

Economics 703: Microeconomics II

Modelling Strategic Behavior¹

George J. Mailath
Department of Economics
University of Pennsylvania

December 7, 2011

¹Copyright December 7, 2011 by George J. Mailath.

Contents

Contents	iii
1 Normal and Extensive Form Games	1
1.1 Normal Form Games	1
1.2 Iterated Deletion of Dominated Strategies	8
1.3 Extensive Form Games	10
1.3.1 The Reduced Normal Form	13
1.4 Problems	14
2 A First Look at Equilibrium	17
2.1 Nash Equilibrium	17
2.1.1 Why Study Nash Equilibrium?	21
2.2 Credible Threats and Backward Induction	22
2.3 Subgame Perfection	25
2.4 Mixing	29
2.4.1 Mixed Strategies and Security Levels	29
2.4.2 Domination and Optimality	31
2.4.3 Equilibrium in Mixed Strategies	34
2.4.4 Behavior Strategies	35
2.5 Dealing with Multiplicity	38
2.5.1 I: Refinements	38
2.5.2 II: Selection	39
2.6 Problems	40
3 Games with Nature	47
3.1 An Introductory Example	47
3.2 Purification	48
3.3 Auctions and Related Games	51

3.4	Games of Incomplete Information	64
3.5	Higher Order Beliefs and Global Games	66
3.6	Problems	71
4	Nash Equilibrium	75
4.1	Existence of Nash Equilibria	75
4.2	Foundations for Nash Equilibrium	81
4.2.1	Boundedly Rational Learning	82
4.2.2	Social Learning (Evolution)	83
4.2.3	Individual learning	89
4.3	Problems	92
5	Dynamic Games	97
5.1	Sequential Rationality	97
5.2	Perfect Bayesian Equilibrium	102
5.3	Sequential Equilibrium	106
5.4	Problems	109
6	Signaling	113
6.1	General Theory	113
6.2	Job Market Signaling	116
6.2.1	Full Information	117
6.2.2	Incomplete Information	117
6.2.3	Refining to Separation	119
6.2.4	Continuum of Types	122
6.3	Problems	123
7	Repeated Games	129
7.1	Basic Structure	129
7.2	Modeling Competitive Agents	139
7.3	The Folk Theorem	144
7.4	Efficiency Wages	149
7.5	Imperfect Public Monitoring	151
7.6	Problems	155
8	Topics in Dynamic Games	165
8.1	Dynamic Games and Markov Perfect Equilibria	165
8.2	Coase Conjecture	170

8.2.1	One and Two Period Example	171
8.2.2	Infinite Horizon	172
8.3	Reputations	174
8.3.1	Infinite Horizon	176
8.4	Problems	180
9	Bargaining	185
9.1	Axiomatic Nash Bargaining	185
9.1.1	The Axioms	185
9.1.2	Nash's Theorem	186
9.2	Rubinstein Bargaining	189
9.2.1	The Stationary Equilibrium	189
9.2.2	All Equilibria	190
9.2.3	Impatience	191
9.3	Outside Options	192
9.4	Exogenous Risk of Breakdown	195
9.5	Problems	197
10	Appendices	199
10.1	Section 2.5.1: Trembling Hand Perfection	199
10.1.1	Existence and Characterization	200
10.2	Problems	203
	References	205

Chapter 1

Normal and Extensive Form Games¹

1.1 Normal Form Games

Example 1.1.1 (Prisoner's Dilemma)

		II	
		Confess	Don't confess
I	Confess	$-6, -6$	$0, -9$
	Don't confess	$-9, 0$	$-1, -1$

Often interpreted as a partnership game: effort E produces an output of 6 at a cost of 4, with output shared equally, and S denotes shirking.

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

★

¹Copyright December 7, 2011 by George J. Mailath

Definition 1.1.1 An n -player normal (or strategic) form game G is an n -tuple

$\{(S_1, u_1), \dots, (S_n, u_n)\}$, where for each i ,

- S_i is a nonempty set, called i 's strategy space, and
- $u_i : \prod_{k=1}^n S_k \rightarrow \mathbb{R}$ is called i 's payoff function.

Equivalently, a normal form game is simply a vector-valued function $u : \prod_{i=1}^n S_i \rightarrow \mathbb{R}^n$.

Notation: $S \equiv \prod_{k=1}^n S_k$,

$s \equiv (s_1, \dots, s_n) \in S$,

$s_{-i} \equiv (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in S_{-i} \equiv \prod_{k \neq i} S_k$.

$(s'_i, s_{-i}) \equiv (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n) \in S$.

Example 1.1.2 (Sealed bid second price auction) 2 bidders, $b_i = i$'s bid, $v_i = i$'s reservation price (willingness to pay).

Then, $n = 2$, $S_i = \mathbb{R}_+$, and

$$u_i(b_1, b_2) = \begin{cases} v_i - b_j, & \text{if } b_i > b_j, \\ \frac{1}{2}(v_i - b_j), & \text{if } b_i = b_j, \\ 0, & \text{if } b_i < b_j. \end{cases} \quad \star$$

Example 1.1.3 (Sealed bid first price auction) 2 bidders, $b_i = i$'s bid, $v_i = i$'s reservation price (willingness to pay).

Then, $n = 2$, $S_i = \mathbb{R}_+$, and

$$u_i(b_1, b_2) = \begin{cases} v_i - b_i, & \text{if } b_i > b_j, \\ \frac{1}{2}(v_i - b_i), & \text{if } b_i = b_j, \\ 0, & \text{if } b_i < b_j. \end{cases} \quad \star$$

Example 1.1.4 (Cournot duopoly) Perfect substitutes, so that market clearing price is given by $P(Q) = \max\{a - Q, 0\}$, $Q = q_1 + q_2$, $C(q_i) = cq_i$, $0 < c < a$, and $n = 2$

Quantity competition: $S_i = \mathbb{R}_+$, and $u_i(q_1, q_2) = (P(q_1 + q_2) - c)q_i$.

★

Example 1.1.5 (Bertrand duopoly) Economic environment is as for example 1.1.4, but price competition. Since perfect substitutes, lowest pricing firm gets the whole market (with the market split in the event of a tie). We again have $S_i = \mathbb{R}_+$, but now

$$u_1(p_1, p_2) = \begin{cases} (p_1 - c) \max\{a - p_1, 0\}, & \text{if } p_1 < p_2, \\ (p_1 - c) \frac{\max\{a - p_1, 0\}}{2}, & \text{if } p_1 = p_2, \\ 0, & \text{if } p_1 > p_2. \end{cases} \quad ★$$

Example 1.1.6 (Voting by veto) Three outcomes: x, y , and z . Player 1 first vetoes an outcome, and then player 2 vetoes one of the remaining outcomes. The non-vetoed outcome results. Suppose 1 ranks outcomes: $x \succ y \succ z$ (i.e., $u_1(x) = 2, u_1(y) = 1, u_1(z) = 0$), and 2 ranks outcomes as: $y \succ x \succ z$ (i.e., $u_2(x) = 1, u_2(y) = 2, u_2(z) = 0$).

1's strategy is an uncontingent veto, so $S_1 = \{x, y, z\}$.

2's strategy is a contingent veto, so $S_2 = \{(abc) : a \in \{y, z\}, b \in \{x, z\}, c \in \{x, y\}\}$.

2

		(yxx)	(yx y)	(yzx)	(yzy)	(zxx)	(zxy)	(zzx)	(zzy)
1	x	0,0	0,0	0,0	0,0	1,2	1,2	1,2	1,2
	y	0,0	0,0	2,1	2,1	0,0	0,0	2,1	2,1
	z	1,2	2,1	1,2	2,1	1,2	2,1	1,2	2,1 ★

Definition 1.1.2 s'_i strictly dominates s''_i if $\forall s_{-i} \in S_{-i}$,

$$u_i(s'_i, s_{-i}) > u_i(s''_i, s_{-i}).$$

s_i is a strictly dominant strategy if s_i strictly dominates every strategy $s''_i \neq s_i, s''_i \in S_i$.

If i has a strictly dominant strategy, then

$$\arg \max_{s_i} u_i(s_i, s_{-i}) = \left\{ s_i : u_i(s_i, s_{-i}) = \max_{s'_i} u_i(s'_i, s_{-i}) \right\}$$

is a singleton and is independent of s_{-i} .

Remark 1.1.1 The definition per se of a normal form game (or strict dominance for that matter) makes no assumption about the knowledge that players have about the game. We will, however, typically assume (at least) that players know the strategy spaces, and their own payoffs as a function of strategy profiles. However, as the large literature on evolutionary game theory in biology suggests (see also Section 4.2), this is not necessary.

The assertion that players will not play a strictly dominant strategy is compelling when players know the strategy spaces and their own payoffs. But (again, see Section 4.2), this is not necessary for the plausibility of the assertion.

◆

Definition 1.1.3 s'_i (weakly) dominates s''_i if $\forall s_{-i} \in S_{-i}$,

$$u_i(s'_i, s_{-i}) \geq u_i(s''_i, s_{-i}),$$

and $\exists s'_{-i} \in S_{-i}$,

$$u_i(s'_i, s'_{-i}) > u_i(s''_i, s'_{-i}).$$

A strategy is said to be *strictly* or *weakly undominated* if it is not strictly or weakly dominated by some other strategy. If the adjective is omitted from dominated (or undominated), weak is typically meant (but not always, unfortunately). A strategy is *weakly dominant* if it weakly dominates every other strategy.

Lemma 1.1.1 *If a weakly dominant strategy exists, it is unique.*

Proof. Suppose s_i is a weakly dominant strategy. Then for all $s'_i \in S_i$, there exists $s_{-i} \in S_{-i}$, such that

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}).$$

But this implies that s'_i cannot weakly dominate s_i , and so s_i is the only weakly dominant strategy. ■

Remark 1.1.2 (Warning) There is also a notion of *dominant* strategy:

Definition 1.1.4 s'_i is a dominant strategy if $\forall s''_i \in S_i, \forall s_{-i} \in S_{-i}$,

$$u_i(s'_i, s_{-i}) \geq u_i(s''_i, s_{-i}).$$

If s'_i is a dominant strategy for i , then $s'_i \in \arg \max_{s_i} u_i(s_i, s_{-i})$, for all s_{-i} ; but $\arg \max_{s_i} u_i(s_i, s_{-i})$ need not be a singleton and it need not be independent of s_{-i} [example?]. If i has only one dominant strategy, then that strategy weakly dominates every other strategy, and so is weakly dominant.

Dominant strategies have played an important role in mechanism design and implementation (see Remark 1.1.3), but not otherwise (since a dominant strategy—when it exists—will typically weakly dominate every other strategy, as in Example 1.1.7). ◆

Remark 1.1.3 (Strategic behavior is ubiquitous) Consider a society consisting of a finite number n of members and a finite set of outcomes X . Suppose each member of society has a strict preference ordering of X , and let Θ be the set of all possible strict orderings on X . A profile $(\theta_1, \dots, \theta_n) \in \Theta^n$ describes a particular society (a preference ordering for each member).

A *social choice rule* or *function* is a mapping $f : \Theta^n \rightarrow X$. For any $\theta_i \in \Theta$, let $t(\theta_i)$ be the top-ranked outcome in X under θ_i . A social choice rule f is *dictatorial* if there is some i such that for all $(\theta_1, \dots, \theta_n) \in \Theta^n$, $f(\theta_1, \dots, \theta_n) = t(\theta_i)$. A social choice rule f is *unanimous* if $f(\theta_1, \dots, \theta_n) = x$ whenever $x = t(\theta_j)$ for all j .

The *direct mechanism* is the normal form game in which all members of society simultaneously announce a preference ordering and the outcome is determined by the social choice rule as a function of the announced preferences.

Theorem 1.1.1 (Gibbard-Satterthwaite) *Suppose $|X| \geq 3$ and f is unanimous. Then, announcing truthfully in the direct mechanism is a dominant strategy for all preference profiles if, and only if, the social choice rule is dictatorial.*

A social choice rule is said to be *strategy proof* if announcing truthfully in the direct mechanism is a dominant strategy for all preference profiles. It is trivial that for any dictatorial social choice rule, it is a dominant strategy to always truthfully report in the direct mechanism. The surprising result is the converse.

♦

Example 1.1.7 (Continuation of example 1.1.2) In the 2nd price auction, each player has a weakly dominant strategy, given by $b_1 = v_1$.

Sufficient to show this for 1. First argue that bidding v_1 is a best response for 1, no matter what bid 2 makes (i.e., it is a dominant strategy). Recall that payoffs are given by

$$u_1(b_1, b_2) = \begin{cases} v_1 - b_2, & \text{if } b_1 > b_2, \\ \frac{1}{2}(v_1 - b_2), & \text{if } b_2 = b_1, \\ 0, & \text{if } b_1 < b_2. \end{cases}$$

Two cases:

1. $b_2 < v_1$: Then $u_1(v_1, b_2) = v_1 - b_2 \geq u_1(b_1, b_2)$.
2. $b_2 \geq v_1$: Then, $u_1(v_1, b_2) = 0 \geq u_1(b_1, b_2)$.

Thus, bidding v_1 is optimal.

Bidding v_1 also weakly dominates every other bid (and so v_1 is weakly dominant). Suppose $b_1 < v_1$ and $b_1 < b_2 < v_1$. Then $u_1(b_1, b_2) = 0 < v_1 - b_2 = u_1(v_1, b_2)$. If $b_1 > v_1$ and $b_1 > b_2 > v_1$, then $u_1(b_1, b_2) = v_1 - b_2 < 0 = u_1(v_1, b_2)$. ★

Example 1.1.8 (Provision of public goods) n people. Agent i values public good at r_i , total cost of public good is C .

Suppose costs are shared uniformly and utility is linear, so agent i 's net utility is $v_i \equiv r_i - \frac{1}{n}C$.

Efficient provision: Public good provided iff $0 \leq \sum v_i$, i.e., $C \leq \sum r_i$.

Eliciting preferences: Agents announce \hat{v}_i and provide if $\sum \hat{v}_i \geq 0$? Gives incentive to overstate if $v_i > 0$ and understate if $v_i < 0$.

Groves-Clarke mechanism: if public good provided, pay agent i amount $\sum_{j \neq i} \hat{v}_j$ (tax if negative) .

$$\text{Agent } i\text{'s payoff}(\hat{v}_i) = \begin{cases} v_i + \sum_{j \neq i} \hat{v}_j, & \text{if } \hat{v}_i + \sum_{j \neq i} \hat{v}_j \geq 0, \\ 0, & \text{if } \hat{v}_i + \sum_{j \neq i} \hat{v}_j < 0. \end{cases}$$

Dominant strategy to announce $\hat{v}_i = v_i$: If $v_i + \sum_{j \neq i} \hat{v}_j > 0$, announcing $\hat{v}_i = v_i$ ensures good is provided, while if $v_i + \sum_{j \neq i} \hat{v}_j < 0$, announcing $\hat{v}_i = v_i$ ensures good is not provided. Moreover, conditional on provision, announcement does not affect payoff—note similarity to second price auction.

No payments if no provision, but payments large if provision: Total payments to agents when provision = $\sum_i (\sum_{j \neq i} \hat{v}_j) = (n - 1) \sum_i \hat{v}_i$. Taxing agent i by an amount independent of i 's behavior has no impact, so tax i the amount $\max\{\sum_{j \neq i} \hat{v}_j, 0\}$. Result is

$$\text{payoff to } i = \begin{cases} v_i, & \text{if } \sum_j \hat{v}_j \geq 0 \text{ and } \sum_{j \neq i} \hat{v}_j \geq 0, \\ v_i + \sum_{j \neq i} \hat{v}_j, & \text{if } \sum_j \hat{v}_j \geq 0 \text{ and } \sum_{j \neq i} \hat{v}_j < 0, \\ -\sum_{j \neq i} \hat{v}_j, & \text{if } \sum_j \hat{v}_j < 0 \text{ and } \sum_{j \neq i} \hat{v}_j \geq 0, \\ 0, & \text{if } \sum_j \hat{v}_j < 0 \text{ and } \sum_{j \neq i} \hat{v}_j < 0. \end{cases}$$

This is the *pivotal mechanism*. Note that i only pays a tax if i changes social decision. Moreover, total taxes are no larger than $\sum_i \max\{\hat{v}_i, 0\}$ if the good is provided and no larger than $\sum_i \max\{-\hat{v}_i, 0\}$ if the good is not provided. ★

Example 1.1.9 (Continuation of example 1.1.4, Cournot) There are no weakly dominating quantities in the Cournot duopoly: Suppose $q_2 < a$. Then $\arg \max_{q_1} u_1(q_1, q_2) = \arg \max (a - c - q_1 - q_2)q_1$. First

order condition implies $a - c - 2q_1 - q_2 = 0$ or

$$q_1(q_2) = \frac{a - c - q_2}{2}.$$

Since $\arg \max u_1$ is unique and a nontrivial function of q_2 , there is no weakly dominating quantity. ★

1.2 Iterated Deletion of Dominated Strategies

Example 1.2.1

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

Delete R and then B , and then L to get (T, M) . ★

Example 1.2.2 (Continuation of example 1.1.6) Apply iterative deletion of weakly dominated strategies to veto game. After one round of deletions,

		2
		(z, z, x)
1	y	2, 1
	z	1, 2

and so 1 vetoes y , and not z ! ★

Remark 1.2.1 For finite games, the order of removal of strictly dominated strategies is irrelevant (see Problem 1.4.4(a)). This is not true for weakly dominated strategies:

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

Both TL and BL can be obtained as the *singleton* profile that remains from the iterative deletion of weakly dominated strategies. In addition, $\{TL, TM\}$ results from a different sequence, and $\{BL, BR\}$ from yet another sequence.

Similarly, the order of elimination may matter for infinite games (see Problem 1.4.4(b)).

Because of this, the procedure of the *iterative deletion of weakly (or strictly) dominated strategies* is often understood to *require* that at each stage, all weakly (or strictly) dominated strategies be deleted. We will follow that understanding in this class (unless explicitly stated otherwise). With that understanding, the iterated deletion of weakly dominated strategies in this example leads to $\{TL, BL\}$.



Remark 1.2.2 The plausibility of the iterated deletion of dominated strategies requires something like players knowing the structure of the game (including that other players know the game), not just their own payoffs. For example in Example 1.2.1, in order for the column player to delete L at the third round, then the column player needs to know that the row player will not play B , which requires the column player to know that the row player knows that the column player will not play R .

As illustrated in Section 4.2, this kind of iterated knowledge is not *necessary* for the plausibility of the procedure (though in many contexts it provides the most plausible foundation).



1.3 Extensive Form Games

Game trees look like decision trees. Role of definition is to make clear who does what, when, knowing what.

Definition 1.3.1 A finite extensive form game consists of:

1. A set of players $\{1, \dots, n\}$ and nature, denoted player 0.
2. A game tree $(T, <)$, where $(T, <)$ is an arborescence: T is a finite set of nodes and $<$ is a binary relation on T denoting “precedence” satisfying
 - (a) $<$ is asymmetric ($t < t' \Rightarrow t' \not< t$),²
 - (b) transitive ($\forall t, t', t'' \in T, t < t', t' < t'' \Rightarrow t < t''$),³
 - (c) if $t < t''$ and $t' < t''$ then either $t < t'$ or $t' < t$, and finally,
 - (d) there is a unique initial node, $t_0 \in T$, i.e., $\{t_0\} = \{t \in T : \nexists t' \in T, t' < t\}$.

Let $Z \equiv \{t \in T : \nexists t' \in T, t < t'\}$, Z is the set of terminal nodes.

NOTE: (a)–(d) implies that $\forall t \neq t_0, \exists!$ path from the initial node to t (see Problem 1.4.5).

3. Assignment of players to nodes, $\iota : T \setminus Z \rightarrow \{0, 1, \dots, n\}$. Define $T_j \equiv \iota^{-1}(j) = \{t \in T \setminus Z : \iota(t) = j\}$, $\forall j \in \{0, 1, \dots, n\}$.
4. Actions: Define $s(t) \equiv \{t' \in T : t < t' \text{ and } \nexists t'', t < t'' < t'\}$, the set of immediate successors of t . Actions lead to (label) immediate successors, i.e., there is a set A and a mapping

$$\alpha : T \setminus \{t_0\} \rightarrow A,$$

such that $\alpha(t') \neq \alpha(t'')$ for all $t', t'' \in s(t)$. Define $A(t) \equiv \alpha(s(t))$, the set of actions available at $t \in T \setminus Z$.

5. Information sets: H_i is a partition of T_i for all $i \neq 0$ (H_i is a collection of subsets of T_i such that (i) $\forall t \in T_i, \exists h \in H_i, t \in h$, and (ii) $\forall h, h' \in H_i, h \neq h' \Rightarrow h \cap h' = \emptyset$). Assume $\forall t, t' \in h$,

²Note that this implies that $<$ is irreflexive: $t \not< t$ for all $t \in T$.

³A binary relation satisfying 2(a) and 2(b) is called a strict partial order.

- (a) $t \not\prec t', t' \not\prec t$,
- (b) $A(t) = A(t') \equiv A(h)$, and
- (c) *perfect recall* (every player knows whatever he knew previously, including own previous actions).

6. Payoffs, $u_i : Z \rightarrow \mathbb{R}$.

7. Prob dsn for nature, $\rho : T_0 \rightarrow \cup_{t \in T_0} \Delta(A(t))$ such that $\rho(t) \in \Delta(A(t))$.

Definition 1.3.2 A strategy for player i is a function

$$s_i : H_i \rightarrow \cup_h A(h) \text{ such that } s_i(h) \in A(h), \forall h \in H_i. \quad (1.3.1)$$

The set of player i 's strategies is i 's strategy space, denoted S_i .
Strategy profile is (s_1, \dots, s_n) .

Definition 1.3.3 Suppose there are no moves of nature.

The outcome path is the sequence of nodes reached by strategy profile, or equivalently, the sequence of specified actions.

The outcome is the unique terminal node reached by the strategy profile s , denoted $z(s)$. In this case, the normal form representation is given by $\{(S_1, U_1), \dots, (S_n, U_n)\}$, where S_i is the set of i 's extensive form strategies, and


$$U_i(s) = u_i(z(s)).$$

Definition 1.3.4 If there are moves of nature, the outcome is the implied probability distribution over terminal nodes, denoted $\pi^s \in \Delta(Z)$. In this case, in the normal form representation given by $\{(S_1, U_1), \dots, (S_n, U_n)\}$, where S_i is the set of i 's extensive form strategies, we have

$$U_i(s) = \sum_z \pi^s(z) u_i(z).$$

For an example, see Example 3.1.1.

Example 1.3.1 The extensive form for example 1.1.6:




The result of the iterative deletion of weakly dominated strategies is (y, zzx) , implying the outcome (terminal node) t_7 .

Note that this outcome also results from the profiles (y, yzx) and (y, zzy) , and (y, yzy) . ★

Definition 1.3.5 A game has perfect information if all information sets are singletons.


Example 1.3.2 (Simultaneous moves) The prisoners' dilemma:



★

1.3.1 The Reduced Normal Form

Example 1.3.3 Consider the extensive form:



Note that payoffs have not been specified, just the terminal nodes. The normal form is (where $u(z) = (u_1(z), u_2(z))$):

	Stop	Go
Stop, Stop ₁	$u(z_1)$	$u(z_1)$
Stop, Go ₁	$u(z_1)$	$u(z_1)$
Go, Stop ₁	$u(z_2)$	$u(z_3)$
Go, Go ₁	$u(z_2)$	$u(z_4)$

I 's strategies of Stop, Stop₁ and Stop, Go₁ are equivalent. ★

Definition 1.3.6 Two strategies $s_i, s'_i \in S_i$ are strategically equivalent if $u_j(s_i, s_{-i}) = u_j(s'_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ and all j .

In the (pure strategy) reduced normal form of a game, every set of strategically equivalent strategies is replaced by a single representative.

Example 1.3.4 (Example 1.3.3 continued) The reduced normal form is

	Stop	Go
Stop	$u(z_1)$	$u(z_1)$
Go, Stop ₁	$u(z_2)$	$u(z_3)$
Go, Go ₁	$u(z_2)$	$u(z_4)$

The strategy *Stop* for player I in the reduced normal form should be interpreted as the *equivalence class* of extensive form strategies $\{\text{Stop}, \text{Stop}_1, \text{Stop}, \text{Go}_1\}$, where the equivalence relation is given by strategic equivalence. ★

When describing the normal form representation of an extensive form, it is common to (and we will typically) use the reduced normal form.

An extensive form strategy *always* has the form given by (1.3.1), while a normal form strategy may represent an equivalence class of extensive form strategies (and a reduced normal form strategy always does).

1.4 Problems

- 1.4.1. Describe a social choice rule that is unanimous, nondictatorial, and strategy proof when there are two alternatives.
- 1.4.2. Consider the following social choice rule over the set $X = \{x, y, z\}$. There is an exogenously specified order $x \succ y \succ z$, and define $S(X') \equiv \{a \in X' : a \succ a' \ \forall a' \in X' \setminus \{a\}\}$. Then,

$$f(\theta) = S(\{a \in X : a = t(\theta_i) \text{ for some } i\}).$$

Prove that f is unanimous and nondictatorial, but not strategy proof.

- 1.4.3. Consider the Cournot duopoly example (Example 1.1.4).
- (a) Characterize the set of strategies that survive iterated deletion of strictly dominated strategies.
 - (b) Formulate the game when there are $n \geq 3$ firms and identify the set of strategies surviving iterated deletion of strictly dominated strategies.
- 1.4.4. (a) Prove that the order of deletion does not matter for the process of iterated deletion of strictly dominated strategies in a finite game (Remark 1.2.1 shows that *strictly* cannot be replaced by *weakly*).

- (b) Show that the order of deletion matters for the process of iterated deletion of strictly dominated strategies for the following game: $S_1 = S_2 = [0, 1]$ and payoffs

$$u_i(s_1, s_2) = \begin{cases} s_i, & \text{if } s_i < 1, \\ 0, & \text{if } s_i = 1, s_j < 1, \\ s_i, & \text{if } s_i = s_j = 1. \end{cases}$$

- 1.4.5. Suppose $(T, <)$ is an arborescence (recall Definition 1.3.1). Say a node t' is an *immediate predecessor* of a node t if $t' < t$ and there are no intervening nodes, that is, there does not exist a node t'' satisfying

$$t' < t'' < t.$$

Prove that every noninitial node has a *unique* immediate predecessor (and so there is a unique path to every noninitial node from the initial node).

- 1.4.6. (a) Prove that each player i 's information sets H_i are strictly partially ordered by the precedence ranking $<^*$, where we define $h' <^* h$ if there exists $t \in h$ and $t' \in h'$ such that $t' < t$.
- (b) Give an example showing that the set of all information sets is not similarly strictly partially ordered. (Perfect recall implies that the set of i 's information sets also satisfy Property 2(c).)
- (c) Prove that if $h' <^* h$ for $h, h' \in H_i$, then for all $t \in h$, there exists $t' \in h'$ such that $t' < t$. (In other words, an individual player's information is refined through play in the game.)

Chapter 2

A First Look at Equilibrium¹

2.1 Nash Equilibrium

Example 2.1.1 (Battle of the sexes)

		Sheila	
		Opera	Ballet
Bruce	Opera	2, 1	0, 0
	Ballet	0, 0	1, 2


★

Definition 2.1.1 $s^* \in S$ is a Nash equilibrium of $G = \{(S_1, u_1), \dots, (S_n, u_n)\}$ if for all i and for all $s_i \in S_i$,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*).$$

Example 2.1.2 A simple extensive form:

¹Copyright December 7, 2011 by George J. Mailath



$$S_1 = \{L, R\}, S_2 = \{\ell\ell', \ell r', r\ell', rr'\}.$$

$U_j(s_1, s_2) = u_j(z)$, where z is terminal node reached by (s_1, s_2) .

$$II$$

		$\ell\ell'$	$\ell r'$	$r\ell'$	rr'
I	L	2,3	2,3	4,2	4,2
	R	1,0	3,1	1,0	3,1

Two Nash equilibria: $(L, \ell\ell')$ and $(R, \ell r')$. Though $\ell r'$ is a best reply to L , $(L, \ell r')$ is not a Nash equilibrium.

Note that the equilibria are strategy profiles, *not* outcomes. The outcome path for $(L, \ell\ell')$ is $L\ell$, while the outcome path for $(R, \ell r')$ is Rr' . In examples where the terminal nodes are not separately labeled, it is common to also refer to the outcome path as simply the outcome—recall that every outcome path reaches a unique terminal node, and conversely, every terminal node is reached by a unique sequence of actions (and moves of nature).

NOTE: (R, rr') is *not* a Nash eq, even though the outcome path associated with it, Rr' , is a Nash outcome path. ★

$\wp(Y) \equiv$ collection of all subsets of Y , the power set of $Y \equiv \{Y' \subset Y\}$.

A function $\phi : X \rightarrow \wp(Y) \setminus \emptyset$ is a *correspondence* from X to Y , sometimes written $\phi : X \rightrightarrows Y$.

Note that $\phi(x)$ is simply a nonempty subset of Y . If $f : X \rightarrow Y$ is a function, then $\phi(x) = \{f(x)\}$ is a correspondence and a singleton-valued correspondence can be naturally viewed as a function.

Definition 2.1.2 *The best reply correspondence for player i is*

$$\begin{aligned}\phi_i(s_{-i}) &= \arg \max_{s_i \in S_i} u_i(s_i, s_{-i}) \\ &= \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}), \forall s'_i \in S_i\}.\end{aligned}$$

Note: without assumptions on S_i and u_i , ϕ_i is not well defined. When it is well defined everywhere, $\phi_i : S_{-i} \Rightarrow S_i$.

If $\phi_i(s_{-i})$ is a singleton for all s_{-i} , then ϕ_i is i 's *reaction function*.

Remark 2.1.1 Defining $\phi : S \Rightarrow S$ by

$$\phi(s) := \prod_i \phi_i(s_{-i}) = \phi_1(s_{-1}) \times \cdots \times \phi_n(s_{-n}),$$

we have that s^* is a Nash equilibrium if, and only if,

$$s^* \in \phi(s^*).$$

◆

Example 2.1.3 (Continuation of example 1.1.9, Cournot) Recall that, if $q_2 < a - c$, $\arg \max u_i(q_1, q_2)$ is unique and given by

$$q_1(q_2) = \frac{a - c - q_2}{2}.$$

More generally, i 's *reaction function* is

$$\phi_i(q_j) = \max\left\{\frac{1}{2}(a - c - q_j), 0\right\}.$$

Nash eq (q_1^*, q_2^*) solves

$$\begin{aligned}q_1^* &= \phi_1(q_2^*), \\ q_2^* &= \phi_2(q_1^*).\end{aligned}$$




Figure 2.1.1: The reaction (or best reply) functions for the Cournot game.

So (ignoring the boundary conditions for a second),

$$\begin{aligned}
 q_1^* &= \frac{1}{2}(a - c - q_2^*) \\
 &= \frac{1}{2}\left(a - c - \frac{1}{2}(a - c - q_1^*)\right) \\
 &= \frac{1}{2}(a - c) - \frac{1}{4}(a - c) + \frac{q_1^*}{4} \\
 &= \frac{1}{4}(a - c) + \frac{q_1^*}{4}
 \end{aligned}$$

and so

$$q_1^* = \frac{1}{3}(a - c).$$

Thus,

$$q_2^* = \frac{1}{2}(a - c) - \frac{1}{6}(a - c) = \frac{1}{3}(a - c).$$

Since

$$p = a - q_1^* - q_2^* = a - \frac{2}{3}(a - c) = \frac{1}{3}a + \frac{2}{3}c > 0.$$

Thus the boundary condition is not binding. Note also that there is no equilibrium with zero prices. ★

Example 2.1.4 (continuation of Example 1.1.7) Suppose $v_1 < v_2$, and the valuations are commonly known. There are many Nash equilibria: Bidding v_i for each i is a Nash equilibrium (of course?). But so is any bidding profile (b_1, b_2) satisfying $b_1 < b_2$, $b_1 \leq v_2$ and $v_1 \leq b_2$ (Why? Make sure you understand why some inequalities are weak and some are strict). Are there any other equilibria? ★

Example 2.1.5 (continuation of example 1.1.3) Suppose $v_1 = v_2 = v$, and the valuations are commonly known. The unique Nash equilibrium is for both bidders to bid $b_i = v$. But this eq is in weakly dominated strategies. But what if bids are in pennies? ★

2.1.1 Why Study Nash Equilibrium?

Nash equilibrium is based on two principles:

1. each player is optimizing given beliefs/predictions about the behavior of the other players; and
2. these beliefs/predictions are correct.

While optimization is not in principle troubling (it is true almost by definition), the consistency of beliefs with the actual behavior is a strong assumption. Where does this consistency come from?

Several arguments have been suggested:

1. preplay communication (but see Section 2.5.2),
2. self-fulfilling prophecy (if a theory did not specify a Nash equilibrium, it would invalidate itself),




Figure 2.2.1: An entry game.

3. focal points (natural way to play),
4. introspection more generally (currently not viewed as persuasive),
5. learning (either individual or social), see Section 4.2, and
6. provides important discipline on modelling.

2.2 Credible Threats and Backward Induction

Example 2.2.1 (Entry deterrence) The entry game illustrated in Figure 2.2.1 has two Nash equilibria: (In, Accommodate) and (Out, Fight). The latter violates backward induction. ★

Example 2.2.2 (The case of the reluctant kidnapper) Kidnapper has two choices after receiving ransom: release or kill victim. After release, victim has two choices: whether or not to reveal identity of kidnapper. Payoffs are illustrated in Figure 2.2.2. Victim is killed in only outcome satisfying backward induction.

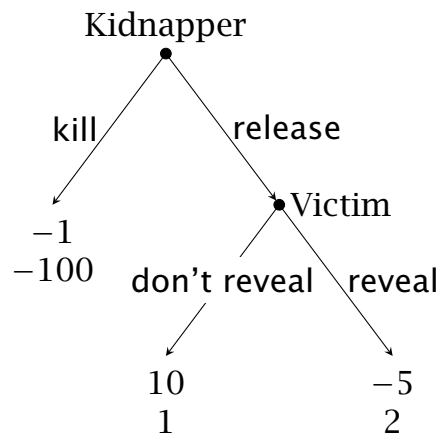


Figure 2.2.2: The Reluctant Kidnapper




Figure 2.2.3: A short centipede game.

★

Example 2.2.3 (Rosenthal's centipede game) The perils of backward induction are illustrated in Figure 2.2.3, with reduced normal form given in Figure 2.2.4.

A longer (and more dramatic version of the centipede is given in Figure 2.2.5.

Backward induction solution in both cases is that both players choose to stop the game at *each* decision point.

★

	Stop	Go
Stop	1, 0	1, 0
Go, Stop ₁	0, 10	100, 1
Go, Go ₁	0, 10	10, 1000

Figure 2.2.4: The reduced normal form for the short centipede in Figure 2.2.3




Figure 2.2.5: A long centipede game.




Figure 2.3.1: (L, T, r) is Nash. Is it plausible?

Theorem 2.2.1 (Zermelo, Kuhn): *A finite game of perfect information has a pure strategy Nash equilibrium.*

2.3 Subgame Perfection

Example 2.3.1 In the game illustrated in Figure 2.3.1, the profile (L, T, r) is Nash. Is it plausible? ★

Define $S(t) \equiv \{t' \in T : t' \prec t\}$.

Definition 2.3.1 *The subset $T^t \equiv \{t\} \cup S(t)$, of T , together with pay-offs, etc. appropriately restricted to T^t , is a subgame if for all information sets h ,*

$$h \cap T^t \neq \emptyset \Rightarrow h \subset T^t.$$

The information set containing the initial node of a subgame is necessarily a singleton (see Problem 2.6.6).

Definition 2.3.2 *The strategy profile s is a subgame perfect equilibrium if s prescribes a Nash equilibrium in every subgame.*

Example 2.3.2 (augmented PD)

	E	S	P
E	2, 2	-1, 3	-1, -1
S	3, -1	0, 0	-1, -1
P	-1, -1	-1, -1	-2, -2

Play game twice and add payoffs.

Nash strategy profile: E in first period, and S in second period as long as opponent also cooperated in first period, and P if opponent didn't exert effort in first period. Every first period action profile describes an information set for each player. Player i 's strategy is

$$s_i^1 = E,$$

$$s_i^2(a_i, a_j) = \begin{cases} S, & \text{if } a_j = E, \\ P, & \text{if } a_j \neq E. \end{cases}$$

Not subgame perfect: Every first period action profile induces a subgame, on which SS must be played. But the profile prescribes SP after ES , for example. Only subgame perfect equilibrium is always S . ★

Example 2.3.3 (A different repeated game) The *stage game* is

	L	C	R
T	4, 6	0, 0	9, 0
M	0, 0	6, 4	0, 0
B	0, 0	0, 0	8, 8

(T, L) and (M, C) are both Nash eq of the stage game. The profile (s_1, s_2) of the once-repeated game with payoffs added is subgame perfect:

$$s_1^1 = B,$$

$$s_1^2(x, y) = \begin{cases} M, & \text{if } x = B, \\ T, & \text{if } x \neq B, \end{cases}$$

$$s_2^1 = R, \\ s_2^2(x, y) = \begin{cases} C, & \text{if } x = B, \text{ and} \\ L, & \text{if } x \neq B. \end{cases}$$

The outcome path induced by (s_1, s_2) is (BR, MC) .

These are strategies of the extensive form of the repeated game, not of the reduced normal form. The reduced form strategies corresponding to (s_1, s_2) are

$$\begin{aligned} \hat{s}_1^1 &= B, \\ \hat{s}_1^2(y) &= M, \text{ for all } y, \\ \hat{s}_2^1 &= R, \\ \hat{s}_2^2(x) &= \begin{cases} C, & \text{if } x = B, \text{ and} \\ L, & \text{if } x \neq B. \end{cases} \end{aligned} \quad \star$$

Example 2.3.4 Consider the extensive form in Figure 2.3.2.

The game has three Nash eq: (RB, r) , (LT, ℓ) , and (LB, ℓ) . Note that (LT, ℓ) , and (LB, ℓ) are distinct extensive form strategy profiles.

The only subgame perfect equilibrium is (RB, r) .

But, (L, ℓ) also subgame perfect in the extensive form in Figure 2.3.3.

Both games have the same reduced normal form, given in Figure 2.3.4. ★

Remark 2.3.1 (Equivalent representations?) A given strategic setting has both a normal form and an extensive form representation. Moreover, the extensive form apparently contains more information (since it in particular contains information about dynamics and information). For example, the application of weak domination to rule out the (Stay out, Fight) equilibrium can be argued to be less




Figure 2.3.2: A game with “nice” subgames.




Figure 2.3.3: An “equivalent” game with no “nice” subgames.

	ℓ	r
L	2, 0	2, 0
T	-1, 1	4, 0
B	0, 0	5, 1

Figure 2.3.4: The reduced form for the games in Figures 2.3.2 and 2.3.3.

compelling than the backward induction (ex post) argument in the extensive form: faced with the fait accompli of Enter, the incumbent “must” Accommodate. As Kreps and Wilson (1982, p. 886) write: “analysis based on normal form representation inherently ignore the role of anticipated actions off the equilibrium path...and in the extreme yields Nash equilibria that are patently implausible.”

But, backward induction and iterated deletion of weakly dominated strategies lead to the same outcomes in finite games of perfect information. Motivated by this and other considerations, a “classical” argument holds that all extensive forms with the same reduced normal form representation are strategically equivalent (Kohlberg and Mertens (1986) is a well known statement of this position; see also Elmes and Reny (1994)). Such a view implies that “good” extensive form solutions should not depend on the extensive form in the way illustrated in Example 2.3.4. For more on this issue, see van Damme (1984) and Mailath, Samuelson, and Swinkels (1993, 1997).



2.4 Mixing

2.4.1 Mixed Strategies and Security Levels

Example 2.4.1 (Matching Pennies) A game with no Nash eq:

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

The greatest payoff that player 1 can guarantee himself may appear to be -1 (the unfortunate result of player 2 correctly anticipating 1's choice).

But suppose that player 1 flips a fair coin so that player 2 cannot anticipate 1's choice. Then, 1 should be able to do better. ★

Definition 2.4.1 Suppose $\{(S_1, u_i), \dots, (S_n, u_n)\}$ is an n -player normal form game. A mixed strategy for player i is a probability distribution over S_i , denoted σ_i . Strategies in S_i are called pure strategies. A strategy σ_i is completely mixed if $\sigma_i(s_i) > 0$ for all $s_i \in S_i$.

In order for the set of mixed strategies to have a nice mathematical structure (such as being metrizable or compact), we need the set of pure strategies to also have a nice structure (often complete separable metric, i.e., Polish). For our purposes here, it will suffice to consider finite sets, or nice subsets of \mathbb{R}^k . More generally, a mixed strategy is a probability *measure* over the set of pure strategies. The set of probability measures over a set A is denoted $\Delta(A)$.

If S_i is finite, $\sigma_i : S_i \rightarrow [0, 1]$ such that $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

Extend u_i to $\prod_{j=1}^n \Delta(S_j)$ by taking expected values, so that u_i is i 's expected payoff under randomization.

If S_i is finite,

$$u_i(\sigma_1, \dots, \sigma_n) = \sum_{s_1 \in S_1} \cdots \sum_{s_n \in S_n} u_i(s_1, \dots, s_n) \sigma_1(s_1) \cdots \sigma_n(s_n).$$

Writing

$$u_i(s_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \prod_{j \neq i} \sigma_j(s_j),$$

we then have

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

Definition 2.4.2 *Player i 's security level is the greatest payoff that i can guarantee himself:*

$$\underline{v}_i = \sup_{\sigma_i \in \Delta(S_i)} \inf_{\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)} u_i(\sigma_i, \sigma_{-i}).$$

If σ_i^ achieves the sup, then σ_i^* is a security strategy for i .*

In matching pennies, each player's security level is 0, guaranteed by the security strategy $\frac{1}{2} \circ H + \frac{1}{2} \circ T$.

