation is useful for understanding Bayesian games with a common prior:

**Observation 4.3.** *A Bayesian game  $BG$  with common prior distribution  $p$  is equivalent to the following extensive form game  $\Gamma_{BG}$ :*

*Stage 0: Nature draws  $t$  according to the common prior  $p \in \Delta T$ . Player  $i$  learns  $t_i$  only.*

*Stage 1: Each player  $i$  chooses an action from  $A_i$ .*

We refer to this observation often below.

There is also a way of representing *any* Bayesian game using a normal form game with  $\sum_{i \in \mathcal{P}} |T_i|$  players—see the end of Section 4.2.

### ***Independent types and private values***

If game  $BG$  has a common prior  $p$  that is a product distribution, so that  $p_i(\cdot | t_i)$  is independent of  $t_i$ , then  $BG$  has *independent types*. Otherwise, types are *dependent*, though the less accurate term *correlated* is often used instead.

If game  $BG$  satisfies  $u_i(a, t) \equiv u_i(a, t_i)$  (i.e., if player  $i$ 's utility does not depend directly on other players' types), then  $BG$  is said to have *private values*. Otherwise,  $BG$  has *interdependent values*.

The combination of independent types and private values is called *independent private values* for short.

### ***Interpreting types: a first pass***

In Bayesian games with independent types (including those in which only one player has multiple types), we can identify a player's type with some payoff-relevant information that he possesses. For instance, imagine that a seller is auctioning off a collection of disparate goods as a bundle. Each bidder receives a private signal about the quality of a distinct good from the bundle, and these signals are independent of one another. Because the goods are sold as a bundle, the signals of a bidder's opponents contain information about the value of the bundle to the bidder himself. This aspect of types is reflected in the fact that player  $i$ 's utility function  $u_i$  may condition on his own type  $t_i$  and on other players' types  $t_{-i}$ .

In Bayesian games with correlated types, a player's type also describes beliefs that he holds but that other players do not know. For instance, consider this environment for *mineral rights auctions* (see also Section 4.4): A seller is auctioning the right to drill for oil on a certain plot of land. Each bidder receives an unbiased private signal about the value of the land (obtained, for instance, by having a team of geologists examine core samples).

In this case, a bidder who receives a high signal about the land's quality is likely to believe that other bidders have also received high signals. These beliefs depend on the bidder's signal, and so are only known to the bidder himself. This aspect of types is reflected in the fact that a player's first-order beliefs  $p_i$  depend on his type  $t_i$ .

Notice that one aspect of a player's type is his beliefs about other players' types, and that one aspect of those types is those players' beliefs about other players' types, and that one aspect of those types is those players' beliefs about other players' types. . . Thus in settings with correlated types, the simple definition of Bayesian games given above is less simple than it seems: in addition to first-order beliefs, we have second-order beliefs, third-order beliefs, and so on. We return to this point in Sections 4.7–4.9.

## 4.2 Bayesian Strategies, Dominance, and Bayesian Equilibrium

### *Bayesian strategies*

$s_i : T_i \rightarrow A_i$  a pure Bayesian strategy

$s_i(t_i) \in A_i$  is the action chosen by type  $t_i$

$\sigma_i : T_i \rightarrow \Delta A_i$  a mixed Bayesian strategy

$\sigma_i(a_i | t_i) \in [0, 1]$  is the probability that a player  $i$  of type  $t_i$  plays action  $a_i$

Let  $S_i$  denote player  $i$ 's set of pure Bayesian strategies. Let  $S_{-i} = \prod_{j \neq i} S_j$  and  $S = \prod_{j \in \mathcal{P}} S_j$ .

In games with an ex ante stage, a Bayesian strategy describes a player's plan for playing the game at that stage. (For other games, see Section 4.4.)

### *Interim payoffs*

If player  $i$  of type  $t_i$  chooses action  $a_i$ , and each opponent  $j \neq i$  plays some Bayesian strategy  $s_j$ , then type  $t_i$ 's *interim payoff* (or *expected payoff*) is

$$(22) \quad U_i(a_i, s_{-i} | t_i) = \sum_{t_{-i} \in T_{-i}} p_i(t_{-i} | t_i) u_i((a_i, s_{-i}(t_{-i})), (t_i, t_{-i})).$$

Why  $s_{-i}(t_{-i})$ ? Other players' types determine their actions.

Why  $(t_i, t_{-i})$ ? All players' types may directly affect  $i$ 's payoffs.

*Example 4.4.* Consider the card game from Example 4.1, in which (i) each player is dealt a card which only he observes, and then (ii) each player simultaneously chooses an action. According to definition (22), when a player with card  $t_i$  is deciding what to do, he uses his posterior beliefs  $p_i(\cdot | t_i)$  as his assessment of the probabilities to the possible cards  $t_j$  of his

opponent. (In this example the correlation in types takes a simple form: the player knows that his opponent's card differs from his own, but puts equal weight on his opponent holding each of the remaining cards.)

Now suppose that (as assumed in equilibrium) player  $i$  correctly anticipates his opponent's strategy  $s_{-i}$ , and hence the action  $s_j(t_j)$  that his opponent would play if she had card  $t_j$ . Definition (22) indicates that player  $i$ 's payoffs are affected by his opponent's card in two distinct ways: there is a direct effect (if both players Play, he will win if his card is better than his opponent's), as well as an indirect effect (his opponent's card determines whether she will Play or Fold).  $\diamond$

### *Dominance*

Action  $a_i$  is *very weakly dominant* for type  $t_i \in T_i$  if

$$(23) \quad U_i(a_i, s_{-i}|t_i) \geq U_i(\hat{a}_i, s_{-i}|t_i) \quad \text{for all } s_{-i} \in S_{-i} \text{ and } \hat{a}_i \in A_i.$$

In words: action  $a_i$  is optimal for  $t_i$  given his beliefs about his opponents' types, regardless of his opponents' Bayesian strategies.

Bayesian strategy  $s_i$  is *very weakly dominant* if for each type  $t_i \in T_i$ , action  $s_i(t_i)$  is very weakly dominant for  $t_i$ . That is:

$$(24) \quad U_i(s_i(t_i), s_{-i}|t_i) \geq U_i(\hat{a}_i, s_{-i}|t_i) \quad \text{for all } s_{-i} \in S_{-i}, \hat{a}_i \in A_i, \text{ and } t_i \in T_i.$$

This is the standard notion of dominance used in mechanism design (see Section 7.1.3).

In Bayesian games with private values, ( $u_i(a, t) \equiv u_i(a, t_i)$ ), the interim condition (23) is equivalent to the simpler ex post condition

$$u_i((a_i, a_{-i}), t_i) \geq u_i((\hat{a}_i, a_{-i}), t_i) \quad \text{for all } a_{-i} \in A_{-i} \text{ and } \hat{a}_i \in A_i.$$

(Verifying this is a good finger exercise.) Thus with private values, (24) simplifies to

$$(25) \quad u_i((s_i(t_i), a_{-i}), t_i) \geq u_i((\hat{a}_i, a_{-i}), t_i) \quad \text{for all } a_{-i} \in A_{-i}, \hat{a}_i \in A_i, \text{ and } t_i \in T_i.$$

We introduce an ex post (but equilibrium) solution concept for games with interdependent values in Section 4.6.

Action  $a_i$  is *weakly dominant* for type  $t_i$  if, in addition to (23), for all  $\hat{a}_i \neq a_i$  there is an  $s_{-i} \in S_{-i}$  for which the inequality in (23) is strict. Bayesian strategy  $s_i$  is *weakly dominant* if it is very weakly dominant and if  $s_i(t_i)$  is weakly dominant for some type  $t_i \in T_i$ .



One can define weak domination, strict dominance, and strict domination in a similar fashion.

### ***Bayesian equilibrium***

Pure strategy profile  $s \in S$  is a *Bayesian equilibrium* of  $BG$  if

$$(26) \quad U_i(s_i(t_i), s_{-i}|t_i) \geq U_i(\hat{a}_i, s_{-i}|t_i) \quad \text{for all } \hat{a}_i \in A_i, t_i \in T_i, \text{ and } i \in \mathcal{P}.$$

In words: for each player  $i$  and each type  $t_i$ , action  $s_i(t_i)$  is optimal given  $t_i$ 's beliefs about his opponents' types, and given his opponents' Bayesian strategies  $s_{-i}$ .

The definition of mixed strategy Bayesian equilibrium is what one would expect, but requires more notation.

If player  $i$  of type  $t_i$  chooses mixed action  $\alpha_i$ , and each opponent  $j \neq i$  plays some mixed Bayesian strategy  $\sigma_j$ , then type  $t_i$ 's expected payoff is

$$U_i(\alpha_i, \sigma_{-i}|t_i) = \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) \left( \sum_{a \in A} \left( \alpha_i(a_i) \cdot \prod_{j \neq i} \sigma_j(a_j | t_j) \right) u_i(a, (t_i, t_{-i})) \right).$$

Mixed strategy profile  $\sigma$  is a *Bayesian equilibrium* of  $BG$  if

$$U_i(\sigma_i(t_i), \sigma_{-i}|t_i) \geq U_i(\hat{\alpha}_i, \sigma_{-i}|t_i) \quad \text{for all } \hat{\alpha}_i \in \Delta A_i, t_i \in T_i, \text{ and } i \in \mathcal{P}.$$

Remarks:

- (i) Since we are considering Bayesian games with simultaneous moves, Bayesian equilibrium need not (and does not) include any notion of sequential rationality. For this reason, the terms *Bayes-Nash equilibrium* or just *Nash equilibrium* are often used in place of Bayesian equilibrium.
- (ii) If Bayesian game  $BG$  has a common prior, then a Bayesian equilibrium of  $BG$  is just a Nash equilibrium of the extensive form game  $\Gamma_{BG}$  from Observation 4.3.

### ***The interim normal form***

A Bayesian game  $BG$  can be represented as a simultaneous move game using the *interim normal form*  $INF(BG)$ . If  $BG$  has player set  $\mathcal{P}$  and type sets  $T_i$ , the set of players of  $INF(BG)$  is  $\{(i, t_i) : i \in \mathcal{P}, t_i \in T_i\}$ : that is, the interim normal form has one player for every type of every player in the original game, and hence  $\sum_{i \in \mathcal{P}} |T_i|$  players in total. Player  $(i, t_i)$

has action set  $A_i$ , and his utility function is the interim payoff function  $U_i(a_i, s_{-i}|t_i)$  from (22). (Notice that in  $INF(BG)$ , player  $(i, t_i)$ 's payoffs are independent of the actions of any opponents  $(i, \hat{t}_i)$  corresponding to the same player in the original game  $BG$ .)

Because Bayesian equilibrium only depends on interim payoffs, the following observation is immediate:

**Observation 4.5.** *A pure (or mixed) strategy profile is a Bayesian equilibrium of  $BG$  if and only if it corresponds to a pure (or mixed) Nash equilibrium of  $INF(BG)$ .*

It follows immediately from this observation and Theorem 1.39 that every Bayesian game with finite type and action sets admits a (possibly mixed) Bayesian equilibrium.

### 4.3 Computing Bayesian Equilibria

The remarks above show that Bayesian equilibrium is really just Nash equilibrium in the context of a Bayesian game. Thus in settings with finite numbers of types and actions (and even some without), computation of Bayesian equilibrium works like computation of Nash equilibrium: first apply iterated dominance arguments, and then check all combinations of strategy profiles that remain. In both cases, equilibria are fixed points, and finding all fixed points generally requires an exhaustive search.

*Example 4.6.* Example 2.51 introduced the extensive-form Ace-King-Queen Poker game  $\Gamma_{AKQ}$ . In this game, the cards are dealt, player 1 chooses to Raise ( $R$ ) or Fold ( $F$ ), and in the event of a raise, player 2 chooses to Call ( $c$ ) or Fold ( $f$ ).

To represent this interaction as a simultaneous move Bayesian game  $BG_{AKQ}$ , we describe the play of the game after the cards are dealt as a simultaneous move game. This game has type spaces  $T_1 = \{A, K, Q\}$ ,  $T_2 = \{a, k, q\}$  and a common prior  $p \in \Delta T$ , specified in the next table, that describes the possible deals of the cards.

		$t_2$		
		$a$	$k$	$q$
$t_1$	$A$	0	$\frac{1}{6}$	$\frac{1}{6}$
	$K$	$\frac{1}{6}$	0	$\frac{1}{6}$
	$Q$	$\frac{1}{6}$	$\frac{1}{6}$	0

Payoffs are described below. Notice that payoffs only depend on types if player 1 raises and player 2 calls, in which case the player with the higher card wins:

$$\begin{array}{c}
2 \\
\begin{array}{cc}
c & f \\
1 \begin{array}{c} R \\ F \end{array} \begin{array}{|c|c|} \hline u((R,c),t) & 1,-1 \\ \hline -1,1 & -1,1 \\ \hline \end{array} \end{array}
\end{array}
\quad \text{where } u((R,c),t) = \begin{cases} (2,-2) & \text{if } t \in \{Ak, Aq, Kq\}, \\ (-2,2) & \text{if } t \in \{Ka, Qa, Qk\}. \end{cases}$$

In Example 2.51, we showed that the unique sequential equilibrium of the extensive form game  $\Gamma_{AKQ}$  is

$$(27) \quad \sigma_1(R|A) = 1, \sigma_1(R|K) = 1, \sigma_1(R|Q) = \frac{1}{3}; \quad \sigma_2(c|a) = 1, \sigma_2(c|k) = \frac{1}{3}, \sigma_2(c|q) = 0.$$

By using sequential equilibrium as the solution concept, we appealed to the principle of sequential rationality, requiring players to act optimally even at unreached information sets.

We now compute the Bayesian equilibria of  $BG_{AKQ}$ . The analysis differs from that of the sequential equilibria of  $\Gamma_{AKQ}$ , because Bayesian equilibrium does not require sequential rationality. For instance, if player 2 of type  $a$  were certain that her opponent would fold, then both calling ( $c$ ) or folding ( $f$ ) would be a best response for her, and so she could play either, despite the fact that  $c$  is the best response against both of her possible opponents (types  $K$  and  $Q$ ). In contrast, in a sequential equilibrium of  $\Gamma_{AKQ}$ , sequential rationality requires player 2 of type  $a$  to play  $c$  regardless of player 1's strategy.

Nevertheless, we now argue that (27) is also the unique Bayesian equilibrium of  $BG_{AKQ}$ . That is, we show that (27) is the only Bayesian strategy profile in which each type of player is choosing an optimal action given the other player's Bayesian strategy. The analysis is as follows:

$R$  is dominant for player  $t_1 = A$ , so he must play this in equilibrium.

Thus type  $t_2 = q$  assigns probability  $\frac{1}{2}$  to facing a type  $A$  who chooses  $R$ . So whether type  $K$  plays  $R$  or  $F$ , type  $q$ 's unique best response is  $f$ ; she must play this in equilibrium.

Continuing, type  $t_1 = K$  assigns probability  $\frac{1}{2}$  to facing a type  $q$  who plays  $f$ . So whether type  $a$  plays  $c$  or  $f$ , type  $K$ 's unique best response is  $R$ , and so he must play this in equilibrium.

Hence type  $t_2 = a$  assigns probability  $\frac{1}{2}$  to facing a type  $K$  who plays  $R$ . So whether type  $Q$  plays  $R$  or  $F$ , type  $a$ 's unique best response is  $c$ , and so she must play this in equilibrium. All of this is as stated in (27).

If we take these equilibrium actions as given, then as we explain in detail below, the remaining types,  $t_1 = Q$  and  $t_2 = k$ , are effectively playing the following  $2 \times 2$  game:

		$t_2 = k$	
		$c$	$f$
$t_1 = Q$	$R$	$-2, 0$	$-\frac{1}{2}, -1$
	$F$	$-1, -\frac{1}{2}$	$-1, 0$

The unique Nash equilibrium of this game has  $t_1 = Q$  play  $\frac{1}{3}R + \frac{2}{3}F$  and has  $t_2 = k$  play  $\frac{1}{3}c + \frac{2}{3}f$ , as stated in (27).

(How do we obtain the payoffs in the table above? First consider the payoffs of  $t_1 = Q$ . This type is equally likely to face a type  $a$ , who we know will play  $c$ , or a type  $k$ . Suppose type  $Q$  raises. If type  $k$  plays  $c$ , then type  $Q$  gets  $-2$  for sure. If instead type  $k$  plays  $f$ , then type  $Q$  will get  $-2$  if player 2 is type  $a$  and  $1$  if player 2 is type  $k$ , for an expected payoff of  $\frac{1}{2}(-2) + \frac{1}{2}(1) = -\frac{1}{2}$ . Finally, if type  $Q$  folds, it gets  $-1$  for sure.

Now consider the payoffs of  $t_2 = k$ . This type is equally likely to face a type  $A$ , who we know will play  $R$ , or a type  $Q$ . Suppose type  $k$  calls. If type  $Q$  plays  $R$ , then type  $k$ 's expected payoff is  $\frac{1}{2}(-2) + \frac{1}{2}(2) = 0$ ; if type  $Q$  plays  $F$ , then type  $k$ 's expected payoff is  $\frac{1}{2}(-2) + \frac{1}{2}(1) = -\frac{1}{2}$ . Now suppose type  $k$  folds. If type  $Q$  plays  $R$ , then type  $k$  gets  $-1$  for sure. If type  $Q$  plays  $F$ , then type  $k$ 's expected payoff is  $\frac{1}{2}(-1) + \frac{1}{2}(1) = 0$ .)  $\diamond$

*Example 4.7.* In the two-player Bayesian game  $BG$ , player  $i$ 's type  $t_i$ , representing his level of productivity, takes values in the finite set  $T_i \subset \{1, 2, \dots\}$ . Types are drawn according to the prior distribution  $p$  on  $T = T_1 \times T_2$ . After types are drawn, each player chooses to be In or Out of a certain project. If player  $i$  chooses Out, his payoff is 0. If player  $i$  chooses In and player  $j$  chooses Out, player  $i$ 's payoff is  $-c$ , where  $c > 0$  is the cost of participating in the project. Finally, if both players choose In, then player  $i$ 's payoff is  $t_i t_j - c$ . Thus, a player who chooses In must pay a cost, but the project only succeeds if both players choose In; in the latter case, the per-player benefit of the project is the product of the players' productivity levels.

Now suppose that the type sets are  $T_1 = \{3, 4, 5\}$  and  $T_2 = \{4, 5, 6\}$ , and that the prior distribution  $p$  is given by the table below.

		$t_2$		
		4	5	6
$t_1$	3	.2	.1	0
	4	.1	.1	.1
	5	0	.1	.3

What are the Bayesian equilibria of  $BG$  when  $c = 15$ ?

Use  $\sigma_i^{t_i}$  to denote the probability with which a player  $i$  of type  $t_i$  chooses In.

When  $c = 15$ , there are three Bayesian equilibria. In all three, types  $t_1 = 3$ ,  $t_1 = 4$ ,  $t_2 = 4$ , and  $t_2 = 5$  choose Out. The possible behaviors of the remaining two types are

$$\begin{aligned}\sigma_1^5 = \sigma_2^6 = 1 & \text{ (both types choose In),} \\ \sigma_1^5 = \sigma_2^6 = 0 & \text{ (both types choose Out),} \\ \sigma_1^5 = \sigma_2^6 = \frac{2}{3} & \text{ (both types mix).}\end{aligned}$$

This is shown as follows: The highest benefit that type  $t_1 = 3$  could obtain from playing In is  $\frac{2}{3} \cdot 12 + \frac{1}{3} \cdot 15 = 13$ . Since this is less than  $c = 15$ , this type stays out in any equilibrium.

Proceeding sequentially, we argue that types  $t_2 = 4$ ,  $t_1 = 4$ , and  $t_2 = 5$  play Out. The highest benefit type  $t_2 = 4$  could obtain from playing In is  $\frac{1}{3} \cdot 16 = 5\frac{1}{3} < 15$ , so this type plays Out; thus, the highest benefit type  $t_1 = 4$  could obtain from playing In is  $\frac{1}{3} \cdot 20 + \frac{1}{3} \cdot 24 = \frac{44}{3} = 14\frac{2}{3} < 15$ , so this type plays Out; and thus, finally, the highest benefit type  $t_2 = 5$  could obtain from playing In is  $\frac{1}{3} \cdot 25 = 8\frac{1}{3} < 15$ , so this type plays Out.

Conditional on this behavior for the low and middle productivity types, the remaining types,  $t_1 = 5$  and  $t_2 = 6$ , are playing a  $2 \times 2$  coordination game, namely

		$t_2 = 6$	
		In	Out
$t_1 = 5$	In	$7\frac{1}{2}, 7\frac{1}{2}$	$-15, 0$
	Out	$0, -15$	$0, 0$

where  $7.5 = \frac{3}{4} \cdot 30 - 15$ . It is an equilibrium for both to be In, and also for both to be Out. For type  $t_1 = 5$  to be willing to randomize, his expected benefit to being in must equal  $c = 15$ ; thus  $\frac{3}{4} \cdot \sigma_2^6 \cdot 30 = 15$ , implying that  $\sigma_2^6 = \frac{2}{3}$ . Virtually the same calculation shows that  $\sigma_1^5 = \frac{2}{3}$  as well.  $\diamond$

## 4.4 More on Interpretation

### *Ex ante and interim interpretations of Bayesian games*

Observation 4.3 says that a Bayesian games  $BG$  with a common prior can be represented as an extensive form game  $\Gamma_{BG}$  that begins with a random draw of types from the common prior distribution. There are applications in which this is just what happens: that is, there really is an ex ante stage at which the players have identical information, followed by a moment at which each player receives his private information. In such applications, we say that the Bayesian game  $BG$  admits an *ex ante interpretation*. Card games (Examples 4.4

and 4.6) are one example of this sort. Another example is that of a *mineral rights auction*: Bidders' begin with publicly available information about possible amounts of a mineral in a given tract of land. Then each bidder sends its own geologist to the tract, and the signals reported to each bidder by its geologist are independent conditional on the tract's actual quality. (See Section 4.6 for an example.) The ex ante interpretation only applies in some applications. But when it does, there is really no difference at all between  $BG$  and  $\Gamma_{BG}$ , and the Bayesian game is really just a particular sort of game with asymmetric information.

On the other hand, many applications of Bayesian games in economics, even ones with a common prior, do not have an ex ante stage. For instance, in the joint production game (Example 4.7), the payoff-relevant information in a player's type is his level of productivity, which is a characteristic that we'd expect the player to have known long before the game is played. In cases like this, the Bayesian game should be given an *interim interpretation*. In this interpretation, each player  $i$  knows from the very start what his type  $t_i$ , and the other types in  $T_i$  are there to allow us to capture other players' uncertainty about player  $i$ 's type. Now in defining Bayesian equilibrium, we require optimal play by all types of player  $i$ : not just his actual type  $t_i$ , but also all other types  $\hat{t}_i$ . This is the only sensible way to define equilibrium in this setting:  $i$ 's opponents don't know what type  $i$  is, so in equilibrium they should anticipate optimal behavior from each possible type; if optimality is not required for type  $\hat{t}_i$ , it is as if  $i$ 's opponents know that this is not  $i$ 's actual type.

To start the section, we said that a strategic interaction has *incomplete information* if when play begins, players possess payoff-relevant information that is not common knowledge. This is so when play begins at the *interim* stage. Thus a Bayesian game with the interim interpretation is a convenient way of representing a strategic interaction with incomplete information. In Section 4.8, we explain why *any* such interaction can be represented by a Bayesian game.

### *Common priors*

In applications of Bayesian games in which the ex ante interpretation holds, it is natural to assume a common prior.

Even in applications that use the interim interpretation, it may still be the case that all players' beliefs can be derived from a common prior. In this case, the Bayesian game  $BG$  is still formally equivalent to the extensive form game  $\Gamma_{BG}$ , even though the ex ante stage in  $\Gamma_{BG}$  is fictional.

There are situations in which assuming a common prior is innocuous. If there are two

players and only one has private information, then the prior is just a description of his opponents' beliefs. Also, if there are two players and neither player's first-order beliefs depend on his type, then beliefs can be derived from a common prior with independent types.

In information economics, the vast majority of models assume a common prior, with or without independent types. There are practical reasons for doing this. Models with a common prior are less complicated to write down. Also, in applications that do not assume a common prior, equilibria can be driven by different players having irreconcilable beliefs about payoff-relevant variables. In most applications this is viewed as undesirable.

But there are settings in which common priors lead to surprising results, most notably the *no trade theorem*: Milgrom and Stokey (1982) consider risk-averse agents with a common prior and a joint allocation that is Pareto efficient relative to this prior. They show that if the agents subsequently receive private information, it cannot be common knowledge that mutually beneficial recontracting is possible. The reason is that if one agent suggests a new contract that he is willing to sign, this fact itself provides information about what this agent learned, making the contract look undesirable to his opponent. If a common prior is not assumed, then this impossibility result no longer holds—see Morris (1994). For further discussion of these and related results, see Morris (1995) and Samuelson (2004).

Despite the prevalence of the common priors in applications, and despite some assertions to the contrary in the literature, there is no reason why a common prior must be assumed under the interim interpretation of Bayesian games. For discussions, see Dekel and Gul (1997), Gul (1998), and Section 4.8 below.

## 4.5 Applications

### 4.5.1 Auctions

Auction models not only have great practical importance, but are also a rich source of examples of Bayesian games with continuous type and action sets. The analyses here introduce ideas that will be useful in later sections.

#### *Auctions in environments with symmetric independent private values*

Basic references: Vickrey (1961, 1962), Myerson (1981), Riley and Samuelson (1981). A good textbook presentation is Krishna (2002, ch. 2).

We consider the following auction environment: The set of players is  $\mathcal{P} = \{1, \dots, n\}$ . Player  $i$ 's valuation for the good being sold,  $v_i \in [0, 1]$ , is his private information. Player  $i$ 's utility

is  $v_i - p$  if he obtains the good and pays  $p$ , and his utility is  $-p$  if he does not obtain the good and pays  $p$ .

Players' valuations are drawn independently of one another, and they are symmetric, in that they are all drawn from the same distribution on  $[0, 1]$ . The cdf and pdf of this distribution are denoted  $F$  and  $f$ . We assume that  $f$  is positive, which implies that  $F$  is increasing.

In auction theory, a *sealed-bid* auction is one in which the bidders submit bids simultaneously. A pure Bayesian strategy for player  $i$ , sometimes called a *bidding strategy*, is a map  $b_i: [0, 1] \rightarrow [0, 1]$ . We will only consider pure strategies.

We consider two sealed-bid auction formats.

In a *second-price auction*, the player who submitted the highest bid wins the object and pays the amount of the second-highest bid. If more than one player submits the highest bid, the object is assigned to one of these players at random, who pays the amount of the bid.

**Proposition 4.8.** *In a second-price auction,  $\beta_i^*(v_i) = v_i$  (i.e., always bidding one's valuation) is a weakly dominant strategy.*

*Proof.* Let  $v_i$  be player  $i$ 's valuation, let  $b_i \neq v_i$  be a possible alternative bid for player  $i$  and let  $b_{-i}^* = \max_{j \neq i} b_j$  be the highest of the opponents' bids. First suppose that  $b_i > v_i$ . If  $b_{-i}^* > b_i$ , then player  $i$  loses whether he bids  $v_i$  or  $b_i$ . If  $b_{-i}^* < v_i$ ,  $i$  wins and pays  $b_{-i}^*$  either way. But if  $b_{-i}^* \in (v_i, b_i)$ , then player  $i$  loses if he bids  $v_i$  but wins if he bids  $b_i$ ; in the latter case, he pays  $b_{-i}^*$ , yielding a payoff of  $v_i - b_{-i}^* < 0$ . Hence, whenever  $v_i$  and  $b_i$  perform differently,  $v_i$  does better. In other words,  $v_i$  weakly dominates  $b_i$ . Similar reasoning shows that this is also true for  $b_i < v_i$ , and so  $\beta_i^*(v_i) = v_i$  is a weakly dominant strategy. ■

Remarks:

- (i) The proof of Proposition 4.8 is identical to the one for the case in which the values are common knowledge. Thus this result does not depend on the assumption that values are independent or identically distributed, or that they have full support on  $[0, 1]$ , but it does depend on the private values assumption.
- (ii) While having all player's bid their valuations is a Bayesian equilibrium in weakly dominant strategies, there are many other Bayesian equilibria that use weakly dominated strategies. For instance, it is a Bayesian equilibrium for player 1 to always bid 1 and the others to always bid 0. For a complete analysis, see Blume and Heidhues (2004).



In a *first-price auction*, the player who submitted the highest bid wins the object and pays the amount of his bid. If more than one player submits the highest bid, the object is assigned to one of these players at random, who pays the amount of the bid.

We only consider symmetric equilibria:  $b_i^*(\cdot) = b^*(\cdot)$  for all  $i \in \mathcal{P}$ . (It can be shown that there are no asymmetric equilibria.)

Let  $V_i$  be a random variable representing player  $i$ 's valuation at the ex ante stage. Let the order statistic  $V_{(k)}^m$  be the  $k$ th lowest of  $V_1, \dots, V_m$ . Then  $V_{(n-1)}^{n-1} = \max\{V_1, \dots, V_{n-1}\}$  has cdf

$$\mathbb{P}(V_{(n-1)}^{n-1} \leq v) = \mathbb{P}(V_1 \leq v, \dots, V_{n-1} \leq v) = \mathbb{P}(V_1 \leq v) \times \dots \times \mathbb{P}(V_{n-1} \leq v) = F(v)^{n-1}.$$

In what follows, it is convenient to write  $F^{n-1}(v)$  for  $F(v)^{n-1}$ .

**Proposition 4.9.** *Any symmetric equilibrium  $b^*$  is increasing and differentiable.*

See Matthews (1995) for a proof. We will see proofs of similar results later, but this proof has a number of picky details.

**Proposition 4.10.** *The unique symmetric equilibrium is  $b^*(v) = \mathbb{E}(V_{(n-1)}^{n-1} | V_{(n-1)}^{n-1} \leq v)$  for  $v > 0$  and  $b^*(0) = 0$ .*

That is, each player bids the expected value of the maximum of his opponents' valuations, conditional on this maximum not exceeding his own valuation.

*Proof.* (Necessity.) Let  $b$  be a symmetric equilibrium. By the previous proposition  $b$  is increasing and differentiable. If a player's opponents choose this strategy, then the expected payoff of a player of type  $v$  who bids as though he were type  $\hat{v}$  is

$$(28) \quad \begin{aligned} (v - b(\hat{v})) \mathbb{P}\left(\max_{j \neq i} b(V_j) < b(\hat{v})\right) &= (v - b(\hat{v})) \mathbb{P}\left(\max_{j \neq i} V_j < \hat{v}\right) \\ &= (v - b(\hat{v})) F^{n-1}(\hat{v}), \end{aligned}$$

where the first equality uses the fact that  $b$  is increasing.

Because  $b$  is a symmetric equilibrium, the player maximizes his payoff by choosing  $\hat{v} = v$ . Thus when  $v \in (0, 1)$ , the following first order condition must hold:

$$\frac{d}{d\hat{v}} \left( b(\hat{v}) F^{n-1}(\hat{v}) \right) \Big|_{\hat{v}=v} = v \cdot \left( \frac{d}{d\hat{v}} F^{n-1}(\hat{v}) \right) \Big|_{\hat{v}=v}.$$

Integrating both sides yields from 0 to  $\tilde{v}$  (using the continuity of  $b$  and  $F^{n-1}$ ) yields

$$b(\tilde{v})F^{n-1}(\tilde{v}) - b(0)F^{n-1}(0) = \int_0^{\tilde{v}} w \, dF^{n-1}(w).$$

Since  $F(0) = 0$ , changing the names of the variables and rearranging yields

$$b(v) = \frac{\int_0^v w \, dF^{n-1}(w)}{F^{n-1}(v)} = \mathbb{E}(V_{(n-1)}^{n-1} \mid V_{(n-1)}^{n-1} < v) = b^*(v).$$

This establishes necessity for  $v \in (0, 1)$ . Since  $b$  is increasing and  $b^*$  is continuous, we also have  $b(0) = 0 = b^*(0)$  and  $b(1) \geq b^*(1)$ . In fact, it must be that  $b(1) = b^*(1)$ , since this bid wins the object for sure.

(Sufficiency.) We need to show that  $b^*$  is a symmetric equilibrium. Suppose that a player's opponents choose this strategy. Clearly the player should never choose a bid above  $b^*(1)$ , since bidding  $b^*(1)$  already ensures that he wins the object. Thus since  $b^*$  has range  $[0, b^*(1)]$ , it is enough to show that a player of type  $v$  is at least as well off choosing  $b^*(v)$  as  $b^*(\hat{v})$  for any  $\hat{v} \in [0, 1]$ .

Substituting  $b^*(\hat{v})$  into (28) shows that type  $v$ 's payoff for placing bid  $b^*(\hat{v})$  is

$$\left( v - \frac{\int_0^{\hat{v}} w \, dF^{n-1}(w)}{F^{n-1}(\hat{v})} \right) F^{n-1}(\hat{v}) = vF^{n-1}(\hat{v}) - \int_0^{\hat{v}} w \, dF^{n-1}(w) = \int_0^{\hat{v}} (v - w) \, dF^{n-1}(w).$$

This is clearly maximized at  $\hat{v} = v$ . ■

Remark: The proof of necessity considered what would happen if a player of type  $v$  acted as though he were of some other type  $\hat{v}$ . This trick is closely related to the *revelation principle* for mechanism design, a fundamental result that we present in Section 7.1.4.

We now compare the performances of the two auction formats.

**Proposition 4.11.** *In both the weakly dominant equilibrium of the second price auction and the symmetric equilibrium of the first price auction,*

- (i) *the expected payment of a bidder of type  $v > 0$  is  $\mathbb{P}(V_{(n-1)}^{n-1} < v) \cdot \mathbb{E}(V_{(n-1)}^{n-1} \mid V_{(n-1)}^{n-1} \leq v)$ , and*
- (ii) *the seller's expected revenue is  $\mathbb{E}V_{(n-1)}^n$ .*

The expected payment in statement (i) is the product of (a) the probability that the maximum of a player's opponents' valuations is lower than  $v$  and (b) the equilibrium bid for

valuation  $v$  in a first price auction. The expected revenue in statement (ii) is the expected value of the second-highest of the players' valuations.

Proposition 4.11 is an instance of a much more general phenomenon: see Theorems 7.11 and 7.12.

*Proof.* (i) For the first price auction, the statement is immediate from Proposition 4.10. For the second price auction, since  $\beta_i^*(v) = v$ , and using the fact that ties occur with probability zero, player  $i$ 's expected payment when he is of type  $v$  is

$$\begin{aligned} & \mathbb{P}\left(\max_{j \neq i} \beta_j^*(V_j) < \beta_i^*(v)\right) \mathbb{E}\left(\max_{j \neq i} \beta_j^*(V_j) \mid \max_{j \neq i} \beta_j^*(V_j) < \beta_i^*(v)\right) \\ &= \mathbb{P}\left(\max_{j \neq i} V_j < v\right) \mathbb{E}\left(\max_{j \neq i} V_j \mid \max_{j \neq i} V_j < v\right) \\ &= \mathbb{P}\left(V_{(n-1)}^{n-1} < v\right) \mathbb{E}\left(V_{(n-1)}^{n-1} \mid V_{(n-1)}^{n-1} < v\right). \end{aligned}$$

(ii) In a symmetric private values environment, the seller's expected revenue can be determined from the expected payments of each bidder type: if the expected payment of type  $v$  is  $\bar{r}(v)$ , then the seller's expected revenue is  $n \mathbb{E}(\bar{r}(V_i)) = n \int_0^1 \bar{r}(v) dF(v)$ . Part (i) says that the function  $\bar{r}(\cdot)$  is the same in the given equilibria of both auction formats, so the expected revenue is the same as well.

To obtain an expression for this expected revenue, it is enough to consider the second price auction. In the weakly dominant equilibrium of this auction, the seller's revenue is the second highest valuation among the  $n$  agents. In ex ante terms this is  $V_{(n-1)}^n$ , so the expected revenue is  $\mathbb{E}V_{(n-1)}^n$ . ■

In contrast to the environment considered here, many auctions arising in practice feature *interdependent values*, meaning that the signals of a bidder's opponents would give the bidder a clearer idea of his value for the good. The key early reference on standard auctions in environments with interdependent values is Milgrom and Weber (1982); see Krishna (2002) and Milgrom (2004) for textbook treatments. We consider an example using a "non-standard" auction format in the next section.

#### 4.5.2 The revelation principle for normal form games

Suppose that the players in a normal form game are sent correlated, payoff-irrelevant signals before the start of play. This augmentation of the normal form game generates a Bayesian game. Every Bayesian strategy profile in this game induces a distribution over

action profiles in the normal form game. According to the *revelation principle for normal form games* (Proposition 4.12 below), the distribution over action profiles induced by any Bayesian equilibrium is a correlated equilibrium of the normal form game. Thus correlated equilibrium describes the distributions over action profiles that are possible if play of a normal form game is preceded by the observation of *any* correlated, payoff-irrelevant signals. See Myerson (1991, Chapter 6) for further discussion.

Let  $G = \{\mathcal{P}, \{A_i\}_{i \in \mathcal{P}}, \{u_i\}_{i \in \mathcal{P}}\}$  be a normal form game. Let  $BG = \{\mathcal{P}, \{A_i\}_{i \in \mathcal{P}}, \{T_i\}_{i \in \mathcal{P}}, p, \{u_i\}_{i \in \mathcal{P}}\}$  be a Bayesian game with common prior distribution  $p \in \Delta T$ , and with the same action sets  $A_i$  and utility functions  $u_i : A \rightarrow \mathbb{R}$  as  $G$ . Note that  $u_i$  does not condition on players' types.

**Proposition 4.12** (The revelation principle for normal form games).

Let  $s^*$  be a Bayesian equilibrium of  $BG$ , and let  $\rho^* \in \Delta A$  be the correlated strategy induced by  $s^*$ :

$$\rho^*(a) = \sum_{t: s^*(t)=a} p(t).$$

Then  $\rho^*$  is a correlated equilibrium of  $G$ .

*Proof.* Assume without loss of generality that  $p(t_i) = \sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i}) > 0$  for all  $t_i \in T_i$  and  $i \in \mathcal{P}$ . Strategy profile  $s^*$  is a Bayesian equilibrium of  $BG$  if and only if

$$(29) \quad \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) u_i(s_i^*(t_i), s_{-i}^*(t_{-i})) \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) u_i(a'_i, s_{-i}^*(t_{-i}))$$

for all  $a'_i \in A_i$ ,  $t_i \in T_i$ , and  $i \in \mathcal{P}$ .

Multiplying (29) through by  $p(t_i)$  shows that it is equivalent to

$$(30) \quad \sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i}) u_i(s_i^*(t_i), s_{-i}^*(t_{-i})) \geq \sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i}) u_i(a'_i, s_{-i}^*(t_{-i})).$$

Let  $T_i^*(a_i) = \{t_i \in T_i : s_i^*(t_i) = a_i\}$  be the set of player  $i$  types who choose action  $a_i$  under strategy  $s^*$ . Then

$$\sum_{t_i \in T_i^*(a_i)} \sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i}) u_i(s_i^*(t_i), s_{-i}^*(t_{-i})) = \sum_{a_{-i} \in A_{-i}} \rho^*(a_i, a_{-i}) u_i(a_i, a_{-i}).$$

Therefore, fixing  $a_i$  and summing inequality (30) over  $t_i \in T_i^*(a_i)$ , we find that

$$(31) \quad \sum_{a_{-i} \in A_{-i}} \rho^*(a_i, a_{-i}) u_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \rho^*(a_i, a_{-i}) u_i(a'_i, a_{-i})$$

for all  $a'_i \in A_i$  and  $i \in \mathcal{P}$ . (31) says that  $\rho^*$  is a correlated equilibrium of  $G$ . ■

The proof is easy to explain in words. Since  $s^*$  is a Bayesian equilibrium, a player  $i$  of type  $t_i$  gets at least as high a payoff from playing action  $s_i^*(t_i)$  as from any other action given that others follow their equilibrium strategies. Interpreting the signals in  $T_i^*(a_i)$  as a recommendation to play  $a_i$ , it follows directly that when player  $i$  gets this recommendation, it is optimal for him to obey given that others obey their recommendations. This is the definition of correlated equilibrium.

## 4.6 Ex Post Equilibrium

In a Bayesian equilibrium, each type  $t_i$  of player  $i$  chooses an action that maximizes his expected payoffs given his opponents' Bayesian strategies  $s_{-i}$ . These expected payoffs (22) depend on type  $t_i$ 's beliefs  $p_i(\cdot|t_i)$  about his opponents' types conditional on his own type. For a more demanding equilibrium concept, one can still assume that player  $i$  correctly anticipates his opponents' Bayesian strategies  $s_{-i}$ , but can require that the action he chooses when his type is  $t_i$  is optimal for *any* realization of his opponents' types (and so for all possible beliefs about their types). This is not a notion of dominance, because equilibrium knowledge is still used; what is dropped is all information about opponents' types. This relaxation is particularly important in auction and mechanism design: when an auction with an ex post equilibrium can be constructed, there is no need for the designer to know the players' beliefs—see Sections 7.5 and 7.6.

