Econ 210B: Game Theory

Simon Rubinstein–Salzedo

Fall 2005

0.1 Introduction

These notes are based on a graduate course on game theory I took from Professor Rod Garratt in the Fall of 2005. The primary textbooks were *Game Theory* by Drew Fudenberg and Jean Tirole and *Game Theory for the Applied Economist* by Robert Gibbons. Also recommended were *A Course in Microeconomic Theory* by David Kreps and *The Winner's Curse* by Richard Thaler.

Chapter 1

Games in Strategic Form

We seek to model and make predictions about the outcome in situations where

1. each decision maker, in a group, makes choices (or takes actions) independently and simultaneously.

2. the payoff to each player is determined by the profile of actions taken.

To model these situations, we introduce the formal notion of a **game**.

Definition. A game in strategic (normal) form consists of three elements:

1. A set of players $\mathcal{P} = \{1, 2, \ldots, I\}$.

2. A strategy set S_i of pure strategies for each player $i \in \mathcal{P}$, e.g. $S_i = \{\text{up, down}\}$.

3. For each $i \in \mathcal{P}$, a function $u_i : S \to \mathbb{R}$, where $S = \prod_{i \in \mathcal{P}} S_i$, which gives player i's von Neumann-Morgenstern utility for each profile $s \in S$.

Example. $\mathcal{P} = \{1, 2\}$, $S_1 = \{\text{up, down}\}$, $S_2 = \{\text{left, right}\}$, $u_1(\text{up, left})$=number, $u_2(\text{up, left})$=number.

We assume that

- each player is an expected utility maximizer (i.e. players are rational).

- the structure of the game and the rationality of the players is common knowledge.

Consider the following game:

	L	R
U	4,11	4,1
D	1,3	1,13

In this game, the row player ought to play U.

Now consider the following game:

	A	B
a	7,7	8,0
b	0,8	9,9

In this game, there are no clearly bad strategies.

1.1 Implications of Common Knowledge of Rationality

Consider the following game:

	L	C	R
u	4,3	5,1	6,2
m	2,1	8,1	3,6
d	3,0	9,6	2,8

The column player ought not to play C as it is dominated by R. Then the row player ought not to play m or d as the remaining possibilities are dominated by u. Then the column player ought not to play R as it is dominated by L. Hence (u, L) is the only rational strategy. This process is called iterated elimination of strictly dominated strategies. It does not always produce a solution, but the order in which we eliminate solutions does not matter.

Now consider the following game:

	L	R
u	2,0	-1,0
m	0,0	0,0
d	-1,0	2,0

This game has no strictly dominated strategies in pure strategies, however, we can mixed strategies as well. Let p, q, r be the probabilities the row player places on

u, m, d, respectively. Then $(p, q, r) = (1/2, 0, 1/2)$ strictly dominates m. Hence m is not a rational play. Thus we may examine the game

	L	R
u	2,0	-1,0
d	-1,0	2,0

instead, and we will not miss any Nash equilibria.

1.2 Mixed Strategies

A mixed strategy is a probability distribution over pure strategies available to a player.

Let Σ_i denote the set of mixed strategies for player i. A generic element is $\sigma_i : S_i \to [0,1]$, where $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$, e.g. if $S_i = \{\text{up,down}\}$, then perhaps $\sigma_i(\text{up}) = 1/4$ and $\sigma_i(\text{down}) = 3/4$. So $\sigma_i(s_i)$ is the probability σ_i assigns to $s_i \in S_i$. Suppose $\mathcal{P} = \{1, 2\}$. A mixed strategy profile is $\sigma = (\sigma_1, \sigma_2)$, e.g. $S_1 = \{\text{up,down}\}$, $S_2 = \{\text{left,right}\}$, $\sigma = ((1/4, 3/4), (1/2, 1/2))$.

A mixed strategy is degenerate if some action is taken with probability 1. A mixed strategy is completely mixed if each action is taken with a strictly positive probability.

Each player is assumed to randomize independently of all other players. Therefore, player i's expected payoff to a given mixed strategy profile σ is

$$u_i(\sigma) = \sum_{s \in S} \left(\prod_{j=1}^{I} \sigma_j(s_j) \right) u_i(s).$$

Consider the following game:

	ℓ: 1/2	r: 1/2
u: 1/4	4,0	3,9
d: 3/4	2,8	1,5

Then $u_1(\sigma) = \sigma_1(\text{up}) \cdot \sigma_2(\text{left}) \cdot 4 + \cdots$.

Note. A pure strategy s_i is **strictly dominated** for player i if there exists $\sigma_i \in \Sigma_i$ such that $\mathcal{P} = \{1, \ldots, I\}$, $u_i(\sigma, s_{-i})$, $s = (s_1, \ldots, s_I)$ is a pure strategy profile, where

$s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_I)$ and $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, where $S_{-i} = \prod_{j \in \mathcal{P} \setminus \{i\}} S_j$.

Unfortunately, few games can be solved by iterated elimination of strictly dominated strategies. Also, to have confidence in the prediction yielded by iterated elimination of strictly dominated strategies, players must be highly confident of the rationality of the players.

Example.

	L	R
u	1.01,5	-500,.4
d	1,.25	.9,.2

Then u, L is the optimal play, but if the row player is not fully confident in the rationality of the column player, then u is a very risky play.

1.3 Nash Equilibria

Definition. A (mixed) strategy profile σ^* is a Nash equilibrium if for every player i, $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*)$ for all $s_i \in S_i$.

In a Nash equilibrium, player i's payoff from following σ_i^* is at least as great as the payoff she would obtain from any one of her pure strategies. In fact, this implies she cannot do better by using some other mixed strategy.

Claim. Let σ^* be such that $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*)$ for all $s_i \in S_i$. Then $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ for all $\sigma_i \in \Sigma_i$.

Proof. Let σ_i be an arbitrary mixed strategy. Then

$$\begin{aligned} u_i(\sigma_i, \sigma_{-i}^*) &= \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i}^*) \\ &\leq \sum_{s_i \in S_i} \sigma_i(s_i) u_i(\sigma_i^*, \sigma_{-i}^*) \\ &= u_i(\sigma_i^*, \sigma_{-i}^*), \end{aligned}$$

as desired. ■

1.4 Pure Strategy Nash Equilibria

For the moment, we focus on pure strategies.

1.4.1 Computing Nash Equilibria

Example.

	FB	B
fb	2,1	0,0
b	0,0	1,2

The pure strategy Nash equilibria are (fb, FB) and (b, B).

Example. Cournot competition. Firms 1 and 2 each simultaneously choose nonnegative outputs q_1 and q_2. There is an inverse market demand $P(q) = 1 - q$, where $q = q_1 + q_2$. The cost is given by $c_i(q_i) = cq_i$, where $c < 1$ is the constant marginal cost. $u_i(q_i, q_{-i}) = q_i P(q) - c_i(q_i) = q_i(1 - q_1 - q_2) - cq_i$. Firm 1 solves

$$\max_{q_1} (1 - q_1 - q_2)q_1 - cq_1.$$

The first-order condition is $1 - 2q_1 - q_2 - c = 0$. Thus

$$q_1 = \begin{cases} \frac{1-q_2-c}{2} & q_2 \leq 1 - c, \\ 0 & \text{otherwise.} \end{cases}$$

We call this the reaction function $r_1(q_2)$. By symmetry,

$$q_2 = \begin{cases} \frac{1-q_1-c}{2} & q_1 \leq 1 - c, \\ 0 & \text{otherwise.} \end{cases}$$

This is $r_2(q_1)$. The Nash equilibrium is the intersection of these two curves, namely $\left(\frac{1-c}{3}, \frac{1-c}{3}\right)$.