2.4.2 Domination and Optimality

Example 2.4.2 In the following game (payoffs are for the row player), M is not dominated by any strategy (pure or mixed) and it is the unique best reply to $\frac{1}{2} \circ L + \frac{1}{2} \circ R$:

	L	R
T	3	0
M	2	2
B	0	3

In the following game (again, payoffs are for row player), M is not dominated by T or B , it is never a best reply, and it is strictly dominated by $\frac{1}{2} \circ T + \frac{1}{2} \circ B$:

	L	R
T	5	0
M	2	2
B	0	5

★

Definition 2.4.3 *The strategy $s'_i \in S_i$ is strictly dominated by the mixed strategy $\sigma_i \in \Delta(S_i)$ if*

$$u_i(\sigma_i, s_{-i}) > u_i(s'_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}.$$

Henceforth, a strategy is strictly (or weakly) undominated if there is no pure *or mixed* strategy that strictly (or weakly, respectively) dominates it.

It is immediate that Definition 2.4.3 is equivalent to Definition 8.B.4 in MWG.

Lemma 2.4.1 *Suppose $n = 2$. The strategy $s'_1 \in S_1$ is not strictly dominated by any other pure or mixed strategy if, and only if, $s'_1 \in \arg \max u_1(s_1, \sigma_2)$ for some $\sigma_2 \in \Delta(S_2)$.*

Proof. We present the proof for finite S_i .

If there exists $\sigma_2 \in \Delta(S_2)$ such that $s'_1 \in \arg \max u_1(s_1, \sigma_2)$, then it is straightforward to show that s'_1 is not strictly dominated by any other pure or mixed strategy (left as exercise).

Suppose s'_1 is a player 1 strategy satisfying

$$s'_1 \notin \arg \max u_1(s_1, \sigma_2) \quad \forall \sigma_2 \in \Delta(S_2). \quad (2.4.1)$$


Define $x(s_1, s_2) = u_1(s_1, s_2) - u_1(s'_1, s_2)$, and observe that for fixed s_2 , we can represent the vector of payoff differences $\{x(s_1, s_2) : s_1 \neq s'_1\}$ as a point in $\mathbb{R}^{|S_1|-1}$. Define

$$X \equiv \text{conv}\{x \in \mathbb{R}^{|S_1|-1} : x_{s_1} = x(s_1, s_2), s_1 \neq s'_1, \text{ some } s_2 \in S_2\}.$$

Denote the closed negative orthant by $\mathbb{R}^{|S_1|-1}_- \equiv \{x \in \mathbb{R}^{|S_1|-1} : x_{s_1} \leq 0, \forall s_1 \neq s'_1\}$. Equation (2.4.1) implies that for all $\sigma_2 \in \Delta(S_2)$, there exists s_1 such that $\sum x(s_1, s_2) \sigma_2(s_2) > 0$, and so $\mathbb{R}^{|S_1|-1}_- \cap X = \emptyset$. Moreover, X is closed, since it is the convex hull of a finite number of vectors. See Figure 2.4.1.

By an appropriate strict separating hyperplane theorem (see, for example, Vohra (2005, Theorem 3.7)), $\exists \lambda \in \mathbb{R}^{|S_1|-1} \setminus \{0\}$ such that $\lambda \cdot x > \lambda \cdot x'$ for all $x \in X$ and all $x' \in \mathbb{R}^{|S_1|-1}_-$. Since $\mathbb{R}^{|S_1|-1}_-$ is unbounded below, $\lambda(s_1) \geq 0 \forall s_1$ (otherwise making $|x'(s_1)|$ large enough for s_1 satisfying $\lambda(s_1) < 0$ ensures $\lambda \cdot x' > \lambda \cdot x$). Define

$$\sigma_1^*(s''_1) = \begin{cases} \lambda(s''_1) / \sum_{s_1 \neq s'_1} \lambda(s_1), & \text{if } s''_1 \neq s'_1, \\ 0, & \text{if } s''_1 = s'_1. \end{cases}$$

Figure 2.4.1: The sets X and $\mathbb{R}_-^{|S_1|-1}$.

We now argue that σ_1^* is a mixed strategy for 1 strictly dominating s'_1 : Since $0 \in \mathbb{R}_-^{|S_1|-1}$, we have $\lambda \cdot x > 0$ for all $x \in X$, and so

$$\sum_{s_1 \neq s'_1} \sigma_1^*(s_1) \sum_{s_2} x(s_1, s_2) \sigma_2(s_2) > 0, \quad \forall \sigma_2,$$

i.e., for all σ_2 ,

$$\begin{aligned} u_1(\sigma_1^*, \sigma_2) &= \sum_{s_1 \neq s'_1, s_2} u_1(s_1, s_2) \sigma_1^*(s_1) \sigma_2(s_2) \\ &> \sum_{s_1 \neq s'_1, s_2} u_1(s'_1, s_2) \sigma_1^*(s_1) \sigma_2(s_2) = u_1(s'_1, \sigma_2). \end{aligned}$$

■

Remark 2.4.1 This proof requires us to strictly separate two disjoint closed convex sets (one bounded), rather than a point from a closed convex set (the standard separating hyperplane theorem). To apply the standard theorem, define $Y \equiv \{y \in \mathbb{R}^{|S_1|-1} : \exists x \in$

$X, y_\ell \geq x_\ell \forall \ell\}$. Clearly Y is closed, convex and $0 \notin Y$. We can proceed as in the proof, since the normal for the separating hyperplane must again have only nonnegative coordinates (use now the unboundedness of Y).

◆

Iterated strict dominance is thus iterated non-best replies—the *rationalizability* notion of Bernheim (1984) and Pearce (1984). But, assumes utilities are disjoint from subjective probabilities—not really in spirit of Savage, see Börgers (1993).

If $n > 2$, issue of correlation across other players. Rationalizability formally does not allow correlation. Above result only holds if beliefs of opponents play can be correlated.

Lemma 2.4.1 holds for mixed strategies (see Problem 2.6.13).

2.4.3 Equilibrium in Mixed Strategies

Definition 2.4.4 Suppose $\{(S_1, u_1), \dots, (S_n, u_n)\}$ is an n -player normal form game. A Nash eq in mixed strategies is a profile $(\sigma_1^*, \dots, \sigma_n^*)$ such that, for all i , for all $\sigma_i \in \Delta(S_i)$,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*). \quad (2.4.2)$$

Equivalently, for all $s_i \in S_i$,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*),$$

since

$$u_i(\sigma_i, \sigma_{-i}^*) = \sum_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \sigma_i(s_i) \leq u_i(s'_i, \sigma_{-i}^*),$$

where $s'_i \in \arg \max u_i(s_i, \sigma_{-i}^*)$.

Lemma 2.4.2 A strategy σ_i^* is a best reply to σ_{-i}^* (i.e., satisfies (2.4.2)) if, and only if,

$$\sigma_i^*(s'_i) > 0 \implies s'_i \in \arg \max_{s_i} u_i(s_i, \sigma_{-i}^*).$$

Proof. Left as an exercise (Problem 2.6.15). ■

Corollary 2.4.1 *A strategy σ_i^* is a best reply to σ_{-i}^* (i.e., satisfies (2.4.2)) if, and only if, for all $s_i \in S_i$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

Example 2.4.3

	<i>L</i>	<i>R</i>
<i>T</i>	2, 1	0, 0
<i>B</i>	0, 0	1, 1

$\Delta(S_1) = \Delta(S_2) = [0, 1]$, $p = \Pr(T)$, $q = \Pr(L)$.
 ϕ is best replies in mixed strategies:

$$\phi_1(q) = \begin{cases} \{1\}, & \text{if } q > \frac{1}{3}, \\ [0, 1], & \text{if } q = \frac{1}{3}, \\ \{0\}, & \text{if } q < \frac{1}{3}. \end{cases}$$

$$\phi_2(p) = \begin{cases} \{1\}, & \text{if } p > \frac{1}{2}, \\ [0, 1], & \text{if } p = \frac{1}{2}, \\ \{0\}, & \text{if } p < \frac{1}{2}. \end{cases}$$

The best replies are graphed in Figure 2.4.2. ★

2.4.4 Behavior Strategies

What does mixing involve in extensive form games? Recall Definition 1.3.2:

Definition 2.4.5 *A pure strategy for player i is a function*

$$s_i : H_i \rightarrow \cup_h A(h) \text{ such that } s_i(h) \in A(h) \ \forall h \in H_i.$$




Figure 2.4.2: The best reply mappings for Example 2.4.3.

Denote i 's set of pure strategies by S_i . Note that $|S_i| < \infty$ for finite extensive form games.

As above:

Definition 2.4.6 A mixed strategy for player i , σ_i , is a probability distribution over S_i , i.e., $\sigma_i \in \Delta(S_i)$.

Definition 2.4.7 A behavior strategy for player i is a function

$$b_i : H_i \rightarrow \cup_h \Delta(A(h)) \text{ such that } s_i(h) \in \Delta(A(h)) \forall h \in H_i.$$

Write $b_i(h)(a)$ for the probability assigned to the action $a \in A(h)$ by the probability distribution $b_i(h)$.

Note that if $|H_i| = 3$ and $|A(h)| = 2 \forall h \in H_i$, then $|S_i| = 8$ and so $\Delta(S_i)$ is a 7-dimensional simplex. On the hand, a behavior strategy in this case requires only 3 numbers (the probability on the first action in each information set).

The behavior strategy corresponding to a pure strategy s_i is given by

$$b_i(h) = \delta_{s_i(h)}, \forall h \in H_i,$$

where $\delta_x \in \Delta(X)$, $x \in X$, is the degenerate distribution (Kronecker's delta),

$$\delta_x(y) = \begin{cases} 1, & \text{if } y = x, \\ 0, & \text{otherwise.} \end{cases}$$

Definition 2.4.8 *Two strategies for a player i are realization equivalent if, fixing the strategies of the other players, the two strategies induce the same distribution over outcomes (terminal nodes).*

Thus, two strategies that are realization equivalent are strategically equivalent (Definition 1.3.6).

Moreover, if two extensive form strategies only differ in the specification of behavior at an information set that one of those strategies had precluded, then the two strategies are realization equivalent. For example the strategies Stop,Stop₁ and Stop,Go₁ in Example 2.2.3 are realization equivalent.

Given a behavior strategy b_i , the realization-equivalent mixed strategy $\sigma_i \in \Delta(S_i)$ is

$$\sigma_i(s_i) = \prod_{h \in H_i} b_i(h)(s_i(h)).$$

Theorem 2.4.1 (Kuhn) *Every mixed strategy has a realization equivalent behavior strategy.*

The behavior strategy realization equivalent to the mixture σ_i can be calculated as follows: Fix an information set h for player i (i.e., $h \in H_i$). Suppose h is reached with strictly positive probability under σ_i , for *some* specification \hat{s}_{-i} . Then, $b_i(h)$ is the distribution over $A(h)$ implied by σ_i *conditional* on h being reached. While this calculation appears to depend on the particular choice of \hat{s}_{-i} , it turns out it does not. (If for all specifications s_{-i} , h is reached with zero probability under σ_i , then $b_i(h)$ can be determined arbitrarily.)

Using behavior strategies, mixing can be easily accommodated in subgame perfect equilibria.

2.5 Dealing with Multiplicity

2.5.1 I: Refinements

Example 2.5.1 Following game has two Nash equilibria (UL and DR), but only DR is plausible. The other eq is in weakly dominated strategies.

	L	R
U	2, 1	0, 0
D	2, 0	1, 1

★

Natural to require eq be “robust” to small mistakes.

Definition 2.5.1 *An equilibrium σ of a finite normal form game G is (normal form) trembling hand perfect if there exists a sequence $\{\sigma^k\}_k$ of completely mixed strategy profiles converging to σ such that σ_i is a best reply to every σ_{-i}^k in the sequence.*

This is NOT the standard definition in the literature (see Subsection 10.1.1).

Every finite normal form game has a trembling hand perfect equilibrium (see Subsection 10.1.1).

Note that weakly dominated strategies cannot be played in a trembling hand perfect equilibrium:

Theorem 2.5.1 *If a strategy profile in a finite normal form game is trembling hand perfect then it is a Nash equilibrium in weakly undominated strategies. If there are only two players, every Nash equilibrium in weakly undominated strategies is trembling hand perfect.*

Proof. The proof of the first statement is straightforward and left as an exercise (Problem 2.6.16). A proof of the second statement can be found in van Damme (1991, Theorem 3.2.2). ■

Can similarly explore the role of trembles in extensive form games, which we will do in Chapter 5.

Some additional material on trembling hand perfect equilibria are collected in the appendix to this chapter (Section 10.1).

2.5.2 II: Selection

Example 2.5.2 (Focal Points)

	A	a	b
A	2, 2	0, 0	0, 0
a	0, 0	0, 0	2, 2
b	0, 0	2, 2	0, 0

★

Example 2.5.3 (payoff dominance)

	ℓ	r
T	2, 2	0, 0
B	0, 0	1, 1

★

Example 2.5.4 (Renegotiation) Compare with example 2.3.3. The stage game is

	L	C	R
T	4, 4	0, 0	9, 0
M	0, 0	6, 6	0, 0
B	0, 0	0, 0	8, 8

(T, L) and (M, C) are both Nash eq of the stage game. The profile (s_1, s_2) of the once-repeated game with payoffs added is subgame perfect:

$$s_1^1 = B,$$

$$s_1^2(x, y) = \begin{cases} M, & \text{if } x = B, \\ T, & \text{if } x \neq B, \end{cases}$$

$$s_2^1 = R,$$

$$s_2^2(x, y) = \begin{cases} C, & \text{if } x = B, \text{ and} \\ L, & \text{if } x \neq B. \end{cases}$$

The outcome path induced by (s_1, s_2) is (BR, MC) . But TL is Pareto dominated by MC , and so players may renegotiate from TL to MC after 1's deviation. But then 1 has no incentive not to play T in the first period. ★

Example 2.5.5 (risk dominance)

	A	B
A	9, 9	0, 5
B	5, 0	7, 7

While AA is the efficient profile (*payoff dominant*), BB has some interest. In particular, B is “less risky” than A : technically, it is *risk dominant* since B is the unique best reply to the uniform lottery over $\{A, B\}$, i.e., to the mixture

$$\frac{1}{2} \circ A + \frac{1}{2} \circ B. \quad \star$$

2.6 Problems

- 2.6.1. Suppose $\{(S_i, U_i)_{i=1}^n\}$ is a normal form game, and $\hat{s}_1 \in S_1$ is a weakly dominated strategy for player 1. Let $S'_1 = S_1 \setminus \{\hat{s}_1\}$, and $S'_i = S_i$ for $i \neq 1$. Suppose s is a Nash equilibrium of $\{(S'_i, U_i)_{i=1}^n\}$. Prove that s is a Nash equilibrium of $\{(S_i, U_i)_{i=1}^n\}$.
- 2.6.2. Suppose $\{(S_i, U_i)_{i=1}^n\}$ is a normal form game, and s is a Nash equilibrium of $\{(S_i, U_i)_{i=1}^n\}$. Let $\{(S'_i, U_i)_{i=1}^n\}$ be the normal form game obtained by the iterated deletion of some or all strictly dominated strategies. Prove that s is a Nash equilibrium of $\{(S'_i, U_i)_{i=1}^n\}$. (Of course, you must first show that $s_i \in S'_i$ for all i .) Give an example showing that this is false if *strictly* is replaced by *weakly*.

- 2.6.3. Consider (again) the Cournot example (Example 1.1.4). What is the Nash Equilibrium of the n -firm Cournot oligopoly? What happens to both individual firm output and total output as n approaches infinity?
- 2.6.4. Consider now the Cournot duopoly where inverse demand is $P(Q) = a - Q$ but firms have asymmetric marginal costs: c_i for firm i , $i = 1, 2$.
- What is the Nash equilibrium when $0 < c_i < a/2$ for $i = 1, 2$? What happens to firm 2's equilibrium output when firm 1's costs, c_1 , increase? Can you give an intuitive explanation?
 - What is the Nash equilibrium when $c_1 < c_2 < a$ but $2c_2 > a + c_1$?
- 2.6.5. Consider the following Cournot duopoly game: The two firms are identical. The cost function facing each firm is denoted by $C(q)$, is continuously differentiable with $C(0) = 0, C'(0) = 0, C'(q) > 0 \forall q > 0$. Firm i chooses $q_i, i = 1, 2$. Inverse demand is given by $p = P(Q)$, where $Q = q_1 + q_2$ is total supply. Suppose P is continuous and there exists $\bar{Q} > 0$ such that $P(Q) > 0$ for $Q \in [0, \bar{Q})$ and $P(Q) = 0$ for $Q \geq \bar{Q}$. Assume firm i 's profits are strictly concave in q_i for all $q_j, j \neq i$.
- Prove that for each value of q_j , firm i ($i \neq j$) has a unique profit maximizing choice. Denote this choice $R_i(q_j)$. Prove that $R_i(q) = R_j(q)$, i.e., the two firms have the same reaction function. Thus, we can drop the subscript of the firm on R .
 - Prove that $R(0) > 0$ and that $R(\bar{Q}) = 0 < \bar{Q}$.
 - We know (from the maximum theorem) that R is a continuous function. Use the Intermediate Value Theorem to argue that this Cournot game has at least one symmetric Nash equilibrium, i.e., a quantity q^* , such that (q^*, q^*) is a Nash equilibrium. [Hint: Apply the Intermediate Value Theorem to the function $f(q) = R(q) - q$. What does $f(q) = 0$ imply?]
 - Give some conditions on C and P that are sufficient to imply that firm i 's profits are strictly concave in q_i for all $q_j, j \neq i$.
- 2.6.6. (easy) Prove that the information set containing the initial node of a subgame is necessarily a singleton.

2.6.7. In the canonical Stackelberg model, there are two firms, I and II , producing the same good. Their inverse demand function is $P = 6 - Q$, where Q is market supply. Each firm has a constant marginal cost of \$4 per unit and a capacity constraint of 3 units (the latter restriction will not affect optimal behavior, but assuming it eliminates the possibility of negative prices). Firm I chooses its quantity first. Firm II , knowing firm I 's quantity choice, then chooses its quantity. Thus, firm I 's strategy space is $S_1 = [0, 3]$ and firm II 's strategy space is $S_2 = \{\tau_2 \mid \tau_2 : S_1 \rightarrow [0, 3]\}$. A strategy profile is $(q_1, \tau_2) \in S_1 \times S_2$, i.e., an action (quantity choice) for I and a specification for *every* quantity choice of I of an action (quantity choice) for II .

- (a) What are the outcome and payoffs of the two firms implied by the strategy profile (q_1, τ_2) ?
- (b) Show that the following strategy profile does not constitute a Nash equilibrium: $(\frac{1}{2}, \tau_2)$, where $\tau_2(q_1) = (2 - q_1)/2$. Which firm(s) is (are) not playing a best response?
- (c) Prove that the following strategy profile constitutes a Nash equilibrium: $(\frac{1}{2}, \hat{\tau}_2)$, where $\hat{\tau}_2(q_1) = \frac{3}{4}$ if $q_1 = \frac{1}{2}$ and $\hat{\tau}_2(q_1) = 3$ if $q_1 \neq \frac{1}{2}$, i.e., II threatens to flood the market unless I produces exactly $\frac{1}{2}$. Is there any other Nash equilibrium which gives the outcome path $(\frac{1}{2}, \frac{3}{4})$? What are the firms' payoffs in this equilibrium?
- (d) Prove that the following strategy profile constitutes a Nash equilibrium: $(0, \tilde{\tau}_2)$, where $\tilde{\tau}_2(q_1) = 1$ if $q_1 = 0$ and $\tilde{\tau}_2(q_1) = 3$ if $q_1 \neq 0$, i.e., II threatens to flood the market unless I produces exactly 0. What are the firms' payoffs in this equilibrium?
- (e) Given $q_1 \in [0, 2]$, specify a Nash equilibrium strategy profile in which I chooses q_1 . Why is it not possible to do this for $q_1 \in (2, 3]$?
- (f) What is the unique backward induction equilibrium of this game?

2.6.8. Consider the extensive form in Figure 2.6.1.

- (a) What is the normal form of this game?
- (b) Describe the pure strategy Nash equilibrium strategies *and* outcomes of the game.




Figure 2.6.1: The game for Problem 2.6.8

- (c) Describe the pure strategy subgame perfect equilibria (there may only be one).
- 2.6.9. Consider the following game G between two players. Player 1 first chooses between A or B , with A giving payoff of 1 to each player, and B giving a payoff of 0 to player 1 and 3 to player 2. After player 1 has publicly chosen between A and B , the two players play the following bimatrix game (with 1 being the row player):

	L	R
U	1, 1	0, 0
D	0, 0	3, 3

Payoffs in the overall game are given by the sum of payoffs from 1's initial choice and the bimatrix game.

- (a) What is the extensive form of G ?
- (b) Describe a subgame perfect equilibrium strategy profile in pure strategies in which 1 chooses B .

- (c) What is the reduced normal form of G ?
 - (d) What is the result of the iterated deletion of weakly dominated strategies?
- 2.6.10. Suppose s is a pure strategy Nash equilibrium of a finite extensive form game, Γ . Suppose Γ' is a subgame of Γ that is on the path of play of s . Prove that s prescribes a Nash equilibrium on Γ' . (It is probably easier to first consider the case where there are no moves of nature.) (The result is also true for mixed strategy Nash equilibria, though the proof is more notationally intimidating.)
- 2.6.11. Suppose the 2×2 normal form game G has a unique Nash equilibrium, and each player's Nash equilibrium strategy and security strategy are both completely mixed.
- (a) Describe the implied restrictions on the payoffs in G .
 - (b) Prove that each player's security level is given by his/her Nash equilibrium payoff.
 - (c) Give an example showing that (in spite of part 2.6.11(b)), the Nash equilibrium profile need not agree with the strategy profile in which each player is playing his or her security strategy. (This is not possible for zero-sum games, see Problem 4.3.1.)
 - (d) For games like you found in part 2.6.11(c), which is the better prediction of play, security strategy or Nash equilibrium?
- 2.6.12. Suppose $\{(S_1, u_1), \dots, (S_n, u_n)\}$ is a finite normal form game. Prove that if $s'_1 \in S_1$ is strictly dominated in the sense of Definition 2.4.3, then it is not a best reply to any belief over S_{-1} . [While you can prove this by contradiction, try to obtain the direct proof, which is more informative.] (This is the contrapositive of the "straightforward" direction of Lemma 2.4.1.)
- 2.6.13. (a) Prove that Lemma 2.4.1 also holds for mixed strategies, i.e., prove that $\sigma_1 \in \Delta(S_1)$ is not a best reply to any mixture $\sigma_2 \in \Delta(S_2)$ if and only if σ_1 is strictly dominated by some other strategy σ'_1 (i.e., $u_1(\sigma'_1, s_2) > u_1(\sigma_1, s_2)$, $\forall s_2 \in S_2$).
- (b) For the game illustrated in Figure 2.6.2, prove that $\frac{1}{2} \circ T + \frac{1}{2} \circ B$ is not a best reply to any mixture over L and R . Describe a strategy that strictly dominates it.

	L	R
T	5, 0	0, 1
C	2, 6	4, 0
B	0, 0	5, 1

Figure 2.6.2: The game for Problem 2.6.13(b).

- 2.6.14. Suppose $\{(S_1, u_1), (S_2, u_2)\}$ is a two player finite normal form game and let \hat{S}_2 be a strict subset of S_2 . Suppose $s'_1 \in S_1$ is not a best reply to any beliefs with support \hat{S}_2 . Prove that there exists $\varepsilon > 0$ such that s'_1 is not a best reply to any beliefs $\mu \in \Delta(S_2)$ satisfying $\mu(\hat{S}_2) > 1 - \varepsilon$. Is the restriction to two players important?
- 2.6.15. Prove Lemma 2.4.2.
- 2.6.16. Is it necessary to assume that σ is a Nash equilibrium in the definition of normal form trembling hand perfection (Definition 2.5.1)? Prove that every trembling hand perfect equilibrium of a finite normal form game is a Nash equilibrium in weakly undominated strategies.

Chapter 3

Games with Nature¹

3.1 An Introductory Example

Example 3.1.1 (Incomplete information version of example 1.1.4)

Firm 1's costs are private information, while firm 2's are public. Nature determines the costs of firm 1 at the beginning of the game, with $\Pr(c_1 = c_H) = \theta \in (0, 1)$. As in example 1.1.4, firm i 's profit is

$$\pi_i(q_1, q_2; c_i) = [(a - q_1 - q_2) - c_i]q_i,$$

where c_i is firm i 's cost. Assume $c_L, c_H, c_2 < a/2$. A strategy for player 2 is a quantity q_2 . A strategy for player 1 is a function $q_1 : \{c_L, c_H\} \rightarrow \mathbb{R}$. For simplicity, write q_L for $q_1(c_L)$ and q_H for $q_1(c_H)$.

Note that for any strategy profile $((q_H, q_L), q_2)$, the associated outcome is

$$\theta \circ (q_H, q_2) + (1 - \theta) \circ (q_L, q_2),$$

that is, with probability θ , the terminal node (q_H, q_2) is realized, and with probability $1 - \theta$, the terminal node (q_L, q_2) is realized.

To find a Nash equilibrium, we must solve for three numbers q_L^* , q_H^* , and q_2^* .

Assume interior solution. We must have:

$$\begin{aligned} (q_H^*, q_L^*) = \arg \max_{q_H, q_L} & \theta[(a - q_H - q_2^*) - c_H]q_H \\ & + (1 - \theta)[(a - q_L - q_2^*) - c_L]q_L. \end{aligned}$$

¹Copyright December 7, 2011 by George J. Mailath

This implies pointwise maximization, i.e.,

$$\begin{aligned} q_H^* &= \arg \max_{q_1} [(a - q_1 - q_2^*) - c_H] q_1 \\ &= \frac{1}{2}(a - q_2^* - c_H). \end{aligned}$$

and

$$q_L^* = \frac{1}{2}(a - q_2^* - c_L).$$

We must also have

$$\begin{aligned} q_2^* &= \arg \max_{q_2} \theta[(a - q_H^* - q_2 - c_2] q_2 \\ &\quad + (1 - \theta)[(a - q_L^* - q_2 - c_2] q_2 \\ &= \arg \max_{q_2} [(a - \theta q_H^* - (1 - \theta) q_L^* - q_2) - c_2] q_2 \\ &= \frac{1}{2}(a - c_2 - \theta q_H^* - (1 - \theta) q_L^*). \end{aligned}$$

Solving,

$$\begin{aligned} q_H^* &= \frac{a - 2c_H + c_2}{3} + \frac{1 - \theta}{6}(c_H - c_L) \\ q_L^* &= \frac{a - 2c_L + c_2}{3} - \frac{\theta}{6}(c_H - c_L) \\ q_2^* &= \frac{a - 2c_2 + \theta c_H + (1 - \theta) c_L}{3} \end{aligned} \quad \star$$

3.2 Purification

Player i 's mixed strategy σ_i in a complete information game G is said to be *purified* if in an incomplete information version of G (with player i 's type space given by T_i), that player's behavior can be written as a pure strategy $s_i : T_i \rightarrow A_i$ such that

$$\sigma_i(a_i) = \Pr\{s_i(t_i) = a_i\},$$

where \Pr is given by the prior distribution over T_i (and so is player $j \neq i$ beliefs over T_i).

Example 3.2.1

	A	B
A	9, 9	0, 5
B	5, 0	7, 7

Mixed strategy equilibrium: Let $p = \Pr \{A\}$, then

$$\begin{aligned}
 9p &= 5p + 7(1 - p) \\
 \Leftrightarrow 9p &= 7 - 2p \\
 \Leftrightarrow 11p &= 7 \Leftrightarrow p = \frac{7}{11}.
 \end{aligned}$$

Trivial purification: give player i payoff irrelevant type t_i where $t_i \sim \mathcal{U}([0, 1])$, and t_1 and t_2 are independent. Then, the mixed strategy eq is purified by many pure strategy eq in the incomplete information game, such as

$$s_i(t_i) = \begin{cases} B, & \text{if } t_i \leq 4/11, \\ A, & \text{if } t_i \geq 4/11. \end{cases}$$

Harsanyi (1973) purification: Consider game of incomplete information, denoted $G(\varepsilon)$, where $t_i \sim \mathcal{U}([0, 1])$ and t_1 and t_2 are independent:

	A	B
A	$9 + \varepsilon t_1, 9 + \varepsilon t_2$	0, 5
B	5, 0	7, 7

Pure strategy for player i is $s_i : [0, 1] \rightarrow \{A, B\}$. Suppose 2 is following a cutoff strategy,

$$s_2(t_2) = \begin{cases} A, & t_2 \geq \bar{t}_2, \\ B, & t_2 < \bar{t}_2, \end{cases}$$

with $\bar{t}_2 \in (0, 1)$.

Type t_1 expected payoff from A is

$$\begin{aligned} U_1(A, t_1, s_2) &= (9 + \varepsilon t_1) \Pr\{s_2(t_2) = A\} \\ &= (9 + \varepsilon t_1) \Pr\{t_2 \geq \bar{t}_2\} \\ &= (9 + \varepsilon t_1)(1 - \bar{t}_2), \end{aligned}$$

while from B is

$$\begin{aligned} U_1(B, t_1, s_2) &= 5 \Pr\{t_2 \geq \bar{t}_2\} + 7 \Pr\{t_2 < \bar{t}_2\} \\ &= 5(1 - \bar{t}_2) + 7\bar{t}_2 \\ &= 5 + 2\bar{t}_2. \end{aligned}$$

Thus, A is optimal iff

$$(9 + \varepsilon t_1)(1 - \bar{t}_2) \geq 5 + 2\bar{t}_2$$

i.e.,

$$t_1 \geq \frac{11\bar{t}_2 - 4}{\varepsilon(1 - \bar{t}_2)}.$$

Thus the best reply to the cutoff strategy s_2 is a cutoff strategy with $\bar{t}_1 = (11\bar{t}_2 - 4)/\varepsilon(1 - \bar{t}_2)$.² Since the game is symmetric, try for a symmetric eq: $\bar{t}_1 = \bar{t}_2 = \bar{t}$. So

$$\bar{t} = \frac{11\bar{t} - 4}{\varepsilon(1 - \bar{t})},$$

or

$$\varepsilon\bar{t}^2 + (11 - \varepsilon)\bar{t} - 4 = 0. \quad (3.2.1)$$

Let $t(\varepsilon)$ denote the value of \bar{t} satisfying (3.2.1). Note first that $t(0) = 4/11$, and that writing (3.2.1) as $g(t, \varepsilon) = 0$, we can apply the implicit function theorem to conclude that for $\varepsilon > 0$ but close to 0, the cutoff type $t(\varepsilon)$ is close to $4/11$, the probability of the mixed strategy eq in the unperturbed game. In other words, for ε small, there is a symmetric equilibrium in cutoff strategies, with $\bar{t} \in (0, 1)$. This equilibrium is not only pure, but almost everywhere strict!

²indeed, even if player 2 were not following a cutoff strategy, player 1's best reply is a cutoff strategy.

The interior cutoff equilibrium of $G(\varepsilon)$ approximates the mixed strategy equilibrium of $G(0)$ in the following sense: Let $p(\varepsilon)$ be the probability assigned to A by the symmetric cutoff equilibrium strategy of $G(\varepsilon)$. Then $p(0) = 7/11$ and $p(\varepsilon) = 1 - t(\varepsilon)$. Since we argued in the previous paragraph that $t(\varepsilon) \rightarrow 4/11$ as $\varepsilon \rightarrow 0$, we have that for all $\eta > 0$, there exists $\varepsilon > 0$ such that

$$|p(\varepsilon) - p(0)| < \eta. \quad \star$$

Harsanyi's (1973) purification theorem is the most compelling justification for mixed equilibria in finite normal form games: Take *any* sequence of incomplete-information games, where each player's payoffs are subject to private shocks, converging to the complete-information normal form game. Harsanyi proved that every equilibrium (pure or mixed) of the original game is the limit of equilibria of these close-by games with incomplete information. Moreover, in the incomplete-information games, players have essentially strict best replies, and so will not randomize. Consequently, a mixed strategy equilibrium can be viewed as a pure strategy equilibrium of any close-by game of incomplete information.

See Govindan, Reny, and Robson (2003) for a modern exposition and generalization of Harsanyi (1973). A brief introduction can also be found in Morris (2008).

3.3 Auctions and Related Games

Example 3.3.1 (First-price sealed-bid auction—private values)

Bidder i 's value for the object, v_i is known only to i . Nature chooses v_i , $i = 1, 2$ at the beginning of the game, with v_i being independently drawn from the interval $[\underline{v}_i, \bar{v}_i]$, with CDF F_i , density f_i . Bidders know F_i (and so f_i).

This is an example of *independent private values*.

Remark 3.3.1 An auction (or similar environment) is said to have *private values* if each buyer's (private) information is sufficient to determine his value (i.e., it is a sufficient statistic for the other buyers' information). The values are *independent* if each buyer's pri-

vate information is stochastically independent of every other bidder's private information.

An auction (or similar environment) is said to have *interdependent values* if the value of the object to the buyers is unknown at the start of the auction, and if a bidder's (expectation of the) value can be affected by the private information of other bidders. If all bidders have the same value, then we have the case of *pure common value*.

◆

Set of possible bids, \mathbb{R}_+ .

Bidder i 's *ex post* payoff as a function of bids b_1 and b_2 , and values v_1 and v_2 :

$$u_i(b_1, b_2, v_1, v_2) = \begin{cases} 0, & \text{if } b_i < b_j, \\ \frac{1}{2}(v_i - b_i), & \text{if } b_i = b_j, \\ v_i - b_i, & \text{if } b_i > b_j. \end{cases}$$

Suppose 2 uses strategy $\sigma_2 : [\underline{v}_2, \bar{v}_2] \rightarrow \mathbb{R}_+$. Then, bidder 1's expected (or *interim*) payoff from bidding b_1 at v_1 is

$$\begin{aligned} U_1(b_1, v_1; \sigma_2) &= \int u_1(b_1, \sigma_2(v_2), v_1, v_2) dF_2(v_2) \\ &= \frac{1}{2}(v_1 - b_1) \Pr\{\sigma_2(v_2) = b_1\} \\ &\quad + \int_{\{v_2: \sigma_2(v_2) < b_1\}} (v_1 - b_1) f_2(v_2) dv_2. \end{aligned}$$

Player 1's *ex ante* payoff from the strategy σ_1 is given

$$\int U_1(\sigma_1(v_1), v_1; \sigma_2) dF_1(v_1),$$

and so for an optimal strategy σ_1 , the bid $b_1 = \sigma_1(v_1)$ must maximize $U_1(b_1, v_1; \sigma_2)$ for almost all v_1 .

Suppose σ_2 is strictly increasing. Then, $\Pr\{\sigma_2(v_2) = b_1\} = 0$ and

$$U_1(b_1, v_1; \sigma_2) = \int_{\{v_2: \sigma_2(v_2) < b_1\}} (v_1 - b_1) f_2(v_2) dv_2$$

$$\begin{aligned}
&= E[v_1 - b_1 \mid \text{winning}] \Pr\{\text{winning}\} \\
&= (v_1 - b_1) \Pr\{\sigma_2(v_2) < b_1\} \\
&= (v_1 - b_1) \Pr\{v_2 < \sigma_2^{-1}(b_1)\} \\
&= (v_1 - b_1) F_2(\sigma_2^{-1}(b_1)).
\end{aligned}$$

Assuming σ_2 is, moreover, differentiable, and an interior maximum, the first order condition is

$$0 = -F_2(\sigma_2^{-1}(b_1)) + (v_1 - b_1) f_2(\sigma_2^{-1}(b_1)) \frac{d\sigma_2^{-1}(b_1)}{db_1}.$$

But

$$\frac{d\sigma_2^{-1}(b_1)}{db_1} = \frac{1}{\sigma_2'(\sigma_2^{-1}(b_1))},$$

so

$$F_2(\sigma_2^{-1}(b_1)) \sigma_2'(\sigma_2^{-1}(b_1)) = (v_1 - b_1) f_2(\sigma_2^{-1}(b_1)),$$

i.e.,

$$\sigma_2'(\sigma_2^{-1}(b_1)) = \frac{(v_1 - b_1) f_2(\sigma_2^{-1}(b_1))}{F_2(\sigma_2^{-1}(b_1))}.$$

Suppose $F_1 = F_2$ and suppose the eq is symmetric, so that $\sigma_1 = \sigma_2 = \tilde{\sigma}$, and $b_1 = \sigma_1(v) \Rightarrow v = \sigma_2^{-1}(b_1)$. Then,

$$\tilde{\sigma}'(v) = \frac{(v - \tilde{\sigma}(v)) f(v)}{F(v)}. \quad (3.3.1)$$

If $\bar{v} = \underline{v} + 1$ and values are uniformly distributed on $[\underline{v}, \bar{v}]$, then

$$\tilde{\sigma}'(v) = \frac{v - \tilde{\sigma}(v)}{v - \underline{v}},$$

i.e.,

$$(v - \underline{v}) \tilde{\sigma}'(v) + \tilde{\sigma}(v) = v.$$

But,

$$\frac{d}{dv} (v - \underline{v}) \tilde{\sigma}(v) = (v - \underline{v}) \tilde{\sigma}'(v) + \tilde{\sigma}(v),$$

so

$$(v - \underline{v}) \tilde{\sigma}(v) = \frac{v^2}{2} + k,$$

where k is a constant of integration. Moreover, evaluating both sides at $v = \underline{v}$ shows that $k = -\underline{v}^2/2$, and so

$$\begin{aligned} (v - \underline{v}) \tilde{\sigma}(v) &= \frac{v^2 - \underline{v}^2}{2} \\ \Rightarrow \tilde{\sigma}(v) &= \frac{1}{2}(v + \underline{v}). \end{aligned}$$

As an illustration of the kind of arguments that are useful, I now argue that every Nash equilibrium must be in nondecreasing strategies.

Lemma 3.3.1 *Suppose (σ_1, σ_2) is a Nash equilibrium of the first price sealed bid auction with independent private values, with CDF F_i on $[\underline{v}_1, \bar{v}_1]$. Suppose type v'_i wins the auction with positive probability. Then, $\sigma_i(v''_i) \geq \sigma_i(v'_i)$ for all $v''_i > v'_i$.*

Proof. Suppose not. Then there exists $v''_1 < v'_1$ with $\sigma_1(v'_1) =: b'_1 > b''_1 := \sigma_1(v''_1)$.

Incentive compatibility implies

$$\begin{aligned} U_1(b'_1, v'_1; \sigma_2) &\geq U_1(b''_1, v'_1; \sigma_2), \\ \text{and} \quad U_1(b'_1, v''_1; \sigma_2) &\leq U_1(b''_1, v''_1; \sigma_2). \end{aligned}$$

Note that

$$\begin{aligned} U_1(b_1, v'_1; \sigma_2) - U_1(b_1, v''_1; \sigma_2) &= \\ \frac{1}{2}(v'_1 - v''_1) \Pr\{\sigma_2(v_2) = b_1\} &+ (v'_1 - v''_1) \Pr\{\sigma_2(v_2) < b_1\}. \end{aligned}$$

Subtracting the second from the first inequality gives

$$U_1(b'_1, v'_1; \sigma_2) - U_1(b'_1, v''_1; \sigma_2) \geq U_1(b''_1, v'_1; \sigma_2) - U_1(b''_1, v''_1; \sigma_2),$$

and so substituting,

$$\begin{aligned} \frac{1}{2}(v'_1 - v''_1) \Pr\{\sigma_2(v_2) = b'_1\} &+ (v'_1 - v''_1) \Pr\{\sigma_2(v_2) < b'_1\} \geq \\ \frac{1}{2}(v'_1 - v''_1) \Pr\{\sigma_2(v_2) = b''_1\} &+ (v'_1 - v''_1) \Pr\{\sigma_2(v_2) < b''_1\}, \end{aligned}$$

and simplifying (and dividing by $(v'_1 - v''_1) < 0$) we get

$$\begin{aligned}
0 &\geq \Pr\{b''_1 \leq \sigma_2(v_2) < b'_1\} \\
&\quad + \frac{1}{2} \{\Pr\{\sigma_2(v_2) = b'_1\} - \Pr\{\sigma_2(v_2) = b''_1\}\} \\
&= \Pr\{b''_1 < \sigma_2(v_2) < b'_1\} \\
&\quad + \frac{1}{2} \{\Pr\{\sigma_2(v_2) = b'_1\} + \Pr\{\sigma_2(v_2) = b''_1\}\}.
\end{aligned}$$

This implies

$$\begin{aligned}
0 &= \Pr\{\sigma_2(v_2) = b'_1\}, \\
0 &= \Pr\{\sigma_2(v_2) = b''_1\}, \\
\text{and} \quad 0 &= \Pr\{b''_1 < \sigma_2(v_2) < b'_1\}.
\end{aligned}$$

That is, bidder 2 does not make a bid between b''_1 and b'_1 , and there are no ties at b'_1 or b''_1 . A bid of b'_1 and b''_1 therefore wins with the same probability. But this implies a contradiction: Since b'_1 wins with positive probability, v'_1 strictly prefers to win with the same probability at the strictly lower bid of b''_1 . ■
★

Example 3.3.2 (independent private values, symmetric n bidders)

Suppose now there n identical bidders, with valuations v_i independently distributed on $[\underline{v}, \bar{v}]$ according to F with density f .

Interested in characterizing the symmetric Nash equilibrium (if it exists). Let σ be the symmetric strategy, and suppose it is strictly increasing. Consequently, the probability of a tie is zero, and so bidder i 's interim payoff from the bid b_i is

$$\begin{aligned}
U_i(b_i, v_i; \sigma) &= E[v_i - b_i \mid \text{winning}] \Pr\{\text{winning}\} \\
&= (v_i - b_i) \Pr\{v_j < \sigma^{-1}(b_i), \forall j \neq i\} \\
&= (v_i - b_i) \prod_{j \neq i} \Pr\{v_j < \sigma^{-1}(b_i)\} \\
&= (v_i - b_i) F^{n-1}(\sigma^{-1}(b_i)).
\end{aligned}$$

As before, assuming σ is differentiable, and an interior solution, the first order condition is

$$0 = -F^{n-1}(\sigma^{-1}(b_i)) + (v_i - b_i)(n-1)F^{n-2}(\sigma^{-1}(b_i))f(\sigma^{-1}(b_i))\frac{d\sigma^{-1}(b_i)}{db_i},$$

and simplifying (similarly to (3.3.1)), we get

$$\begin{aligned} \sigma'(v)F^{n-1}(v) + \sigma(v)(n-1)F^{n-2}(v)f(v) \\ = v(n-1)F^{n-2}(v)f(v), \end{aligned}$$

that is,

$$\frac{d}{dv}\sigma(v)F^{n-1}(v) = v(n-1)F^{n-2}(v)f(v),$$

or (where the constant of integration is zero, since $F(v) = 0$),

$$\sigma(v) = \frac{1}{F^{n-1}(v)} \int_v^v x dF^{n-1}(x).$$

Remark 3.3.2 (Order statistics) Given n independent draws from a common distribution F , denoted v_1, \dots, v_n , let $y_1^{(n)}, y_2^{(n)}, \dots, y_n^{(n)}$ denote the rearrangement satisfying $y_1^{(n)} \geq y_2^{(n)} \geq \dots \geq y_n^{(n)}$. The statistic $y_k^{(n)}$ is the k^{th} -order statistic. (**Warning:** Some authors reverse the inequalities.)