Pure Bayesian strategy profile  $s$  is an *ex post equilibrium* if

$$(32) \quad u_i((s_i(t_i), s_{-i}(t_{-i})), t) \geq u_i((\hat{a}_i, s_{-i}(t_{-i})), t) \quad \text{for all } t \in T, \hat{a}_i \in A_i, \text{ and } i \in \mathcal{P}.$$

The name “ex post equilibrium” refers to the fact that if there were an ex post stage at which all players' types are revealed, then no player of any type would prefer to switch strategies at that stage. (This is typically not true in a Bayesian equilibrium. For an example, consider the Bayesian equilibrium of a first price auction (Proposition 4.10.)) Ex post equilibrium is sometimes called *belief-free equilibrium*, a reference to the fact that beliefs  $p_i(\cdot|t_i)$  are irrelevant in definition (32). (This term is more commonly used in the context of repeated games with private monitoring, where it means something different.) For either of these reasons, ex post equilibrium is well-defined in a “pre-Bayesian game” in which beliefs are not specified at all.

### *An auction environment with interdependent values*

In auction environments with *interdependent values*, the signals of a bidder's opponents would give the bidder himself a clearer idea of his value for the good.

Consider the following environment: There are  $n$  bidders for a good. Each bidder  $i$  receives a private signal  $t_i \in [0, 1]$  that provides information about the good's properties. When the profile of signals is  $t = (t_1, \dots, t_n)$ , player  $i$ 's valuation for the good is

$$(33) \quad v_i(t) = t_i + \gamma \sum_{j \neq i} t_j \text{ for some } \gamma \in [0, 1].$$

Player  $i$ 's utility for obtaining the good when the type profile is  $t$  and paying  $p$  is  $v_i(t) - p$ , and his utility for not obtaining the good and paying  $p$  is  $-p$ .

When  $\gamma = 0$ , (33) is a private values environment. When  $\gamma = 1$ , (33) is a *common value* environment: conditional on any profile of signals  $t$ , everyone's (expected) valuation for the good is the same. Intermediate values of  $\gamma$  can be used, e.g., when each player's signal is most informative about the properties of the good he cares about most. (Note: Some authors use the term "common values" to refer to all environments with interdependent values.)

To complete the definition of the Bayesian game, we need to specify the joint distribution of signals. But because we are considering ex post equilibrium, what we specify plays no role in the analysis.

(The mathematically simplest specification is that of independent signals. This assumption can be justified when the "good" being sold is a bundle, and each agent's signal concerns a distinct item in the bundle. But economic environments with interdependent values typically have dependent signals. For instance, in a mineral rights auction, after conditioning on the quality of the tract of land the signals may be independent. However, the quality is precisely what the bidders do not know, and if we do not condition on this, the signals are dependent. A standard model for this case is that of so-called *affiliated* signals; see Milgrom and Weber (1982).)

Maskin (1992) introduced the following generalization of a second-price auction for the environment above: Each agent makes a report  $\hat{t}_i \in [0, 1]$ . If bidder  $i$  has the unique highest report, then he receives the object, and he pays the highest of his opponents' reports plus  $\gamma$  times the sum of his opponents' reports. In the event of a tied highest report, one of the bidders who made this report is selected at random to receive the object and to make the payment described above.

**Proposition 4.13.** *In this auction, it is an ex post equilibrium for all bidders to report truthfully.*

*Proof.* Suppose that the type profile is  $t$  and that player  $i$ 's opponents report truthfully. If bidder  $i$ 's report is higher than any other agent's report, then agent  $i$ 's payoff is

$$\left( t_i + \gamma \sum_{j \neq i} t_j \right) - \left( \max_{k \neq i} t_k + \gamma \sum_{j \neq i} t_j \right) = t_i - \max_{k \neq i} t_k.$$

(Notice that this payoff from being allocated the good does not depend on agent  $i$ 's report.) Thus the fact that it is always optimal for player  $i$  to tell the truth follows directly from the analysis of the second price auction. (To check your understanding, make sure you see why truthful reporting is not a weakly dominant strategy. This is related to how we used the supposition that player  $i$ 's opponents report truthfully.) ■

## 4.7 Correlated Types and Higher-Order Beliefs

In the examples we have seen so far, we have named a player's types after the payoff-relevant information he possesses: his card, his productivity level, his valuation for a good. But as we noted at the end of Section 4.1, when types are not independent, any privately known elements of a player's beliefs are also part of his type, and these beliefs may concern other players' beliefs about other players' beliefs. . .

*Example 4.14.* In Example 4.7, the prior distribution  $p$  was

		$t_2$		
		4	5	6
	3	.2	.1	0
$t_1$	4	.1	.1	.1
	5	0	.1	.3

Evidently,  $p_2(t_1 = 3 | t_2 = 4) = \frac{2}{3}$ ,  $p_2(t_1 = 3 | t_2 = 5) = \frac{1}{3}$ , and  $p_2(t_1 = 3 | t_2 = 6) = 0$ . Thus

$$p_1(\text{player 2 assigns probability } \frac{1}{3} \text{ to } t_1 = 3 | t_1 = 3) = p_1(t_2 = 5 | t_1 = 3) = \frac{1}{3}.$$

(Soon we will introduce notation for referring to beliefs about beliefs—see Section 4.8.) ◇

When types are correlated, one needs to be careful about specifying beliefs. In some settings, doing so naively can have dramatic and unexpected consequences. The leading

example here is the possibility of full surplus extraction in mechanism design. We explore this topic in detail in Section 7.6.

Here we focus on another consequence of types describing privately-known beliefs. In the next example, there is only one piece of payoff-relevant information, and it is revealed to player 1 at the interim stage. But both players have many types, reflecting each player's beliefs about the other player's beliefs. . . about what player 1 knows.

*Example 4.15. The Email Game (Rubinstein (1989)).*

In this game, there is probability  $\frac{2}{3}$  that the payoff matrix is  $G_L$ , and probability  $\frac{1}{3}$  that the payoff matrix is  $G_R$ . In  $G_L$ ,  $A$  is a strictly dominant strategy.  $G_R$  is a coordination game in which the players want to coordinate on action  $B$ . If they coordinate on the wrong action, both get 0; if they miscoordinate, the player who chose  $B$  gets punished.

$G_L$ :	2	$G_R$ :	2
	A	B	
1	A	2, 2	0, -3
	B	-3, 0	-1, -1
	A	B	
1	A	0, 0	0, -3
	B	-3, 0	2, 2

Only player 1 observes whether the payoff matrix is  $G_L$  or  $G_R$ . If it is  $G_R$ , player 1's computer automatically sends a message to player 2's computer, a message that arrives with probability  $1 - \varepsilon$ , where  $\varepsilon \in (0, 1)$ . If player 2's computer receives this message, it automatically sends a confirmation to player 1's computer, a confirmation that arrives with probability  $1 - \varepsilon$ . If player 1's computer receives this confirmation, then it automatically sends another confirmation. . . Each player knows how many messages his own computer sent.

Let  $m \in \{0, 1, 2, \dots\}$  be the total number of messages sent.

Then the probability that no messages are sent is  $\pi_0 = \frac{2}{3}$ , and the probability that  $m > 0$  messages are sent is  $\pi_m = \frac{1}{3} \cdot \varepsilon(1 - \varepsilon)^{m-1}$ .

A player's type  $t_i \in T_i = \{0, 1, 2, \dots\}$  is the number of messages he sends. The player's types are related to the total number of messages sent by  $t_1 = \lceil \frac{m}{2} \rceil$  and  $t_2 = \lfloor \frac{m}{2} \rfloor$ .

The common prior on  $T = T_1 \times T_2$  is determined by these relations and  $\pi$ :



		$t_2$				
		0	1	2	3	
$t_1$	0	$\frac{2}{3}$	0	0	0	...
	1	$\frac{1}{3}\varepsilon$	$\frac{1}{3}(1-\varepsilon)\varepsilon$	0	0	...
	2	0	$\frac{1}{3}(1-\varepsilon)^2\varepsilon$	$\frac{1}{3}(1-\varepsilon)^3\varepsilon$	0	...
	3	0	0	$\frac{1}{3}(1-\varepsilon)^4\varepsilon$	$\frac{1}{3}(1-\varepsilon)^5\varepsilon$	...
		$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

**Proposition 4.16.** *If  $\varepsilon > 0$ , the unique Bayesian equilibrium has both players play A regardless of type.*

If  $\varepsilon > 0$  is small, and the payoff matrix turns out to be  $G_R$ , it is very likely that both players know that the payoff matrix is  $G_R$ . But the players still play A, even though (B, B) is a strict equilibrium in  $G_R$ .

*Proof.* If the payoff matrix is  $G_L$ , player 1 knows this, and plays A, which is dominant for him in this matrix.

Suppose that player 2's type is  $t_2 = 0$ .

Then her posterior probability that no messages were sent is

$$\frac{\pi_0}{\pi_0 + \pi_1} = \frac{\frac{2}{3}}{\frac{2}{3} + \frac{1}{3}\varepsilon} = \frac{2}{2 + \varepsilon} > \frac{2}{3}.$$

Since player 1 plays A when no messages are sent, player 2's expected payoff to playing A is more than  $\frac{2}{3} \times 2 = \frac{4}{3}$ , while her expected payoff to choosing B is less than  $\frac{2}{3} \times (-3) + \frac{1}{3} \times 2 = -\frac{4}{3}$ .

Therefore, when her type is  $t_2 = 0$ , player 2 should choose A.

Now suppose that player 1's type is  $t_1 = 1$ .

In this case, player 1 knows that the payoff matrix is  $G_R$ , and so that his payoff from playing A is 0.

His posterior probability that exactly 1 message was sent is

$$\frac{\pi_1}{\pi_1 + \pi_2} = \frac{\frac{1}{3}\varepsilon}{\frac{1}{3}\varepsilon + \frac{1}{3}(1-\varepsilon)\varepsilon} = \frac{1}{1 + (1-\varepsilon)} > \frac{1}{2}.$$

Therefore, player 1's expected payoff to playing B is less than  $\frac{1}{2} \times (-3) + \frac{1}{2} \times 2 = -\frac{1}{2}$ .

Therefore, when his type is  $t_1 = 1$ , player 1 should play A.

Continuing, suppose that player 2 is of type 1. In this case, *both* players know that the payoff matrix is  $G_R$ . But player 2 will assign probability greater than  $\frac{1}{2}$  to 2 messages having been sent, and so, knowing that player 1 plays  $A$  in this case, will play  $A$  herself.

And so on. ■

Remarks:

1. The previous example and the next one illustrate the idea of *contagion*, which refers to iterated dominance arguments that arise in certain Bayesian games with “chains” of correlation in the common prior distribution. These examples illustrate the important role that higher-order beliefs can play in equilibrium analysis. Versions of these examples have been used to model bank runs and currency crises—see Morris and Shin (2003).
2. The email game raises questions about what should be considered “almost common knowledge”, and the relationship between almost common knowledge and robustness of equilibrium. To address these questions, one needs to think about notions of “closeness” on type spaces. See Monderer and Samet (1989), Weinstein and Yildiz (2007), Dekel et al. (2006), Chen et al. (2010), and Section 4.9.

*Example 4.17. A global game (Carlsson and van Damme (1993)).*

Consider the following normal form game  $G(r)$ :

		2	
		$I$	$N$
1	$I$	$r, r$	$r - 1, 0$
	$N$	$0, r - 1$	$0, 0$

In this game, strategy  $I$  represents investing, and strategy  $N$  represents not investing. Investing yields a payoff of  $r$  or  $r - 1$  according to whether the player’s opponent invests or not. Not investing yields a certain payoff of 0. If  $r < 0$ , then the unique NE is  $(N, N)$ ; if  $r = 0$ , the NE are  $(N, N)$  and  $(I, I)$ ; if  $r \in (0, 1)$ , the NE are  $(N, N)$ ,  $(I, I)$ , and  $((1 - r)I + rN, (1 - r)I + rN)$ ; if  $r = 1$ , the NE are  $(N, N)$  and  $(I, I)$ ; and if  $r > 1$ , then the unique NE is  $(I, I)$ .

Now consider a Bayesian game  $BG$  in which payoffs are given by the above payoff matrix, but in which the value of  $r$  is the realization of a random variable that is uniformly distributed on  $[-2, 3]$ . In addition, each player  $i$  only observes a noisy signal  $t_i$  about the value of  $r$ . Specifically,  $t_i$  is defined by  $t_i = r + \varepsilon_i$ , where  $\varepsilon_i$  is uniformly distributed on  $[-\frac{1}{10}, \frac{1}{10}]$ , and  $r$ ,  $\varepsilon_1$ , and  $\varepsilon_2$  are independent of one another. This construction is known as a *global game*; see Carlsson and van Damme (1993) and Morris and Shin (2003).

Pure strategies for player  $i$  are of the form  $s_i: [-\frac{21}{10}, \frac{31}{10}] \rightarrow \{I, N\}$ . To define a pure Bayesian equilibrium, let  $u_i(a, r)$  be player  $i$ 's payoff under action profile  $a$  under payoff parameter  $r$  (which we can read from the payoff matrix). Then define the payoff to a player  $i$  of type  $t_i$  for playing action  $a_i$  against an opponent playing strategy  $s_j$  as

$$U_i(a_i, s_j, t_i) = \int_{t_j, r} u_i((a_i, s_j(t_j)), r) d\mu_i(t_j, r | t_i),$$

where  $\mu_i(\cdot | t_i)$  represents  $i$ 's beliefs about  $(t_j, r)$  when he is of type  $t_i$ . Then for  $s$  to be a pure Bayesian equilibrium, it must be that

$$s_i(t_i) \in \operatorname{argmax}_{a_i} U(a_i, s_j, t_i)$$

for all  $t_i$  and  $i \in \{1, 2\}$ .

Surprisingly, equilibrium behavior in BG is essentially unique.

**Proposition 4.18.** *In any Bayesian equilibrium of BG, player  $i$  strictly prefers not to invest when his signal is less than  $\frac{1}{2}$ , and strictly prefers to invest when his signal is above  $\frac{1}{2}$ .*

Even though the underlying normal form game has multiple strict equilibria when  $r \in (0, 1)$ , the possibility that one's opponent may have a dominant strategy, combined with uncertainty about others' preferences and beliefs, lead almost every type to have a unique equilibrium action. The choice of  $\varepsilon = \frac{1}{10}$  is not important here; any positive  $\varepsilon$  would lead to the same result.

*Proof.* Observe that if  $t_i \in [-\frac{19}{10}, \frac{29}{10}]$ , then player  $i$ 's posterior belief about  $r$  conditional on  $t_i$  is uniform on  $[t_i - \frac{1}{10}, t_i + \frac{1}{10}]$ , and his posterior belief about  $t_j$  conditional on  $t_i$  is the triangular distribution with support  $[t_i - \frac{1}{5}, t_i + \frac{1}{5}]$  (that is, the conditional density equals 0 at  $t_j = t_i - \frac{1}{5}$ , 5 at  $t_j = t_i$ , and 0 at  $t_j = t_i + \frac{1}{5}$ , and is linear on  $[t_i - \frac{1}{5}, t_i]$  and on  $[t_i, t_i + \frac{1}{5}]$ ).

In any Bayesian equilibrium, if the value of player  $i$ 's signal  $t_i$  is less than 0, then player  $i$  strictly prefers not to invest. Indeed, if the random variable  $R$  represents the payoff  $r$  at the ex ante stage, then  $\mathbb{E}(R | t_i) = t_i$  when  $t_i \in [-\frac{19}{10}, 0]$  and that  $\mathbb{E}(R | t_i) < -\frac{18}{20}$  when  $t_i < -\frac{19}{20}$ . It follows that if  $t_i < 0$ , the expected payoff to  $I$  is negative, and so that  $N$  is a strict best response for a player  $i$ .

Next, we show that if the value of player  $i$ 's signal is less than  $\frac{1}{20}$ , then player  $i$  strictly prefers not to invest. We have just seen that if  $t_2 < 0$ , then player 2 will play  $N$ . Now focus on the most demanding case, in which  $t_1 = \frac{1}{20}$ . In this case, player 1's posterior beliefs about player 2's type are described by a triangular distribution with support  $[-\frac{3}{20}, \frac{1}{4}]$ . Since the density of this distribution reaches its maximum of 5 at  $t_2 = \frac{1}{20}$ , the probability that player 1 assigns to player 2 having a type below 0 is the area of a triangle with base  $\frac{3}{20}$  and height  $\frac{3}{4} \cdot 5$ . This area is  $\frac{9}{32}$ ; by the previous paragraph, it is a lower bound on the probability that player 1 assigns to player 2 playing  $N$ . Therefore, player 1's expected

payoff to playing  $I$  is

$$U_1(I, s_2, \frac{1}{20}) = \frac{1}{20} - \mu_1(s_2(t_2) = N | t_1 = \frac{1}{20}) \leq \frac{1}{20} - \frac{9}{32} < 0.$$

It is thus a strict best response for player 1 to play  $N$  when  $t_1 = \frac{1}{20}$ , and *a fortiori* when  $t_1 < \frac{1}{20}$ .

In a similar fashion, we can show that in any Bayesian equilibrium, if the value of player  $i$ 's signal is less than  $\frac{1}{10}$ , then player  $i$  strictly prefers not to invest. Again we focus on the most demanding case, in which  $t = \frac{1}{10}$ . In this case player 1's beliefs about player 2's type are centered at  $\frac{1}{10}$ , and so have support  $[-\frac{1}{10}, \frac{3}{10}]$ . But we know from the previous paragraph that player 2 will play  $N$  whenever  $t_2 < \frac{1}{20}$ , so essentially the same calculation as before shows that player 1 again must assign a probability of at least  $\frac{9}{32}$  to player 1 playing  $N$ . Thus, since  $U_1(I, s_2, \frac{1}{10}) \leq \frac{1}{10} - \frac{9}{32} < 0$ , player 1's strict best response when  $t_1 \leq \frac{1}{10}$  is to play  $N$ .

Rather than iterate this argument further, let  $\tau$  be the supremum of the set of types that can be shown by such an iterative argument to play  $N$  in any equilibrium. If  $\tau$  were less than  $\frac{1}{2}$ , then since a player whose type is less than  $\tau$  will play  $N$ , a player whose type is  $\tau$  obtains an expected payoff of  $\tau - \frac{1}{2} < 0$  from playing  $N$ . By continuity, this is also true of a player whose type is slightly larger than  $\tau$ , contradicting the definition of  $\tau$ . We therefore conclude that  $\tau$  is at least  $\frac{1}{2}$ . This establishes that player  $i$  strictly prefers not to invest when his signal is less than  $\frac{1}{2}$ . A symmetric argument shows that player  $i$  strictly prefers to invest when his signal is above  $\frac{1}{2}$ . ■

In these examples, the introduction of a “small” amount of higher-order uncertainty to a baseline game with multiple equilibria leads to the selection of a unique equilibrium. But it is important to be aware that which equilibrium is selected is sensitive to the exact form that the higher-order uncertainty takes—see Section 4.9 below.

## 4.8 Incomplete Information, Hierarchies of Beliefs, and Bayesian Games

### *Informal explanation*

Harsanyi (1967–1968) initiated the study of *games of incomplete information*, where at the start of play, players possess payoff-relevant information that is not common knowledge. To describe such a situation directly, one starts by introducing a set of payoff-relevant states  $\Theta = \Theta_0 \times \Theta_1 \times \cdots \times \Theta_n$ . Taking the interim point of view, the actual payoff state is some  $\theta \in \Theta$ , and each player  $i$  knows his  $\theta_i$ , which is called his *payoff type*. The remaining  $\theta_0 \in \Theta_0$  is the *state of nature*, representing payoff-relevant information that no one observes. (We did not include states of nature in our definition of Bayesian games from Section 4.1, because one need not include it explicitly to define the notions of dominance and equilibrium introduced in Section 4.2. (Some of our examples actually had a state of

nature: for instance, the mineral rights auction described in Section 4.4, and the global game of Example 4.17.) However, states of nature play a more important role when defining rationalizability for Bayesian games—see Section 4.9.)

Harsanyi observed that to describe an environment with incomplete information, one needs to specify a player's beliefs about his opponents' payoff types, his beliefs about his opponents' beliefs, and so on. Why? Player  $i$ 's beliefs about opponents' payoff types  $\theta_{-i}$  help determine his preferred action, because his opponent's actions depend on their types. But this implies that player  $j$  should care about player  $i$ 's beliefs about  $\theta_{-i}$ , because player  $j$ 's preferred action depends on player  $i$ 's action, which depends on player  $i$ 's beliefs about  $\theta_{-i}$ . But this in turn implies that player  $k$  should care about player  $j$ 's beliefs about player  $i$ 's beliefs about  $\theta_{-i}$  . . . And so on.

The direct approach to addressing this has one specify infinite *hierarchies of beliefs* for each player. However, games with infinite hierarchies of beliefs are difficult to analyze. Harsanyi proposed *Bayesian games* as a tractable alternative. The fundamental notion here is that of a *type space*, a term which refers to the collection of all players' type sets and first-order beliefs. As usual, a player's type encodes both his payoff-relevant information (i.e., his payoff type) and information about his beliefs; and a player's first-order beliefs directly specify his beliefs about opponents' types and the state of nature as a function of his own type. The players' first-order beliefs implicitly determine their higher-order beliefs.

The simplification provided by Bayesian games comes from the fact that each player's beliefs are defined by a single probability distribution over opponents' types rather than by an infinite sequence of probability distributions. The games should be interpreted from a suitable interim perspective, in which player  $i$ 's type corresponds to the payoff type and hierarchy of beliefs we started with.

Every Bayesian game corresponds to a game with hierarchies of beliefs, because the players' first-order beliefs implicitly determine their higher-order beliefs. But can every game with hierarchies of beliefs be represented as a Bayesian game? This question was answered in the affirmative under some natural consistency conditions by Mertens and Zamir (1985) (see also Brandenburger and Dekel (1993)), through their construction of *universal type spaces*. (Aside: This construction may or may not lead to a Bayesian game satisfying the common prior assumption.)

In conclusion, using Bayesian games to model strategic environments with incomplete information is without loss of generality. Of course, complicated hierarchies of beliefs may need to be represented by type spaces that are much nastier than those in the examples we have seen. As a practical matter, Bayesian games appearing in most applications have fairly simple type sets, even though beliefs of arbitrarily high order may matter. But certain questions—in particular, questions about the robustness of predictions to slight misspecifications of beliefs—require one to work with universal type spaces; see Section 4.9.

### Some details

In the rest of this section, we assume that there are just two players, and that  $\Theta_0$  is a singleton (meaning that it can be ignored). The general case only requires more notation. We thus begin with the following objects:

$$\begin{aligned} \Theta &= \Theta_1 \times \Theta_2 && \text{set of payoff-type profiles} \\ A &= A_1 \times A_2 && \text{set of action profiles} \\ u_i &: A \times \Theta \rightarrow \mathbb{R} && \text{player } i\text{'s utility function} \end{aligned}$$

To describe a setting with incomplete information (and hence one understood from the interim perspective), we introduce *hierarchies of beliefs*  $\{\pi_i^k\}_{k=1}^\infty$  for each player  $i \in \{1, 2\}$ . The probability distribution  $\pi_i^k$  is player  $i$ 's  $k$ th-order belief. To explain what kind of object this is, we describe the sets in which each player's first three orders of beliefs live:

$$\begin{aligned} \pi_1^1 &\in \Delta\Theta_2 & \pi_2^1 &\in \Delta\Theta_1 \\ \pi_1^2 &\in \Delta(\Theta_2 \times \Delta\Theta_1) & \pi_2^2 &\in \Delta(\Theta_1 \times \Delta\Theta_2) \\ \pi_1^3 &\in \Delta(\Theta_2 \times \Delta\Theta_1 \times \Delta(\Theta_1 \times \Delta\Theta_2)) & \pi_2^3 &\in \Delta(\Theta_1 \times \Delta\Theta_2 \times \Delta(\Theta_2 \times \Delta\Theta_1)) \end{aligned}$$

Each player's first-order belief concerns his opponent's payoff type. Each player's second-order belief concerns his opponent's payoff type and first-order belief, and (naturally) allows for correlation between them. Each player's third-order belief concerns his opponent's payoff type, first-order belief, and second-order belief, again allowing for correlation among them. We can continue similarly to higher-order beliefs.

Notice that the marginal distribution of second-order belief  $\pi_1^2$  on  $\Theta_2$  can be viewed as a first-order belief, which we will call  $\hat{\pi}_1^1$ . Likewise, the marginal distribution of third-order belief  $\pi_1^3$  on  $\Theta_2 \times \Delta\Theta_1$  can be viewed as a second-order belief, which we will call  $\hat{\pi}_1^2$ . For player 1's hierarchy of beliefs to be sensible, different orders of belief should not contradict one another: we should have  $\hat{\pi}_1^1 = \pi_1^1$  and  $\hat{\pi}_1^2 = \pi_1^2$ . If these agreements across levels always hold for player  $i$ , that is, if  $\hat{\pi}_i^k = \pi_i^k$  for all  $k \geq 1$ , then player  $i$ 's hierarchy of beliefs is said to be *coherent*. These are the only hierarchies we should care about.

In addition, in a multi-player model of rational behavior, it is also natural to require player  $i$  to be certain that (i) player  $j$ 's hierarchy of beliefs is coherent, and that (ii) player  $j$  is certain that player  $i$ 's hierarchy of beliefs is coherent, and so on. When this is true, we say that player  $i$ 's belief hierarchy is *collectively coherent*. We denote the set of such hierarchies by  $H_i$ . (The set  $H_i$  can be defined inductively.)

Use the notation  $\pi_i^* = \{\pi_i^k\}_{k=1}^\infty$  as a shorthand for a belief hierarchy of player  $i$ . A *game with collectively coherent hierarchies of beliefs* is a collection  $GH = \{A_i, \Theta_i, \pi_i^*, u_i\}_{i \in \{1,2\}}$ , where  $\pi_i^* \in H_i$ . Collective coherence ensures that neither player can experience inconsistencies of any kind when reasoning about beliefs. This makes games with collectively coherent hierarchies of beliefs acceptable for analysis, in that we can define solution concepts for these games without running into loose ends. (The situation is somewhat analogous to that of perfect recall in extensive form games: although we can write down games that

fail perfect recall, it is not clear how to analyze them.)

While in principle we can analyze games of form  $GH$ , the infinite hierarchies of beliefs can make the analysis difficult to carry out. This raises the question of whether games of form  $GH$  can always be represented as Bayesian games—that is, by specifying appropriate type spaces and first-order beliefs for each player. Mertens and Zamir (1985) proved that this is always possible, thus establishing that Bayesian games provide a completely general approach to modeling strategic interactions with incomplete information. Brandenburger and Dekel (1993) provided a much simpler proof under different technical assumptions.

Here are the details. The *universal type space* is a collection  $\{(T_i^*, p_i^*)\}_{i \in \{1,2\}}$ . Player  $i$ 's set of types is  $T_i^* = \Theta_i \times H_i$ , where  $\Theta_i$  is his set of payoff types, and  $H_i$  is his set of collectively coherent hierarchies of beliefs. (Thus there is nothing that is really new so far.) The function  $p_i^*: T_i^* \rightarrow \Delta T_j^*$  is player  $i$ 's first order belief function; it describes his beliefs about player  $j$ 's type as a function of his own type. (This  $p_i^*$  is the new part. Let us note that in general there is no reason that  $p_1^*$  and  $p_2^*$  should be derivable from a common prior.)

The key step in the construction is defining the first-order belief functions  $p_i^*$ . To do so, Mertens and Zamir (1985) show that there is a homeomorphism (a continuous function with continuous inverse) between  $H_i$  and  $\Delta T_j^*$ . This function, which we denote by  $p_i$ , assigns to each hierarchy of beliefs  $\pi_i^*$  a probability distribution  $p_i(\cdot | \pi_i^*)$  on  $T_j^*$ . A key property of  $p_i$  is that for every  $k \geq 1$ , the marginal of distribution  $p_i(\cdot | \pi_i^*)$  that describes player  $i$ 's  $k$ th order belief is none other than  $\pi_i^k$ , the  $k$ th-order belief from hierarchy  $\pi_i^*$ . To sum up, this result tells us that the set  $H_i$  of collectively coherent hierarchies of beliefs is equivalent (via  $p_i$ ) to the set  $\Delta T_j^*$  of distributions over the opponent's types, and that  $p_i$  preserves the meaning of the hierarchies of beliefs.

To finish the definition of the universal type space, we set  $p_i^*(\theta_i, \pi_i^*) = p_i(\pi_i^*)$ , so that  $i$ 's posterior beliefs only depend on the hierarchy part of his type, and capture this hierarchy appropriately (by virtue of the key property of  $p_i$ ).

Putting this all together, we can conclude that any game  $GH$  with payoff-type profiles  $\Theta$  and collectively coherent hierarchies of beliefs is equivalent to a Bayesian game  $BG = \{A_i, T_i^*, p_i^*, u_i\}_{i \in \{1,2\}}$  whose type space  $\{(T_i^*, p_i^*)\}_{i \in \{1,2\}}$  is the universal type space based on  $\Theta$ , with  $BG$  being given the appropriate interim interpretation.

Let us emphasize this last point.  $GH$  represents an incomplete information interaction in which each player  $i$  has a particular payoff type  $\theta_i$  and hierarchy of beliefs  $\pi_i^*$ .  $BG$  introduces a large set of possible types  $T_i^* = \Theta_i \times H_i$  in order to include all types that may arise when considering the players' beliefs. Nevertheless, player  $i$ 's actual type is the pair  $(\theta_i, \pi_i^*)$  we started with.

Technical remark 1: For it to be meaningful to call the function  $p_i: H_i \rightarrow \Delta T_j^*$  continuous, topologies (i.e., notions of "closeness", or more specifically, convergence and continuity) must be specified for  $H_i$  and  $\Delta T_j^*$ . The central point here is that in Mertens and Zamir (1985) and Brandenburger and Dekel (1993), the space of hierarchies  $H_j$  is assigned the *product topology*, under which a sequence of hierarchies  $\pi_j^k$  converges if for each  $k \geq 1$ , the  $k$ th-order beliefs  $\pi_j^k$  converge in distribution (i.e., in the sense used in the central limit

theorem) to their limit point. Informally, convergence in the product topology requires that beliefs of any given finite order eventually settle down. There are some contexts in which more demanding notions of closeness are considered—see the last discussion point in the next section.

Technical remark 2: The central ingredient of Brandenburger and Dekel’s (1993) short proof is a result from probability theory called the *Kolmogorov extension theorem*. This theorem guarantees the existence of a probability distribution on infinite sequences whose marginals are any prespecified collection of distributions on finite subsequences that is “coherent” (i.e., that could possibly fit together as the marginals of a single distribution). This result seems plausible enough, but it turns out that constructing a distribution on an infinite-dimensional space that captures an infinite collection of finite-dimensional distributions is a complicated business.

## 4.9 Rationalizability in Bayesian Games

We now turn to analyses of Bayesian games based on common knowledge of rationality. Observation 4.3 noted that a Bayesian game  $BG$  with a common prior can be represented as an extensive form game  $\Gamma_{BG}$  in which an initial move by Nature determines the type profile.  $\Gamma_{BG}$  captures the ex ante interpretation of the Bayesian game. To define a notion of rationalizability appropriate to the ex ante interpretation, one can write down the normal form of  $\Gamma_{BG}$  and apply the definition of normal form rationalizability from Section 1.3.

### *Interim correlated rationalizability*

We now turn to the interim interpretation. In this context it is natural to consider Bayesian games which explicitly include states of nature  $\theta_0 \in \Theta_0$ , representing payoff-relevant information that no player observes. Compared to the definitions in Section 4.1, the state of nature must be incorporated into first-order beliefs and utility functions appropriately:

$$\begin{aligned} p_i : T_i &\rightarrow \Delta(T_{-i} \times \Theta_0) && \text{the first-order belief function of player } i \\ u_i : A \times T \times \Theta_0 &\rightarrow \mathbb{R} && \text{the utility function of player } i \end{aligned}$$

Expression (22) for the expected payoff of type  $\theta_i$  is replaced by

$$(34) \quad U_i(a_i, s_{-i}|t_i) = \sum_{t_{-i} \in T_{-i}} \sum_{\theta_0 \in \Theta_0} p_i(t_{-i}, \theta_0|t_i) u_i((a_i, s_{-i}(t_{-i})), t, \theta_0).$$

Taking this change into account, the condition (26) for a pure strategy profile  $s$  to be a Bayesian equilibrium remains the same:

$$(26) \quad U_i(s_i(t_i), s_{-i}|t_i) \geq U_i(\hat{a}_i, s_{-i}|t_i) \quad \text{for all } \hat{a}_i \in A_i, t_i \in T_i, \text{ and } i \in \mathcal{P}.$$

Recall that in the interim interpretation of Bayesian games, there is no ex ante stage. Each player knows his type from the start, and the other types in his type set are only there



to capture other players' uncertainty, as well as higher-order uncertainty. When defining rationalizability in this context, we allow a player's different types to hold different conjectures about the others' actions and the state of nature. In addition, we allow conjectures to include not only correlation between different opponents' actions, but between those actions and the state.

In this setting, a *conjecture* for type  $t_i$  is a probability distribution  $\mu_i(\cdot|t_i) \in \Delta(A_{-i} \times T_{-i} \times \Theta_0)$ . The conjecture must allow correlation between opponents' actions and types, since different types of player  $j$  may choose different actions. To not contradict earlier definitions, a type's conjecture must agree with his first-order beliefs. Formally, we say that type  $t_i$ 's conjecture  $\mu_i(\cdot|t_i)$  *agrees with first-order beliefs* if

$$(35) \quad \sum_{a_{-i} \in A_{-i}} \mu_i(a_{-i}, t_{-i}, \theta_0|t_i) = p_i(t_{-i}, \theta_0|t_i).$$

*Interim correlated rationalizability (ICR)* (Battigalli and Siniscalchi (2003), Dekel et al. (2007)) is defined by iteratively eliminating actions of each type that are not a best response to any of that type's conjectures that are still feasible and that agree with first-order beliefs. Since each type is considered separately, one refers to the ICR actions of a particular type  $t_i$  rather than of a player  $i$ .

Our analysis of the e-mail game (Example 4.15) used just such an iterated elimination procedure. Thus we can restate our conclusion there as saying that  $A$  is the unique ICR action of all types of both players.

Some discussion and interesting results about ICR:

- (i) As with rationalizability for normal form games, many game theorists find it most natural to allow for correlations of all sorts in the context of Bayesian games, even the novel form of correlation between opponents' actions and an unobservable state of nature—see e.g. p. 24 of Dekel et al. (2007).
- (ii) We argued in Section 4.8 that the primitive description of an incomplete information environment is a game  $GH$  defined by directly specifying each player's payoff type and hierarchy of beliefs. It is possible to specify two different Bayesian games,  $BG$  and  $BG'$ , each of which has a type,  $t_i$  and  $t'_i$ , whose payoff type and hierarchy of beliefs agrees with that of player  $i$  in  $GH$ . A desirable invariance property for an interim solution concept is that it make the same predictions of play for  $t_i$  and  $t'_i$ , since from a more primitive point of view these types should be regarded as identical. Dekel et al. (2007) show that ICR satisfies this property. In contrast, notions of interim rationalizability that require beliefs about others' actions and the state of nature to be independent do not have this invariance property. (Example 1 of Dekel et al. (2007) explains the issues here clearly.)
- (iii) It is desirable to use solution concepts that are robust to slight misspecifications of higher-order beliefs, since these may be hard for a modeler to know. Weinstein and Yildiz (2007) show that if one allows enough flexibility in these slight misspecifications, then ICR is robust, and, more importantly, no stronger interim solution concept is robust. "Enough flexibility" is captured by a *richness assumption*, which

requires that for each player  $i$  and each action  $a_i$  of that player, player  $i$  has a payoff type for whom action  $a_i$  is dominant. (More on this below.)

For intuition we again consider the e-mail game (Example 4.15). This example included a normal form game  $G_R$  with strict equilibria  $(A, A)$  and  $(B, B)$ , with  $(B, B)$  being Pareto dominant. The full e-mail game  $BG$  included types  $t_i$  whose beliefs up to high orders asserted that  $G_R$  describes the actual payoffs. However, still-higher-order beliefs of  $t_i$  allowed that payoffs might be described by another normal form game,  $G_L$ , in which  $A$  is strictly dominant. In the end,  $A$  turned out to be type  $t_i$ 's unique ICR action.

Weinstein and Yildiz (2007) observe that the selection of equilibrium  $(A, A)$  is fragile. It is also possible to construct a Bayesian game  $BG'$  with an “approximately complete information” type  $t'_i$  whose unique ICR action is  $B$ . Thus, if all we know is that agent  $i$ 's type is approximately a complete information type, there is no basis for selecting between actions  $A$  and  $B$ .

More generally, Weinstein and Yildiz (2007) show the following: Starting from any Bayesian game in which some type  $t_i$  has multiple ICR actions, one can always slightly alter beliefs to create a type  $t'_i$  that is nearly identical to  $t_i$ , but has any type  $t_i$  ICR action you like as its unique ICR action. It follows that no refinement of ICR is robust to small changes in higher-order beliefs.

(Two further results tell us more about the ICR correspondence: i.e., the map from types in the universal type space to ICR actions. Weinstein and Yildiz (2007) show that “nearly all” types in the universal type space have a unique ICR action, where “nearly all” is defined in a natural topological sense. A complementary result of Dekel et al. (2007) shows that the ICR correspondence is upper hemicontinuous.)

The richness assumption stated above is not sensible in all applications. For instance, if one starts with a first price auction, it seems odd to introduce types who find it dominant to place a bid higher than the highest value that any bidder in the original auction assigns to the good. Likewise, one might want to impose other common knowledge assumptions about the game's structure that restrict what slight modifications of the game are allowable. Penta (2013) shows that Weinstein and Yildiz's (2007) results extend as far as one could reasonably expect to settings with such common knowledge restrictions.

- (iv) The foregoing results follow Mertens and Zamir (1985) in using the product topology on hierarchies of beliefs to define what is meant by a “slight” perturbation. One can ask whether there is a stronger topology on hierarchies of beliefs (i.e., a more demanding notion of closeness) under which Weinstein and Yildiz's (2007) results no longer hold. This question is studied by Dekel et al. (2006) and Chen et al. (2010). The latter paper establishes continuity properties of the ICR correspondence under a suitably defined *uniform topology*, in which a sequence of hierarchies of beliefs  $\pi_i^*$  converges if beliefs of *all* orders *simultaneously* become close to their limiting values.

### *Weakening belief restrictions*

Section 4.6 introduced ex post equilibrium (= belief-free equilibrium), an equilibrium concept for Bayesian games that makes no use of agents' beliefs. One can proceed in a similar spirit with rationalizability. To define *belief-free rationalizability*, one drops condition (35) that players' conjectures agree with their first-order beliefs. The agnosticism inherent in this solution concept makes it quite weak in most games. As in the equilibrium case, belief-free rationalizability is well-defined in "pre-Bayesian games" in which beliefs are not specified, only the sets of payoff types and states of nature. Given how complicated beliefs can turn out to be, we can now appreciate that this is a dramatic simplification!

The definitions of ICR and belief-free rationalizability can be viewed as two extremes: in the first, the beliefs of each type are specified exactly, while in the second, no information about beliefs need be specified at all. One can also consider intermediate cases. For a general approach, one can assume that it is common knowledge that the conjectures of payoff type  $\theta_i$  lie in some prespecified subset  $\Delta_{\theta_i}$  of  $\Delta(A_{-i} \times T_{-i} \times \Theta_0)$ . To obtain the intermediate cases just noted, one can choose  $\Delta_{\theta_i}$  to embody restrictions on  $\theta_i$ 's beliefs about opponents' payoff types and the state of nature. More broadly,  $\Delta_{\theta_i}$  can capture joint restrictions on opponents' payoff types and actions (for instance, that opponents' bid functions are monotone). The class of solution concepts of this sort is introduced by Battigalli and Siniscalchi (2003) under the name  $\Delta$ -rationalizability.

Solution concepts with weak belief restrictions are basic tools in robust mechanism design, as we describe in Section 7.6.4.

## Interlude: Information Economics

The basic models:

1. *Hidden actions/moral hazard*: one player takes an action that is not perfectly observed.

*The basic principal-agent problem with hidden actions*: The player whose actions are perfectly observed, called the *principal*, chooses a contract. Then the player whose action is not perfectly observed, called the *agent*, chooses whether to accept the contract, and then chooses an action. The action generates an observable signal on which payoffs depend.

2. *Hidden information/adverse selection*: one player has private information about his personal characteristics (preferences, productivity). In the basic case, this information is obtained before play begins.

(a) *Signalling*: the informed player moves first.

- i. one uninformed player: basic signalling games à la Cho and Kreps (1987).
  - ii. competing uninformed players: Spence (1973) job market signalling.
- (b) *Screening*: the uninformed player moves first.
- i. one uninformed player: *the basic principal-agent problem with hidden information/monopolistic screening*. The player without private information, the *principal*, chooses a menu of contracts. Then the privately informed player, the *agent*, decides which contract to accept.
  - ii. competing uninformed players: Akerlof (1970) market for lemons, Rothschild and Stiglitz (1976) market for insurance
3. *Mechanism design*. A group of privately informed agents plays a game designed either by an uninformed party (e.g., the seller of a good) or by the agents themselves. In the former case, the mechanism can be viewed as a principal's solution to a *multiagent screening problem*. In the latter, the mechanism can be viewed as a solution to a *collective choice problem* faced by the agents.