1.4.2 Interpretations of Nash Equilibria

1. Nash equilibria can be interpreted as self-enforcing agreements (equilibrium after the fact).

2. Nash equilibria can also be interpreted as self-confirmatory predictors (equilibria in beliefs). Suppose that player 1 predicts that player 2 will play FB and player 2 predicts player 1 will play fb. Then player 1 optimally chooses fb, and player 2 optimally chooses FB. Each player's prediction is correct. Thus the predictions are self-fulfilling.

3. Nash equilibria are also the result of adaptive learning. If we start with an arbitrary point on the reaction curve and choose the optimal strategy and then continually refine the process, we converge to a Nash equilibrium.

1.5 Mixed Strategy Nash Equilibria

1. Many games have a mixed strategy Nash equilibrium that is nondegenerate, e.g. the Battle of the Sexes.

2. Some games have no pure strategy Nash equilibria, e.g. Matching pennies.

	h	t
H	1,−1	−1,1
T	−1,1	1,−1

We can find the mixed-strategy Nash equilibrium using reaction functions. The curves $r_1(\sigma_2(h))$ and $r_2(\sigma_1(H))$ intersect at $(1/2, 1/2)$, so this is the Nash equilibrium.

1.5.1 Finding Mixed-Strategy Nash Equilibria

The following fact is very useful for finding mixed-strategy Nash equilibria.

Claim. If $u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i$, then $u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i$ with $\sigma_i(s_i) > 0$.

This means that you must be indifferent in order to randomize.

Proof. (By contradiction.) Suppose $u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i$, but that $u_i(\sigma_i, \sigma_{-i}) > u_i(s_i', \sigma_{-i})$ for some $s_i' \in S_i$ with $\sigma_i(s_i') > 0$. Then

$$\begin{aligned} u_i(\sigma_i, \sigma_{-i}) &= \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i}) \\ &, \sum_{s_i \in S_i} \sigma_i(s_i) u_i(\sigma_i, \sigma_{-i}) \\ &= u_i(\sigma_i, \sigma_{-i}), \end{aligned}$$

a contradiction. ∎

Example.

	FB	B
fb	2,1	0,0
b	0,0	1,2

Suppose $\sigma_1(fb) = p$ and $\sigma_2(FB) = q$. To get player 1 to randomize, we require $2q + 0(1-q) = 0q + 1(1-q)$, or $q = \frac{1}{3}$. For player 2 to randomize, we require $p \cdot 1 + (1-p) \cdot 0 = p \cdot 0 + (1-p)2$, or $p = \frac{2}{3}$. The induced probability distribution over strategy profiles is

	FB	B
fb	pq	$p(1-q)$
b	$(1-p)q$	$(1-p)(1-q)$

=

2/9	4/9
1/9	2/9

The expected payoff to each player is $\frac{2}{9} \cdot 2 + \frac{2}{9} \cdot 1 = \frac{2}{3}$.

Exercise. Find all the Nash equilibria of the following game:

	L	R
u	1,0	−1,0
d	0,1	0,0

(u, L) is the unique pure strategy Nash equilibrium. There is also a mixed strategy Nash equilibrium.

Other facts.

1. Theorem. (Nash, 1950) Every finite strategic (normal) form game has a Nash equilibrium.

2. Theorem (Wilson, 1971) Almost all finite strategic form games have a finite and odd number of Nash equilibria.

	L	R
u	1,1	0,0
d	0,0	0,0

   (u, L) and (d, R) are the only Nash equilibria. Thus the above game is non-generic. Note that one can think of a 2×2 game (for example) as a point in $\mathbb{R}^8$.

3. If a unique strategy profile s^* survives iterated elimination of strictly dominated strategies, then s^* is the unique Nash equilibrium.

4. Only those strategies which survive iterated elimination of strictly dominated strategies are played with positive probability in a Nash equilibrium.

1.6 Correlated Equilibria

Premise. Suppose that the players of a game can build a "signaling device" which makes to each player a non-binding recommendation of what action to choose.

Example.

	L q	R
p u	5,1	0,0
d	4,4	1,5

(u, L) and (d, R) are Nash equilibria. There is also a mixed-strategy Nash equilibrium with $p = \frac{1}{2}$ and $q = \frac{1}{2}$; i.e. $\sigma_1^*(u) = \sigma_2^*(L) = \frac{1}{2}$. Moreover, $u_i(\sigma_1^*, \sigma_2^*) = 2.5$ for $i = 1, 2$.

Denote a signaling device by ρ which maps strategy profiles into $[0, 1]$.

Device 1: $\rho(u, L) = \rho(d, R) = \frac{1}{2}$. Consider player 1. Suppose he is recommended to play u. Then he can calculate $\rho(L \mid u) = \frac{\rho(L|u)}{\rho(u)} = 1$. Likewise, $\rho(R \mid u) = 0$. Thus his expected utility of playing u (or d) is $u_1(u \mid u) = 5$, $u_1(d \mid u) = 4$ assuming the other player follows her recommendation! Hence player 1 optimally obeys his recommendation. One can show that the same holds for player 2. Device 1 is therefore a correlated equilibrium.

Device 2: $\rho(u, L) = \rho(d, R) = \rho(d, L) = \frac{1}{3}$. Consider player 1. Suppose he is told to play d. Then he calculates $\rho(L \mid d) = \frac{1}{2}$, $\rho(R \mid d) = \frac{1}{2}$. Thus, assuming player 2 follows her recommendation, then player 1 calculates $u_1(d \mid d) = 2.5$, $u_1(u \mid d) = 2.5$, $u_1(u \mid u) = 5$, and $u_1(d \mid u) = 4$. Hence the device works. The expected payoff to each player under this correlated equilibrium is $\frac{10}{3}$.

Device 3: $\rho(d, L) = 1$. This doesn't work.

Definition. A correlated equilibrium is any probability distribution $\rho(\cdot)$ over pure strategy profiles in S such that for each player i and every $s_i \in S_i$ with $\rho(s_i) > 0$,

$$\sum_{s_{-i} \in S_{-i}} \rho(s_{-i} \mid s_i) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \rho(s_{-i} \mid s_i) u_i(s_i', s_{-i})$$

for all $s_i' \in S_i$.

Remark. The probability distribution over strategy profiles that is produced by a Nash equilibrium of a game is a correlated equilibrium.

Chapter 2

Games in Extensive Form

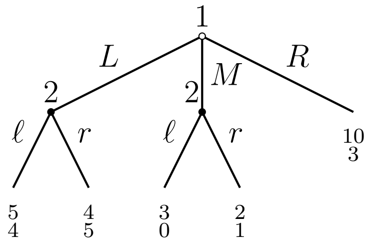

This game is one of perfect information, which means that whenever a player moves, she knows exactly where she is in the game tree.