The distribution of $y_1^{(n)}$ is $\Pr\{y_1^{(n)} \leq y\} = F^n(y)$.

◆

If σ is a symmetric Nash equilibrium, then

$$\sigma(v) = E[y_1^{(n-1)} \mid y_1^{(n-1)} \leq v].$$

That is, each bidder bids the expectation of the maximum of all the other bidders' valuation, conditional on that valuation being less than his (i.e., conditional on his value being the highest). Equivalently, the bidder bids the expected value of the second order statistic of values, conditional on his value being the first order statistic. ★

Example 3.3.3 (First-price sealed-bid auction—pure common values)

Each bidder receives a private signal about the value of the object, t_i , with $t_i \in T_i = [0, 1]$, uniformly independently distributed. The common (to both players) value of the object is $v = t_1 + t_2$.

Ex post payoffs are given by

$$u_i(b_1, b_2, t_1, t_2) = \begin{cases} t_1 + t_2 - b_i, & \text{if } b_i > b_j, \\ \frac{1}{2}(t_1 + t_2 - b_i), & \text{if } b_i = b_j, \\ 0, & \text{if } b_i < b_j. \end{cases}$$

Suppose 2 uses strategy $\sigma_2 : T_2 \rightarrow \mathbb{R}_+$. Suppose σ_2 is strictly increasing. Then, t_1 's expected payoff from bidding b_1 is

$$\begin{aligned} U_1(b_1, t_1; \sigma_2) &= E[t_1 + t_2 - b_1 \mid \text{winning}] \Pr\{\text{winning}\} \\ &= E[t_1 + t_2 - b_1 \mid t_2 < \sigma_2^{-1}(b_1)] \Pr\{t_2 < \sigma_2^{-1}(b_1)\} \\ &= (t_1 - b_1)\sigma_2^{-1}(b_1) + \int_0^{\sigma_2^{-1}(b_1)} t_2 dt_2 \\ &= (t_1 - b_1)\sigma_2^{-1}(b_1) + (\sigma_2^{-1}(b_1))^2/2. \end{aligned}$$

If σ_2 is differentiable, the first order condition is

$$0 = -\sigma_2^{-1}(b_1) + (t_1 - b_1) \frac{d\sigma_2^{-1}(b_1)}{db_1} + \sigma_2^{-1}(b_1) \frac{d\sigma_2^{-1}(b_1)}{db_1},$$

and so

$$\sigma_2^{-1}(b_1) \sigma_2'(\sigma_2^{-1}(b_1)) = (t_1 + \sigma_2^{-1}(b_1) - b_1).$$

Suppose $F_1 = F_2$ and suppose the eq is symmetric, so that $\sigma_1 = \sigma_2 = \sigma$, and $b_1 = \sigma_1(t) \Rightarrow t = \sigma_2^{-1}(b_1)$. Then,

$$t\sigma'(t) = 2t - \sigma(t).$$

Integrating,

$$t\sigma(t) = t^2 + k,$$

where k is a constant of integration. Evaluating both sides at $t = 0$ shows that $k = 0$, and so

$$\sigma(t) = t.$$

Note that this is NOT the profile that results from the analysis of the private value auction when $v = 1/2$ (the value of the object

in the common value auction, conditional on *only* t_1 , is $E[t_1 + t_2 \mid t_1] = t_1 + 1/2$). In particular, letting $v' = t + \frac{1}{2}$, we have

$$\sigma_{\text{private}}(t) = \tilde{\sigma}(v') = \frac{v' + 1/2}{2} = \frac{t + 1}{2} > t = \sigma_{\text{public}}(t).$$

This illustrates the winner's curse: $E[v \mid t_1] > E[v \mid t_1, \text{winning}]$. In particular, in the equilibrium just calculated,

$$\begin{aligned} E[v \mid t_1, \text{winning}] &= E[t_1 + t_2 \mid t_1, t_2 < t_1] \\ &= t_1 + \frac{1}{\Pr\{t_2 < t_1\}} \int_0^{t_1} t_2 dt_2 \\ &= t_1 + \frac{1}{t_1} \left[(t_2)^2 / 2 \right]_0^{t_1} = \frac{3t_1}{2}, \end{aligned}$$

while $E[v \mid t_1] = t_1 + \frac{1}{2} > 3t_1/2$ (recall $t_1 \in [0, 1]$).

Notice moreover, that the bidder bids as if it is a private values auction and the value of the object is $E[v_1 \mid t_1, t_2 = t_1]$. ★

Example 3.3.4 (War of attrition)

Action spaces $S_i = \mathbb{R}_+$.

Private information (type) $t_i \in T_i \equiv \mathbb{R}_+$, CDF F_i , density f_i , with f_i strictly positive on T_i .

Ex post payoffs

$$u_i(s_1, s_2, t_1, t_2) = \begin{cases} t_i - s_j, & \text{if } s_j < s_i, \\ -s_i, & \text{if } s_j \geq s_i. \end{cases}$$

Suppose 2 uses strategy $\sigma_2 : T_2 \rightarrow S_2$. Then, t_1 's expected (or *interim*) payoff from stopping at s_1 is

$$\begin{aligned} U_1(s_1, t_1; \sigma_2) &= \int u_1(s_1, \sigma_2(t_2), t) dF_2(t_2) \\ &= -s_1 \Pr\{\sigma_2(t_2) \geq s_1\} + \int_{\{t_2: \sigma_2(t_2) < s_1\}} (t_1 - \sigma_2(t_2)) dF_2(t_2). \end{aligned}$$

Any Nash equilibrium is *sequentially rational* on the equilibrium path: Suppose $\tau < \sigma_1(t_1)$ is reached (i.e., 2 has not yet dropped

out) and $\Pr\{\sigma_2(t_2) > \tau\} > 0$ (so that such an event has positive probability). Is stopping at $\sigma_1(t_1)$ still optimal? Suppose that, conditional on play reaching τ , the stopping time $\hat{s}_1 > \tau$ yields a higher payoff than the original stopping time $s_1 = \sigma_1(t_1)$, i.e.,

$$E_{t_2}[u_1(s_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) > \tau] < E_{t_2}[u_1(\hat{s}_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) > \tau]$$

Then,

$$\begin{aligned} U_1(s_1, t_1; \sigma_2) &= E_{t_2}[u_1(s_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) \leq \tau] \Pr\{\sigma_2(t_2) \leq \tau\} \\ &\quad + E_{t_2}[u_1(s_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) > \tau] \Pr\{\sigma_2(t_2) > \tau\} \\ &< E_{t_2}[u_1(s_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) \leq \tau] \Pr\{\sigma_2(t_2) \leq \tau\} \\ &\quad + E_{t_2}[u_1(\hat{s}_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) > \tau] \Pr\{\sigma_2(t_2) > \tau\} \\ &= E_{t_2}[u_1(\hat{s}_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) \leq \tau] \Pr\{\sigma_2(t_2) \leq \tau\} \\ &\quad + E_{t_2}[u_1(\hat{s}_1, \sigma_2(t_2), t) \mid \sigma_2(t_2) > \tau] \Pr\{\sigma_2(t_2) > \tau\} \\ &= U_1(\hat{s}_1, t_1; \sigma_2), \end{aligned}$$

and so s_1 cannot have been the unconditionally optimal stopping time. This is an application of the principle of Problem 2.6.10 to an infinite game.

Define $\bar{s}_i \equiv \inf\{s_i : \Pr\{\sigma_i(t_i) \leq s_i\} = 1\} = \inf\{s_i : \Pr\{\sigma_i(t_i) > s_i\} = 0\}$, where $\inf\{\emptyset\} = \infty$. It can be shown that in any Nash eq with $\bar{s}_1, \bar{s}_2 > 0$, $\bar{s}_1 = \bar{s}_2$. (If $\bar{s}_2 < \bar{s}_1$, then for sufficiently large types for player 2, there are late stopping times that are profitable deviations—see Problem 3.6.3).

Lemma 3.3.2 *Suppose (σ_1, σ_2) is a Nash eq. profile. Then, σ_i is nondecreasing for $i = 1, 2$.*

Proof. We use a standard revealed preference argument. Let $s'_1 = \sigma_1(t'_1)$ and $s''_1 = \sigma_1(t''_1)$, with $s'_1, s''_1 \leq \bar{s}_1$. If σ_1 is a best reply to σ_2 ,

$$U_1(s'_1, t'_1; \sigma_2) \geq U_1(s''_1, t'_1; \sigma_2)$$

and

$$U_1(s''_1, t''_1; \sigma_2) \geq U_1(s'_1, t''_1; \sigma_2).$$

Thus,

$$U_1(s'_1, t'_1; \sigma_2) - U_1(s'_1, t''_1; \sigma_2) \geq U_1(s''_1, t'_1; \sigma_2) - U_1(s''_1, t''_1; \sigma_2).$$

Since,

$$U_1(s_1, t'_1; \sigma_2) - U_1(s_1, t''_1; \sigma_2) = (t'_1 - t''_1) \Pr\{t_2 : \sigma_2(t_2) < s_1\}$$

we have

$$(t'_1 - t''_1) \Pr\{t_2 : \sigma_2(t_2) < s'_1\} \geq (t'_1 - t''_1) \Pr\{t_2 : \sigma_2(t_2) < s''_1\},$$

i.e.,

$$(t'_1 - t''_1) [\Pr\{t_2 : \sigma_2(t_2) < s'_1\} - \Pr\{t_2 : \sigma_2(t_2) < s''_1\}] \geq 0.$$

Suppose $t'_1 > t''_1$. Then, $\Pr\{t_2 : \sigma_2(t_2) < s'_1\} \geq \Pr\{t_2 : \sigma_2(t_2) < s''_1\}$. If $s'_1 < s''_1$, then $\Pr\{t_2 : s'_1 \leq \sigma_2(t_2) < s''_1\} = 0$. That is, 2 does not stop between s'_1 and s''_1 .

The argument to this point has only used the property that σ_1 is a best reply to σ_2 . To complete the argument, we appeal to the fact that $\bar{s}_1 \leq \bar{s}_2$ (an implication of σ_2 being a best reply to σ_1), which implies $\Pr\{\sigma_2(t_2) \geq s''_1\} > 0$, and so stopping earlier (at a time $s_1 \in (s'_1, s''_1)$) is a profitable deviation for t'_1 . Thus, $s'_1 = \sigma_1(t'_1) \geq s''_1 = \sigma_1(t''_1)$. ■

It can also be shown that in any Nash eq, σ_i is a strictly increasing and continuous function. Thus,

$$\begin{aligned} U_1(s_1, t_1; \sigma_2) &= -s_1 \Pr\{t_2 \geq \sigma_2^{-1}(s_1)\} + \int_{\{t_2 < \sigma_2^{-1}(s_1)\}} (t_1 - \sigma_2(t_2)) f_2(t_2) dt_2 \\ &= -s_1 (1 - F_2(\sigma_2^{-1}(s_1))) + \int_0^{\sigma_2^{-1}(s_1)} (t_1 - \sigma_2(t_2)) f_2(t_2) dt_2. \end{aligned}$$

Assuming σ_2 is, moreover, differentiable, the first order condition is

$$0 = -(1 - F_2(\sigma_2^{-1}(s_1)))$$

$$+ s_1 f_2 (\sigma_2^{-1} (s_1)) \frac{d\sigma_2^{-1} (s_1)}{ds_1} + (t_1 - s_1) f_2 (\sigma_2^{-1} (s_1)) \frac{d\sigma_2^{-1} (s_1)}{ds_1}.$$

But

$$\frac{d\sigma_2^{-1} (s_1)}{ds_1} = 1 / \sigma_2' (\sigma_2^{-1} (s_1)) ,$$

so

$$\{1 - F_2 (\sigma_2^{-1} (s_1))\} \sigma_2' (\sigma_2^{-1} (s_1)) = t_1 f_2 (\sigma_2^{-1} (s_1)) ,$$

i.e.,

$$\sigma_2' (\sigma_2^{-1} (s_1)) = \frac{t_1 f_2 (\sigma_2^{-1} (s_1))}{\{1 - F_2 (\sigma_2^{-1} (s_1))\}}.$$

Suppose $F_1 = F_2$ and suppose the eq is symmetric, so that $\sigma_1 = \sigma_2 = \sigma$, and $s_1 = \sigma_1 (t) \Rightarrow t = \sigma_2^{-1} (s_1)$. Then,

$$\sigma' (t) = \frac{t f (t)}{1 - F (t)}.$$

Since $\sigma (0) = 0$,

$$\sigma (t) = \int_0^t \frac{\tau f (\tau)}{1 - F (\tau)} d\tau.$$

If $f (t) = e^{-t}$, then $F (t) = 1 - e^{-t}$, and

$$\sigma (t) = \int_0^t \frac{\tau e^{-\tau}}{e^{-\tau}} d\tau = t^2 / 2.$$

Note that $\sigma (t) > t$ for $t > 2$!

If we extend the strategy space to allow for never stopping, i.e., $S_i = \mathbb{R}_+ \cup \{\infty\}$ and allow payoffs to take on the value $-\infty$, then there are also two asymmetric equilibria, in which one player drops out immediately, and the other never drops out. ★

Example 3.3.5 (Double Auction) Let $v_s \sim U [0, 1]$, $v_b \sim U [0, 1]$, v_s and v_b independent. If sale occurs at price p , buyer receives $v_b - p$, seller receives $p - v_s$. Seller and buyer simultaneously propose prices $p_s \in [0, 1]$ and $p_b \in [0, 1]$ respectively. Trade at $\frac{1}{2} (p_s + p_b)$ if $p_s \leq p_b$; otherwise no trade.

Buyer's strategy is a function $\tilde{p}_b : [0, 1] \rightarrow [0, 1]$, seller's strategy is a function $\tilde{p}_s : [0, 1] \rightarrow [0, 1]$. We check for interim optimality (i.e., optimality of a strategy conditional on a type). (With a continuum of strategies, ex ante optimality formally requires only that the strategy is optimal for *almost all* types.)

Fix seller's strategy, $\tilde{p}_s : [0, 1] \rightarrow [0, 1]$, buyer's valuation v_b and his bid p_b . Buyer's (conditional) expected payoff is:

$$\begin{aligned} \mathcal{U}_b(p_b, v_b; \tilde{p}_s) &= \int_{\{v_s : p_b \geq \tilde{p}_s(v_s)\}} \left(v_b - \frac{1}{2} (\tilde{p}_s(v_s) + p_b) \right) dv_s \\ &= \Pr(\{v_s : p_b \geq \tilde{p}_s(v_s)\}) \\ &\quad \times \left(v_b - \frac{1}{2} p_b - \frac{1}{2} E(\tilde{p}_s(v_s) \mid \{v_s : p_b \geq \tilde{p}_s(v_s)\}) \right). \end{aligned}$$

Suppose that seller's strategy is linear in his valuation, i.e. $\tilde{p}_s(v_s) = a_s + c_s v_s$, with $a_s \geq 0$, $c_s > 0$ and $a_s + c_s \leq 1$. Then,

$$\Pr(\{v_s : p_b \geq \tilde{p}_s(v_s)\}) = \Pr\left(\left\{v_s : v_s \leq \frac{p_b - a_s}{c_s}\right\}\right).$$

So,

$$\Pr(\{v_s : p_b \geq \tilde{p}_s(v_s)\}) = \begin{cases} 0, & \text{if } p_b \leq a_s, \\ \frac{p_b - a_s}{c_s}, & \text{if } a_s \leq p_b \leq a_s + c_s, \\ 1, & \text{if } p_b \geq a_s + c_s, \end{cases}$$

and so

$$E(\tilde{p}_s(v_s) \mid \{v_s : p_b \geq \tilde{p}_s(v_s)\}) = a_s + c_s E\left(v_s \mid \left\{v_s : v_s \leq \frac{p_b - a_s}{c_s}\right\}\right)$$

(and if $a_s \leq p_b \leq a_s + c_s$)

$$= a_s + c_s \frac{1}{2} \frac{p_b - a_s}{c_s} = \frac{p_b + a_s}{2}.$$

So,

$$E(\tilde{p}_s(v_s) \mid \{v_s : p_b \geq \tilde{p}_s(v_s)\}) = \begin{cases} \text{not defined,} & \text{if } p_b < a_s, \\ \frac{p_b + a_s}{2}, & \text{if } a_s \leq p_b \leq a_s + c_s, \\ a_s + \frac{1}{2} c_s, & \text{if } p_b \geq a_s + c_s. \end{cases}$$

and

$$u_b(p_b, v_b; \tilde{p}_s) = \begin{cases} 0, & \text{if } p_b \leq a_s, \\ \left(\frac{p_b - a_s}{c_s}\right) \left(v_b - \frac{3}{4}p_b - \frac{1}{4}a_s\right), & \text{if } a_s \leq p_b \leq a_s + c_s, \\ v_b - \frac{1}{2}p_b - \frac{1}{2}a_s - \frac{1}{4}c_s, & \text{if } p_b \geq a_s + c_s. \end{cases}$$

The interior expression is maximized by solving the FOC

$$0 = \frac{1}{c_s} \left[v_b - \frac{3}{4}p_b - \frac{1}{4}a_s - \frac{3}{4}p_b + \frac{3}{4}a_s \right],$$

and so

$$\begin{aligned} \frac{6}{4}p_b &= v_b - \frac{a_s}{4} + \frac{3}{4}a_s \\ &= v_b + \frac{1}{2}a_s \\ \Rightarrow p_b &= \frac{1}{3}a_s + \frac{2}{3}v_b. \end{aligned}$$

Thus any bid less than a_s is optimal if $v_b \leq a_s$, $\frac{1}{3}a_s + \frac{2}{3}v_b$ is the unique optimal bid if $a_s < v_b \leq a_s + \frac{3}{2}c_s$ and $a_s + c_s$ is the unique optimal bid if $v_b \geq a_s + \frac{3}{2}c_s$. Thus strategy $\tilde{p}_b(v_b) = \frac{1}{3}a_s + \frac{2}{3}v_b$ is a best response to $\tilde{p}_s(v_s) = a_s + c_s v_s$ as long as $1 \leq a_s + \frac{3}{2}c_s$.

Symmetric argument shows that if $\tilde{p}_b(v_s) = a_b + c_b v_b$, then seller's optimal bid (if interior) is $\tilde{p}_s(v_s) = \frac{2}{3}v_s + \frac{1}{3}(a_b + c_b)$. Thus a linear equilibrium must have $a_b = \frac{1}{3}a_s$, $a_s = \frac{1}{3}(a_b + c_b)$, $c_b = \frac{2}{3}$ and $c_s = \frac{2}{3}$, so $a_s = \frac{1}{4}$ and $a_b = \frac{1}{12}$. There is then a linear equilibrium with:

$$\begin{aligned} \tilde{p}_s(v_s) &= \frac{1}{4} + \frac{2}{3}v_s \\ \tilde{p}_b(v_b) &= \frac{1}{12} + \frac{2}{3}v_b \end{aligned}$$

Efficient trade requires trade if $v_s < v_b$, and no trade if $v_s > v_b$.

Under linear equilibrium, trade occurs iff $\tilde{p}_s(v_s) \leq \tilde{p}_b(v_b)$, which requires

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

Thus, for valuations in the set $\{(\nu_s, \nu_b) \mid \nu_s < \nu_b < \nu_s + \frac{1}{4}\}$, trade is efficient but does not occur in equilibrium.

Note: there are other equilibria (see Problem 3.6.5). ★

3.4 Games of Incomplete Information

Example 3.4.1 Suppose payoffs of a two player two action game are given by one of the two bimatrices:

	<i>H</i>	<i>T</i>		<i>H</i>	<i>T</i>
<i>H</i>	1, 1	0, 0	<i>H</i>	1, 0	0, 1
<i>T</i>	0, 1	1, 0	<i>T</i>	0, 0	1, 1

Either player *I* has dominant strategy to play *H* or a dominant strategy to play *T*. Suppose that *I* knows his own payoffs but player *II* thinks there is probability α that payoffs are given by the first matrix, probability $1 - \alpha$ that they are given by the second matrix. Say that player *II* is of type 1 if payoffs are given by the first matrix, type 2 if payoffs are given by the second matrix. Clearly equilibrium must have: *II* plays *H* if type 1, *T* if type 2; *I* plays *H* if $\alpha > \frac{1}{2}$, *T* if $\alpha < \frac{1}{2}$. But how to analyze this problem in general? ★

Definition 3.4.1 (Harsanyi) A game of incomplete information or Bayesian game is the collection $\{(A_i, T_i, p_i, u_i)_{i=1}^n\}$, where

- A_i is *i*'s action space,
- T_i is *i*'s type space,
- $p_i : T_i \rightarrow \Delta(\prod_{j \neq i} T_j)$ is *i*'s subjective beliefs about the other players' types, given *i*'s type and
- $u_i : \prod_j A_j \times \prod_j T_j \rightarrow \mathbb{R}$ is *i*'s payoff function.

A strategy for i is

$$s_i : T_i \rightarrow A_i.$$

Let $s(t) \equiv (s_1(t_1), \dots, s_n(t_n))$, etc.

Definition 3.4.2 *The profile $(\hat{s}_1, \dots, \hat{s}_n)$ is a Bayesian (or Bayes-Nash) equilibrium if, for all i and all $t_i \in T_i$,*

$$E_{t_{-i}}[u_i(\hat{s}(t), t)] \geq E_{t_{-i}}[u_i(a_i, \hat{s}_{-i}(t_{-i}), t)], \quad \forall a_i \in A_i, \quad (3.4.1)$$

where the expectation over t_{-i} is taken with respect to the probability distribution $p_i(t_i)$.

If the type spaces are finite, then the probability i assigns to the vector $t_{-i} \in \prod_{j \neq i} T_j \equiv T_{-i}$ when his type is t_i can be denoted $p_i(t_{-i} | t_i)$, and (3.4.1) can be written as

$$\sum_{t_{-i}} u_i(\hat{s}(t), t) p_i(t_{-i} | t_i) \geq \sum_{t_{-i}} u_i(a_i, \hat{s}_{-i}(t_{-i}), t) p_i(t_{-i} | t_i), \quad \forall a_i \in A_i.$$

Definition 3.4.3 *The subjective beliefs are consistent or are said to satisfy the Common Prior Assumption (CPA) if there exists a probability distribution $p \in \Delta(\prod_i T_i)$ such that $p_i(t_i)$ is the probability distribution on T_{-i} conditional on t_i implied by p .*

If the type spaces are finite, this is equivalent to

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

If beliefs are consistent, Bayesian game can be interpreted as having an initial move by nature, which selects $t \in T$ according to p . The Common Prior Assumption is controversial, sometimes viewed as a mild assumption (Aumann, 1987) and sometimes not (Gul, 1998). Nonetheless, in applications it is standard to assume it.

For simplicity of notation only, suppose type spaces are finite. Viewed as a game of complete information, a profile \hat{s} is a Nash equilibrium if, for all i ,

$$\sum_t u_i(\hat{s}(t), t) p(t) \geq \sum_t u_i(s_i(t_i), \hat{s}_{-i}(t_{-i}), t) p(t), \quad \forall s_i : T_i \rightarrow A_i.$$

This inequality can be rewritten as (where $p_i^*(t_i) \equiv \sum_{t_{-i}} p(t_{-i}, t_i)$)

$$\sum_{t_i} \left\{ \sum_{t_{-i}} u_i(\hat{s}(t), t) p_i(t_{-i}|t_i) \right\} p_i^*(t_i) \geq \sum_{t_i} \left\{ \sum_{t_{-i}} u_i(s_i(t_i), \hat{s}_{-i}(t_{-i}), t) p_i(t_{-i}|t_i) \right\} p_i^*(t_i),$$

$\forall s_i : T_i \rightarrow A_i.$

If $p_i^*(t_i) \neq 0$, this is then equivalent to the definition of a Bayesian eq.

3.5 Higher Order Beliefs and Global Games

Example 3.5.1

	A	B
A	θ, θ	$\theta - 9, 5$
B	$5, \theta - 9$	$7, 7$

For $\theta = 9$, this is the game studied in examples 3.2.1 and 2.5.5.

Suppose, as in example 3.2.1, that there is incomplete information about payoffs. However, now the information will be correlated. In particular, suppose $\theta \in \{4, 9\}$, with prior probability $\Pr\{\theta = 9\} > 7/9$.

If players have no information about θ , then there are two pure strategy Nash eq, with (A, A) Pareto dominating (B, B) . Now suppose players have some private information represented as follows: Underlying state space is $\Omega \equiv \{\omega_1, \omega_2, \dots, \omega_K\}$, K odd. In all states ω_k , $k \leq K-1$, $\theta = 9$, while in state ω_K , $\theta = 4$. Player 1 has a partition on Ω given by $\{\{\omega_1\}, \{\omega_2, \omega_3\}, \dots, \{\omega_{K-1}, \omega_K\}\}$; we denote $\{\omega_1\}$ by t_1^0 , and $\{\omega_{2\ell}, \omega_{2\ell+1}\}$ by t_1^ℓ . Player 2 has a partition $\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \dots, \{\omega_{K-2}, \omega_{K-1}\}, \{\omega_K\}\}$, and we denote $\{\omega_{2\ell-1}, \omega_{2\ell}\}$ by t_2^ℓ and $\{\omega_K\}$ by $t_2^{(K+1)/2}$. Finally, the probability distribution on Ω is uniform. Figure 3.5.1 illustrates Ω for $K = 9$.




Figure 3.5.1: The dashed lines describe player 1's information sets, while the solid lines describe player 2's.

In state ω_5 , both players know $\theta = 9$, player 1 knows that player 2 knows that $\theta = 9$, but player 2 assigns probability $1/2$ to ω_6 , and so to the event that player 1 does not know that $\theta = 9$.

In state ω_4 , both players know $\theta = 9$, both players know that both players know that $\theta = 9$, player 2 knows that player 1 knows that player 2 knows that $\theta = 9$, but player 1 does not know that player 2 knows that player 1 knows that $\theta = 9$.

A pure strategy for player 1 is

$$s_1 : \{t_1^0, t_1^1, \dots, t_1^{(K-1)/2}\} \rightarrow \{A, B\},$$

while a pure strategy for player 2 is

$$s_2 : \{t_2^1, t_2^2, \dots, t_2^{(K+1)/2}\} \rightarrow \{A, B\}.$$

Before we begin examining eq, note that $\Pr \{\theta = 9\} = (K - 1) / K$, and this converges to 1 as $K \rightarrow \infty$.

This game of asymmetric information has a unique equilibrium (idea similar to Rubinstein's (1989) email game).

Claim 3.5.1 *This game has a unique Nash equilibrium (\hat{s}_1, \hat{s}_2) , and in this equilibrium, both players necessarily choose B, i.e., $\hat{s}_1(t_1) = \hat{s}_2(t_2) = B$ for all t_1 and t_2 .*

Proof. (by induction) Let s^* be a Nash equilibrium. Note first that $s_2^*(t_2^{(K+1)/2}) = B$ (since B is the unique best response at $t_2^{(K+1)/2}$).

Suppose $s_2^*(t_2^{\ell+1}) = B$. Then, since

$$\begin{aligned} \Pr \{t_2^{\ell+1} | t_1^\ell\} &= \Pr \{\{\omega_{2\ell+1}, \omega_{2\ell+2}\} | \{\omega_{2\ell}, \omega_{2\ell+1}\}\} \\ &= \frac{\Pr \{\omega_{2\ell+1}\}}{\Pr \{\omega_{2\ell}, \omega_{2\ell+1}\}} = \frac{1/K}{2/K} = \frac{1}{2}, \end{aligned}$$

$s_1^*(t_1^\ell) = B$ (the probability that 2 plays B is at least $1/2$). Moreover, if $s_1^*(t_1^\ell) = B$, then since

$$\begin{aligned} \Pr \{t_1^\ell | t_2^\ell\} &= \Pr \{\{\omega_{2\ell}, \omega_{2\ell+1}\} | \{\omega_{2\ell-1}, \omega_{2\ell}\}\} \\ &= \frac{\Pr \{\omega_{2\ell}\}}{\Pr \{\omega_{2\ell-1}, \omega_{2\ell}\}} = \frac{1}{2}, \end{aligned}$$

we also have $s_2^*(t_2^\ell) = B$. ■

The proof actually proves something a little stronger, that only profile that survives the iterated deletion of strictly dominated strategies involves both players always choosing B . Since B is the unique best response at $t_2^{(K+1)/2}$, any strategy s_2' satisfying $s_2'(t_2^{(K+1)/2}) = A$ is strictly dominated by the strategy \hat{s}_2' given by $\hat{s}_2'(t_2^{(K+1)/2}) = B$ and $\hat{s}_2'(t_2) = s_2'(t_2)$ for all other t_2 . We now proceed by iteratively deleting strictly dominated strategies. ★

The complete information version of the game with $\theta = 9$ has two strict equilibria. None the less, by making a small perturbation to the game by introducing a particular form of incomplete information, the result is stark, with only BB surviving, even in a state like ω_1 , where both players know the state to an some large finite order.

Carlsson and van Damme (1993) introduced the term *global games* to emphasize the importance of viewing the benchmark complete information game in a broader (global), i.e., perturbed, context. The term *global game* is now commonly understood to refer to a modelling that incorporates both a *richness* assumption on the uncertainty (so that each action is dominant for at least one value of the uncertainty) and *small noise* (as illustrated next).

Example 3.5.2 (Global Games) The stage game is as in example 3.5.1. We change the information structure: We now assume θ is uniformly distributed on the interval $[0, 20]$. For $\theta < 5$, B is strictly dominant, while if $\theta > 16$, A is strictly dominant.

Each player i receives a signal x_i , with x_1 and x_2 independently and uniformly drawn from the interval $[\theta - \varepsilon, \theta + \varepsilon]$ for $\varepsilon > 0$. A

pure strategy for player i is a function

$$s_i : [-\varepsilon, 20 + \varepsilon] \rightarrow \{A, B\}.$$

First observe that, for $x_i \in [\varepsilon, 20 - \varepsilon]$, player i 's posterior on θ is uniform on $[x_i - \varepsilon, x_i + \varepsilon]$. This is most easily seen as follows: Letting g be the density of θ and h be the density of x given θ , we immediately have $g(\theta) = \frac{1}{20}$ for all $\theta \in [0, 20]$ and

$$h(x | \theta) = \begin{cases} \frac{1}{2\varepsilon}, & \text{if } x \in [\theta - \varepsilon, \theta + \varepsilon], \\ 0, & \text{otherwise.} \end{cases}$$

Since

$$h(x | \theta) = \frac{f(x, \theta)}{g(\theta)},$$

where f is the joint density, we have

$$f(x, \theta) = \begin{cases} \frac{1}{40\varepsilon}, & \text{if } x \in [\theta - \varepsilon, \theta + \varepsilon] \text{ and } \theta \in [0, 20], \\ 0, & \text{otherwise.} \end{cases}$$

The marginal density for $x \in [\varepsilon, 20 - \varepsilon]$ is thus simply the constant function $\frac{1}{20}$, and so the density of θ conditional on an $x \in [\varepsilon, 20 - \varepsilon]$ is the constant function $\frac{1}{2\varepsilon}$ on the interval $[x - \varepsilon, x + \varepsilon]$.

Similar considerations show that for $x_i \in [\varepsilon, 20 - \varepsilon]$, player i 's posterior on x_j is symmetric around x_i with support $[x_i - 2\varepsilon, x_i + 2\varepsilon]$. Hence, $\Pr\{x_j > x_i | x_i\} = \Pr\{x_j < x_i | x_i\} = \frac{1}{2}$.

Claim 3.5.2 *For ε small, the game has an essentially unique Nash equilibrium (s_1^*, s_2^*) , given by*

$$s_i^*(x_i) = \begin{cases} A, & \text{if } x_i \geq 10\frac{1}{2}, \\ B, & \text{if } x_i < 10\frac{1}{2}. \end{cases}$$

Proof. We again apply iterated deletion of dominated strategies. Suppose $x_i < 5$. Then, player i 's conditional expected payoff from A is less than that from B irrespective of player j 's action, and so i plays B for $x_i < 5$ (as does j for $x_j < 5$). But then at $x_i = 5$, player i assigns at least probability $\frac{1}{2}$ to j playing B , and so i strictly

prefers B . Let x_i^* be the largest signal for which B is implied by iterated dominance (i.e., $x_i^* = \sup\{x'_i \mid B \text{ is implied by iterated strict dominance for all } x_i < x'_i\}$). By symmetry, $x_1^* = x_2^* = x^*$. At $x_i = x^*$, player i cannot strictly prefer B to A (otherwise, we can expand the set of signals for which iterated dominance implies B), and he assigns at least probability $\frac{1}{2}$ to j playing B . Hence, $x^* \geq 10\frac{1}{2}$.

Similarly, for $x_i > 16$, player i 's conditional expected payoff from A is greater than that from B irrespective of player j 's action, and so i plays A for $x_i > 16$ (as does j for $x_j > 16$). Let x_i^{**} be the smallest signal for which A is implied by iterated dominance (i.e., $x_i^{**} = \inf\{x'_i \mid A \text{ is implied by iterated strict dominance for all } x_i > x'_i\}$). By symmetry, $x_1^{**} = x_2^{**} = x^{**}$. At $x_i = x^{**}$, player i cannot strictly prefer A to B , and he assigns at least probability $\frac{1}{2}$ to j playing A . Hence, $x^{**} \leq 10\frac{1}{2}$.

But then

$$10\frac{1}{2} \leq x^* \leq x^{**} \leq 10\frac{1}{2}.$$

■

The iterated deletion argument connecting x_i in the dominance regions to values not in the dominance regions is often called an *infection argument*.

This idea is not dependent on the particular distributional assumptions made here. See Morris and Shin (2003) for details. ★

Remark 3.5.1 (CAUTION) Some people have interpreted the global games literature as solving the multiplicity problem, at least in some settings. There is in fact a stronger result: Weinstein and Yildiz (2007) show that “almost all” games have a unique rationalizable outcome (which of course implies a unique Nash equilibrium)!

Does this mean that we don't need to worry about multiplicity? Of course not: This is a result about robustness. The uniqueness of the rationalizable outcome is driven by similar ideas to that in example 3.5.1—“almost all” simply means that all information structures can be approximated by information structures allowing an infection argument. In order for a modeler to be confident that he

knows the unique rationalizable outcome, he needs to be confident of the information structure.



3.6 Problems

- 3.6.1. There are two firms, 1 and 2, producing the same good. The inverse demand curve is given by $P = \theta - q_1 - q_2$, where $q_i \in \mathbb{R}_+$ is firm i 's output. (Note that we are allowing negative prices.) There is demand uncertainty with nature determining the value of θ , assigning probability $\alpha \in (0, 1)$ to $\theta = 3$, and complementary probability $1 - \alpha$ to $\theta = 4$. Firm 2 knows (is informed of) the value of θ , while firm 1 is not. Finally, each firm has zero costs of production. As usual, assume this description is common knowledge. Suppose the two firms choose quantities simultaneously. Define a strategy profile for this game. Describe the Nash equilibrium behavior and interim payoffs (which may be unique).
- 3.6.2. Consider the following variant of a sealed bid auction in a setting of independent private values. The highest bidder wins, and pays a price determined as the weighted average of the highest bid and second highest bid, with weight $\alpha \in (0, 1)$ on the highest bid (ties are resolved by a fair coin). Suppose there are two bidders, with bidder i 's value v_i randomly drawn from the interval $[\underline{v}_i, \bar{v}_i]$ according to the distribution function F_i , with density f_i .
- What are the interim payoffs of player i ?
 - Suppose (σ_1, σ_2) is a Nash equilibrium of the auction, and assume σ_i is a strictly increasing and differentiable function, for $i = 1, 2$. Describe the pair of differential equations the strategies must satisfy.
 - Suppose v_1 and v_2 are uniformly and independently distributed on $[0, 1]$. Describe the differential equation a symmetric increasing and differentiable equilibrium bidding strategy must satisfy.
 - Solve the differential equation found in part 3.6.2(c). **[Hint: Conjecture a functional form.]**

- (e) For the assumptions under part 3.6.2(c), prove the strategy found in part 3.6.2(d) is a symmetric equilibrium strategy.
- 3.6.3. This question asks you to prove a claim made in Example 3.3.4 as follows:
- (a) Suppose $\bar{s}_2 < \bar{s}_1$, and set $\delta \equiv \Pr\{\bar{s}_2 < \sigma_1(t_1) \leq \bar{s}_1\} > 0$. Prove that there exists \tilde{s}_1 satisfying $\Pr\{\sigma_1(t_1) > \tilde{s}_1\} < \delta/2$. [**Hint:** This is trivial if $\bar{s}_1 < \infty$ (why?). The case where $\bar{s}_1 = \infty$ uses a basic continuity property of probability.]
 - (b) Show that a deviation by type $t_2 > 2\tilde{s}_1$ to a stopping time $s_2 > \tilde{s}_1$ (which implies that t_2 wins the war of attrition with probability of at least $\delta/2$) satisfying $s_2 < t_2/2$ is strictly profitable.
- 3.6.4. This question concerns Example 3.1.1, the Cournot game with incomplete information. The idea is to capture the possibility that firm 2 may know that firm 1 has low costs, c_L . This can be done as follows: Firm 1's space of uncertainty (types) is, as before, $\{c_L, c_H\}$, while firm 2's is $\{t_L, t_U\}$. Nature determines the types according to the distribution
- $$\Pr(t_1, t_2) = \begin{cases} \alpha, & \text{if } (t_1, t_2) = (c_L, t_L), \\ \beta, & \text{if } (t_1, t_2) = (c_L, t_U), \\ 1 - \alpha - \beta, & \text{if } (t_1, t_2) = (c_H, t_U), \end{cases}$$
- where $\alpha, \beta \in (0, 1)$ and $1 - \alpha - \beta > 0$. Firm 2's type, t_L or t_U , does not affect his payoffs (in particular, his cost is c_2 , as in Example 3.1.1). Firm 1's type is just his cost, c_1 .
- (a) What is the probability firm 2 assigns to $c_1 = c_L$ when his type is t_L ? When his type is t_U ?
 - (b) What is the probability firm 1 assigns to firm 2 knowing firm 1's cost? [This may depend on 1's type.]
 - (c) What is the Nash equilibrium of this game. Compare your analysis to that of Example 3.1.1.
- 3.6.5. The linear equilibrium of Example 3.3.5 is not the only equilibrium of the double auction.
- (a) Fix a price $p \in (0, 1)$. Show that there is an equilibrium at which, if trade occurs, then it occurs at the price p .

- (b) What is the probability of trade?
 - (c) At what p is the probability of trade maximized?
 - (d) Compare the expected gains from trade under these “fixed price” equilibria with the linear equilibrium of Example 3.3.5.
- 3.6.6. In Example 3.5.1, suppose that in the state ω_1 , $\theta = 20$, while in states ω_k , $2 \leq k \leq K - 1$, $\theta = 9$ and in state ω_K , $\theta = 4$. Suppose the information partitions are as in the example. In other words, apart from the probability distribution over Ω (which we have not yet specified), the only change is that in state ω_1 , $\theta = 20$ rather than 9.
- (a) Suppose the probability distribution over Ω is uniform (as in the lecture notes). What is the unique Nash equilibrium, and why? What is the unconditional probability of both players choosing A in the equilibrium?
 - (b) Suppose now the probability of ω_k is $4^{10-k}\alpha$, for $k = 1, \dots, 9$, where α is chosen so that $\sum_{k=1}^9 4^{10-k}\alpha = 1$. What is the unique Nash equilibrium, and why? What is the unconditional probability of both players choosing A in the equilibrium?

Chapter 4

Existence and Foundations for Nash Equilibrium¹

4.1 Existence of Nash Equilibria

Recall Nash eq are fixed points of best reply correspondence:

$$s^* \in \phi(s^*).$$

When does ϕ have a fixed point?

Theorem 4.1.1 (Kakutani's fixed point theorem) *Suppose $X \subset \mathbb{R}^m$ for some m and $F : X \rightrightarrows X$. Suppose*

1. *X is nonempty, compact, and convex;*
2. *F has nonempty convex-values (i.e., $F(x)$ is a convex set and $F(x) \neq \emptyset \forall x \in X$); and*
3. *F has closed graph: $(x^k, \hat{x}^k) \rightarrow (x, \hat{x}), \hat{x}^k \in F(x^k) \Rightarrow \hat{x} \in F(x)$.*

Then F has a fixed point.

Remark 4.1.1 The closed graph property is sometimes called upperhemicontinuity of the correspondence. But note that the closed

¹Copyright December 7, 2011 by George J. Mailath

graph property does not imply continuity of functions: the function $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = 1/x$ for $x > 0$ and $f(0) = 0$ has a closed graph but is not continuous (for more on this, see Ok (2007, §E.2)).

◆

Theorem 4.1.2 *Given a normal form game $G = \{(S_i, u_i) : i = 1, \dots, n\}$, if for all i ,*

1. S_i is a nonempty, convex, and compact subset of \mathbb{R}^k for some k , and
2. $u_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ is continuous in $s \in S_1 \times \dots \times S_n$ and quasiconcave in s_i ,

then G has a Nash equilibrium strategy profile.

Proof. Since u_i is continuous, the Maximum Theorem (MWG Theorem M.K.6) implies that ϕ_i has a closed graph.