One can also organize information economics according to the tools used to analyze them, and by their historical development:

The earliest models are those with competing uninformed players (Akerlof, Spence, Rothschild-Stiglitz). The competition can be represented using a partial equilibrium framework, as they were originally: price-taking behavior by the uninformed players leads to a "zero profit" condition, which must be satisfied simultaneously with optimization by each type of informed agent. These models are easily recast in game-theoretic terms, which have the advantage of specifying what happens out of equilibrium.

Signalling models are most naturally presented in a game-theoretic framework.

The basic principal-agent problems use only rudimentary game theory; the analyses amount to solving constrained optimization problems. Also, according to the *revelation principle*, it is often enough to restrict attention to *direct mechanisms*, under which agents report their private information directly to an "administrator", and are provided with incentives to do so truthfully. This device converts basic mechanism design problems into optimization or feasibility problems. Moving beyond these basic models or looking at "natural" mechanisms (like auctions) rather than direct mechanisms leads one to use game theory in a more serious way.

The outline above describes only the most basic models, and there is a vast theory that elaborates on these models in various directions.

## 5. Signalling and Screening with Competing Uninformed Players

### 5.1 The Market for Lemons (Akerlof (1970))

The players are two buyers and a seller.

A seller has a used car whose quality  $\theta \in [0, 1]$  is his private information. This  $\theta$  represents his benefit from owning the car. It is drawn from cdf  $F$  that satisfies  $F(\theta) > 0$  for all  $\theta > 0$ . We can view this as the distribution of qualities among sellers in the population, with the actual seller being drawn at random.

There are two potential buyers. Each would obtain benefit  $b(\theta) \geq \theta$  from owning a car of quality  $\theta$ . Thus the buyers always value the car at least as much as the seller.

The game proceeds as follows:

[0] The seller learns the quality of her car.

[1] The buyers simultaneously make price offers  $p_i, p_j \geq 0$ .

[2] The seller can accept one of these offers or reject both.

Buyer's utility:  $b(\theta) - p$  for buying at price  $p$ ; 0 for not buying.

Seller's utility:  $p - \theta$  for selling at price  $p$ ; 0 for not selling.

Since the uninformed parties (the buyers) move first, this is a screening model.

We look for subgame perfect equilibrium.

The seller will only consider the higher price. If the higher price is  $p$ , sellers with quality  $\theta \leq p$  sell. (What choice we specify for sellers with  $\theta = p$  is not so important.)

A buyer who chooses price 0 obtains an expected payoff of 0, since the probability that such an offer is accepted is 0. If buyer  $i$  offers  $p_i = p > p_j \geq 0$ , his expected benefit conditional on obtaining the car is

$$B(p) = \begin{cases} \frac{\int_0^p b(\theta) dF(\theta)}{F(p)} & \text{if } p \leq 1, \\ \int_0^1 b(\theta) dF(\theta) & \text{if } p > 1. \end{cases}$$

(If both buyers choose the same price, we need to say which seller types sell to each buyer—see Example 5.2.)

*Example 5.1.* Suppose  $F$  is *uniform*(0, 1) and  $b(\theta) = \beta\theta$  with  $\beta \in (1, 2)$ . Then for  $p \in (0, 1]$ ,

$$B(p) = \frac{\int_0^p \beta\theta \, d\theta}{p} = \frac{\frac{1}{2}\beta p^2}{p} = \frac{\beta}{2}p < p.$$

Since being the sole buyer at a positive price generates a negative expected payoff, buyer  $i$  does not want to choose  $p_i > p_j$ . If  $p_i = p_j = p > 0$  then the overall expected benefit is less than  $p$ , so at least one buyer is unhappy regardless of the seller's strategy (see the final paragraph). Thus in any subgame perfect equilibrium,  $p_i = p_j = 0$ , and trade occurs with probability 0. This is despite the fact that trade is always efficient! ♦

*Example 5.2.* Suppose  $F$  is *uniform*(0, 1) and  $b(\theta) = \alpha + \beta\theta$  with  $\alpha \in (0, 1)$  and  $\beta \in (1 - \alpha, 2 - 2\alpha)$ . Then for  $p \in (0, 1]$ ,

$$B(p) = \frac{\int_0^p (\alpha + \beta\theta) \, d\theta}{p} = \frac{\alpha p + \frac{1}{2}\beta p^2}{p} = \alpha + \frac{\beta}{2}p.$$

There is a unique price  $p^*$  satisfying  $B(p^*) = p^*$ , namely  $p^* = \frac{2\alpha}{2-\beta}$ .  $B(\cdot)$  crosses the 45° degree line from above at this point. (That  $\beta < 2 - 2\alpha$  ensures that  $p^* < 1$ , and that  $\beta > 1 - \alpha$  ensures that  $b(\theta) > \theta$  for all  $\theta \in [0, 1]$ .)

It cannot be that  $\bar{p} \equiv \max\{p_i, p_j\} > p^*$ : since  $B(\bar{p}) < \bar{p}$ , at least one buyer choosing this price obtains a negative expected payoff. It also cannot be that  $\bar{p} < p^*$ : If  $p_i < p_j < p^*$ , then buyer  $i$  can deviate to a price  $p$  just above  $p_j$  and change his expected payoff from 0 to  $F(p)(B(p) - p) > 0$ . If  $p_i < p_j = p^*$ , then buyer  $j$  could improve his expected payoff from 0 to something positive by slightly reducing her price. If  $p_i = p_j = p < p^*$ , then the sum of the agents' expected payoffs is  $F(p)(B(p) - p) > 0$ ; thus there is at least one agent who would improve his expected payoff by increasing his price slightly to  $\hat{p}$  and obtaining  $F(\hat{p})(B(\hat{p}) - \hat{p})$  (again, see the final paragraph).

The remaining possibility is that  $p_i = p_j = p^*$ . Then the overall expected benefit is  $B(p^*) = p^*$ . If the seller's strategy equates the buyer's expected benefits (i.e., makes them both equal to  $B(p^*) = p^*$ ), then we have a subgame perfect equilibrium. For example, seller strategies under which all types  $\theta \leq p^*$  respond to price  $p^*$  by choosing buyer 1 with probability  $q \in (0, 1)$  support  $p_i = p_j = p^*$  as an equilibrium. Notice that in equilibrium, only cars whose quality is in  $[0, p^*]$  are traded; higher quality cars are not traded despite trade being efficient.

(Combining  $p_i = p_j = p^*$  with seller strategies that do not equate the buyers' expected benefits does not lead to an equilibrium. For example, if the seller sells to buyer 1 when  $\theta \in [p, p^*]$  (with  $p > 0$ ) and sells to buyer 2 otherwise, then buyer 2's expected payoff is

$B(p) - p^* < B(p^*) - p^* < 0$ , so buyer 2 would prefer to deviate to a lower price and never obtain the car.) ♦

Remarks:

- (i) Akerlof's original model used a form of competitive equilibrium that imposes a zero profit condition on buyers. As above the equilibrium prices are defined as fixed points of  $B(\cdot)$ , but what happens out of equilibrium is not specified. In the game-theoretic model above, price competition among buyers results in zero profits without appealing to a reduced form. The game-theoretic model also shows that having two buyers is enough to drive the profits of each to zero.
- (ii) In cases where  $B(\cdot)$  has multiple fixed points, then all fixed points are competitive equilibria à la Akerlof, but typically only the largest fixed point is a subgame perfect equilibrium price. (See the problem set.)

## 5.2 Job Market Signalling (Spence (1973))

### *The model*

The players are a worker and two firms.

The worker's type  $\theta \in \Theta = \{\theta_L, \theta_H\}$ ,  $0 < \theta_L < \theta_H$ , is his ability level. He is type  $\theta_H$  with probability  $p_H$ , and thus type  $\theta_L$  with probability  $1 - p_H$ .

The game proceeds as follows:

[0] The worker learns his type.

[1] The worker chooses an education level  $e \in \mathbb{R}_+$ . (This is *signalling*: the informed party moves first.)

[2] The two firms observe the worker's education level and simultaneously make wage offers  $w_1, w_2 \in [\theta_L, \theta_H]$ .

[3] The worker then chooses which offer to accept, if any.

Firm  $i$ 's payoff for hiring a worker of type  $\theta_A \in \Theta$  at wage  $w$  is  $u_i(w, \theta_A) = \theta_A - w$ . Its payoff for not hiring a worker is 0. Thus we are assuming here that education has no direct value; its only purpose is to signal ability. Other assumptions are also possible.

The payoff of a type  $\theta_A \in \Theta$  worker is  $u_A(w, e) = w - c_A(e)$ .


The function  $c_A(\cdot)$  describes type  $\theta_A$ 's costs of acquiring an education.

We assume  $c_A(\cdot)$  is differentiable, increasing, and strictly convex with  $c_A(0) = 0$ , and that the *single crossing condition* holds:

$$c'_L(e) > c'_H(e) \text{ for } e > 0.$$

Thus at every positive education level, the high ability types have a lower marginal cost of education.

Note that all of type  $\theta_A$ 's indifference curves have slope  $c'_A(e)$  at education level  $e$ ; this slope does not depend on the wage. (Let  $\bar{u}_A \equiv w - c_A(e)$  be a type  $\theta_A$  indifference curve. Expressing  $w$  as a function of  $e$  yields  $\bar{w}_A(e) \equiv c_A(e) + \bar{u}_A$ , so differentiating yields  $\bar{w}'_A(e) \equiv c'_A(e)$ .) Thus a given pair of indifference curves, one for each type, crosses at most once; if they do cross, then the  $\theta_H$  indifference curve crosses the  $\theta_L$  indifference curve from above.



Also, since getting no education is free for both types,

$$(36) \quad c_L(e) = 0 + \int_0^e c'_L(\hat{e}) d\hat{e} > 0 + \int_0^e c'_H(\hat{e}) d\hat{e} = c_H(e).$$

Thus type  $\theta_L$  has a higher absolute cost than type  $\theta_H$  of attaining any given education level.

We will see that versions of the single crossing condition play a fundamental role throughout information economics. For general treatments, see Milgrom and Shannon (1994) or Milgrom (2004, ch. 4).



## Analysis

Since there is imperfect information, our first thought is to look for sequential equilibria of this model. However, it is not clear how to define sequential equilibrium when there are continuous sets of actions and/or types. (The problem is that in perturbed strategy profiles, almost all strategies will have probability zero.)

The usual practice is to use weak sequential equilibrium and to directly impose some of the implications of consistency. Refinements of this sort are typically called *perfect Bayesian equilibrium* (see Fudenberg and Tirole (1991a,b)). Exactly what this term means varies from paper to paper. Here we assume that the firms have common beliefs  $\mu: \mathbb{R}_+ \rightarrow \Delta\Theta$  about the worker's type conditional on his education level. We write  $\mu(e) = (\mu_L(e), \mu_H(e))$  for the beliefs given education level  $e$ . (Aside: if we write down the game tree, firm 2 will act after firm 1 but without viewing firm 1's choice, and so firm 2's beliefs must also address firm 1's choice. In effect, we are assuming that firm 2's beliefs about the worker's type and firm 1's choice are independent of one another.)

In a pure perfect Bayesian equilibrium with common beliefs,  $(e_L, e_H, w_1(\cdot), w_2(\cdot), \mu(\cdot))$ :

- (i) In the last stage, the worker chooses optimally among wage offers (i.e., works for a firm whose wage offer is highest).
- (ii) In the second-to-last stage, having observed the worker's education choice  $e$ , the firms choose wages  $w_1(e)$  and  $w_2(e)$  optimally given their common belief  $\mu(e)$ . In view of the competition between the firms, both will choose wages equal to the worker's expected ability. (The details of this argument are the same as in the Akerlof lemons model.)
- (iii) Each worker type chooses his education level optimally given (i) and (ii). Thus the education choice  $e_A$  of a worker of type  $\theta_A \in \Theta$  must satisfy

$$e_A \in \operatorname{argmax}_{e \in \mathbb{R}_+} \left( \max\{w_1(e), w_2(e)\} - c_A(e) \right).$$

- (iv) Beliefs  $\mu(e)$  are given by conditional probabilities when possible.

Mixed perfect Bayesian equilibrium is defined similarly, but we will not consider such equilibria here.

Forward induction refinements à la Cho and Kreps (1987) are used to refine the set of perfect Bayesian equilibria.

There are two kinds of pure equilibrium: *separating* (the two types choose different education levels), *pooling* (the two types choose the same education level). There are also

*semi-separating* mixed equilibria (one type mixes, sometimes choosing the same education level as the other type).

*Separating equilibrium:* The two types choose different education levels,  $e_L \neq e_H$ .

$$\Rightarrow \mu_L(e_L) = 1 \text{ and } \mu_H(e_H) = 1$$

$$\Rightarrow w_1(e_L) = w_2(e_L) = \theta_L \text{ and } w_1(e_H) = w_2(e_H) = \theta_H$$

$$\Rightarrow e_L = 0$$

(Why must  $e_L$  be zero? In a separating equilibrium, the firms recognize type  $\theta_L$  by his education choice and pay him wage  $\theta_L$  for sure. Given this, type  $\theta_L$  cannot benefit from choosing a positive education level.)

We must also check that neither type wants to choose the other type's education level. He would then be mistaken for the other type, and thus get that type's equilibrium wage.

$$\begin{aligned} \theta_L - \overbrace{c_L(0)}^{=0} &\geq \theta_H - c_L(e_H) &\Leftrightarrow & c_L(e_H) \geq \theta_H - \theta_L \\ \theta_H - c_H(e_H) &\geq \theta_L - \underbrace{c_H(0)}_{=0} &\Leftrightarrow & c_H(e_H) \leq \theta_H - \theta_L \end{aligned}$$

That is, equilibrium requires that

$$c_L(e_H) \geq \theta_H - \theta_L \geq c_H(e_H).$$


(Loosely speaking: For the low ability type, the cost of obtaining education level  $e_H$  exceeds the resulting increase in wages; for the high ability type, it does not.)

At  $e = 0$ , the first inequality fails but the second is satisfied.

Now increase  $e$  until reaching the value  $\underline{e}$  at which the first inequality binds. Since  $c_L(e) > c_H(e)$  for all  $e > 0$  (by (36)), the second inequality is satisfied strictly at  $\underline{e}$ . As we continue to increase  $e$ , both inequalities are satisfied strictly until we reach the value  $\bar{e}$  at which the right inequality binds. For  $e > \bar{e}$ , the right inequality fails.

Thus the values of  $e_H$  that are possible in equilibrium are those in the interval  $[\underline{e}, \bar{e}]$ .

The figure at left defines  $\underline{e}$  and  $\bar{e}$  and shows the types' utility levels in the  $e_H = \bar{e}$  equilibrium, at which type  $\theta_H$  is indifferent between education levels  $e_H$  and 0. The figure at right shows the types' utility levels in the  $e_H = \underline{e}$  equilibrium, at which type  $\theta_L$  is indifferent. Of course, conditional on separation,  $\theta_H$  is better off when he is able to choose a lower education level.



To complete the specification of an equilibrium with effort levels  $e_L = 0$  and  $e_H \in [e, \bar{e}]$ , we need to specify the firms' beliefs after observing education levels other than  $e_L$  and  $e_H$ . These beliefs will determine the wage they offer after the observing these unused signals, and so will determine whether either type of worker wants to deviate. The simplest specification of beliefs ensuring that neither worker type deviates is  $\mu_L(e) = 1$  for any  $e \notin \{0, e_H\}$ , so that  $w_i(e) = \theta_L$  for all such  $e$ . Many other choices of beliefs work as well. Our forward induction refinements place additional structure on out of equilibrium beliefs

We can refine the set of equilibria by applying a weak form of forward induction based on removal of conditionally dominated strategies. If firms only choose wages in  $[\theta_L, \theta_H]$ , then  $\theta_L$  must be better off playing  $e = 0$  and obtaining a payoff of at least  $\theta_L$  than playing  $e > \underline{e}$  and obtaining a payoff of at most

$$\theta_H - c_L(e) < \theta_H - c_L(\underline{e}) = \theta_L.$$

Thus a weak forward induction requirement—one that does not fix a particular equilibrium in advance—is to require  $\mu_L(e) = 0$  for all  $e > \underline{e}$  that are not dominated for type  $\theta_H$ . This rules out every  $e_H > \underline{e}$ , since if type  $\theta_H$  were to choose  $e_H$ , he could deviate to  $e'_H \in (\underline{e}, e_H)$ , pay a lower cost, but still receive wage  $\theta_H$ . Thus only the separating equilibria with  $e_H^* = \underline{e}$  survives. This is sometimes called the *Riley outcome*, after Riley (1979).

In summary:

**Proposition 5.3.** *There are perfect Bayesian equilibria with  $e_L = 0$  and  $e_H = e$  for every  $e \in [\underline{e}, \bar{e}]$ , where  $c_L(\underline{e}) = \theta_H - \theta_L = c_H(\bar{e})$ . Only equilibria generating the Riley outcome, i.e., with  $e_H = \underline{e}$ , survive the removal of conditionally dominated strategies.*

*Pooling equilibrium:* Both types choose effort level  $e^*$ .

Thus  $\mu_H(e^*) = p_H$  and  $w_1(e^*) = w_2(e^*) = (1 - p_H)\theta_L + p_H\theta_H \equiv \bar{\theta}$ .

(Note: As in the Akerlof model, a full specification of equilibrium must have workers choosing whom to work for in a way that gives both firms zero expected profits (e.g., randomizing uniformly). Otherwise, the firm earning negative expected profits would prefer to deviate.)

Whether pooling at  $e^*$  is optimal for both worker types depends on both equilibrium and out-of-equilibrium wages. For the best chance of making this work, specify  $\mu_L(e) = 1$  for  $e \neq e^*$ , so  $w_i(e) = \theta_L$  for such  $e$ . Then pooling equilibrium requires

$$(37) \quad \bar{\theta} - c_L(e^*) \geq \theta_L - c_L(0) \quad \Leftrightarrow \quad c_L(e^*) \leq \bar{\theta} - \theta_L$$

$$(38) \quad \bar{\theta} - c_H(e^*) \geq \theta_L - c_H(0) \quad \Leftrightarrow \quad c_H(e^*) \leq \bar{\theta} - \theta_L$$

Let  $\tilde{e} > 0$  make (37) bind. Then (37) is satisfied when  $e^* \in [0, \tilde{e}]$ , and by the single crossing property so is (38). Thus  $[0, \tilde{e}]$  is the set of education levels arising in pooling equilibria. (Note that since  $c_L(\underline{e}) = \theta_H - \theta_L > \bar{\theta} - \theta_L = c_L(\tilde{e})$ , it follows that  $\tilde{e} < \underline{e}$ .)

We impose forward induction here by way of the Cho-Kreps criterion (Section 2.5.2).

Fix a perfect Bayesian equilibrium with payoffs  $u_L^*$  and  $u_H^*$ . For  $T \subseteq \Theta$ , define  $W_T(e)$  to be the set of equilibrium wages that are possible given education level  $e$  if the firms' beliefs satisfy  $\sum_{\theta_A \in T} \mu_A(e) = 1$ . Then

$$D(e) = \left\{ \theta_A : u_A^* > \max_{w \in W_{\Theta}(e)} u_A(w, e) \right\}$$

is the set of types for whom education level  $e$  is equilibrium dominated. If for some  $e \neq e^*$  with  $D(e) \neq \Theta$  and some type  $\theta_B \in \Theta$ , we have

$$u_B^* < \min_{w \in W_{\Theta - D(e)}(e)} u_B(w, e),$$

then the component fails the Cho-Kreps criterion.

Fix a pooling equilibrium with common effort level  $e^* \in [0, \tilde{e}]$ , so that  $\bar{\theta} - c_L(e^*) \geq \theta_L$ .

Define  $e' > e^*$  by  $\bar{\theta} - c_L(e^*) = \theta_H - c_L(e')$  and  $e'' > e^*$  by  $\bar{\theta} - c_H(e^*) = \theta_H - c_H(e'')$ . The single crossing property implies that  $e'' > e'$ . (See the figures below for examples.)

We now argue that given the firms' anticipated reactions, only type  $\theta_H$  workers will prefer

to deviate to any  $e \in (e', e'')$ , breaking the pooling equilibrium. For such  $e$ ,

$$u_L^* > \theta_H - c_L(e), \text{ but}$$

$$(\dagger) \quad u_H^* < \theta_H - c_H(e).$$

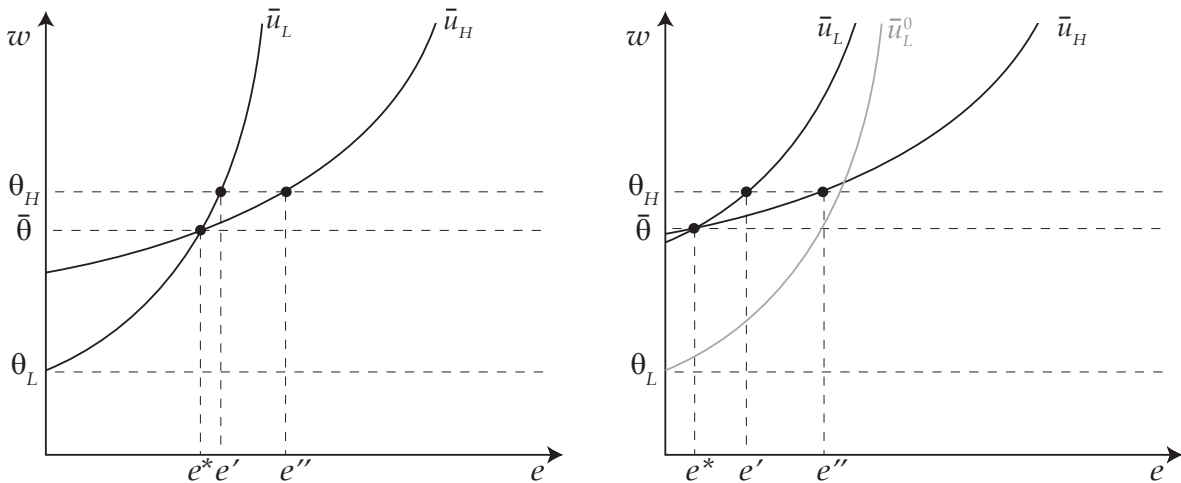
Thus  $D(e) = \{\theta_L\}$  and  $\Theta - D(e) = \{\theta_H\}$ . But if  $\mu_H(e) = 1$ , then both firms choose  $w_i(e) = \theta_H$ . Thus by  $(\dagger)$ , type  $\theta_H$  prefers to deviate to  $e$ .

In summary:

**Proposition 5.4.** *There are perfect Bayesian equilibria with  $e_L = e_H = e^*$  for every  $e^* \in [0, \bar{e}]$ , where  $c_L(\bar{e}) = \bar{\theta} - \theta_L$ . No such equilibrium satisfies the Cho-Kreps criterion.*

Some pooling equilibria can be ruled out using conditional dominance arguments, and so without appealing to the full strength of the Cho-Kreps criterion. In the figure at left, choices by type  $\theta_L$  of effort levels above  $e'$  are strictly dominated. Thus the simple forward induction argument used to refine the set of separating equilibria applies equally well here.

But in the figure at right, no effort level  $e \in (e', e'')$  is strictly dominated for type  $\theta_L$ , since the points  $(e, \theta_H)$  are all above  $\theta_L$ 's indifference curve through  $(0, \theta_L)$ . However, since the points  $(e, \theta_H)$  are below  $\theta_L$ 's indifference curve through  $(e^*, \bar{\theta})$ , an equilibrium dominance argument rules out such choices by  $\theta_L$ . Thus the Cho-Kreps criterion rules out the pooling equilibrium pictured here.



Remarks:

- (i) The Riley outcome Pareto dominates the outcomes of the other separating equilibria.

- (ii) Relative to the Riley outcome, type  $\theta_L$  workers would be better off if signalling were unavailable, since they would then obtain wage  $(1 - p_H)\theta_L + p_H\theta_H$ .
- (iii) Type  $\theta_H$  workers might also be better off if signalling were unavailable. They prefer the best pooling equilibrium to the Riley outcome when

$$(1 - p_H)\theta_L + p_H\theta_H > \theta_H - c_H(e) \Leftrightarrow c_H(e) > (1 - p_H)(\theta_H - \theta_L).$$

This will be true if  $p_H$  is high enough, since in this case the no-signal wage is nearly as large as  $\theta_H$ 's separating equilibrium wage.

- (iv) Another way to put (ii) and (iii) is to say that when  $p_H$  is close to 1, the Riley outcome is Pareto inferior to the outcome of the pooling equilibria with  $e^* = 0$ . But if the pooling equilibrium breaks down, the firms' beliefs change from  $\mu_H(0) = p_H$  to  $\mu_H(0) = 0$ , and there is no reason for them to change back.

### 5.3 Screening in Insurance Markets (Rothschild and Stiglitz (1976))

The presentation mainly follows Jehle and Reny (2011, Sec. 8.1.3).

#### *The model*

The players are two firms and a consumer.

The consumer's initial wealth is  $w$ . She may incur a loss of cost  $\ell$ .

The consumer's type  $\theta \in \Theta \in \{\theta_L, \theta_H\}$ ,  $0 < \theta_L < \theta_H$ , is her probability of a loss. (Thus  $\theta_L$ , the low risk type, is now the "good" type.) She is type  $\theta_L$  with probability  $\alpha_L \in (0, 1)$ .

An insurance policy is a pair  $(b, p)$ , representing the benefit in the event of a loss and the price of the policy.

The game proceeds as follows:

[0] The consumer learns her type.

[1] Each firm  $i$  offers a pair of policies,  $((b_i^L, p_i^L), (b_i^H, p_i^H))$ . Policy  $(b_i^P, p_i^P)$  is intended for type  $\theta_P$ , although we will need to verify that this type will actually choose the policy. There is no benefit to offering more than two policies, and, firms can in effect offer fewer than two policies by offering the null policy  $(0, 0)$ . (Each firm's pure strategy space is large!)

[2] The consumer chooses among the available policies, or chooses not to be insured—in effect, the null policy is always available.

Firms maximize expected profit.

The utility of a type  $\theta_P$  consumer from choosing policy  $(b, p)$  is

$$u_P(b, p) = (1 - \theta_P) \underbrace{v(w - p)}_{x_0} + \theta_P \underbrace{v(w - p - \ell + b)}_{x_1},$$

where  $v$  is differentiable, increasing, and strictly concave.

Thus  $u_P$  is differentiable, increasing in  $b$ , and decreasing in  $p$ .

Suppose we draw an indifference curve  $u_P(p, b) = \bar{u}_P$  of type  $\theta_P$  in  $(b, p)$  space. Expressing  $p$  as a function of  $b$  and differentiating yields

$$\frac{\partial u_P}{\partial b}(b, p) + \frac{\partial u_P}{\partial p}(b, p) \frac{dp_P}{db}(b) = 0.$$


Thus the slope of type  $P$ 's indifference curve is

$$(39) \quad \frac{dp_P}{db}(b) = -\frac{\frac{\partial u_P}{\partial b}(b, p)}{\frac{\partial u_P}{\partial p}(b, p)} = \frac{\theta_P v'(x_1)}{(1 - \theta_P)v'(x_0) + \theta_P v'(x_1)} = \frac{1}{1 + \frac{(1 - \theta_P)v'(x_0)}{\theta_P v'(x_1)}} > 0.$$

And  $v$  is strictly concave  $\Rightarrow v'(\cdot)$  is decreasing  $\Rightarrow \frac{dp_P}{db}(\cdot)$  is decreasing  $\Rightarrow p_P(\cdot)$  is strictly concave.

Implications:

- (i) The single crossing property holds: since  $\frac{dp_P}{db}(b)$  is increasing in  $\theta_P$ , type  $\theta_H$ 's indifference curves cross type  $\theta_L$ 's from below.
- (ii) A *full insurance* policy  $(b, p)$  is one with benefit  $b = \ell$ , implying that  $x_1 = x_0$ . Equation (39) shows that at such a policy, type  $\theta_P$ 's indifference curve has slope  $\theta_P$ .
- (iii) An *actuarially fair* policy  $(b^P, p^P)$  for type  $\theta_P$  is one with  $p = \theta_P b$ . Selling such a policy to its intended type earns a firm zero expected profit. By (ii), a type  $\theta^P$  consumer who could choose any actuarially fair policy would choose the full insurance policy  $(b^P, p^P) = (\ell, \theta_P \ell)$ . (See the figure.)



### Analysis

We look for pure strategy subgame perfect equilibria. (Since the uninformed party moves first, this game has stagewise perfect information, so subgame perfection is enough. We comment on mixed equilibria at the end.)

**Lemma 5.5.** *In any pure subgame perfect equilibrium, both firms earn zero expected profit.*

*Proof.* Each firm's expected profits are non-negative, since a firm can always offer a pair of null policies. So suppose expected profits are  $\pi_i > 0$  and  $\pi_j \in [0, \pi_i]$ .

If both types choose policy  $(b, p)$  (that is offered by at least one firm), then  $j$  can offer  $(b + \varepsilon, p)$  for some small  $\varepsilon > 0$  and improve his profits to  $\approx \pi_i + \pi_j$ .  $\Downarrow$

So suppose  $\theta_L$  chooses  $(b^L, p^L)$  and  $\theta_H$  chooses  $(b^H, p^H)$ . Then

$$(\dagger) \quad u_L(b^L, p^L) \geq u_L(b^H, p^H) \quad \text{and}$$

$$(\ddagger) \quad u_H(b^H, p^H) \geq u_H(b^L, p^L),$$

and the single crossing property implies that at least one inequality is strict. We suppose that  $(\dagger)$  is strict; the proof if  $(\ddagger)$  is strict is similar. If firm  $j$  offers  $(b^H + \varepsilon^H, p^H)$  for some  $\varepsilon^H > 0$  and  $\theta_H$  consumers will strictly prefer this to  $(b^H, p^H)$ :

$$u_H(b^H + \varepsilon^H, p^H) > u_H(b^H, p^H).$$

If  $\varepsilon^H$  is small enough, then the strictness of  $(\dagger)$  implies that

$$u_L(b^L, p^L) > u_L(b^H + \varepsilon^H, p^H).$$




For this fixed  $\varepsilon^H$ , the last two inequalities still hold if  $b^L$  is replaced by  $b^L + \varepsilon^L$  for  $\varepsilon^L > 0$  small enough. So if firm  $j$  chooses  $((b^L + \varepsilon^L, p^L), (b^H + \varepsilon^H, p^H))$ , it earns  $\approx \pi_i + \pi_j$ .  $\downarrow$  ■

**Lemma 5.6.** *There is no pure subgame perfect equilibrium with pooling (i.e., in which both consumer types select the same policy  $(b_*, p_*)$ ).*

*Proof.* Since each firm's expected profits are zero, so are total expected profits:

$$(40) \quad \alpha_L(p_* - \theta_L b_*) + (1 - \alpha_L)(p_* - \theta_H b_*) = 0$$

Suppose that  $b_* > 0$ . Then (40) implies that  $p_* > 0$ , and in particular that  $p_* - \theta_L b_* > 0$ .




If firm  $i$  offers  $(b_*, p_*)$ , then firm  $j$  can offer a policy in the shaded region in the figure above and sell only to type  $\theta_L$ . (This is known as *cream skimming*.) If this policy is close enough to  $(b_*, p_*)$ , firm  $j$  earns  $\approx \alpha_L(p_* - \theta_L b_*) > 0$ .  $\downarrow$

Suppose instead that  $b_* = 0$ . Then (40) implies that  $p_* = 0$ . Since  $(0, 0)$  is actuarially fair,  $\theta_H$  strictly prefers  $(\ell, \theta_H \ell)$  to  $(0, 0)$ , and so strictly prefers  $(\ell, \theta_H \ell + \varepsilon)$  to  $(0, 0)$  when  $\varepsilon > 0$  is small enough. Moreover, policy  $(\ell, \theta_H \ell + \varepsilon)$  generates positive expected profits if it is sold to type  $\theta_H$  or to type  $\theta_L$  (or both). Thus either firm can profitably deviate to this policy.  $\downarrow$

■

**Proposition 5.7.** *Suppose there is pure subgame perfect equilibrium with separation, meaning that the policy  $(b_*^L, p_*^L)$  chosen by type  $\theta_L$  differs from the policy  $(b_*^H, p_*^H)$  chosen by type  $\theta_H$ . Then*

- (i) *Each policy is actuarially fair for its intended type:  $p_*^H = \theta_H b_*^H$  and  $p_*^L = \theta_L b_*^L$ .*
- (ii) *Type  $\theta_H$ 's policy provides full insurance:  $b_*^H = \ell$ .*
- (iii) *Type  $\theta_H$  is indifferent between  $(b_*^H, p_*^H)$  and  $(b_*^L, p_*^L)$ .*
- (iv) *Type  $\theta_L$ 's policy provides less than full insurance:  $b_*^L < \ell$ .*




*Proof.* Step 1:  $u_H(b_*^H, p_*^H) \geq u_H(\ell, \theta_H \ell)$  ( $\theta_H$  gets at least its actuarially fair full insurance payoff).

Why? If not, then a firm can profitably offer an almost actuarially fair policy and make positive profit (from just  $\theta_H$  or from both types).

Step 2:  $p_*^H - \theta_H b_*^H \leq 0$  (firms generate nonpositive expected profits from  $\theta_H$ ).


Why? Since  $\theta_H$ 's favorite fair policy is full insurance, any policy  $\theta_H$  likes at least as much as fair full insurance generates a nonpositive profit—see the figure.



Step 3:  $p_*^L - \theta_L b_*^L = 0$  (firms generate zero expected profits from  $\theta_L$ , and so from  $\theta_H$  too).

Why? Since expected profits are zero by Lemma 5.5, Step 2 implies that expected profits from type  $\theta_L$  are nonnegative:  $p_*^L - \theta_L b_*^L \geq 0$ .


Suppose  $p_*^L - \theta_L b_*^L$  is positive. Then  $p_*^L > 0$ , and so  $b_*^L > 0$ ; otherwise  $\theta_L$  would prefer to be uninsured. Now some firm  $i$  offers policy  $(b_*^H, p_*^H)$ . Consider what happens if firm  $j \neq i$  only offers policies that are both close to  $(b_*^L, p_*^L)$  and in the shaded region in the figure below. Since these policies are close to  $(b_*^L, p_*^L)$ , they earn positive expected profits when sold to  $\theta_L$ . Since they are in the shaded region,  $\theta_L$  will buy them rather than  $(b_*^L, p_*^L)$ , and  $\theta_H$  likes them less than  $(b_*^L, p_*^L)$ , and hence less than  $(b_*^H, p_*^H)$ . Thus firm  $j$  will earn a positive expected profit, contradicting Lemma 5.5.



Step 4:  $(b_*^H, p_*^H) = (\ell, \theta_H \ell)$

Why? The inequality in step 1 holds, and step 3 implies that the inequality in step 2 binds;  $(\ell, \theta_H \ell)$  is the only policy consistent with these two facts.

Step 5:  $(b_*^L, p_*^L) = (\underline{b}, \theta_L \underline{b})$  where  $u_H(\underline{b}, \theta_L \underline{b}) = u_H(\ell, \theta_H \ell)$  and  $\underline{b} < \ell$ .



Why? We know from step 3 that  $p_*^L = \theta_L b_*^L$ .

Suppose there is a  $\bar{b} > \ell$  such that  $u_H(\bar{b}, \theta_L \bar{b}) = u_H(\ell, \theta_H \ell)$ . Then applying the single-crossing property at  $(\ell, \theta_H \ell)$  shows that  $\theta_L$ 's indifference curve through this contract

crosses  $p = \theta_L b$  from above at a point southwest of  $(\bar{b}, \theta_L \bar{b})$ . Thus  $\theta_L$  strictly prefers  $(\ell, \theta_H \ell)$  to any  $(b_*^L, \theta_L b_*^L)$  with  $b_*^L \geq \bar{b}$ .  $\Downarrow$

If  $b_*^L > \underline{b}$  (and  $b_*^L < \bar{b}$  if  $\bar{b}$  exists), then  $\theta_H$  strictly prefers  $(b_*^L, p_*^L)$  to  $(b_*^H, p_*^H) = (\ell, \theta_H \ell)$ .  $\Uparrow$

If  $b_*^L < \underline{b}$  and  $i$  offers  $(b_*^L, p_*^L)$ , then since  $\theta_L$ 's indifference curve through  $(b_*^L, p_*^L)$  crosses the actuarially fair line from below,  $j$  can cream-skim by slightly increasing  $b$  and  $p$ . (The claim about the indifference curve is true because the function  $p_L(\cdot)$  describing  $\theta_L$ 's indifference curve is concave and has slope  $\theta_L$  at  $b = \ell$ , and thus slope greater than  $\theta_L$  at  $b = b_*^L$ .)  $\Downarrow$  ■

Note the similarity of the equilibrium outcome here to that of the Riley outcome of the Spence model.

Each model has an "indicator" (benefit, education level) and a transfer (price, wage).

In each model,


- (i) The firm obtains zero expected profit from each type.
- (ii) The "bad" type gets its preferred indicator: since it does not benefit from being identified, it will not accept an inferior indicator to make this happen.
- (iii) The "good" type's indicator-transfer pair makes the "bad" type indifferent: because of the competition between the firms, the "good" type sacrifices only as much as it has to to achieve separation.

While we looked at signalling in a labor market and screening in an insurance market, we could have reversed either of these:

- (i) In a labor market screening game (firms offer wage-effort pairs), if a pure subgame perfect equilibrium exists it generates the Riley (1979) outcome. (See MWG ch. 13.D.)
- (ii) In an insurance signaling game (consumers propose policies, firms accept or reject), the Cho-Kreps criterion selects the "Riley outcome" (See Jehle-Reny sec. 8.1.2.)

Pure subgame perfect equilibria of this model need not exist. One possibility for a failure of existence is illustrated in the figure below, where policies in the shaded region are accepted by both types and generate positive expected profits. Such a region exists whenever  $\alpha_L$  is close enough to 1 (i.e., whenever the probability of a low risk type is high enough).

Dasgupta and Maskin (1986b) prove that mixed subgame perfect equilibrium always exists in this model.



## 6. Principal-Agent Models

In a *principal-agent model*, an uninformed party (the *principal*) makes a take-it-or-leave-it offer to an informed party (the *agent*). This setup gives the principal all of the bargaining power, but the agent's informational advantage gives her some power as well.

### 6.1 Moral Hazard

In the *principal-agent model with moral hazard*, the form of the informational advantage is that the principal cannot observe the agent's action choice, but only a stochastic output that is influenced by this choice.

Play proceeds as follows:

[1] The principal offers a contract that specifies a *wage schedule*: for each output level  $x_j$ , it prescribes a wage  $w_j$ .

[2] The agent decides whether to accept the contract, and, if she accepts it, makes an unobservable effort choice  $e$ . The output is then drawn from a distribution that depends on  $e$ . The principal keeps the output and pays the agent the appropriate wage.

The agent's utility is  $\underline{u} \in \mathbb{R}$  for rejecting the contract, and  $u(w) - e$  for choosing effort  $e$  and earning wage  $w$ , where  $u(\cdot)$  is twice differentiable with  $u(0) = 0$  and  $u'(w) > 0$  and  $u''(w) < 0$  for  $w > 0$ . Thus effort levels are named according to their costs.

The principal maximizes expected profit. (A common variation is to assume the principal is risk-averse.)

Analysis: We look for subgame perfect equilibrium.

In stage 2, the agent will either choose her preferred effort level or choose to reject the contract.

Thus for each possible effort level  $e$ , the principal determines the wage schedule that induces choice  $e$  while earning the principal the highest expected profit for the principal while inducing the agent to choose  $e$ . The principal then selects optimally among these wage schedules and the option to shut down.

One basic question concerns the form of the wage schedules: under what conditions are wages nondecreasing in output?

Some key early papers: Mirrlees (1999), Holmström (1979). Complete analysis: Grossman and Hart (1983). The presentation here borrows from the textbook treatment of Salanié (2005, Sec. 5.1–5.2).

### 6.1.1 Two effort levels, two output levels

#### *Notation*

effort levels/costs:  $e \in \{e_L, e_H\}$ ,  $0 < e_L < e_H$ .

output levels:  $x \in \{x_F, x_S\}$  (failure, success),  $x_F < x_S$ .

output law:  $q_K = \text{Prob}(x_S|e_K)$ ,  $0 < q_L < q_H < 1$ .


#### *Analysis*

Inducing  $e = e_L$ :

$$\begin{aligned} & \max_{w_F, w_S} q_L(x_S - w_S) + (1 - q_L)(x_F - w_F) \quad \text{subject to} \\ \text{(I}_L\text{)} \quad & q_L u(w_S) + (1 - q_L)u(w_F) - e_L \geq q_H u(w_S) + (1 - q_H)u(w_F) - e_H, \\ & \left( \text{equiv.: } u(w_S) - u(w_F) \leq \frac{e_H - e_L}{q_H - q_L} \right) \\ \text{(R}_L\text{)} \quad & q_L u(w_S) + (1 - q_L)u(w_F) - e_L \geq \underline{u}. \end{aligned}$$

Here we implicitly assume that when she is indifferent, the agent acts in the way that the principal prefers. This must be the case in equilibrium, for the same reason that in the unique equilibrium of alternating offer bargaining games, each player's offers make the opponent indifferent between accepting and rejecting.