Alternatively, we can represent games of imperfect information using non-singleton information sets.

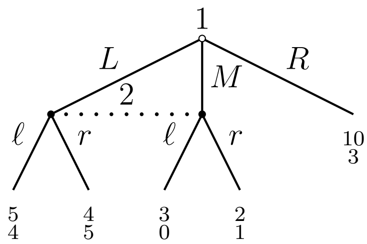

2.1 Subgame Perfect Equilibria

Definition. A **subgame** G of an extensive form game T consists of a single node and all of its successor nodes in T with the property that if $x' \in G$ and $x', x'' \in h$ (an information set), then $x'' \in G$.

Definition. A **pure behavior strategy** for player i is a map $s_i : H_i \to A_i$ (where H_i is the collection of information sets available to i and A_i is the set of actions to player i).

2.2 Pure Behavioral Strategies

A strategy for a player is a plan of action for every information set controlled by the player. Let H_i denote the set of player i's information sets, and let $A_i = \bigcup_{h_i \in H_i} A(h_i)$, where $A(h_i) = A(x)$ is the set of actions available at node x for $x \in h_i$.

Definition. A pure (behavioral) strategy for player i is a map $s_i : H_i \to A_i$ with $s_i(h_i) \in A(h_i)$ for all $h_i \in H_i$.

Definition. A pure (behavioral) strategy profile s of an extensive form game is a subgame perfect equilibrium if the restriction of s to G is a Nash equilibrium of G for every G.

2.3 Games in Extensive Form (Formal Treatment)

The elements of extensive form games are the following:

1. A (finite) set of players $\mathcal{P} = \{1, \ldots, I\}$ and the player nature, denoted by N.
2. An arborescence $(X, \mapsto)$, where X is a finite set of nodes and $\mapsto$ is a binary operator denoting precedence. $(X, \mapsto)$ satisfies:
   (a) $\mapsto$ is asymmetric and transitive.
   (b) If $x \mapsto x''$ and $x' \mapsto x''$, then either $x \mapsto x'$ or $x' \mapsto x$.

   Note. We assume there is a single node that is a predecessor to all other nodes.

   Definition. x' is an immediate predecessor of x'' if for any $x \mapsto x''$ with $x \neq x'$, then $x \mapsto x'$.

3. A player partition given by the function $i : X \to \{N, 1, 2, \ldots, I\}$. For every $x \in X$, $i(x)$ gives the name of the player who moves at decision node x. Nature moves at the initial node.

4. For each $x \in X$, a (finite) set of actions $A(x)$ and a function $\ell_x : s(x) \to A(x)$, where $s(x)$ is the successor of x. We require that ℓ_x be a bijection.

5. An information partition H of the decision nodes of X. Let $h \in H$ be an information set. The interpretation of h is that if two nodes x and x' belong to h, the player moving at h cannot distinguish between x and x'. We require that if $x, x' \in h$, then

   (a) neither node is a predecessor of the other,

   (b) $i(x) = i(x')$,

   (c) $A(x) = A(x')$.

   Information sets are indicated in the game trees by dashed lines.

2.4 Perfect Recall Assumption

Informally, no player ever forgets any information he once knew, and all players know the actions they have chosen previously. We may add the following condition to (5) above:

5. (d) if $x'' \in h(x')$, $x \mapsto x'$, and $i(x) = i(x')$, then there is an $\widehat{x}$ (possibly x itself) with $\widehat{x} \in h(x)$, $\widehat{x} \mapsto x''$, $\ell_x(x') = \ell_{\widehat{x}}(x'')$.

Example. (Fudenberg and Tirole, page 81.)

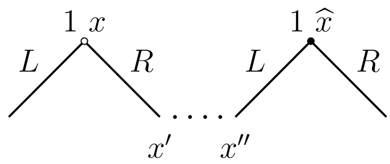

Example.

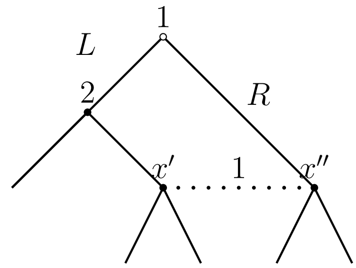

6. For each $i \in \mathcal{P}$, a function $u_i : Z \to \mathbb{R}$ (where Z is the set of terminal nodes) which gives the payoff to player i at each terminal node.

7. For each x such that $i(x) = N$, a probability distribution over $A(x)$, e.g.

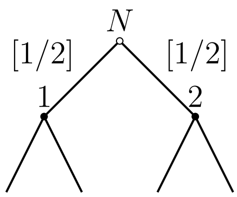

2.5 Strategies

Definition. A pure behavior strategy for player i is a map $s_i : H_i \to A_i$ with $s_i(h_i) \in A(h_i)$ for all $h_i \in H_i$. A **path** of s is the set of information sets reached with positive probability under s, e.g.

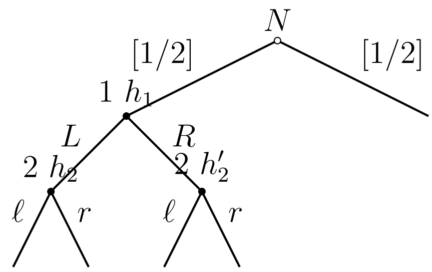

$s_1(h_1) = L$, $s_2(h_2) = \ell$, $s_2(h_2') = r$. The path of s is $\{h_1, h_2\}$.

2.6 Normal Form Representation of Extensive Form Games

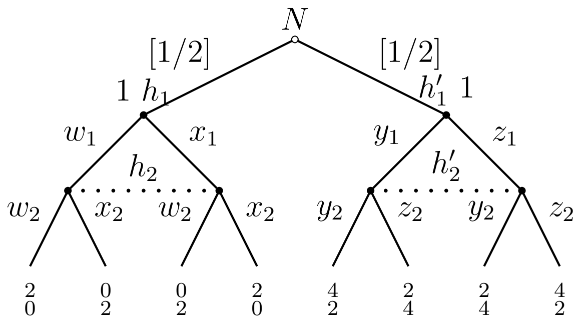

The normal form representation is

	w_2y_2	w_2z_2	x_2y_2	x_2z_2
w_1y_1	3,1	2,2	2,2	1,3
w_1z_1	2,2	3,1	1,3	2,2
x_1y_1	2,2	1,3	3,1	2,2
x_1z_1	1,3	2,2	2,2	3,1

It is easy to see that $A = (\sigma_1 = (0, 1/2, 1/2, 0),\ \sigma_2 = (0, 1/2, 1/2, 0))$ is a Nash equilibrium. It is also easy to see that $B = (\sigma'_1 = (1/2, 0, 0, 1/2),\ \sigma'_2 = (1/2, 0, 0, 1/2))$ is a Nash equilibrium. In fact, any convex combination of A and B is also a Nash equilibrium. All these Nash equilibria are equivalent in the following sense: consider the strategy $\widetilde{\sigma}_1 = (1/4, 1/4, 1/4, 1/4)$. What is the probability that player 1's move is w_1 at information set h_1? $\text{Prob}[w_1 \mid h_1] = \frac{\text{Prob}[w_1, h_1]}{\text{Prob}[h_1]} = \frac{1/2 \cdot 1/2}{1/2} = \frac{1}{2}$. Similarly, $\text{Prob}[y_1 \mid h_1] = \frac{1}{2}$. Thus we can view the strategy $\widetilde{\sigma}_1$ as player 1 playing each of his actions with equal probability at each information set. In fact, all Nash equilibrium strategies for player 1 generate the same "behavioral strategy" $\sigma_1(w_1 \mid h_1) = \sigma(y_1 \mid h'_1) = \frac{1}{2}$. In this game, there is a *unique* Nash equilibrium in behavioral strategies.