The quasiconcavity of u_i implies that ϕ_i is convex-valued: For fixed $s_{-i} \in S_{-i}$, suppose $s'_i, s''_i \in \arg \max u_i(s_i, s_{-i})$. Then, from the quasiconcavity of u_i , for all $\alpha \in [0, 1]$,

$$u_i(\alpha s'_i + (1 - \alpha)s''_i, s_{-i}) \geq \min\{u_i(s'_i, s_{-i}), u_i(s''_i, s_{-i})\},$$

and so

$$u_i(\alpha s'_i + (1 - \alpha)s''_i, s_{-i}) \geq \max u_i(s_i, s_{-i})$$

so that $\alpha s'_i + (1 - \alpha)s''_i \in \arg \max u_i(s_i, s_{-i})$.

The theorem then follows from Kakutani's fixed point theorem by taking $X = S_1 \times \dots \times S_n$ and $F = (\phi_1, \dots, \phi_n)$. ■

Theorem 4.1.3 (Nash) *Every finite normal form game, $(S_i, u_i)_i$, has a mixed strategy Nash equilibrium.*

Proof. Define $X_i = \Delta(S_i)$ and $X \equiv \prod_i X_i$. Then $X \subset \mathbb{R}^{\sum |S_i|}$ and X is nonempty, convex and compact.

In this case, rather than appealing to the maximum theorem, it is an easy (and worthwhile!) exercise to prove that ϕ_i is convex-valued and has a closed graph. ■

In some applications, we need an infinite dimensional version of Kakutani's fixed point theorem.

Theorem 4.1.4 (Fan-Glicksberg fixed point theorem) *Suppose X is a nonempty compact convex subset of a locally convex Hausdorff space, and suppose $F : X \rightrightarrows X$. Suppose*

1. *F has nonempty convex-values (i.e., $F(x)$ is a convex set and $F(x) \neq \emptyset \forall x \in X$); and*
2. *F has closed graph: $(x^k, \hat{x}^k) \rightarrow (x, \hat{x}), \hat{x}^k \in F(x^k) \Rightarrow \hat{x} \in F(x)$.*

Then F has a fixed point.

In particular, every normed vector space is locally convex Hausdorff. Locally convex Hausdorff spaces generalize many of the nice properties of normed vector spaces. This generalization is needed in the following theorem, since the spaces are typically not normed.

Theorem 4.1.5 *Suppose X is a compact metric space. Then the space of probability measures on X is a nonempty compact convex subset of a locally convex Hausdorff space. Moreover, if $f : X \rightarrow \mathbb{R}$ is a continuous function, then $\int f d\mu$ is a continuous function of μ .*

Corollary 4.1.1 *Suppose S_i is a compact subset of a finite dimensional Euclidean space \mathbb{R}^{m_i} , and suppose $u_i : S \rightarrow \mathbb{R}$ is continuous. Then $\{(S_i, u_i)_i\}$ has a Nash equilibrium in mixed strategies.*

Proof. The proof mimics that of Theorem 4.1.3. ■

Example 4.1.1 (An example of mixed strategies) Return to first price sealed bid auction with independent private values (Example 3.3.1), but assume that the value for each bidder is drawn independently from a uniform distribution on the two point set $\{\underline{v}, \bar{v}\}$, with $\underline{v} < \bar{v}$.

This game has no equilibrium in pure strategies (see problem 4.3.3). Even though the game has discontinuous payoffs (why?), it does have an equilibrium in mixed strategies.

Suppose bidder 2 follows the strategy σ_2 of bidding \underline{v} if $v_2 = \underline{v}$ (why is that a reasonable “guess”?), and of bidding according to the distribution function $F_2(b)$ if $v_2 = \bar{v}$. Then, assuming there are no atoms in F_2 , player 1 has interim payoffs from $b > \underline{v}$ given by

$$\begin{aligned} U_1(b, \bar{v}, \sigma_2) &= \frac{1}{2}(\bar{v} - b) + \frac{1}{2}(\bar{v} - b)F_2(b) \\ &= \frac{1}{2}(\bar{v} - b)(1 + F_2(b)). \end{aligned}$$

Note that the minimum of the support of F_2 is given by \underline{v} (why?). Denote the maximum of the support by \bar{b} .

Suppose 1 is also randomizing over the set of bids, $(\underline{v}, \bar{b}]$. The indifference condition requires that 1 is indifferent over all $b \in (\underline{v}, \bar{b}]$. The bid $b = \underline{v}$ is excluded because there is a positive probability of a tie at \underline{v} (from the low value bidder) and so it cannot be optimal for \bar{v} to bid \underline{v} . That is, for all $\varepsilon > 0$,

$$\frac{1}{2}(\bar{v} - b)(1 + F_2(b)) = U_1(\underline{v} + \varepsilon, \bar{v}, \sigma_2),$$

and

$$U_1(\underline{v} + \varepsilon, \bar{v}, \sigma_2) = \frac{1}{2}(\bar{v} - \underline{v} - \varepsilon)(1 + F_2(\underline{v} + \varepsilon)).$$

Since $\lim_{\varepsilon \rightarrow 0} F_2(\underline{v} + \varepsilon) = F_2(\underline{v}) = 0$ (where the first equality follows from the continuity of probabilities and the second equality follows from the assumption of no atoms), and so

$$(\bar{v} - b)(1 + F_2(b)) = \bar{v} - \underline{v},$$

yielding

$$F_2(b) = \frac{b - \underline{v}}{\bar{v} - \underline{v}}, \quad \forall b \in (0, \bar{b}].$$

Note that $F_2(\underline{v}) = 0$ as required for a distribution function. Moreover, $F_2(\bar{b}) = 1$ implies $\bar{b} = (\bar{v} + \underline{v})/2$.

It is straightforward to verify that the symmetric profile in which each bidder bids \underline{v} if $v = \underline{v}$, and according to the distribution function $F(b) = (b - \underline{v})/(\bar{v} - \underline{b})$ if $v = \bar{v}$. ★

Remark 4.1.2 (More formal treatment of mixed strategies) For continuum action spaces (such as auctions), a mixed strategy for a player i is a probability distribution on \mathbb{R} (which we can denote F_i). Player 1's expected payoff from an action b_1 is

$$\int u_1(s_1, s_2) dF_2(s_2).$$

As an aside, note that this notation (which may not be familiar to all of you) covers all relevant possibilities: If the mixed strategy of player 2 has a countable support $\{s^k\}$ with action s^k having probability $\sigma_2(s^k) > 0$ (the distribution is said to be *discrete* in this case), we have

$$\int u_1(s_1, s_2) dF_2(s_2) = \sum_{s^k} u_1(s_1, s^k) \sigma_2(s^k).$$

Note that $\sum_{s^k} \sigma_2(s^k) = 1$. Any single action receiving strictly positive probability is called an *atom*. If the distribution function describing player 2's behavior has a density f_2 , then

$$\int u_1(s_1, s_2) dF_2(s_2) = \int u_1(s_1, s_2) f_2(s_2) ds_2.$$

Finally, combinations of distributions with densities on part of the support and atoms elsewhere, as well as more esoteric possibilities (that are almost never relevant) are also covered.

Suppose F_1 is a best reply for player 1 to player 2's strategy F_2 . Then

$$\iint u_1(s_1, s_2) dF_2(s_2) dF_1(s_1) = \max_{s_1} \int u_1(s_1, s_2) dF_2(s_2).$$

Observe first that if F_1 is discrete with support $\{s^k\}$ and action s^k having probability $\sigma_1(s^k) > 0$, then

$$\iint u_1(s_1, s_2) dF_2(s_2) dF_1(s_1) = \sum_{s^k} \int u_1(s^k, s_2) dF_2(s_2) \sigma_1(s^k),$$

and so we immediately have

$$\sum_{s^k} \int u_1(s^k, s_2) dF_2(s_2) \sigma_1(s^k) = \max_{s_1} \int u_1(s_1, s_2) dF_2(s_2)$$

and so, for all s^k ,

$$\int u_1(s^k, s_2) dF_2(s_2) = \max_{s_1} \int u_1(s_1, s_2) dF_2(s_2).$$

This is just the familiar statement that player 1 is indifferent over all actions in his support, and each such action maximizes his payoff against F_2 .

What is the appropriate version of this statement for general F_1 ? The key observation is that zero probability sets don't matter. Thus, the statement is: Let \hat{S}_1 be the set of actions that are suboptimal against F_2 , i.e.,

$$\hat{S}_1 = \left\{ \hat{s}_1 : \int u_1(\hat{s}_1, s_2) dF_2(s_2) < \max_{s_1} \int u_1(s_1, s_2) dF_2(s_2) \right\}.$$

Then, the set \hat{S}_1 is assigned zero probability by any best response F_1 . [If such a set received positive probability under some F_1 , then F_1 could not be a best reply, since expected payoffs are clearly increased by moving probability weight off the set \hat{S}_1 .]

In most applications, the set \hat{S}_1 is disjoint from the support of F_1 , in which case player 1 is indeed indifferent over all actions in his support, and each such action maximizes his payoff against F_2 . However, Example 4.1.1 is an example where \hat{S}_1 includes one point of the support.

◆

Example 4.1.2 (Examples of nonexistence) The following example of nonexistence in a game of complete information is due to Sion and Wolfe (1957): $S_1 = S_2 = [0, 1]$ and payoff are

$$u_1(s_1, s_2) = \begin{cases} -1, & \text{if } s_1 < s_2 < s_1 + \frac{1}{2}, \\ 0, & \text{if } s_1 = s_2 \text{ or } s_2 = s_1 + \frac{1}{2}, \\ +1, & \text{otherwise,} \end{cases}$$

and $u_2(s_1, s_2) = -u_1(s_1, s_2)$ (i.e., the game is *zero sum*). The proof of nonexistence is beyond the scope of this course.

I am almost certain that a simpler example of nonexistence in games of incomplete information is provided by the following example of a private value first price auction with ties broken using a fair coin. Bidder 1 has value 3 and bidder 2 has value 3 with probability $\frac{1}{2}$ and value 4 with probability $\frac{1}{2}$. An intuition for nonexistence can be obtained by observing that since there is positive probability that 2 has value 3, bidder 1 will not bid over 3 (since he may win). Moreover, bidder 1 cannot randomize on bids below 3 because there is positive probability of bidder 2 having value 3 (and so responding to bidder 1's randomization). But if bidder 1 bids 3, then bidder of value 4 does not have a best reply. Existence is restored if ties are broken in favor of bidder 2.

Suppose now we assume that there is a joint distribution over bidder valuations, so that

- with prob $\frac{1}{2}$, $v_1 = 3$ and $v_2 = 3$ with prob $\frac{1}{2}$ and 4 with prob $\frac{1}{2}$; and
- with prob $\frac{1}{2}$, $v_2 = 3$ and $v_1 = 3$ with prob $\frac{1}{2}$ and 4 with prob $\frac{1}{2}$.

Note that the valuations are *not* independent. Moreover, existence would only be restored by using a tie breaking rule that awarded the item to the highest value bidder. For (much) more on this, see Jackson, Simon, Swinkels, and Zame (2002) (note that the similar example given in Jackson, Simon, Swinkels, and Zame (2002) is incorrect). ★

Existence results that replace continuity assumptions on payoffs with complementarity or supermodularity assumptions are of increasing importance. For an excellent introduction, see chapter 7 in Vohra (2005).

4.2 Foundations for Nash Equilibrium

A good book length treatment of the general topic covered here is Fudenberg and Levine (1998). For much more on evolution, see

Samuelson (1997) and Weibull (1995) (or Mailath (1998) for a longer nontechnical introduction).

Consider two players repeatedly playing a 2×2 matrix game.

Neither player knows how opponent will play.

Suppose each player is Bayesian in the following way: each believes the opponent plays the first strategy in each period with the same fixed probability, and the prior on this probability is uniform on $[0, 1]$.

After each round of play, players update to posterior.

This behavior, while Bayesian, is not “rational”: players make assumptions on opponents that they should know, by introspection, are false.

Rational Learning: Bayesian learning with a prior and likelihood that encompasses the “truth.”

Must good Bayesians be rational learners?

If so, not all Savage-Bayesians are good Bayesians.

4.2.1 Boundedly Rational Learning

Motivated by concerns over

- informational (knowledge) requirements of traditional analyses.
- Coordination: beliefs about endogenous variables are correct. where does information about endogenous variables (that leads to coordinated behavior) come from?
- rationality (computational) requirements and
- (?) Bayesian paradigm.
 - Can agents formulate a sufficiently general model to encompass the “truth”? Updating on small (prior) probability events. At which point should agents question their view of the world?

Rationality and coordination are distinct and independent ideas.

Common thread:

- behavior is myopic,
- dynamic analysis, focusing on asymptotic properties, and
- focus on interaction of learning with evolution of system.

4.2.2 Social Learning (Evolution)

Example 4.2.1 Large population of players randomly and repeatedly matched (paired) to play the same game:

	<i>A</i>	<i>B</i>
<i>A</i>	1	0
<i>B</i>	0	1

No *role identification*, so the payoff represents the payoff to a player who chooses the row action, when facing the column action. If α is fraction of population playing *A*, then

$$u(A; \alpha) = \alpha,$$


and $u(B; \alpha) = 1 - \alpha.$

Then

$$\dot{\alpha} = \frac{d\alpha}{dt} > 0 \quad \text{iff} \quad u(A; \alpha) > u(B; \alpha) \quad \text{iff} \quad \alpha > \frac{1}{2},$$

and

$$\dot{\alpha} < 0 \quad \text{iff} \quad u(A; \alpha) < u(B; \alpha) \quad \text{iff} \quad \alpha < \frac{1}{2}.$$



The symmetric mixed strategy equilibrium $\frac{1}{2} \circ A + \frac{1}{2} \circ B$ is unstable.

★

Example 4.2.2 Large population paired to play:

	A	B
A	1	2
B	2	1

No role identification (so that AB and BA are infeasible). If α is fraction of population playing A , then

$$u(A; \alpha) = \alpha + 2(1 - \alpha) = 2 - \alpha,$$


and $u(B; \alpha) = 2\alpha + 1 - \alpha = 1 + \alpha.$

Then

$$\dot{\alpha} > 0 \quad \text{iff} \quad 2 - \alpha > 1 + \alpha \quad \text{iff} \quad \frac{1}{2} > \alpha,$$

and

$$\dot{\alpha} < 0 \quad \text{iff} \quad 2 - \alpha < 1 + \alpha \quad \text{iff} \quad \frac{1}{2} < \alpha.$$



★

Let S denote a finite set of strategies in a symmetric game. In the above examples, $S = \{A, B\}$. Payoff to playing the strategy $s \in S$ against an opponent who plays $r \in S$ is $u(s, r)$.

State of society is $\sigma \in \Delta(S)$. Expected payoff to s when state of society is σ is

$$u(s, \sigma) = \sum_r u(s, r) \sigma(r).$$

Dynamics:

$$F : \Delta(S) \times \mathbb{R}_+ \rightarrow \Delta(S),$$

with

$$F(\sigma, t' + t) = F(F(\sigma, t'), t).$$

Definition 4.2.1 A state σ^* is a rest (or stationary) point of F if

$$\sigma^* = F(\sigma^*, t) \quad \forall t.$$

A rest point σ^* is asymptotically stable under F if there exists $\varepsilon > 0$ such that if $|\sigma' - \sigma^*| < \varepsilon$, then $F(\sigma', t) \rightarrow \sigma^*$.

Assume F is continuously differentiable in all its arguments (on boundaries, the appropriate one-sided derivatives exist and are continuous).

Interpretation: if population strategy profile is σ' at t' , then at time $t' + t$ it will be $F(\sigma', t)$. Write $\dot{\sigma}$ for $\partial F(\sigma, t)/\partial t|_{t=0}$.

Note that

$$\sum_s \sigma(s) = 1 \Rightarrow \sum_s \dot{\sigma}(s) = 0.$$

Definition 4.2.2 F is a myopic adjustment dynamic if $\forall \sigma, s, r \in S$ satisfying $\sigma(s), \sigma(r) > 0$,

$$u(s, \sigma) > u(r, \sigma) \Rightarrow \dot{\sigma}(s) > \dot{\sigma}(r).$$

Theorem 4.2.1 Suppose F is a myopic adjustment dynamic.

1. If σ^* is asymptotically stable under F , then it is a symmetric Nash equilibrium.
2. If σ^* is a strict Nash equilibrium, then σ^* is asymptotically stable under F .

Proof.

1. Left as an exercise.
2. Suppose σ^* is a strict Nash equilibrium. Then σ^* is a pure strategy s and $u(s, s) > u(r, s)$ for all $r \neq s$. This implies that there exists $\bar{\varepsilon} > 0$ such that for all σ satisfying $\sigma(s) > 1 - \varepsilon$,

$$u(s, \sigma) > u(r, \sigma), \quad \forall r \neq s.$$

Suppose $1 > \sigma(s) > 1 - \varepsilon$ and $\sigma(r) > 0$ for all r . Myopic adjustment implies $\dot{\sigma}(s) > \max\{\dot{\sigma}(r) : r \neq s\}$, and so $\dot{\sigma}(s) > 0$ (since $\sum_{r \in S} \dot{\sigma}(r) = 0$).

Consider now σ satisfying $1 > \sigma(s) > 1 - \varepsilon$ with $\sigma(r) = 0$ for some r . Since $\dot{\sigma}$ is continuous in σ (since F is continuously differentiable, including on the boundaries), $\dot{\sigma}(s) \geq 0$ and $\dot{\sigma}(s) \geq \dot{\sigma}(r)$. Suppose $\dot{\sigma}(s) = 0$ (so that $\dot{\sigma}(r) \leq 0$). Then, $\dot{\sigma}(r') < 0$ for all r' satisfying $\sigma(r') > 0$ and so $\dot{\sigma}(r) > 0$, a contradiction.

Hence, if $1 > \sigma(s) > 1 - \varepsilon$, then $\dot{\sigma}(s) > 0$. Defining $\sigma^t \equiv F(\sigma, t) \in \Delta(S)$, this implies $\sigma^t(s) > 1 - \bar{\varepsilon}$ for all t , and so $\sigma^t(s) \rightarrow 1$.

■

There are examples of myopic adjustment dynamics that do *not* eliminate strategies that are iteratively strictly dominated. Stronger conditions (such as *aggregate monotonicity*) are needed—see Fudenberg and Levine (1998). These conditions are satisfied by the *replicator dynamic*:

Biological model, payoffs are reproductive fitness (normalize payoffs so $u(s, r) > 0$ for all $s, r \in S$). At the end of each period, each agent is replaced by a group of agents who play the same strategy, with the size of the group given by the payoff (fitness) of the agent. Let $x_s(t)$ be the size of the population playing s in period t . Then,

$$x_s(t+1) = x_s(t)u(s, \sigma^t),$$

where

$$\sigma^t(s) = \frac{x_s(t)}{\sum_r x_r(t)} \equiv \frac{x_s(t)}{\bar{x}(t)}.$$

Then, since $\bar{x}(t+1) = \bar{x}(t)u(\sigma^t, \sigma^t)$,

$$\sigma^{t+1}(s) = \sigma^t(s) \frac{u(s, \sigma^t)}{u(\sigma^t, \sigma^t)},$$

so the difference equation is

$$\sigma^{t+1}(s) - \sigma^t(s) = \sigma^t(s) \frac{u(s, \sigma^t) - u(\sigma^t, \sigma^t)}{u(\sigma^t, \sigma^t)}.$$

Thus, $\sigma^{t+1}(s) > (<) \sigma^t(s)$ iff $u(s, \sigma^t) > (<) u(\sigma^t, \sigma^t)$. In continuous time, this is

$$\dot{\sigma}(s) = \sigma(s) \frac{u(s, \sigma) - u(\sigma, \sigma)}{u(\sigma, \sigma)}.$$

This has the same trajectories as

$$\dot{\sigma}(s) = \sigma(s) [u(s, \sigma) - u(\sigma, \sigma)].$$

Note that under the replicator dynamic, every pure strategy profile is a rest point: if $\sigma(s) = 0$ then $\dot{\sigma}(s) = 0$ *even when* $u(s, \sigma) > u(\sigma, \sigma)$.

Idea extends in straightforward fashion to games with role identification. In that case, we have

$$\dot{\sigma}_i(s_i) = \sigma_i(s_i) [u_i(s_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i})].$$

Example 4.2.3 (Domination) The game is:

	<i>L</i>	<i>R</i>
<i>T</i>	1, 1	1, 0
<i>B</i>	1, 1	0, 0

Let p^t be the fraction of row players choosing *T*, while q^t is the fraction of column players choosing *L*. The replicator dynamics are

$$\begin{aligned} \dot{p} &= p(1-p)(1-q) \\ \text{and} \quad \dot{q} &= q(1-q). \end{aligned}$$

The phase diagram is illustrated in Figure 4.2.1.² No rest point is asymptotically stable. ★

Example 4.2.4 (Simplified ultimatum game) In the simplified ultimatum game, the proposer offer either an equal split, or a small payment. The responder only responds to the small payment (he must accept the equal split). The extensive form is given in Figure 4.2.2. The normal form is

²Note that the phase diagram in Mailath (1998, Figure 11) is incorrect.




Figure 4.2.1: The phase diagram for the domination example.




Figure 4.2.2: The extensive form of the simplified ultimatum game.

	N	Y
equal split	50, 50	50, 50
small offer	0, 0	80, 20

Let p be the fraction of row players choosing equal split, while q is the fraction of column players choosing N .

The subgame perfect profile is $(0, 0)$. There is another Nash outcome, given by the row player choosing equal division. The set of Nash equilibrium yielding this outcome is $N = \{(1, q) : 3/8 \leq q\}$.

The replicator dynamics are

$$\begin{aligned} \dot{p} &= p(1-p)(80q-30) \\ \text{and} \quad \dot{q} &= -20q(1-q)(1-p). \end{aligned}$$

Note that $\dot{p} > 0$ if $q > 3/8$ and $\dot{p} < 0$ if $q < 3/8$, while $\dot{q} < 0$ for all (p, q) .

The phase diagram is illustrated in Figure 4.2.3. The subgame perfect equilibrium B is the only asymptotically stable rest point.

In the presence of drift, the dynamics are now given by

$$\begin{aligned} \dot{p} &= p(1-p)(80q-30) + \delta_1\left(\frac{1}{2} - p\right) \\ \text{and} \quad \dot{q} &= -20q(1-q)(1-p) + \delta_2\left(\frac{1}{2} - q\right). \end{aligned}$$

For small δ_i , with $\delta_1 \ll \delta_2$, the dynamics have two asymptotically stable rest points, one near B and one near A (see Samuelson (1997, chapter 5)). ★

4.2.3 Individual learning

Fix n -player finite strategic form game, $G = (S, u)$, $S \equiv S_1 \times \cdots \times S_n$, $u : S \rightarrow \mathbb{R}^n$.

Players play G^∞ .

History $h^t \equiv (s^0, \dots, s^{t-1}) \in H^t \equiv S^t$.

Assessments $\mu_i^t : H^t \rightarrow \Delta(S_{-i})$.

Behavior rule $\phi_i^t : H^t \rightarrow \Delta(S_i)$.




Figure 4.2.3: The phase diagram for the simplified ultimatum example. A is the nonsubgame perfect equilibrium $(1, 3/8)$, and B is the subgame perfect equilibrium.

Definition 4.2.3 ϕ_i is myopic with respect to μ_i if, $\forall t$ and h^t , $\phi_i^t(h^t)$ maximizes $u_i(\sigma_i, \mu_i^t(h^t))$.

Definition 4.2.4 μ_i is adaptive if, $\forall \varepsilon > 0$ and t , $\exists T(\varepsilon, t)$ s.t. $\forall t' > T(\varepsilon, t)$ and $h^{t'}$, $\mu_i^{t'}(h^{t'})$ puts no more than ε probability on pure strategies not played by $-i$ between t and t' in $h^{t'}$.

Examples:

- fictitious play (play best reply to empirical dsn of history)
- Cournot dynamics
- exponential weighting of past plays

Rules out rationalizability-type sophisticated analysis. Notion of “adaptive” does not impose any restrictions on the relative weight on strategies that are not excluded.

Definition 4.2.5 $h \equiv (s^0, s^1, \dots)$ is compatible with ϕ if s_i^t is in the support of $\phi_i^t(h^t)$, $\forall i$ and t .

Theorem 4.2.2 Suppose (s^0, s^1, \dots) is compatible with behavior that is myopic with respect to an adaptive assessment.

1. There exists T s.t. $s^t \in \bar{S} \forall t \geq T$, where \bar{S} is the result of the iterative deletion of all strictly dominated strategies (= rationalizable if $n = 2$).
2. If $\exists T$ s.t. $s^t = s^* \forall t > T$, then s^* is a (pure-strategy) Nash equilibrium of G .

Proof.

1. Let S_i^k denote the set of player i 's strategies after k rounds of deletions of strictly dominated strategies. Since S is finite, there exists $K < \infty$ such that $\bar{S} = S^K$. Proof proceeds by induction.

There exists T s.t. $s^t \in S^1 \forall t \geq T$: Any $s_i \notin S_i^1$ is not a best reply to any beliefs, myopia implies never chosen.

Suppose $\exists T$ s.t. $s^t \in S^k \forall t \geq T$. If $s_i \notin S_i^{k+1}$, then s_i is not a best reply to any beliefs with support in S_{-i}^k . But then $\exists \varepsilon > 0$ s.t. s_i is not a best reply to any belief $\mu_i \in \Delta(S_{-i})$ satisfying $\mu_i(S_{-i}^k) > 1 - \varepsilon$. (Exercise: calculate the bound on ε .) Since assessments are adaptive, $\exists T' > T$ s.t. $\mu_i^t(h^t)(S_{-i}^k) > 1 - \varepsilon$ for all $t > T$. Since behavior is myopic, $s^t \in S^{k+1} \forall t \geq T'$.

2. Suppose $\exists T$ s.t. $s^t = s^* \forall t > T$ and s^* is not a Nash equilibrium of G . Then $\exists i$ and $s'_i \in S_i$ s.t. $u_i(s'_i, s_{-i}^*) > u_i(s^*)$. There exists $\varepsilon > 0$ s.t. $u_i(s'_i, \sigma_{-i}) > u_i(s_i^*, \sigma_{-i})$ if $\sigma_{-i}(s^*) > 1 - \varepsilon$. But then adaptive assessments with myopic behavior implies $s_i^t \neq s_i^*$ for t large, a contradiction. ■

Stronger results on convergence (such as to mixed strategy equilibria) require more restrictions on assessments. For more, see Fudenberg and Levine (1998).

Convergence of beliefs need *not* imply convergence in behavior. For example, in matching pennies, empirical distribution converges to $(\frac{1}{2}, \frac{1}{2})$, but always play pure strategy.

4.3 Problems

- 4.3.1. A two-player normal form game G is *zero sum* if $u_1(s) = -u_2(s)$ for all $s \in S$. Suppose G is a finite zero-sum game.

- (a) Suppose $f : X \times Y \rightarrow \mathbb{R}$ is a continuous function and X and Y are compact subsets of \mathbb{R} . Prove that

$$\max_{x \in X} \min_{y \in Y} f(x, y) \leq \min_{y \in Y} \max_{x \in X} f(x, y).$$

Give an example showing that the inequality can hold strictly (it suffices to do this for X and Y each only containing two points—recall matching pennies from Section 2.4.1).

- (b) von Neumann's celebrated *Minmax Theorem* states the following equality:

$$\max_{\sigma_1 \in \Delta(S_1)} \min_{\sigma_2 \in \Delta(S_2)} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Delta(S_2)} \max_{\sigma_1 \in \Delta(S_1)} u_1(\sigma_1, \sigma_2).$$

Deduce this equality from Theorem 4.1.3.³

- (c) Prove that (σ_1^*, σ_2^*) is a Nash equilibrium of G if and only if σ_i^* is a security strategy for player i , and that player i 's security level \underline{v}_i is given by i 's payoff in any Nash equilibrium. (Compare with problem 2.6.11.)
 - (d) Prove the following generalization of Problem 2.6.11(b): Suppose a two-player normal form game has a unique Nash equilibrium, and each player's Nash equilibrium strategy and security strategy are both completely mixed. Prove that each player's security level is given by his/her Nash equilibrium payoff.
- 4.3.2. Consider a first price sealed bid auction with private values. There are two bidders with values $v_1 < v_2$. These values are common knowledge. Prove that this auction has no pure strategy equilibrium. Characterize the set of mixed strategy equilibria. [**Hint:** In these equilibria, bidder 2 plays a pure strategy and wins with probability 1.]
- 4.3.3. This question asks you to fill in the details of Example 4.1.1.
- (a) Prove that in any equilibrium, any bidder with value \underline{v} must bid \underline{v} .
 - (b) Prove that there is no equilibrium in pure strategies.
 - (c) Prove that in any mixed strategy equilibrium, the minimum of the support of F_2 is given by \underline{v} .
 - (d) Prove that it is not optimal for \bar{v} to bid \underline{v} .
 - (e) Prove that the symmetric profile in which each bidder bids \underline{v} if $v = \underline{v}$, and according to the distribution function $F(b) = (b - \underline{v})/(\bar{v} - \underline{v})$ if $v = \bar{v}$ is a Nash equilibrium.
- 4.3.4. Consider the following variant of a sealed bid auction: There are two bidders who each value the object at v , and simultaneously submit bids. As usual, the highest bid wins and in the event of a tie, the

³ von Neumann's original argument (1928), significantly predates Nash's existence theorem, and the result is true more generally. There are elementary proofs of the minmax theorem (based on the basic separating hyperplane theorem) that do not rely on a fixed point theorem. See, for example, Owen (1982, §II.4) for the finite dimensional case, and Ben-El-Mechaiekh and Dimand (2011) for the general case.

object is awarded on the basis of a fair coin toss. But now *all* bidders pay their bid. (This is an *all-pay auction*.)

- (a) Formulate this auction as a normal form game.
- (b) Show that there is no equilibrium in pure strategies.
- (c) This game has an equilibrium in mixed strategies. What is it? (You should verify that the strategies do indeed constitute an equilibrium).

4.3.5. **[Hard]** Prove (or provide a counterexample to the claim that) the following game does not have an equilibrium in pure or mixed strategies: The game is a private-value sealed bid auction of incomplete information with ties broken using a fair coin. Bidder 1 has value 3 for sure and bidder 2 has value 3 with probability $\frac{1}{2}$ and value 4 with probability $\frac{1}{2}$.

4.3.6. Prove that the phase diagram for example 4.2.3 is as portrayed in Figure 4.2.1. [This essentially asks you to give an expression for dq/dp .]

4.3.7. Prove that if σ^* is asymptotically stable under a myopic adjustment dynamic defined on a game with no role identification, then it is a symmetric Nash equilibrium.

4.3.8. Suppose $F : \Delta(S) \times \mathbb{R}_+ \rightarrow \Delta(S)$ is a dynamic on the strategy simplex with F continuously differentiable (including on the boundaries). Suppose that if

$$\eta < \sigma(s) < 1,$$

for some $\eta \in (0, 1)$, then

$$\dot{\sigma}(s) > 0,$$

where

$$\dot{\sigma} \equiv \left. \frac{\partial F(\sigma, t)}{\partial t} \right|_{t=0}.$$

Fix σ^0 satisfying $\sigma^0(s) > \eta$. Prove that

$$\sigma^t(s) \rightarrow 1,$$

where $\sigma^t \equiv F(\sigma^0, t)$.

- 4.3.9. Suppose a large population of players are randomly paired to play the game (where the payoffs are to the row player)

	<i>A</i>	<i>B</i>	<i>C</i>
<i>A</i>	1	1	0
<i>B</i>	0	1	1
<i>C</i>	0	0	1

(such a game is said to have *no role identification*). Let α denote the fraction of the population playing *A*, and γ denote the fraction of the population playing *C* (so that $1 - \alpha - \gamma$ is the fraction of the population playing *B*). Suppose the state of the population adjusts according to the continuous time replicator dynamic.

- Give an expression for $\dot{\alpha}$ and for $\dot{\gamma}$.
- Describe all the rest points of the dynamic.
- Describe the phase diagram in the space $\{(\alpha, \gamma) \in \mathbf{R}_+^2 : \alpha + \gamma \leq 1\}$. Which of the rest points are asymptotically stable?

Chapter 5

Dynamic Games and Sequential Equilibria¹

5.1 Sequential Rationality

Example 5.1.1 (Selten's horse)




Figure 5.1.1: Selten's horse.

¹Copyright December 7, 2011 by George J. Mailath

Let (p_1, p_2, p_3) denote the mixed strategy profile where

$$\begin{aligned}\Pr(I \text{ plays } A) &= p_1, \\ \Pr(II \text{ plays } a) &= p_2, \text{ and} \\ \Pr(III \text{ plays } L) &= p_3.\end{aligned}$$

Consider the Nash equilibrium profile $(0, 1, 0)$ (i.e., DaR). This profile is subgame perfect, and yet player II is not playing *sequentially rationally*. It is also *not* trembling hand perfect (Definition 2.5.1): playing a is not optimal against any mixture close to DR .

The only trembling hand perfect equilibrium outcome is Aa . The set of Nash equilibria with this outcome is $\{(1, 1, p_3) : \frac{3}{4} \leq p_3 \leq 1\}$. In these equilibria, player III 's information set is not reached, and so the profile cannot be used to obtain beliefs for III . However, each Nash equilibrium is trembling hand perfect: Fix an equilibrium $(1, 1, p_3)$. Suppose first that $p_3 \in [\frac{3}{4}, 1)$ (so that $p_3 \neq 1$!) and consider the completely mixed profile

$$\begin{aligned}p_1^n &= 1 - \frac{1}{n}, \\ p_2^n &= 1 - \frac{2}{(n-1)}, \\ \text{and } p_3^n &= p_3.\end{aligned}$$

Note that $p_1^n, p_2^n \rightarrow 1$ as $n \rightarrow \infty$. Suppose $n \geq 4$. Easy to verify that both I and II are playing optimally against the mixed profile in $(1, 1, p_3)$. What about III ? The probability that III is reached is

$$\frac{1}{n} + \frac{(n-1)}{n} \times \frac{2}{(n-1)} = 3/n,$$

and so the induced beliefs for III at his information set assign probability $\frac{1}{3}$ to the left node and $\frac{2}{3}$ to the right. Player III is therefore indifferent and so willing to randomize.

The same argument shows that $(1, 1, 1)$ is trembling hand perfect, using the trembles

$$p_1^n = 1 - \frac{1}{n},$$

$$p_2^n = 1 - \frac{2}{(n-1)},$$

and $p_3^n = 1 - \frac{1}{n}.$

Indeed, any sequence of trembles satisfying $p_1^n \rightarrow 1$, $p_2^n \rightarrow 1$, and $p_3^n \rightarrow 1$ will work, providing

$$\limsup_{n \rightarrow \infty} \frac{(1 - p_1^n)}{(1 - p_1^n p_2^n)} \leq \frac{1}{3}.$$

Note in particular it is not even necessary for $(1 - p_1^n)/(1 - p_1^n p_2^n)$ to have a well-defined limit. ★

Definition 5.1.1 *A system of beliefs μ in a finite extensive form is a specification of a probability distribution over the decision nodes in every information set, i.e., $\mu : X \rightarrow [0, 1]$ such that*

$$\sum_{x \in h} \mu(x) = 1, \quad \forall h.$$

Note that $\mu \in \prod_{h \in \cup_i H_i} \Delta(h)$, a compact set.

We interpret μ as describing player beliefs, in particular, if h is player i 's information set, then μ describes i 's beliefs over the nodes in h .

Let $\mathbf{P}^{\rho, b}$ denote the probability distribution on Z implied by the behavior profile b (and nature ρ).

Example 5.1.2 Consider the profile (LR, UD) in the game displayed in Figure 5.1.2. The label $[p]$ indicates that the player owning that information set assigns probability p to the labeled node. The induced distribution $\mathbf{P}^{\rho, b}$ is $p \circ z_1 + (1 - p) \circ z_8$. ★

The expected payoff to player i is (recalling Definition 1.3.1)

$$E[u_i | b] \equiv \sum_{z \in Z} u_i(z) \mathbf{P}^{\rho, b}(z).$$

Let $Z(h) = \{z \in Z : \exists x \in H, x < z\}$. Let $\mathbf{P}^{\mu, b}(\cdot | h)$ denote the probability distribution on $Z(h)$ implied by $\mu \in \Delta(h)$ and the behavior profile b (interpreted as describing behavior at information




Figure 5.1.2: Game for Example 5.1.2.

set h and any that could be reached from h . By setting $\mathbf{P}^{\mu,b}(z|h) = 0$ for all $z \notin Z(h)$, $\mathbf{P}^{\mu,b}(\cdot|h)$ can be interpreted as the distribution on Z , conditional on h being reached. Then, player i 's expected payoff conditional on h is

$$E^{\mu,b}[u_i|h] \equiv \sum_{z \in Z} u_i(z) \mathbf{P}^{\mu,b}(z|h).$$

Definition 5.1.2 A behavior strategy profile \hat{b} in a finite extensive form is sequentially rational at $h \in H_i$, given a system of beliefs μ , if

$$E^{\mu,\hat{b}}[u_i | h] \geq E^{\mu,(b_i,\hat{b}_{-i})}[u_i | h],$$

for all b_i .

A behavior strategy profile \hat{b} in an extensive form is sequentially rational, given a system of beliefs μ , if for all players i and all information sets $h \in H_i$, \hat{b} is sequentially rational at h .

A behavior strategy profile \hat{b} in an extensive form is sequentially rational if it is sequentially rational given some system of beliefs.

Definition 5.1.3 A one-shot deviation by player i from \hat{b} is a strategy b'_i with the property that there exists a (necessarily unique) information set $h' \in H_i$ such that $\hat{b}_i(h) = b'_i(h)$ for all $h \neq h'$, $h \in H_i$, and $\hat{b}_i(h') \neq b'_i(h')$.




Figure 5.1.3: The game for Example 5.1.3.

A one-shot deviation b'_i (from b , given a system of beliefs μ) is profitable if

$$E^{\mu, (b'_i, b_{-i})}[u_i | h'] > E^{\mu, b}[u_i | h'],$$

where $h' \in H_i$ is the information set for which $b'_i(h') \neq b_i(h')$.

Example 5.1.3 Consider the profile $((\text{Stop}, \text{Go}_1), \text{Go})$ in the game in Figure 5.1.3. Player I is not playing sequentially rationally at his first information set h , but does not have a profitable one-shot deviation there. Player I does have a profitable one-shot deviation at his second information set h' . Player II also has a profitable one-shot deviation. ★

The following result is obvious.

Lemma 5.1.1 If \hat{b} is sequentially rational given μ , then there are no profitable one-shot deviations.

Without further restrictions on μ (see Theorems 5.1.1 and 5.3.2), a profile may fail to be sequentially rational and yet have no profitable one-shot deviations (Problem 5.4.1).

Recall from Definition 1.3.5 that a game of perfect information has singleton information sets. In such a case, the system of beliefs is trivial, and sequential rationality is equivalent to subgame perfection.

Theorem 5.1.1 A strategy profile b in a finite game of perfect information is subgame perfect if and only if it is sequentially rational. A strategy profile b in a finite game of perfect information is

sequentially rational if and only if there are no profitable one-shot deviations.

Proof. The equivalence of subgame perfection and sequential rationality for finite games of perfect information is immediate.

It is immediate that a sequentially rational strategy profile has no profitable one-shot deviations.

The proof of the other direction is left as an exercise (Problem 5.4.3). ■

5.2 Perfect Bayesian Equilibrium

Without some restrictions connecting beliefs to behavior, even Nash equilibria need *not* be sequentially rational. For any distribution $\mathbf{P} \in \Delta(Z)$, for any $x \in X$, define

$$\mathbf{P}(x) = \sum_{\{z \in Z: x < z\}} \mathbf{P}(z).$$

Definition 5.2.1 *The information set h in a finite extensive form game is reached with positive probability under b , or is on the path-of-play, if*

$$\mathbf{P}^{\rho, b}(h) = \sum_{x \in h} \mathbf{P}^{\rho, b}(x) > 0.$$

Theorem 5.2.1 *The behavior strategy profile b of a finite extensive form game is Nash if and only if it is sequentially rational at every information set on the path of play, given a system of beliefs μ obtained using Bayes' rule at those information sets, i.e., for all h on the path of play,*

$$\mu(x) = \frac{\mathbf{P}^{\rho, b}(x)}{\mathbf{P}^{\rho, b}(h)} \quad \forall x \in h.$$

The proof of Theorem 5.2.1 is left as an exercise (Problem 5.4.4).

Example 5.2.1 Recall the extensive form from Example 2.3.4, reproduced in Figure 5.2.1. The label $[p]$ indicates that the player




Figure 5.2.1: Game for Example 5.2.1

owning that information set assigns probability p to the labeled node. The profile RBr (illustrated) is Nash and satisfies the conditions of the theorem. ★

Note that Theorem 5.2.1 implies Problem 2.6.10.

In Theorem 5.2.1, sequential rationality is only imposed at information sets on the path of play. Strengthening this to all information sets yields:

Definition 5.2.2 *A strategy profile b of a finite extensive form game is a weak perfect Bayesian equilibrium if there exists a system of beliefs μ such that*

1. b is sequentially rational given μ , and
2. for all h on the path of play,

$$\mu(x) = \frac{\mathbf{P}^{\rho, b}(x)}{\mathbf{P}^{\rho, b}(h)} \quad \forall x \in h.$$

Remark 5.2.1 Note that a strategy profile b is a weak perfect Bayesian equilibrium, if and only if, it is a Nash equilibrium that is sequentially rational.



Using Bayes' rule "where possible" yields something even stronger. The phrase "where possible" is meant to suggest that we apply Bayes' rule in a conditional manner. We first need (recall from Problem 1.4.6 that information sets are not partially ordered by precedence):

Definition 5.2.3 *The information set h follows h' if for all $x \in h$, there exists $x' \in h'$ such that $x' \prec x$.*

An information set h (following h') is reached with positive probability from h' under (μ, b) if

$$\mathbf{P}^{\mu, b}(h \mid h') = \sum_{x \in h} \mathbf{P}^{\mu, b}(x \mid h') > 0.$$

Note that for any two information sets owned by the same player, $h, h' \in H_i$, h follows h' in the sense of Definition 5.2.3 if, and only if, $h' \prec^* h$ (see Problem 1.4.6).