Solution: Consider the relaxed problem that omits constraint (I<sub>L</sub>). In this problem,



- (i)  $w_S$  and  $w_F$  should be chosen to make  $(R_L)$  bind.
- (ii)  $w_S = w_F$ . (See the figure. The principal wants to choose the wage pair that is on the isoprofit line closest to the origin among those that give the agent her reservation utility of  $\underline{u}$ . Implicit differentiation yields the slopes of the isoprofit lines and indifference curves. Tangency requires  $u'(w_F) = u'(w_S)$ , so strict concavity of  $u$  implies that  $w_S = w_F$ .) Then (i) implies that  $w_L^* = w_S = w_F$  should satisfy  $u(w_L^*) = \underline{u} + e_L$ .

But if wages are independent of output, constraint  $(I_L)$  is satisfied. Thus  $w_S = w_F = w_L^*$  solves the original problem.

Inducing  $e = e_H$ :

$$\max_{w_F, w_S} q_H(x_S - w_S) + (1 - q_H)(x_F - w_F) \quad \text{subject to}$$


$$(I_H) \quad u(w_S) - u(w_F) \geq \frac{e_H - e_L}{q_H - q_L},$$

$$(R_H) \quad q_H u(w_S) + (1 - q_H)u(w_F) - e_H \geq \underline{u}.$$

Solution: Wages should be chosen to make  $(R_H)$  bind, since if it did not bind, reducing  $u(w_S)$  and  $u(w_F)$  equally would increase profits without affecting  $(I_H)$ .

Constraint  $(I_H)$  must also bind. (See the figure. In the relaxed program without constraint  $(I_H)$  the principal would choose the black dot. But the constraint curve  $u(w_S) - u(w_F) = d > 0$  is above the 45° line, and points satisfying the constraint are on or above this curve. Thus

the black dot is not feasible. The feasible set is the shaded region above the gray dot, so the best the principal can do is to choose the gray dot.)



Thus simultaneously solving the binding versions of  $(I_H)$  and  $(R_H)$  gives us the equilibrium utilities (and implicitly the equilibrium wages):

$$u(w_F^*) = \underline{u} + \frac{q_H e_L - q_L e_H}{q_H - q_L}, \quad u(w_S^*) = \underline{u} + \frac{(1 - q_L)e_H - (1 - q_H)e_L}{q_H - q_L}.$$

Thus since  $u(w_L^*) = \underline{u} + e_L$ , a calculation verifies that  $w_F^* < w_L^* < w_S^*$ .

Which effort level (if either) should be induced?

$$\begin{aligned} \Pi_H^* \geq \Pi_L^* &\Leftrightarrow q_H(x_S - w_S^*) + (1 - q_H)(x_F - w_F^*) \geq q_L(x_S - w_L^*) + (1 - q_L)(x_F - w_L^*) \\ &\Leftrightarrow (q_H - q_L)(x_S - x_F) \geq q_H w_S^* + (1 - q_H)w_F^* - w_L^*. \end{aligned}$$

Recall that  $w_L^*$ ,  $w_F^*$ , and  $w_S^*$  do not depend on  $x_S$  or  $x_F$ . Thus the greater is  $x_S - x_F$ , the more attractive inducing effort becomes.

### 6.1.2 Finite sets of effort and output levels

#### *Notation*

effort levels/costs:  $e \in E = \{e_1, \dots, e_n\}$ ,  $0 < e_1 < \dots < e_n$ .



output levels:  $x \in X = \{x_1, \dots, x_m\}$ ,  $x_1 < \dots < x_m$ .

output law:  $p_{ij} = \text{Prob}(x_j|e_i) > 0$ .

### Analysis

In stage 2, the agent will choose his preferred effort level or to reject the contract. Thus as before, the principal determines his maximal expected profit conditional on inducing each effort level  $e_i$ , and then optimizes over these possibilities and the option to shut down.

The following program determines the maximal expected profit and optimal wage schedule for inducing effort level  $e_i$ :

$$\begin{aligned}
 (\text{P}_i) \quad & \max_{w \in \mathbb{R}^m} \sum_{j=1}^m p_{ij}(x_j - w_j) \quad \text{subject to} \\
 (\text{I}_{h|i}) \quad & \sum_{j=1}^m p_{ij}u(w_j) - e_i \geq \sum_{j=1}^m p_{hj}u(w_j) - e_h \quad (\text{for all } h \neq i) \\
 (\text{R}_i) \quad & \sum_{j=1}^m p_{ij}u(w_j) - e_i \geq \underline{u}
 \end{aligned}$$

We solve this program using the Kuhn-Tucker method. The Lagrangian is

$$\mathcal{L}_i(w, \lambda, \mu) = \sum_{j=1}^m p_{ij}(x_j - w_j) + \sum_{h \neq i} \lambda_h \left( \sum_{j=1}^m (p_{ij} - p_{hj})u(w_j) - e_i + e_h \right) + \mu \left( \sum_{j=1}^m p_{ij}u(w_j) - e_i - \underline{u} \right).$$

The Kuhn-Tucker conditions are

$$\begin{aligned}
 (\text{KT1}) \quad & -p_{ij} + \sum_{h \neq i} \lambda_h (p_{ij} - p_{hj})u'(w_j) + \mu p_{ij}u'(w_j) = 0 \quad \text{for all } j \in \{1, \dots, m\} \\
 & \left( \text{equiv.: } \frac{1}{u'(w_j)} = \mu + \sum_{h \neq i} \lambda_h \left( 1 - \frac{p_{hj}}{p_{ij}} \right) \quad \text{for all } j \in \{1, \dots, m\} \right) \\
 (\text{KT2}) \quad & \lambda_h \geq 0; \quad (\text{I}_{h|i}) \text{ binds or } \lambda_h = 0 \text{ or both} \quad (\text{for all } h \neq i) \\
 (\text{KT3}) \quad & \mu \geq 0; \quad (\text{R}_i) \text{ binds or } \mu = 0 \text{ or both}
 \end{aligned}$$

Although  $(\text{P}_i)$  is not a concave program, conditions  $(\text{KT1})$ – $(\text{KT3})$  are necessary and sufficient for maximization: Suppose we change the choice variables from the wages  $w_j$  to the utilities  $u(w_j)$ . Then  $(\text{P}_i)$  becomes a program with a strictly concave objective function and linear constraints. The Kuhn-Tucker conditions for this new program are necessary and sufficient for maximization, and in fact they are equivalent to  $(\text{KT1})$ – $(\text{KT3})$ .

We first consider inducing the lowest effort level.

**Proposition 6.1.** *The optimal contract for inducing effort level  $e_1$  has a flat wage schedule:  $w_1 = \dots = w_m = \underline{w}^*$ , where  $u(\underline{w}^*) = \underline{u} + e_1$ .*

*Proof.* Suppose we set  $\lambda_2 = \dots = \lambda_n = 0$ . Then (KT1) implies that  $u'(w_j)$  is constant over  $j$ , and thus that  $w_j$  is constant over  $j$ . Let  $w^*$  be this common wage. Since  $u'(w^*) > 0$  by assumption, (KT1) also implies that  $\mu > 0$ . Thus (KT3) and (R<sub>1</sub>) imply that  $u(w^*) = \underline{u} + e_1$ , and so that  $w^* = \underline{w}^*$ . Finally, for each  $h \neq 1$ , (I<sub>h|1</sub>) is satisfied since  $e_h > e_1$ . ■

Setting all  $\lambda_j$ s to 0 is tantamount to assuming that none of the (I<sub>h|1</sub>) constraints bind. So in essence, the proof above solves the relaxed program and then checks that the solution is feasible in the original program.

So far, we have made no assumptions at all about the output law  $p$ . We now introduce conditions on the output law that ensure that optimal wage schedules are nondecreasing.

We say that the output law satisfies the *monotone likelihood ratio property (MLRP)* if

$$\text{if } g < i, \text{ then } \frac{p_{ij}}{p_{gj}} \text{ is nondecreasing in } j.$$

That is: higher outcomes have higher likelihood ratios in favor of higher actions. To match (KT1), we can write the MLRP in the following way:

$$(41) \quad \text{if } g < i, \text{ then } 1 - \frac{p_{gj}}{p_{ij}} \text{ is nondecreasing in } j.$$

The next lemma is a consequence of the MLRP for cases in which the principal does not need to actively dissuade the agent from higher efforts.

**Lemma 6.2.** *If the MLRP holds and the optimal solution to (KT1)–(KT3) satisfies  $\lambda_h = 0$  for  $h > i$ , then wages are nondecreasing in output:  $k \geq j$  implies that  $w_k \geq w_j$ .*

*Proof.* When  $\lambda_h = 0$  for  $h > i$ , the MLRP (41) and (KT1) imply that increasing  $j$  weakly increases the right-hand side of the restated version of (KT1). To maintain the equality,  $u'(w)$  must weakly decrease; thus since  $u$  is concave,  $w$  must weakly increase. ■

Of course, the condition on the multipliers is trivially true for the highest effort level. Thus Lemma 6.2 immediately implies

**Proposition 6.3.** *Under the MLRP, the optimal wage schedule for inducing the maximal effort  $e_n$  is nondecreasing.*

If there are more than two actions, what can we say about optimal contracts to induce intermediate effort levels? Here clean results require an additional assumption, which we state in terms of the decumulative probabilities

$$\tilde{P}_{ij} = \sum_{k>j} p_{ik} = \text{Prob}(x > x_j | e_i).$$

We first state an implication of the MLRP whose proof is left as an exercise.

**Lemma 6.4.** *If the MLRP holds, then the outcome distributions  $p_i$  satisfy first order stochastic dominance: for all  $j$ ,  $\tilde{P}_{ij}$  is nondecreasing in  $i$ .*

If  $e_g < e_i < e_h$ , then there is an  $\alpha \in (0, 1)$  such that  $e_i = (1 - \alpha)e_g + \alpha e_h$ . We say that the output law satisfies *concavity of decumulative probabilities in effort (CDE)* if

$$(42) \quad \tilde{P}_{ij} \geq (1 - \alpha)\tilde{P}_{gj} + \alpha\tilde{P}_{hj} \quad \text{for all } e_g < e_i < e_h.$$

Lemma 6.4 says that as effort  $e_i$  increases, the probability  $\tilde{P}_{ij}$  of an outcome better than  $x_j$  weakly increases. CDE adds the requirement that this probability increase at a weakly decreasing rate. (For instance, if  $e_i = i$  for all  $i$ , then (42) says that  $\tilde{P}_{(i+1)j} - \tilde{P}_{ij}$ , the increase in the probability of an output above  $x_j$  obtained by increasing effort from  $i$  to  $i + 1$ , is nonincreasing in  $i$ .) Thus CDE is a form of decreasing marginal returns to effort.

**Proposition 6.5.** *Under the MLRP and CDE, the optimal wage schedule for inducing any effort level  $e_i$  is nondecreasing in output.*

*Proof.* We have already seen that the result is true for  $i = 1$  and  $i = n$ .

**Lemma 6.6.** *Let  $i > 1$ . In the optimal solution  $(w, \lambda_{-i}, \mu)$  to  $(P_i)$ ,  $\lambda_g > 0$  (and hence  $(I_{g|i})$  binds) for some  $g < i$ .*

*Proof.* Suppose to the contrary that  $\lambda_g = 0$  for all  $g < i$ . Then this and the fact that  $(w, \lambda_{-i}, \mu)$  satisfies (KT1)–(KT3) for  $(P_i)$  imply that  $(w, \lambda_{-i}, \mu)$  also satisfies (KT1)–(KT3) for the relaxed problem  $(P_i^+)$  in which only efforts in  $\{e_i, \dots, e_n\}$  are available. In  $(P_i^+)$ ,  $e_i$  is the smallest effort, so Proposition 6.1 implies that  $w_1 = \dots = w_m$ . But then  $(I_{g|i})$  is violated for all  $g < i$ , implying that  $(w, \lambda_{-i}, \mu)$  is not feasible in  $(P_i)$ .  $\updownarrow$  ■

Now consider the relaxed problem  $(P_i^-)$  of inducing  $e_i$  when only efforts in  $\{e_1, \dots, e_i\}$  are available. By Proposition 6.3, the optimal wages  $w_1, \dots, w_m$  for this problem are nondecreasing. To complete the proof of the proposition, we show that these wages remain feasible in  $(P_i)$ , and so are optimal in  $(P_i)$ .

To do so, fix  $h > i$ ; we want to show that  $(I_{hi})$  is satisfied. By Lemma 6.6, there is a  $g < i$  such that  $(I_{g|i})$  binds. The intuition now is that since effort levels  $e_g$  and  $e_i$  generate the same expected payoffs under wage schedule  $w$ , decreasing marginal returns to effort implies that the expected payoffs from effort level  $e_h$  must be lower under  $w$ . In other words, CDE implies that effort levels higher than the principal intends are unattractive to the agent.

To show this formally, write  $e_i = (1 - \alpha)e_g + \alpha e_h$ , sum by parts, and apply CDE (42) to obtain

$$\begin{aligned}
\sum_{j=1}^m p_{ij}u(w_j) - e_i &= u(w_1) + \sum_{j=1}^{m-1} \tilde{P}_{ij} (u(w_{j+1}) - u(w_j)) - e_i \\
&\geq (1 - \alpha) \left( u(w_1) + \sum_{j=1}^{m-1} \tilde{P}_{gj} (u(w_{j+1}) - u(w_j)) - e_g \right) \\
&\quad + \alpha \left( u(w_1) + \sum_{j=1}^{m-1} \tilde{P}_{hj} (u(w_{j+1}) - u(w_j)) - e_h \right) \\
&= (1 - \alpha) \left( \sum_{j=1}^m p_{gj}u(w_j) - e_g \right) + \alpha \left( \sum_{j=1}^m p_{hj}u(w_j) - e_h \right) \\
&= (1 - \alpha) \left( \sum_{j=1}^m p_{ij}u(w_j) - e_i \right) + \alpha \left( \sum_{j=1}^m p_{hj}u(w_j) - e_h \right),
\end{aligned}$$

where the final equality follows from the fact that  $(I_{g|i})$  binds. Rearranging the inequality between the first and last expressions yields  $(I_{hi})$ . ■

## 6.2 Adverse Selection (Monopolistic Screening)

Here the (uninformed) principal offers a (privately informed) agent a menu of contracts. *Adverse selection* refers to the fact that in designing the menu of contracts, the principal must ensure that each type will select the contract intended for it rather than the contract intended for some other type.

In the basic setting considered here, the principal's payoff does not depend directly on the agent's type. This assumption is appropriate for modeling sales of consumption goods,

but not for insurance contracting.

We describe the model as one of screening via product quality. One can also interpret the model as one of screening via quantity purchased—that is, one of *nonlinear pricing*. For the latter interpretation to be appropriate in a given application, the principal must be able to ensure that agents cannot circumvent the pricing scheme, either by making many small purchases, or by teaming up to make group purchases that are then divided among the team. See Wilson (1993) for a comprehensive treatment.

Key early references: Mirrlees (1971), Mussa and Rosen (1978), Baron and Myerson (1982), Maskin and Riley (1984).

### 6.2.1 Two types

The agent's type,  $\theta \in \{\theta_\ell, \theta_h\}$ ,  $0 < \theta_\ell < \theta_h$ , represents his marginal return to quality. The probability that he is type  $\theta_h$  is  $\pi_h \in (0, 1)$ .

A type  $\theta$ 's utility for purchasing a good of quality  $q \geq 0$  at price  $p \in \mathbb{R}$  is

$$u(q, p, \theta) = \theta q - p.$$

In addition to obvious monotonicity properties, we have the single-crossing property

$$\frac{\partial^2}{\partial q \partial \theta} u(q, p, \theta) > 0.$$

If the agent does not purchase a good, her utility is 0.

The principal's utility for selling a good of quality  $q$  at price  $p$  is  $p - c(q)$ , where  $c(0) = c'(0) = 0$ ,  $c'(q), c''(q) > 0$  for  $q > 0$ , and  $\lim_{q \rightarrow \infty} c'(q) = \infty$ .

The game proceeds as follows:

[0] The agent learns her type.

[1] The principal offers a menu of contracts,  $((q_\ell, p_\ell), (q_h, p_h)) \in \mathbb{R}_+^{2 \times 2}$ .

[2] The agent either chooses one or none of the contracts offered.

If the principal could observe the agent's type, he could extract all of the agent's surplus. The principal solves

$$\max_{q \geq 0, p} p - c(q) \quad \text{subject to } \theta q - p \geq 0.$$

Since it is clearly optimal to make the constraint bind, this problem is equivalent to

$$\max_{q \geq 0} \theta q - c(q)$$

The solution is to choose the  $\hat{q}$  satisfying  $c'(\hat{q}) = \theta$ , and to set price  $\hat{p} = \theta \hat{q}$ .

What if the agent's type is unobservable? Suppose the principal offered the menu  $((\hat{q}_\ell, \theta_\ell \hat{q}_\ell), (\hat{q}_h, \theta_h \hat{q}_h))$  consisting of each type's optimal contract when types are observable. Since

$$\theta_h \hat{q}_\ell - \hat{p}_\ell = \theta_h \hat{q}_\ell - \theta_\ell \hat{q}_\ell > 0,$$

the high type prefers the low type's contract. What should the principal do?

The principal's problem is to choose a pair of contracts,  $((q_\ell, p_\ell), (q_h, p_h))$ , that solves

$$\begin{aligned} & \max_{q_\ell, q_h \geq 0; p_\ell, p_h} (1 - \pi_h)(p_\ell - c(q_\ell)) + \pi_h(p_h - c(q_h)) \quad \text{subject to} \\ (\text{IC}_\ell) & \quad \theta_\ell q_\ell - p_\ell \geq \theta_\ell q_h - p_h \\ (\text{IC}_h) & \quad \theta_h q_h - p_h \geq \theta_h q_\ell - p_\ell \\ (\text{IR}_\ell) & \quad \theta_\ell q_\ell - p_\ell \geq 0 \\ (\text{IR}_h) & \quad \theta_h q_h - p_h \geq 0 \end{aligned}$$

Constraints  $(\text{IC}_\ell)$  and  $(\text{IC}_h)$ , which say that each type prefers the contract intended for it, are called *incentive compatibility* constraints. Constraints  $(\text{IR}_\ell)$  and  $(\text{IR}_h)$ , which say that each type prefers its contract to no contract, are called *individual rationality* (or *participation*) constraints.


The null contract pair  $((0, 0), (0, 0))$  is feasible in the principal's problem. What contract pair is optimal?

**Proposition 6.7.** *At the optimal solution  $((q_\ell^*, p_\ell^*), (q_h^*, p_h^*))$  of the principal's problem,*

- (i)  $(\text{IC}_h)$  binds. (The high type is indifferent between the contracts.)
- (ii)  $(\text{IR}_\ell)$  binds. (The low type gets zero surplus.)
- (iii)  $c'(q_h^*) = \theta_h$ . (The high type gets its efficient quality.)
- (iv)  $c'(q_\ell^*) = \max\{\theta_\ell - \frac{\pi_h}{1-\pi_h}(\theta_h - \theta_\ell), 0\}$ . (The low type gets less than its efficient quality, and may not be served at all.)

The next figure illustrates a possible optimal contract pair.  $(\text{IR}_\ell)$  binds because the type

$\theta_\ell$  contract is on the type  $\theta_\ell$  indifference curve through the origin.  $(IC_h)$  binds because both contracts are on the same type  $\theta_h$  indifference curve. Quality  $q_h^*$  is determined by the marginal cost requirement  $c'(q_h^*) = \theta_h$ , and in the case illustrated, quality  $q_\ell^*$  is determined by the marginal cost requirement  $c'(q_\ell^*) = \theta_\ell - \frac{\pi_h}{1-\pi_h}(\theta_h - \theta_\ell)$ .



*Proof.*

Observe first that constraint  $(IR_h)$  is redundant: If  $(IC_h)$  and  $(IR_\ell)$  hold, then

$$(43) \quad \theta_h q_h - p_h \geq \theta_h q_\ell - p_\ell \geq \theta_\ell q_\ell - p_\ell \geq 0.$$

We now establish that properties (i) and (ii) must hold for a feasible menu  $((q_\ell, p_\ell), (q_h, p_h))$  to be optimal.

(i) Constraint  $(IC_h)$  binds (i.e.,  $p_h = \theta_h(q_h - q_\ell) + p_\ell$ ): If  $(IC_h)$  is loose, then (43) implies that  $(IR_h)$  is loose as well, so one can increase  $p_h$  without violating these constraints. Increasing  $p_h$  also makes  $(IC_\ell)$  easier to satisfy. Thus increasing  $p_h$  is feasible and increases the principal's payoffs.

(ii) Constraint  $(IR_\ell)$  binds (i.e.,  $p_\ell = \theta_\ell q_\ell$ ): If  $(IR_\ell)$  is loose, one can increase  $p_h$  and  $p_\ell$  equally without violating  $(IR_\ell)$ ,  $(IC_\ell)$ , or  $(IC_h)$  (or the redundant  $(IR_h)$ ). Thus doing so is feasible and increases profits.

When  $(IC_h)$  binds, we can substitute the expression for  $p_h$  from (i) into  $(IC_\ell)$  to obtain

$$(44) \quad (IC_\ell) \Leftrightarrow (\theta_h - \theta_\ell)(q_h - q_\ell) \geq 0 \Leftrightarrow q_h \geq q_\ell.$$

Therefore, since  $(IR_h)$  is redundant, (i), (ii), and (44) reduce the principal's problem to

$$(45) \quad \begin{aligned} & \max_{q_\ell, q_h \geq 0} (1 - \pi_h)(\theta_\ell q_\ell - c(q_\ell)) + \pi_h(\theta_\ell q_\ell + \theta_h(q_h - q_\ell) - c(q_h)) \\ & \text{subject to } q_h \geq q_\ell \end{aligned}$$

To solve (45), we consider the relaxed problem in which the constraint  $q_h \geq q_\ell$  is omitted. If we find that the solution to this relaxed problem satisfies  $q_h \geq q_\ell$ , it is then a solution to (45), and hence (with the expressions for  $p_h$  and  $p_\ell$  above) a solution to the original problem.

To solve this relaxed version of (45), note that its objective function is strictly concave and separable in  $q_\ell$  and  $q_h$ . The first order condition for  $q_h$  is  $\theta_h - c'(q_h) = 0$ , giving us (iii). (Since  $c'(\cdot)$  is increasing with  $c'(0) = 0$  and  $\lim_{q \rightarrow \infty} c'(q) = \infty$  this equation has a unique solution.)

To obtain (iv), compute the partial derivative of the objective function with respect to  $q_\ell$ :

$$(46) \quad (1 - \pi_h)(\theta_\ell - c'(q_\ell)) - \pi_h(\theta_h - \theta_\ell) = (1 - \pi_h) \left[ \theta_\ell - \frac{\pi_h}{1 - \pi_h}(\theta_h - \theta_\ell) - c'(q_\ell) \right].$$

It follows that if  $\theta_\ell - \frac{\pi_h}{1 - \pi_h}(\theta_h - \theta_\ell)$  is nonnegative, then setting  $c'(q_\ell)$  equal to it yields the optimal choice of  $q_\ell$ , which is less than  $q_h$  because  $\theta_h > \theta_\ell > \theta_\ell - \frac{\pi_h}{1 - \pi_h}(\theta_h - \theta_\ell)$  and  $c'(\cdot)$  is increasing. If instead  $\theta_\ell - \frac{\pi_h}{1 - \pi_h}(\theta_h - \theta_\ell)$  is negative, then (46) is negative for all  $q_\ell \geq 0$ , implying that it is optimal to choose  $q_\ell = 0$ . Either way,  $q_h \geq q_\ell$  as required. ■

Remarks:

- (i) Type  $\theta_h$ 's equilibrium payoff is  $\theta_h q_h - p_h = (\theta_h - \theta_\ell)q_\ell$ . This is sometimes called type  $\theta_h$ 's *information rent*, since it is the part of her surplus that the principal is unable to extract.
- (ii) Examining (45) shows that the principal chooses  $q_\ell$  to maximize

$$(47) \quad \theta_\ell q_\ell - \frac{\pi_h}{1 - \pi_h}(\theta_h - \theta_\ell)q_\ell - c(q_\ell).$$

The first two terms of (47) are called the *virtual utility* of type  $\theta_\ell$ , and (47) as a whole is called the *virtual surplus* from type  $\theta_\ell$ . The first term in (47) is type  $\theta_\ell$ 's actual utility for obtaining quality  $q_\ell$ , and the last term is the principal's cost of providing this quality. Subtracted from this is the information rent of type  $\theta_h$ , multiplied by that type's relative probability. This deduction accounts for the fact that increasing  $q_\ell$  increases the information rent that the principal will have to pay to type  $\theta_h$  to



keep her choosing the contract intended for her. It is the reason that the low type's quality is distorted downward from the quality she would be offered if types were observable.

### 6.2.2 A continuum of types

The techniques used to study the continuum-of-types model are fundamental tools in mechanism design.

$\Theta = [0, 1]$	set of types
$F, f$	cdf and pdf of type distribution; $f > 0$
$t$	transfer paid by the agent to the principal
$u(q, \theta)$	agent's consumption utility
	twice differentiable
	$u(0, \theta) = 0$
	$\frac{\partial}{\partial q} u(q, \theta) \geq 0, > 0$ when $\theta > 0$
	$\frac{\partial}{\partial \theta} u(q, \theta) \geq 0, > 0$ when $\theta > 0$ , bounded on compact sets
	$\frac{\partial^2}{\partial q \partial \theta} u(q, \theta) > 0$ for $q, \theta > 0$ (single crossing)
$u(q, \theta) - t$	agent's total utility
$c(\cdot)$	principal's cost function; nondecreasing, $c(0) = 0$
$t - c(q)$	principal's utility

#### ***The principal's problem***

The principal offers a menu of contracts  $(q(\cdot), t(\cdot)) = \{(q(\theta), t(\theta))\}_{\theta \in \Theta}$ .

$$\max_{q(\cdot), t(\cdot)} \int_0^1 (t(\theta) - c(q(\theta))) f(\theta) d\theta \quad \text{subject to}$$

$$(IC) \quad u(q(\theta), \theta) - t(\theta) \geq u(q(\hat{\theta}), \theta) - t(\hat{\theta}) \quad \text{for all } \theta, \hat{\theta} \in \Theta$$

$$(IR) \quad u(q(\theta), \theta) - t(\theta) \geq 0 \quad \text{for all } \theta \in \Theta$$

We sometimes write  $(IC_{\hat{\theta}|\theta})$  and  $(IR_{\theta})$  for type  $\theta$ 's constraints.

Note that  $q$  and  $t$  are being used both as variables and as functions. For the latter case, the notation  $q(\cdot)$  and  $t(\cdot)$  will help avoid ambiguity. We call refer to these as the *menu of qualities* and the *menu of transfers*, respectively.

**Payoff equivalence: characterizing incentive compatible menus**

Theorem 6.11 characterizes the incentive compatible menus of contracts. It follows from three lemmas.

**Lemma 6.8.** *Fix  $q(\cdot)$ . If there is a  $t(\cdot)$  such that  $(q(\cdot), t(\cdot))$  satisfies (IC), then  $q(\cdot)$  is nondecreasing.*

*Proof.* Let  $\theta_1 > \theta_0$ . By (IC),

$$(48) \quad u(q(\theta_1), \theta_1) - t(\theta_1) \geq u(q(\theta_0), \theta_1) - t(\theta_0),$$

$$(49) \quad -(u(q(\theta_1), \theta_0) - t(\theta_1)) \geq -(u(q(\theta_0), \theta_0) - t(\theta_0)).$$

Add:

$$(50) \quad u(q(\theta_1), \theta_1) - u(q(\theta_1), \theta_0) \geq u(q(\theta_0), \theta_1) - u(q(\theta_0), \theta_0).$$

Rewrite in integral form:

$$(51) \quad \int_{\theta_0}^{\theta_1} \frac{\partial}{\partial \theta} u(q(\theta_1), \hat{\theta}) d\hat{\theta} \geq \int_{\theta_0}^{\theta_1} \frac{\partial}{\partial \theta} u(q(\theta_0), \hat{\theta}) d\hat{\theta}.$$

If  $q(\theta_1) < q(\theta_0)$ , then single crossing would imply that  $\frac{\partial}{\partial \theta} u(q(\theta_1), \hat{\theta}) < \frac{\partial}{\partial \theta} u(q(\theta_0), \hat{\theta})$  for all  $\hat{\theta}$ , contradicting (51). ■

Let

$$(52) \quad U(\theta) = u(q(\theta), \theta) - t(\theta)$$

be type  $\theta$ 's payoff from her intended contract—in other words, type  $\theta$ 's information rent.

**Lemma 6.9.** *If  $(q(\cdot), t(\cdot))$  satisfies (IC), then for some  $U(0) \in \mathbb{R}$ ,*

$$(53) \quad U(\theta) = U(0) + \int_0^\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} \text{ for all } \theta \in \Theta, \text{ or equivalently}$$

$$(54) \quad t(\theta) = u(q(\theta), \theta) - \int_0^\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} - U(0) \text{ for all } \theta \in \Theta.$$

Lemma 6.9 says that incentive compatibility and the menu of qualities determine all types' payoffs / information rents (or equivalently, all types' transfers) up to the choice of the additive constant  $U(0)$ .

It also shows that payoffs are nondecreasing in type:  $U'(\theta) = \frac{\partial}{\partial \theta} u(q(\theta), \theta) \geq 0$ .

In fact, up to technicalities (see below), equation (53) says that

$$(55) \quad U'(\theta) = \frac{\partial}{\partial \theta} u(q(\theta), \theta).$$

Here is a non-rigorous derivation of (55): By definition (52) and (IC),

$$U(\theta) = \max_{\hat{\theta} \in \Theta} (u(q(\hat{\theta}), \theta) - t(\hat{\theta})) = u(q(\hat{\theta}^*(\theta)), \theta) - t(\hat{\theta}^*(\theta)) = u(q(\theta), \theta) - t(\theta).$$

So by the envelope theorem,  $U'(\theta) = \frac{\partial}{\partial \theta} u(q(\theta), \theta)$ .

What was not rigorous? First, to apply the usual envelope theorem at value  $\theta$ , we need for various smoothness and continuity conditions to hold at  $\theta$ , but this is not always the case. (For instance, the usual envelope theorem requires that  $\frac{\partial}{\partial \theta} (u(q(\hat{\theta}), \theta) - t(\hat{\theta}))$  be continuous in  $(\theta, \hat{\theta})$ . Lemma 6.8 tells us that  $q(\cdot)$  is nondecreasing, which implies that it is differentiable almost everywhere, but it may have discontinuities, and these may be passed on to  $\frac{\partial}{\partial \theta} (u(q(\hat{\theta}), \theta) - t(\hat{\theta}))$ .) Second, the integration requires the fundamental theorem of calculus to hold. The necessary and sufficient condition for this is that  $U(\cdot)$  be absolutely continuous (see, e.g., Folland (1999)); this property of  $U(\cdot)$  must be established, either directly or indirectly.

The integral envelope theorem of Milgrom and Segal (2002) (see also Milgrom (2004, ch. 3)) gives conditions under which the integral formula (53) for  $U(\cdot)$  must hold, establishing absolute continuity of  $U(\cdot)$  along the way.

Here is a direct proof of Lemma 6.9:

*Proof.* Since  $\frac{\partial}{\partial \theta} u(q, \theta)$  is bounded on compact sets, there is a constant  $M < \infty$  such that

$$0 \leq \frac{\partial}{\partial \theta} u(q, \theta) \leq M \text{ for } (q, \theta) \in [q(0), q(1)] \times \Theta.$$

Since  $q(\cdot)$  is nondecreasing by Lemma 6.8,  $q(\theta) \in [q(0), q(1)]$  for all  $\theta \in \Theta$ , and so

$$(56) \quad 0 \leq \frac{\partial}{\partial \theta} u(q(\theta), \theta) \leq M \text{ for all } \theta \in \Theta.$$

Let  $\theta_1 > \theta_0$ . Since  $U(\theta_1)$  is the left-hand side of (48), and since  $-U(\theta_0)$  is the right-hand side of (49), we can augment (50) as follows:

$$u(q(\theta_1), \theta_1) - u(q(\theta_1), \theta_0) \geq U(\theta_1) - U(\theta_0) \geq u(q(\theta_0), \theta_1) - u(q(\theta_0), \theta_0).$$

Thus the mean value theorem implies that for some  $\bar{\theta}_0, \bar{\theta}_1 \in (\theta_0, \theta_1)$ ,

$$\frac{\partial}{\partial \theta} u(q(\theta_1), \bar{\theta}_1)(\theta_1 - \theta_0) \geq U(\theta_1) - U(\theta_0) \geq \frac{\partial}{\partial \theta} u(q(\theta_0), \bar{\theta}_0)(\theta_1 - \theta_0).$$

Bound (56) then implies that  $U$  is Lipschitz continuous, and thus absolutely continuous. Furthermore,

$$(57) \quad \frac{\partial}{\partial \theta} u(q(\theta_1), \bar{\theta}_1) \geq \frac{U(\theta_1) - U(\theta_0)}{\theta_1 - \theta_0} \geq \frac{\partial}{\partial \theta} u(q(\theta_0), \bar{\theta}_0).$$

Since  $q(\cdot)$  is nondecreasing by Lemma 6.8, it is continuous almost everywhere. If  $\theta$  is a point of continuity of  $q(\cdot)$ , then as  $\theta_0$  and  $\theta_1$  approach  $\theta$ , the left and right expressions in (57) approach  $\frac{\partial}{\partial \theta} u(q(\theta), \theta)$ . Thus the limit of the middle expression exists and equals this same value. This yields (55) for almost all  $\theta$ . Integrating then yields (53). ■

**Lemma 6.10.** *If  $q(\cdot)$  is nondecreasing and  $t(\cdot)$  is given by (54) for some  $U(0) \in \mathbb{R}$ , then  $(q(\cdot), t(\cdot))$  satisfies (IC).*

*Proof.* Since  $t(\cdot)$  is given by (54), definition (52) of  $U(\theta)$  implies that  $U(\theta)$  satisfies (53).

If  $\theta_1 > \theta_0$ , then (53) and the single crossing property imply that

$$U(\theta_1) - U(\theta_0) = \int_{\theta_0}^{\theta_1} \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} \geq \int_{\theta_0}^{\theta_1} \frac{\partial}{\partial \theta} u(q(\theta_0), \hat{\theta}) d\hat{\theta} = u(q(\theta_0), \theta_1) - u(q(\theta_0), \theta_0).$$

Thus

$$U(\theta_1) \geq u(q(\theta_0), \theta_1) - (u(q(\theta_0), \theta_0) - U(\theta_0)) = u(q(\theta_0), \theta_1) - t(\theta_0),$$

which is (IC) <sub>$\theta_0|\theta_1$</sub> . A careful reading reveals that the argument remains valid if  $\theta_1 < \theta_0$ . Thus (IC) holds. ■

In summary, Lemma 6.8 says that in solving the principal's problem, we need only consider nondecreasing  $q(\cdot)$ , Lemma 6.9 specifies the implied information rents  $U(\cdot)$  up to the choice of the constant  $U(0)$ , and Lemma 6.10 says that we must consider all nondecreasing  $q(\cdot)$ . Putting these together, we have:

**Theorem 6.11** (Payoff equivalence).

*Menu  $(q(\cdot), t(\cdot))$  satisfies (IC) if and only if*

- (i)  $q(\cdot)$  is nondecreasing, and

$$(ii) \quad U(\theta) = U(0) + \int_0^\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} \text{ for all } \theta \in \Theta.$$

Equivalently,  $(q(\cdot), t(\cdot))$  satisfies (IC) if and only if (i) holds and transfers are given by

$$(ii') \quad t(\theta) = u(q(\theta), \theta) - \int_0^\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} - U(0) \text{ for all } \theta \in \Theta.$$

### ***Simplifying the principal's problem***

We now reintroduce the individual rationality constraints.

**Lemma 6.12.** *If  $(q(\cdot), t(\cdot))$  satisfies (IC) and  $(IR_0)$ , then it satisfies  $(IR_\theta)$  for all  $\theta > 0$ .*

*Proof.* Use  $(IC_{0|\theta})$  and the fact that  $\frac{\partial}{\partial \theta} u(q, \theta) \geq 0$ :

$$u(q(\theta), \theta) - t(\theta) \geq u(q(0), \theta) - t(0) \geq u(q(0), 0) - t(0) \geq 0. \blacksquare$$

After taking advantage of the lemmas, the principal's problem reduces to

$$(58) \quad \begin{aligned} & \max_{q(\cdot), U(0)} \int_0^1 \left( u(q(\theta), \theta) - \int_0^\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} - U(0) - c(q(\theta)) \right) f(\theta) d\theta \\ & \text{subject to } q(\cdot) \text{ nondecreasing} \\ & \quad U(0) \geq 0 \end{aligned}$$

Clearly  $U(0) = 0$ : the lowest type gets no information rent.

To put the objective function in a more pliable form, reverse the order of integration in the middle term:

$$(59) \quad \begin{aligned} \int_0^1 \left( \int_0^\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta} \right) f(\theta) d\theta &= \int_0^1 \left( \int_{\hat{\theta}}^1 f(\theta) d\theta \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) \right) d\hat{\theta} \\ &= \int_0^1 (1 - F(\hat{\theta})) \frac{\partial}{\partial \theta} u(q(\hat{\theta}), \hat{\theta}) d\hat{\theta}. \end{aligned}$$

Thus since  $f(\cdot)$  is positive, we can rewrite the principal's problem again as

$$(60) \quad \begin{aligned} & \max_{q(\cdot)} \int_0^1 \left( u(q(\theta), \theta) - \frac{1 - F(\theta)}{f(\theta)} \frac{\partial}{\partial \theta} u(q(\theta), \theta) - c(q(\theta)) \right) f(\theta) d\theta \\ & \text{subject to } q(\cdot) \text{ nondecreasing} \end{aligned}$$

The function  $\frac{f(\cdot)}{1-F(\cdot)}$  is called the *hazard rate* of distribution  $F$ . Its reciprocal is the *inverse*

hazard rate.


The expression

$$(61) \quad u(q, \theta) - \frac{1 - F(\theta)}{f(\theta)} \frac{\partial}{\partial \theta} u(q, \theta)$$

from the integrand is called the *virtual utility* of type  $\theta$  (cf. (47)). In the relaxed problem without the requirement that  $q(\cdot)$  be nondecreasing, the principal chooses  $q(\theta)$  to maximize the difference between (61) and the production cost  $c(q(\theta))$ . (This difference is known as the *virtual surplus* from type  $\theta$ .) The second term in (61) accounts for the effect of the choice of  $q(\theta)$  on the information rents that must be paid to types above  $\theta$ , as discussed next.

*Remark 6.13.* To interpret virtual utility (61), it is easiest to imagine that the agent is drawn from a population in which the type density is  $f$ .

We first interpret the reversal in the order of integration in (59). This is an accounting trick that reassigns information rents from the types who earn them to the types who necessitate their payment. In the initial expression in (59), the expression in parentheses is the information rent that an agent of type  $\theta$  obtains because of the provision of quality to *lower* types  $\hat{\theta} < \theta$  (the vertical line in the figure at left). The outer integral totals this over all types  $\theta$ , so that the region of integration is the shaded triangle. After we reverse the order of integration, the integrand of the outer integral (in parentheses) is the information rent that provision of quality to agents of type  $\hat{\theta}$  generates for agents of *higher* types  $\theta > \hat{\theta}$  (the horizontal line in the figure at right). When we substitute the final expression in (59) into the objective function in (60), we change the name of the variable of integration from  $\hat{\theta}$  to  $\theta$ . Thus consumption benefits are indexed by the type that *receives* them, but information rents are indexed by the type that *generates* them.



We next factor out  $f(\theta)$  from (60). There is already an  $f(\theta)$  in the consumption benefit term, since this term is indexed by who receives the benefit, and there are  $f(\theta)$  agents of

type  $\theta$ . But as explained above, the information rent term is indexed by who generates the rent. Here the density term, which is tied to the agents receiving the rent, has been integrated away. To factor out an  $f(\theta)$  that is not there, we divide by it. Doing so converts  $1 - F(\theta)$ , the information rents attributed to agents of type  $\theta$  as a group, into  $(1 - F(\theta))/f(\theta)$ , the information rent attributed to each agent of type  $\theta$ . This factoring out  $f(\theta)$  gives us the virtual utility (61).

***Solving the principal's problem when utility is linear in type***

Problem (60) requires optimization over a set of functions, and so is an optimal control problem. However, if choosing each  $q(\theta)$  to maximize the integrand leads to a non-decreasing  $q(\cdot)$ , this choice of  $q(\cdot)$  is optimal. We consider environments in which the agent's utility is linear in her type.

$$\begin{aligned}
 & q \in Q, \text{ where } Q = [0, 1] \text{ or } Q = [0, \infty) \\
 & u(q, \theta) = v(q)\theta, \\
 & v(0) = 0, v'(\cdot) \text{ positive and nonincreasing (with } v'(0) < \infty) \\
 & c(0) = 0, c'(\cdot) \text{ nonnegative and nondecreasing} \\
 & \psi(\theta) \equiv \theta - \frac{1 - F(\theta)}{f(\theta)} \text{ increasing and continuous}
 \end{aligned}$$

(When  $u(q, \theta) = q\theta$ ,  $\psi(\theta)$  is called type  $\theta$ 's *virtual valuation*; see Section 7.3.)