2.7 Behavioral Strategies More Generally

Let $\Delta(A(h_i))$ be the set of probability distributions on $A(h_i)$.

Definition. A behavioral (mixed) strategy for player i is an element of

$$\prod_{h_i \in H_i} \Delta(A(h_i)).$$

Example. Suppose $H_1 = \{h_1, h_1', h_1''\}$. Then $\prod_{h_i \in H_i} \Delta(A(h_i))$ denotes the set of all triples $\langle \sigma_1(\cdot \mid h_1), \sigma_1(\cdot \mid h_1'), \sigma_1(\cdot \mid h_1'')\rangle$, where $\sigma_1(\cdot \mid h_1) \in \Delta(A(h_1))$, $\sigma_1(\cdot \mid h_1') \in \Delta(A(h_1'))$, and $\sigma_1(\cdot \mid h_1'') \in \Delta(A(h_1''))$. Typically a Nash equilibrium is more compactly represented by a profile of behavioral strategies.

Note. Not all Nash equilibria of extensive form games are reasonable.

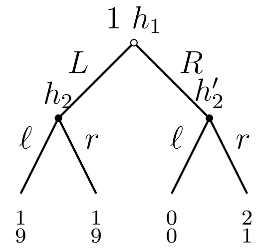

This is a game of perfect information. The Nash equilibria are $(L, \ell\ell)$, $(R, \ell r)$, $(L, r\ell)$, and (R, rr). Not all of these are reasonable. By backward induction, (R, rr) and $(R, \ell r)$ are the solutions.

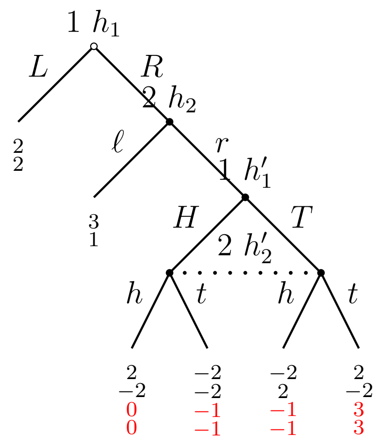

If we use the red payoffs, then there are multiple subgame perfect equilibria. A Nash equilibrium in behavioral strategies is $\sigma_1(L \mid h_1) = 1$, $\sigma_1(H \mid h_1') = \frac{1}{2}$, $\sigma_2(r \mid h_2) = 1$, $\sigma_2(h \mid h_2') = \frac{1}{2}$. However, the move r by play 2 in this Nash equilibrium is not reasonable.

2.8 The Idea of Subgame Perfection

At any point during play of the game, each player should expect that play proceeds according to *some* Nash equilibrium of the subgame that remains. Given a behavioral strategy profile σ of the original game T, we can define a restriction of σ to any subgame G in the obvious way. Let σ_i be a behavioral strategy for player i, and let $\widehat{H}_i$ be the set of information sets for player i in the subgame G.

Definition. A restriction of σ_i to the subgame G is $\widehat{\sigma}_i(\cdot \mid h_i) = \sigma_i(\cdot \mid h_i)$ for every $h_i \in \widehat{H}_i$.

Example. In the game we just considered, we have $\sigma_i(L \mid h_1) = 1$, $\sigma_1(H \mid h_1') = \frac{1}{2}$. The restriction of σ_1 to the smallest subgame is $\widehat{\sigma}_1(H \mid h_1') = \frac{1}{2}$.

Definition. A behavioral strategy profile σ of an extensive form game T is a subgame perfect equilibrium if the restriction of σ to G is a Nash equilibrium for every G.

Again looking at the example: consider the smallest subgame G. The unique Nash equilibrium is

$$\sigma_2(H \mid h_1') = \frac{1}{2}, \qquad \sigma_2(h \mid h_2') = \frac{1}{2}. \tag{$*$}$$

Now consider the second smallest subgame G'. The unique Nash equilibrium that is consistent with $(*)$ is

$$\sigma_1(H \mid h_1') = \frac{1}{2}, \qquad \sigma_1(h \mid h_2') = \frac{1}{2}, \qquad \sigma_2(\ell \mid h_2) = 1.$$

2.9 Stackelburg Competition Revisited

Firm 1 chooses an output q_1. Firm 2 observes q_1 and then chooses its own output q_2. Payoffs are given by $\prod_i(q_i, q_{-i}) = q_i(1 - q_1 - q_2) - cq_i$. A strategy for player 1 is a number $q_i \in [0, Q_1]$. A strategy for player 2 is a function $r_2 : [0, Q_1] \to [0, Q_2]$. Consider subgame q_1: Firm 2 solves $\max_{q_2} q_2(1 - q_1 - q_2) - cq_2$, so

$$r_2(q_1) = \begin{cases} \frac{1-q_1-c}{2} & q_1 \le 1 - c, \\ 0 & q_1 > 1 - c. \end{cases}$$

Now we can consider the overall game. Given $q_1 \le 1 - c$, firm 1's profit is

$$\max_{q_1} q_1 \left(1 - q_1 - \frac{1 - q_1 - c}{2}\right) - cq_1,$$

and the solution is $q_1^* = \frac{1-c}{2}$, $q_2^* = \frac{1-c}{4}$. Hence the subgame perfect equilibrium is

$$q_1^* = \frac{1-c}{2}, \qquad r(q_1) = \begin{cases} \frac{1-q_1-c}{2} & q_1 \le 1 - c, \\ 0 & q_1 > 1 - c. \end{cases}$$

Perfect information means you always know where you are in the game tree (no non-singleton information sets). Complete information means you know your opponent, payoffs, etc.

Chapter 3

Static Games of Incomplete Information

Incomplete means some players do not know the payoffs (or types) of other players.

Example. Cournot competition under asymmetric information. Two firms simultaneously choose levels of output q_1 and q_2. $P(q) = a - q$, where $q = q_1 + q_2$. The cost of firm 1 is $c_1(q_1) = cq_1$. The cost of firm 2 is $c_2(q_2) = \begin{cases} c_H q_2 & \text{with probability } \gamma, \\ c_L q_2 & \text{with probability } 1-\gamma, \end{cases}$ $c_H > c_L$.