Definition 5.2.4 *A strategy profile b of a finite extensive form game is an almost perfect Bayesian equilibrium if there exists a system of beliefs μ such that*

1. b is sequentially rational given μ , and
2. for any information set h' and following information set h reached with positive probability from h' under (μ, b) ,

$$\mu(x) = \frac{\mathbf{P}^{\mu, b}(x \mid h')}{\mathbf{P}^{\mu, b}(h \mid h')} \quad \forall x \in h.$$

Theorem 5.2.2 *Every almost perfect Bayesian equilibrium is sub-game perfect.*

The proof is left as an exercise (Problem 5.4.5).




Figure 5.2.2: The profile $LB\ell$ (illustrated) is weak perfect Bayesian, but not almost perfect Bayesian.

Example 5.2.2 Continuing with the extensive form from Example 2.3.4 displayed in Figure 5.2.2: The profile $LB\ell$ (illustrated) is weak perfect Bayesian, but not almost perfect Bayesian. Note that $LT\ell$ is *not* weak perfect Bayesian. The only subgame perfect eq is RBr . ★

We are not yet at perfect Bayesian equilibrium, because we still need to address the phenomenon illustrated by MWG Example 9.C.4. While it is straightforward to directly deal with the example, the conditions that deal with the general phenomenon are complicated and hard to interpret. It is rare for the complicated conditions to be used in practice.

The term *perfect Bayesian equilibrium* (or *PBE*) is often used in applications to describe the collections of restrictions on the system of beliefs that “do the right/obvious thing,” and as such is one of the more abused notions in the literature. I will similarly abuse the term.

5.3 Sequential Equilibrium

A natural way of restricting the system of beliefs without simply adding one seemingly ad hoc restriction after another is to use Bayes’ rule on completely mixed profiles as follows:

Definition 5.3.1 *In a finite extensive form game, a system of beliefs μ is consistent with the strategy profile b if there exists a sequence of completely mixed sequence of behavior strategy profiles $\{b^k\}_k$ converging to b such that the associated sequence of system of beliefs $\{\mu^k\}_k$ obtained via Bayes’ rule converges to μ .*

A strategy profile b is a sequential equilibrium if, for some consistent system of beliefs μ , b is sequentially rational at every information set.

To illustrate the type of restrictions that consistency places on beliefs, consider the game in Figure 5.3.1. The strategy profile b is indicated by the double thickness arrows, and the tremble probabilities are in parentheses (the superscript k ’s have been omitted for clarity). Note that player II ’s trembles must be equal at the




Figure 5.3.1: Illustrating sequential equilibrium.

three nodes x_2 , x_3 , and x_4 , since these three nodes are in the same information set. We then have

$$\begin{aligned}\mu(x_2) &= \frac{\varepsilon}{\varepsilon + \eta + \zeta}, \\ \mu(x_3) &= \frac{\eta}{\varepsilon + \eta + \zeta}, \\ \text{and } \mu(x_4) &= \frac{\zeta}{\varepsilon + \eta + \zeta},\end{aligned}$$

and so any distribution over $\{x_2, x_3, x_4\}$ can be achieved in the limit by letting the three tremble probabilities go to zero at appropriate rates (verify this!). Note that $\lim \mu(x_4) > 0$ requires η/ζ not go to $+\infty$ (it also requires ε/ζ not go to $+\infty$, but that does not play a role in determining the limit behavior of μ at III 's information set).

Turning to player III 's information set, this information set is reached with probability $\eta\alpha + \zeta$ and so we have

$$\mu(x_5) = \frac{\eta\alpha}{\eta\alpha + \zeta},$$

$$\mu(x_6) = \frac{\zeta(1 - \alpha)}{\eta\alpha + \zeta},$$

and $\mu(x_7) = \frac{\zeta\alpha}{\eta\alpha + \zeta} < \alpha,$

and so

$$\lim_{\eta, \zeta, \alpha \rightarrow 0} \mu(x_7) = 0$$

(note that this is independent of the relative rates at which η, ζ, α go to zero). Moreover, we have

$$\begin{aligned} \frac{\eta\alpha}{\zeta} \rightarrow 0 & \text{ implies } \mu(x_6) \rightarrow 1, \\ \frac{\eta\alpha}{\zeta} \rightarrow +\infty & \text{ implies } \mu(x_6) \rightarrow 0, \\ \text{and } \frac{\eta\alpha}{\zeta} \rightarrow M > 0 & \text{ implies } \mu(x_6) \rightarrow \frac{1}{M+1}. \end{aligned}$$

In particular, $\lim \mu(x_4) > 0$ implies $\lim \mu(x_6) = 1$.

Note that it is also possible for $\lim \mu(x_6) > 0$ and $\lim \mu(x_4) = 0$ (this arises when $\eta\alpha/\zeta \rightarrow M > 0$, since $\alpha \rightarrow 0$ requires $\eta/\zeta \rightarrow +\infty$).

Theorem 5.3.1 *A sequential equilibrium is almost perfect Bayesian.*

Proof. Obvious (but make sure you understand why!). ■

Theorem 5.3.2 *In a finite extensive form game, suppose μ is consistent with a profile b . The profile b is sequentially rational given μ (and so a sequential equilibrium) if and only if there are no profitable one-shot deviations from b (given μ).*

Proof. Lemma 5.1.1 is the easy direction.

Suppose b is not sequentially rational given μ . Then there is a player, denoted i , with a profitable deviation. Denote the profitable deviation (by player i) by b'_i and the information set h' . Player i information sets H_i are strictly partially ordered by precedence in the obvious way (see Problem 1.4.6). Let $H_i(h')$ denote the finite (since the game is finite) collection of information sets that follow h' . Let K be the length of the longest chain in $H_i(h')$, and say

an information set $h \in H_i(h')$ is of level k if the successor chain from h' to h has k links (h' is 0-level and its immediate successors are all 1-level). If i has a profitable deviation from b_i at any K -level information set, then that deviation is a profitable one-shot deviation, and we are done.

Suppose i does not have a profitable deviation from b_i at any K -level information set. Define a strategy $b_i^{(K)}$ by

$$b_i^{(K)}(h) = \begin{cases} b_i(h), & \text{if } h \text{ is a } K\text{-level information set or } h \notin H_i(h'), \\ b'_i(h), & \text{if } h \text{ is a } k\text{-level information set, } k = 0, \dots, K-1. \end{cases}$$

Then $E^{\mu, (b_i^{(K)}, b_{-i})}[u_i|h'] \geq E^{\mu, (b'_i, b_{-i})}[u_i|h']$. (This requires proof, which is left as an exercise, see Problem 5.4.11. This is where consistency is important.)

But this implies that, like b'_i , the strategy $b_i^{(K)}$ is a profitable deviation at h' . We now induct on k . Either there is profitable one-shot deviation from b_i at a $(K-1)$ -level information set (in which case we are again done), or we can define a new strategy $b_i^{(K-1)}$ that is a profitable deviation at h' and which agrees with b_i on the $(K-1)$ -level as well as the K -level information sets.

Proceeding in this way, we either find a profitable one-shot deviation at some k -level information set, or the action specified at h' by b'_i is a profitable one-shot deviation. ■

5.4 Problems

- 5.4.1. Give an example of a game with a profile failing to be sequentially rational given a system of beliefs and yet with no profitable one-shot deviations.
- 5.4.2. This problem concerns the game given in Figure 5.4.1.
 - (a) Show that $(\text{GoStop}_1\text{Stop}_2, \text{StopGo}_1)$ is a Nash equilibrium.
 - (b) Identify all of the profitable one-shot deviations.
 - (c) Does player I choose Go in any subgame perfect equilibrium?
- 5.4.3. Complete the proof of Theorem 5.1.1.




Figure 5.4.1: The game for Problem 5.4.2.

- 5.4.4. Prove Theorem 5.2.1 (recall Problem 2.6.10).
- 5.4.5. Prove Theorem 5.2.2.
- 5.4.6. Let $p \in \Delta(\{x_2, x_3, x_4\})$ denote beliefs over player II's information set in the game of Figure 5.3.1. Prove that p can be obtained as the limit beliefs for some sequence of trembles.
- 5.4.7. Show that (A, a, L) is a sequential equilibrium of Selten's horse (Figure 5.1.1) by exhibiting the sequence of converging completely mixed strategies and showing that the profile is sequentially rational with respect to the limit beliefs.
- 5.4.8. Prove by direct verification that the only sequential equilibrium of the first extensive form in Example 2.3.4 is (RB, R) , but that (L, ℓ) is a sequential equilibrium of the second extensive form.
- 5.4.9. We return to the environment of Problem 3.6.1, but with one change. Rather than the two firms choosing quantities simultaneously, firm 1 is a Stackelberg leader: Firm 1 chooses its quantity, q_1 , first. Firm 2, knowing firm 1's quantity choice then chooses its quantity. Describe a strategy profile for this dynamic game. What is the appropriate equilibrium notion for this game and why? Describe an equilibrium of this game. Is it unique?
- 5.4.10. Fix a finite extensive form game. Suppose μ is consistent with b . Suppose for some player i there are two information sets $h, h' \in H_i$ with $h \prec^* h'$ and $\mathbf{P}^{(\mu, b)}(h|h') = 0$. Prove that if there exists another strategy \hat{b}_i for player i with the property that $\mathbf{P}^{(\mu, (\hat{b}_i, b_{-i}))}(h|h') > 0$, then

$$\mu(x) = \frac{\mathbf{P}^{(\mu, (\hat{b}_i, b_{-i}))}(x|h')}{\mathbf{P}^{(\mu, (\hat{b}_i, b_{-i}))}(h|h')}, \quad \forall x \in h.$$

5.4.11. Complete the proof of Theorem 5.3.2 by showing that

$$E^{\mu, (b_i^{(K)}, b_{-i})} [u_i | h'] \geq E^{\mu, (b'_i, b_{-i})} [u_i | h'].$$

Be sure to explain the role of consistency. (Hint: use Problem 5.4.10).

Chapter 6

Signaling¹

6.1 General Theory

Sender (informed player) types $t \in T \subset \mathbb{R}$. T may be finite. Probability distribution $\rho \in \Delta(T)$.

Sender chooses $m \in M \subset \mathbb{R}$. M may be finite.

Responder chooses $r \in R \subset \mathbb{R}$. R may be finite.

Payoffs: $u(m, r, t)$ for sender and $v(m, r, t)$ for responder.

Strategy for sender, $\tau : T \rightarrow M$.

Strategy for responder, $\sigma : M \rightarrow R$.

Definition 6.1.1 *A perfect Bayesian equilibrium of the signaling game is a strategy profile $(\hat{\tau}, \hat{\sigma})$ such that*

1. *for all $t \in T$,*

$$\hat{\tau}(t) \in \arg \max_{m \in M} u(m, \hat{\sigma}(m), t),$$

2. *for all m , there exists some $\mu \in \Delta(T)$ such that*

$$\hat{\sigma}(m) \in \arg \max_{r \in R} E^{\mu}[v(m, r, t)],$$

where E^{μ} denotes expectation with respect to μ , and

¹Copyright December 7, 2011 by George J. Mailath




Figure 6.1.1: A signaling game

3. for $m \in \hat{\tau}(T)$, μ in part 2 is given by

$$\mu(t) = \rho\{t \mid m = \hat{\tau}(t)\}.$$

Since the different information sets for player II are not ordered by $<^*$ (recall Problem 1.4.6), consistency places no restrictions on beliefs at different information sets of player II . This implies the following result (which Problem 6.3.2 asks you to prove).

Theorem 6.1.1 *Suppose T , M , and R are finite. A profile is a perfect Bayesian equilibrium if, and only if, it is a sequential equilibrium.*

Example 6.1.1 (Separating equilibria) In the game given in Figure 6.1.1, (bq, fr) is a separating eq. ★

Example 6.1.2 (Beer-quiche) In the game in Figure 6.1.2, (bb, rf) and (qq, fr) are both pooling eq.

The eq in which the types pool on q is often argued to be un-intuitive: Would the w type ever “rationally” deviate to b . In this pooling eq, w receives 0, and this is strictly larger than his payoff from b no matter how II responds. On the other hand, if by deviating to B , s can “signal” that he is indeed s , he is strictly better off, since II ’s best response is r , yielding payoff of 0. This is




Figure 6.1.2: The Beer-Quiche game.

an example of the *intuitive criterion*, or of the *test of equilibrium domination*. ★

Let $BR(T', m)$ be the set of best replies to m of the responder for some beliefs over T' , i.e.,

$$\begin{aligned} BR(T', m) &= \{r \in R : \exists \mu \in \Delta(T'), r \in \arg \max_{r' \in R} E^\mu[v(m, r', t)]\} \\ &= \bigcup_{\mu \in \Delta(T')} \arg \max_{r' \in R} E^\mu[v(m, r', t)]. \end{aligned}$$

Suppose $(\hat{\tau}, \hat{\sigma})$ is a perfect Bayesian equilibrium, and let $\hat{u}(t) = u(\hat{\tau}(t), \hat{\sigma}(\hat{\tau}(t)), t)$. Define $D(m) \subset T$ as the set of types satisfying

$$\hat{u}(t) > \max_{r \in BR(T, m)} u(m, r, t).$$

Definition 6.1.2 The equilibrium $(\hat{\tau}, \hat{\sigma})$ fails the intuitive criterion if there exists m' (necessarily not in $\hat{\tau}(T)$, i.e., an unsent message) and a type t' (necessarily not in $D(m)$) such that

$$\hat{u}(t') < \min_{r \in BR(T \setminus D(m'), m')} u(m', r, t').$$

Remark 6.1.1 As defined, the test concerns equilibrium outcomes, and not the specification of behavior after out-of-equilibrium messages. However, for messages m' that satisfy $\emptyset \neq D(m') \subsetneq T$, it is

in the spirit of the the test to require out-of-equilibrium responses r to m' to satisfy $r \in BR(T \setminus D(m'), m')$.

♦

6.2 Job Market Signaling

Worker with private ability $\theta \in \Theta$.

Worker can signal ability through choice of level of education, $e \in \mathbb{R}_+$.

Worker utility

$$w - c(e, \theta),$$

w is wage, and c is disutility of education. Assume c is C^2 and satisfies *single-crossing*:

$$\frac{\partial^2 c(e, \theta)}{\partial e \partial \theta} < 0.$$

Also assume $c(e, \theta) \geq 0$, $c_e(e, \theta) \equiv \partial c(e, \theta) / \partial e \geq 0$, $c_e(0, \theta) = 0$, $c_{ee}(e, \theta) > 0$, and $\lim_{e \rightarrow \infty} c_e(e, \theta) = \infty$.

Two identical firms competing for worker. Each firm values worker of type θ with education e at $f(e, \theta)$. In any discretization of the game, in any almost perfect Bayesian equilibrium, after any e , firms have identical beliefs about worker ability (see Problem 6.3.4). Consequently, the two firms are effectively playing a sealed bid common value first price auction, and so both firms bid their value $Ef(e, \theta)$. To model as a game, replace the two firms with a single uninformed receiver (the “market”) with payoff

$$-(f(e, \theta) - w)^2.$$

Strategy for worker, $e : \Theta \rightarrow \mathbb{R}_+$.

Strategy for “market”, $w : \mathbb{R}_+ \rightarrow \mathbb{R}_+$.

Assume f is C^2 . Assume $f(e, \theta) \geq 0$, $f_e(e, \theta) \equiv \partial f(e, \theta) / \partial e \geq 0$, $f_{\theta}(e, \theta) > 0$, $f_{ee}(e, \theta) \leq 0$, and $f_{e\theta}(e, \theta) \geq 0$.

Unproductive education is $f(e, \theta) = \theta$.

Productive education is $f_e(e, \theta) > 0$.

If market believes worker has ability $\hat{\theta}$, firm pays wage $f(e, \hat{\theta})$. The result is a signaling game as described in Section 6.1, and so we can apply equilibrium notion of perfect Bayesian as defined there.

6.2.1 Full Information

If firm *knows* worker has ability θ , worker chooses e to maximize

$$f(e, \theta) - c(e, \theta). \quad (6.2.1)$$

For each θ there is a unique e^* maximizing (6.2.1). That is,

$$e^*(\theta) = \arg \max_{e \geq 0} f(e, \theta) - c(e, \theta).$$

Assuming $f_e(0, \theta) > 0$ (together with the assumption on c above) is sufficient to imply that $e^*(\theta)$ is interior for all θ and so

$$\frac{de^*}{d\theta} = -\frac{f_{e\theta}(e, \theta) - c_{e\theta}(e, \theta)}{f_{ee}(e, \theta) - c_{ee}(e, \theta)} > 0.$$

6.2.2 Incomplete Information

Define

$$U(\theta, \hat{\theta}, e) \equiv f(e, \hat{\theta}) - c(e, \theta).$$

Note that

$$e^*(\theta) = \arg \max_{e \geq 0} U(\theta, \theta, e). \quad (6.2.2)$$

Suppose (\hat{e}, \hat{w}) is a separating perfect Bayesian equilibrium. The associated outcome is

$$(\hat{e}(\theta), \hat{w}(\hat{e}(\theta)))_{\theta \in \Theta}.$$

If $e' = \hat{e}(\theta')$ for some $\theta' \in \Theta$, then $\hat{w}(e') = f(e', (\hat{e})^{-1}(e')) = f(e', \theta')$, and so the payoff to the worker of type θ is

$$\hat{w}(e') - c(e', \theta) = f(e', \theta') - c(e', \theta) = U(\theta, \theta', e').$$




Figure 6.2.1: Indifference curves in $\hat{\theta}-e$ space. Note $k' = U(\theta', \theta', e^*(\theta'))$, $k'' = U(\theta'', \theta'', e'')$, and $k_0 = U(\theta'', \theta'', e^*(\theta''))$, and that incentive compatibility is satisfied at the indicated points: $U(\theta'', \theta'', e'') \geq U(\theta'', \theta', e^*(\theta'))$ and $U(\theta', \theta', e^*(\theta')) \geq U(\theta', \theta'', e'')$. For any $e < e''$, firms believe $\theta = \theta'$, and for any $e \geq e''$, firms believe $\theta = \theta''$.

In equilibrium, no type strictly benefits from mimicking another type, i.e., for all $\theta', \theta'' \in \Theta$,

$$U(\theta', \theta', \hat{e}(\theta')) \geq U(\theta', \theta'', \hat{e}(\theta'')). \quad (6.2.3)$$

This is called *incentive compatibility*. See Figure 6.2.1.

Let $\underline{\theta} = \min \Theta$. Sequential rationality implies

$$\hat{e}(\underline{\theta}) = e^*(\underline{\theta}). \quad (6.2.4)$$

(Why?)

The set of separating perfect Bayesian equilibrium outcomes is illustrated in Figure 6.2.2.

The Riley outcome is $((e^*(\theta'), f(e^*(\theta'), \theta')), (e_1'', f(e_1'', \theta'')))$; it is the separating outcome that minimizes the distortion— θ' is indifferent between $((e^*(\theta'), f(e^*(\theta'), \theta'))$ and $(e_1'', f(e_1'', \theta''))$). Any lower education level for θ'' violates (6.2.3).

(6.2.3) can be rewritten as

$$U(\theta', \theta', \hat{e}(\theta')) \geq U(\theta', (\hat{e})^{-1}(e), e) \quad \forall e \in \hat{e}(\Theta).$$

That is, the function $\hat{e} : \Theta \rightarrow \mathbb{R}_+$ satisfies the functional equation

$$\hat{e}(\theta') \in \arg \max_{e \in \hat{e}(\Theta)} U(\theta', (\hat{e})^{-1}(e), e), \quad \forall \theta' \in \Theta. \quad (6.2.5)$$

Note that (6.2.2) and (6.2.5) differ in two ways: the set of possible maximizers and how e enters into the objective function.

6.2.3 Refining to Separation

Suppose two types, θ', θ'' .

Suppose f is *affine in* θ , so that $Ef(e, \theta) = f(e, E\theta)$ (but see Problem 6.3.3(a)).

The pooling outcome in Figure 6.2.3 is a perfect Bayesian outcome, but is ruled out by the intuitive criterion: For two types, the intuitive criterion selects the “Riley” separating outcome, i.e., the separating outcome that minimizes the signaling distortion.

With three types needs a much stronger refinement (D1, see Cho and Kreps (1987)).




Figure 6.2.2: Separating equilibria. The set of separating perfect Bayesian equilibrium outcomes is given by $\{((e^*(\theta'), f(e^*(\theta'), \theta')), (e'', f(e'', \theta'')) : e'' \in [e''_1, e''_2]\}$. Note that θ'' cannot receive a lower payoff than $\max_e U(\theta'', \theta', e)$.




Figure 6.2.3: A pooling outcome at $e = e^p$. $k'_p = U(\theta', E\theta, e^p)$, $k''_p = U(\theta'', E\theta, e^p)$. Note that $E[\theta|e]$, firms' beliefs after potential deviating e 's must lie below the θ' and θ'' indifference curves indexed by k'_p and k''_p , respectively.

6.2.4 Continuum of Types

Suppose $\Theta = [\underline{\theta}, \bar{\theta}]$ (so that there is a continuum of types), and suppose \hat{e} is differentiable.

Then the first derivative of the objective function in (6.2.5) w.r.t. e is

$$\begin{aligned} & U_{\hat{\theta}}(\theta', (\hat{e})^{-1}(e), e) \frac{d(\hat{e})^{-1}(e)}{de} + U_e(\theta', (\hat{e})^{-1}(e), e) \\ &= U_{\hat{\theta}}(\theta', (\hat{e})^{-1}(e), e) \left(\frac{d\hat{e}(\theta)}{d\theta} \Big|_{\theta=(\hat{e})^{-1}(e)} \right)^{-1} + U_e(\theta', (\hat{e})^{-1}(e), e). \end{aligned}$$

The first order condition is obtained by evaluating this derivative at $e' = \hat{e}(\theta')$ (so that $(\hat{e})^{-1}(e') = \theta'$) and setting the result equal to 0:

$$U_{\hat{\theta}}(\theta', \theta', e') \left(\frac{d\hat{e}(\theta')}{d\theta} \right)^{-1} + U_e(\theta', \theta', e') = 0.$$

The result is a differential equation characterizing \hat{e} ,

$$\frac{d\hat{e}(\theta')}{d\theta} = - \frac{U_{\hat{\theta}}(\theta', \theta', e')}{U_e(\theta', \theta', e')} = - \frac{f_{\theta}(e', \theta')}{f_e(e', \theta') - c_e(e', \theta')}.$$

Together with (6.2.4), we have an initial value problem that characterizes the unique separating perfect Bayesian equilibrium strategy for the worker.

Note that because of (6.2.4), as $\theta \rightarrow \underline{\theta}$, $d\hat{e}(\theta)/d\theta \rightarrow +\infty$, and that for $\theta > \underline{\theta}$, $\hat{e}(\theta) > e^*(\theta)$, that is, there is necessarily a signalling distortion.

Remark 6.2.1 The above characterization of separating strategies works for any signaling game, given $U(\theta, \hat{\theta}, e)$, the payoff to the informed player of type θ , when the uninformed player best responds to a belief that the type is $\hat{\theta}$, and e is chosen by the informed player. See Problem 6.3.3 for a description of the canonical signaling model.

◆




Figure 6.3.1: Game for problem 6.3.1. The probability that player I is type t_1 is $1/2$ and the probability that he is type t_2 is $1/2$. The first payoff is player I 's payoff, and the second is player II 's.

6.3 Problems

- 6.3.1. Show that in the game illustrated in Figure 6.3.1, for all values of x , the outcome in which both types of player I play L is sequential by explicitly describing the converging sequence of completely mixed behavior strategy profiles and the associated system of beliefs. For what values of x does this equilibrium pass the intuitive criterion?
- 6.3.2. Prove Theorem 6.1.1.
- 6.3.3. The canonical signaling game has a sender with private information, denoted $\theta \in \Theta \subset \mathbb{R}$ choosing a message $m \in \mathbb{R}$, where Θ is compact. A receiver, observing m , but not knowing θ then chooses a response $r \in \mathbb{R}$. The payoff to the sender is $u(m, r, \theta)$ while the payoff to the receiver is $v(m, r, \theta)$. Assume both u and v are C^2 . Assume v is strictly concave in r , so that $v(m, r, \theta)$ has a unique maximizer in r for all (m, θ) , denoted $\xi(m, \theta)$. Define

$$U(\theta, \hat{\theta}, m) = u(m, \xi(m, \hat{\theta}), \theta).$$

Assume u is strictly increasing in r , and ξ is strictly increasing in $\hat{\theta}$, so that U is also strictly increasing in $\hat{\theta}$. Finally, assume that for all $(\theta, \hat{\theta})$, $U(\theta, \hat{\theta}, m)$ is bounded above (and so has a well-defined maximum).

- (a) Given a message m^* and a belief F over θ , suppose r^* maximizes the receiver's expected payoff. Prove there exists $\hat{\theta}$ such that $r^* = \xi(m^*, \hat{\theta})$. Moreover, if the support of F is a continuum, $[\underline{\theta}, \bar{\theta}]$, prove that $\hat{\theta}$ is in the support of F .

Assume u satisfies the *single-crossing condition*:

If $\theta < \theta'$ and $m < m'$, then $u(m, r, \theta) \leq u(m', r', \theta)$ implies $u(m, r, \theta') < u(m', r', \theta')$.

(Draw the indifference curves for different types in $m - r$ space to see that they can only cross once.)

- (b) Provide restrictions on the productive education case covered in Section 6.2 so that the sender's payoff satisfies the single-crossing condition as defined here.
- (c) Prove that U satisfies an analogous version of the single-crossing condition: If $\theta < \theta'$ and $m < m'$, then $U(\theta, \hat{\theta}, m) \leq U(\theta, \hat{\theta}', m')$ implies $U(\theta', \hat{\theta}, m) < U(\theta', \hat{\theta}', m')$.
- (d) Prove that the messages sent by the sender in any separating Nash equilibrium are strictly increasing in type.
- (e) Prove that in any separating perfect Bayesian equilibrium, type $\underline{\theta} \equiv \min \Theta$ chooses the action \underline{m} maximizing $u(m, \xi(m, \underline{\theta}), \underline{\theta})$ (recall (6.2.4)). How is this implication of separating perfect Bayesian equilibrium changed if u is strictly decreasing in r ? If ξ is strictly decreasing in $\hat{\theta}$?
- 6.3.4. Prove that in any discretization of the job market signaling game, in any almost perfect Bayesian equilibrium, after any e , firms have identical beliefs about worker ability.
- 6.3.5. Suppose that, in the incomplete information model of Section 6.2, the payoff to a firm from hiring a worker of type θ with education e at wage w is

$$f(e, \theta) - w = 3e\theta - w.$$

The utility of a worker of type θ with education e receiving a wage w is

$$w - c(e, \theta) = w - \frac{e^3}{\theta}.$$

Suppose the support of the firms' prior beliefs ρ on θ is $\Theta = \{1, 3\}$.

- (a) Describe a perfect Bayesian equilibrium in which both types of worker choose their full information education level. Be sure to verify that all the incentive constraints are satisfied.
- (b) Are there other separating perfect Bayesian equilibria? What are they? Do they depend on the prior distribution ρ ?

Now suppose the support of the firms' prior beliefs on θ is $\Theta = \{1, 2, 3\}$.

- (c) Why is it no longer consistent with a separating perfect Bayesian equilibrium to have $\theta = 3$ choose his full information education level $e^*(3)$? Describe the Riley outcome (the separating equilibrium outcome that minimizes the distortion), and verify that it is indeed the outcome of a perfect Bayesian equilibrium.
- (d) What is the largest education level for $\theta = 2$ consistent with separating perfect Bayesian equilibrium? Prove that any separating equilibrium in which $\theta = 2$ chooses that level of education fails the intuitive criterion. [Hint: consider the out-of-equilibrium education level $e = 3$.]
- (e) Describe the separating perfect Bayesian equilibria in which $\theta = 2$ chooses $e = 2.5$. Some of these equilibria fail the intuitive criterion and some do not. Give an example of one of each (i.e., an equilibrium that fails the intuitive criterion, and an equilibrium that does not fail).

6.3.6. The owner of a small firm is contemplating selling all or part of his firm to outside investors. The profits from the firm are risky and the owner is risk averse. The owner's preferences over x , the fraction of the firm the owner retains, and p , the price "per share" paid by the outside investors, are given by

$$u(x, \theta, p) = \theta x - x^2 + p(1 - x),$$

where θ is the value of the firm (i.e., expected profits). The quadratic term reflects the owner's risk aversion. The outside investors are risk neutral, and so the payoff to an outside investor of paying p per share for $1 - x$ of the firm is then

$$\theta(1 - x) - p(1 - x).$$

There are at least two outside investors, and the price is determined by a first price sealed bid auction: The owner first chooses the fraction of the firm to sell, $1 - x$; the outside investors then bid, with

the $1 - \alpha$ fraction going to the highest bidder (ties are broken with a coin flip).

- (a) Suppose θ is public information. What fraction of the firm will the owner sell, and how much will he receive for it?
- (b) Suppose now θ is privately known by the owner. The outside investors have common beliefs, assigning probability $\alpha \in (0, 1)$ to $\theta = \theta_1 > 0$ and probability $1 - \alpha$ to $\theta = \theta_2 > \theta_1$. Characterize the separating perfect Bayesian equilibria. Are there any other perfect Bayesian equilibria?
- (c) Maintaining the assumption that θ is privately known by the owner, suppose now that the outside investors' beliefs over θ have support $[\theta_1, \theta_2]$, so that there is a continuum of possible values for θ . What is the initial value problem (differential equation plus initial condition) characterizing separating perfect Bayesian equilibria?

6.3.7. Firm 1 is an incumbent firm selling widgets in a market in two periods. In the first period, firm 1 is a monopolist, facing a demand curve $P^1 = A - q_1^1$, where $q_1^1 \in \mathbb{R}_+$ is firm 1's output in period 1 and P^1 is the first period price. In the second period, a second firm, firm 2, will enter the market, observing the first period quantity choice of firm 1. In the second period, the two firms choose quantities simultaneously. The inverse demand curve in period 2 is given by $P^2 = A - q_1^2 - q_2^2$, where $q_i^2 \in \mathbb{R}_+$ is firm i 's output in period 2 and P^2 is the second period price. Negative prices are possible (and will arise if quantities exceed A). Firm i has a constant marginal cost of production $c_i > 0$. Firm 1's overall payoff is given by

$$(P^1 - c_1)q_1^1 + (P^2 - c_1)q_1^2,$$

while firm 2's payoff is given by

$$(P^2 - c_2)q_2^2.$$

Firm 2's marginal cost, c_2 , is common knowledge (i.e., each firm knows the marginal cost of firm 2), and satisfies $c_2 < A/2$.

- (a) Suppose c_1 is also common knowledge (i.e., each firm knows the marginal cost of the other firm), and also satisfies $c_1 < A/2$. What are the subgame perfect equilibria and why?

- (b) Suppose now that firm 1's costs are private to firm 1. Firm 2 does not know firm 1's costs, assigning prior probability $p \in (0, 1)$ to cost c_1^L and complementary probability $1 - p$ to cost c_1^H , where $c_1^L < c_1^H < A/2$.
- Define a pure strategy almost perfect Bayesian equilibrium for *this* game of incomplete information. What restrictions on second period quantities must be satisfied in any pure strategy almost perfect Bayesian equilibrium? [Make the game finite by considering discretizations of the action spaces. Strictly speaking, this is not a signaling game, since firm 1 is choosing actions in both periods, so the notion from Section 6.1 does not apply.]
 - What do the equilibrium conditions specialize to for *separating* pure strategy almost perfect Bayesian equilibria?
- (c) Suppose now that firm 2's beliefs about firm 1's costs have support $[c_1^L, c_1^H]$; i.e., the support is now an interval and not two points. What is the direction of the signaling distortion in the separating pure strategy almost perfect Bayesian equilibrium? What differential equation does the function describing first period quantities in that equilibrium satisfy?
- 6.3.8. Suppose that in the setting of Problem 3.6.1, firm 2 is a Stackelberg leader, i.e., we are reversing the order of moves from Problem 5.4.9.
- Illustrate the preferences of firm 2 in q_2 - $\hat{\theta}$ space, where q_2 is firm 2's quantity choice, and $\hat{\theta}$ is firm 1's belief about θ .
 - There is a separating perfect Bayesian equilibrium in which firm 2 chooses $q_2 = \frac{1}{2}$ when $\theta = 3$. Describe it, and prove it is a separating perfect Bayesian equilibrium (the diagram from part 6.3.8(a) may help).
 - Does the equilibrium from part 6.3.8(b) pass the intuitive criterion? Why or why not? If not, describe a separating perfect Bayesian equilibrium that does.
- 6.3.9. We continue with the setting of Problem 3.6.1, but now suppose that firm 2 is a Stackelberg leader who has the option of *not* choosing before firm 1: Firm 2 either chooses its quantity, q_2 , first, or the action W (for wait). If firm 2 chooses W , then the two firms simultaneously choose quantities, *knowing* that they are doing so. If firm 2 chooses

its quantity first (so that it did not choose W), then firm 1, knowing firm 2's quantity choice then chooses its quantity.

- (a) Describe a strategy profile for this dynamic game. Following the practice in signaling games, say a strategy profile is *perfect Bayesian* if it satisfies the conditions implied by sequential equilibrium in discretized versions of the game. (In the current context, a discretized version of the game restricts quantities to some finite subset.) What conditions must a perfect Bayesian equilibrium satisfy, and why?
- (b) For which parameter values is there an equilibrium in which firm 2 waits for all values of θ .
- (c) Prove that the outcome in which firm 2 does not wait for any θ , and firms behave as in the separating outcome of question 6.3.8(b) is *not* an equilibrium outcome of this game.

Chapter 7

Repeated Games¹

7.1 Basic Structure

Stage game $G \equiv \{(A_i, u_i)\}$:

Action space for i is A_i , with typical action $a_i \in A_i$. An action profile is denoted $a = (a_1, \dots, a_n)$.

Discount factor $\delta \in (0, 1)$.

Play G at each date $t = 0, 1, \dots$.

At the end of each period, all players observe the action profile a chosen. Actions of every player are *perfectly monitored* by all other players.

History up to date t : $h^t \equiv (a^0, \dots, a^{t-1}) \in A^t \equiv H^t$; $H^0 \equiv \{\emptyset\}$.

Set of all possible histories: $H \equiv \cup_{t=0}^{\infty} H^t$.

Strategy for player i — $s_i : H \rightarrow A_i$. Often written $s_i = (s_i^0, s_i^1, s_i^2, \dots)$, where $s_i^t : H^t \rightarrow A_i$. Since $H^0 = \{\emptyset\}$, we have $s^0 \in A$, and so can write a^0 for s^0 .

Note distinction between

- actions $a_i \in A_i$ and
- strategies $s_i : H \rightarrow A_i$.

Given strategy profile $s \equiv (s_1, s_2, \dots, s_n)$, outcome path induced by s is $a(s) = (a^0, a^1, a^2, \dots)$, where

$$a^0 = (s_1^0, s_2^0, \dots, s_n^0),$$

¹Copyright December 7, 2011 by George J. Mailath

$$\begin{aligned}
a^1 &= (s_1^1(a^0), s_2^1(a^0), \dots, s_n^1(a^0)), \\
a^2 &= (s_1^2(a^0, a^1), s_2^2(a^0, a^1), \dots, s_n^2(a^0, a^1)), \\
&\vdots
\end{aligned}$$

Payoffs of $G(\infty)$ are

$$U_i(s) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t(s)).$$

Definition 7.1.1 *The strategy profile \hat{s} is a Nash equilibrium of $G(\infty)$ if, for all i and all $\tilde{s}_i : H \rightarrow S_i$,*

$$U_i(\hat{s}_i, \hat{s}_{-i}) \geq U_i(\tilde{s}_i, \hat{s}_{-i}).$$

Definition 7.1.2 *Player i 's pure strategy minmax utility is*

$$\underline{v}_i^p = \min_{a_{-i}} \max_{a_i} u_i(a_i, a_{-i}).$$

The profile $\hat{a}_{-i} \in \arg \min_{a_{-i}} \max_{a_i} u_i(a_i, a_{-i})$ minmaxes player i . The set of (pure strategy) strictly individually rational payoffs in $\{(S_i, u_i)\}$ is $\{v \in \mathbb{R}^n : v_i > \underline{v}_i^p\}$. The set of feasible payoffs in $\{(S_i, u_i)\}$ is $\text{conv}\{v \in \mathbb{R}^n : \exists a \in S, v = u(a)\}$. Define $\mathcal{F}^{p} \equiv \{v \in \mathbb{R}^n : v_i > \underline{v}_i^p\} \cap \text{conv}\{v \in \mathbb{R}^n : \exists a \in S, v = u(a)\}$.*

Theorem 7.1.1 *Suppose s^* is a pure strategy Nash equilibrium. Then,*

$$U_i(s^*) \geq \underline{v}_i^p.$$

Proof. Let \hat{s}_i be a strategy satisfying

$$\hat{s}_i(h^t) \in \arg \max_{a_i} u_i(a_i, s_{-i}^*(h^t)), \quad \forall h^t \in H^t$$

(if the $\arg \max$ is unique for some history h^t , $\hat{s}_i(h^t)$ is uniquely determined, otherwise make a selection from the $\arg \max$). Since

$$U_i(s^*) \geq U_i(\hat{s}_i, s_{-i}^*),$$

and since in every period \underline{v}_i^p is a lower bound for the flow payoff received under the profile (\hat{s}_i, s_{-i}^*) , we have

$$U_i(s^*) \geq U_i(\hat{s}_i, s_{-i}^*) \geq (1 - \delta) \sum_{t=0}^{\infty} \delta^t \underline{v}_i^p = \underline{v}_i^p.$$

■

Remark 7.1.1 In some settings it is necessary to allow players to randomize. For example, in matching pennies, the set of pure strategy feasible and individually rational payoffs is empty.

Definition 7.1.3 *Player i 's mixed strategy minmax utility is*

$$\underline{v}_i = \min_{\alpha_{-i} \in \prod_{j \neq i} \Delta(A_j)} \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \alpha_{-i}).$$

The profile $\hat{\alpha}_{-i} \in \arg \min_{\alpha_{-i}} \max_{\alpha_i} u_i(\alpha_i, \alpha_{-i})$ minmaxes player i . The set of (mixed strategy) strictly individually rational payoffs in $\{(S_i, u_i)\}$ is $\{v \in \mathbb{R}^n : v_i > \underline{v}_i\}$. Define $\mathcal{F}^* \equiv \{v \in \mathbb{R}^n : v_i > \underline{v}_i\} \cap \text{conv}\{v \in \mathbb{R}^n : \exists a \in S, v = u(a)\}$.

The Minmax Theorem (Problem 4.3.1) implies that \underline{v}_i is i 's security level (Definition 2.4.2).

A proof essentially identical to that proving Theorem 7.1.1 (applied to the behavior strategy profile realization equivalent to σ^*) proves the following:

Theorem 7.1.2 *Suppose σ^* is a (possibly mixed) Nash equilibrium. Then,*

$$U_i(\sigma^*) \geq \underline{v}_i.$$

Since $\underline{v}_i \leq \underline{v}_i^p$ (with a strict inequality in some games, such as matching pennies), lower payoffs often can be enforced using mixed strategies. The possibility of enforcing lower payoffs allows higher payoffs to be enforced in subgame perfect equilibria.

◆

Given $h^{t'} = (a^0, \dots, a^{t'-1}) \in H^{t'}$ and $\bar{h}^t = (\bar{a}^0, \dots, \bar{a}^{t-1}) \in H^t$, the history $(a^0, \dots, a^{t'-1}, \bar{a}^0, \dots, \bar{a}^{t-1}) \in H^{t'+t}$ is the *concatenation* of $h^{t'}$

followed by \bar{h}^t , denoted by $(h^{t'}, \bar{h}^t)$. Given s_i , define $s_i|_{h^{t'}}: H \rightarrow S_i$ as follows:

$$s_i^t|_{h^{t'}}(\bar{h}^t) = s_i^{t'+t}(h^{t'}, \bar{h}^t).$$

Definition 7.1.4 *The strategy profile \hat{s} is a subgame perfect equilibrium of $G(\infty)$ if, for all histories, $h^t \in H^t$, $\hat{s}|_{h^t} = (\hat{s}_i|_{h^t}, \dots, \hat{s}_n|_{h^t})$ is a Nash equilibrium of $G(\infty)$.*

Example 7.1.1 (Grim trigger in the repeated PD)

	<i>E</i>	<i>S</i>
<i>E</i>	2, 2	-1, 3
<i>S</i>	3, -1	0, 0

Grim trigger: player i 's strategy is given by

$$\hat{s}_i^0 = E,$$

and for $t \geq 1$,

$$\hat{s}_i^t(a^0, \dots, a^{t-1}) = \begin{cases} E, & \text{if } a^{t'} = EE \text{ for all } t' = 0, 1, \dots, t-1, \\ S, & \text{otherwise.} \end{cases}$$

Payoff to I from (\hat{s}_1, \hat{s}_2) is: $(1-\delta) \sum 2 \times \delta^t = 2$. Payoff from deviating in period $t = 0$ (most profitable deviation) is $(1-\delta)3$. Nash if

$$2 \geq 3(1-\delta) \iff \delta \geq \frac{1}{3}.$$

Strategy profile is subgame perfect: only need to check subgame after *one* deviation. On such, \hat{s} specifies SS in every period and that is clearly Nash. ★

Represent strategy profiles by *automata*, $(\mathcal{W}, w^0, f, \tau)$, where

- \mathcal{W} is set of states,
- w^0 is initial state,

- $f : \mathcal{W} \rightarrow A$ is output function (decision rule),² and
- $\tau : \mathcal{W} \times A \rightarrow \mathcal{W}$ is transition function.

Any automaton $(\mathcal{W}, w^0, f, \tau)$ induces a pure strategy profile as follows: First, extend the transition function from the domain $\mathcal{W} \times A$ to the domain $\mathcal{W} \times (H \setminus \{\emptyset\})$ by recursively defining

$$\tau(w, h^t) = \tau(\tau(w, h^{t-1}), a^{t-1}).$$

With this definition, the strategy s induced by the automaton is given by $s(\emptyset) = f(w^0)$ and

$$s(h^t) = f(\tau(w^0, h^t)), \forall h^t \in H \setminus \{\emptyset\}.$$

Conversely, it is straightforward that any strategy profile can be represented by an automaton. Take the set of histories H as the set of states, the null history \emptyset as the initial state, $f(h^t) = s(h^t)$, and $\tau(h^t, a) = h^{t+1}$, where $h^{t+1} \equiv (h^t, a)$ is the concatenation of the history h^t with the action profile a .