Under these assumptions, program (60) becomes

$$\begin{aligned}
 (62) \quad & \max_{q(\cdot)} \int_0^1 (v(q(\theta))\psi(\theta) - c(q(\theta))) f(\theta) d\theta \\
 & \text{subject to } q(\cdot) \text{ nondecreasing}
 \end{aligned}$$

The assumption that  $\psi(\cdot)$  is increasing holds if the hazard rate  $\frac{f(\cdot)}{1-F(\cdot)}$  is nondecreasing. This *monotone hazard rate* condition is satisfied by many common distributions; see Bagnoli and Bergstrom (2005).

Since  $\psi(\cdot)$  is increasing and continuous with  $\psi(0) < 0 < 1 = \psi(1)$ , it has a unique zero  $\theta^* \in (0, 1)$ .

**Proposition 6.14.** *Suppose that  $Q = [0, 1]$ . Then the optimal menu of qualities satisfies*

$$\begin{aligned}
 & \text{if } \psi(\theta) < \frac{c'(0)}{v'(0)}, \text{ let } q(\theta) = 0, \\
 & \text{if } \psi(\theta) \in \left[ \frac{c'(0)}{v'(0)}, \frac{c'(1)}{v'(1)} \right], \text{ let } q(\theta) \text{ be a solution to } \psi(\theta) = \frac{c'(q(\theta))}{v'(q(\theta))}, \\
 & \text{if } \psi(\theta) > \frac{c'(1)}{v'(1)}, \text{ let } q(\theta) = 1.
 \end{aligned}$$

(Multiple optimal choices of  $q(\theta)$  are possible if there is an interval of qualities  $q$  over which  $\frac{c'(q)}{v'(q)}$  is constant and equal to  $\psi(\theta)$ .)

*Proof.* By the assumptions on  $c'(\cdot)$  and  $v'(\cdot)$ ,

$$(63) \quad \frac{d}{dq} (v(q)\psi(\theta) - c(q)) = v'(q)\psi(\theta) - c'(q) \text{ is nonincreasing in } q.$$

Thus for a type  $\theta$  with  $\psi(\theta) < \frac{c'(0)}{v'(0)}$ , the integrand from (62) is maximized at  $q = 0$ .

If  $\psi(\theta) \in [\frac{c'(0)}{v'(0)}, \frac{c'(1)}{v'(1)}]$ , a  $q(\theta)$  that satisfies  $\psi(\theta) = \frac{c'(q(\theta))}{v'(q(\theta))}$  also satisfies the first-order condition for maximizing the integrand from (62), and this condition is sufficient for maximization because of (63). Also, since  $\frac{c'(\cdot)}{v'(\cdot)}$  is nondecreasing,  $q(\theta)$  is increasing in  $\theta$  over this range.

Finally, if  $\psi(\theta) > \frac{c'(1)}{v'(1)}$ , then the derivative in (63) is still positive at  $q = 1$ , leading to the corner solution  $q(\theta) = 1$ . ■

*Example 6.15.* Suppose that  $Q = [0, 1]$ ,  $c(q) \equiv 0$ , and  $v(q) = q$ . Then since  $\frac{c'(q)}{v'(q)} \equiv 0$ , the optimal menu of qualities is

$$q(\theta) = \begin{cases} 0 & \text{if } \psi(\theta) < 0, \\ 1 & \text{if } \psi(\theta) > 0. \end{cases}$$

(One can choose  $q(\psi^{-1}(0))$  arbitrarily.) The corresponding menu of transfers is

$$t(\theta) = q(\theta)\theta - \int_0^\theta q(\hat{\theta}) d\hat{\theta} = q(\theta)\theta - \int_{\psi^{-1}(0)}^\theta 1 d\hat{\theta} = \begin{cases} 0 & \text{if } \psi(\theta) < 0, \\ \psi^{-1}(0) & \text{if } \psi(\theta) > 0. \end{cases} \quad \blacklozenge$$

**Proposition 6.16.** Suppose that  $Q = [0, \infty)$  and that  $c'(\bar{q}) \geq v'(\bar{q})$  for some  $\bar{q} \geq 0$ . Then the optimal menu of qualities satisfies

$$\begin{aligned} &\text{if } \psi(\theta) < \frac{c'(0)}{v'(0)}, \text{ let } q(\theta) = 0, \\ &\text{if } \psi(\theta) \geq \frac{c'(0)}{v'(0)}, \text{ let } q(\theta) \text{ be a solution to } \psi(\theta) = \frac{c'(q(\theta))}{v'(q(\theta))}. \end{aligned}$$

*Proof.* Because  $\frac{c'(\cdot)}{v'(\cdot)}$  is continuous and nondecreasing, there are two ways that the condition on  $c'(\cdot)$  can hold: either  $c'(0) > v'(0)$ , or  $c'(\bar{q}) = v'(\bar{q})$  for some  $\bar{q} \geq 0$ . In the former case it is optimal to choose  $q(\theta) \equiv 0$ .

If instead  $c'(\bar{q}) = v'(\bar{q})$ , then since  $\psi(1) = 1$ , setting  $q(1) = \bar{q}$  is optimal, and because  $\frac{c'(\cdot)}{v'(\cdot)}$  is nondecreasing and  $\psi(\cdot)$  is increasing, the condition also ensures that optimal quantities  $q(\theta)$  exist for all types  $\theta < 1$  and that  $q(\cdot)$  is increasing in  $\theta$  once  $\psi(\theta) \geq \frac{c'(0)}{v'(0)}$ . ■



Remarks:

- (i) If the agent's type  $\theta > 0$  were observable, the principal would maximize  $\theta v(q) - aq$  by solving  $v'(q) = a/\theta$ . Here he instead either solves  $v'(q) = a/(\theta - \frac{1-F(\theta)}{f(\theta)})$  (if  $\theta > \theta^*$ ) or chooses  $q = 0$  (otherwise). Thus only the highest type receives her optimal quality, all other types above  $\theta^*$  receive lower-than-optimal qualities, and types  $\theta^*$  and below are not served.
- (ii) If the assumption on the hazard rate fails, the allocation function  $q(\cdot)$  obtained by pointwise maximization of virtual surplus may not be nondecreasing. In this case, one must use further arguments to find the optimal  $q(\cdot)$ . This  $q(\cdot)$  will have flat spots, reflecting *bunching* of intervals of types. This can happen when the hazard rate  $f(\cdot)/(1 - F(\cdot))$  has a decreasing segment: if there are relatively few agents with types near  $\theta$ , it may be optimal to extract less from these types in order to lower the information rents of higher types. The term *ironing* is sometimes used to describe the form of the optimal  $q(\cdot)$ .

To address these problems, Myerson (1981) (see also Baron and Myerson (1982) and Toikka (2011)) introduce a *generalized virtual utility* that incorporates the distortions generated by the allocation to type  $\theta$  on the allocations to other types generated by the monotonicity requirement on the allocation function  $q(\cdot)$ . Using convex analysis arguments, they show that pointwise maximization of generalized virtual surplus generates the optimal monotone allocation function.

## 7. Mechanism Design

Mechanism design considers the use of games to elicit information from groups of privately informed agents, whether to extract surplus from the agents, to ensure efficient social choices, or to achieve other ends. Instead of optimal choice by a single agent, the designer anticipates (some kind of) equilibrium behavior among the multiple agents. (The designer himself is not a part of this equilibrium; instead he is assumed to have the power to commit to the mechanism of his choice.)

Some basic questions:

- (i) characterization of implementable social choice functions (where "social choice function" and "implementable" will be defined shortly)
- (ii) revenue maximization (often called "optimality")
- (iii) ensuring allocative efficiency

Revenue maximization models have strong commonalities with the principal-agent prob-

lem with adverse selection, and it is accurate to describe them as “principal - many agent problems”. Models focusing on allocative efficiency have a different flavor: for instance, it is more natural to think of the mechanism designer as a planner or as the agents as a group rather than as a self-interested principal. But we shall see that there is significant overlap in how problems with different objectives are analyzed.

The father of mechanism design is Hurwicz (1960, 1972). Much of our analysis focuses on *Bayesian implementation*, under which agents play a Bayesian equilibrium of the Bayesian game introduced by the planner. Myerson (1979, 1981) are key early references.

## 7.1 Mechanism Design Environments and the Revelation Principle

### 7.1.1 Mechanism design environments and social choice functions

A *mechanism design environment* (or *Bayesian collective choice environment*)  $\mathcal{E}$  is defined by the following objects:

- $\mathcal{A} = \{1, \dots, n\}$  set of agents
- $\Theta_i$  set of agent  $i$ 's types
- $\mu \in \Delta\Theta$  common prior distribution
- $\mathcal{X}$  set of *social alternatives* (or *outcomes*)
- $v_i: \mathcal{X} \times \Theta \rightarrow \mathbb{R}$  agent  $i$ 's utility function

Environment  $\mathcal{E}$  has *private values* if we can write  $v_i(\chi, \theta) = v_i(\chi, \theta_i)$ , so that each agent's utilities do not depend on other agents' types.

A *social choice function*  $g: \Theta \rightarrow \mathcal{X}$  assigns to each type profile  $\theta$  a social alternative  $g(\theta) \in \mathcal{X}$ . Social choice functions can be used to describe the designer's aims at the ex ante stage.

*Example 7.1. Assignment of distinct objects without monetary transfers.* There is a set  $Z$  consisting of  $n$  distinct prizes. Agent  $i$ 's type  $\theta_i$  is an element of  $\mathbb{R}^n$ , where  $\theta_{iz}$  is agent  $i$ 's benefit from receiving prize  $z$ . The distribution  $\mu$  is the common prior over type profiles.

A social alternative  $\chi \in \mathcal{X}$  is an assignment of one prize to each agent. Agent  $i$ 's utility is  $v_i(\chi, \theta) = \theta_{iz}$  when he receives prize  $z$  under  $\chi$ . Thus this environment has private values.

A social choice function  $g$  specifies an assignment  $g(\theta) \in \mathcal{X}$  for each type profile  $\theta \in \Theta$ . For example,  $g(\theta)$  could be an assignment that maximizes the sum of the agents' utilities given their types. ♦

### Quasilinear environments

Starting in Section 7.2, we will focus on *quasilinear environments*. These environments incorporate monetary transfers whose effects on utility are additively separable.

$$\mathcal{X} = X \times \mathbb{R}^n, X \text{ finite}$$

$$\chi = (x, t)$$

$$x \in X \quad \text{an allocation (broadly defined)}$$

$$t_i \in \mathbb{R} \quad \text{agent } i\text{'s monetary transfer}$$

$$u_i(x, \theta) \quad \text{agent } i\text{'s allocation utility (or consumption utility)}$$

$$v_i(\chi, \theta) = u_i(x, \theta) - t_i$$

Allocation  $x$  is *efficient* for type profile  $\theta$  if

$$x \in \operatorname{argmax}_{y \in X} \sum_{i \in \mathcal{A}} u_i(y, \theta).$$

(We will worry about transfers later—see Section 7.4.)

In quasilinear environments, social choice functions are of the form  $g(\cdot) = (x(\cdot), t(\cdot))$ , where  $x: \Theta \rightarrow X$  is an *allocation function*, and  $t: \Theta \rightarrow \mathbb{R}^n$  is a *transfer function*. Allocation function  $x^*(\cdot)$  is *ex post efficient* if for each type profile  $\theta \in \Theta$ ,  $x^*(\theta)$  is efficient for  $\theta$ .

*Example 7.2. Allocation of an indivisible good.* The set of allocations is  $X = \mathcal{A}$  (who gets the good), or  $X = \mathcal{A} \cup \{0\}$  (allowing the good not to be allocated). We assume private values, with types representing valuations for the good. Thus agent  $i$ 's consumption utility is described by  $u_i(i, \theta_i) = \theta_i$  and  $u_i(j, \theta_i) = 0$  for  $j \neq i$ , and his total utility from allocation  $\chi = (j, t)$  is  $v_i(\chi, \theta) = u_i(j, \theta_i) - t_i$ . Allocation function  $x$  is *ex post efficient* if  $x(\theta) \in \operatorname{argmax}_j \theta_j$  for all  $\theta \in \Theta$ , so that the good is assigned to someone who values it most. ♦

We will see that analyses are often made easier by allowing *randomized allocations*:

$$\mathcal{X} = \Delta X \times \mathbb{R}^n, X \text{ finite}$$

$$\chi = (q, t)$$

$$q_x \in [0, 1] \quad \text{the probability of allocation } x$$

$$v_i(\chi, \theta) = \sum_{x \in X} q_x u_i(x, \theta) - t_i$$

Randomized allocation  $q$  is *efficient* for type profile  $\theta$  if

$$q \in \operatorname{argmax}_{p \in \Delta X} \sum_{i \in \mathcal{A}} \sum_{x \in X} p_x u_i(x, \theta).$$

For the usual reasons (cf. Proposition 1.14), randomized allocation  $q$  is efficient for  $\theta$  if and only if every allocation  $x$  in its support is efficient for  $\theta$ .

Now social choice functions are of the form  $g(\cdot) = (q(\cdot), t(\cdot))$ , where  $q: \Theta \rightarrow \Delta X$ . Allocation function  $q^*(\cdot)$  is *ex post efficient* if for each type profile  $\theta \in \Theta$ ,  $q^*(\theta)$  is efficient for  $\theta$ .

*Example 7.3. Randomized allocation of an indivisible good.* As in Example 7.2,  $X = \mathcal{A}$  or  $X = \mathcal{A} \cup \{0\}$ ,  $u_i(i, \theta_i) = \theta_i$ , and  $u_i(j, \theta_i) = 0$  for  $j \neq i$ . Agent  $i$ 's expected utility from randomized allocation  $\chi = (q, t)$  is  $v_i(\chi, \theta_i) = q_i \theta_i - t_i$ . An allocation function is *ex post efficient* if for all  $\theta \in \Theta$ ,  $q_k(\theta) > 0$  implies that  $k \in \operatorname{argmax}_j \theta_j$ , so that the good is always assigned to someone who values it most. ♦

*Example 7.4. Linear utility.* Private good allocation (Example 7.3) and public good provision (Example 7.22) are instances of quasilinear environments with private values and *linear utility*, which here means linearity in own type:

$$u_i(x, \theta_i) = a_i(x)\theta_i + c_i(x).$$

Thus an agent's expected utility from randomized allocation  $\chi = (q, t)$  is

$$v_i(\chi, \theta_i) = \sum_{x \in X} q_x (a_i(x)\theta_i + c_i(x)) - t_i. \quad \blacklozenge$$

### ***Independent private values environments***

Often we will consider environments with *independent private values*. This means that different agents' types are drawn independently ( $\mu$  is a product distribution), and that each agent's payoff does not depend on other agents' types ( $v_i(\chi, \theta) = v_i(\chi, \theta_i)$ ).

The independent private values environment provides a natural baseline, and it allows for many powerful results. But in many applications, types are correlated, or values are interdependent, or both. Results about optimality and efficiency in these settings sometimes differ markedly from those in the independent private values case. We focus on independent private values in Sections 7.2–7.4. We then consider interdependent values and correlated types in Sections 7.5 and 7.6.

#### 7.1.2 Mechanisms

Given a mechanism design environment  $\mathcal{E} = \{\mathcal{A}, \{\Theta_i\}_{i \in \mathcal{A}}, \mu, \mathcal{X}, \{v_i\}_{i \in \mathcal{A}}\}$ , a *mechanism* is a collection  $\mathcal{M} = \{\{A_i\}_{i \in \mathcal{A}}, \gamma\}$ , where  $A_i$  is agent  $i$ 's *action set* and  $\gamma: A \rightarrow \mathcal{X}$  is a *decision*

function. Together,  $\mathcal{E}$  and  $\mathcal{M}$  define a Bayesian game with Bernoulli utility functions

$$u_i(a, \theta) = v_i(\gamma(a), \theta).$$

We typically take  $\mathcal{E}$  for granted and identify the Bayesian game with the mechanism  $\mathcal{M}$ .

*Example 7.5. First price auctions.* In Example 7.3, suppose that agents' types are independent draws from a distribution with density  $f$  on  $[0, 1]$ . This is the independent private values environment for auctions from Section 4.5.1.

Since preferences are quasilinear, a mechanism in this environment has a decision function of the form  $\gamma(\cdot) = (\rho(\cdot), \tau(\cdot))$ , where  $\rho: A \rightarrow \Delta X$  and  $\tau: A \rightarrow \mathbb{R}^n$ .

To express a first price auction using the definitions above, we let  $A_i = [0, 1]$  be the set of agent  $i$ 's bids. Ignoring the case of a tied high bid, we can define the decision function  $\gamma(\cdot)$  as follows:

$$\text{if } \{i\} = \operatorname{argmax}_{j \in \mathcal{A}} a_j, \text{ then } \rho_i(a) = 1, \tau_i(a) = a_i, \text{ and } \tau_j(a) = 0 \text{ for all } j \neq i.$$

The utility functions for the Bayesian game are


$$u_i(a, \theta) = v_i(\gamma(a), \theta) = \rho_i(a)\theta_i - \tau_i(a).$$

This is just as in Example 7.3, except that the allocation probabilities and transfers are now determined as functions of the bid profile. Combining the last two displays shows that

$$\text{if } \{i\} = \operatorname{argmax}_{j \in \mathcal{A}} a_j, \text{ then } u_i(a, \theta) = \theta_i - a_i, \text{ and } u_j(a, \theta) = 0 \text{ for all } j \neq i. \blacklozenge$$

As usual, a pure Bayesian strategy for agent  $i$  under mechanism  $\mathcal{M}$  is a map  $s_i: \Theta_i \rightarrow A_i$ .

The following diagram compares a social choice function  $g$  to a mechanism  $\mathcal{M}$ :



A social choice function  $g: \Theta \rightarrow \mathcal{X}$  sends type profiles to social alternatives.

But under the mechanism  $\mathcal{M}$ , going from  $\Theta$  to  $\mathcal{X}$  takes two steps. First, the agents choose strategy profile  $s: \Theta \rightarrow A$ , which for each type profile  $\theta$  determines an action profile  $s(\theta)$ .

Then, the decision function  $\gamma: A \rightarrow \mathcal{X}$  uses the action profile  $s(\theta)$  to determine the social alternative  $\gamma(s(\theta))$ .

### 7.1.3 Implementation


We now introduce two notions of implementation of social choice functions.

Mechanism  $\mathcal{M}$  *Bayesian implements* social choice function  $g$  if there is a Bayesian equilibrium  $s^*$  of  $\mathcal{M}$  such that for every type profile  $\theta \in \Theta$ ,  $\gamma(s^*(\theta)) = g(\theta)$ .

Mechanism  $\mathcal{M}$  *dominant strategy implements* social choice function  $g$  if there is a profile  $s^*$  of very weakly dominant strategies (see (23)) of  $\mathcal{M}$  such that for every type profile  $\theta \in \Theta$ ,  $\gamma(s^*(\theta)) = g(\theta)$ .

Clearly, the latter notion of implementation is more demanding than the former.

If  $\mathcal{M} = \{A_i\}_{i \in \mathcal{A}}, \gamma\}$  implements  $g$  in either sense, then the specified equilibrium strategy profile  $s^*$  makes the following diagram commute (meaning that either path from  $\theta \in \Theta$  to  $\mathcal{X}$  leads to the same  $\chi$ ):



*Example 7.6. First price auctions revisited.* As in Example 7.5, consider a first price auction in an environment with symmetric independent private values drawn from  $\Theta_i = [0, 1]$ . Proposition 4.10 showed that in the unique symmetric equilibrium, all agents use the increasing bidding function  $b^*(v) = \mathbb{E}(V_{(n-1)}^{n-1} | V_{(n-1)}^{n-1} \leq v)$ . Because  $b^*(\cdot)$  is increasing, the good is allocated to a bidder who values it most. Thus (again ignoring ties), the first-price auction Bayesian implements a social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$  satisfying

$$\text{if } \{i\} = \underset{j \in \mathcal{A}}{\operatorname{argmax}} \theta_j, \text{ then } (q_i(\theta), t_i(\theta)) = (1, b^*(\theta_i)), \text{ and } (q_j(\theta), t_j(\theta)) = (0, 0) \text{ for } j \neq i.$$

Thus the allocation function  $q(\cdot)$  implemented by a first price auction is ex post efficient.

◆

Remarks:

- (i) The definition of Bayesian implementation requires that there be *some* equilibrium of the mechanism that sustains the social choice function; there may be other equilibria that do not. We can think of the designer both proposing the mechanism and suggesting the equilibrium to be played.

- (ii) A stronger requirement is that *every* equilibrium of  $\mathcal{M}$  sustains the social choice function. This requirement is called *full implementation*, and this general topic is known (somewhat confusingly) as *implementation theory*. Two key references here are Postlewaite and Schmeidler (1986) and Jackson (1991). See Palfrey (1992) for a survey of the early literature. We discuss some recent results in Section 7.6.4.
- (iii) Starting in Section 7.2 we will focus on quasilinear environments, which suffice for many economic applications. Another branch of the literature instead asks about implementation when general preferences over the set of alternatives  $\mathcal{X}$  are allowed. The basic result here is the *Gibbard-Satterthwaite theorem* (Gibbard (1973), Satterthwaite (1975)) which states that if each preference over alternatives is possible for every agent, and if the social choice function  $g: \Theta \rightarrow \mathcal{X}$  is onto (i.e., if each alternative is assigned under some type profile), then  $g$  is dominant strategy implementable if and only if it is *dictatorial*, meaning that there is a single agent  $i$  whose favorite alternative is always chosen. This result is a descendant of the *Arrow impossibility theorem* (Arrow (1951)), the foundational result of social choice theory.
- (iv) Here we consider implementation in Bayesian environments, where each agent is uncertain about others' types when choosing his action. A complementary branch of the literature considers implementation in environments with *complete information*, meaning that the agents (though not the principal) know one another's types when choosing their actions. The basic notion of implementation supposes that agents play a Nash equilibrium of the realized complete information game. With three or more agents, any social choice function can be implemented in the weak sense (as described in (i)), so the interesting question is whether full implementation (as in (ii)) is possible. This is commonly referred to as *Nash implementation*. The fundamental work here is due to (Maskin (1999) (first version 1977)), who introduced a condition on social choice functions (now called (*Maskin*) *monotonicity*) which is necessary and almost sufficient for Nash implementation. For a survey of the early literature, see Moore (1992).

#### 7.1.4 Direct mechanisms and the revelation principle


We now show that one can focus on mechanisms of a particularly simple kind.

A *direct mechanism* (or *revelation mechanism*) is a mechanism of the form  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$ . Thus the actions are type announcements, often denoted  $\hat{\theta}_i$ . The decision function  $g: \Theta \rightarrow \mathcal{X}$  is therefore a social choice function. However, when the mechanism is run this function

is applied to the profile of type announcements.

We sometimes refer to a social choice function  $g$  as a mechanism, and when we do so we mean the direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$ . This will prove quite convenient, since by virtue of Proposition 7.7 below, our analyses in later sections will largely focus on direct mechanisms.

Writing  $s_i: \Theta_i \rightarrow \Theta_i$  for a Bayesian strategy in  $\mathcal{M}^d$ , we can describe  $\mathcal{M}^d$  using the following diagram:



Direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$  truthfully Bayesian implements  $g$  if the truth-telling strategy profile  $\{\tau_i\}_{i \in \mathcal{A}}$ , defined by  $\tau_i(\theta_i) = \theta_i$ , is a Bayesian equilibrium of  $\mathcal{M}^d$ . We express this more briefly by saying that  $\mathcal{M}^d$  or  $g$  is Bayesian incentive compatible.

The importance of direct mechanisms is due to the following result.

**Proposition 7.7** (The revelation principle (for Bayesian implementation)).

*If there is a mechanism  $\mathcal{M} = \{\{A_i\}_{i \in \mathcal{A}}, \gamma\}$  that Bayesian implements  $g$ , then the direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$  truthfully Bayesian implements  $g$ . Thus a social choice function is Bayesian implementable by some mechanism if and only if it is Bayesian incentive compatible.*

*Proof.* (Using notation for finite  $\Theta_i$ .) By assumption, there is a Bayesian equilibrium  $s^*$  of  $\mathcal{M}$  such that  $\gamma(s^*(\theta)) = g(\theta)$  for all  $\theta \in \Theta$ .

Suppose agent  $i$  is of type  $\theta_i$ . Under  $\mathcal{M}$ , if her opponents play  $s_{-i}^*$ , it is optimal for agent  $i$  to play  $s_i^*(\theta_i)$ . That is,

$$\sum_{\theta_{-i}} \mu(\theta_{-i}|\theta_i) v_i(\underbrace{\gamma(s_i^*(\theta_i), s_{-i}^*(\theta_{-i}))}_{g(\theta)}, \theta) \geq \sum_{\theta_{-i}} \mu(\theta_{-i}|\theta_i) v_i(\underbrace{\gamma(\hat{a}_i, s_{-i}^*(\theta_{-i}))}_{\text{some } \chi \in \mathcal{X}}, \theta) \text{ for all } \hat{a}_i \in A_i.$$

In particular this is true when  $\hat{a}_i = s_i^*(\hat{\theta}_i)$  for some  $\hat{\theta}_i \in \Theta_i$ , so that  $\gamma(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})) = g(\hat{\theta}_i, \theta_{-i})$ . Substituting into the previous expression thus yields

$$(64) \quad \sum_{\theta_{-i}} \mu(\theta_{-i}|\theta_i) v_i(g(\theta_i, \theta_{-i}), \theta) \geq \sum_{\theta_{-i}} \mu(\theta_{-i}|\theta_i) v_i(g(\hat{\theta}_i, \theta_{-i}), \theta) \text{ for all } \hat{\theta}_i \in \Theta_i.$$



This says that it is optimal for  $i$  to be truthful in  $\mathcal{M}^d$  when others are truthful, or in other words, that  $g$  is Bayesian incentive compatible. ■

The inequalities in (64) are known as *Bayesian incentive compatibility constraints*.

It is easy to explain in words why the revelation principle is true. Imagine that mechanism  $\mathcal{M}$  is used, but that each agent  $i$  must have her action enacted by a proxy. The proxy knows  $s_i^*$  from the start. Just before actions must be chosen, the agent reports a type  $\hat{\theta}_i$  to the proxy, causing it to play action  $s_i^*(\hat{\theta}_i)$ . (The proxy is like a machine—it has no incentives, and its only role is to mechanically carry out strategy  $s_i^*$ .)

Now suppose the other agents report their types truthfully to their proxies. Since  $s^*$  is an equilibrium of  $\mathcal{M}$ , it must be optimal for  $i$  to report her type truthfully; otherwise,  $s_i^*$  was not optimal for  $i$  in the first place. This game of reporting types to proxies is the direct mechanism  $\mathcal{M}^d$ , and we have just shown that truth-telling is an equilibrium of this mechanism.

Remarks:

- (i) In a direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$ , elements of  $\Theta$  have two interpretations: as profiles of types, and as profiles of type announcements. For example, when asking whether allocation  $x(\theta)$  is efficient, we think of  $\theta$  as a profile of types. But when asking whether agent  $i$  has an incentive to report truthfully, we think of  $\theta_i \in \Theta_i$  as  $i$ 's actual type and  $\hat{\theta}_i \in \Theta_i$  as  $i$ 's type announcement. Analyses often hinge on the interplay between these two roles—for instance, see Section 7.4.1. (Beware! Since truth-telling means that  $\hat{\theta}_i = \theta_i$ , both notations can refer either to types or announcements.)
- (ii) The revelation principle shows that for the purpose of determining which social choice functions are implementable, direct mechanisms are enough. This is a tremendous simplification. First of all, it implies that there is no need to construct clever mechanisms. Second, it replaces computations of Bayesian equilibria of groups of games with checking of collections of constraints of form (64). The latter problem often has a simple structure, as we will soon see.
- (iii) There are a variety of reasons for considering more general mechanisms:
  - (a) The revelation principle concerns the existence of an equilibrium that sustains  $g(\cdot)$ . There may be other equilibria that do not. One role of general mechanisms is to do away with such equilibria, so such mechanisms are common in work on full implementation (see Section 7.1.3).
  - (b) In general an agent's type may contain a great deal of information, both payoff-

relevant information (e.g., in combinatorial auction environments) and beliefs, beliefs about beliefs, etc. It may be quite burdensome for the agents to reveal this information to the planner and for the planner to process this information. It therefore may be preferable to use less burdensome mechanisms (e.g., simple auction formats) that achieve similar aims. Related reasons for preferring simple mechanisms are put forward in the *Wilson doctrine*—see Section 7.6.

There is also a revelation principle for dominant strategy implementation, and it follows the same lines as above. Given a direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$  with social choice function  $g$ , we say that  $\mathcal{M}^d$  or  $g$  is *dominant strategy incentive compatible* if truth-telling is a very weakly dominant strategy for all agents.

**Proposition 7.8** (The revelation principle (for dominant strategy implementation)).

*If there is a mechanism  $\mathcal{M} = \{\{A_i\}_{i \in \mathcal{A}}, \gamma\}$  that implements  $g$  in dominant strategies, then the direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$  truthfully implements  $g$  in dominant strategies. Thus a social choice function is dominant strategy implementable by some mechanism if and only if it is dominant strategy incentive compatible.*

Proving Proposition 7.8 is a good exercise.

*Example 7.9. Second price auctions revisited.* Consider a second price auction in an environment with symmetric independent private values drawn from  $\Theta_i = [0, 1]$ . Under this mechanism, each agent places a bid, and (ignoring ties) the good is awarded to the agent who values it most, who then pays the second highest bid; other agents pay nothing. In our present language, a second price auction is a direct mechanism with social choice function  $g(\cdot) = (x^*(\cdot), t(\cdot))$ , where (again ignoring ties)

$$x^*(\hat{\theta}) = i \text{ when } \{i\} = \underset{j \in \mathcal{A}}{\operatorname{argmax}} \hat{\theta}_j, \text{ and } t_i(\hat{\theta}) = \begin{cases} \max_{j \neq i} \hat{\theta}_j & \text{if } x^*(\hat{\theta}) = i, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 4.8 showed that truthful reporting is a weakly dominant strategy. Thus the second price auction is a dominant strategy incentive compatible direct mechanism that implements the ex post efficient allocation function. ♦

Except in Section 7.5 we will focus on private values environments ( $v_i(\chi, \theta) = v_i(\chi, \theta_i)$ ). In these environments, dominant strategy incentive compatibility can be expressed in a simple ex post form. In Section 4.2, we noted that in a Bayesian game with private-value payoff functions  $u_i: A \times \Theta_i \rightarrow \mathbb{R}$ , a Bayesian strategy  $s_i: \Theta_i \rightarrow A_i$  is very weakly dominant

if the following ex post optimality condition holds (cf. (25)):

$$u_i((s_i(\theta_i), a_{-i}), \theta_i) \geq u_i((\hat{a}_i, a_{-i}), \theta_i) \quad \text{for all } a_{-i} \in A_{-i}, \hat{a}_i \in A_i, \text{ and } \theta_i \in \Theta_i$$

Under a direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$  for private values environment  $\mathcal{E}$ , action sets are  $A_i = \Theta_i$ , payoff functions are  $u_i((\hat{\theta}_i, \theta_{-i}), \theta_i) = v_i(g(\hat{\theta}_i, \theta_{-i}), \theta_i)$ , and the truth-telling strategy is  $s_i(\theta_i) = \theta_i$ . We therefore have

**Observation 7.10.** *Let  $g$  be a social choice function for the private values environment  $\mathcal{E}$ . Then  $g$  is dominant strategy incentive compatible if and only if for all agents  $i \in \mathcal{A}$ ,*

$$(65) \quad v_i(g(\theta_i, \hat{\theta}_{-i}), \theta_i) \geq v_i(g(\check{\theta}_i, \hat{\theta}_{-i}), \theta_i) \quad \text{for all } \theta_i, \check{\theta}_i \in \Theta_i \text{ and } \hat{\theta}_{-i} \in \Theta_{-i}.$$

Comparing the previous two equations reveals that in the dominant strategy incentive compatibility constraint (65),  $\hat{\theta}_{-i}$  represents the *announcements* of  $i$ 's opponents, not their actual types. This is because the announcements determine which social alternative is chosen, even though the actual types determine, e.g., whether a social alternative is efficient.

(Here we use the notations  $\check{\theta}_i$  and  $\hat{\theta}_{-i}$  to distinguish between a (possibly) false announcement of type  $\theta_i$  and arbitrary announcements by the agents besides  $i$ . The  $\hat{\theta}_{-i}$  notation is not needed when we consider Bayesian incentive compatibility: there the equilibrium assumption of truth-telling means we only need consider truthful announcements by  $i$ 's opponents, which we can denote by  $\theta_{-i}$ .)

Prior beliefs are irrelevant in condition (65). While in settings with interdependent values dominant strategy incentive compatibility is generally too demanding to be useful, the notion of *ex post incentive compatibility* also makes no use of prior beliefs, but it does assume equilibrium knowledge of opponents' (truth-telling) strategies. See Section 7.5.

The revelation principle for dominant strategy implementation appeared first, and was introduced by a number of authors, including Gibbard (1973) and Green and Laffont (1977). The revelation principle for Bayesian implementation was also developed by many authors, including Dasgupta et al. (1979) and Myerson (1979, 1981, 1982).

## 7.2 Incentive Compatibility, Payoff Equivalence, and Revenue Equivalence

In this section and the next, we consider mechanisms for allocation problems  $\mathcal{E}$  with quasilinear utility and independent private values. The analysis follows Myerson (1981).

$\mathcal{A} = \{1, \dots, n\}$	set of agents
$\Theta_i = [0, 1]$	set of agent $i$ 's types
$F_i, f_i$	cdf and pdf of agent $i$ 's type distribution; $f_i$ positive different agents' types are independent
$\mathcal{X} = \underline{\Delta}\mathcal{A} \times \mathbb{R}^n$	set of social alternatives
$\underline{\Delta}\mathcal{A} = \{q \in \mathbb{R}_+^n: \sum_i q_i \leq 1\}$	assignment probabilities, allowing for non-assignment
$u_i(q, \theta) - t_i = \theta_i q_i - t_i$	agent $i$ 's utility.

(Nothing in this section or the next would change if the type sets  $\Theta_i$  differed across players.)

Since types are drawn independently, we can write  $f(\theta) = \prod_{j \in \mathcal{A}} f_j(\theta_j)$  for the joint distribution of types, and  $f_{-i}(\theta_{-i}) = \prod_{j \neq i} f_j(\theta_j)$  for the joint distribution of the types of  $i$ 's opponents; the latter does not depend on the value of  $\theta_i$ .

A social choice function is a map  $g(\cdot) = (q(\cdot), t(\cdot))$  with  $q: \Theta \rightarrow \underline{\Delta}\mathcal{A}$  and  $t: \Theta \rightarrow \mathbb{R}^n$ . (We again use  $q$  and  $t$  for scalars and  $q(\cdot)$  and  $t(\cdot)$  for functions.)

### ***Bayesian incentive compatibility and payoff equivalence***

When is  $g$  Bayesian incentive compatible? In other words, when is truth-telling a Bayesian equilibrium of the direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$ ?

Suppose that agent  $i$  reports  $\hat{\theta}_i$  and others report truthfully. Then

$$(66) \quad \bar{q}_i(\hat{\theta}_i) = \int_{\Theta_{-i}} q_i(\hat{\theta}_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

is the interim expected probability that  $i$  receives the good, and

$$\bar{t}_i(\hat{\theta}_i) = \int_{\Theta_{-i}} t_i(\hat{\theta}_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

is her expected transfer. If in addition agent  $i$ 's actual type is  $\theta_i$ , then her expected utility is

$$\theta_i \bar{q}_i(\hat{\theta}_i) - \bar{t}_i(\hat{\theta}_i).$$

And if she reports truthfully ( $\hat{\theta}_i = \theta_i$ ), her expected utility is

$$\bar{U}_i(\theta_i) = \theta_i \bar{q}_i(\theta_i) - \bar{t}_i(\theta_i)$$

Notice the similarities between this and the single-agent model from Section 6.2.2. If we

suppress the dependence of  $\bar{q}_i$  on  $\theta_i$ , then the map  $(\theta_i, \bar{q}_i, \bar{t}_i) \mapsto \theta_i \bar{q}_i - \bar{t}_i$  has the single-crossing property in  $\theta_i$  and  $\bar{q}_i$ , in the sense that  $\frac{\partial^2}{\partial \theta_i \partial \bar{q}_i}(\theta_i \bar{q}_i - \bar{t}_i) = 1$ .

By definition, social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$  is Bayesian incentive compatible if under the direct mechanism for  $g(\cdot)$ , every type of every agent finds it optimal to report truthfully:

$$(IC) \quad \theta_i \bar{q}_i(\theta_i) - \bar{t}_i(\theta_i) \geq \theta_i \bar{q}_i(\hat{\theta}_i) - \bar{t}_i(\hat{\theta}_i) \text{ for all } \theta_i, \hat{\theta}_i \in \Theta_i \text{ and } i \in \mathcal{A}.$$

The following result, known as *payoff equivalence*, provides a characterization of Bayesian incentive compatibility.

**Theorem 7.11** (Payoff equivalence).

*Social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$  is Bayesian incentive compatible if and only if*

- (i)  $\bar{q}_i(\cdot)$  is nondecreasing for all  $i \in \mathcal{A}$ , and
- (ii)  $\bar{U}_i(\theta_i) = \bar{U}_i(0) + \int_0^{\theta_i} \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i$  for all  $\theta_i \in \Theta_i$  and  $i \in \mathcal{A}$ .

*Equivalently,  $g(\cdot) = (q(\cdot), t(\cdot))$  is Bayesian incentive compatible if and only if (i) holds and transfers are given by*

$$(ii') \quad \bar{t}_i(\theta_i) = \theta_i \bar{q}_i(\theta_i) - \bar{U}_i(\theta_i) = \theta_i \bar{q}_i(\theta_i) - \int_0^{\theta_i} \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i - \bar{U}_i(0) \text{ for all } \theta_i \in \Theta_i \text{ and } i \in \mathcal{A}.$$

The proposition tells us that Bayesian incentive compatibility and equilibrium expected payoffs only depend on the *interim* allocation probabilities  $\bar{q}_i(\cdot)$  and the lowest type's expected payoff  $\bar{U}_i(0)$ . In particular, the more detailed information about *ex post* allocation probabilities contained in  $q(\cdot)$  has no bearing on these questions.

Theorem 7.11 is a direct analogue of Theorem 6.11 from the single-agent setting. There we had the more general utility function  $u(q, \theta) - t$  with various assumptions on  $u$ , but here we have multiple agents. However, in considering a single agent's decision under a direct mechanism, our focus on truth-telling lets us integrate out the effects of opponents' types, essentially reducing the analysis here to the analysis of the single-agent setting.

One can prove the necessity of (ii) using either of the arguments used to prove Lemma 6.9. Here we prove the necessity of (i) and (ii) using a convexity argument that takes advantage of the linearity of utility in types. Thus the result extends easily to other environments with linear utility (Example 7.4). We take advantage of this fact in Section 7.4.4.

*Proof.* (Necessity) We can express the (IC) constraints for type  $\theta_i$  as

$$(67) \quad \bar{U}_i(\theta_i) = \max_{\hat{\theta}_i \in \Theta_i} (\theta_i \bar{q}_i(\hat{\theta}_i) - \bar{t}_i(\hat{\theta}_i)).$$

Since  $\bar{U}_i(\cdot)$  is the pointwise maximum of the collection of affine functions  $\{v_{\hat{\theta}_i}(\cdot)\}_{\hat{\theta}_i \in \Theta}$ , where  $v_{\hat{\theta}_i}(\theta_i) = \theta_i \bar{q}_i(\hat{\theta}_i) - \bar{t}_i(\hat{\theta}_i)$ ,  $\bar{U}_i(\cdot)$  is itself a convex function. (This can be confirmed by drawing a picture, or see Theorem 5.5 of Rockafellar (1970).)


Next, since

$$\begin{aligned} \bar{q}_i(\hat{\theta}_i)\theta_i - \bar{t}_i(\hat{\theta}_i) &= \bar{q}_i(\hat{\theta}_i)\hat{\theta}_i - \bar{t}_i(\hat{\theta}_i) + \bar{q}_i(\hat{\theta}_i)(\theta_i - \hat{\theta}_i) \\ &= \bar{U}_i(\hat{\theta}_i) + \bar{q}_i(\hat{\theta}_i)(\theta_i - \hat{\theta}_i), \end{aligned}$$

we can rewrite (IC $_{\hat{\theta}_i|\theta_i}$ ) as

$$(68) \quad \bar{U}_i(\theta_i) \geq \bar{U}_i(\hat{\theta}_i) + \bar{q}_i(\hat{\theta}_i)(\theta_i - \hat{\theta}_i), \text{ with equality when } \theta_i = \hat{\theta}_i.$$

Consider each side of inequality (68) as a function of  $\theta_i$ . The left hand side,  $\theta_i \mapsto \bar{U}_i(\theta_i)$ , describes type  $\theta_i$ 's expected payoff under  $(q(\cdot), t(\cdot))$ . The right hand side,  $\theta_i \mapsto \bar{U}_i(\hat{\theta}_i) + \bar{q}_i(\hat{\theta}_i)(\theta_i - \hat{\theta}_i)$ , represents the expected payoff to type  $\theta_i$  from reporting  $\hat{\theta}_i$ , viewed as a function of  $\theta_i$  for some fixed  $\hat{\theta}_i$ ; in other words, it describes each type's benefit from announcing  $\hat{\theta}_i$ . This right hand side function is affine, passes through  $(\hat{\theta}_i, \bar{U}_i(\hat{\theta}_i))$ , and has slope  $\bar{q}_i(\hat{\theta}_i) \in [0, 1]$ . Thus (68) says that this function supports  $\bar{U}_i(\cdot)$  at  $\hat{\theta}_i$ .