3.1 Bayesian Equilibrium or Bayes-Nash Equilibrium

Firm 2 knows its true cost function. Firm 1 knows Firm 2's costs is c_H with probability γ, etc. Both firms know Firm 1's cost function. We first solve Firm 2's problem. If Firm 2 is high cost, then $q_2 H^*$ solves

$$\max_{q_2} [a - q_1 - q_2 - c_H] q_2.$$

Hence $q_2^{H*} = \frac{a-q_1-c_H}{2}$. Similarly, $q_2^{L*} = \frac{a-q_1-c_L}{2}$. We can now solve Firm 1's problem. Firm 1 has to work out

$$\max_{q_1} \{\gamma[(a - q_1 - q_2^{H*}) - c]q_1 + (1-\gamma)[(a - q_1 - q_2^{L*}) - c]q_1\}.$$

The first-order necessary condition is therefore

$$\gamma[a - 2q_1 - q_2^{H*} - c] + (1-\gamma)[a - 2q_1 - q_2^{L*} - c] = 0.$$

Hence

$$q_1^* = \frac{a - [\gamma q_2^{H*} + (1-\gamma) q_2^{L*}] - c}{2}.$$

So we have three equations and three variables.

$$q_1^* = \frac{2}{3}\left[a + \frac{\gamma(c_H - c_L)}{2} - \frac{a - c_L}{2} - c\right],$$

etc. These values are the Bayesian equilibrium.

3.2 Bayesian Games

1. A finite set of players $\mathcal{P} = \{1, \ldots, I\}$.
2. For each $i \in \mathcal{P}$, a finite set Θ_i of types for player i. Let $\theta_i \in \Theta_i$ denote a generic element, and let $\Theta = \prod_{i \in \mathcal{P}} \Theta_i$. So $\theta \in \Theta$ is a type profile.
3. A probability distribution $p : \Theta \to \mathbb{R}$ over type profiles.

   Note. p is the same for all players, i.e. players have **common priors**.

   Aside. Let $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$. Assume that for each $i \in \mathcal{P}$,

   $$p(\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_i, \theta_{-i}) > 0 \text{ for all } \theta_i.$$

   Then $p_i(\theta_{-i} \mid \theta_i) = \frac{p(\theta_1, \ldots, \theta_I)}{p(\theta_i)}$ is the probability that player i assigns to the event that the profile of types for everyone else is θ_{-i} conditional on his own type being θ_i.
4. For each $i \in \mathcal{P}$, a set S_i of actions of player i. (S_i might depend on type.)
5. For each $i \in \mathcal{P}$, a function $u_i : S \times \Theta \to \mathbb{R}$ that gives player i's von Neumann-Morgenstern utility for each profile s of strategies and θ of types.

Remark. In the Cournot example, we had $\Theta_1 = \{c\}$, $\Theta_2 = \{c_H, c_L\}$, $p(c, c_H) = \gamma$, $p(c, c_L) = 1 - \gamma$, $S_1 = S_2 = [0, \infty)$. $u_1(q_1, q_2, \theta_1, \theta_2) = [a - q_1 - q_2 - \theta_1]q_1$, $u_2(q_1, q_2, \theta_1, \theta_2) = [a - q_1 - q_2 - \theta_2]q_2$.

Definition. A pure strategy for player i is a function $s_i : \Theta_i \to S_i$, where for each $\theta_i \in \Theta_i$, $s_i(\theta_i)$ gives the action in S_i taken by player i when of type θ_i.

Let $s(\cdot)$ denote a strategy profile. Let $s_{-i}(\theta_{-i}) = (s_1(\theta_1), \ldots, s_{i-1}(\theta_{i-1}), s_{i+1}(\theta_{i+1}), \ldots, s_I(\theta_I))$. Assume that for each player, each of her types occurs with strictly positive probability.

Definition. A strategy profile $s(\cdot)$ is a Bayesian equilibrium if for each player i and all $\theta \in \Theta$,

$$s_i(\theta_i) \in \arg\max_{s_i \in S_i} \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i, s_{-i}(\theta_{-i}), \theta).$$

In the Cournot example, $s_1(c) = q_1^*$, $s_2(c_H) = q_2^{H*}$, $s_2(c_L) = q_2^{L*}$.

3.3 Samuelson-Chatterjee Sealed Bid Auction

The buyer and seller simultaneously propose prices P_b and P_s, respectively. If $P_b > P_s$, then they trade at the price of $\frac{1}{2}(P_b + P_s)$. If $P_b < P_s$, then they don't trade.

Assume that the buyer and seller have values $\widetilde{V}_1$ and $\widetilde{V}_2$ that are drawn independently from the uniform distribution on $[0, 1]$. Formally, the game is given by

1. Players: $\mathcal{P} = \{b, s\}$.
2. Strategies: A strategy for the buyer is a mapping from his valuation to his bid, i.e. $P_b(\cdot) : [0, 1] \to [0, 1]$. Similarly for the seller.
3. Payoffs:

$$u_b(P_b, P_s, V_s, V_b) = \begin{cases} V_b - \frac{P_b + P_s}{2} & P_b \geq P_s, \\ 0 & P_b < P_s, \end{cases}$$

$$u_s(P_b, P_s, V_s, V_b) = \begin{cases} \frac{P_b + P_s}{2} - V_s & P_b \geq P_s, \\ 0 & P_b < P_s. \end{cases}$$

A Bayesian equilibrium is a pair $\{P_b(\cdot), P_s(\cdot)\}$ such that

1. For each $V_b \in [0, 1]$, $P_b(r_b)$ solves

$$\max_{P_b} \text{Prob}[P_s(\widetilde{V}_s) < P_b] \left[V_b - \frac{P_b + E[P_s(\widetilde{V}_s) \mid P_s(\widetilde{V}_s) \leq P_b]}{2} \right].$$

2. For each $V_s \in [0,1]$, $P_s(V_s)$ solves

$$\max_{P_s} \text{Prob}[P_b(\widetilde{V}_b) > P_s] \left[\frac{P_s + E[P_b(\widetilde{V}_b) \mid P_b(\widetilde{V}_b) \geq P_s]}{2} - V_s \right].$$

We look for an equilibrium in linear strategies, i.e. $P_s(V_s) = \alpha_s + c_s V_s$ and $P_b(V_b) = \alpha_b + c_b V_b$.