This representation leaves us in the position of working with the full set of histories \mathcal{H} . However, strategy profiles can often be represented by automata with finite sets \mathcal{W} . The set \mathcal{W} is then a partition on \mathcal{H} , grouping together those histories that prompt identical continuation strategies.

Advantage of automaton representation clearest when \mathcal{W} can be chosen finite, but also has conceptual advantages.


Example 7.1.2 (Grim trigger in the repeated PD, cont.)

Grim trigger profile has automata representation, $(\mathcal{W}, w^0, f, \tau)$, with $\mathcal{W} = \{w_{EE}, w_{SS}\}$, $w^0 = w_{EE}$, $f(w_{EE}) = EE$ and $f(w_{SS}) = SS$, and

$$\tau(w, a) = \begin{cases} w_{EE}, & \text{if } w = w_{EE} \text{ and } a = EE, \\ w_{SS}, & \text{otherwise.} \end{cases}$$

The automaton can be represented as:

²A profile of behavior strategies (b_1, \dots, b_n) , $b_i : H \rightarrow \Delta(A_i)$, can also be represented by an automaton. The output function now maps into profiles of mixtures over action profiles, i.e., $f : \mathcal{W} \rightarrow \prod_i \Delta(A_i)$.



★

If s is represented by $(\mathcal{W}, w^0, f, \tau)$, the continuation strategy profile after a history h^t , $s|_{h^t}$ is represented by the automaton $(\mathcal{W}, \tau(w^0, h^t), f, \tau)$, where $\tau(w^0, h^t)$ is the result of recursively applying τ to h^t , i.e., if $h^t = (h^{t-1}, a^{t-1})$, then $\tau(w^0, h^t) = \tau(\tau(w^0, h^{t-1}), a^{t-1})$.

Lemma 7.1.1 *The strategy profile with representing automaton $(\mathcal{W}, w^0, f, \tau)$ is a subgame perfect equilibrium iff for all $w \in \mathcal{W}$ (satisfying $w = \tau(w^0, h^t)$ for some $h^t \in H$) the strategy profile represented by $(\mathcal{W}, w, f, \tau)$ is a Nash eq of the repeated game.*

Given an automaton $(\mathcal{W}, w^0, f, \tau)$, let $V_i(w)$ be i 's value from being in the state $w \in \mathcal{W}$, i.e.,

$$V_i(w) = (1 - \delta)u_i(f(w)) + \delta V_i(\tau(w, f(w))).$$

Note that if \mathcal{W} is finite, V_i solves a finite set of linear equations (see Problem 7.6.3).

Compare the following definition with Definition 5.1.3, and the proofs of Theorem 7.1.3 with that of Theorem 5.3.2.

Definition 7.1.5 *Player i has a profitable one-shot deviation from $(\mathcal{W}, w^0, f, \tau)$, if there is some history h^t and some action $a_i \in A_i$ such that (where $w = \tau(w^0, h^t)$)*

$$V_i(w) < (1 - \delta)u_i(a_i, f_{-i}(w)) + \delta V_i(\tau(w, (a_i, f_{-i}(w)))).$$

Theorem 7.1.3 *A strategy profile is subgame perfect iff there are no profitable one-shot deviations.*

Proof. Clearly, if a strategy profile is subgame perfect, then there are no profitable deviations.

We need to argue that if a profile is not subgame perfect, then there is a profitable one-shot deviation.

Suppose (s_1, \dots, s_n) (with representing automaton $(\mathcal{W}, w^0, f, \tau)$) is not subgame perfect. Then there exists some history $\tilde{h}^{t'}$ and player i such that $s_i|_{\tilde{h}^{t'}}$ is not a best reply to $s_{-i}|_{\tilde{h}^{t'}}$. That is, there exists \hat{s}_i such that

$$0 < U_i(\hat{s}_i, s_{-i}|_{\tilde{h}^{t'}}) - U_i(s_i|_{\tilde{h}^{t'}}, s_{-i}|_{\tilde{h}^{t'}}) \equiv \varepsilon.$$

For simplicity, define $\tilde{s}_j = s_j|_{\tilde{h}^{t'}}$. Defining $M \equiv 2 \max_{i,a} |u_i(a)|$, suppose T is large enough so that $\delta^T M < \varepsilon/2$, and consider the strategy for i defined by

$$\bar{s}_i(h^t) = \begin{cases} \hat{s}_i(h^t), & t < T, \\ \tilde{s}_i(h^t), & t \geq T. \end{cases}$$

Then,

$$|U_i(\bar{s}_i, \tilde{s}_{-i}) - U_i(\hat{s}_i, \tilde{s}_{-i})| \leq \delta^T M < \varepsilon/2,$$

so that

$$U_i(\bar{s}_i, s_{-i}|_{\tilde{h}^{t'}}) - U_i(s_i|_{\tilde{h}^{t'}}, s_{-i}|_{\tilde{h}^{t'}}) \equiv U_i(\bar{s}_i, \tilde{s}_{-i}) - U_i(\tilde{s}_i, \tilde{s}_{-i}) > \varepsilon/2 > 0.$$

Note that \bar{s}_i is a profitable “ T -period” deviation from \tilde{s}_i .

Let \bar{h}^{T-1} be the outcome path up to and including period $T-1$ (history) induced by $(\bar{s}_i, \tilde{s}_{-i})$, and let $\bar{w} = \tau(w^0, \tilde{h}^{t'} \bar{h}^{T-1})$. Note that

$$U_i(\tilde{s}_i|_{\bar{h}^{T-1}}, \tilde{s}_{-i}|_{\bar{h}^{T-1}}) = V_i(\tau(w^0, \tilde{h}^{t'} \bar{h}^{T-1})) = V_i(\bar{w})$$

and

$$\begin{aligned} & U_i(\bar{s}_i|_{\bar{h}^{T-1}}, \tilde{s}_{-i}|_{\bar{h}^{T-1}}) \\ &= (1 - \delta)u_i(\bar{s}_i(\bar{h}^{T-1}), f_{-i}(\bar{w})) + \delta V_i(\tau(\bar{w}, (\bar{s}_i(\bar{h}^{T-1}), f_{-i}(\bar{w}))). \end{aligned}$$

Hence, if $U_i(\bar{s}_i|_{\bar{h}^{T-1}}, \tilde{s}_{-i}|_{\bar{h}^{T-1}}) > U_i(\tilde{s}_i|_{\bar{h}^{T-1}}, \tilde{s}_{-i}|_{\bar{h}^{T-1}})$, then we are done, since player i has a profitable one-shot deviation from $(\mathcal{W}, w^0, f, \tau)$.

Suppose not, i.e., $U_i(\bar{s}_i|_{\bar{h}^{T-1}}, \tilde{s}_{-i}|_{\bar{h}^{T-1}}) \leq U_i(\tilde{s}_i|_{\bar{h}^{T-1}}, \tilde{s}_{-i}|_{\bar{h}^{T-1}})$. For the strategy \check{s}_i defined by

$$\check{s}_i(h^t) = \begin{cases} \bar{s}_i(h^t), & t < T-1, \\ \tilde{s}_i(h^t), & t \geq T-1, \end{cases}$$

we have

$$\begin{aligned}
 U_i(\tilde{s}_i, \tilde{s}_{-i}) &= (1 - \delta) \sum_{t=0}^{T-2} \delta^t u_i(a^t(\tilde{s}_i, \tilde{s}_{-i})) + \delta^{T-1} U_i(\tilde{s}_i|_{\tilde{h}^{T-1}}, \tilde{s}_{-i}|_{\tilde{h}^{T-1}}) \\
 &\geq (1 - \delta) \sum_{t=0}^{T-2} \delta^t u_i(a^t(\tilde{s}_i, \tilde{s}_{-i})) + \delta^{T-1} U_i(\tilde{s}_i|_{\tilde{h}^{T-1}}, \tilde{s}_{-i}|_{\tilde{h}^{T-1}}) \\
 &= U_i(\tilde{s}_i, \tilde{s}_{-i}) > U_i(\tilde{s}_i, \tilde{s}_{-i}).
 \end{aligned}$$

That is, the $(T - 1)$ -period deviation is profitable. But then either the one-shot deviation in period $T - 1$ is profitable, or the $(T - 2)$ -shot deviation is profitable. Induction completes the argument. ■

See Problem 7.6.4 for an alternative (and perhaps more enlightening) proof.

Corollary 7.1.1 *Suppose the strategy profile s is represented by $(\mathcal{W}, w^0, f, \tau)$. Then s is subgame perfect if, and only if, for all $w \in \mathcal{W}$ (satisfying $w = \tau(w^0, h^t)$ for some $h^t \in H$), $f(w)$ is a Nash eq of the normal form game with payoff function $g^w : A \rightarrow \mathbb{R}^n$, where*

$$g_i^w(a) = (1 - \delta)u_i(a) + \delta V_i(\tau(w, a)).$$

Definition 7.1.6 *An action profile $a' \in A$ is enforced by the continuation promises $\gamma : A \rightarrow \mathbb{R}^n$ if a' is a Nash eq of the normal form game with payoff function $g^w : A \rightarrow \mathbb{R}^n$, where*

$$g_i^w(a) = (1 - \delta)u_i(a) + \delta \gamma_i(a).$$

A payoff v is decomposed on a set of payoffs \mathcal{V} if there exists an action profile a' enforced by some continuation promises $\gamma : A \rightarrow \mathcal{V}$ satisfying, for all i ,

$$v_i = (1 - \delta)u_i(a') + \delta \gamma_i(a').$$

Example 7.1.3 (continuation of grim trigger) We clearly have $V_1(w_{EE}) = 2$ and $V_1(w_{SS}) = 0$, so that the normal form associated with w_{EE} is

	E	S
E	2, 2	$-(1 - \delta), 3(1 - \delta)$
S	$3(1 - \delta), -(1 - \delta)$	0, 0

while the normal form for w_{SS} is

	E	S
E	$2(1 - \delta), 2(1 - \delta)$	$-(1 - \delta), 3(1 - \delta)$
S	$3(1 - \delta), -(1 - \delta)$	$0, 0$


As required EE is a (but not the only!) Nash eq of the w_{EE} normal form, while SS is a Nash eq of the w_{SS} normal form. ★

Example 7.1.4 Stage game:

	A	B	C
A	4, 4	3, 2	1, 1
B	2, 3	2, 2	1, 1
C	1, 1	1, 1	-1, -1

Stage game has a unique Nash eq: AA . Suppose $\delta \geq \frac{2}{3}$. Then there is a subgame perfect equilibrium of $G(\infty)$ with outcome path $(BB)^\infty$: $(\mathcal{W}, w^0, f, \tau)$, where $\mathcal{W} = \{w_{BB}, w_{CC}\}$, $w^0 = w_{BB}$, $f_i(w_a) = a_i$, and

$$\tau(w, a) = \begin{cases} w_{BB}, & \text{if } w = w_{BB} \text{ and } a = BB, \text{ or } w = w_{CC} \text{ and } a = CC, \\ w_{CC}, & \text{otherwise.} \end{cases}$$



Values of the states are

$$V_i(w_{BB}) = (1 - \delta)2 + \delta V_i(w_{BB}),$$

$$\text{and } V_i(w_{CC}) = (1 - \delta) \times (-1) + \delta V_i(w_{BB}).$$

Solving,

$$\begin{aligned} V_i(w_{BB}) &= 2, \\ \text{and } V_i(w_{CC}) &= 3\delta - 1. \end{aligned}$$

Player 1's payoffs in the normal form associated with w_{BB} are

	A	B	C
A	$4(1 - \delta) + \delta(3\delta - 1)$	$3(1 - \delta) + \delta(3\delta - 1)$	$1 - \delta + \delta(3\delta - 1)$
B	$2(1 - \delta) + \delta(3\delta - 1)$	2	$1 - \delta + \delta(3\delta - 1)$
C	$1 - \delta + \delta(3\delta - 1)$	$1 - \delta + \delta(3\delta - 1)$	$-(1 - \delta) + \delta(3\delta - 1)$

and since the game is symmetric, BB is a Nash eq of this normal form only if

$$2 \geq 3(1 - \delta) + \delta(3\delta - 1),$$

i.e.,

$$0 \geq 1 - 4\delta + 3\delta^2 \Leftrightarrow 0 \geq (1 - \delta)(1 - 3\delta),$$

or $\delta \geq \frac{1}{3}$.

Player 1's payoffs in the normal form associated with w_{CC} are

	A	B	C
A	$4(1 - \delta) + \delta(3\delta - 1)$	$3(1 - \delta) + \delta(3\delta - 1)$	$1 - \delta + \delta(3\delta - 1)$
B	$2(1 - \delta) + \delta(3\delta - 1)$	$2(1 - \delta) + \delta(3\delta - 1)$	$1 - \delta + \delta(3\delta - 1)$
C	$1 - \delta + \delta(3\delta - 1)$	$1 - \delta + \delta(3\delta - 1)$	$-(1 - \delta) + \delta 2$

and since the game is symmetric, CC is a Nash eq of this normal form only if

$$-(1 - \delta) + \delta 2 \geq 1 - \delta + \delta(3\delta - 1),$$

i.e.,

$$0 \geq 2 - 5\delta + 3\delta^2 \Leftrightarrow 0 \geq (1 - \delta)(2 - 3\delta),$$

or $\delta \geq \frac{2}{3}$.

★


7.2 Modeling Competitive Agents (Small/Short-Lived Players)

Example 7.2.1 (Product-choice game)

	h	ℓ
H	3, 3	0, 2
L	4, 0	2, 1

Player I (row player) is long-lived, player II (the column player) is short lived.

Subgame perfect equilibrium described by the two state automaton:



The action profile $L\ell$ is a static Nash equilibrium, and since $w_{L\ell}$ is an absorbing state, we trivially have that $L\ell$ is a Nash equilibrium of the associated one-shot game, $g^{w_{L\ell}}$.

Note that $V_1(w_{Hh}) = 3$ and $V_1(w_{L\ell}) = 2$. Since player 2 is short-lived, he must myopically optimize in each period. The one-shot game from Corollary 7.1.1 has only one player. The one-shot game $g^{w_{Hh}}$ associated with w_{Hh} is given by

	h
H	$(1 - \delta)3 + \delta 3$
L	$(1 - \delta)4 + \delta 2$

and player I finds H optimal if $3 \geq 4 - 2\delta$, i.e., if $\delta \geq 1/2$.


Thus the profile is a subgame perfect equilibrium if, and only if, $\delta \geq 1/2$. ★

Example 7.2.2

	h	ℓ
H	3, 3	2, 2
L	4, 0	0, 1

The action profile $L\ell$ is no longer a static Nash equilibrium, and so Nash reversion cannot be used to discipline player I 's behavior.

Subgame perfect equilibrium described by the two state automaton:



Since player 2 is short-lived, he must myopically optimize in each period, and he is.

Note that $V_1(w_{Hh}) = 3$ and $V_1(w_{Ll}) = (1 - \delta)0 + \delta 3 = 3\delta$. There are two one shot games we need to consider. The one-shot game $g^{w_{Hh}}$ associated with w_{Hh} is given by

	h
H	$(1 - \delta)3 + \delta 3$
L	$(1 - \delta)4 + 3\delta^2$

and player I finds H optimal if $3 \geq 4 - 4\delta + 3\delta^2 \Leftrightarrow 0 \geq (1 - \delta)(1 - 3\delta) \Leftrightarrow \delta \geq 1/3$.

The one-shot game $g^{w_{Ll}}$ associated with w_{Ll} is given by

	ℓ
H	$(1 - \delta)2 + 3\delta^2$
L	$(1 - \delta)0 + 3\delta$

and player I finds L optimal if $3\delta \geq 2 - 2\delta + 3\delta^2 \Leftrightarrow 0 \geq (1 - \delta)(2 - 3\delta) \Leftrightarrow \delta \geq 2/3$.

Thus the profile is a subgame perfect equilibrium if, and only if, $\delta \geq 2/3$. ★

Example 7.2.3 Stage game: Seller chooses quality, “ H ” or “ L ”, and announces price.

Cost of producing $H = c_H = 2$.

Cost of producing $L = c_L = 1$.

Demand:

$$x(p) = \begin{cases} 10 - p, & \text{if } H, \text{ and} \\ 4 - p, & \text{if } L. \end{cases}$$

If L , $\max_p (4 - p)(p - c_L) \Rightarrow p = \frac{5}{2} \Rightarrow x = \frac{3}{2}, \pi^L = \frac{9}{4}$.

If H , $\max_p (10 - p)(p - c_H) \Rightarrow p = 6 \Rightarrow x = 4, \pi^H = 16$.

Quality is only observed after purchase.

Model as a game: Strategy space for seller $\{(H, p), (L, p') : p, p' \in \mathbb{R}_+\}$.

Continuum of (long-lived) consumers of mass 10, each consumer buys zero or one unit of good. Consumer $i \in [0, 10]$ values one unit of good as follows

$$v_i = \begin{cases} i, & \text{if } H, \text{ and} \\ \max\{0, i - 6\}, & \text{if } L. \end{cases}$$

Strategy space for consumer i is $\{s : \mathbb{R}_+ \rightarrow \{0, 1\}\}$, where 1 is buy and 0 is not buy.

Strategy profile is $((Q, p), \xi)$, where $\xi(i)$ is consumer i 's strategy. Write ξ_i for $\xi(i)$. Consumer i 's payoff function is

$$u_i((Q, p), \xi) = \begin{cases} i - p, & \text{if } Q = H \text{ and } \xi_i(p) = 1, \\ \max\{0, i - 6\} - p, & \text{if } Q = L \text{ and } \xi_i(p) = 1, \text{ and} \\ 0, & \text{if } \xi_i(p) = 0. \end{cases}$$

Firm's payoff function is

$$\pi((Q, p), \xi) = (p - c_Q)\hat{x}(p, \xi)$$

$$\begin{aligned}
&\equiv (p - c_Q) \int_0^{10} \xi_i(p) di \\
&= (p - c_Q) \lambda\{i \in [0, 10] : \xi_i(p) = 1\},
\end{aligned}$$

where λ is Lebesgue measure. [Note that we need to assume that ξ is measurable.]

Assume firm only observes $\hat{x}(p, \xi)$ at the end of the period, so that consumers are *anonymous*.

Note that $\hat{x}(p, \xi)$ is independent of Q , and that the choice (L, p) strictly dominates (H, p) whenever $\hat{x}(p, \xi) \neq 0$.

If consumer i believes the firm has chosen Q , then i 's best response to p is $\xi_i(p) = 1$ only if $u_i((Q, p), \xi) \geq 0$. Let $\xi_i^Q(p)$ denote the maximizing choice of consumer i when the consumer observes price p and believes the firm also chose quality Q . Then,

$$\xi_i^H(p) = \begin{cases} 1, & \text{if } i \geq p, \text{ and} \\ 0, & \text{if } i < p, \end{cases}$$

so $x(p, \xi^H) = \int_p^{10} di = 10 - p$. Also,

$$\xi_i^L(p) = \begin{cases} 1, & \text{if } i \geq p + 6, \text{ and} \\ 0, & \text{if } i < p + 6, \end{cases}$$

so $\hat{x}(p, \xi^L) = \int_{p+6}^{10} di = 10 - (p + 6) = 4 - p$.

Unique *subgame perfect* equilibrium of stage game is $((L, \frac{5}{2}), \xi^L)$.

Why isn't the *outcome path* $((H, 6), \xi^H(6))$ consistent with subgame perfection? Note that there are two distinct deviations by the firm to consider: an *unobserved* deviation to $(L, 6)$, and an *observed* deviation involving a price different from 6. In order to deter an observed deviation, we specify that consumer's believe that, in response to any price different from 6, the firm had chosen $Q = L$, leading to the best response $\tilde{\xi}_i$ given by

$$\tilde{\xi}_i(p) = \begin{cases} 1, & \text{if } p = 6 \text{ and } i \geq p, \text{ or } p \neq 6 \text{ and } i \geq p + 6, \\ 0, & \text{otherwise,} \end{cases}$$

implying aggregate demand

$$\hat{x}(p, \tilde{\xi}) = \begin{cases} 4, & \text{if } p = 6, \\ \max\{0, 4 - p\}, & p \neq 6. \end{cases}$$




Figure 7.2.1: Grim trigger in the quality game. Note that the transitions are only a function of Q .

Clearly, this implies that observable deviations by the firm are not profitable. Consider then the profile $((H, 6), \tilde{\xi})$: the unobserved deviation to $(L, 6)$ is profitable, since profits in this case are $(10 - 6)(6 - 1) = 20 > 16$. Note that for the deviation to be profitable, firm must still charge 6 (not the best response to ξ^H).

Eq with high quality: buyers believe H will be produced as long as H has been produced in the past. If ever L is produced, then L is expected to always be produced in future. See Figure 7.2.1.

It only remains to specify the decision rules:

$$f_1(w) = \begin{cases} (H, 6), & \text{if } w = w_H, \text{ and} \\ (L, \frac{5}{2}), & \text{if } w = w_L. \end{cases}$$

and

$$f_2(w) = \begin{cases} \tilde{\xi}, & \text{if } w = w_H, \text{ and} \\ \xi^L, & \text{if } w = w_L. \end{cases}$$

Since the transitions are independent of price, the firm's price is myopically optimal in each state.

Since the consumers are small and myopically optimizing, in order to show that the profile is subgame perfect, it remains to verify that the firm is behaving optimally in each state. The firm value in each state is $V_1(w_Q) = \pi^Q$. Trivially, L is optimal in w_L . Turning to w_H , we have

$$(1 - \delta)20 + \delta \frac{9}{4} \leq 16 \Leftrightarrow \delta \geq \frac{16}{71}.$$

There are many other equilibria. ★

Remark 7.2.1 (Short-lived player) Can model above as a game between one long-lived player and one short-lived player. In the stage game, the firm chooses p , and then the firm and consumer simultaneously choose quality $Q \in \{L, H\}$, and quantity $x \in [0, 10]$. If the good is high quality, the consumer receives a utility of $10x - x^2/2$ from consuming x units. If the good is of low quality, his utility is reduced by 6 per unit, giving a utility of $4x - x^2/2$.³ The consumer's utility is linear in money, so his payoffs are

$$u_c(Q, p) = \begin{cases} (4 - p)x - \frac{x^2}{2}, & \text{if } Q = L, \text{ and} \\ (10 - p)x - \frac{x^2}{2}, & \text{if } Q = H. \end{cases}$$

Since the period t consumer is *short-lived* (a new consumer replaces him next period), if he expects L in period t , then his best reply is to choose $x = x^L(p) \equiv \max\{4 - p, 0\}$, while if he expects H , his best reply is choose $x = x^H(p) \equiv \max\{10 - p, 0\}$. In other words, his behavior is just like the aggregate behavior of the continuum of consumers.

This is in general true: a short-lived player can typically represent a continuum of long-lived anonymous players.



7.3 The Folk Theorem (Cooperation from Long-Run Interactions and an Embarrassment of Riches)

Theorem 7.3.1 (The Folk Theorem) *Suppose \mathcal{F}^* has nonempty interior in \mathbb{R}^n . For all $v \in \mathcal{F}^*$, there exists a sufficiently large discount factor δ' , such that for all $\delta \geq \delta'$, there is a subgame perfect equilibrium of the infinitely repeated game whose average discounted value is v .*

³Note that for $x > 4$, utility is declining in consumption. This can be avoided by setting his utility equal to $4x - x^2/2$ for $x \leq 4$, and equal to 8 for all $x > 4$. This does not affect any of the relevant calculations.

Example 7.3.1 (Symmetric folk theorem for PD) Suppose restrict attention to *strongly symmetric* strategies, i.e., for all $w \in \mathcal{W}$, $f_1(w) = f_2(w)$. When is $\{(v, v) : v \in [0, 2]\}$ a set of eq payoffs? Since interested in strongly symmetric equilibria, will drop player subscripts. Note that the set of strongly symmetric equilibrium payoffs cannot be any larger, since $[0, 2]$ is the largest set of feasible symmetric payoffs.

Two preliminary calculations (important to note that these preliminary calculations make no assumptions about $[0, 2]$ being a set of eq payoffs):

1. Let \mathcal{W}^{EE} be the set of player 1 payoffs that could be decomposed on $[0, 2]$ using *EE* (i.e., \mathcal{W}^{EE} is the set of player 1 payoffs that could enforceably achieved by *EE* followed by appropriate symmetric continuations in $[0, 2]$). Then $v \in \mathcal{W}^{EE}$ iff

$$\begin{aligned} v &= 2(1 - \delta) + \delta\gamma(EE) \\ &\geq 3(1 - \delta) + \delta\gamma(SE), \end{aligned}$$

for some $\gamma(EE), \gamma(SE) \in [0, 2]$. The largest value for $\gamma(EE)$ is 2, while the incentive constraint implies the smallest is $(1 - \delta)/\delta$, so that $\mathcal{W}^{EE} = [3(1 - \delta), 2]$. See Figure 7.3.1 for an illustration.

2. Let \mathcal{W}^{SS} be the set of player 1 payoffs that could be decomposed on $[0, 2]$ using *SS*. Then $v \in \mathcal{W}^{SS}$ iff

$$\begin{aligned} v &= 0 \times (1 - \delta) + \delta\gamma(SS) \\ &\geq (-1)(1 - \delta) + \delta\gamma(ES), \end{aligned}$$

for some $\gamma(SS), \gamma(ES) \in [0, 2]$. Since the inequality is satisfied by setting $\gamma(SS) = \gamma(ES)$, the largest value for $\gamma(SS)$ is 2, while the smallest is 0, and so $\mathcal{W}^{SS} = [0, 2\delta]$.

Observe that

$$[0, 2] \supset \mathcal{W}^{SS} \cup \mathcal{W}^{EE} = [0, 2\delta] \cup [3(1 - \delta), 2].$$

Lemma 7.3.1 (Necessity) Suppose $[0, 2]$ is the set of strongly symmetric strategy equilibrium payoffs. Then,

$$[0, 2] \subset \mathcal{W}^{SS} \cup \mathcal{W}^{EE}.$$




Figure 7.3.1: An illustration of the folk theorem. The continuations that enforce EE are labelled y^{EE} , while those that enforce SS are labelled y^{SS} . The value v^0 is the average discounted value of the equilibrium whose current value/continuation value is described by the cycle $0 - 1 - 2 - 0$. In this cycle, play follows EE, EE, SS, EE, \dots . The Figure was drawn for $\delta = \frac{2}{3}$; $v^0 = \frac{30}{19}$.

Proof. Suppose v is the payoff of some strongly symmetric strategy equilibrium s . Then either $s^0 = EE$ or SS . Since the continuation equilibrium payoffs must lie in $[0, 2]$, we immediately have that if $s^0 = EE$, then $v \in \mathcal{W}^{EE}$, while if $s^0 = SS$, then $v \in \mathcal{W}^{SS}$. But this implies $v \in \mathcal{W}^{SS} \cup \mathcal{W}^{EE}$. So, if $[0, 2]$ is the set of strongly symmetric strategy equilibrium payoffs, we must have

$$[0, 2] \subset \mathcal{W}^{SS} \cup \mathcal{W}^{EE}.$$

■

So, when is

$$[0, 2] \subset \mathcal{W}^{SS} \cup \mathcal{W}^{EE}?$$

This holds iff $2\delta \geq 3(1 - \delta)$ (i.e., $\delta \geq \frac{3}{5}$).

Lemma 7.3.2 (Sufficiency) *If*

$$[0, 2] = \mathcal{W}^{SS} \cup \mathcal{W}^{EE},$$

then $[0, 2]$ is the set of strongly symmetric strategy equilibrium payoffs.

Proof. Fix $v \in [0, 2]$, and define a recursion as follows: set $y^0 = v$, and

$$y^{t+1} = \begin{cases} (y^t - 2(1 - \delta))/\delta & \text{if } y^t \in \mathcal{W}^{EE} = [3(1 - \delta), 2], \text{ and} \\ y^t/\delta & \text{if } y^t \in \mathcal{W}^{SS} \setminus \mathcal{W}^{EE} = [0, 3(1 - \delta)). \end{cases}$$

Since $[0, 2] \subset \mathcal{W}^{SS} \cup \mathcal{W}^{EE}$, this recursive definition is well defined for all t . Moreover, since $\delta \geq \frac{3}{5}$, $y^t \in [0, 2]$ for all t . The recursion thus yields a *bounded* sequence of continuations $\{y^t\}_t$. Associated with this sequence of continuations is the outcome path $\{\tilde{a}^t\}_t$:

$$\tilde{a}^t = \begin{cases} EE & \text{if } y^t \in \mathcal{W}^{EE}, \text{ and} \\ SS & \text{if } y^t \in \mathcal{W}^{SS} \setminus \mathcal{W}^{EE}. \end{cases}$$

Observe that, by construction,

$$y^t = (1 - \delta)(u_i(\tilde{a}^t) + \delta y^{t+1}).$$

Consider the automaton $(\mathcal{W}, w^0, f, \tau)$ where

- $\mathcal{W} = [0, 2]$;
- $w^0 = v$;
- the output function is

$$f(w) = \begin{cases} EE & \text{if } w \in \mathcal{W}^{EE}, \text{ and} \\ SS & \text{if } w \in \mathcal{W}^{SS} \setminus \mathcal{W}^{EE}, \text{ and} \end{cases}$$

- the transition function is

$$\tau(w, a) = \begin{cases} (w - 2(1 - \delta))/\delta & \text{if } w \in \mathcal{W}^{EE} \text{ and } a = f(w), \\ w/\delta & \text{if } w \in \mathcal{W}^{SS} \setminus \mathcal{W}^{EE} \text{ and } a = f(w), \text{ and} \\ 0, & \text{if } a \neq f(w). \end{cases}$$

The outcome path implied by this strategy profile is $\{\tilde{a}^t\}_t$. Moreover,

$$\begin{aligned} v = y^0 &= (1 - \delta)u_i(\tilde{a}^0) + \delta y^1 \\ &= (1 - \delta)u_i(\tilde{a}^0) + \delta \{(1 - \delta)u_i(\tilde{a}^1) + \delta y^2\} \\ &= (1 - \delta) \sum_{t=0}^{T-1} \delta^t u_i(\tilde{a}^t) + \delta^T y^T \\ &= (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(\tilde{a}^t) \end{aligned}$$

(where the last equality is an implication of $\delta < 1$ and the sequence $\{y^t\}_T$ being bounded). Thus, the payoff of this outcome path is exactly v , that is, v is the payoff of the strategy profile described by the automaton $(\mathcal{W}, w^0, f, \tau)$ with initial state $w^0 = v$.

Thus, there is no profitable one-deviation from this automaton (this is guaranteed by the constructions of \mathcal{W}^{SS} and \mathcal{W}^{EE} and $w \in \mathcal{W}^{SS} \setminus \mathcal{W}^{EE}$ for $w \in [0, 2]$). Consequently the associated strategy profile is subgame perfect. ■

See Mailath and Samuelson (2006, chapter 2) for much more on this. ★

7.4 Efficiency Wages


A slight modification of Section 2.3.D in Gibbons.⁴ In the stage game, the worker (player *I*) first decides whether to be self-employed or to exert effort (*E*) for the firm (player *II*), or to shirk (*S*) for the firm. Effort yields output y for sure, while shirking yields output y with probability p , and output 0 with probability $1 - p$. The firm chooses a wage w that period if the worker turns up. At the end of the period, the firm does *not* observe effort, but does observe output.

Suppose

$$y - e > \max\{w_0, py\}$$


so it is efficient for the worker to exert effort.

Payoffs: self-employed worker receives $w_0 > 0$, effort has a disutility of e :



Consider the profile described by the automaton,

⁴The modification ensures that the firm faces an intertemporal trade-off.



with output function

$$f(H) = (e, w^*) \quad \text{and} \quad f(L) = (self, 0),$$

where w^* remains to be determined.

The value functions are

$$V_1(L) = w_0, \quad V_1(H) = w^* - e,$$

$$V_2(L) = 0, \quad V_2(H) = y - w^*,$$

In the absorbing state L , play is the unique eq of the stage game, and so incentives are trivially satisfied.

The worker does not wish to deviate in H if

$$V_1(H) \geq (1 - \delta)w^* + \delta\{pV_1(H) + (1 - p)w^0\},$$

i.e.,

$$\delta(1 - p)(w^* - w_0) \geq (1 - \delta p)e$$

or

$$w^* \geq w_0 + \frac{1 - \delta p}{\delta(1 - p)}e = w_0 + e + \frac{1 - \delta}{\delta(1 - p)}e.$$

Note that this also implies $w^* - e \geq w_0$, and so choosing *self* is also not a profitable deviation.

The firm does not wish to deviate in H if

$$V_2(H) \geq (1 - \delta)y,$$

i.e.,

$$y - w^* \geq (1 - \delta)y \Leftrightarrow \delta y \geq w^*.$$

So, the profile is an “equilibrium” if

$$\delta y \geq w^* \geq w_0 + e + \frac{1 - \delta}{\delta(1 - p)}e.$$

In fact, it is implication of the next section that the profile is a perfect public equilibrium.

7.5 Imperfect Public Monitoring

As before, action space for i is A_i , with typical action $a_i \in A_i$. An action profile is denoted $a = (a_1, \dots, a_n)$.

At the end of each period, rather than observing a , all players observe a *public* signal y taking values in some space Y according to the distribution $\Pr\{y | (a_1, \dots, a_n)\} \equiv \rho(y | a)$.

Since the signal y is a possibly noisy signal of the action profile a in that period, the actions are *imperfectly monitored* by the other players. Since the signal is public (and so observed by all players), the game is said to have *public monitoring*.

Public history up to date t : $h^t \equiv (a^0, \dots, a^{t-1}) \in A^t \equiv H^t$; $H^0 \equiv \{\emptyset\}$.

Assume Y is finite.

$u_i^* : A_i \times Y \rightarrow \mathbb{R}$, i 's *ex post* or realized payoff.

Stage game (*ex ante*) payoffs:

$$u_i(a) \equiv \sum_{y \in Y} u_i^*(a_i, y) \rho(y | a).$$

Public histories:

$$H \equiv \cup_{t=0}^{\infty} Y^t,$$

with $h^t \equiv (y^0, \dots, y^{t-1})$ being a t period history of public signals ($Y^0 \equiv \{\emptyset\}$).

Public strategies:

$$s_i : H \rightarrow A_i.$$

Definition 7.5.1 A perfect public equilibrium is a profile of public strategies s that, after observing any public history h^t , specifies a Nash equilibrium for the repeated game, i.e., for all t and all $h^t \in Y^t$, $\sigma|_{h^t}$ is a Nash equilibrium.

If $\rho(y|a) > 0$ for all y and a , every public history arises with positive probability, and so every Nash equilibrium in public strategies is a perfect public equilibrium.

Automaton representation of public strategies: $(\mathcal{W}, w^0, f, \tau)$, where

- \mathcal{W} is set of states,

- w^0 is initial state,
- $f : \mathcal{W} \rightarrow A$ is output function (decision rule), and
- $\tau : \mathcal{W} \times Y \rightarrow \mathcal{W}$ is transition function.

As before, $V_i(w)$ is i 's value of being in state w .

Lemma 7.5.1 *Suppose the strategy profile s is represented by $(\mathcal{W}, w^0, f, \tau)$. Then s is a perfect public eq if, and only if, for all $w \in \mathcal{W}$ (satisfying $w = \tau(w^0, h^t)$ for some $h^t \in H$), $f(w)$ is a Nash eq of the normal form game with payoff function $g^w : A \rightarrow \mathbb{R}^n$, where*

$$g_i^w(a) = (1 - \delta)u_i(a) + \delta \sum_y V_i(\tau(w, y))\rho(y|a).$$

See Problem 7.6.12 for the proof.

Example 7.5.1 (PD as a partnership) Effort determines output $\{\underline{y}, \bar{y}\}$ stochastically:

$$\Pr\{\bar{y} | a\} \equiv \rho(\bar{y} | a) = \begin{cases} p, & \text{if } a = EE, \\ q, & \text{if } a = SE \text{ or } ES, \\ r, & \text{if } a = SS, \end{cases}$$

where $0 < q < p < 1$ and $0 < r < p$.

Ex post payoffs (u_i^*):

	\bar{y}	\underline{y}
E	$\frac{(3-p-2q)}{(p-q)}$	$-\frac{(p+2q)}{(p-q)}$
S	$\frac{3(1-r)}{(q-r)}$	$-\frac{3r}{(q-r)}$


so that ex ante payoffs (u_i) are:

	E	S
E	2, 2	-1, 3
S	3, -1	0, 0

★

Example 7.5.2 (One period memory) Two state automaton: $\mathcal{W} = \{w_{EE}, w_{SS}\}$, $w^0 = w_{EE}$, $f(w_{EE}) = EE$, $f(w_{SS}) = SS$, and

$$\tau(w, y) = \begin{cases} w_{EE}, & \text{if } y = \bar{y}, \\ w_{SS}, & \text{if } y = \underline{y}. \end{cases}$$



Value functions (I can drop player subscripts by symmetry):

$$V(w_{EE}) = (1 - \delta) \cdot 2 + \delta \{pV(w_{EE}) + (1 - p)V(w_{SS})\}$$

and

$$V(w_{SS}) = (1 - \delta) \cdot 0 + \delta \{rV(w_{EE}) + (1 - r)V(w_{SS})\}.$$

This is eq if

$$V(w_{EE}) \geq (1 - \delta) \cdot 3 + \delta \{qV(w_{EE}) + (1 - q)V(w_{SS})\}$$

and

$$V(w_{SS}) \geq (1 - \delta) \cdot (-1) + \delta \{qV(w_{EE}) + (1 - q)V(w_{SS})\}.$$

Rewriting the incentive constraint at w_{EE} ,

$$\begin{aligned} (1 - \delta) \cdot 2 + \delta \{pV(w_{EE}) + (1 - p)V(w_{SS})\} \\ \geq (1 - \delta) \cdot 3 + \delta \{qV(w_{EE}) + (1 - q)V(w_{SS})\} \end{aligned}$$

or

$$\delta(p - q)\{V(w_{EE}) - V(w_{SS})\} \geq (1 - \delta).$$

We can obtain an expression for $V(w_{EE}) - V(w_{SS})$ without solving for the value functions separately by differencing the value recursion equations, yielding

$$\begin{aligned} V(w_{EE}) - V(w_{SS}) &= (1 - \delta) \cdot 2 + \delta \{pV(w_{EE}) + (1 - p)V(w_{SS})\} \\ &\quad - \delta \{rV(w_{EE}) + (1 - r)V(w_{SS})\} \\ &= (1 - \delta) \cdot 2 + \delta(p - r)\{V(w_{EE}) - V(w_{SS})\}, \end{aligned}$$

so that

$$V(w_{EE}) - V(w_{SS}) = \frac{2(1 - \delta)}{1 - \delta(p - r)},$$

and so

$$\delta \geq \frac{1}{3p - 2q - r}.$$

Turning to w_{SS} , we have

$$\begin{aligned} &\delta \{rV(w_{EE}) + (1 - r)V(w_{SS})\} \\ &\geq (1 - \delta) \cdot (-1) + \delta \{qV(w_{EE}) + (1 - q)V(w_{SS})\} \end{aligned}$$

or

$$(1 - \delta) \geq \delta(q - r)\{V(w_{EE}) - V(w_{SS})\},$$

requiring

$$\delta \leq \frac{1}{p + 2q - 3r}.$$

The two bounds on δ are consistent if

$$p \geq 2q - r.$$

Solving for the value functions,

$$\begin{aligned} \begin{bmatrix} V(w_{EE}) \\ V(w_{SS}) \end{bmatrix} &= (1 - \delta) \begin{bmatrix} 1 - \delta p & -\delta(1 - p) \\ -\delta r & 1 - \delta(1 - r) \end{bmatrix}^{-1} \begin{bmatrix} 2 \\ 0 \end{bmatrix} \\ &= \frac{(1 - \delta)}{(1 - \delta p)(1 - \delta(1 - r)) - \delta^2(1 - p)r} \times \\ &\quad \begin{bmatrix} 1 - \delta(1 - r) & \delta(1 - p) \\ \delta r & 1 - \delta p \end{bmatrix} \begin{bmatrix} 2 \\ 0 \end{bmatrix} \end{aligned}$$

$$\begin{aligned}
&= \frac{(1-\delta)}{(1-\delta)(1-\delta(p-r))} \begin{bmatrix} 2(1-\delta(1-r)) \\ 2\delta r \end{bmatrix} \\
&= \frac{1}{1-\delta(p-r)} \begin{bmatrix} 2(1-\delta(1-r)) \\ 2\delta r \end{bmatrix}.
\end{aligned}$$

Moreover, for fixed p and r ,

$$\lim_{\delta \rightarrow 1} V(w_{EE}) = \lim_{\delta \rightarrow 1} V(w_{SS}) = \frac{2r}{1-p+r},$$

and, for $r > 0$,

$$\lim_{p \rightarrow 1} \lim_{\delta \rightarrow 1} V(w_{EE}) = 2. \quad \star$$

Remark 7.5.1 The notion of PPE only imposes ex ante incentive constraints. If the stage game has a non-trivial dynamic structure, such as Problem 7.6.15, then it is natural to impose additional incentive constraints.

◆

7.6 Problems

- 7.6.1. Suppose $G \equiv \{(A_i, u_i)\}$ is an n -person normal form game and G^T is its T -fold repetition (with payoffs evaluated as the average). Let $A \equiv \prod_i A_i$. The strategy profile, s , is *history independent* if for all i and all $1 \leq t \leq T-1$, $s_i(h^t)$ is independent of $h^t \in A^t$ (i.e., $s_i(h^t) = s_i(\hat{h}^t)$ for all $h^t, \hat{h}^t \in A^t$). Let $N(1)$ be the set of Nash equilibria of G . Suppose s is history independent. Prove that s is a subgame perfect equilibrium *if and only if* $s(h^t) \in N(1)$ for all t , $0 \leq t \leq T-1$ and all $h^t \in A^t$ ($s(h^0)$ is of course simply s^0). Provide examples to show that the assumption of history independence is needed in both directions.
- 7.6.2. Prove the infinitely repeated game with stage game given by matching pennies does not have a pure strategy Nash equilibrium for any δ .