If  $\bar{U}_i(\cdot)$  is differentiable, then  $\bar{U}'_i(\hat{\theta}_i)$  equals  $\bar{q}_i(\hat{\theta}_i)$ , the slope of the support function from (68) at  $\hat{\theta}_i$ . Thus the convexity of  $\bar{U}_i(\cdot)$  implies that  $\bar{q}_i(\cdot)$  is nondecreasing, and the integral formula (ii) follows from the fundamental theorem of calculus. In general, these results follow from the convexity of  $\bar{U}_i(\cdot)$  and results from convex analysis (see Rockafellar (1970, Theorem 24.2)).

(Sufficiency) Suppose that  $\theta_1 > \theta_0$ . Then since  $\bar{q}_i(\cdot)$  is nondecreasing,

$$\bar{U}_i(\theta_1) = \bar{U}_i(\theta_0) + \int_{\theta_0}^{\theta_1} \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i \geq \bar{U}_i(\theta_0) + \int_{\theta_0}^{\theta_1} \bar{q}_i(\theta_0) d\hat{\theta}_i = \bar{U}_i(\theta_0) + \bar{q}_i(\theta_0)(\theta_1 - \theta_0),$$

which is  $(IC_{\theta_0|\theta_1})$  as expressed in (68). A careful reading of the argument shows that the same inequality holds if  $\theta_1 < \theta_0$ . ■

### Revenue equivalence

Returning to allocation problems, Theorem 7.11 and the revelation principle easily yield the following result:

**Theorem 7.12** (Revenue equivalence). *Fix an allocation problem  $\mathcal{E}$ , and suppose that mechanisms  $\mathcal{M}$  and  $\mathcal{M}^\dagger$  Bayesian implement social choice functions  $g(\cdot) = (q(\cdot), t(\cdot))$  and  $g^\dagger(\cdot) = (q^\dagger(\cdot), t^\dagger(\cdot))$ , where the latter have*

- (I) *the same interim allocation probabilities ( $\bar{q}_i(\cdot) = \bar{q}_i^\dagger(\cdot)$  for all  $i \in \mathcal{A}$ ), and*
- (II) *the same expected utilities of the lowest types of each agent ( $\bar{U}_i(0) = \bar{U}_i^\dagger(0)$  for all  $i \in \mathcal{A}$ ).*

*Then these mechanisms generate the same expected payoffs for each type of each agent and the same expected revenue.*

*Proof.* Fix an allocation problem  $\mathcal{E}$  and a mechanism  $\mathcal{M}$  that Bayesian implements social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$ . By the revelation principle,  $g(\cdot)$  is Bayesian incentive compatible. Thus by payoff equivalence, agent  $i$ 's expected payoff function  $\bar{U}_i(\cdot)$  and expected transfer function  $\bar{t}_i(\cdot)$  under (the relevant equilibrium of)  $\mathcal{M}$  are given by formulas (ii) and (ii'), whose right-hand sides only depend on  $g(\cdot)$  by way of  $\bar{q}_i(\cdot)$  and  $\bar{U}_i(0)$ . The expected transfer functions in turn determine the designer's expected revenue, which is  $\sum_{i \in \mathcal{A}} \int_0^1 \bar{t}_i(\theta_i) f_i(\theta_i) d\theta_i$ . Thus interim expected payoffs and expected revenue under  $\mathcal{M}$  only depend on  $g(\cdot)$  by way of  $\bar{q}_i(\cdot)$  and  $\bar{U}_i(0)$ . ■

*Example 7.13. Revenue equivalence in auctions.* Theorem 7.12 implies that any auction format which Bayesian implements a social choice function that

- (I') awards the good to the bidder with the highest valuation, and
- (II) gives each bidder's lowest type an expected payoff of zero

generates the same expected revenue.

(Notice that (I') is a more demanding sort of requirement than (I) from Theorem 7.12: (I') is a condition on *ex post* allocation probabilities, whereas (I) is a condition on *interim* allocation probabilities.)

When the environment is symmetric (i.e., all agents values are drawn from the same distribution), there are many auction formats with equilibria satisfying (I') and (II).

In asymmetric environments, conditions (I') and (II) do not hold for all standard auction formats. For instance, in a second price auction, bidding truthfully is still weakly

dominant, and so leads to an efficient allocation, but in a first price auction the equilibrium allocation need not be efficient. In fact, either format may generate higher expected revenue than the other depending on the type distributions; see Section 4.3.2 of Krishna (2002). ♦

### *Payoff equivalence for dominant strategy incentive compatibility*

There is also a payoff equivalence theorem for dominant strategy incentive compatibility (Laffont and Maskin (1979)). Dominant strategy incentive compatibility requires the reports of each type  $\theta_i$  of each agent  $i$  to be optimal regardless of his opponents' Bayesian strategies (i.e., maps from types to announcements). By Observation 7.10, this requirement amounts to the ex post incentive compatibility constraints

$$(69) \quad \theta_i q_i(\theta_i, \theta_{-i}) - t_i(\theta_i, \theta_{-i}) \geq \theta_i q_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) \quad \text{for all } \theta_i, \hat{\theta}_i \in \Theta_i, \theta_{-i} \in \Theta_{-i}, \text{ and } i \in \mathcal{A}.$$

For each fixed  $\theta_{-i}$ , the constraints (69) are an instance of the incentive compatibility constraints from the principal-agent problem from Section 6.2.2. Thus Theorem 7.14 below, which characterizes dominant strategy incentive compatibility, follows immediately from Theorem 6.11, though it too can be established using the convexity argument used to prove Theorem 7.11.

To state this result, let

$$U_i(\theta) = \theta_i q_i(\theta) - t_i(\theta)$$

be agent  $i$ 's ex post utility under social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$ .

#### **Theorem 7.14** (Payoff equivalence).

*Social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$  is dominant strategy incentive compatible if and only if*

- (i)  $q_i(\cdot, \theta_{-i})$  is nondecreasing for all  $\theta_{-i} \in \Theta_{-i}$  and  $i \in \mathcal{A}$ , and
- (ii)  $U_i(\theta_i, \theta_{-i}) = U_i(0, \theta_{-i}) + \int_0^{\theta_i} q_i(\hat{\theta}_i, \theta_{-i}) d\hat{\theta}_i$  for all  $\theta_{-i} \in \Theta_{-i}$ ,  $\theta_i \in \Theta_i$ , and  $i \in \mathcal{A}$ .

*Equivalently,  $g(\cdot)$  is dominant strategy incentive compatible if and only if (i) holds and*

$$(70) \quad t_i(\theta) = \theta_i q_i(\theta) - U_i(\theta) = \theta_i q_i(\theta) - \int_0^{\theta_i} q_i(\hat{\theta}_i, \theta_{-i}) d\hat{\theta}_i - U_i(0, \theta_{-i}) \quad \text{for all } \theta \in \Theta \text{ and } i \in \mathcal{A}.$$



### 7.3 Revenue Maximization and Optimal Auctions

Following Myerson (1981), we now consider expected revenue maximization in allocation problems. Here the mechanism design problem is naturally viewed as a principal / many agents problem.

We consider Bayesian implementation (but see the remarks after Theorem 7.11). By the revelation principle, we can focus on Bayesian incentive compatible direct mechanisms. Furthermore, payoff equivalence gives us necessary and sufficient conditions for a social choice function  $g(\cdot) = (q(\cdot), t(\cdot))$  to be Bayesian incentive compatible.

For revenue maximization to be a sensible question, we must require (*interim*) *individual rationality*:

$$(71) \quad \bar{U}_i(\theta_i) \geq 0 \text{ for all } \theta_i \in \Theta_i \text{ and } i \in \mathcal{A}$$

By payoff equivalence, (71) reduces to the requirement that  $\bar{U}_i(0) \geq 0$  for all  $i \in \mathcal{A}$ . (Were individual rationality not imposed, one could obtain arbitrarily high revenues by taking each  $\bar{U}_i(0)$  to  $-\infty$ .)

Now, recalling that  $\Theta_i = [0, 1]$ , we can write the principal's problem as follows

$$(72) \quad \max_{q(\cdot), \bar{U}(0)} \sum_{i \in \mathcal{A}} \int_0^1 \left( \theta_i \bar{q}_i(\theta_i) - \int_0^{\theta_i} \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i - \bar{U}_i(0) \right) f_i(\theta_i) d\theta_i$$

subject to  $q(\theta) \in \underline{\Delta} \mathcal{A}$

$\bar{q}_i(\cdot)$  nondecreasing for all  $i \in \mathcal{A}$

$\bar{U}_i(0) \geq 0$  for all  $i \in \mathcal{A}$

It is instructive to compare problem (72) to the principal's problem (58) in the setting of adverse selection with a continuum of types (Section 6.2.2). If in the latter we assume that utilities are of the linear form  $u(q, \theta) = \theta q$ , we obtain

$$(58) \quad \max_{q(\cdot), U(0)} \int_0^1 \left( \theta q(\theta) - \int_0^\theta q(\hat{\theta}) d\hat{\theta} - U(0) - c(q(\theta)) \right) f(\theta) d\theta$$

subject to  $q(\cdot)$  nondecreasing

$U(0) \geq 0$

Thus in this linear case, (58) closely resembles (72). In the former problem, the quality  $q(\theta)$  is a scalar, and the objective function includes a production cost term that is absent

from (72). For its part, (72) is a multiagent problem, with the allocation probabilities  $q(\theta) \in \mathbb{R}^n$  required to satisfy the constraint  $q(\theta) \in \underline{\Delta}\mathcal{A}$ , and with the objective function and monotonicity constraints stated in terms of the interim objects  $\bar{q}_i(\cdot)$  and  $\bar{U}_i(0) \geq 0$ . The analysis to follow will address these novelties.

We assume that the *virtual valuation* of type  $\theta_i$ , defined by

$$(73) \quad \psi_i(\theta_i) = \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)},$$

is increasing in  $\theta_i$ .

**Theorem 7.15** (Revenue maximization).

*Suppose that virtual valuations  $\psi_i(\theta_i)$  are increasing in  $\theta_i$  for all  $i \in \mathcal{A}$ . Then any Bayesian incentive compatible, individually rational mechanism that*

- (i) *only assigns the good if some agent has a nonnegative virtual valuation, and*
- (ii) *if so, assigns the good to an agent whose virtual valuation (and hence type) is highest*  
*maximizes the principal's expected revenue among all such mechanisms.*

*Proof.* To start, reversing the order of integration in the double integral in (72) yields

$$(74) \quad \int_0^1 \left( \int_0^{\theta_i} \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i \right) f_i(\theta_i) d\theta_i = \int_0^1 \left( \int_{\hat{\theta}_i}^1 f_i(\theta_i) d\theta_i \right) \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i = \int_0^1 (1 - F_i(\hat{\theta}_i)) \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i.$$

(Compare display (59).)

Choosing  $\bar{U}_i(0) = 0$  is clearly optimal in the principal's problem. Thus substituting (74) into the objective function (72) (after renaming the variable of integration) yields

$$(75) \quad \begin{aligned} & \sum_{i \in \mathcal{A}} \int_0^1 (\theta_i f_i(\theta_i) - (1 - F_i(\theta_i))) \bar{q}_i(\theta_i) d\theta_i \\ &= \sum_{i \in \mathcal{A}} \int_0^1 \left( \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) \bar{q}_i(\theta_i) f_i(\theta_i) d\theta_i \\ &= \sum_{i \in \mathcal{A}} \int_{\Theta_i} \psi_i(\theta_i) \bar{q}_i(\theta_i) f_i(\theta_i) d\theta_i \end{aligned}$$

Then substituting in definition (66) of  $\bar{q}_i(\theta_i)$  and rearranging the result lets us express the

objective function as

$$(76) \quad \sum_{i \in \mathcal{A}} \int_{\Theta_i} \psi_i(\theta_i) \left( \int_{\Theta_{-i}} q_i(\theta_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i} \right) f_i(\theta_i) d\theta_i = \int_{\Theta} \left( \sum_{i \in \mathcal{A}} \psi_i(\theta_i) q_i(\theta) \right) f(\theta) d\theta.$$

If we ignore the requirement that  $\bar{q}_i(\cdot)$  be nondecreasing, then under the constraint  $q(\theta) \in \underline{\Delta}\mathcal{A}$ , maximization requires that

$$(77) \quad q_i(\theta) > 0 \Rightarrow \psi_i(\theta_i) = \max\{\psi_1(\theta_1), \dots, \psi_n(\theta_n), 0\}.$$

To see that the monotonicity requirement is moot, suppose that  $\hat{\theta}_i < \theta_i$ . Then  $\psi_i(\hat{\theta}_i) < \psi_i(\theta_i)$ , so (77) implies that  $q_i(\hat{\theta}_i, \theta_{-i}) \leq q_i(\theta_i, \theta_{-i})$  for all  $\theta_{-i} \in \Theta_{-i}$ . Thus integrating yields  $\bar{q}_i(\hat{\theta}_i) \leq \bar{q}_i(\theta_i)$  as required. ■

This proof provides an interpretation of the virtual valuation  $\psi_i(\theta_i) = \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}$ . The first term is an agent's actual valuation. The second term is the information rent *attributed to* an agent  $i$  of type  $\theta_i$ , but this rent is earned by *higher types* of agent  $i$ . Remark 7.16 explains why.

Roughly speaking, the virtual valuation  $\psi_i(\theta_i)$  represents the surplus that the principal can extract from agent  $i$  of type  $\theta_i$  if information rents to higher types of agent  $i$  are properly accounted for. This explains why the principal never allocates the good to an agent  $i$  of a type  $\theta_i$  with a negative virtual valuation. Ex post the principal may prefer to do so, but ex ante it is a bad idea: if the principal employs a mechanism that does so, his payoff when type  $\theta_i$  (or a nearby type) is realized will be outweighed by the expected rent that the mechanism must pay to higher types of agent  $i$ .

*Remark 7.16.*

- (i) The interpretation of virtual valuations, which first appear in (75), is essentially the same as to the interpretation of virtual utilities in the monopolistic screening problem (see Remark 6.13). Indeed, since  $u_i(q, \theta) = q\theta$ , virtual valuations and virtual utility are equal here.
- (ii) To proceed from (75) to the solution of the principal's problem, we rearrange (75) to obtain the right-hand side of (76), an integral over type profiles. This is useful because it allows us to try maximizing the expression in parentheses at each type profile  $\theta$ , subject to the constraint  $q(\theta) \in \underline{\Delta}\mathcal{A}$ , after which we can check the each agent's monotonicity condition. The independence of types, and implicitly the assumption of private values, makes this rearrangement possible. These assumptions likewise ensure that the quantity  $\psi_i(\theta_i)$  used to determine whether agent  $i$  gets the object at type profile  $\theta = (\theta_i, \theta_{-i})$  only reflects information rents paid to higher types of agent  $i$ , and not information rents paid to other agents.

(iii) To interpret virtual valuations, Bulow and Roberts (1989) imagine a monopolist who faces inverse demand curve  $D(p) = 1 - F_i(p)$ . (This means that the 90th percentile of willingness to pay,  $p_{.90}$ , is implicitly defined by  $1 - F_i(p_{.90}) = .10$ , and so by  $F_i(p_{.90}) = .90$ .) If this monopolist has no production costs, it should choose a price  $p$  to maximize revenue  $R(p) = p(1 - F_i(p))$ . The marginal change in revenue from reducing the price is  $-R'(p) = pf_i(p) - (1 - F_i(p)) = f_i(p)\psi_i(p)$ . This captures the tradeoff between the two effects of a price reduction: selling to more consumers on the margin versus reducing the price paid by existing consumers.

We now introduce transfer function  $t(\cdot)$  to complete the definition of the mechanism. In view of payoff equivalence,  $t(\cdot)$  must be chosen so that expected transfers satisfy

$$\bar{t}_i(\theta_i) = \theta_i \bar{q}_i(\theta_i) - \int_0^{\theta_i} \bar{q}_i(\hat{\theta}_i) d\hat{\theta}_i$$

We can make this property of interim transfers hold by making it hold ex post—that is, by defining

$$(78) \quad t_i(\theta_i, \theta_{-i}) = \theta_i q_i(\theta_i, \theta_{-i}) - \int_0^{\theta_i} q_i(\hat{\theta}_i, \theta_{-i}) d\hat{\theta}_i.$$

If  $q_i(\theta_i, \theta_{-i}) = 0$ , then  $q_i(\hat{\theta}_i, \theta_{-i}) = 0$  for  $\hat{\theta}_i < \theta_i$  (by Theorem 7.15), so (78) implies that  $t_i(\theta_i, \theta_{-i}) = 0$ . In other words, an agent who is not assigned the object pays nothing.

To evaluate other cases, let

$$(79) \quad \tau_i(\theta_{-i}) = \inf \left\{ \theta_i : \psi_i(\theta_i) = \max \{ \psi_1(\theta_1), \dots, \psi_n(\theta_n), 0 \} \right\}.$$

Then (again ignoring ties)

$$(80) \quad q_i(\theta_i, \theta_{-i}) = \begin{cases} 1 & \text{if } \theta_i > \tau_i(\theta_{-i}), \\ 0 & \text{if } \theta_i < \tau_i(\theta_{-i}). \end{cases}$$

So if  $\theta_i > \tau_i(\theta_{-i})$ , then (78) implies that

$$t_i(\theta_i, \theta_{-i}) = \theta_i - (\theta_i - \tau_i(\theta_{-i})) = \tau_i(\theta_{-i}).$$

In words: if agent  $i$  is assigned the good, her payment is the lowest valuation she could have had and still have been assigned the good.

Summing up, we have

**Proposition 7.17.** *Suppose that virtual valuations are increasing. Then one mechanism that maximizes revenue among Bayesian incentive compatible, individually rational mechanisms is the direct mechanism*

$$(q_i(\theta), t_i(\theta)) = \begin{cases} (1, \tau_i(\theta_{-i})) & \text{if } \theta_i > \tau_i(\theta_{-i}), \\ (0, 0) & \text{if } \theta_i < \tau_i(\theta_{-i}), \end{cases}$$

where  $\tau_i(\theta_i)$  is defined by (79). In fact, this mechanism is dominant strategy incentive compatible.

The final claim is easy to check, but still quite surprising. We were looking to maximize revenue while satisfying Bayesian incentive compatibility, but found we were able to do so while satisfying the much more demanding requirement of dominant strategy incentive compatibility, which does not require agents to correctly anticipate opponents' behavior. (We could have seen this coming: The statement of Theorem 7.15 noted (and it is easy to see directly) that the allocation function (80) is “ex post monotone”, in that for each agent  $i$  and each profile  $\theta_{-i}$  of opponents types,  $q_i(\cdot, \theta_{-i})$  is nondecreasing. The transfers (78) are the transfers (70) with  $U_i(0, \theta_{-i}) \equiv 0$ . Thus Theorem 7.14 implies that  $(q_i(\cdot), t_i(\cdot))$  is dominant strategy incentive compatible. We return to this point shortly.)

Theorem 7.15 shows that under revenue maximizing mechanisms, the seller may wind up keeping the good. This is clearly not ex post efficient. In addition, the foregoing arguments show that even when the good is sold, it may not be sold to the agent who values it most, since the agent with the highest virtual valuation  $\psi_i(\theta_i)$  need not have the highest valuation  $\theta_i$ . However, if the environment is symmetric, with all agents' types being drawn from the same distribution, the latter inefficiency is absent, as we now discuss.

### *Symmetric environments*

Now suppose further that the environment is symmetric, so that  $\psi_j(\cdot) = \psi(\cdot)$  for all  $j$ . Then (79) becomes

$$\begin{aligned} \tau_i(\theta_{-i}) &= \inf \left\{ \theta_i : \psi(\theta_i) = \max\{\psi(\theta_1), \dots, \psi(\theta_n), 0\} \right\} \\ &= \inf \left\{ \theta_i : \theta_i = \max\{\theta_1, \dots, \theta_n, \psi^{-1}(0)\} \right\}. \end{aligned}$$

Given the formulas for  $q_i(\cdot)$  and  $t_i(\cdot)$  above, we have

**Proposition 7.18.** *Consider a symmetric independent private values environment with increasing virtual valuations. Then a second price auction with reserve price  $\psi^{-1}(0)$  is revenue maximizing*

among Bayesian incentive compatible, individually rational mechanisms. Moreover, this mechanism is dominant strategy incentive compatible.

Thus in symmetric environments with increasing virtual valuations, revenue maximization is achieved using a standard auction format, and in (very weakly) dominant strategies. Moreover, to employ this mechanism, the only thing the designer needs to know about the distribution of types is which type has virtual valuation zero. This is a striking result.

### *More on Bayesian and dominant strategy incentive compatibility*

In the allocation problem studied above, we were able to find dominant strategy incentive compatible mechanisms that were optimal among all Bayesian incentive compatible mechanisms. Returning to the focus of Section 7.2, suppose that the social choice function  $(q(\cdot), t(\cdot))$  is Bayesian incentive compatible. Under what conditions on the allocation function  $q(\cdot)$  can we find an alternative transfer function  $t^\dagger(\cdot)$  such that  $(q(\cdot), t^\dagger(\cdot))$  is dominant strategy incentive compatible? Theorem 7.14 implies that this is possible if and only if  $q_i(\cdot, \theta_{-i})$  is nondecreasing for all  $\theta_{-i} \in \Theta_{-i}$  and  $i \in \mathcal{A}$ . For results on this question in more general environments, see Mookherjee and Reichelstein (1992).

One can also ask the following less demanding question: if a direct mechanism  $g(\cdot) = (q(\cdot), t(\cdot))$  is Bayesian incentive compatible, when is it possible to find another direct mechanism  $g^\dagger(\cdot) = (q^\dagger(\cdot), t^\dagger(\cdot))$  with the same *interim* allocation probabilities ( $\bar{q}_i^\dagger(\cdot) = \bar{q}_i(\cdot)$ ) and *expected* transfers ( $\bar{t}_i^\dagger(\cdot) = \bar{t}_i(\cdot)$ ), but that is dominant strategy incentive compatible? Remarkably, Manelli and Vincent (2010) show that this is *always* possible! Of course, the ex post allocation probabilities  $q^\dagger(\cdot)$  under the new mechanism generally differ from the allocation probabilities  $q(\cdot)$  under the original mechanism; for instance, even if the original allocation function  $q(\cdot)$  was ex post efficient, the new allocation function  $q^\dagger(\cdot)$  generally will not be.

Gershkov et al. (2013) extend Manelli and Vincent's (2010) result to all environments with linear utility (Example 7.4), and provide a more direct proof. In the context of allocation problems, the analysis boils down to the following non-obvious fact: given any "joint distribution"  $q(\cdot)$  whose "marginal distributions"  $\bar{q}_i(\cdot)$  are nondecreasing, one can construct a new "joint distribution"  $q^\dagger(\cdot)$  with the same "marginal distributions"  $\bar{q}_i^\dagger(\cdot) = \bar{q}_i(\cdot)$ , but whose "conditional distributions"  $q_i^\dagger(\cdot, \theta_{-i})$  are all nondecreasing.

## 7.4 Allocative Efficiency

We now turn to the question of implementing efficient allocations. Here it is most natural to think of the mechanism being chosen by a social planner, or designed by the agents themselves (rather than as a principal / many agents problem as in Section 7.3). Applications include public good provision and allocation problems, in particular bilateral trading problems.

(We have seen that revenue maximizing allocation mechanisms are generally inefficient: the principal may wind up keeping the good, and in asymmetric environments, a good may be sold to an agent whose valuation is not highest.)

We consider mechanisms design problems  $\mathcal{E}$  with quasilinear utility and independent private values.

$\mathcal{A} = \{1, \dots, n\}$	set of agents
$\Theta_i$	set of agent $i$ 's types
$F_i$	cdf of agent $i$ 's type distribution
	agents' types are drawn independently
$\mathcal{X} = X \times \mathbb{R}^n$ or $\Delta X \times \mathbb{R}^n$	set of social alternatives
	$X$ is finite
$u_i(x, \theta_i) - t_i$	agent $i$ 's utility

We would like to implement social choice functions  $g(\cdot) = (x^*(\cdot), t(\cdot))$  whose *allocation functions*  $x^*(\cdot)$  are *ex post efficient*, as defined at the start of the section:

$$(81) \quad x^*(\theta) \in \operatorname{argmax}_{x \in X} \sum_{i \in \mathcal{A}} u_i(x, \theta_i) \quad \text{for all } \theta \in \Theta.$$

If all actors in the economy are agents in the mechanism, then true efficiency also requires that transfers not be burned—see Section 7.4.2.

#### 7.4.1 Groves mechanisms and the VCG mechanism

References: Vickrey (1961), Clarke (1971), Groves (1973).

Idea: Use transfers to make each agent account for other agents' payoffs. If the allocation function is ex post efficient, this will provide agents with incentives for truthful reporting.

Let  $x^*(\cdot)$  be an ex post efficient allocation function. Let

$$\beta_i^*(\theta) = \sum_{j \neq i} u_j(x^*(\theta), \theta_j)$$

be the total utility of agents besides  $i$  at allocation  $x^*(\theta)$  when the type profile is  $\theta$ .

The direct mechanism for  $g(\cdot) = (x^*(\cdot), t^G(\cdot))$ , where the transfer function  $t^G(\cdot)$  is of the form

$$t_i^G(\theta) = -\beta_i^*(\theta) + h_i(\theta_{-i})$$

for some  $h_i: \Theta_{-i} \rightarrow \mathbb{R}$  is called a *Groves mechanism*. For future reference, we call the Groves

mechanism with  $h_i(\theta_{-i}) \equiv 0$  the *plain Groves mechanism*.

**Proposition 7.19.** *Groves mechanisms are dominant strategy incentive compatible.*

*Proof.* If agent  $i$  is of type  $\theta_i$ , his payoff from reporting  $\check{\theta}_i$  when others report  $\hat{\theta}_{-i}$  is

$$(82) \quad \begin{aligned} & u_i(x^*(\check{\theta}_i, \hat{\theta}_{-i}), \theta_i) - (-\beta_i^*(\check{\theta}_i, \hat{\theta}_{-i}) + h_i(\hat{\theta}_{-i})) \\ & = u_i(x^*(\check{\theta}_i, \hat{\theta}_{-i}), \theta_i) + \sum_{j \neq i} u_j(x^*(\check{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j) - h_i(\hat{\theta}_{-i}). \end{aligned}$$

Since  $x^*(\cdot)$  is ex post efficient, it follows that for any  $\hat{\theta}_{-i}$ , this function is maximized by choosing  $\check{\theta}_i = \theta_i$ . Thus truth-telling is a very weakly dominant strategy (cf. Observation 7.10). ■

While this proof is very short, it is not transparent, so it is worth going through slowly. Consider the sum of the first two terms in (82). Since  $x^*(\cdot)$  is ex post efficient, allocation  $x^*(\theta_i, \hat{\theta}_{-i})$  maximizes this sum when the second arguments of  $u_1, \dots, u_n$  are given by the components of  $(\theta_i, \hat{\theta}_{-i})$ . By definition, this is true whether any given component of  $(\theta_i, \hat{\theta}_{-i})$  is an agent's actual type or merely a type announcement.

So suppose that agent  $i$  is of type  $\theta_i$  and that the others report  $\hat{\theta}_{-i}$ . If agent  $i$  announces  $\theta_i$ , then the allocation is  $x^*(\theta_i, \hat{\theta}_{-i})$ , which maximizes

$$(83) \quad u_i(\cdot, \theta_i) + \sum_{j \neq i} u_j(\cdot, \hat{\theta}_j).$$

This is (82) (ignoring the  $h_i(\hat{\theta}_{-i})$  term). If the agent announces  $\check{\theta}_i \neq \theta_i$ , then the allocation is  $x^*(\check{\theta}_i, \hat{\theta}_{-i})$ , which maximizes

$$u_i(\cdot, \check{\theta}_i) + \sum_{j \neq i} u_j(\cdot, \hat{\theta}_j),$$

but which need not maximize (83). So for any type reports  $\hat{\theta}_{-i}$  of the others, it is a best response for an agent  $i$  of type  $\theta_i$  to report  $\theta_i$ . In other words, truth-telling is a very weakly dominant strategy.

Remark: An important property of Groves mechanisms (which follows immediately from their definition) is that an agent's report only affects his transfer indirectly, by way of his report's effect on the allocation. That is, for any fixed reports  $\hat{\theta}_{-i}$  of  $i$ 's opponents, any



report by agent  $i$  that leads to a given allocation  $x$  will also require  $i$  to pay the same transfers. We will observe this property in the examples below.

Let  $x^{-i}: \Theta_{-i} \rightarrow X$  be an “allocation function” that is ex post efficient if agent  $i$ 's payoffs are ignored:

$$x^{-i}(\theta_{-i}) \in \operatorname{argmax}_{x \in X} \sum_{j \neq i} u_j(x, \theta_j).$$

Then the Groves mechanism with

$$(84) \quad t_i^V(\theta) = \sum_{j \neq i} u_j(x^{-i}(\theta_{-i}), \theta_j) - \sum_{j \neq i} u_j(x^*(\theta), \theta_j)$$

is called the *Vickrey-Clarke-Groves* (or *VCG*) *mechanism*. The mechanism is defined so as to make each agent pay for the externalities his announcement creates. In particular:

**Observation 7.20.** *Transfers under the VCG mechanism are nonnegative, and agent  $i$ 's transfer is positive only if his announcement affects the allocation.*

*Example 7.21. Allocation of an indivisible private good.* As in Example 7.2, let  $X = \mathcal{A}$ ,  $u_i(i, \theta_i) = \theta_i$  and  $u_i(j, \theta_i) = 0$  for  $j \neq i$ . Then ignoring ties,

$$x^*(\theta) = i \text{ when } \{i\} = \operatorname{argmax}_{j \in \mathcal{A}} \theta_j \text{ and } x^{-i}(\theta_{-i}) = k \text{ when } \{k\} = \operatorname{argmax}_{j \neq i} \theta_j, \text{ so}$$

$$t_i^V(\theta) = \begin{cases} \max_{j \neq i} \theta_j & \text{if } x^*(\theta) = i, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, the VCG mechanism is equivalent to a second price auction! ♦

*Example 7.22. Public good provision.* A community of  $n$  agents is deciding whether to build a monument. If it is built, the construction cost  $c$  will be paid for by levying a tax of  $\frac{c}{n}$  on each agent. (The agents are compelled to pay this tax—there is no opting out (cf. Section 7.4.3).) The value that agent  $i$  assigns to building the monument,  $\theta_i$ , is his private information. What are the monetary transfers under the VCG mechanism?

Let  $X = \{0, 1\}$ , where 1 represents the decision to build the monument and levy the tax. Then  $u_i(1, \theta_i) = \theta_i - \frac{c}{n}$  and  $u_i(0, \theta_i) = 0$ . Thus up to the choice of how to break ties, we have

$$x^*(\theta) = \begin{cases} 1 & \text{if } \sum_{i \in \mathcal{A}} \theta_i - c \geq 0, \\ 0 & \text{otherwise,} \end{cases} \text{ and } x^{-i}(\theta_{-i}) = \begin{cases} 1 & \text{if } \sum_{j \neq i} \theta_j - \frac{n-1}{n}c \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Thus

$$t_i^V(\theta) = \begin{cases} \frac{n-1}{n}c - \sum_{j \neq i} \theta_j & \text{if } x^*(\theta) = 1 \text{ and } x^{-i}(\theta_{-i}) = 0, \\ \sum_{j \neq i} \theta_j - \frac{n-1}{n}c & \text{if } x^*(\theta) = 0 \text{ and } x^{-i}(\theta_{-i}) = 1, \\ 0 & \text{if } x^*(\theta) = x^{-i}(\theta_{-i}). \end{cases}$$

(The transfer payment  $t_i^V(\cdot)$  is in addition to the tax  $\frac{c}{n}$  that the agent must pay if the monument is built. The latter is already included in the the definition of allocation  $1 \in X$ ; in particular, it appears in the agents' utilities from that allocation.)

If agent  $i$ 's report causes the monument to be built, his VCG transfer  $t_i^V(\theta)$  is the difference between the others' contributions to the construction cost and the sum of the others' reported valuations. Since the monument would not have been built if agent  $i$  were absent, this transfer is nonnegative. If agent  $i$ 's report prevents the monument from being built, his VCG transfer is the difference between the sum of the others' reported valuations and their contributions to the construction cost; again, this transfer is nonnegative. If agent  $i$ 's report does not affect the public decision, his VCG transfer is zero.

For some type profiles  $\theta$ , the monument will be built ( $x^*(\theta) = 1$ ) even though some agent  $i$ 's valuation is less than his contribution to the construction cost ( $\theta_i - \frac{c}{n} < 0$ ). Under the VCG mechanism, agent  $i$  still prefers to tell the truth. If he does so, then his transfer is zero, so his payoff is  $\theta_i - \frac{c}{n} < 0$ . If instead he reports a low enough type that the monument is not built, then his VCG transfer exceeds  $|\theta_i - \frac{c}{n}|$ , so he is even worse off. ♦

Example 7.22 can be interpreted as a *collective choice problem*, with the agents themselves devising the mechanism in order to ensure an ex post efficient outcome. In such contexts, some mechanisms that satisfy incentive compatibility constraints may not be useful in practice. If agents have property rights, they may be unwilling to participate in a mechanism in which they always lose (as above). In most cases the mechanism cannot run an ex post deficit, and if agents are unwilling to "burn money" it cannot run an ex post surplus either. The remainder of this section considers what can be achieved when these additional constraints must be met.

#### 7.4.2 Budget balance and the AGV mechanism

Also known as the *expected externality mechanism*. References: d'Aspremont and Gérard-Varet (1979), Arrow (1979).

The VCG mechanism implements ex post efficient allocation functions in weakly dominant

strategies. Transfers under this mechanism are nonnegative.

Suppose instead that we would like to implement an ex post efficient social choice function with a transfer function that satisfies (*ex post*) *budget balance*:

$$\sum_{i \in \mathcal{A}} t_i(\theta) = 0 \text{ for all } \theta \in \Theta.$$

This means that money is not burned. Budget balance is necessary for full ex post efficiency when everyone in the economy participates in the mechanism.

Budget balance is easily achieved if there is one agent whose preferences are known, as we can have this agent receive the others' payments. What can be done if this is not the case? A natural idea is to have each agent make a payment equal to his VCG transfer, and then to redistribute this payment among the other agents. But this raises the possibility of an agent misrepresenting his type in order to increase the redistributions he receives.

We now show that if types are independent, both ex post efficiency and budget balance can be achieved. However, in addition to assuming independence, we must weaken the notion of implementation from dominant strategy to Bayesian implementation.

Let allocation function  $x^*(\cdot)$  be ex post efficient. Using the notation for continuous type spaces, and using the assumption that types are independent, let

$$(85) \quad \bar{t}_i^V(\theta_i) = \int_{\Theta_{-i}} t_i^V(\theta_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

denote type  $\theta_i$ 's expected VCG transfer when his opponents are truthful. (This is also called type  $\theta_i$ 's "expected externality".)

The direct mechanism for  $g(\cdot) = (x^*(\cdot), t^A(\cdot))$ , where

$$(86) \quad t_i^A(\theta) = \bar{t}_i^V(\theta_i) - \frac{1}{n-1} \sum_{j \neq i} \bar{t}_j^V(\theta_j)$$

is called the *AGV mechanism*. In words, an agent of type  $\theta_i$  makes a *payment* equal to  $\bar{t}_i^V(\theta_i)$ , his expected transfer under the VCG mechanism, and he receives a share of each other agent's payment. Dividing each agent's payment among the others ensures budget balance. The fact that each agent's payment only depends on his type, and hence that each opponent's payment only depends on that opponent's type, ensures that misrepresenting his type cannot change the payments that an agent receives from his opponents. (It is not so important that the sharing is equal under (86); what matters is that each agent's base

payment is given to the other agents in a manner that the others cannot influence.)

*Example 7.23. Allocation of an indivisible private good.* Once again let  $X = \mathcal{A}$ ,  $u_i(i, \theta_i) = \theta_i$  and  $u_i(j, \theta_i) = 0$  for  $j \neq i$ . We saw in Example 7.21 that the VCG mechanism is equivalent to a second price auction:

$$x^*(\theta) = i \text{ when } \theta_i > \max_{j \neq i} \theta_j;$$

$$t_i^V(\theta) = \begin{cases} \max_{j \neq i} \theta_j & \text{if } \theta_i > \max_{j \neq i} \theta_j, \\ 0 & \text{otherwise.} \end{cases}$$

Thus under the AGV mechanism, the payment made by an agent  $i$  of type  $\theta_i$  is

$$\begin{aligned} \bar{t}_i^V(\theta_i) &= \int_{\Theta_{-i}} t_i^V(\theta_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i} \\ &= \int_{\{\theta_{-i}: \theta_i > \max_{j \neq i} \theta_j\}} \left( \max_{j \neq i} \theta_j \right) f_{-i}(\theta_{-i}) d\theta_{-i} \\ &= \int_0^{\theta_i} v dF^{n-1}(v), \end{aligned}$$

since  $F^{n-1}(v)$  ( $= F(v)^{n-1}$ ) is the distribution of the highest valuation of the other agents. (Verifying the final equality rigorously is a good exercise.)

If types are uniformly distributed, then  $F^{n-1}(v) = v^{n-1}$  and  $f^{n-1}(v) = (n-1)v^{n-2}$ , so the payment of an agent  $i$  of type  $\theta_i$  is

$$\bar{t}_i^V(\theta_i) = \int_0^{\theta_i} v \cdot (n-1)v^{n-2} dv = (n-1) \int_0^{\theta_i} v^{n-1} dv = \frac{n-1}{n} (\theta_i)^n.$$

Thus an agent  $i$  of type  $\theta_i$  pays each of the others  $\frac{1}{n}(\theta_i)^n$ , and receives his shares of their payments. ♦

**Proposition 7.24.** *The AGV mechanism is Bayesian incentive compatible and budget balanced.*

*Proof.* Budget balance is clear, so we need only check Bayesian incentive compatibility. For the latter, we can ignore the redistribution term (i.e., the second term) in the AGV transfers (86), since it is independent of agent  $i$ 's announcement, and so does not affect  $i$ 's incentives. (This amounts to ignoring identical terms that would appear on each side of each inequality below.) It is thus enough to show that the mechanism  $(x^*(\cdot), \bar{t}^V(\cdot))$  based on the expected VCG transfers (85) is Bayesian incentive compatible.

To do so, we show that if agent  $i$ 's opponents are truthful, then his Bayesian incentive compatibility constraints under  $(x^*(\cdot), \bar{t}^V(\cdot))$  are averages of his dominant strategy incentive compatibility constraints from the VCG mechanism. Type  $\theta_i$ 's VCG incentive compatibility constraints require that for every announcement  $\hat{\theta}_i \in \Theta_i$  and every announcement profile  $\theta_{-i} \in \Theta_{-i}$  of  $i$ 's opponents,

$$u_i(x^*(\theta_i, \theta_{-i}), \theta_i) - t_i^V(\theta_i, \theta_{-i}) \geq u_i(x^*(\hat{\theta}_i, \theta_{-i}), \theta_i) - t_i^V(\hat{\theta}_i, \theta_{-i}).$$

Averaging over  $\theta_{-i}$  with weights  $f_{-i}(\theta_{-i})$  shows that

$$\int_{\Theta_{-i}} \left( u_i(x^*(\theta_i, \theta_{-i}), \theta_i) - t_i^V(\theta_i, \theta_{-i}) \right) f_{-i}(\theta_{-i}) d\theta_{-i} \geq \int_{\Theta_{-i}} \left( u_i(x^*(\hat{\theta}_i, \theta_{-i}), \theta_i) - t_i^V(\hat{\theta}_i, \theta_{-i}) \right) f_{-i}(\theta_{-i}) d\theta_{-i},$$

or equivalently

$$\int_{\Theta_{-i}} u_i(x^*(\theta_i, \theta_{-i}), \theta_i) f_{-i}(\theta_{-i}) d\theta_{-i} - \bar{t}_i^V(\theta_i) \geq \int_{\Theta_{-i}} u_i(x^*(\hat{\theta}_i, \theta_{-i}), \theta_i) f_{-i}(\theta_{-i}) d\theta_{-i} - \bar{t}_i^V(\hat{\theta}_i)$$

for all announcements  $\hat{\theta}_i \in \Theta_i$ . This is type  $\theta_i$ 's Bayesian incentive compatibility constraint under  $(x^*(\cdot), \bar{t}^V(\cdot))$ . (Understanding why the independence assumption is needed in this proof is a very good exercise.) ■

Remark: While we constructed the AGV mechanism by modifying the VCG mechanism, we can also obtain a Bayesian incentive compatible, budget balanced mechanism by modifying any other Groves mechanism in the same fashion. It is easy to check that the resulting transfers  $t(\cdot)$  are of the form  $t_i(\theta) = t_i^A(\theta) + c_i$ , where the constants  $c_i$  satisfy  $\sum_{i \in \mathcal{A}} c_i = 0$  (as required for budget balance). The values of these constants determine, for instance, whether the mechanism also satisfies individual rationality constraints, a topic we return to in Sections 7.4.3 and 7.4.4.