Aside. Computing $E[P_s(\widetilde{V}_s) \mid P_s(\widetilde{V}_s) \leq P_b]$. For ease of notation, let $\widetilde{P}_s = P_s(\widetilde{V}_s)$. First we need the appropriate density function.

$$E[\widetilde{P}_s \mid \widetilde{P}_s \leq P_b] = \int_{\alpha_s}^{P_b} P_s \cdot \frac{1}{P_b - \alpha_s}\, dP_s = \frac{P_b + \alpha_s}{2}.$$

We also need

$$\text{Prob}[\widetilde{P}_s \leq P_b] = \text{Prob}[\alpha_s + c_s \widetilde{V}_s \leq P_b] = \text{Prob}\left[\widetilde{V}_s \leq \frac{P_b - \alpha_s}{c_s}\right] = \frac{P_b - \alpha_s}{c_s}.$$

So, with linear strategies, we have

$$\max_{P_b} \frac{P_b - \alpha_s}{c_s} \left[V_b - \frac{P_b + \frac{P_b + \alpha_s}{2}}{2} \right] \iff \max_{P_b} \frac{P_b - \alpha_s}{c_s} \left[V_b - \frac{3P_b + \alpha_s}{4} \right].$$

Take a derivative with respect to P_b and set equal to 0. $P_b = \frac{2}{3}V_b + \frac{1}{3}V_s$. Similarly for the seller, her best response is $P_s = \frac{2}{3}V_s + \frac{1}{3}(c_b - \alpha_b)$. Since $P_b = \alpha_b + c_b V_b$ and $P_s = \alpha_s + c_s V_s$, we have $\alpha_s = \frac{1}{4}$ and $\alpha_b = \frac{1}{12}$. We have found a Bayesian equilibrium in linear strategies:

$$P_b(V_b) = \frac{2}{3}V_b + \frac{1}{12},$$
$$P_s(V_s) = \frac{2}{3}V_s + \frac{1}{4}.$$

Definition. A Bayesian equilibrium $(P_b(\cdot), P_s(\cdot))$ is expost efficient if $P_b(V_b) > P_s(V_s)$ if and only if $V_b > V_s$, i.e. it is expost efficient if the good trades whenever $V_b > V_s$ and conversely.

In this case, they trade if and only if $V_b \geq V_s + \frac{1}{4}$. Hence this Bayesian equilibrium is not expost efficient.

Myerson and Sattherthwaite showed that there is no expost efficient Bayesian equilibrium for this game.

3.4 First-Price Sealed Bid Auction

We have two bidders, $i = 1, 2$. $\widetilde{V}_i$ is bidder i's value, distributed independently and uniformly on $[0, 1]$. b_i denotes player i's bid, $b_i \in [0, 1]$. Each bidder knows her own value, but not the other player's value. However, the distribution function that determines each bidder's value is common knowledge. The highest bid wins the auction and pays the bid.

We analyze this setup as a Bayesian game.

The payoffs are

$$u_i(b_1, b_2, V_1, V_2) = \begin{cases} v_i - b_i & b_i > b_{3-i}, \\ \frac{v_i - b_i}{2} & b_i = b_{3-i}, \\ 0 & b_i < b_{3-i}. \end{cases}$$

The strategies are functions $b_i(\cdot) : [0, 1] \to [0, 1]$. $(b_1(\cdot), b_2(\cdot))$ is a Bayesian equilibrium if for each bidder i and $V_i \in [0, 1]$, $b_i(V_i)$ solves

$$\max_{b_i} (V_i - b_i) \cdot \text{Prob}[b_{-i}(\widetilde{V}_{-i}) < b_i] + \frac{1}{2}(V_i - b_i) \cdot \text{Prob}[b_{-i}(\widetilde{V}_{-i}) = b_i].$$

Claim. There is a unique symmetric Bayesian equilibrium in strictly increasing and differentiable strategies.

Proof. Suppose $(b(\cdot), b(\cdot))$ is a Bayesian equilibrium and $b(\cdot)$ is strictly increasing and differentiable. If $b(\cdot)$ is part of a Bayesian equilibrium, then for each player i and each $V_i \in [0, 1]$, $b(B_i)$ solves

$$\max_{b_i} (V_i - b_i) \cdot \text{Prob}[b(V_{-i}) < b_i].$$

Now

$$\text{Prob}[b(\widetilde{V}_{-i}) < b_i] = \text{Prob}[\widetilde{V}_{-i} < b^{-1}(b_i)] = b^{-1}(b_i)$$

by uniformity. Thus player i's problem (when her value is V_i) is

$$\max_{b_i} (V_i - b_i) b^{-1}(b_i).$$

The first-order necessary condition for an interior solution is

$$-v^{-1}(b_i) + (V_i - b_i)\frac{1}{b'(b^{-1}(b_i))} = 0. \qquad (*)$$

This is an implicit solution for player i's best response to $b(\cdot)$. If $(b(\cdot), b(\cdot))$ is a Bayesian equilibrium, we require that $b_i = b(V_i)$ satisfy $(*)$. Substituting $b(\cdot)$ into $(*)$ gives

$$\begin{aligned} -b^{-1}(b(V_i)) + (V_i - b(V_i))\frac{1}{b'(b^{-1}(b(V_i)))} &= 0 \\ -V_i + (V_i - b(V_i))\frac{1}{b'(V_i)} &= 0. \end{aligned}$$

We now rewrite this as

$$V_i b'(V_i) + b(V_i) = V_i. \qquad (**)$$

Observe that

$$\frac{d[b(V_i) - V_i]}{dV_i} = b'(V_i)V_i + b(V_i),$$

so

$$\frac{d[b(V_i) - V_i]}{dV_i} = V_i.$$

Integrating both sides, we have $b(V_i)V_i = \frac{V_i^2}{2} + K$, or $b(V_i) = \frac{V_i}{2} + \frac{K}{V_i}$. Suppose $K > 0$. Then for V_i small enough, $b(V_i) > V_i$, which is never optimal. Suppose $K < 0$. Then for V_i small enough, $b(V_i) < 0$, which is not allowed. Hence $K = 0$, and so $b(V_i) = \frac{V_i}{2}$ is the solution, i.e. $(b(\cdot), b(\cdot)) = (1/2V_1, 1/2V_2)$ is the unique symmetric Bayesian equilibrium in increasing and differentiable strategies. ■

3.5 Second-Price Auction

High bidder wins but pays the second-highest bid. Let $b^{(1)}_{-i}$ denote $\max\{b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n\}$.

	$b^{(1)}_{-i} > V_i$	$b^{(1)}_{-i} = V_i$	$b^{(1)}_{-i} < V_i$
$b_i > V_i$	win: $V_i - b^{(1)}_{-i} < 0$ lose: 0	0	$V_i - b^{(1)}_{-i} > 0$
$b_i = V_i$	0	0	$V_i - b^{(1)}_{-i} > 0$
$b_i < V_i$	0	0	win: $V_i - b^{(1)}_{-i} > 0$ lose: 0

We see that $b_i = V_i$ is a dominating strategy.

3.6 The Revelation Principle and Mechanism Design Economic Environment

1. A finite set of players $\mathcal{P} = \{1, \dots, I\}$.
2. For each $i \in \mathcal{P}$, a (finite) set Θ_i of types for player i.
3. A probability distribution ρ over type profiles.
4. A set $Y = X \times \mathbb{R}^I$ of feasible allocations, where X is the set of feasible allocations and $\mathbb{R}^I$ is the set of monetary transfers.
5. For each player i, a function $u_i : Y \times \Theta \to \mathbb{R}$ which gives player i's utility for each allocation $y \in Y$ and type profile $\theta \in \Theta$.
6. For each player i, a reservation utility $\bar{u}_i$.

Example. A trading game.

1. $\mathcal{P} = \{1, 2\}$.
2. $\Theta_s = \{h, \ell\}$, $\Theta_b = \{\theta_b\}$.
3. $\rho(h, \theta_b) = .2$, $\rho(\ell, \theta_b) = .8$.
4. $Y = \{0, 1\} \times \mathbb{R}$.
5. $u_s(x, t, h) = t - 40x$, $u_s(x, t, \ell) = t - 20x$, $u_b(x, t, h) = 50x - t$, $u_s(x, t, \ell) = 30x - t$.