7.6.3. Suppose $(\mathcal{W}, w^0, f, \tau)$ is a (pure strategy representing) finite automaton with $|\mathcal{W}| = K$. Label the states from 1 to K , so that $\mathcal{W} = \{1, 2, \dots, K\}$, $f : \{1, 2, \dots, K\} \rightarrow A$, and $\tau : \{1, 2, \dots, K\} \times A \rightarrow \{1, 2, \dots, K\}$. Consider the function $\Phi : \mathbb{R}^K \rightarrow \mathbb{R}^K$ given by $\Phi(v) = (\Phi_1(v), \Phi_2(v), \dots, \Phi_K(v))$, where

$$\Phi_k(v) = (1 - \delta)u_i(f(k)) + \delta v_{\tau(k, f(k))}, \quad k = 1, \dots, K.$$

- (a) Prove that Φ has a unique fixed point. [**Hint:** Show that Φ is a contraction.]
- (b) Given an explicit equation for the fixed point of Φ .
- (c) Interpret the fixed point.

7.6.4. A different (and perhaps more enlightening) proof of Theorem 7.1.3 is the following:

Suppose \mathcal{W} and A_i are finite. Let $\tilde{V}_i(w)$ be player i 's payoff from the best response to $(\mathcal{W}, w, f_{-i}, \tau)$ (i.e., the strategy profile for the other players specified by the automaton with initial state w).

Prove that

$$\tilde{V}_i(w) = \max_{a_i \in A_i} \left\{ (1 - \delta)u_i(a_i, f_{-i}(w)) + \delta \tilde{V}_i(\tau(w, (a_i, f_{-i}(w)))) \right\}.$$

Note that $\tilde{V}_i(w) \geq V_i(w)$ for all w .

Let a_i^w be the action solving the above maximization, and define

$$V_i^\dagger(w) = (1 - \delta)u_i(a_i^w, f_{-i}(w)) + \delta V_i(\tau(w, (a_i^w, f_{-i}(w)))).$$

Prove that if s is not subgame perfect, then there exists a player i and a state w satisfying

$$V_i^\dagger(w) > V_i(w).$$

What is the profitable one-shot deviation?

Extend the argument to infinite \mathcal{W} and A_i .

7.6.5. Suppose two players play the infinitely repeated prisoners' dilemma:

	E	S
E	1, 1	$-\ell, 1 + g$
S	$1 + g, -\ell$	0, 0

where $\ell > 0$ and $g > 0$.

- (a) For what values of the discount factor δ is grim trigger a subgame perfect equilibrium?
- (b) Describe a simple automaton representation of the behavior in which player *I* alternates between *E* and *S* (beginning with *E*), player *II* always plays *E*, and any deviation results in permanent *SS*. For what parameter restrictions is this a subgame perfect equilibrium?
- (c) For what parameter values of ℓ , g , and δ is tit-for-tat a subgame perfect equilibrium?

7.6.6. Suppose the following game is infinitely repeated:

	<i>L</i>	<i>R</i>
<i>U</i>	2, 2	$x, 0$
<i>D</i>	0, 5	1, 1

Let δ denote the common discount factor for both players and consider the strategy profile that induces the outcome path DL, UR, DL, UR, \dots , and that, after any unilateral deviation by the row player specifies the outcome path DL, UR, DL, UR, \dots , and after any unilateral deviation by the column player, specifies the outcome path UR, DL, UR, DL, \dots (simultaneous deviations are ignored. i.e., are treated as if neither player had deviated).

- (a) What is the simplest automaton that represents this strategy profile?
- (b) Suppose $x = 5$. For what values of δ is this strategy profile subgame perfect?
- (c) Suppose now $x = 4$. How does this change your answer to part 7.6.6(b)?
- (d) Suppose $x = 5$ again. How would the analysis in part 7.6.6(b) be changed if the column player were short-lived (lived for only one period)?

7.6.7. Fix a stage game $G = \{(A_i, u_i)\}$ and discount factor δ . Let $\mathcal{E}^p(\delta) \subset \mathcal{F}^{p*}$ be the set of pure strategy subgame perfect equilibrium payoffs. Suppose $\gamma : A \rightarrow \mathcal{E}^p(\delta)$ enforces the action profile a . Describe a pure strategy profile in which a is played in the first period, and prove that is subgame perfect equilibrium.

7.6.8. Consider the prisoner's dilemma:

	E	S
E	2, 2	-1, 3
S	3, -1	0, 0

Suppose the game is infinitely repeated with perfect monitoring. Recall that a *strongly symmetric* strategy profile (s_1, s_2) satisfies $s_1(h^t) = s_2(h^t)$ for all h^t . Equivalently, its automaton representation satisfies $f_1(w) = f_2(w)$ for all w . Let $\mathcal{W} = \{\delta v, v\}$, $v > 0$ to be determined, be the set of continuation promises. Describe a strongly symmetric strategy profile (equivalently, automaton) whose continuation promises come from \mathcal{W} which is a subgame perfect equilibrium for some values of δ . Calculate the appropriate bounds on δ and the value of v (which may or may not depend on δ).

7.6.9. Describe the four state automaton that yields v^0 as a strongly symmetric equilibrium payoff with the indicated cycle in Figure 7.3.1.

7.6.10. Consider the following (asymmetric) prisoner's dilemma:

	E	S
E	1, 2	-1, 3
S	2, -4	0, 0

Suppose the game is infinitely repeated with perfect monitoring. Prove that for $\delta < \frac{1}{2}$, the maximum (average discounted) payoff to player 1 in any pure strategy subgame perfect equilibrium is 0, while for $\delta = \frac{1}{2}$, there are equilibria in which player 1 receives a payoff of 1. **[Hint:** First prove that, if $\delta \leq \frac{1}{2}$, in any pure strategy subgame perfect equilibrium, in any period, if player 2 chooses E then player 1 chooses E in that period.]

7.6.11. Consider the stage game where player 1 is the row player and 2, the column player (as usual):

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

- (a) Suppose the game is infinitely repeated, with perfect monitoring. Players 1 and 2 are both long-lived, and have the same discount factor, $\delta \in (0, 1)$. Construct a three state automaton that for large δ is a subgame perfect equilibrium, and yields a payoff to player 1 that is close to $2\frac{1}{2}$. Prove that the automaton has the desired properties. (**Hint:** One state is only used off the path-of-play.)
- (b) Now suppose that player 2 is short-lived (but maintain the assumption of perfect monitoring, so that the short-lived player in period t knows the entire history of actions up to t). Prove that player 1's payoff in any pure strategy subgame perfect equilibrium is no greater than 2 (the restriction to pure strategy is not needed—can you prove the result without that restriction?). For which values of δ is there a pure strategy subgame perfect equilibrium in which player 1 receives a payoff of precisely 2?
- 7.6.12. Fix a repeated finite game of imperfect public monitoring. Say that a player has a *profitable one-shot deviation* from the public strategy $(\mathcal{W}, w^0, f, \tau)$ if there is some history $h^t \in Y^t$ and some action $a_i \in A_i$ such that (where $w = \tau(w^0, h^t)$)
- $$V_i(w) < (1 - \delta)u_i(a_i, f_{-i}(w)) + \delta \sum_{\mathcal{Y}} V_i(\tau(w, y))\rho(y \mid (a_i, f_{-i}(w))).$$
- (a) Prove that a public strategy profile is a perfect public equilibrium if and only if there are no profitable one-shot deviations.
- (b) Prove Lemma 7.5.1.
- 7.6.13. Consider the prisoners' dilemma game in Example 7.5.1.
- (a) The grim trigger profile is described by the automaton $(\mathcal{W}, w^0, f, \tau)$, where $\mathcal{W} = \{w_{EE}, w_{SS}\}$, $w^0 = w_{EE}$, $f(w_a) = a$, and
- $$\tau(w, y) = \begin{cases} w_{EE}, & \text{if } w = w_{EE} \text{ and } y = \bar{y}, \\ w_{SS}, & \text{otherwise.} \end{cases}$$
- For what parameter values is the grim-trigger profile an equilibrium?

- (b) An example of a forgiving grim trigger profile is described by the automaton $(\widehat{W}, \hat{w}^0, \hat{f}, \hat{\tau})$, where $\widehat{W} = \{\hat{w}_{EE}, \hat{w}'_{EE}, \hat{w}_{SS}\}$, $\hat{w}^0 = \hat{w}_{EE}$, $\hat{f}(w_a) = a$, and

$$\hat{\tau}(w, y) = \begin{cases} \hat{w}_{EE}, & \text{if } w = \hat{w}_{EE} \text{ or } \hat{w}'_{EE}, \text{ and } y = \bar{y}, \\ \hat{w}'_{EE}, & \text{if } w = \hat{w}_{EE} \text{ and } y = \underline{y}, \\ \hat{w}_{SS}, & \text{otherwise.} \end{cases}$$

For what parameter values is this forgiving grim-trigger profile an equilibrium? Compare the payoffs of grim trigger and this forgiving grim trigger when both are equilibria.

- 7.6.14. Player 1 (the row player) is a firm who can exert either high effort (H) or low effort (L) in the production of its output. Player 2 (the column player) is a consumer who can buy either a high-priced product, h , or a low-priced product ℓ . The actions are chosen simultaneously, and payoffs are given by:

	h	ℓ
H	4, 3	0, 2
L	$x, 0$	3, 1

Player 1 is infinitely lived, discounts the future with discount factor δ , and plays the above game in every period with a different consumer (i.e., each consumer lives only one period). The game is one of *public monitoring*: while the actions of the consumers are public, the actions of the firm are not. Both the high-priced and low-priced products are *experience* goods of random quality, with the distribution over quality determined by the effort choice. The consumer learns the quality of the product after purchase (consumption). Denote by \bar{y} the event that the product purchased is high quality, and by \underline{y} the event that it is low quality (in other words, $y \in \{\bar{y}, \underline{y}\}$ is the quality signal). Assume the distribution over quality is independent of the price of the product:

$$\Pr(\bar{y} \mid a) = \begin{cases} p, & \text{if } a_1 = H, \\ q, & \text{if } a_1 = L, \end{cases}$$

with $0 < q < p < 1$.

- (a) Describe the ex post payoffs for the consumer. Why can the ex post payoffs for the firm be taken to be the ex ante payoffs?
 - (b) Suppose $x = 5$. Describe a perfect public equilibrium in which the patient firm chooses H infinitely often with probability one, and verify that it is an equilibrium. [**Hint:** This can be done with one-period memory.]
 - (c) Suppose now $x \geq 8$. Is the one-period memory strategy profile still an equilibrium? If not, can you think of an equilibrium in which H is still chosen with positive probability?
- 7.6.15. A financial manager undertakes an infinite sequence of trades on behalf of a client. Each trade takes one period. In each period, the manager can invest in one of a large number of risky assets. By exerting effort ($a = E$) in a period (at a cost of $e > 0$), the manager can identify the most profitable risky asset for that period, which generates a high return of $R = H$ with probability p and a low return $R = L$ with probability $1 - p$. In the absence of effort ($a = S$), the manager cannot distinguish between the different risky assets. For simplicity, assume the manager then chooses the wrong asset, yielding the low return $R = L$ with probability 1; the cost of no effort is 0. In each period, the client chooses the level of the fee $x \in [0, \bar{x}]$ to be paid to the financial manager for that period. Note that there is an exogenous upper bound \bar{x} on the fee that can be paid in a period. The client and financial manager are risk neutral, and so the client's payoff in a period is

$$u_c(x, R) = R - x,$$

while the manager's payoff in a period is

$$u_m(x, a) = \begin{cases} x - e, & \text{if } a = E, \\ x, & \text{if } a = S. \end{cases}$$

The client and manager have a common discount factor δ . The client observes the return on the asset prior to paying the fee, but does not observe the manager's effort choice.

- (a) Suppose the client cannot sign a binding contract committing him to pay a fee (contingent or not on the return). Describe the unique sequentially rational equilibrium when the client uses the manager for a single transaction. Are there any other Nash equilibria?

- (b) Continue to suppose there are no binding contracts, but now consider the case of an infinite sequence of trades. For a range of values for the parameters (δ , \bar{x} , e , p , H , and L), there is a perfect public equilibrium in which the manager exerts effort on behalf of the client in *every* period. Describe it and the restrictions on parameters necessary and sufficient for it to be an equilibrium.
- (c) Compare the fee paid in your answer to part 7.6.15(b) to the fee that would be paid by a client for a single transaction,
 - i. when the client *can* sign a legally binding commitment to a fee schedule as a function of the return of that period, and
 - ii. when the client can sign a legally binding commitment to a fee schedule as a function of effort.
- (d) Redo question 7.6.15(b) assuming that the client's choice of fee level and the manager's choice of effort are simultaneous, so that the fee paid in period t cannot depend on the return in period t . Compare your answer with that to question 7.6.15(b).

7.6.16. In this question, we revisit the partnership game of Example 7.5.1. Suppose $3p - 2q > 1$. This question asks you to prove that for sufficiently large δ , any payoff in the interval $[0, \bar{v}]$, is the payoff of some strongly symmetric PPE equilibrium, where

$$\bar{v} = 2 - \frac{(1-p)}{(p-q)},$$

and that no payoff larger than \bar{v} is the payoff of some strongly symmetric PPE equilibrium. Strong symmetry implies it is enough to focus on player 1, and the player subscript will often be omitted.

- (a) The action profile SS is trivially *enforced* by any constant continuation $y \in [0, \bar{y}]$ independent of y . Let \mathcal{W}^{SS} be the set of values that can be obtained by SS and a constant continuation $y \in [0, \bar{y}]$, i.e.,

$$\mathcal{W}^{SS} = \{(1-\delta)u_1(SS) + \delta y : y \in [0, \bar{y}]\}.$$

Prove that $\mathcal{W}^{SS} = [0, \delta\bar{y}]$. [This is almost immediate.]

- (b) Recalling Definition 7.1.6, say that v is *decomposed* by EE on $[0, \bar{y}]$ if there exists $y^{\bar{y}}, y^{\mathcal{Y}} \in [0, \bar{y}]$ such that

$$v = (1-\delta)u_1(EE) + \delta\{py^{\bar{y}} + (1-p)y^{\mathcal{Y}}\} \quad (7.6.1)$$

$$\geq (1 - \delta)u_1(SE) + \delta\{q\gamma^{\bar{y}} + (1 - q)\gamma^{\underline{y}}\}. \quad (7.6.2)$$

(That is, EE is *enforced* by the continuation promises $\gamma^{\bar{y}}, \gamma^{\underline{y}}$ and implies the value v .) Let \mathcal{W}^{EE} be the set of values that can be decomposed by EE on $[0, \bar{y}]$. It is clear that $\mathcal{W}^{EE} = [\gamma', \gamma'']$, for some γ' and γ'' . Calculate γ' by using the smallest possible choices of $\gamma^{\bar{y}}$ and $\gamma^{\underline{y}}$ in the interval $[0, \bar{y}]$ to enforce EE . (This will involve having the inequality (7.6.2) holding with equality.)

- (c) Similarly, give an expression for γ'' (that will involve \bar{y}) by using the largest possible choices of $\gamma^{\bar{y}}$ and $\gamma^{\underline{y}}$ in the interval $[0, \bar{y}]$ to enforce EE . Argue that $\delta\bar{y} < \gamma''$.
- (d) As in Example 7.3.1, we would like all continuations in $[0, \bar{y}]$ to be themselves decomposable using continuations in $[0, \bar{y}]$, i.e., we would like

$$[0, \bar{y}] \subset \mathcal{W}^{SS} \cup \mathcal{W}^{EE}.$$

Since $\delta\bar{y} < \gamma''$, we then would like $\bar{y} \leq \gamma''$. Moreover, since we would like $[0, \bar{y}]$ to be the largest such interval, we have $\bar{y} = \gamma''$. What is the relationship between γ'' and \bar{v} ?

- (e) For what values of δ do we have $[0, \bar{y}] = \mathcal{W}^{SS} \cup \mathcal{W}^{EE}$?
- (f) Let $(\mathcal{W}, w^0, f, \tau)$ be the automaton given by $\mathcal{W} = [0, \bar{v}]$, $w^0 \in [0, \bar{v}]$,

$$f(w) = \begin{cases} EE, & \text{if } w \in \mathcal{W}^{EE}, \\ SS, & \text{otherwise,} \end{cases}$$

and

$$\tau(w, \gamma) = \begin{cases} \gamma^{\underline{y}}(w), & \text{if } w \in \mathcal{W}^{EE}, \\ w/\delta, & \text{otherwise,} \end{cases}$$

where $\gamma^{\underline{y}}(w)$ solves (7.6.1)–(7.6.2) for $w = v$ and $\gamma = \bar{y}, \underline{y}$. For our purposes here, assume that $V(w) = w$, that is, the value to a player of being in the automaton with initial state w is precisely w . (From the argument of Lemma 7.3.2, this should be intuitive.) Given this assumption, prove that the automaton describes a PPE with value w^0 .

Chapter 8

Topics in Dynamic Games¹

8.1 Dynamic Games and Markov Perfect Equilibria

Set of players: $\{1, \dots, n\}$.

Action space for i is A_i .

Set of states S , with typical state $s \in S$.

Payoffs for each i :

$$u_i : S \times A \rightarrow \mathbb{R},$$

with future flow payoffs discounted at rate $\delta \in (0, 1)$.

State transitions:

$$q : S \times A \rightarrow S,$$

and initial state $s_0 \in S$. (More generally, can have random transitions from $S \times A$ into $\Delta(S)$, but deterministic transitions suffice for an introduction.)

Example 8.1.1 Suppose players 1 and 2 fish from a common area (pool). In each period t , the pool contains a stock of fish of size $s^t \in \mathbb{R}_+$. This is the state of the game.

In period t , player i attempts to extract $a_i^t \geq 0$ units of fish. In particular, if player i attempts to extract a_i^t , then player i actually

¹Copyright December 7, 2011 by George J. Mailath

extracts

$$\hat{a}_i^t = \begin{cases} a_i^t, & \text{if } a_1^t + a_2^t \leq s^t, \\ \frac{a_i^t}{a_1^t + a_2^t} s^t, & \text{if } a_1^t + a_2^t > s^t. \end{cases}$$

It will turn out that in equilibrium, $a_1^t + a_2^t < s^t$, so we can ignore the rationing rule (see Problem 8.4.1(a)) and assume that i derives payoff

$$u_i(s, a) = \log(a_i)$$

from (s, a) for all values of (s, a) .

The transition rule is

$$q(s, a) = 2 \max\{0, s - a_1 - a_2\},$$

that is, it is deterministic and doubles any leftover stock after extraction. The initial stock is fixed at some value s^0 . ★

State is public and perfect monitoring of actions, so history to period t is

$$h^t = (s^0, a^0, s^1, a^1, \dots, s^{t-1}, a^{t-1}, s^t) \in (S \times A)^t \times S.$$

Let H^t denote the set of all feasible t -period histories (so that s^τ is consistent with $(s^{\tau-1}, a^{\tau-1})$ for all $1 \leq \tau \leq t$). A pure strategy for i is a mapping

$$\sigma_i : \cup_t H^t \rightarrow A_i.$$

For any history h^t , write the function that identifies the last state s^t by $s(h^t)$. Let $G(s)$ denote the dynamic game with initial state s . As usual, we have:

Definition 8.1.1 *The profile σ is a subgame perfect equilibrium if for all h^t , $\sigma|_{h^t} := (\sigma_1|_{h^t}, \dots, \sigma_n|_{h^t})$ is a Nash equilibrium of the dynamic game $G(s(h^t))$.*

Different histories that lead to the same state are effectively “payoff equivalent.” Loosely, a strategy is said to be Markov if at different histories that are effectively payoff equivalent, the strategy specifies identical behavior. See Maskin and Tirole (2001) for a discussion of why this may be a reasonable restriction.

Definition 8.1.2 A strategy $\sigma_i : \cup_t H^t \rightarrow A_i$ is Markov if for all histories h^t and \hat{h}^t , if $s(h^t) = s(\hat{h}^t)$, then

$$\sigma_i(h^t) = \sigma_i(\hat{h}^t).$$

If the above holds for histories h^t and \hat{h}^τ of possibly different length (so that $t \neq \tau$ is allowed), the strategy is stationary.

Restricting equilibrium behavior to Markov strategies:

Definition 8.1.3 A Markov perfect equilibrium is a strategy profile σ that is a subgame perfect equilibrium, and for which each σ_i is Markov.

Note that while there is a superficial similarity between Markov states s and automata states used in the theory of repeated games, they are very different. In particular, a repeated game has a trivial set of Markov states, and the only Markov perfect equilibria involve specifying static Nash equilibria in each period.

Example 8.1.2 (Example 8.1.1 continued) Fix a symmetric MPE. Let $V(s)$ denote the common equilibrium value from the state s (in an MPE, this must be independent of other aspects of the history).

The common eq strategy is $a_1(s) = a_2(s)$.

One-shot deviation principle holds here, and so for each player i , $a_i(s)$ solves, for any $s \in S$, the Bellman equation:

$$a_i(s) \in \arg \max_{\tilde{a}_i \in A_i} (1 - \delta) \log(\tilde{a}_i) + \delta V(2(s - \tilde{a}_i - a_j(s))).$$

Assuming V is differentiable and imposing $a_1 = a_2$ after differentiating, the implied first order condition is

$$\frac{(1 - \delta)}{a_i(s)} = 2\delta V'(2(s - 2a_i(s))).$$

To find an equilibrium, suppose

$$a_i(s) = ks,$$

for some k . Then we have

$$s^{t+1} = 2(s^t - 2ks^t) = 2(1 - 2k)s^t.$$

Given an initial stock s , in period t , $a_i^t = k[2(1 - 2k)]^t s$, and so

$$\begin{aligned} V(s) &= (1 - \delta) \sum_{t=0}^{\infty} \delta^t \log\{k[2(1 - 2k)]^t s\} \\ &= (1 - \delta) \sum_{t=0}^{\infty} \delta^t \log\{k[2(1 - 2k)]^t\} + \log s. \end{aligned}$$

This implies V is indeed differentiable, with $V'(s) = 1/s$.

Solving the first order condition, $k = \frac{1-\delta}{2-\delta}$, and so

$$a_i(s) = \frac{1 - \delta}{2 - \delta} s.$$

★

Example 8.1.3 (Asynchronous move games) Consider the repeated prisoners' dilemma, but where player 1 moves in odd periods only and player 2 moves in even periods only. The game starts with E_1 exogenously and publicly specified for player 1. The stage game is ($x > 0$):

	E_2	S_2
E_1	2, 2	$-x, 3$
S_2	3, $-x$	0, 0

This fits into the above formulation of a repeated game: $S = \{E_1, S_1, E_2, S_2\}$, $s^0 = E_1$,

$$q(s, a) = \begin{cases} a_1, & \text{if } s \in \{E_2, S_2\}, \\ a_2, & \text{if } s \in \{E_1, S_1\}, \end{cases}$$

$$u_1(s, a) = \begin{cases} g_1(a_1, s), & \text{if } s \in \{E_2, S_2\}, \\ g_1(s, a_2), & \text{if } s \in \{E_1, S_1\}, \end{cases}$$

where g_i describes the stage game payoffs from the PD, and

$$u_2(s, a) = \begin{cases} g_2(s, a_2), & \text{if } s \in \{E_1, S_1\}, \\ g_2(a_1, s), & \text{if } s \in \{E_2, S_2\}. \end{cases}$$

In particular, when the current state is player 1's action (i.e., we are in an even period), 1's choice is irrelevant and can be ignored.

Grim Trigger:

$$\sigma_i^{GT}(h^t) = \begin{cases} E_i, & \text{if } h^t = E_1 \text{ or always } E, \\ S_i, & \text{otherwise.} \end{cases}$$

Need to check two classes of information sets: when players are supposed to play E_i , and when they are supposed to play S_i :

1. Optimality of E_i after all E 's:

$$\begin{aligned} 2 &\geq 3(1 - \delta) + \delta \times 0 \\ \Leftrightarrow \delta &\geq \frac{1}{3}. \end{aligned}$$

2. The optimality of S_i after any S_1 or S_2 is trivially true for all δ :

$$0 \geq (-x)(1 - \delta) + \delta \times 0.$$

This equilibrium is not an MPE.

Supporting effort using Markov pure strategies requires a “tit-for-tat” like behavior:

$$\hat{\sigma}_i(h^t) = \begin{cases} E_i, & \text{if } s^t = E_j, \\ S_i, & \text{if } s^t = S_j. \end{cases} \quad (8.1.1)$$

For $t \geq 1$, everything is symmetric. The value when the current state is E_j is

$$V_i(E_j) = 2,$$

while the payoff from a one-shot deviation is

$$3(1 - \delta) + \delta \times 0 = 3(1 - \delta),$$

and so the deviation is not profitable if (as before) $\delta \geq \frac{1}{3}$.

The value when the current state is S_j is

$$V_i(S_j) = 0,$$

while the payoff from a one-shot deviation is (since under the Markov strategy, a deviation to E_i triggers perpetual E_1E_2 ; the earlier deviation is “forgiven”)

$$-x(1 - \delta) + \delta \times 2 = (2 + x)\delta - x.$$

The deviation is not profitable if

$$\begin{aligned} (2 + x)\delta - x &\leq 0 \\ \Leftrightarrow \delta &\leq \frac{x}{2 + x}. \end{aligned}$$

Note that

$$\frac{x}{2 + x} \geq \frac{1}{3} \Leftrightarrow x \geq 1.$$

Thus, $\hat{\sigma}$ is an MPE (inducing the outcome path $(E_1E_2)^\infty$) if $x \geq 1$ and

$$\frac{1}{3} \leq \delta \leq \frac{x}{2 + x}. \quad (8.1.2)$$

★

8.2 Disappearance of Monopoly Power and the Coase Conjecture

Uninformed seller selling to informed buyer—one sided offers, seller makes all the offers.

Seller’s cost (value) is zero.

Buyer values good at v . Assume v is uniformly distributed on $[0, 1]$, so that there are buyers with valuations arbitrarily close to the seller’s valuation.

8.2.1 One and Two Period Example

Buyer accepts a take-it-or-leave-it offer of p if $v > p$ and rejects if $v < p$.

Seller chooses p to maximize

$$p \Pr\{\text{sale}\} = p(1 - p),$$

i.e., chooses $p = 1/2$, for a payoff of $1/4$. This is the optimal seller mechanism (this can be easily shown using standard mechanism design techniques).

Suppose now two periods, with common discount factor $\delta \in (0, 1)$. If seller chose $p^0 = 1/2$ in the first period, and buyers with $v > 1/2$ buy in period 0, then buyers with value $v \in [0, 1/2]$ are left, and then optimal for seller to price $p^1 = 1/4$. But then buyer $v = 1/2$ strictly prefers to wait till period 1 (and so by continuity so do some buyers with $v > 1/2$).

Suppose seller makes offers p^t in period t , $t = 0, 1$, and buyers $v < \kappa$ don't buy in period 0. Then, $p^1 = \kappa/2$. If $\kappa < 1$, then κ should be indifferent between purchasing in period 0 and period 1, so that

$$\begin{aligned} \kappa - p^0 &= \delta(\kappa - p^1) \\ &= \delta\kappa/2 \\ \Rightarrow p^0 &= \kappa(1 - \delta/2) < \kappa. \end{aligned}$$

The seller's payoff (as a function of κ) is

$$\kappa(1 - \delta/2)(1 - \kappa) + \delta\kappa^2/4.$$

The first order condition is

$$\begin{aligned} (1 - \delta/2) - 2(1 - \delta/2)\kappa + \delta\kappa/2 &= 0 \\ \Rightarrow \kappa &= \frac{2 - \delta}{4 - 3\delta} (< 1) \\ \Rightarrow p^0 &= \frac{(2 - \delta)^2}{8 - 6\delta} < \frac{1}{2}. \end{aligned}$$

The resulting payoff is

$$\frac{(2 - \delta)^2}{4(4 - 3\delta)} < \frac{1}{4}.$$

8.2.2 Infinite Horizon

Seller makes offers p^t in period $t = 0, 1, \dots, T, T \leq \infty$.

After each offer, buyer Accepts or Rejects.

If agreement in period t at price p^t , payoff to seller is

$$u_s = \delta^t p^t,$$

and payoff to buyer is

$$u_b = \delta^t (v - p^t).$$

Interested in equilibrium of following form:

If p^t offered in period t , types $v \geq \lambda p^t$ accept and types $v < \lambda p^t$ reject, where $\lambda > 1$.

If at time t , seller's posterior beliefs are uniform on $[0, \kappa]$, seller offers $p^t(\kappa) = \gamma \kappa$, where $\gamma < 1$.

Natural to treat κ as a state variable.

Under this profile, $p^0 = \gamma$ and seller's posterior entering period 1 is $[0, \gamma\lambda]$, so in order for profile to be well defined, $\gamma\lambda < 1$. Thus, $p^1 = \gamma(\gamma\lambda) = \gamma^2\lambda$ and seller's posterior entering period 2 is $[0, \gamma^2\lambda^2]$. Prices are thus falling exponentially, with $p^t = \gamma^{t+1}\lambda^t$.

Let $U_s(\kappa)$ be the discounted expected value to the seller, when his posterior beliefs are uniform on $[0, \kappa]$. Then

$$U_s(\kappa) = \max_p \left\{ \frac{(\kappa - \lambda p)}{\kappa} \times p + \delta \frac{\lambda p}{\kappa} U_s(\lambda p) \right\},$$

or

$$W_s(\kappa) = \max_p (\kappa - \lambda p) p + \delta W_s(\lambda p), \quad (8.2.1)$$

where $W_s(\kappa) = \kappa U_s(\kappa)$. If W_s is differentiable, then $p(\kappa)$ solves the first order condition,

$$\kappa - 2\lambda p(\kappa) + \delta \lambda W'_s(\lambda p(\kappa)) = 0.$$

The envelope theorem applied to (8.2.1) gives

$$W'_s(\kappa) = p(\kappa) = \gamma \kappa,$$

so that

$$W'_s(\lambda p(\kappa)) = p(\lambda p(\kappa)) = y\lambda p(\kappa) = \lambda y^2 \kappa.$$

Substituting,

$$\kappa - 2\lambda y \kappa + \delta \lambda^2 y^2 \kappa = 0,$$

or

$$1 - 2\lambda y + \delta \lambda^2 y^2 = 0. \quad (8.2.2)$$

Turning to the buyer's optimality condition, a buyer with valuation $v = \lambda p$ must be indifferent between accepting and rejecting, so

$$\lambda p - p = \delta (\lambda p - y\lambda p),$$

or

$$\lambda - 1 = \delta \lambda (1 - y). \quad (8.2.3)$$

Solving (8.2.2) for λy yields

$$y\lambda = \frac{2 \pm \sqrt{4 - 4\delta}}{2\delta} = \frac{1 \pm \sqrt{1 - \delta}}{\delta}.$$

Since we know $y\lambda < 1$, take the negative root,

$$y\lambda = \frac{1 - \sqrt{1 - \delta}}{\delta}.$$

Substituting into (8.2.3),

$$\lambda = \delta \lambda + \sqrt{1 - \delta},$$

or

$$\lambda = \frac{1}{\sqrt{1 - \delta}},$$

so that

$$y = \sqrt{1 - \delta} \times \frac{(1 - \sqrt{1 - \delta})}{\delta} = \frac{\sqrt{1 - \delta} - (1 - \delta)}{\delta}.$$

Note that in this equilibrium, there is skimming: higher valuation buyers buy before lower valuation buyers.

Equilibrium is not unique. It is the only stationary equilibrium.

Let τ denote *real* time, Δ the length of a period, and r the rate of time discount, so that $\delta = e^{-r\Delta}$. If buyer with valuation v buys at or after τ , his utility is no more than $e^{-r\tau}v$. Buying in period 0, he earns $v - \gamma(\delta)$, and so for δ close to 1, buyer v buys before τ . Note that this is not a uniform statement (since for all τ and all δ there exists v such that v purchases after τ).

The Coase conjecture is:

As time between offers shrinks, price charged in first period converges to competitive price, and trade becomes efficient.

8.3 Reputations

Recall the stage game of the the chain store paradox from example 2.2.1, reproduced in Figure 8.3.1.




Figure 8.3.1: The stage game for the chain store.

Two Nash equilibria: (In, Accommodate) and (Out, Fight). Latter violates backward induction.

Chain store, play the game twice, against two different entrants (E_1 and E_2), with the second entrant E_2 observing outcome of first interaction. Incumbent receives total payoffs.

“Chain store paradox” only backward induction (subgame perfect) outcome is that both entrants enter (play In), and incumbent




Figure 8.3.2: A signaling game representation of the subgame reached by E_1 entering.

always accommodates.

But, now suppose incumbent could be *tough*, ω_0 : such an incumbent receives a payoff of 2 from fighting and only 1 from accommodating. Other incumbent is *normal*, ω_n . Both entrants' prior assigns prob $\rho \in (0, 1/2)$ to the incumbent being ω_t . In second market, normal incumbent accommodates and tough fights. Conditional on entry in the first market, result is the signaling game illustrated in Figure 8.3.2.

Note first that there are no pure strategy equilibria.

There is a unique mixed strategy equilibrium: ω_n plays $\alpha \circ F + (1 - \alpha) \circ A$, ω_0 plays F for sure. E_2 enters for sure after A , and plays $\beta \circ E + (1 - \beta) \circ S$ after F .

E_2 is willing to randomize only if his posterior after F that the incumbent is ω_t equals $1/2$. Since that posterior is given by

$$\begin{aligned} \Pr\{\omega_0 \mid F\} &= \frac{\Pr\{F \mid \omega_0\} \Pr\{\omega_0\}}{\Pr\{F\}} \\ &= \frac{\rho}{\rho + (1 - \rho)\alpha}, \end{aligned}$$

solving

$$\frac{\rho}{\rho + (1 - \rho)\alpha} = \frac{1}{2}$$

gives

$$\alpha = \frac{\rho}{1 - \rho},$$

where $\alpha < 1$ since $\rho < 1/2$.

Type ω_n is willing to randomize if

$$\underbrace{4}_{\text{Payoff from } A} = \underbrace{\beta 3 + 5(1 - \beta)}_{\text{Payoff from } F},$$

i.e.,

$$\beta = \frac{1}{2}.$$

Entrant E_1 thus faces a probability of F given by

$$\rho + (1 - \rho)\alpha = 2\rho.$$

Hence, if $\rho < 1/4$, E_1 faces F with sufficiently small probability that he enters. However, if $\rho \in (1/4, 1/2)$, E_1 faces F with sufficiently high probability that he stays out.

8.3.1 Infinite Horizon

Suppose now infinite horizon with incumbent discounting at rate $\delta \in (0, 1)$ and a new potential entrant in each period.

Remark 8.3.1 (Complete Information) Note first that in the complete information game, the outcome in which all entrants enter (play In) and the incumbent accommodates in every period is an equilibrium. Moreover, the profile in which all entrants stay out, any entry is met with F is a subgame perfect equilibrium, supported by the “threat” that play switches to the always-enter/always-accommodate equilibrium if the incumbent ever responds with A . The automaton representation is given in Figure 8.3.3.

Note that the relevant incentive constraint for the incumbent is conditional on I in state w_{OF} (since I does not make a decision when the entrant chooses O), i.e.,

$$(1 - \delta) + \delta 4 \geq 2,$$




Figure 8.3.3: Automaton representation of Nash reversion.

i.e.,

$$\delta \geq \frac{1}{3}.$$

◆

We now consider the *reputation game*, where the incumbent may be normal or tough.

The profile in which all entrants stay out, any entry is met with F is a subgame perfect equilibrium, supported by the “threat” that the entrants believe that the incumbent is normal and play switches to the always-enter/always-accommodate equilibrium if the incumbent ever responds with A .

Theorem 8.3.1 *Suppose the incumbent is either of type ω_n or type ω_0 , and that type ω_0 has prior probability less than $1/2$. Type ω_n must receive a payoff of at least $(1 - \delta) \times 1 + \delta \times 4 = 1 + 3\delta$ in any pure strategy Nash equilibrium in which ω_t always plays F .*

If type ω_0 has prior probability greater than $1/2$, trivially there is never any entry and the normal has payoff 4.

Proof. In the pure strategy Nash equilibrium, either the incumbent always plays F , (in which case, the entrants always stay out and the incumbent’s payoff is 4), or there is a first period (say τ) in which the normal type accommodates, revealing to future entrants that he is the normal type (since the tough type plays F in every period). In such an equilibrium, entrants stay out before τ (since both types of incumbent are choosing F), and there is entry in period τ . After observing F in period τ , entrants conclude the firm is the t type,

and there is no further entry. An easy lower bound on the normal incumbent's equilibrium payoff is then obtained by observing that the normal incumbent's payoff must be at least the payoff from mimicking the t type in period τ . The payoff from such behavior is at least as large as

$$\begin{aligned}
 & \underbrace{(1 - \delta) \sum_{\tau'=0}^{\tau-1} \delta^{\tau'} 4}_{\text{payoff in } \tau' < \tau \text{ from pooling}} + \underbrace{(1 - \delta) \delta^{\tau} \times 1}_{\text{payoff in } \tau \text{ from playing } F \text{ when } A \text{ is myopically optimal}} \\
 & + \underbrace{(1 - \delta) \sum_{\tau'=\tau+1}^{\infty} \delta^{\tau'} 4}_{\text{payoff in } \tau' > \tau \text{ from being treated as the } \omega_0 \text{ type}} \\
 & = (1 - \delta^{\tau})4 + (1 - \delta)\delta^{\tau} + \delta^{\tau+1}4 \\
 & = 4 - 3\delta^{\tau}(1 - \delta) \\
 & \geq 4 - 3(1 - \delta) = 1 + 3\delta.
 \end{aligned}$$

■

For $\delta > 1/3$, the outcome in which all entrants enter and the incumbent accommodates in every period is thus eliminated.

In the reputation literature (see Mailath and Samuelson (2006) for an extensive introduction), it is standard to model the tough type as a *behavioral* type. In that case, the tough type is constrained to necessarily choose F . Then, the result is that in *any* equilibrium, $1 + 3\delta$ is the lower bound on the normal type's payoff.

In fact, irrespective of the presence of other types, if the entrants assign positive probability to the incumbent being a tough behavioral type, for δ close to 1, player I 's payoff in *any* Nash equilibrium is close to 4 (this is an example of a *reputation effect*):

Suppose there is a set of behavioral types Ω . One type is $\omega_0 \in \Omega$ is the Stackelberg, or tough, type, who always plays F . The normal type is ω_n . Other types may include ω_k , who plays F in every

period before k and A afterwards. Suppose the prior beliefs over Ω are given by μ .

Lemma 8.3.1 *Fix a Nash equilibrium. Let h^t be a positive probability period- t history in which every entry results in F . The number of periods in h^t in which an entrant entered is no larger than*

$$k^* := -\frac{\log \mu_0}{\log 2}.$$

Proof. Denote by q_τ the probability that the incumbent plays F in period τ conditional on h^τ if entrant τ plays I . Then, if in equilibrium, entrant τ does play I ,

$$q_\tau \leq \frac{1}{2}.$$

An upper bound on the number of periods in h^t in which an entrant entered is thus

$$k(t) := \#\{\tau : q_\tau \leq \frac{1}{2}\},$$

the number of periods in h^t where $q_\tau \leq \frac{1}{2}$.

Let $\mu_\tau := \Pr\{\omega_0 | h^\tau\}$ be the posterior probability assigned to ω_0 after h^τ , where $\tau < t$ (so that h^τ is an initial segment of h^t). If entrant τ does not enter, $\mu_{\tau+1} = \mu_\tau$. If entrant τ does enter in h^t , then the incumbent fights and²

$$\begin{aligned} \mu_{\tau+1} = \Pr\{\omega_0 | h^\tau, F\} &= \frac{\Pr\{\omega_0, F | h^\tau\}}{\Pr\{F | h^\tau\}} \\ &= \frac{\Pr\{F | \omega_0, h^\tau\} \Pr\{\omega_0 | h^\tau\}}{\Pr\{F | h^\tau\}} \\ &= \frac{\mu_\tau}{q_\tau}. \end{aligned}$$

Defining

$$\tilde{q}_\tau = \begin{cases} q_\tau, & \text{if there is entry in period } \tau, \\ 1, & \text{if there is no entry in period } \tau, \end{cases}$$

²Since the entrant's action is a function of h^τ only, it is uninformative about the incumbent and so can be ignored in the conditioning.

we have, for all $\tau \leq t$,

$$\mu_\tau = \tilde{q}_\tau \mu_{\tau+1},$$

Note that $\tilde{q}_\tau < 1 \implies \tilde{q}_\tau = q_\tau \leq \frac{1}{2}$.

Then,

$$\begin{aligned} \mu_0 &= \tilde{q}_0 \mu_1 = \tilde{q}_0 \tilde{q}_1 \mu_2 \\ &= \mu_t \prod_{\tau=0}^{t-1} \tilde{q}_\tau \\ &= \mu_t \prod_{\{\tau: \tilde{q}_\tau \leq \frac{1}{2}\}} \tilde{q}_\tau \\ &\leq \left(\frac{1}{2}\right)^{k(t)}. \end{aligned}$$

Taking logs, $\log \mu_0 \leq k(t) \log \frac{1}{2}$, and so

$$k(t) \leq -\frac{\log \mu_0}{\log 2}.$$

■

The key intuition here is that since the entrants assign prior positive probability (albeit small) to the Stackelberg type, they cannot be surprised too many times (in the sense of assigning low prior probability to F and then seeing F). Note that the upper bound is independent of t and δ , though it is unbounded in μ_0 .

The normal type can guarantee histories of the form h^t by always playing F when an entrant enters, so we immediately have the lower bound on the normal types payoff of

$$\sum_{\tau=0}^{k^*-1} (1-\delta)\delta^\tau 1 + \sum_{\tau=k^*}^{\infty} (1-\delta)\delta^\tau 4 = 1 - \delta^{k^*} + 4\delta^{k^*} = 1 + 3\delta^{k^*},$$

which can be made arbitrarily close to 4 by choosing δ close to 1.