### 7.4.3 The KPW mechanism

Our aim in this section and the next is to answer the following question: in independent private values environments, when can we find a mechanism that satisfies Bayesian incentive compatibility, ex post efficiency, budget balance, and interim individual rationality. The last requirement asks that each agent be willing to participate in the mechanism after learning her own type.

### *Interim individual rationality*

What happens if an agent declines to participate in the mechanism? We now suppose that if any agent opts not to participate, then no transfers are paid, and a (predesignated) *default allocation*  $x^\dagger \in X$  obtains.

The default allocation may reflect some agent's property rights. For instance, if the agents are a buyer (or multiple buyers) and a seller, then a seller who declines to participate in the mechanism keeps the good; how much this is worth to him depends on his privately known valuation for the good. In public good provision problems, the default option is typically non-provision. An agent's payoff to non-provision may be type-dependent, if his benefits from the existing resources he uses in place of the public good are themselves type-dependent.

In such environments, all agents are willing to participate in the mechanism if it satisfies the following (*interim*) *individual rationality* constraints:

$$(87) \quad \bar{U}_i(\theta_i) \geq u_i(x^\dagger, \theta_i) \quad \text{for all } \theta_i \in \Theta_i \text{ and } i \in \mathcal{A},$$

where as usual (using notation for continuous type spaces),

$$\bar{U}_i(\theta_i) = \int_{\Theta_{-i}} \left( u_i(x^*(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) \right) f_{-i}(\theta_{-i}) d\theta_{-i}$$

denotes type  $\theta_i$ 's expected utility when all agents report truthfully.

### *The IR Groves mechanism*

It is easy to introduce subsidies to Groves mechanisms to ensure that individual rationality constraints are satisfied. We will do so starting from the plain Groves mechanism, but we will soon see that this choice is not essential (Proposition 7.27).

The plain Groves mechanism  $g(\cdot) = (x^*(\cdot), t^P(\cdot))$  has the transfer function

$$(88) \quad t_i^P(\theta) = - \sum_{j \neq i} u_j(x^*(\theta), \theta_j).$$

Type  $\theta_i$ 's expected utility under the plain Groves mechanism when all agents report truthfully is

$$\bar{U}_i^P(\theta_i) = \int_{\Theta_{-i}} \left( u_i(x^*(\theta_i, \theta_{-i}), \theta_i) - t_i^P(\theta_i, \theta_{-i}) \right) f_{-i}(\theta_{-i}) d\theta_{-i}$$

$$(89) \quad = \sum_{j \in \mathcal{A}} \int_{\Theta_{-i}} u_j(x^*(\theta_i, \theta_{-i}), \theta_j) f_{-i}(\theta_{-i}) d\theta_{-i}.$$

To construct a mechanism that satisfies interim individual rationality, we introduce type-independent rebates for each agent. Define agent  $i$ 's *rebate* by

$$(90) \quad r_i^P = \max_{\theta_i \in \Theta_i} (u_i(x^\dagger, \theta_i) - \bar{U}_i^P(\theta_i))$$

(We implicitly assume that this maximum exists, as is true under suitable continuity assumptions.) We call a type  $\theta_i^\dagger$  that achieves the maximum in (90) a *most tempted type* of agent  $i$ . The direct mechanism for  $g(\cdot) = (x^*(\cdot), t^I(\cdot))$ , where

$$(91) \quad t_i^I(\theta) = t_i^P(\theta) - r_i^P.$$

is called the *individually rational Groves mechanism*.

**Observation 7.25.** *The individually rational Groves mechanism is dominant strategy incentive compatible. It is also interim individually rational: each type's expected utility,  $\bar{U}_i^I(\theta_i) = \bar{U}_i^G(\theta_i) + r_i^G$ , is at least  $u_i(x^\dagger, \theta_i)$ , and  $\bar{U}_i^I(\theta_i^\dagger) = u_i(x^\dagger, \theta_i^\dagger)$  for any most tempted type  $\theta_i^\dagger$ .*

### **The KPW mechanism**

We now consider Bayesian incentive compatibility, interim individual rationality, ex post efficiency, and budget balance all at once.

Some perspective: Section 7.4.1 introduced Groves mechanisms, which implement the efficient allocation function  $x^*$  in dominant strategies. Section 7.4.2 introduced the AGV mechanism, under which each agent of type  $\theta_i$  makes a payment equal to his expected VCG transfer, which is then divided among the other agents. The fact that payments are disbursed in this way ensures budget balance. The use of *expected* VCG transfers combined with the independence of agents' types ensures *Bayesian* incentive compatibility. The same trick can be applied starting from any Groves mechanism to obtain a Bayesian incentive compatible, budget balanced mechanism.

Just now we introduced the IR Groves mechanism, which implemented  $x^*$  in dominant strategies and satisfied interim individual rationality. We now introduce a mechanism that, under a suitable sufficient condition, implements  $x^*$  while ensuring both budget balance and interim individual rationality. The trick is again to make payments equal to expected transfers of the simpler mechanism and to rely on independence of types. The averaging again replaces dominant strategy incentive compatibility constraints with

Bayesian ones. To achieve both budget balance and interim individual rationality, we must ensure that the agents do not find opting out of the mechanism too appealing. We present a simple sufficient condition for this to be so, and show that in allocation problems, this condition is also necessary, and thus cannot be improved upon.

Building on the previous section, let  $g(\cdot) = (x^*(\cdot), t^l(\cdot))$  be an individually rational Groves mechanism, let  $\bar{t}_i^l(\theta_i)$  be type  $\theta_i$ 's expected transfer under this mechanism, and let  $\bar{\bar{t}}_i^l$  be agent  $i$ 's ex ante expected transfer under this mechanism, assuming truth-telling in each case.

$$\begin{aligned}\bar{t}_i^l(\theta_i) &= \int_{\Theta_{-i}} t_i^l(\theta_i, \theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i} \text{ and} \\ \bar{\bar{t}}_i^l &= \int_{\Theta_i} \bar{t}_i^l(\theta_i) f_i(\theta_i) d\theta_i = \int_{\Theta} t_i^l(\theta) f(\theta) d\theta.\end{aligned}$$

By Proposition 7.27,  $\bar{t}_i^l(\theta_i)$  and  $\bar{\bar{t}}_i^l$  are the same for every individually rational Groves mechanism.

Define the *KPW mechanism* to be the direct mechanism for  $g(\cdot) = (x^*(\cdot), t^*(\cdot))$ , where

$$(92) \quad t_i^*(\theta) = \bar{t}_i^l(\theta_i) - \frac{1}{n} \bar{\bar{t}}_i^l - \frac{1}{n-1} \sum_{j \neq i} \left( \bar{t}_j^l(\theta_j) - \frac{1}{n} \bar{\bar{t}}_j^l \right).$$

Under (92), agent  $i$  of type  $\theta_i$  pays  $\bar{t}_i^l(\theta_i) - \frac{1}{n} \bar{\bar{t}}_i^l$ , giving an equal share to each other agent; he also receives his share of the other agents' payments. (The mechanism is named for Krishna and Perry (2000) and Williams (1999)—see Section 7.4.4.)

**Theorem 7.26.** *If an IR Groves mechanism  $(x^*(\cdot), t^l(\cdot))$  generates nonnegative expected revenue, then the KPW mechanism is Bayesian incentive compatible, interim individually rational, and budget balanced.*

*Proof.* To start, note that the assumption that an individually rational Groves mechanism generates nonnegative expected revenue is expressed as  $\sum_{j=1}^n \bar{\bar{t}}_j^l \geq 0$ .

Now use the equality  $\frac{1}{n} \cdot \frac{1}{n-1} = \frac{1}{n-1} - \frac{1}{n}$  to rewrite  $t^*(\cdot)$  as

$$(93) \quad t_i^*(\theta) = \bar{t}_i^l(\theta_i) - \frac{1}{n-1} \sum_{j \neq i} \bar{t}_j^l(\theta_j) + \frac{1}{n-1} \sum_{j \neq i} \bar{\bar{t}}_j^l - \frac{1}{n} \sum_{j \in \mathcal{A}} \bar{\bar{t}}_j^l.$$

Consider the direct mechanisms for  $(x^*(\cdot), t(\cdot))$ , with  $t(\cdot)$  defined by the first  $k = 1, 2, 3, 4$  terms of (93).



- (1) If  $t(\cdot)$  is defined by the first term in (93), then  $(x^*(\cdot), t(\cdot))$  satisfies IC and IR. (Averaging  $t_i^I(\theta_i, \theta_{-i})$  over  $\theta_{-i}$  does not affect  $\theta_i$ 's IR constraints, and because types are independent, it does not affect  $\theta_i$ 's IC constraints either (see the proof of Proposition 7.24).)
- (2) If  $t(\cdot)$  is defined by the first two terms in (93), then  $(x^*(\cdot), t(\cdot))$  satisfies IC and BB (since the first term has been redistributed), but it may not satisfy IR (since  $\bar{t}_j^I(\theta_j)$  may be negative).
- (3) If  $t(\cdot)$  is defined by the first three terms in (93), then  $(x^*(\cdot), t(\cdot))$  satisfies IC and IR (since  $\bar{t}_j^I$  is the expected value of  $t_j^I(\theta_j)$ ), but it may not satisfy BB.
- (4) If  $t(\cdot) = t^*(\cdot)$  is defined by the entirety of (93), then  $(x^*(\cdot), t^*(\cdot))$  satisfies IC, BB (since the third term has been redistributed), and IR (since  $\sum_j \bar{t}_j^I \geq 0$ ). ■

One can verify that the KPW mechanism is symmetric in the following sense: the slack in the individual rationality constraint of each agent  $i$ 's most tempted type is equal to the individually rational Groves mechanism's expected revenue divided by  $n$ .

### *Computing the IR Groves expected transfers and the KPW transfers*

We derived the KPW mechanism and the sufficient condition from Theorem 7.26 using the plain Groves mechanism as a starting point. We now show that starting from *any* Groves mechanism would have led us to the same conclusion. This observation is often useful when the VCG transfers take a simple form—see Example 7.31.

Recall that a Groves mechanism has transfers of the form

$$(94) \quad t_i^G(\theta) = - \sum_{j \neq i} u_j(x^*(\theta), \theta_j) + h_i(\theta_{-i}) = t_i^P(\theta) + h_i(\theta_{-i}).$$

for some  $h_i: \Theta_{-i} \rightarrow \mathbb{R}$ . Let

$$(95) \quad \bar{h}_i = \int_{\Theta_{-i}} h_i(\theta_{-i}) f_{-i}(\theta_{-i}) d\theta_{-i}$$

the expected value of the type-independent part of agent  $i$ 's transfer, which of course does not depend on agent  $i$ 's type. Type  $\theta_i$ 's expected utility under the Groves mechanism  $g(\cdot) = (x^*(\cdot), t^G(\cdot))$  when all agents report truthfully is

$$(96) \quad \begin{aligned} \bar{U}_i^G(\theta_i) &= \int_{\Theta_{-i}} \left( u_i(x^*(\theta_i, \theta_{-i}), \theta_i) - t_i^G(\theta_i, \theta_{-i}) \right) f_{-i}(\theta_{-i}) d\theta_{-i} \\ &= \int_{\Theta_{-i}} \left( u_i(x^*(\theta_i, \theta_{-i}), \theta_i) - (t_i^P(\theta_i, \theta_{-i}) + h_i(\theta_{-i})) \right) f_{-i}(\theta_{-i}) d\theta_{-i} \end{aligned}$$

$$= \sum_{j \in \mathcal{A}} \int_{\Theta_{-i}} u_j(x^*(\theta_i, \theta_{-i}), \theta_j) f_{-i}(\theta_{-i}) d\theta_{-i} - \bar{h}_i.$$

We can use introduce rebates to construct a dominant strategy incentive compatible, interim individually rational mechanism from this Groves mechanism. The rebates take the form

$$\begin{aligned} r_i^G &= \max_{\theta_i \in \Theta_i} (u_i(x^\dagger, \theta_i) - \bar{U}_i^G(\theta_i)) \\ &= \max_{\theta_i \in \Theta_i} (u_i(x^\dagger, \theta_i) - (\bar{U}_i^P(\theta_i) - \bar{h}_i)) \\ &= r_i^P + \bar{h}_i. \end{aligned}$$

(Notice that the identities of the most tempted types do not depend on the functions  $h_i$ .) The new mechanism with rebates,  $g(\cdot) = (x^*(\cdot), t^{Lh}(\cdot))$ , has transfer functions

$$\begin{aligned} (97) \quad t_i^{Lh}(\theta) &= t_i^G(\theta) - r_i^G \\ &= (t_i^P(\theta) + h_i(\theta_{-i})) - (r_i^P + \bar{h}_i) \\ &= t_i^I(\theta) + h_i(\theta_{-i}) - \bar{h}_i. \end{aligned}$$

**Proposition 7.27.**

- (i) All Groves mechanisms with rebates generate the same interim transfers for all types of all agents:

$$\bar{t}_i^{Lh}(\theta_i) = \bar{t}_i^I(\theta_i).$$

- (ii) Transfers (92) under the KPW mechanism can be computed starting from any Groves mechanism using the same formula:

$$t_i^*(\theta) = \bar{t}_i^{Lh}(\theta_i) - \frac{1}{n} \bar{t}_i^{Lh} - \frac{1}{n-1} \sum_{j \neq i} \left( \bar{t}_j^{Lh}(\theta_j) - \frac{1}{n} \bar{t}_j^{Lh} \right).$$

*Proof.* Part (i) follows from the definition of interim transfers and equations (95) and (97). Part (ii) follows from part (i) and definition (92) of the KPW transfers. ■

**An example**

Here we ask whether the KPW mechanism satisfies the desired properties in a specific Bayesian collective choice problem by computing expected transfers under the IR Groves

mechanism. In doing so, it is useful to separate out the rebate terms by writing  $\bar{t}_i^I(\theta_i) = \bar{t}_i^G(\theta_i) - r_i^G$  and  $\bar{t}_i^I = \bar{t}_i^G - r_i^G$ .

*Example 7.28. A collective choice problem.* Agent 1 (a zoologist) and agent 2 are deciding whether to jointly adopt a pet. They can adopt an alligator ( $a$ ) or a bunny ( $b$ ), or they can not adopt ( $d$ ), which is the default option. Each agent's type  $\theta_i \in [0, 1]$  represents his bravery; types are independent and uniformly distributed.

Both agent's utilities from adoption are described by  $u_i(a, \theta_i) = \theta_i$  and  $u_i(b, \theta_i) = \frac{4}{9}$ . If the agents do adopt, agent 1 must care for the pet, which costs him  $\frac{2}{3}$  for either sort of pet. To represent this, we make his utility from the default option  $u_1(d, \theta_1) = \frac{2}{3}$ . (We define his utility function this way for convenience; alternatively, we could also subtract  $\frac{2}{3}$  from his utility at every allocation.) Agent 2's utility from the default option is  $u_2(d, \theta_2) = 0$ .

Can we construct a Bayesian incentive compatible, interim individually rational, ex post efficient, and budget balanced mechanism for this problem?

The efficient allocation function  $x^*(\cdot)$  is

$$x^*(\theta) = \begin{cases} b & \text{if } \theta_1 + \theta_2 < \frac{8}{9}, \\ a & \text{if } \theta_1 + \theta_2 > \frac{8}{9}. \end{cases}$$

(Either  $a$  or  $b$  can be specified when  $\theta_1 + \theta_2 = \frac{8}{9}$ ; since all of our functions are continuous and since ties occur with probability zero we can safely ignore such cases.) Not adopting is never efficient. Thus, since  $u_i(a, \theta_i) = \theta_i$  and  $u_i(b, \theta_i) = \frac{4}{9}$  are the same for both agents, each agent  $i$ 's ex post consumption benefit under the efficient allocation function is given by

$$u_i^*(x^*(\theta), \theta_i) = \begin{cases} \frac{4}{9} & \text{if } \theta_1 + \theta_2 < \frac{8}{9}, \\ \theta_i & \text{if } \theta_1 + \theta_2 > \frac{8}{9}. \end{cases}$$

Thus the expected consumption benefit of an agent  $i$  of type  $\theta_i$  is

$$\begin{aligned} \bar{u}_i^*(\theta_i) &\equiv \int_0^1 u_i(x^*(\theta_i, \theta_j), \theta_i) d\theta_j \\ &= \begin{cases} \int_0^{8/9-\theta_i} \frac{4}{9} d\theta_j + \int_{8/9-\theta_i}^1 \theta_i d\theta_j & \text{if } \theta_i < \frac{8}{9}, \\ \int_0^1 \theta_i d\theta_j & \text{if } \theta_i > \frac{8}{9}, \end{cases} \end{aligned}$$

$$(98) \quad = \begin{cases} \theta_i^2 - \frac{1}{3}\theta_i + \frac{32}{81} & \text{if } \theta_i < \frac{8}{9}, \\ \theta_i & \text{if } \theta_i > \frac{8}{9}. \end{cases}$$

Now consider agent  $i$ 's transfers and expected transfers under the plain Groves mechanism ( $h_i(\theta_{-i}) \equiv 0$ ). Transfers under this mechanism are defined as

$$t_i^G(\theta) = -u_j(x^*(\theta), \theta_j) = \begin{cases} -\frac{4}{9} & \text{if } \theta_i + \theta_j < \frac{8}{9}, \\ -\theta_j & \text{if } \theta_i + \theta_j > \frac{8}{9}. \end{cases}$$

Type  $\theta_i$ 's interim expected transfer is thus

$$(99) \quad \begin{aligned} \bar{t}_i^G(\theta_i) &= \int_0^1 t_i^G(\theta_i, \theta_j) d\theta_j \\ &= \begin{cases} \int_0^{8/9-\theta_i} (-\frac{4}{9}) d\theta_j + \int_{8/9-\theta_i}^1 (-\theta_j) d\theta_j & \text{if } \theta_i < \frac{8}{9}, \\ \int_0^1 (-\theta_j) d\theta_j & \text{if } \theta_i > \frac{8}{9}, \end{cases} \\ &= \begin{cases} \frac{1}{2}\theta_i^2 - \frac{4}{9}\theta_i - \frac{1}{2} & \text{if } \theta_i < \frac{8}{9}, \\ -\frac{1}{2} & \text{if } \theta_i > \frac{8}{9}. \end{cases} \end{aligned}$$

Thus agent  $i$ 's ex ante expected transfer is

$$\begin{aligned} \bar{\bar{t}}_i^G &= \int_0^{8/9} (\frac{1}{2}\theta_1^2 - \frac{4}{9}\theta_1 - \frac{1}{2}) d\theta_1 + \int_{8/9}^1 (-\frac{1}{2}) d\theta_1 \\ &= (\frac{1}{6}(\frac{8}{9})^3 - \frac{2}{9}(\frac{8}{9})^2 - \frac{1}{2}(\frac{8}{9})) - \frac{1}{18} = \frac{8^3 - 2 \cdot 8^2 \cdot 6 - 8 \cdot 3^5 - 3^5}{2 \cdot 3^7} = -\frac{2443}{4374}. \end{aligned}$$

Using (98) and (99), we find that type  $\theta_i$ 's expected utility under the plain Groves mechanism is

$$\bar{U}_i^G(\theta_i) = \bar{u}_i^*(\theta_i) - \bar{t}_i^G(\theta_i) = \begin{cases} \frac{1}{2}\theta_i^2 + \frac{1}{9}\theta_i + \frac{145}{162} & \text{if } \theta_i < \frac{8}{9}, \\ \theta_i + \frac{1}{2} & \text{if } \theta_i > \frac{8}{9}. \end{cases}$$

Since this function is increasing in  $\theta_i$ , and since the agents' utilities from the default allocation do not depend on their types, their rebates are given by

$$\begin{aligned} r_1^G &= u_1(d, 0) - \bar{U}_1^G(0) = \frac{2}{3} - \frac{145}{162} = -\frac{37}{162} \quad \text{and} \\ r_2^G &= u_2(d, 0) - \bar{U}_2^G(0) = -\frac{145}{162}. \end{aligned}$$

Thus the expected revenue of the individually rational Groves mechanism is

$$\bar{t}_1^G + \bar{t}_2^G - r_1^G - r_2^G = 2 \cdot \left(-\frac{2443}{4374}\right) - \left(-\frac{37}{162} - \frac{145}{162}\right) = -\frac{2443}{3^7} + \frac{2457}{3^7} = \frac{14}{2187}.$$

Since the individually rational Groves mechanism runs a surplus, Theorem 7.26 implies that the KPW mechanism is Bayesian incentive compatible, interim individually rational, and budget balanced. Under this mechanism, an agent 1 of type  $\theta_1$  pays agent 2

$$\begin{aligned} (\bar{t}_1^G(\theta_1) - r_1) - \frac{1}{2}(\bar{t}_1^G - r_1) &= \bar{t}_1^G(\theta_1) - \frac{1}{2}\bar{t}_1^G - \frac{1}{2}r_1 \\ &= \bar{t}_1^G(\theta_1) - \frac{1}{2} \cdot \left(-\frac{2443}{4374}\right) - \frac{1}{2} \cdot \left(-\frac{37}{162}\right) \\ &= \begin{cases} \frac{1}{2}\theta_1^2 - \frac{4}{9}\theta_1 - \frac{233}{2187} & \text{if } \theta_1 < \frac{8}{9}, \\ -\frac{233}{2187} & \text{if } \theta_1 > \frac{8}{9}, \end{cases} \end{aligned}$$

and an agent 2 of type  $\theta_2$  pays agent 1

$$\begin{aligned} (\bar{t}_2^G(\theta_2) - r_2) - \frac{1}{2}(\bar{t}_2^G - r_2) &= \bar{t}_2^G(\theta_2) - \frac{1}{2}\bar{t}_2^G - \frac{1}{2}r_2 \\ &= \bar{t}_2^G(\theta_2) - \frac{1}{2} \cdot \left(-\frac{2443}{4374}\right) - \frac{1}{2} \cdot \left(-\frac{145}{162}\right) \\ &= \begin{cases} \frac{1}{2}\theta_2^2 - \frac{4}{9}\theta_2 + \frac{496}{2187} & \text{if } \theta_2 < \frac{8}{9}, \\ \frac{496}{2187} & \text{if } \theta_2 > \frac{8}{9}. \end{cases} \end{aligned}$$

One can check that agent 1's "payment" is always negative, and that agent 2's payment is always positive. ♦

#### 7.4.4 The Krishna-Perry-Williams and Myerson-Satterthwaite theorems

Theorem 7.26 shows that when an individually rational Groves mechanism runs an expected surplus, it is possible to find a mechanism that is Bayesian incentive compatible, interim individually rational, ex post efficient, and budget balanced. We now show that in the environment of the payoff and revenue equivalence theorems, the converse statement is also true. The characterization that results is due to Krishna and Perry (2000) and Williams (1999).

As in Section 7.2, we consider allocation problems with independent types and continuous type sets, and we allow for random mechanisms. (With quasilinear utility, there is no advantage in using random transfers, so we only need consider random allocations.) Since we are interested in efficiency, we require the good to be allocated to one of the agents.

$\mathcal{A} = \{1, \dots, n\}$	set of agents
$\Theta_i = [0, 1]$	set of agent $i$ 's types
	different agents' types are independent
$F_i, f_i$	cdf and pdf of agent $i$ 's type distribution
$\mathcal{X} = \Delta \mathcal{A} \times \mathbb{R}^n$	set of social alternatives
$u_i(q, \theta) - t_i = \theta_i q_i - t_i$	agent $i$ 's utility

Allocation function  $q^*(\cdot)$  is *ex post efficient* if  $k \notin \operatorname{argmax}_j \theta_j$  implies that  $q_k(\theta) = 0$ .

**Theorem 7.29** (Krishna-Perry-Williams).

Suppose that allocation function  $q^*(\cdot)$  is *ex post efficient*. There is a transfer function  $t(\cdot)$  such that  $g^*(\cdot) = (q^*(\cdot), t(\cdot))$  is Bayesian incentive compatible, interim individually rational, and budget balanced if and only if an individually rational Groves mechanism  $(q^*(\cdot), t^1(\cdot))$  generates nonnegative expected revenue.

*Proof.* Sufficiency follows from Theorem 7.26 (which extends immediately to random mechanisms). To establish necessity, suppose that a mechanism with the desired properties exists. Budget balance implies that the expected revenue of this mechanism is zero. Thus the following proposition, which drops the budget balance constraint but retains the others, implies that individually rational Groves mechanisms have a nonnegative expected revenue.

**Proposition 7.30.** Among Bayesian incentive compatible, interim individually rational mechanisms with *ex post efficient* allocation function  $q^*(\cdot)$ , an individually rational Groves mechanism maximizes the expected transfer from each type of each agent, and so maximizes expected revenue.

*Proof.* By payoff equivalence (Theorem 7.11), any Bayesian incentive compatible mechanism  $(q^*(\cdot), t(\cdot))$  that implements allocation function  $q^*(\cdot)$  has expected utility functions that satisfy  $\bar{U}_i(\theta_i) = \bar{U}_i^I(\theta_i) + c_i$ .

Observation 7.25 says that the interim individual rationality constraint binds for some type of agent  $i$  under an individually rational Groves mechanism. Thus the new mechanism satisfies interim individual rationality if and only if  $c_i \geq 0$ . As the expected transfer of type  $\theta_i$  under the new mechanism is

$$\bar{t}_i(\theta_i) = \theta_i \bar{q}_i^*(\theta_i) - \bar{U}_i(\theta_i) = \theta_i \bar{q}_i^*(\theta_i) - (\bar{U}_i^I(\theta_i) + c_i)$$

we can do no better than choosing  $c_i = 0$ , as under an individually rational Groves mechanism. ■

Remark: In the proof above, the restriction to allocation problems was used in the appeal to payoff equivalence (Theorem 7.11). Since the conclusion of that theorem remains true in all linear utility environments (Example 7.4), so does that of Theorem 7.29. So, for example, if in the collective choice problem from Example 7.28 we had found that an individually rational Groves mechanism generated negative expected revenue, it would have followed that no mechanism satisfying all four desiderata exists.

As an application of Theorem 7.29, we prove the Myerson-Satterthwaite (1983) theorem on the nonexistence of ex post efficient bilateral trading mechanisms.

*Example 7.31. Bilateral trade.* The owner and potential seller of a good values it at  $\theta_s$ . A potential buyer values it at  $\theta_b$ . These valuations are drawn from distributions on  $[0, 1]$  with positive densities  $f_s$  and  $f_b$ . Thus there is a positive probability that there are gains from trade, and also a positive probability that there are not, and either of these probabilities may be arbitrarily close to 1.

When the allocation probabilities and transfers are  $q = (q_b, q_s)$  and  $t = (t_b, t_s)$ , the buyer's utility is  $u_b(q, \theta_b) - t_b = q_b\theta_b - t_b$ , and the seller's is  $u_s(q, \theta_s) - t_s = q_s\theta_s - t_s$ .

Since the seller owns the good at the start, the default allocation is  $q^\dagger = (q_b^\dagger, q_s^\dagger) = (0, 1)$ , so that  $u_b(q^\dagger, \theta_b) = 0$  and  $u_s(q^\dagger, \theta_s) = \theta_s$ .

The two agents would like to design a mechanism ensuring that trade occurs whenever it is ex post efficient. The mechanism must be Bayesian incentive compatible and interim individually rational, and must not require payments to or from a third party.

The following result shows that no such mechanism exists.

**Theorem 7.32** (Myerson-Satterthwaite).

*There is no Bayesian incentive compatible, interim individually rational, ex post efficient, and budget balanced bilateral trading mechanism.*

*Proof.* We apply Theorem 7.29 and Proposition 7.27 (to take the VCG transfers as our starting point). Ignoring ties, efficient allocation requires

$$q_b^*(\theta) = \begin{cases} 1 & \text{if } \theta_b > \theta_s, \\ 0 & \text{if } \theta_b < \theta_s, \end{cases} \quad \text{and} \quad q_s^*(\theta) = \begin{cases} 0 & \text{if } \theta_b > \theta_s, \\ 1 & \text{if } \theta_b < \theta_s. \end{cases}$$

Thus the VCG transfers are

$$t_b^V(\theta) = \sum_{j \neq b} \theta_j (q_j^{-i}(\theta_{-i}) - q_j^*(\theta)) = \theta_s(1 - q_s^*(\theta)) = \theta_s q_b^*(\theta) \quad \text{and}$$

$$t_s^V(\theta) = \theta_b q_s^*(\theta).$$

Notice that  $t_b^V(\theta) + t_s^V(\theta) = \min\{\theta_b, \theta_s\}$  for all  $\theta \in \Theta$ . It follows that

$$\begin{aligned} (100) \quad \bar{t}_b^V + \bar{t}_s^V &= \int_{\Theta} (t_b^V(\theta) + t_s^V(\theta)) f(\theta) d\theta \\ &= \int_{\Theta} \min\{\theta_b, \theta_s\} f(\theta) d\theta \\ &< \int_{\Theta} \theta_b f(\theta) d\theta \\ &= \int_{\Theta_b} \theta_b f_b(\theta_b) d\theta_b \\ &\equiv \bar{\theta}_b. \end{aligned}$$

The strict inequality reflects the fact that  $\theta_b > \theta_s$  with positive probability.

Now observe that if the buyer is of type  $\theta_b = 0$ , then (with probability 1) he does not receive the good and pays no transfer ( $q_b^*(0, \theta_s) = 0$  and  $t_b^V(0, \theta_s) = \theta_s q_b^*(0, \theta_s) = 0$  whenever  $\theta_s > 0$ ), so his interim expected utility under the VCG mechanism is  $\bar{U}_b^V(0) = 0$ .

Also, if the seller is of type  $\theta_s = 1$ , then (with probability 1) she keeps the good and pays a transfer of  $\theta_b$  ( $q_s^*(\theta_b, 1) = 1$  and  $t_s^V(\theta_b, 1) = \theta_b q_s^*(\theta_b, 1) = \theta_b$  whenever  $\theta_b < 1$ ), so her interim expected utility under the VCG mechanism is  $\bar{U}_s^V(1) = 1 - \bar{t}_s^V(1) = 1 - \bar{\theta}_b$ .

It then follows from (96) and (90) that the rebates in the individually rational VCG mechanism must satisfy

$$\begin{aligned} r_b &\geq u_b(q^+, 0) - \bar{U}_b^V(0) = 0 - 0 = 0 \quad \text{and} \\ r_s &\geq u_s(q^+, 1) - \bar{U}_s^V(1) = 1 - (1 - \bar{\theta}_b) = \bar{\theta}_b. \end{aligned}$$

Combining these inequalities with (100) shows that the expected surplus of the individually rational VCG mechanism is

$$\bar{t}_b^V + \bar{t}_w^V - (r_b + r_s) < \bar{\theta}_b - (0 + \bar{\theta}_b) = 0.$$

Thus by Theorem 7.29, the desired bilateral trading mechanism does not exist. ■

Remark: Before the development of information economics, it was sometimes taken for granted that with property rights and enforceable contracts, allocations of goods among individuals would be more-or-less efficient. After all, when allocations are inefficient,



mutually beneficial trades are available, and so will be undertaken voluntarily unless the inefficiency is negligible.

The Myerson-Satterthwaite theorem shows that when individuals possess private information, this claim is simply not true.

## 7.5 Allocative Efficiency with Interdependent Values

Section 7.4 considered the implementation of efficient allocations in private values environments. Here we study the implementation of efficient allocations when each agent's utility may depend directly on other agents' private information. We restrict attention to the problem of allocating an indivisible private good. Until the end of the section, we assume that each agent's type, here interpreted as a signal about the quality of the good, is one-dimensional.

$\mathcal{A} = \{1, \dots, n\}$	set of agents = set of allocations
$\Theta_i = [0, 1]$	set of agent $i$ 's types
$F$	cdf of the joint distribution of agents' types
$X = \mathcal{A}$	set of allocations
$v_i(\theta)$	agent $i$ 's valuation for the good (differentiable)
$v_i(\theta) - t_i$	agent $i$ 's utility if he receives the good and pays $t_i$
$-t_i$	agent $i$ 's utility if he does not receive the good and pays $t_i$

Notice that agent  $i$ 's valuation for the good may depend on all agents' signals. We assume throughout that  $i$ 's valuation for the good is increasing in his own signal; in particular,

$$(101) \quad \frac{\partial v_i}{\partial \theta_i}(\theta) > 0 \quad \text{for all } i \in \mathcal{A} \text{ and } \theta \in \Theta.$$

An allocation function  $x(\cdot)$  is *ex post efficient* if for each  $\theta \in \Theta$ , an agent with the highest valuation  $v_i(\theta)$  is allocated the good.

The following example shows that ex post efficiency cannot always be achieved:

*Example 7.33.* Suppose that there are two agents, and that only agent 1 receives a signal. The agents' value functions are  $v_1(\theta_1) = \theta_1 + 1$  and  $v_2(\theta_1) = 3\theta_1$ . Ignoring boundary cases, the efficient allocation function has  $x^*(\theta_1) = 1$  when  $\theta_1 < \frac{1}{2}$  and  $x^*(\theta_1) = 2$  when  $\theta_1 > \frac{1}{2}$ .

Only player 1 has private information, so we need only specify the transfer function  $t_1(\cdot)$ . Moreover, the direct mechanism for  $g(\cdot) = (x^*(\cdot), t_1(\cdot))$  is a single-agent decision problem, so the only notion of implementation is to require optimal choices by each type of agent 1.

There is no transfer scheme  $t_1(\cdot)$  that implements  $(x^*(\cdot), t_1(\cdot))$ . This follows immediately from Lemma 6.8, but it is easy to show directly. Let  $\theta_1^\ell < \frac{1}{2} < \theta_1^h$ . Then incentive compatibility requires that  $\theta_1^\ell + 1 - t_1(\theta_1^\ell) \geq -t_1(\theta_1^h)$  and that  $-t_1(\theta_1^h) \geq \theta_1^h + 1 - t_1(\theta_1^\ell)$ . Adding these inequalities yields  $\theta_1^\ell \geq \theta_1^h$ , a contradiction.  $\diamond$

What prevents efficient implementation in Example 7.33 is that while increasing agent 1's signal makes the good more valuable to him, it can also make it efficient to allocate the good to agent 2. To obtain a positive result we must exclude this possibility. This can be accomplished by assuming that *own signals matter most*:

$$(102) \quad \frac{\partial v_i}{\partial \theta_i}(\theta) > \frac{\partial v_j}{\partial \theta_i}(\theta) \quad \text{for all } i \in \mathcal{A}, j \neq i, \text{ and } \theta \in \Theta.$$

In words: increasing agent  $i$ 's signal increases his valuation for the good more than it increases other agents' valuations for the good.

If own signals matter most, efficient implementation is possible, not only in the Bayesian sense, but also in the ex post sense. Recall from Section 4.6 that a Bayesian strategy profile  $s^*$  is an *ex post equilibrium* if no type  $\theta_i$  of any agent  $i$  would benefit from deviating from his specified action  $s_i^*(\theta_i)$  regardless of the realizations  $\theta_{-i}$  of the other agents' types (see equation (32)).

Given a direct mechanism  $\mathcal{M}^d = \{\{\Theta_i\}_{i \in \mathcal{A}}, g\}$  with social choice function  $g$ , we say that  $\mathcal{M}^d$  or  $g$  is *ex post incentive compatible* if truth-telling is an ex post equilibrium. The revelation principle also holds for ex post implementation, and takes the same form as those for Bayesian and dominant strategy implementation (see Section 7.1.4).

Let  $x^*(\cdot)$  be an ex post efficient allocation function, and let

$$(103) \quad m_i(\theta_{-i}) = \inf \left\{ \underline{\theta}_i : v_i(\underline{\theta}_i, \theta_{-i}) \geq \max_{j \in \mathcal{A}} v_j(\underline{\theta}_i, \theta_{-i}) \right\}$$

be the lowest type that agent  $i$  could be and still value the good the most, given that the others' types are  $\theta_{-i}$ . (Note that if  $i$  could not value the good the most (in which case  $x^*(\theta) \neq i$ ), then  $m_i(\theta_{-i}) = \inf \emptyset = +\infty$ .) The *generalized VCG mechanism* is the direct mechanism for  $g(\cdot) = (x^*(\cdot), t(\cdot))$  with transfer function

$$t_i(\theta) = \begin{cases} v_i(m_i(\theta_{-i}), \theta_{-i}) & \text{if } x^*(\theta) = i, \\ 0 & \text{if } x^*(\theta) \neq i. \end{cases}$$

In words, if agent  $i$  is assigned the good when the others' types are  $\theta_{-i}$ , he pays the

valuation he would have had for the good if he were the lowest type consistent with him receiving the good. This mechanism becomes a second-price auction when values are private, and it also generalizes the interdependent value auction studied in Proposition 4.13.

**Proposition 7.34** (Maskin (1992)).

*Suppose that own signals matter most. Then the generalized VCG mechanism is ex post incentive compatible.*

*Proof.* Fix any  $\theta_{-i}$ , and assume that agent  $i$ 's opponents report truthfully. We need to show that it is optimal for agent  $i$  to report truthfully regardless of his type.

Suppose first that  $\theta_i$  is such that  $x^*(\theta) = i$ . Then agent  $i$  values the good at least as much as all opponents, so (102) and (103) imply that  $\theta_i \geq m_i(\theta_{-i})$ . If  $i$  reports truthfully, he receives the good and pays

$$(104) \quad v_i(m_i(\theta_{-i}), \theta_{-i}) = \max_{j \in \mathcal{A}} v_j(m_i(\theta_{-i}), \theta_{-i}),$$

where the equality follows from (103). If  $\theta_i = m_i(\theta_{-i})$ , then  $i$ 's payoff is zero; thus if  $\theta_i > m_i(\theta_{-i})$ , then (102) implies that  $i$ 's payoff is positive. Now if  $i$  reports  $\hat{\theta}_i > m_i(\theta_{-i})$ , then (102) and (103) imply that

$$v_i(\hat{\theta}_i, \theta_{-i}) > \max_{j \in \mathcal{A}} v_j(\hat{\theta}_i, \theta_{-i}),$$

so  $i$  still receives the good and pays (104), so he is indifferent between reporting  $\hat{\theta}_i$  and  $\theta_i$ . If  $i$  reports  $\check{\theta}_i < m_i(\theta_{-i})$ , then (102) and (103) imply that

$$v_i(\check{\theta}_i, \theta_{-i}) < \max_{j \in \mathcal{A}} v_j(\check{\theta}_i, \theta_{-i});$$

thus  $i$  does not receive the good and earns a payoff of zero. If  $i$  reports  $m_i(\theta_{-i})$ , then his payoff is zero whether or not he receives the good. Thus truthful reporting is optimal.

Now suppose that  $\theta_i$  is such that  $x^*(\theta) \neq i$ . Then  $i$  has an opponent  $j$  who values the good at least as much as  $i$  does, and so (102) and (103) imply that  $\theta_i \leq m_i(\theta_{-i})$ . If  $i$  reports truthfully, or reports any  $\check{\theta}_i < m_i(\theta_{-i})$ , his payoff is zero. If  $i$  reports  $\hat{\theta}_i > m_i(\theta_{-i})$ , then  $i$  receives the good and pays  $v_i(m_i(\theta_{-i}), \theta_{-i})$ ; thus his payoff is 0 if  $\theta_i = m_i(\theta_{-i})$ , and is negative if  $\theta_i < m_i(\theta_{-i})$  (by (101)). Thus truthful reporting is again optimal. ■

Remarks:

- (i) As in a second price auction, an agent's report here helps determine whether he receives the good, but does not affect what he pays if he does receive it. The latter property is key for truthful reporting to be optimal.
- (ii) The assumption (102) that own signals matter most is stronger than necessary. It is enough that the inequality in (102) hold at signal profiles  $\theta$  for which multiple agents have the highest valuation, and then the inequality need only hold for those agents. See Dasgupta and Maskin (2000).
- (iii) The generalized VCG mechanism requires the designer to have detailed information about the agents' valuation functions. Dasgupta and Maskin (2000) and Perry and Reny (2002) introduce mechanisms that require the planner to have far less information. However, these mechanisms do require the agents to have detailed information about one another's preferences.

The analysis above considered allocation of a single good in an interdependent values environment with one-dimensional signals. If multiple goods are to be allocated, one might expect agents' signals to be multidimensional. In such cases, the possibilities for efficient implementation with interdependent values are much more grim. Jehiel and Moldovanu (2001) consider Bayesian implementation with multidimensional signals under the assumption that different agents' signals are independent (for reasons that will become clear in Section 7.6). They show that efficient implementation is impossible except in borderline cases. Jehiel et al. (2006) consider the more demanding notion of ex post implementation, and show that beyond borderline cases, only *constant* social choice functions—those under which the same allocation is chosen regardless of the type profile—are implementable when values are interdependent.

## 7.6 Correlated Types, Full Surplus Extraction, and the Wilson Doctrine

### 7.6.1 Full surplus extraction via the Crémer-McLean mechanism

So far, all of the results that ensure Bayesian implementation of social choice functions have relied on the assumption of independent types. (By definition, the results ensuring dominant strategy implementation under private values (VCG, IR-VCG) and ex post implementation (generalized VCG) did not require this assumption.) We will now see that what can be achieved through Bayesian implementation expands radically when correlation of types is introduced. We will then explain why this change takes advantage of a naive aspect of standard mechanism design models.

We consider a private values environment like that in Section 7.4, except that we allow for

correlated types. We also assume that type spaces are finite.