The planner's problem is to choose a mechanism that achieves some objective, e.g. efficiency or revenue maximization. A mechanism consists of

1. For each player i, a set $\mathfrak{M}_i$ of messages.
2. A function $y : \mathfrak{M} \to Y$ which maps messages into allocations.

Remark. Since players' types are private information, the function y depends on θ only insofar as the messages depend on θ.

The economic environment and the mechanism yield a Bayesian game. For each player i, denote a pure strategy by $\mu_i : \Theta_i \to \mathfrak{M}_i$. Given a strategy profile $\mu = (\mu_1, \ldots, \mu_I)$, player i's expected utility when her type profile is θ_i is

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i} \mid \theta_i) u_i(y < \mu(\theta_i, \theta_{-i}), (\theta_i, \theta_{-i})).$$

3.7 The Revelation Principle

Suppose that a mechanism with message space $\mathfrak{M} = \mathfrak{M}_1 \times \cdots \times \mathfrak{M}_I$ and allocation function y has a Bayesian equilibrium $\mu^* = (\mu_1^*, \ldots, \mu_I^*)$. Then there is an equivalent (for any type profit, it produces the same allocation) mechanism with message space $\Theta = \Theta_1 \times \cdot \times \Theta_I$ in which individuals are willing to participate in the mechanism, and in the Bayesian equilibrium individuals report their types truthfully.

Proof. Define the new allocation rule $\bar{y} : \Theta \to Y$ by $\bar{y}(\theta) = y(\mu^*(\theta))$. Suppose that $\bar{y}$ is not a Bayesian equilibrium allocation, i.e. suppose for some player i and types θ_i and $\theta_i' \neq \theta_i$, that

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i} \mid \theta_i) u_i(\bar{y}(\theta_{-i}, \theta_i'), (\theta_{-i}, \theta_i)) > \sum_{\theta_{-i} \in \Theta_{-i}} \rho_i(\theta_{-i} \mid \theta_i) u_i(\bar{y}(\theta_{-i}, \theta_i), (\theta_{-i}, \theta_i)).$$

Then let $m_i' \in \mathfrak{M}_i$ be such that $\mu_i^*(\theta_i') = m_i'$. Then we have

$$\sum_{\theta_{-i} \in \Theta_{-i}} \rho_i(\theta_{-i} \mid \theta_i) u_i(y(\mu_{-i}^*(\theta_{-i}), m_i'), (\theta_{-i}, \theta_i))$$
$$> \sum_{\theta_{-i} \in \Theta_{-i}} \rho(\theta_{-i} \mid \theta_i) u_i(y(\mu_{-i}^*(\theta_{-i}), \mu_i^*(\theta_i)), (\theta_{-i}, \theta_i)),$$

which contradicts the fact that μ^* is a Bayesian equilibrium of the original mechanism. ■

We now return to the trading game. What can we say about the mechanisms we could potentially design?

By the revelation principle, we can restrict attention to a direct mechanism in which it is a Bayesian equilibrium for each player to reveal her type truthfully and every player is willing to participate in the mechanism. Set $\mathfrak{M}_i = \Theta_i$. Let t_{θ_s} denote the transfer from buyer to seller when seller reports type θ_s. Let x_{θ_s} be the probability that the good is exchanged when the seller reports type θ_s.

We have the following incentive compatibility constraints:

$$\text{IC—}h\text{: } t_h - 40x_h \geq t_\ell - 40x_\ell,$$
$$\text{IC—}\ell\text{: } t_\ell - 20x_\ell \geq t_h - 20x_h.$$

We also have the following individual rationality constraints:

$$\text{IR—}h\text{: } t_h - 40x_h \geq 0,$$
$$\text{IR—}\ell\text{: } t_\ell - 20x_\ell \geq 0,$$
$$\text{IR—}\theta_B\text{: } .2[50x_h - t_h] + .8[30x_\ell - t_\ell] \geq 0.$$

Proposition. There is no incentive compatible and individually rational mechanism that is expost efficient.

Proof. Expost efficiency requires $x_h = x_\ell = 1$.

$$\left.\begin{array}{ll} \text{IC—}h & t_h \geq t_\ell \\ \text{IC—}\ell & t_\ell \geq t_h \\ \text{IR—}h & t_h \geq 40 \\ \text{IR—}\ell & t_\ell \geq 20 \end{array}\right\} \Rightarrow t_h = t_\ell \Rightarrow t_h = t_\ell \geq 40.$$

We also require

$$\text{IR—}\theta_B\text{: } .2[50 - t_h] + .8[30 - t_\ell] \geq 0,$$

or

$$34 \geq .2t_h + .8t_\ell = t_h.$$

This is not possible, hence there is no Bayesian equilibrium of the direct mechanism which is expost efficient. Hence, by the revelation principle, there is no Bayesian equilibrium of any mechanism that is expost efficient.

Chapter 4

Perfect Bayesian Equilibria

Consider the following game of complete but imperfect information:

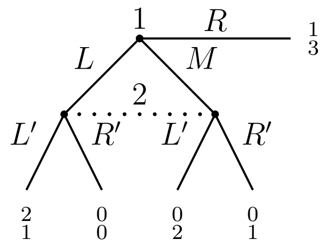

There are two pure strategy Nash equilibria, namely (L, L') and (R, R'). Both are also subgame perfect equilibria, as there is only one subgame. However, (R, R') is unreasonable, in the sense that if player 2 gets to move she would always go L'. We rule out such "unreasonable" Nash equilibria or subgame perfect equilibria by adding the following requirements:

1. At each information set, the player who moves must have a belief about which node in the information set has been reached. (A belief at an information set is a probability distribution over the nodes in the information set.)

2. Given their beliefs, the players' strategies must be *sequentially rational*, which means that for each information set controlled by a player, the action taken by the player must be optimal given the players' beliefs and the other players' strategies.

There are no beliefs for player 2 for which playing R' would be optimal.

Definition. For a given equilibrium, an information set is on the equilibrium path if it is reached with positive probability when the game is played according to the equilibrium strategies.

Definition. An information set is off the equilibrium path if it is not reached with positive probability when the game is played according to the equilibrium strategies.

3. At an information set on the equilibrium path, beliefs are determined by Bayes's Rule and the players' equilibrium strategies.

Suppose that $\sigma_1(L) = q_1$, $\sigma_1(M) = q_2$, and $\sigma_1(R) = 1 - q_1 - q_2$. Then by Bayes's Rule, $p = \frac{q_1}{q_1+q_2}$.

4. At information sets off the equilibrium path, beliefs are determined by Bayes's Rule and the players' equilibrium strategies where possible.

Definition. A **perfect Bayesian equilibrium** consists of behavior strategies and beliefs satisfying requirements 1–4.