8.4 Problems

- 8.4.1. (a) Suppose (σ_1, σ_2) is an MPE of the fisheries game from Example 8.1.1 satisfying $\sigma_1(s) + \sigma(s) < s$ for all s . Prove that the profile

remains an MPE of the dynamic game where payoffs are given by

$$u_i(s, a) = \begin{cases} \log a_i, & \text{if } a_1 + a_2 \leq s, \\ \log \left\{ \frac{a_i}{a_1 + a_2} s \right\}, & \text{if } a_1 + a_2 > s. \end{cases}$$

(b) Prove that

$$a_i(s) = \frac{1 - \delta}{2 - \delta} s, \quad i = 1, 2,$$

does indeed describe an MPE of the fisheries game described in Example 8.1.1.

8.4.2. What is the symmetric MPE for the fisheries game of Example 8.1.2 when there are n players, and the transition function is given by

$$q(s, a) = \alpha \max \{0, s - \sum_i a_i\},$$

where $\alpha > 1$?

- 8.4.3. (a) In the MPE calculated in Example 8.1.2, for what values of the discount factor does the stock of fish grow without bound, and for which values does the stock decline to extinction?
- (b) This MPE is inefficient, involving excess extraction. To see this, calculate the largest symmetric payoff profile that can be achieved when the firms choose identical Markov strategies, and prove that the efficient solution extracts less than does the MPE.
- (c) Describe an efficient subgame perfect equilibrium for this game (it is necessarily non-Markov).
- 8.4.4. Consider the asynchronous move prisoners' dilemma from Section 8.1.
- (a) Suppose $\alpha = -\frac{1}{2}$. For some values of δ , there is a Markov perfect equilibrium in which players randomize at E between E and S , and play S for sure at S . Identify the bounds on δ and the probability of randomization for which the described behavior is an MPE.
- (b) Suppose that the initial action of player 1 is not exogenously fixed. The game now has three states, the initial null state and E and S . At the initial state, both players choose an action, and then thereafter player 1 chooses an action in odd periods and player 2 in even periods. Suppose $\alpha > 1$ and δ satisfies (8.1.2).

coalition	1's payoff	2's payoff	3's payoff
$\{1, 2\}$	9	3	0
$\{2, 3\}$	0	9	3
$\{1, 3\}$	3	0	9

Figure 8.4.1: Payoffs to players in each pairwise coalition for Problem 8.4.5. The excluded player receives a payoff of 0.

Prove that there is no pure strategy MPE in which the players choose E .

8.4.5. (A simplification of Livshits (2002).) There are three players. In the initial period, a player i is selected randomly and uniformly to propose a coalition with one other player j , who can accept or reject. If j accepts, the game is over with payoffs given in Figure 8.4.1. If j rejects, play proceeds to the next period, with a new proposer randomly and uniformly selected. The game continues with a new proposer randomly and uniformly selected in each period until a proposal is accepted. Thus, the game is potentially of infinite horizon, and if no coalition is formed (i.e., there is perpetual rejection), all players receive a payoff of 0.

- Suppose $\delta < 3/4$. Describe a stationary pure strategy Markov perfect equilibrium. [Hint: in this equilibrium, every proposal is immediately accepted.]
- Suppose $\delta > 3/4$. Prove there is no Markov perfect equilibrium in stationary pure strategies. There is a stationary Markov perfect equilibrium in behavior strategies. What is it? [Hint: The randomization is on the part of the responder.]
- Suppose $3/4 < \delta < \sqrt{3/4}$. There are two nonstationary pure strategy Markov equilibria. What are they? [Hint: if $\delta < \sqrt{3/4}$, then $\delta^2 < 3/4$.]
- Suppose $3/4 < \delta < \sqrt{3/4}$. Construct a non-Markov perfect

equilibrium in which in the first period, if 1 is selected, then 1 chooses 3.

- 8.4.6. Consider the model of Section 8.2, but assume the buyer's valuation v can only take on two values, 2 and 3. Moreover, the seller's beliefs assign probability α to the value 2. The seller's cost (value) is zero, and the buyer and seller have a common discount factor $\delta \in (0, 1)$.
- What is the unique perfect Bayesian equilibrium of the one period model (in this model, the seller makes a take-it-or-leave-it offer to the buyer)?
 - Suppose $\alpha = \frac{1}{2}$. The two period model has a unique perfect Bayesian equilibrium. What is it? [You may assume that any rejection in the first period results in a posterior that assigns at least probability $\frac{1}{2}$ to $v = 2$.]
- 8.4.7. As in the model of Section 8.2, there is an uninformed seller with cost zero facing a buyer with value uniformly distributed on $[0, 1]$. Suppose the seller has a rate of continuous time discounting of r_S (so the seller's discount factor is $\delta_S = e^{-r_S \Delta}$, where $\Delta > 0$ is the time between offers), while the buyer has a rate of continuous time discounting of r_B (so the buyer's discount factor is $\delta_B = e^{-r_B \Delta}$). Solve for an equilibrium of the infinite horizon game in which the uninformed seller makes all the offers. What happens to the initial offer as $\Delta \rightarrow 0$?
- 8.4.8. Reconsider the two period reputation example (illustrated in Figure 8.3.2) with $\rho > \frac{1}{2}$. Describe all of the equilibria. Which equilibria survive the intuitive criterion?
- 8.4.9. Describe the equilibria of the three period version of the reputation example.
- 8.4.10. Consider the following stage game where player 1 is the row player and 2, the column player (as usual). Player 1 is one of two types ω_n and ω_0 . Payoffs are:

	L	R
T	2, 3	0, 2
B	3, 0	1, 1
	ω_n	

	L	R
T	3, 3	1, 2
B	2, 0	0, 1
	ω_0	

The stage game is played twice, and player 2 is short-lived: a different player 2 plays in different periods, with the second period player 2 observing the action profile chosen in the first period. Describe all the equilibria of the game. Does the intuitive criterion eliminate any of them?

- 8.4.11. This is a continuation of Problem 7.6.11. Suppose now that the game with the long-lived player 1 and short-lived player 2's is a game of incomplete information. With prior probability $\rho \in (0, 1)$, player 1 is a *behavioral type* who chooses T in every period, and with probability $1 - \rho$, he is a strategic or normal type as described above. Suppose $\rho > \frac{1}{2}$. Describe an equilibrium in which the normal type of player 1 has a payoff strictly greater than 2 for large δ .

Chapter 9

Bargaining¹

9.1 Axiomatic Nash Bargaining

A bargaining problem is a pair $\langle S, d \rangle$, $S \subset \mathbb{R}^2$ compact and convex, $d \in S$ and $\exists s \in S$ such that $s_i > d_i$ for $i = 1, 2$. Let \mathcal{B} denote the collection of bargaining problems. While d is often interpreted as a disagreement point, this is not the role it plays in the axiomatic treatment. It only plays a role in INV (where its role has the flavor of a normalization constraint) and in SYM. The appropriate interpretation is closely linked to noncooperative bargaining. It is *not* the value of an outside option!

Definition 9.1.1 A bargaining solution is a function $f : \mathcal{B} \rightarrow \mathbb{R}^2$ such that $f(S, d) \in S$.

9.1.1 The Axioms

1. INV (Invariance to Equivalent Utility Representations)

Given $\langle S, d \rangle$, let $\langle S', d' \rangle$ be the bargaining problem given by, for some $(\alpha_i, \beta_i)_{i=1}^2$, $\alpha_i > 0$,

$$S' = \{(\alpha_1 s_1 + \beta_1, \alpha_2 s_2 + \beta_2) : (s_1, s_2) \in S\}$$

and

$$d'_i = \alpha_i d_i + \beta_i, \quad i = 1, 2.$$

¹Copyright December 7, 2011 by George J. Mailath

Then

$$f_i(S', d') = \alpha_i f_i(S, d) + \beta_i, \quad i = 1, 2.$$

2. SYM (Symmetry)

If $d_1 = d_2$ and $(s_1, s_2) \in S \Rightarrow (s_2, s_1) \in S$, then

$$f_1(S, d) = f_2(S, d).$$

3. IIA (Independence of Irrelevant Alternatives)

If $S \subset T$ and $f(T, d) \in S$, then

$$f(S, d) = f(T, d).$$

4. PAR (Pareto Efficiency)

If $s \in S$, $t \in S$, $t_i > s_i$, $i = 1, 2$, then

$$f(S, d) \neq s.$$

9.1.2 Nash's Theorem

Theorem 9.1.1 (Nash) *If $f : \mathcal{B} \rightarrow \mathbb{R}^2$ satisfies INV, SYM, IIA, and PAR, then*

$$f(S, d) = \arg \max_{(d_1, d_2) \leq (s_1, s_2) \in S} (s_1 - d_1)(s_2 - d_2) \equiv f^N(S, d).$$

If $s_1 + s_2 \leq 1$, then player 1's Nash share is

$$s_1^* = \frac{1 + d_1 - d_2}{2}$$

(which is (9.4.1)).

Proof. Leave as an exercise that f^N satisfies the four axioms.

Suppose that f satisfies the four axioms. Fix $\langle S, d \rangle$.

Step 1: Let $z = f^N(S, d)$. Then $z_i > d_i$, $i = 1, 2$. Apply the following affine transformations to move d to the origin and z to $(1/2, 1/2)$:

$$\alpha_i = \frac{1}{2(z_i - d_i)}; \quad \beta_i = \frac{-d_i}{2(z_i - d_i)}.$$




Figure 9.1.1: Illustration of step 2.

Denote the transformed problem $\langle S', 0 \rangle$.

INV implies

$$f_i(S', 0) = \alpha_i f_i(S, d) + \beta_i$$

and

$$f_i^N(S', 0) = \alpha_i f_i^N(S, d) + \beta_i = \frac{1}{2}.$$


Note that $f_i(S, d) = f_i^N(S, d)$ if and only if $f_i(S', 0) = 1/2$.

Step 2: Claim - $\nexists (s'_1, s'_2) \in S'$ such that $s'_1 + s'_2 > 1$.

Suppose not. Then convexity of S' implies $t = (1 - \varepsilon)(1/2, 1/2) + \varepsilon s' \in S'$. Moreover, for ε small, $t_1 t_2 > 1/4$, contradicting $f^N(S', 0) = (1/2, 1/2)$ (see Figure 9.1.1).

Step 3: Let $T = \{(s_1, s_2) \in \mathbb{R}^2 : s_1 + s_2 \leq 1, |s_i| \leq \max\{|s'_1|, |s'_2| : s' \in S'\}\}$ (see Figure 9.1.2). Then, SYM and PAR, $f(T, 0) = (1/2, 1/2)$.

Step 4: Since $S' \subset T$, IIA implies $f(S', 0) = (1/2, 1/2)$. ■

Figure 9.1.2: The bargaining set T .

9.2 Rubinstein (1982) Bargaining

Two agents bargain over $[0, 1]$. Time is indexed by t , $t = 1, 2, \dots$. A proposal is a division of the pie $(x, 1 - x)$, $x \geq 0$. The agents take turns to make proposals. Player I makes proposals on odd t and II on even t . If the proposal $(x, 1 - x)$ is agreed to at time t , I 's payoff is $\delta_1^{t-1}x$ and II 's payoff is $\delta_2^{t-1}(1-x)$. Perpetual disagreement yields a payoff of $(0, 0)$. Impatience implies $\delta_i < 1$.

Histories are $h^t \in [0, 1]^{\tau-1}$.

Strategies for player I , $\tau_I^1 : \cup_{t \text{ odd}} [0, 1]^{t-1} \rightarrow [0, 1]$, $\tau_I^2 : \cup_{t \text{ even}} [0, 1]^t \rightarrow \{A, R\}$, and for player II , $\tau_{II}^1 : \cup_{t \geq 2 \text{ even}} [0, 1]^{t-1} \rightarrow [0, 1]$ and $\tau_{II}^2 : \cup_{t \text{ odd}} [0, 1]^t \rightarrow \{A, R\}$.

Need to distinguish between histories in which all proposals have been rejected, and those in which all but the last have been rejected and the current one is being considered.

9.2.1 The Stationary Equilibrium

All the subgames after different even length histories of rejected proposals are strategically identical. A similar comment applies to different odd length histories of rejected proposals. Finally, all the subgames that follow different even length histories of rejected proposals followed by the *same* proposal on the table are strategically identical. Similarly, all the subgames that follow different odd length histories of rejected proposals followed by the *same* proposal on the table are strategically identical.

Consider first equilibria in history independent (or stationary) strategies. Recall that a strategy for player I is a pair of mappings, (τ_I^1, τ_I^2) . The strategy τ_I^1 is stationary if, for all $h^t \in [0, 1]^{t-1}$ and $\hat{h}^t \in [0, 1]^{\hat{t}-1}$, $\tau_I^1(h^t) = \tau_I^1(\hat{h}^t)$ (and similarly for the other strategies). Thus, if a strategy profile is a stationary equilibrium (with agreement), there is a pair (x^*, z^*) , such that I expects x^* in any subgame in which I moves first and expects z^* in any subgame in which II moves first. In order for this to be an equilibrium, I 's claim should make II indifferent between accepting and rejecting: $1 - x^* = \delta_2(1 - z^*)$, and similarly I is indifferent, so $z^* = \delta_1 x^*$. [Consider the first indifference. Player I won't make a claim that

II strictly prefers to $1 - z^*$ next period, so $1 - x^* \leq \delta_2(1 - z^*)$. If II strictly prefers $(1 - z^*)$ next period, she rejects and gets $1 - z^*$ next period, leaving I with z^* . But I can offer II a share $1 - z^*$ this period, avoiding the one period delay.] Solving yields

$$x^* = (1 - \delta_2)/(1 - \delta_1\delta_2),$$

and

$$z^* = \delta_1(1 - \delta_2)/(1 - \delta_1\delta_2).$$

The stationary subgame perfect equilibrium (note that backward induction is not well defined for the infinite horizon game) is for I to always claim x^* and accept any offer $\geq z^*$, and for II to always offer z^* and always accept any claim $\leq x^*$.

9.2.2 All Equilibria

While in principal, there could be nonstationary equilibria, it turns out that there is only one subgame perfect eq.

Denote by i/j the game in which i makes the initial proposal to j . Define

$$M_i = \sup \{i\text{'s discounted expected payoff} \\ \text{in any subgame perfect eq of } i/j\}$$

and

$$m_i = \inf \{i\text{'s discounted expected payoff} \\ \text{in any subgame perfect eq of } i/j\}.$$

Claim 9.2.1 $m_j \geq 1 - \delta_i M_i$.

Proof. Note first that i must, in equilibrium, accept any offer $> \delta_i M_i$. Suppose $m_j < 1 - \delta_i M_i$. Then there would exist an eq yielding a payoff $u_j < 1 - \delta_i M_i$ to j . But this is impossible, since j has a profitable deviation in such an eq: offer $\delta_i M_i + \varepsilon$, ε small. Player i must accept, giving j a payoff of $1 - \delta_i M_i - \varepsilon > u_j$, for ε sufficiently small. ■

Claim 9.2.2 $M_j \leq 1 - \delta_i m_i$.

Proof. i only accepts an offer if it is at least $\delta_i m_i$.

If i does reject, then i offers no more than $\delta_j M_j$. So,

$$\begin{aligned} M_j &\leq \max \left\{ \underbrace{1 - \delta_i m_i}_{\text{if } i \text{ accepts}}, \underbrace{\delta_j^2 M_j}_{\text{if } i \text{ rejects}} \right\} \\ \Rightarrow M_j &\leq 1 - \delta_i m_i. \end{aligned}$$

■

The first claim implies

$$\begin{aligned} M_j &\leq 1 - \delta_i (1 - \delta_j M_j) \\ \Rightarrow M_j &\leq \frac{(1 - \delta_i)}{(1 - \delta_i \delta_j)}, \quad M_i \leq \frac{(1 - \delta_j)}{(1 - \delta_i \delta_j)}. \end{aligned}$$

This implies

$$m_i \geq 1 - \delta_j \frac{(1 - \delta_i)}{(1 - \delta_i \delta_j)} = \frac{(1 - \delta_j)}{(1 - \delta_i \delta_j)}$$

and so

$$m_i = M_i = \frac{(1 - \delta_j)}{(1 - \delta_i \delta_j)}.$$

9.2.3 Impatience

In order to investigate the impact of reducing the bargaining friction intrinsic in impatience, we do the following:

Time is continuous, with each round of bargaining taking Δ units of time. If player i has discount rate r_i ,

$$\delta_i = e^{-r_i \Delta}.$$

Player 1's share is then

$$x^*(\Delta) = \frac{1 - \delta_2}{1 - \delta_1 \delta_2} = \frac{1 - e^{-r_2 \Delta}}{1 - e^{-(r_1 + r_2) \Delta}}$$

and so

$$\begin{aligned}
 \lim_{\Delta \rightarrow 0} x^*(\Delta) &= \lim_{\Delta \rightarrow 0} \frac{1 - e^{-r_2 \Delta}}{1 - e^{-(r_1 + r_2) \Delta}} \\
 &= \lim_{\Delta \rightarrow 0} \frac{r_2 e^{-r_2 \Delta}}{(r_1 + r_2) e^{-(r_1 + r_2) \Delta}} \\
 &= \frac{r_2}{r_1 + r_2},
 \end{aligned}$$

where l'Hopital's rule was used to get to the second line.

Note that the first mover advantage has disappeared (as it should). The bargaining is determined by relative impatience.

9.3 Outside Options

Player *II* has an outside option of value $(0, b)$. Suppose player *II* can only select outside option when rejecting *I*'s proposal, and receives b in that period. See Figure 9.3.1 for the extensive form.

Claim 9.3.1 $m_2 \geq 1 - \delta_1 M_1$.

Proof. Same argument as Claim 9.2.1. ■

Claim 9.3.2 $M_1 \leq 1 - b, M_1 \leq 1 - \delta_2 m_2$.

Proof. $M_1 \leq 1 - b$ (since *II* can always opt out). $M_1 \leq 1 - \delta_2 m_2$ follows as in case without outside option (Claim 9.2.2). ■

Claim 9.3.3 $m_1 \geq 1 - \max \{b, \delta_2 M_2\}, M_2 \leq 1 - \delta_1 m_1$.

Proof. If $b \leq \delta_2 M_2$, then the argument from Claim 9.2.1 shows that $m_1 \geq \delta_2 M_2$. If $b > \delta_2 M_2$, then *II* takes the outside option rather than rejecting and making a counterproposal. Thus, *II*'s acceptance rule is accept any proposal of a share $> b$, and take the outside option for any proposal $< b$. Thus, *I*'s payoffs is $1 - b$.

$M_2 \leq 1 - \delta_1 m_1$ follows as in case without outside option (Claim 9.2.2). ■




Figure 9.3.1: The first two periods when II can opt out only after rejecting I 's proposal.

Claim 9.3.4 $b \leq \delta_2(1-\delta_1)/(1-\delta_1\delta_2) \Rightarrow m_i \leq (1-\delta_j)/(1-\delta_1\delta_2) \leq M_i$.

Proof. Follows from the Rubinstein shares being equilibrium shares.

■

Claim 9.3.5 $b \leq \delta_2(1-\delta_1)/(1-\delta_1\delta_2) \Rightarrow m_i = (1-\delta_j)/(1-\delta_1\delta_2) = M_i$

Proof. From Claim 9.3.2, $1 - M_1 \geq \delta_2 m_2$, and so from claim 9.3.1, $1 - M_1 \geq \delta_2(1 - \delta_1 M_1)$, and so $M_1 \leq (1 - \delta_2)/(1 - \delta_1\delta_2)$ and so we have equality.

From Claim 9.3.1, $m_2 \geq 1 - \delta_1(1 - \delta_2)/(1 - \delta_1\delta_2) = (1 - \delta_1)/(1 - \delta_1\delta_2)$, and so equality again.

From Claim 9.3.4, $\delta_2 M_2 \geq \delta_2(1 - \delta_1)/(1 - \delta_1\delta_2) \geq b$, and so by Claim 9.3.3, $m_1 \geq 1 - \delta_2 M_2 \geq 1 - \delta_2(1 - \delta_1 m_1)$. Thus, $m_1 \geq (1 - \delta_2)/(1 - \delta_1\delta_2)$, and so equality.

Finally, from Claim 9.3.3,

$$M_2 \leq 1 - \delta_1 m_1 = 1 - \delta_1(1 - \delta_2)/(1 - \delta_1\delta_2) = (1 - \delta_1)/(1 - \delta_1\delta_2).$$

■

Thus, if $b \leq \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$, equilibrium payoffs are uniquely determined. If $b < \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$, then the subgame perfect equilibrium profile is also uniquely determined (player *II* never takes the outside option). If $b = \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$, then there are multiple subgame perfect equilibrium profiles, which differ in whether player *II* takes the outside option or not after an unacceptable offer.

Claim 9.3.6 $b > \delta_2(1-\delta_1)/(1-\delta_1\delta_2) \Rightarrow m_1 \leq 1 - b \leq M_1, m_2 \leq 1 - \delta_1(1 - b) \leq M_2$.

Proof. Follows from the following being an equilibrium: *I* always proposes $1 - b$, and accepts any offer of at least $\delta_1(1 - b)$; *II* always proposes $\delta_1(1 - b)$ and accepts any claim of no more than $1 - b$, opting out if the claim is more than $1 - b$.

■

Claim 9.3.7 $b > \delta_2(1 - \delta_1) / (1 - \delta_1\delta_2) \Rightarrow m_1 = 1 - b = M_1, m_2 = 1 - \delta_1(1 - b) = M_2.$

Proof. From Claim 9.3.2, $1 - M_1 \geq b$, i.e., $M_1 \leq 1 - b$, and so we have equality.

From Claim 9.3.1, $m_2 \geq 1 - \delta_1(1 - b)$, and so we have equality.

From Claim 9.3.6, $1 - b \geq m_1$ and so $(1 - \delta_2) / (1 - \delta_1\delta_2) > m_1$.

We now argue that $\delta_2 M_2 \leq b$. If $\delta_2 M_2 > b$, then $m_1 \geq 1 - \delta_2 M_2 \geq 1 - \delta_2(1 - \delta_1 m_1)$ and so $m_1 \geq (1 - \delta_2) / (1 - \delta_1\delta_2)$, a contradiction. Thus, $\delta_2 M_2 \leq b$.

From Claim 9.3.6 and 9.3.3, $1 - b \geq m_1 \geq 1 - \max\{b, \delta_2 M_2\} = 1 - b$ and so $m_1 = 1 - b$.

Finally, this implies $M_2 \leq 1 - \delta_1 m_1 = 1 - \delta_1(1 - b)$, and so equality. ■

Thus, if $b > \delta_2(1 - \delta_1) / (1 - \delta_1\delta_2)$, equilibrium payoffs are uniquely determined. Moreover, the subgame perfect equilibrium profile is also uniquely determined (player II always takes the outside option after rejection):

$$b > \delta_2(1 - \delta_1) / (1 - \delta_1\delta_2) \Leftrightarrow b > \delta_2[1 - \delta_1(1 - b)].$$

Remark 9.3.1 If II can only select outside option after I rejects (receiving b in that period), then there are multiple equilibria. The equilibrium construction is a little delicate in this case. In fact, there is no pure strategy Markov perfect equilibrium. There is, however, a behavior strategy Markov perfect equilibrium. ♦

9.4 Exogenous Risk of Breakdown

Suppose that after any rejection, there is a probability $1 - \theta$ of breakdown and the outcome (d_1, d_2) is implemented. With probability θ , bargaining continues to the next round. No discounting. Note that since always rejecting is a feasible strategy, $d_i \leq m_i \leq M_i$.

Claim 9.4.1 $m_j \geq 1 - \theta M_i - (1 - \theta) d_i.$

Proof. Note first that i must, in equilibrium, accept any offer $> \theta M_i + (1 - \theta) d_i$. Suppose $m_j < 1 - \theta M_i - (1 - \theta) d_i$. Then there would exist an eq yielding a payoff $u_j < 1 - \theta M_i - (1 - \theta) d_i$ to j . But j can deviate in such an eq, and offer $\theta M_i + (1 - \theta) d_i + \varepsilon$, ε small, which i accepts. This gives j a payoff of $1 - \theta M_i - (1 - \theta) d_i - \varepsilon > u_j$, for ε sufficiently small. ■

Claim 9.4.2 $M_j \leq 1 - \theta m_i - (1 - \theta) d_i$.

Proof. In eq, i rejects any offer $< \theta m_i + (1 - \theta) d_i$ and then i offers no more than $\theta M_j + (1 - \theta) d_j$. So,

$$\begin{aligned} M_j &\leq \max \{1 - \theta m_i - (1 - \theta) d_i, \theta [\theta M_j + (1 - \theta) d_j] + (1 - \theta) d_j\} \\ \Rightarrow M_j &\leq 1 - \theta m_i - (1 - \theta) d_i, \end{aligned}$$

$$\text{since } M_j < \theta^2 M_j + (1 - \theta^2) d_j \Leftrightarrow M_j < d_j \quad \blacksquare$$

The first claim implies

$$\begin{aligned} M_j &\leq 1 - \theta (1 - \theta M_j - (1 - \theta) d_j) - (1 - \theta) d_i \\ \Rightarrow M_j &\leq \frac{(1 + \theta d_j - d_i)}{(1 + \theta)}, \quad M_i \leq \frac{(1 + \theta d_i - d_j)}{(1 + \theta)}. \end{aligned}$$

This implies

$$m_i \geq 1 - \theta \frac{(1 + \theta d_j - d_i)}{(1 + \theta)} - (1 - \theta) d_j = \frac{1 + \theta d_i - d_j}{(1 + \theta)} = M_i$$

and so

$$m_i = M_i = \frac{1 + \theta d_i - d_j}{(1 + \theta)}.$$

Now, we are interested in the payoffs as $\theta \rightarrow 1$, and

$$m_i \rightarrow \frac{1 + d_i - d_j}{2},$$

so that I 's share is

$$x^* = \frac{1 + d_1 - d_2}{2}. \quad (9.4.1)$$

For much more on bargaining, see Osborne and Rubinstein (1990).

9.5 Problems

- 9.5.1. Two agents bargain over $[0, 1]$. Time is indexed by t , $t = 1, 2, \dots, T$, T finite. A proposal is a division of the pie $(x, 1 - x)$, $x \geq 0$. The agents take turns to make proposals. Player I makes proposals on odd t and II on even t . If the proposal $(x, 1 - x)$ is agreed to at time t , I 's payoff is $\delta_1^{t-1}x$ and II 's payoff is $\delta_2^{t-1}(1 - x)$. Perpetual disagreement yields a payoff of $(0, 0)$. Impatience implies $\delta_i < 1$. The game ends in period T if all previous proposals have been rejected, with each receiving a payoff of zero.

- (a) Suppose T odd, so that I is the last player to make a proposal. If $T = 1$, the player I makes a take-it-or-leave-it offer, and so in equilibrium demands the entire pie and II accepts. Prove that in the unique backward induction equilibrium, if there are k periods remaining, where k is odd and $k \geq 3$, I 's proposal is given by

$$x_k = (1 - \delta_2) \sum_{r=0}^{\tau-1} (\delta_1 \delta_2)^r + (\delta_1 \delta_2)^\tau, \quad \tau = (k - 1)/2.$$

[Hint: First calculate x_1 (the offer in the last period), x_2 , and x_3 . Then write out the recursion, and finally verify that the provided expression satisfies the appropriate conditions.]

- (b) What is the limit of x_T as $T \rightarrow \infty$?
- (c) Suppose now that T is even, so that II is the last player to make a proposal. Prove that in the unique backward induction equilibrium, if there are k periods remaining, where k is even and $k \geq 2$, I 's proposal is given by

$$y_k = (1 - \delta_2) \sum_{r=0}^{\tau-1} (\delta_1 \delta_2)^r, \quad \tau = k/2.$$

- (d) What is the limit of y_T as $T \rightarrow \infty$?

- 9.5.2. (a) Give the details of the proof of Claim 9.3.4. (A few sentences explaining why it works is sufficient.)
- (b) Give the details of the proof of Claim 9.3.6. (A few sentences explaining why it works is sufficient.)

9.5.3. We will use the finite horizon bargaining result from question 9.5.1 to give an alternative proof of uniqueness in the Rubinstein model.

(a) Prove that in any subgame perfect equilibrium of the game in which I offers first, I 's payoff is no more than x_k , for all k odd.

[**Hint:** Prove by induction (the result is clearly true for $k = 1$).]

(b) Prove that in any subgame perfect equilibrium of the game in which I offers first, I 's payoff is no less than y_k , for all k even.

(c) Complete the argument.

9.5.4. There is a single seller who has a single object to sell (the seller's reservation utility is zero). There are two potential buyers, and they each value the object at 1. If the seller and buyer i agree to a trade at price p in period t , then the seller receives a payoff of $\delta^{t-1}p$, buyer i a payoff of $\delta^{t-1}(1-p)$, and buyer $j \neq i$ a payoff of zero. Consider alternating offer bargaining, with the seller choosing a buyer to make an offer to (name a price to). If the buyer accepts, the game is over. If the buyer rejects, then play proceeds to the next period, when the buyer who received the offer in the preceding period makes a counter-offer. If the offer is accepted, it is implemented. If the offer is rejected, then the seller makes a new proposal to either the same buyer or to the other buyer. Thus, the seller is free to switch the identity of the buyer he is negotiating with after every rejected offer from the buyer.

(a) Suppose that the seller can only make proposals in odd-numbered periods. Prove that the seller's subgame perfect equilibrium payoff is unique, and describe it. Describe the subgame perfect equilibria. The payoffs to the buyers are not uniquely determined. Why not?

(b) Now consider the following alternative. Suppose that if the seller rejects an offer from the buyer, he can either wait one period to make a counteroffer to this buyer, or he can *immediately* make an offer to the other buyer. Prove that the seller's subgame perfect equilibrium payoff is unique, and describe it. Describe the subgame perfect equilibria. [Cf. Shaked and Sutton (1984).]

Chapter 10

Appendices¹

10.1 Section 2.5.1: Trembling Hand Perfection

Definition 10.1.1 *An equilibrium $b \equiv (b_1, \dots, b_n)$ of a finite extensive form game Γ is extensive form trembling hand perfect if there exists a sequence $\{b^k\}_k$ of completely mixed behavior strategy profiles converging to b such that for all players i and information sets $h \in H_i$, conditional on reaching h , for all k , $b_i^k(h)$ maximizes player i 's expected payoff, given b_{-i}^k and $b_i^k(h')$ for $h' \neq h$.*

Theorem 10.1.1 *If b is an extensive form trembling hand perfect equilibrium of the finite extensive form game Γ , then it is subgame perfect.*

Remark 10.1.1 When each player only has one information set (such as in Selten's horse), trembling hand perfect in the normal and extensive form coincide. In general, they differ. Moreover, trembling hand perfect in the normal form need not imply subgame perfection (see problem 10.2.1).



¹Copyright December 7, 2011 by George J. Mailath

Remark 10.1.2 Note that trembling hand perfect in the extensive form requires players to be sensitive not only to the possibility of trembles by the other players, but also their own trembles (at other information sets). The notion of *quasi-perfect* drops the latter requirement:

Definition 10.1.2 (van Damme (1984)) *A behavior strategy profile b is quasi-perfect if it is the limit of a sequence of completely mixed behavior profiles b^n , and if for each player i and information set h owned by that player, conditional on reaching h , b_i is a best response to b_{-i}^n for all n .*

◆

Theorem 10.1.2 *Suppose b is a extensive form trembling hand perfect equilibrium of a finite extensive form game Γ . Then, b is sequentially rational given some consistent system of beliefs μ , and so is sequential.*

Proof. Suppose b is trembling hand perfect in the extensive form. Let $\{b^k\}_k$ be the trembles, i.e., the sequence of completely mixed behavior strategy profiles converging to b . Let μ^k be the system of beliefs implied by Bayes' rule by b^k (since b^k is completely mixed, μ^k is well-defined). Since the collection of systems of beliefs, $\prod_{h \in \cup_i H_i} \Delta(h)$, is compact, the sequence $\{\mu^k\}$ has a convergent subsequence with limit μ , and so μ is consistent. Moreover, a few minutes of reflection reveals that for each i , (b_i, b_{-i}^k) is sequentially rational at every information set owned by i , i.e., for all $h \in H_i$, given μ^k . Since best replies are hemicontinuous, for each i , b is sequentially rational at every information set owned by i , i.e., for all $h \in H_i$, given μ . That is, b is sequentially rational given μ .

■

10.1.1 Existence and Characterization

This subsection outlines Selten's (1975) original definition and proves the equivalence between the different definitions.

Let $\eta : \cup S_i \rightarrow (0, 1)$ be a function satisfying $\sum_{s_i \in S_i} \eta(s_i) < 1$ for all i . The *associated perturbed game*, denoted (G, η) , is the normal form game $\{(R_1^\eta, v_1), \dots, (R_n^\eta, v_n)\}$ where

$$R_i^\eta = \{\sigma_i \in \Delta(S_i) : \sigma_i(s_i) \geq \eta(s_i), \quad \forall s_i \in S_i\}$$

and v_i is expected payoffs. Note that σ is a Nash equilibrium of (G, η) if and only if for all $i, s_i, s'_i \in S_i$,

$$v_i(s_i, \sigma_{-i}) < v_i(s'_i, \sigma_{-i}) \Rightarrow \sigma_i(s_i) = \eta(s_i).$$

Definition 10.1.3 (Selten (1975)) *An equilibrium σ of a normal form game G is (normal form) trembling hand perfect if there exists a sequence $\{\eta_k\}_k$ such that $\eta_k(s_i) \rightarrow 0 \quad \forall s_i$ as $k \rightarrow \infty$ and an associated sequence of mixed strategy profiles $\{\sigma^k\}_k$ with σ^k a Nash equilibrium of (G, η_k) such that $\sigma^k \rightarrow \sigma$ as $k \rightarrow \infty$.*

Theorem 10.1.3 *Every finite normal form game has a trembling hand perfect equilibrium.*

Proof. Clearly (G, η) has a Nash equilibrium for all η . Suppose $\{\eta_m\}$ is a sequence such that $\eta_m(s_i) \rightarrow 0 \quad \forall s_i$ as $m \rightarrow \infty$. Let σ^m be an equilibrium of (G, η_m) . Since $\{\sigma^m\}$ is a sequence in the compact set $\prod_i \Delta(S_i)$, it has a convergent subsequence. Its limit is a trembling hand perfect equilibrium of G . ■

Remark 10.1.3 If σ is not trembling hand perfect, then there exists $\varepsilon > 0$ such that for all sequences $\{\eta_k\}_k$ satisfying $\eta_k \rightarrow 0$, eventually all Nash equilibria of the associated perturbed games are bounded away from σ by at least ε (i.e., $\exists K$ so that $\forall k \geq K$ and $\forall \sigma^k$ equilibrium of (G, η_k) , $|\sigma^k - \sigma| \geq \varepsilon$). ◆

Definition 10.1.4 (Myerson (1978)) *The mixed strategy profile σ is an ε -perfect equilibrium of G if it is completely mixed ($\sigma_i(s_i) > 0 \quad \forall s_i \in S_i$) and satisfies*

$$s_i \notin BR_i(\sigma_{-i}) \Rightarrow \sigma_i(s_i) \leq \varepsilon.$$

Theorem 10.1.4 *Suppose σ is a strategy profile of the normal form game G . The following are equivalent:*

1. σ is a trembling hand perfect equilibrium of G ;
2. there exists a sequence $\{\varepsilon_k : \varepsilon_k \rightarrow 0\}$ and an associated sequence of ε_k -perfect equilibria converging to σ ; and
3. there exists a sequence $\{\sigma^k\}$ of completely mixed strategy profiles converging to σ such that σ_i is a best reply to σ_{-i}^k , for all k .

Proof. (1) \Rightarrow (2). Take $\varepsilon_k = \max_{s_i \in S_i, i} \eta_k(s_i)$.

(2) \Rightarrow (3). Let $\{\sigma^k\}$ be the sequence of ε_k -perfect equilibria. Suppose s_i receives positive probability under σ_i . Need to show that s_i is a best reply to σ_{-i}^k . Since $\varepsilon_k \rightarrow 0$ and $\sigma_i^k(s_i) \rightarrow \sigma_i(s_i) > 0$, there exists $k^*(s_i)$ such that $k > k^*(s_i)$ implies $\sigma_i^k(s_i) > \sigma_i(s_i)/2 > \varepsilon_k$. But σ^k is ε_k -perfect, so $s_i \in BR_i(\sigma_{-i}^k)$. The desired sequence is $\{\sigma^k : k > \bar{k}\}$, where $\bar{k} \equiv \max \{k^*(s_i) : \sigma_i(s_i) > 0, i\}$.

(3) \Rightarrow (1). Define η_k as

$$\eta_k(s_i) = \begin{cases} \sigma_i^k(s_i), & \text{if } \sigma_i(s_i) = 0, \\ 1/k, & \text{if } \sigma_i(s_i) > 0. \end{cases}$$

Since $\sigma^k \rightarrow \sigma$, there exists k' such that $k > k'$ implies $\sum_{s_i \in S_i} \eta_k(s_i) < 1$ for all i .

Let $m = \min \{\sigma_i(s_i) : \sigma_i(s_i) > 0, i\}$. There exists k'' such that, for all $k > k''$, $|\sigma_i^k(s_i) - \sigma_i(s_i)| < m/2$. Suppose $k > \max \{k', k'', 2/m\}$. Then $\sigma_i^k(s_i) \geq \eta_k(s_i)$. [If $\sigma_i(s_i) = 0$, then immediate. If $\sigma_i(s_i) > 0$, then $\sigma_i^k(s_i) > \sigma_i(s_i) - m/2 > m/2 > 1/k$.]

Since σ_i is a best reply to σ_{-i}^k , if s_i is not a best reply to σ_{-i}^k , then $\sigma_i(s_i) = 0$. But this implies that σ^k is an equilibrium of (G, η_k) (since s_i is played with minimum probability). ■

10.2 Problems

- 10.2.1. By explicitly presenting the completely mixed trembles, show that the profile $L\ell$ is *normal form* trembling hand perfect in the following game (this is the normal form from Example 2.3.4):

	ℓ	r
L	2, 0	2, 0
T	-1, 1	4, 0
B	0, 0	5, 1

Show that there is no *extensive form* trembling hand perfect equilibrium with that outcome in the first extensive form presented in Example 2.3.4.

Bibliography

- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55(1), 1-18.
- BEN-EL-MECHAIEKH, H., AND R. W. DIMAND (2011): "A Simpler Proof of the Von Neumann Minimax Theorem," *American Mathematical Monthly*, 118(7), 636-641.
- BERNHEIM, B. D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007-1028.
- BÖRGERS, T. (1993): "Pure Strategy Dominance," *Econometrica*, 61(2), 423-430.
- CARLSSON, H., AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61(5), 989-1018.
- CHO, I.-K., AND D. KREPS (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179-221.
- ELMES, S., AND P. J. RENY (1994): "On the Strategic Equivalence of Extensive Form Games," *Journal of Economic Theory*, 62(1), 1-23.
- FUDENBERG, D., AND D. K. LEVINE (1998): *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- GOVINDAN, S., P. J. RENY, AND A. J. ROBSON (2003): "A Short Proof of Harsanyi's Purification Theorem," *Games and Economic Behavior*, 45(2), 369-374.
- GUL, F. (1998): "A Comment on Aumann's Bayesian View," *Econometrica*, 66(4), 923-927.

- HARSANYI, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2(1), 1-23.
- JACKSON, M. O., L. K. SIMON, J. M. SWINKELS, AND W. R. ZAME (2002): "Communication and Equilibrium in Discontinuous Games of Incomplete Information," *Econometrica*, 70(5), 1711-1740.
- KOHLBERG, E., AND J.-F. MERTENS (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54(5), 1003-1037.
- KREPS, D., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50(4), 863-894.
- LIVSHITS, I. (2002): "On Non-Existence of Pure Strategy Markov Perfect Equilibrium," *Economics Letters*, 76(3), 393-396.
- MAILATH, G. J. (1998): "Do People Play Nash Equilibrium? Lessons From Evolutionary Game Theory," *Journal of Economic Literature*, 36, 1347-1374.
- MAILATH, G. J., AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York, NY.
- MAILATH, G. J., L. SAMUELSON, AND J. M. SWINKELS (1993): "Extensive Form Reasoning in Normal Form Games," *Econometrica*, 61, 273-302.
- (1997): "How Proper Is Sequential Equilibrium?," *Games and Economic Behavior*, 18, 193-218, Erratum, 19 (1997), 249.
- MASKIN, E., AND J. TIROLE (2001): "Markov Perfect Equilibrium I. Observable Actions," *Journal of Economic Theory*, 100(2), 191-219.
- MORRIS, S. (2008): "Purification," in *The New Palgrave Dictionary of Economics Second Edition*, ed. by S. Durlauf, and L. Blume, pp. 779-782. Macmillan Palgrave.

- MORRIS, S., AND H. S. SHIN (2003): "Global Games: Theory and Applications," in *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, ed. by L. H. M. Dewatripont, and S. Turnovsky. Cambridge University Press.
- MYERSON, R. B. (1978): "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 7, 73–80.
- OK, E. A. (2007): *Real Analysis with Economic Applications*. Princeton University Press.
- OSBORNE, M. J., AND A. RUBINSTEIN (1990): *Bargaining and Markets*. Academic Press, Inc., San Diego, CA.
- OWEN, G. (1982): *Game Theory*. Academic Press, second edn.
- PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–50.
- RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50(1), 97–109.
- (1989): "The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge," *American Economic Review*, 79(3), 385–391.
- SAMUELSON, L. (1997): *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA.
- SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 22–55.
- SHAKED, A., AND J. SUTTON (1984): "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model," *Econometrica*, 52(6), 1351–1364.
- SION, M., AND P. WOLFE (1957): "On a Game Without a Value," *Contributions to the Theory of Games*, 3, 299–306.

- 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.
- (1991): *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin, second, revised and enlarged edn.
- VOHRA, R. V. (2005): *Advanced Mathematical Economics*. Routledge, London and New York.
- WEIBULL, J. W. (1995): *Evolutionary Game Theory*. MIT Press, Cambridge.
- WEINSTEIN, J., AND M. YILDIZ (2007): "A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements," *Econometrica*, 75(2), 365-400.