$\mathcal{A} = \{1, \dots, n\}$	set of agents
$\Theta_i$	finite set of agent $i$ 's types
$p \in \Delta(\Theta)$	common prior distribution
$X$	finite set of allocations
$\mathcal{X} = X \times \mathbb{R}^n$	set of social alternatives
$u_i(x, \theta_i) - t_i$	agent $i$ 's utility

This section presents the *full surplus extraction* theorem of Crémer and McLean (1985, 1988). In essence, this result shows that when there is any correlation among agents' types, there is a dominant strategy mechanism which implements the efficient allocation function, while delivering *all* of the agents' benefits from the interaction to the designer. For instance, in allocation problems, the designer is able to achieve perfect ex post price discrimination despite only knowing the prior distribution on agents' types!

Let  $x^*(\cdot)$  be an ex post efficient allocation function (see (81)). In Section 7.4.1, we saw that the VCG mechanism, the direct mechanism with allocation function  $x^*(\cdot)$  and transfer function  $t^V(\cdot)$  (see (84)), is dominant strategy incentive compatible. Player  $i$ 's interim expected utility under this mechanism is

$$(105) \quad \bar{u}_i^V(\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i} | \theta_i) (u_i(x^*(\theta), \theta_i) - t_i^V(\theta)).$$

Let the matrix  $P_i \in \mathbb{R}^{\Theta_i \times \Theta_{-i}}$ , with elements  $p_i(\theta_{-i} | \theta_i)$ , describe agent  $i$ 's interim beliefs. Thus the  $\theta_i$ th row of  $P_i$  describes type  $\theta_i$ 's interim beliefs. If types were independent, then the rows of  $P_i$  would be identical. In contrast, the full extraction result requires the rows of  $P_i$  to be linearly independent.

**Theorem 7.35** (Crémer and McLean (1985, 1988)).

*Suppose that the interim belief matrices  $P_i$  have full row rank. Then there is a dominant strategy incentive compatible mechanism with ex post efficient allocation function  $x^*(\cdot)$  under which the interim expected utility of every type of every agent is zero.*

Strikingly, Theorem 7.35 reveals that the slightest departure from independent types allows for full surplus extraction via a dominant strategy mechanism. We call the direct mechanism  $g(\cdot) = (x^*(\cdot), t^{CM}(\cdot))$  posited in the theorem the *Crémer-McLean mechanism*.

*Proof.* Let  $\bar{u}_i^V \in \mathbb{R}^{\Theta_i}$  be the column vector whose elements are agent  $i$ 's interim expected utilities under the VCG mechanism (105). Since  $P_i \in \mathbb{R}^{\Theta_i \times \Theta_{-i}}$  has full row rank, there is a

column vector  $\tau_i \in \mathbb{R}^{\Theta_{-i}}$  satisfying

$$(106) \quad P_i \tau_i = \bar{u}_i^V.$$

Now define the Crémer-McLean transfers by

$$(107) \quad t_i^{CM}(\theta) = t_i^V(\theta) + \tau_i(\theta_{-i}).$$

Since truth-telling is dominant under the VCG mechanism, and since the  $\tau_i$  term in (107) does not depend on  $\theta_i$ , truth-telling is also dominant under the Crémer-McLean mechanism. Furthermore, equations (107) and (106) imply that under this mechanism, the interim expected utility of an agent  $i$  of type  $\theta_i$  is

$$\bar{u}_i^{CM}(\theta_i) = \bar{u}_i^V(\theta_i) - \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \tau_i(\theta_{-i}) = \bar{u}_i^V(\theta_i) - \bar{u}_i^V(\theta_i) = 0. \quad \blacksquare$$

The proof of Theorem 7.35 divulges what  $P_i$  having full-row rank actually buys us. For *any* function  $c_i: \Theta_i \rightarrow \mathbb{R}$ , the planner can design transfers  $\tau_i: \Theta_{-i} \rightarrow \mathbb{R}$  with the following property: If agent  $i$ 's opponents report truthfully, then for each type  $\theta_i \in \Theta_i$ , the expected cost of the transfers is  $c_i(\theta_i)$ . Thus by setting  $c_i(\cdot)$  equal to  $\bar{u}_i^V(\cdot)$ , the principal is able to make each type's expected  $\tau_i(\cdot)$  payment exactly equal to its expected VCG surplus. This makes full surplus extraction possible.

(In this argument, we are allowed to suppose that  $i$ 's opponents report truthfully because we are addressing expected transfers, which are evaluated at the mechanism's equilibrium. We could not have supposed this when checking for dominant strategy incentive compatibility, but this property already follows from the Crémer-McLean mechanism being a Groves mechanism.)

Remarks:

- (i) An analogue of Theorem 7.35 holds for settings with interdependent values, if we assume as in Proposition 7.34 that own signals matter most. In this case, starting with the generalized VCG mechanism and proceeding as in the proof of Theorem 7.35 shows that there is an ex post incentive compatible mechanism that achieves full surplus extraction.
- (ii) When types are "almost independent", so that  $P_i$  "almost fails" to have full rank, some of the transfer terms  $\tau_i(\theta_{-i})$  may need to be very large. This makes sense given the fact that when types are independent, full surplus extraction is no longer possible. Robert (1991) shows that if one imposes a *limited liability* constraint,

placing an upper bound on the amount that agents may be required to pay, then the planner's optimal expected payoffs are continuous in the distribution of valuations, implying that full surplus extraction cannot be achieved when the types are "almost independent".

- (iii) The Crémer-McLean mechanism requires the designer to have detailed knowledge about agents' beliefs, and in particular about the correlations between different agents' types. We return to this point below.

### 7.6.2 Full surplus extraction via menus of lotteries

To understand what makes full surplus extraction possible, we introduce a mechanism due to Neeman (2004) that uses the more demanding requirement of Bayesian implementation, but achieves full *ex post* surplus extraction.

Let  $Z \in \mathbb{R}^{\Theta_i \times \Theta_i}$  be the matrix whose diagonal elements equal 0 and whose off-diagonal elements equal  $-K$ , where  $K > 0$  is large. By the full row rank condition on  $P_i$ , there is a matrix  $L_i \in \mathbb{R}^{\Theta_{-i} \times \Theta_i}$  that satisfies  $P_i L_i = Z$ . Each column  $L_i^{\hat{\theta}_i} \in \mathbb{R}^{\Theta_{-i}}$  of  $L_i$  can be interpreted as a "lottery" whose outcome  $L_i^{\hat{\theta}_i}(\theta_{-i}) \in \mathbb{R}$  depends on the announcements of agent  $i$ 's opponents.

Suppose that the designer asks each agent  $i$  to choose a lottery  $L_i^{\hat{\theta}_i}$  by reporting a type  $\hat{\theta}_i$ . If the agent's actual type is  $\theta_i$ , his opponents report truthfully, and he reports  $\hat{\theta}_i$ , then his expected payoff from the lottery is

$$\sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i} | \theta_i) L_i^{\hat{\theta}_i}(\theta_{-i}) = Z_{\theta_i \hat{\theta}_i}.$$

This equals 0 if  $\theta_i = \hat{\theta}_i$  and equals  $-K$  otherwise. Thus it is a Bayesian equilibrium for all agents to report truthfully, and if they do so their expected transfers are zero.

Once all agents report their types truthfully, the planner can determine the efficient allocation  $x^*(\theta)$  and each agent's benefit  $u_i(x^*(\theta), \theta_i)$  at this allocation. He can thus implement allocation  $x^*(\theta)$ , and require each agent to pay his benefit  $u_i(x^*(\theta), \theta_i)$ . If the penalty  $K$  from the lotteries is large enough, the consequences of the reports made in the "first stage" for payoffs in the "second stage" will not dissuade the agents from reporting truthfully.

We thus have

**Proposition 7.36.** *Suppose that the interim belief matrices  $P_i$  have full row rank. Then the Neeman mechanism Bayesian implements the ex post efficient allocation function  $x^*(\cdot)$  and gives every type of every agent an ex post utility of zero.*

Under the Neeman mechanism, the way that the designer learns each agent's preferences over social alternatives is split into two distinct pieces. The agent's choice of lottery reveals his beliefs about his opponents' types. By the full row rank condition, these beliefs in turn reveal the agent's type, and thus his preferences over the social alternatives.

### 7.6.3 The beliefs-determine-preferences property and the Wilson doctrine

#### *The beliefs-determine-preferences property*

Neeman (2004) develops these last observations into a compelling criticism of the relevance of full extraction results. This criticism returns us to our initial discussion of Bayesian games in Section 4.1. There we noted that a player's type captures two separate aspects of his private information: his payoff-relevant information, and his beliefs about others' types. If types are independent, then interim beliefs are common knowledge, so only the payoff-relevant aspect of types remains private information. But if types are correlated, then neither the payoff-relevant information nor interim beliefs are common knowledge. The Crémer-McLean mechanism works by exploiting this joint uncertainty. It is able to do so because of a peculiarity in the standard specification of mechanism design problems.

Traditionally, mechanism design problems are constructed by first specifying a set of "payoff types", which describe each agent's payoff relevant information (see Section 4.8), and then specifying the beliefs of each payoff type about opponents' payoff types. Theorem 7.35 and the discussion that follows show that this approach to defining Bayesian games can have unexpected consequences.

To make this idea precise, Neeman (2004) notes that the full rank condition on the interim belief matrix  $P_i$  imposes the *beliefs-determine-preferences property*: it rules out the possibility that there are two types of player  $i$  who have the same beliefs about opponents' types, but have different payoff types. For instance, in an allocation problem, it is plausible to have two types of agent  $i$  that assign different values to the good, but that share the same beliefs about opponents' private information. Given any menu of lotteries whose outcomes depend on the realization of  $\theta_{-i}$ , these two types will make the same choice, but this choice does not reveal the chooser's preferences. This precludes full surplus extraction.

#### *The Wilson doctrine*

Wilson (1985, 1987) proposes two desirable properties that mechanisms should possess in order to have practical relevance. Wilson (1985) suggests that mechanisms should



not require the designer to have very precise information about the agents' beliefs and preferences. (Dasgupta and Maskin (2000) call mechanisms with this property *detail-free*.) The standard auction formats have this property, as do the interdependent value auctions of Dasgupta and Maskin (2000) and Perry and Reny (2002) mentioned in Section 7.5; the Crémer-McLean mechanism does not. In a similar vein, Wilson (1987) proposes that useful mechanisms should not rely on unrealistically strong assumptions about what is common knowledge among the players. Either or both of these criteria are referred to as the *Wilson doctrine*.

#### 7.6.4 Robust mechanism design

The *robust mechanism design* literature is motivated both by Neeman's (2004) criticism of full extraction mechanisms and by the Wilson doctrine. Broadly speaking, its aims are to see what social choice functions can be implemented under weak assumptions about players' beliefs. Implementation is often based on weak solution concepts, though as explained below, this sometimes comes as a conclusion of the analysis rather than as an assumption. We describe some specific contributions below. (The discussion may be easier to follow after reviewing Sections 4.8 and 4.9.) See Bergemann and Morris (2012) for a survey of the literature.

A natural starting point for robust mechanism design is a "pre mechanism design environment" in which agents' preferences and payoff types are specified, but their beliefs are not. (This is a direct analogue of a pre-Bayesian game—see Section 4.8). Bergemann and Morris (2005) seek conditions that describe when a social choice function is Bayesian implementable under any type spaces and beliefs that can be introduced to turn the pre mechanism design environment into a fully-specified mechanism design environment. The necessary and sufficient condition is that the social choice function be *ex post* implementable—a solution concept that is well-defined for the pre mechanism design environment. In the quasilinear environments considered here, one might also ask about implementation of social choice correspondences that fix an allocation function (e.g., the efficient allocation function) but allow any transfer functions. Bergemann and Morris (2005) show that the previous result holds in this context as well. These results have the happy consequence that requiring robustness over all possible beliefs leads us to a solution concept that is relatively easy to check.

The larger part of this literature looks at *robust full implementation*, which in addition to robustness requires that the social choice function to be implemented is the unique equilibrium outcome regardless of the specification of beliefs. (Previous work on full implementation (without robustness) was noted briefly in Section 7.1.3.) The general description above can be specified precisely in a variety of ways. For instance, Bergemann and Morris (2008) consider *ex post implementation*, characterizing social choice functions  $g$  for which there are mechanisms whose unique *ex post* equilibrium outcome is  $g$ .

Bergemann and Morris (2009, 2011) consider a more demanding requirement that they

dub *robust implementation*. Starting from a pre mechanism design environment, it requires not only that the social choice function  $g$  be an ex post equilibrium of the mechanism, but also that for every possible specification of beliefs, all Bayesian equilibrium outcomes are close to  $g$ . They prove that this requirement is nearly equivalent to  $g$  being implementable using the solution concept of belief-free rationalizability, which is obtained by applying iterated dominance without imposing any common knowledge restrictions on agents' beliefs (see Section 4.9).

Oury and Tercieux (2012) introduce what they term *continuous implementation*, which can be roughly described as a “local” analogue of the approach of Bergemann and Morris (2009, 2011). They fix a mechanism design environment, and consider the implications of a social choice function being Bayesian implementable not only for that problem, but for all problems in which agents' beliefs are close to those in the original problem. They establish that the preceding requirement is tightly linked to the notion of full implementation using the solution concept of interim correlated rationalizability. The logic behind this is similar to the reasoning behind Weinstein and Yildiz's (2007) results on the robustness of ICR (see Section 4.9).

The conclusions of Bergemann and Morris (2009, 2011) are stated in terms of a solution concept that imposes no restrictions on agents' beliefs, while those of Oury and Tercieux (2012) use a solution concept for a fully specified Bayesian game. Taking a middle route, Ollár and Penta (2015) consider full implementation in settings with common knowledge of some belief restrictions. That is, in addition to the pre mechanism design environment, certain restrictions on the beliefs held by each payoff type  $\theta_i$  of agent  $i$  about other agents' payoff types  $\theta_j$  are common knowledge. (For instance, rather than assuming that it is common knowledge that payoff types are i.i.d. with a uniform distribution on the unit interval (as one might in a usual mechanism design environment), it might only be assumed common knowledge that the conditional expectation of each payoff type of agent  $i$  about the payoff type of agent  $j \neq i$  is  $\frac{1}{2}$ .) Full implementation under such belief restrictions is defined using Battigalli and Siniscalchi's (2003) notion of  $\Delta$ -rationalizability (see Section 4.9). Ollár and Penta (2015) show that in a variety of applications, mild belief restrictions are often enough to significantly expand the set of implementable social choice functions relative to the belief-free approach studied in Bergemann and Morris (2009, 2011).

## References

- Abreu, D., Dutta, P. K., and Smith, L. (1994). The folk theorem for repeated games: A NEU condition. *Econometrica*, 62:939–48.
- Abreu, D., Pearce, D., and Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58:1041–63.

- Akerlof, G. A. (1970). The market for 'lemons': quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84:488–500.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica*, 21:503–546.
- Anscombe, F. J. and Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205.
- Arrow, K. J. (1951). *Social Choice and Individual Values*. Wiley, New York.
- Arrow, K. J. (1979). The property rights doctrine and demand revelation under incomplete information. In Boskin, M., editor, *Economics and Human Welfare*, pages 23–39. Academic Press, New York.
- Asheim, G. B. (2001). Proper rationalizability in lexicographic beliefs. *International Journal of Game Theory*, 30:453–478.
- Asheim, G. B. (2002). On the epistemic foundation for backward induction. *Mathematical Social Sciences*, 44:121–144.
- Asheim, G. B. (2006). *The Consistent Preferences Approach to Deductive Reasoning in Games*. Springer, Dordrecht.
- Asheim, G. B. and Perea, A. (2005). Sequential and quasi-perfect rationalizability in extensive games. *Games and Economic Behavior*, 53:15–42.
- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96.
- Aumann, R. J. (1990). Nash equilibria are not self-enforcing. In Gabsewicz, J.-J., Richard, J.-F., and Wolsey, L., editors, *Economic Decision-Making: Games, Econometrics, and Optimisation*, pages 201–206. North Holland, Amsterdam.
- Avis, D., Rosenberg, G. D., Savani, R., and von Stengel, B. (2010). Enumeration of Nash equilibria for two-player games. *Economic Theory*, 42:9–37.
- Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory*, 26:445–469.
- Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica*, 55:647–661.
- Baron, D. P. and Myerson, R. B. (1982). Regulating a monopolist with unknown costs. *Econometrica*, 50:911–930.
- Battigalli, P. (1996). Strategic independence and perfect Bayesian equilibrium. *Journal of Economic Theory*, 70:201–234.

- Battigalli, P. (1997). On rationalizability in extensive games. *Journal of Economic Theory*, 74:40–61.
- Battigalli, P. and Guaitoli, D. (1997). Conjectural equilibria and rationalizability in a game with incomplete information. In Battigalli, P., Montesano, A. M., and Panunzi, F., editors, *Decisions, Games, and Markets*, pages 97–124. Springer.
- Battigalli, P. and Siniscalchi, M. (1999). Interactive beliefs, epistemic independence, and strong rationalizability. *Research in Economics*, 53:247–273.
- Battigalli, P. and Siniscalchi, M. (2002). Strong belief and forward induction reasoning. *Journal of Economic Theory*, 106:356–391.
- Battigalli, P. and Siniscalchi, M. (2003). Rationalization and incomplete information. *B. E. Journal of Theoretical Economics (Advances)*, 3(1):article 3.
- Bellman, R. (1957). *Dynamic Programming*. Princeton University Press, Princeton.
- Ben-Porath, E. (1997). Rationality, Nash equilibrium and backwards induction in perfect-information games. *Review of Economic Studies*, 64:23–46.
- Benoît, J.-P. and Krishna, V. (1985). Finitely repeated games. *Econometrica*, 53:905–922.
- Bergemann, D. and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73:1771–1813.
- Bergemann, D. and Morris, S. (2008). Ex post implementation. *Games and Economic Behavior*, 63:527–566.
- Bergemann, D. and Morris, S. (2009). Robust implementation in direct mechanisms. *Review of Economic Studies*, 76:1175–1206.
- Bergemann, D. and Morris, S. (2011). Robust implementation in general mechanisms. *Games and Economic Behavior*, 71:261–281.
- Bergemann, D. and Morris, S. (2012). An introduction to robust mechanism design. *Foundations and Trends in Microeconomics*, 8:169–230.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52:1007–1028.
- Blume, A. and Heidhues, P. (2004). All equilibria of the Vickrey auction. *Journal of Economic Theory*, 114:170–177.
- Blume, L. E. and Zame, W. R. (1994). The algebraic geometry of perfect and sequential equilibrium. *Econometrica*, 62:783–794.
- Börgers, T. (1994). Weak dominance and approximate common knowledge. *Journal of Economic Theory*, 64:265–276.

- Brandenburger, A. (1992). Lexicographic probabilities and iterated admissibility. In Dasgupta, P. et al., editors, *Economic Analysis of Markets and Games*, pages 282–290. MIT Press, Cambridge.
- Brandenburger, A. and Dekel, E. (1987). Rationalizability and correlated equilibria. *Econometrica*, 55:1391–1402.
- Brandenburger, A. and Dekel, E. (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59:189–198.
- Brandenburger, A., Friedenberg, A., and Keisler, H. J. (2008). Admissibility in games. *Econometrica*, 76:307–352.
- Brouwer, L. E. J. (1912). Über Abbildung von Mannigfaltigkeiten. *Mathematische Annalen*, 71:97–115.
- Brown, G. W. (1951). Iterative solutions of games by fictitious play. In Koopmans, T. C. et al., editors, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, New York.
- Brown, G. W. and von Neumann, J. (1950). Solutions of games by differential equations. In Kuhn, H. W. and Tucker, A. W., editors, *Contributions to the Theory of Games I*, volume 24 of *Annals of Mathematics Studies*, pages 73–79. Princeton University Press, Princeton.
- Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 97:1060–1090.
- Carlsson, H. and van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61:989–1018.
- Chen, X., Deng, X., and Teng, S.-H. (2009). Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 36: article 14.
- Chen, Y.-C., Di Tillio, A., Faingold, E., and Xiong, S. (2010). Uniform topologies on types. *Theoretical Economics*, 5:445–478.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102:179–221.
- Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11:19–33.
- Coste, M. (2002). An introduction to semialgebraic geometry. Lecture notes, Université de Rennes I.
- Crémer, J. and McLean, R. P. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53:345–361.

- Crémer, J. and McLean, R. P. (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56:1247–1257.
- Dasgupta, P., Hammond, P., and Maskin, E. (1979). The implementation of social choice rules: some results on incentive compatibility. *Review of Economic Studies*, 46:185–216.
- Dasgupta, P. and Maskin, E. (1986a). The existence of equilibrium in discontinuous economic games, I: Theory. *Review of Economic Studies*, 53:1–26.
- Dasgupta, P. and Maskin, E. (1986b). The existence of equilibrium in discontinuous economic games, II: Applications. *Review of Economic Studies*, 53:27–41.
- Dasgupta, P. and Maskin, E. (2000). Efficient auctions. *Quarterly Journal of Economics*, 115:341–388.
- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39:195–259.
- d’Aspremont, C. and Gérard-Varet, L. A. (1979). Incentives and incomplete information. *Journal of Public Economics*, 11:25–45.
- Dekel, E. and Fudenberg, D. (1990). Rational behavior with payoff uncertainty. *Journal of Economic Theory*, 52:243–267.
- Dekel, E., Fudenberg, D., and Levine, D. K. (1999). Payoff information and self-confirming equilibrium. *Journal of Economic Theory*, 89:165–185.
- Dekel, E., Fudenberg, D., and Levine, D. K. (2002). Subjective uncertainty over behavior strategies: A correction. *Journal of Economic Theory*, 104:473–478.
- Dekel, E., Fudenberg, D., and Morris, S. (2006). Topologies on types. *Theoretical Economics*, 1:275–309. Correction by Y.-C. Chen and S. Xiong, 3 (2008), 283–285.
- Dekel, E., Fudenberg, D., and Morris, S. (2007). Interim correlated rationalizability. *Theoretical Economics*, 2:15–40.
- Dekel, E. and Gul, F. (1997). Rationality and knowledge in game theory. In Kreps, D. M. and Wallis, K. F., editors, *Advances in Economics and Econometrics: Theory and Applications; Seventh World Congress*, volume 1. Cambridge University Press.
- Dufwenberg, M. and Stegeman, M. (2002). Existence and uniqueness of maximal reductions under iterated strict dominance. *Econometrica*, 70:2007–2023.
- Dutta, P. K. (1995). A folk theorem for stochastic games. *Journal of Economic Theory*, 66:1–32.
- Elmes, S. and Reny, P. J. (1994). On the strategic equivalence of extensive form games. *Journal of Economic Theory*, 62:1–23.

- Ewerhart, C. (2000). Chess-like games are dominance solvable in at most two steps. *Games and Economic Behavior*, 33:41–47.
- Ewerhart, C. (2002). Backward induction and the game-theoretic analysis of chess. *Games and Economic Behavior*, 39:206–214.
- Fan, K. (1952). Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences*, 38:121–126.
- Farrell, J. and Rabin, M. (1996). Cheap talk. *Journal of Economic Perspectives*, 10:103–118.
- Flood, M. M. (1952). Some experimental games. Report RM-789-1, The Rand Corporation.
- Folland, G. B. (1999). *Real Analysis: Modern Techniques and Their Applications*. Wiley, New York, second edition.
- Foster, D. P. and Vohra, R. (1997). Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21:40–55.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38:1–12.
- Friedman, J. W. (1985). Cooperative equilibria in finite horizon non-cooperative supergames. *Journal of Economic Theory*, 52:390–398.
- Fudenberg, D. and Levine, D. K. (1993). Self-confirming equilibrium. *Econometrica*, 61:523–545.
- Fudenberg, D., Levine, D. K., and Maskin, E. (1994). The folk theorem with imperfect public information. *Econometrica*, 62:997–1039.
- Fudenberg, D. and Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54:533–554.
- Fudenberg, D. and Maskin, E. (1991). On the dispensability of public randomization in discounted repeated games. *Journal of Economic Theory*, 53:428–438.
- Fudenberg, D. and Tirole, J. (1991a). *Game Theory*. MIT Press, Cambridge.
- Fudenberg, D. and Tirole, J. (1991b). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53:236–260.
- Gershkov, A., Goeree, J. K., Kushnir, A., Moldovanu, B., and Shi, X. (2013). On the equivalence of Bayesian and dominant strategy implementation. *Econometrica*, 81:197–220.
- Gibbard, A. (1973). Manipulation of voting schemes: a general result. *Econometrica*, 41:587–601.

- Gilboa, I. (2009). *Theory of Decision under Uncertainty*. Cambridge University Press, Cambridge.
- Gilboa, I. (2010). *Rational Choice*. MIT Press, Cambridge.
- Glicksberg, I. L. (1952). A further generalization of the Kakutani Fixed Point Theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3:170–174.
- González-Díaz, J., García-Jurado, I., and Fiestras-Janeiro, M. G. (2010). *An Introductory Course on Mathematical Game Theory*. American Mathematical Society, Providence, RI.
- Govindan, S. and Wilson, R. (2006). Sufficient conditions for stable equilibria. *Theoretical Economics*, 1:167–206.
- Govindan, S. and Wilson, R. (2009). On forward induction. *Econometrica*, 77:1–28.
- Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 52:87–100.
- Green, J. R. and Laffont, J.-J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438.
- Grossman, S. J. and Hart, O. (1983). An analysis of the principal-agent problem. *Econometrica*, 51:7–45.
- Groves, T. (1973). Incentives in teams. *Econometrica*, 41:617–631.
- Gul, F. (1998). A comment on Aumann’s Bayesian view. *Econometrica*, 66:923–928.
- Halpern, J. Y. (2001). Substantive rationality and backward induction. *Games and Economic Behavior*, 37:425–435.
- Harsanyi, J. C. (1967, 1968). Games with incomplete information played by Bayesian players. *Management Science*, 14:159–182, 320–334, 486–502.
- Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2:1–23.
- Hart, S. (2005). Adaptive heuristics. *Econometrica*, 73:1401–1430.
- Hart, S. and Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150.
- Hart, S. and Mas-Colell, A. (2003). Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93:1830–1836.
- Hart, S. and Schmeidler, D. (1989). Existence of correlated equilibria. *Mathematics of Operations Research*, 14:18–25.



- Hendon, E., Jacobsen, H. J., and Sloth, B. (1996). The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior*, 12:274–282.
- Hillas, J. (1998). How much of forward induction is implied by backward induction and ordinality? Unpublished manuscript, University of Auckland.
- Hillas, J. and Kohlberg, E. (2002). Conceptual foundations of strategic equilibrium. In Aumann, R. J. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 42, pages 1597–1663. Elsevier, New York.
- Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). *Fundamentals of Convex Analysis*. Springer, Berlin.
- Hofbauer, J. and Sandholm, W. H. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6:341–377.
- Hofbauer, J. and Swinkels, J. M. (1996). A universal Shapley example. Unpublished manuscript, University of Vienna and Northwestern University.
- Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10:74–91.
- Hopkins, E. and Seymour, R. M. (2002). The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43:1157–1190.
- Howard, R. A. (1960). *Dynamic Programming and Markov Processes*. MIT Press, Cambridge.
- Hurwicz, L. (1960). Optimality and informational efficiency in resource allocation processes. In Arrow, K. J., Karlin, S., and Suppes, P., editors, *Mathematical Methods in the Social Sciences*. Stanford University Press.
- Hurwicz, L. (1972). On informationally decentralized systems. In McGuire, C. and Radner, R., editors, *Decision and Organization*. North Holland, Amsterdam.
- Jackson, M. O. (1991). Bayesian implementation. *Econometrica*, 59:461–477.
- Jackson, M. O., Simon, L. K., Swinkels, J. M., and Zame, W. R. (2002). Communication and equilibrium in discontinuous games of incomplete information. *Econometrica*, 70:1711–1740.
- Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123:81–104.
- Jehiel, P., Meyer-ter-Vehn, M., Moldovanu, B., and Zame, W. R. (2006). The limits of ex post implementation. *Econometrica*, 74:585–610.
- Jehiel, P. and Moldovanu, B. (2001). Efficient design with interdependent valuations. *Econometrica*, 69:1237–1259.

- Jehle, G. A. and Reny, P. J. (2011). *Advanced Microeconomic Theory*. Financial Times/Prentice Hall/Pearson, Harlow, England, third edition.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–291.
- Kakutani, S. (1941). A generalization of Brouwer’s fixed point theorem. *Duke Mathematical Journal*, 8:457–459.
- Kohlberg, E. and Mertens, J.-F. (1986). On the strategic stability of equilibria. *Econometrica*, 54:1003–1037.
- Kohlberg, E. and Reny, P. J. (1997). Independence on relative probability spaces and consistent assessments in game trees. *Journal of Economic Theory*, 75:280–313.
- Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational cooperation in the finitely repeated Prisoner’s Dilemma. *Journal of Economic Theory*, 27:245–252.
- Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50:863–894.
- Krishna, V. (2002). *Auction Theory*. Academic Press, San Diego.
- Krishna, V. and Perry, M. (2000). Efficient mechanism design. Unpublished manuscript, Penn State University and Hebrew University.
- Kuhn, H. W. (1953). Extensive games and the problem of information. In Kuhn, H. W. and Tucker, A. W., editors, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 193–216. Princeton University Press, Princeton.
- Laffont, J.-J. and Maskin, E. (1979). A differential approach to expected utility maximizing mechanisms. In Laffont, J.-J., editor, *Aggregation and Revelation of Preferences*, pages 289–308. North Holland.
- Lahkar, R. (2011). The dynamic instability of dispersed price equilibria. *Journal of Economic Theory*, 146:1796–1827.
- Lemke, C. E. and Howson, J. T. (1964). Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12:413–423.
- Luce, R. D. and Raiffa, H. (1957). *Games and Decisions: Introduction and Critical Survey*. Wiley, New York.
- Mailath, G. J. and Samuelson, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, Oxford.
- Mailath, G. J., Samuelson, L., and Swinkels, J. M. (1993). Extensive form reasoning in normal form games. *Econometrica*, 61:273–302.

- Mailath, G. J., Samuelson, L., and Swinkels, J. M. (1997). How proper is sequential equilibrium? *Games and Economic Behavior*, 18:193–218.
- Manelli, A. M. and Vincent, D. R. (2010). Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica*, 78:1905–1938.
- Marx, L. M. and Swinkels, J. M. (1997). Order independence for iterated weak dominance. *Games and Economic Behavior*, 18:219–245. Corrigendum, 31 (2000), 324–329.
- Maskin, E. (1992). Auctions and privatization. In Siebert, H., editor, *Privatization*. Institut für Weltwirtschaft an der Universität Kiel.
- Maskin, E. (1999). Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66:23–38. Working paper circulated in 1977.
- Maskin, E. and Riley, J. G. (1984). Monopoly with incomplete information. *RAND Journal of Economics*, 15:171–196.
- Maskin, E. and Tirole, J. (2001). Markov perfect equilibrium I: Observable actions. *Journal of Economic Theory*, 100:191–219.
- Matthews, S. A. (1995). A technical primer on auction theory I: Independent private values. Unpublished manuscript, Northwestern University.
- McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the Centipede game. *Econometrica*, 60:803–836.
- Mertens, J.-F. (1989). Stable equilibria—a reformulation. I. Definition and basic properties. *Mathematics of Operations Research*, 14:575–625.
- Mertens, J.-F. (1991). Stable equilibria—a reformulation. II. Discussion of the definition, and further results. *Mathematics of Operations Research*, 16:694–753.
- Mertens, J.-F. (1995). Two examples of strategic equilibrium. *Games and Economic Behavior*, 8:378–388.
- Mertens, J.-F. and Zamir, S. (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.
- Milgrom, P. and Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, 70:583–601.
- Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62:157–180.
- Milgrom, P. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122.

- Milgrom, P. R. and Stokey, N. L. (1982). Information, trade, and common knowledge. *Journal of Economic Theory*, 26:17–27.
- Milnor, J. (1954). Games against Nature. In Thrall, R. M., Coombs, C. H., and Davis, R. L., editors, *Decision Processes*, pages 49–59. Wiley, New York.
- Mirrlees, J. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38:175–208.
- Mirrlees, J. (1999). The theory of moral hazard and unobservable behavior, part I. *Review of Economic Studies*, 66:3–21. Working paper circulated in 1975.
- Monderer, D. and Samet, D. (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1:170–190.
- Mookherjee, D. and Reichelstein, S. (1992). Dominant strategy implementation of Bayesian incentive compatible allocation rules. *Journal of Economic Theory*, 56:378–399.
- Moore, J. (1992). Implementation, contracts, and renegotiation in environments with complete information. In Laffont, J.-J., editor, *Advances in Economic Theory: Sixth World Congress*, volume I. Cambridge University Press, Cambridge.
- Morris, S. (1994). Trade with heterogeneous prior beliefs and asymmetric information. *Econometrica*, 62:1327–1347.
- Morris, S. (1995). The common prior assumption in economic theory. *Economics and Philosophy*, 11:227–253.
- Morris, S. and Shin, H. S. (2003). Global games: Theory and applications. In Dewatripont, M., Hansen, L. P., and Turnovsky, S. J., editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, volume 1, pages 56–114. Cambridge University Press, Cambridge.
- Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18:301–317.
- Myerson, R. B. (1978). Refinements of the Nash equilibrium concept. *International Journal of Game Theory*, 7:73–80.
- Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica*, 47:61–73.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6:58–73.
- Myerson, R. B. (1982). Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, 10:67–81.
- Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge.

- Myerson, R. B. and Reny, P. J. (2015). Sequential equilibria of multi-stage games with infinite sets of types and actions. Unpublished manuscript, University of Chicago.
- Myerson, R. B. and Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281.
- Nash, J. F. (1950). *Non-Cooperative Games*. PhD thesis, Princeton University. Reprinted in *The Essential John Nash*, H. W. Kuhn and S. Nasar, Eds., p. 51-83, Princeton University Press, Princeton (2002).
- Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics*, 54:287–295.
- Neeman, Z. (2004). The relevance of private information in mechanism design. *Journal of Economic Theory*, 117:55–77.
- Ollár, M. and Penta, A. (2015). Full implementation and belief restrictions. *American Economic Review*, 108:2243–2277.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press, Cambridge.
- Østerdal, L. P. (2005). Iterated weak dominance and subgame dominance. *Journal of Mathematical Economics*, 41:637–645.
- Oury, M. and Tercieux, O. (2012). Continuous implementation. *Econometrica*, 80:1605–1637.
- Palfrey, T. R. (1992). Implementation in Bayesian equilibrium: The multiple equilibrium problem in mechanism design. In Laffont, J.-J., editor, *Advances in Economic Theory: Sixth World Congress*, volume I. Cambridge University Press, Cambridge.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52:1029–1050.
- Penta, A. (2013). On the structure of rationalizability for arbitrary spaces of uncertainty. *Theoretical Economics*, 8:405–430.
- Perea, A. (2011). An algorithm for proper rationalizability. *Games and Economic Behavior*, 72:510–525.
- Perry, M. and Reny, P. J. (2002). An efficient auction. *Econometrica*, 70:1199–1212.
- Piccione, M. and Rubinstein, A. (1997). The absent-minded driver’s paradox: Synthesis and responses. *Games and Economic Behavior*, 20:121–130.
- Postlewaite, A. and Schmeidler, D. (1986). Implementation in differential information economies. *Journal of Economic Theory*, 39:14–33.
- Reny, P. J. (1992). Backward induction, normal form perfection, and explicable equilibria. *Econometrica*, 60:627–649.

- Reny, P. J. (1999). On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67:1029–1056.
- Reny, P. J. (2016). Nash equilibrium in discontinuous games. *Economic Theory*, 61:553–569.
- Riley, J. G. (1979). Informational equilibrium. *Econometrica*, 47:331–359.
- Riley, J. G. and Samuelson, W. (1981). Optimal auctions. *American Economic Review*, 71:381–392.
- Ritzberger, K. (2002). *Foundations of Non-Cooperative Game Theory*. Oxford University Press, Oxford.
- Robert, J. (1991). Continuity in auction design. *Journal of Economic Theory*, 55:169–179.
- Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press, Princeton.
- Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25:92–100.
- Rothschild, M. and Stiglitz, J. E. (1976). Equilibrium in competitive insurance markets: An essay in the economics of imperfect information. *Quarterly Journal of Economics*, 80:629–649.
- Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109.
- Rubinstein, A. (1989). The electronic mail game: Strategic behavior under “almost common knowledge”. *American Economic Review*, 79:385–391.
- Rubinstein, A. (1991). Comments on the interpretation of game theory. *Econometrica*, 59:909–924.
- Salanié, B. (2005). *The Economics of Contracts*. MIT Press, second edition.
- Samuelson, L. (1992). Dominated strategies and common knowledge. *Games and Economic Behavior*, 4:284–313.
- Samuelson, L. (2004). Modeling knowledge in economic analysis. *Journal of Economic Literature*, 42:367–403.
- Sandholm, W. H. (2009). Evolutionary game theory. In Meyers, R. A., editor, *Encyclopedia of Complexity and Systems Science*, pages 3176–3205. Springer, Heidelberg.
- Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.
- Sandholm, W. H., Izquierdo, S. S., and Izquierdo, L. R. (2016). Best experienced payoff dynamics and cooperation in the Centipede game. Unpublished manuscript, University of Wisconsin, Universidad de Valladolid, and Universidad de Burgos.

- Sannikov, Y. (2007). Games with imperfectly observable actions in continuous time. *Econometrica*, 75:1285–1329.
- Satterthwaite, M. A. (1975). Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217.
- Savage, L. J. (1954). *The Foundations of Statistics*. Wiley, New York.
- Schaeffer, J. et al. (2007). Checkers is solved. *Science*, 317:1518–1522.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge.
- Schmidt, K. M. (1993). Reputation and equilibrium characterization in repeated games of conflicting interests. *Econometrica*, 61:325–351.
- Schuhmacher, F. (1999). Proper rationalizability and backward induction. *International Journal of Game Theory*, 28:599–615.
- Sekiguchi, T. (1997). Efficiency in repeated prisoner's dilemma with private monitoring. *Journal of Economic Theory*, 76:345–361.
- Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrage-  
tragheit. *Zeitschrift für die Gesamte Staatswissenschaft*, 121:301–324 and 667–689.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in  
extensive games. *International Journal of Game Theory*, 4:25–55.
- Shaked, A. and Sutton, J. (1984). Involuntary unemployment as a perfect equilibrium in  
a bargaining model. *Econometrica*, 52:1351–1364.
- Shapley, L. S. (1964). Some topics in two person games. In Dresher, M., Shapley, L. S.,  
and Tucker, A. W., editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics  
Studies*, pages 1–28. Princeton University Press, Princeton.
- Shapley, L. S. and Snow, R. (1950). Basic solutions of discrete games. In Kuhn, H. W.  
and Tucker, A. W., editors, *Contributions to the Theory of Games I*, volume 24 of *Annals of  
Mathematics Studies*, pages 27–35. Princeton University Press, Princeton.
- Simon, L. K. and Zame, W. R. (1990). Discontinuous games and endogenous sharing rules.  
*Econometrica*, 58:861–872.
- Sion, M. (1958). On general minimax theorems. *Pacific Journal of Mathematics*, 8:171–176.
- Sorin, S. (1986). On repeated games with complete information. *Mathematics of Operations  
Research*, 11:147–160.
- Spence, A. M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87:355–374.

- Ståhl, I. (1972). *Bargaining Theory*. Stockholm School of Economics, Stockholm.
- Stalnaker, R. (1998). Belief revision in games: forward and backward induction. *Mathematical Social Sciences*, 36:31–56.
- Stokey, N. L. and Lucas, R. E. (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge.
- Stoye, J. (2011). Statistical decisions under ambiguity. *Theory and Decision*, 70:129–148.
- Thompson, F. (1952). Equivalence of games in extensive form. RM 759, The Rand Corporation. Reprinted in *Classics in Game Theory*, H. W. Kuhn, Ed., p. 36-45, Princeton University Press, Princeton (1997).
- Toikka, J. (2011). Ironing without control. *Journal of Economic Theory*, 146:2510–2526.
- Tranæs, T. (1998). Tie-breaking in games of perfect information. *Games and Economic Behavior*, 22:148–161.
- van Damme, E. (1984). A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. *International Journal of Game Theory*, 13:1–13.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37.
- Vickrey, W. (1962). Auctions and bidding games. In *Recent Advances in Game Theory*. Princeton University Press, Princeton.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Prentice-Hall, Princeton.
- Weinstein, J. and Yildiz, M. (2007). A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75:365–400.
- Williams, S. R. (1999). A characterization of efficient, Bayesian incentive compatible mechanisms. *Economic Theory*, 14:155–180.
- Wilson, R. (1971). Computing equilibria of  $n$ -person games. *SIAM Journal on Applied Mathematics*, 21:80–87.
- Wilson, R. (1985). Incentive efficiency of double auctions. *Econometrica*, 53:1101–1115.
- Wilson, R. (1987). Game theoretic analysis of trading processes. In Bewley, T. F., editor, *Advances in Economic Theory: Fifth World Congress*. Cambridge University Press.
- Wilson, R. (1993). *Nonlinear Pricing*. Oxford University Press.
- Yamamoto, Y. (2010). The use of public randomization in discounted repeated games. *International Journal of Game Theory*, 39:431–443.



Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford University Press, Oxford.

Zeeman, E. C. (1980). Population dynamics from game theory. In Nitecki, Z. and Robinson, C., editors, *Global Theory of Dynamical Systems (Evanston, 1979)*, number 819 in *Lecture Notes in Mathematics*, pages 472–497, Berlin. Springer.