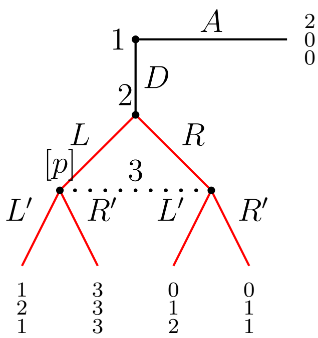

The Nash equilibrium of the red subgame is (L, R'). Thus (D, L, R') is a subgame perfect equilibrium. If we add belief $p = 1$, then (D, L, R') and $p = 1$ is a perfect Bayesian equilibrium.

Now consider (A, L, L'). This is a Nash equilibrium. If we add the belief $p = 0$, then (A, L, L') satisfies requirements 1–3 but violates 4.

Example

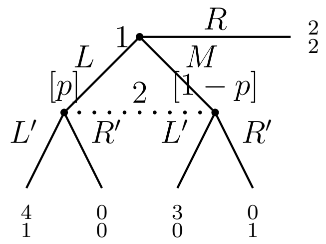

First we will look for a perfect Bayesian equilibrium in which $\sigma_1(R) = 1$.

If $\sigma_1(R) = 1$, then P is not pinned down by Bayes's Rule. For player 1 to choose $\sigma_1(R) = 1$ in a perfect Bayesian equilibrium, we require

$$\left.\begin{array}{l} 2 \geq \sigma_2(L') \cdot 4 \\ 2 \geq \sigma_2(L') \cdot 3 \end{array}\right\} \Rightarrow \sigma_2(L') \leq \frac{1}{2}.$$

We assign belief $p = 0$. Then $\sigma_2(L') = 0$, $\sigma_1(R) = 1$, $p = 0$ is a perfect Bayesian equilibrium. $\sigma_1(R) = 1$, $\sigma_2(L') = \frac{1}{2}$, $p = \frac{1}{2}$ is also a perfect Bayesian equilibrium.

Now we look for a perfect Bayesian equilibrium in which $\sigma_1(R) < 1$. In this case $\sigma_2(L') > 0$ since this is the only way player 1 can gain from choosing L or M. But if $\sigma_2(L') > 0$, then $\sigma_1(M) = 0$ since $4\sigma_2(L') > 3\sigma_2(L')$ (i.e. L is always better than M for player 1). Since $\sigma_1(R) < 1$ and $\sigma_1(M) = 0$, we have $\sigma_1(L) > 0$. But if $\sigma_1(L) > 0$ and $\sigma_1(M) = 0$, then by Bayes's Rule, $p = \frac{\sigma_1(L)}{\sigma_1(L)+\sigma_1(M)} = 1$, which means that $\sigma_2(L') =$ (sequential rationality). But then $\sigma_1(L) = 1$ (also by sequential rationality). Hence $\sigma_1(L) = 1$, $\sigma_2(L') = 1$, $p = 1$ is a perfect Bayesian equilibrium.

4.1 Signaling Games

Consider a dynamic game involving two players: the sender (S) and the receiver (R). Nature draws a type t_i for the sender from $T = \{t_1, \ldots, t_I\}$ according to a probability distribution $p : T \to [0, 1]$. The sender observes t_i and chooses a message m_j from $M = \{m_1, \ldots, m_J\}$. The receiver observes m_j (but not t_i) and then chooses an action a_k from $\{a_1, \ldots, a_K\}$. Payoffs are given by $u_S(t_i, m_j, a_k)$ and $u_R(t_i, m_j, a_k)$.

Consider a $2 \times 2 \times 2$ structure, with $T = \{t_1, t_2\}$, $M = \{m_1, m_2\}$, $A = \{a_1, a_2\}$, and $\text{Prob}[t_1] = p$.

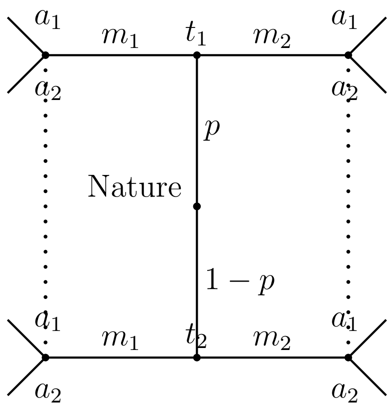

A strategy for the sender is a function $m(t_i) = m_j$. A strategy for the receiver is a function $a(m_j) = a_k$.

We now restate the conditions for perfect Bayesian equilibria in such a way that they apply to signaling games.

1. After observing any message $m_j \in M$, the receiver must have a belief about which types could have sent m_j. (The belief is a probability distribution $\mu(\cdot \mid m_j)$, where $\mu(t_i \mid m_j) \geq 0$ for each $t_i \in T$, and $\sum_{t_i \in T} \mu(t_i \mid m_j) = 1$.)

$2R$. For each $m_j \in M$, the receiver's action $a^*(m_j)$ must maximize her expected utility given the belief $\mu(\cdot \mid m_j)$, i.e. $a^*(m_j)$ solves

$$\max_{a_k \in A} \sum_{t_i \in T} \mu(t_i \mid m_j) u_R(t_i, m_j, a_k).$$

$2S$. For each $t_i \in T$, the sender's message $m^*(t_i)$ must maximize his expected utility given the receiver's strategy $a^*(m_j)$, i.e. $m^*(t_i)$ solves

$$\max_{m_j \in M} u_S(t_i, m_j, a^*(m_j)).$$

Given the sender's strategy $m^*(\cdot)$, let T_j denote the set of types that send the message m_j, i.e. $t_i \in T_j$ iff $m^*(t_i) = m_j$.

3. For each $m_j \in M$, if there exists $t_i \in T$ such that $m^*(t_i) = m_j$ (i.e. T_j is not empty), then the receiver's belief at the information set corresponding to m_j must follow from Bayes's Rule and the sender's strategy, i.e.

$$\mu(t_i \mid m_j) = \frac{p(t_i)}{\sum_{t_i \in T_j} p(t_i)}.$$

Example.

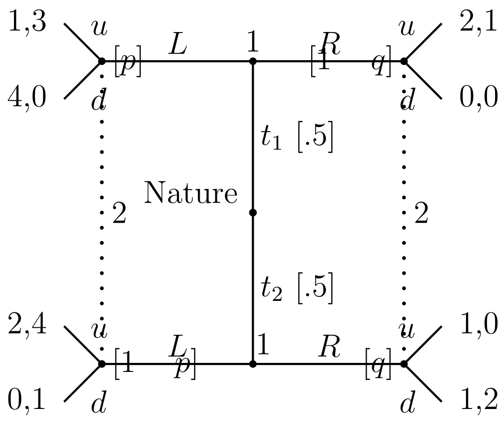

There are four possible types of equilibrium:

1. pooling on L,
2. pooling on R,
3. separating with t_1 playing L,
4. separating with t_1 playing R.

1. Suppose there is a perfect Bayesian equilibrium in which the sender's strategy is $(m(t_1), m(t_2)) = (L, L)$. Then $p = \mu(t_1 \mid L) = \frac{p(t_1)}{p(t_1)+p(t_2)} = .5$. The receiver then recognizes that u is better than d. Hence the sender of types t_1 and t_2 receives payoffs 1 and 2, respectively. We need the sender to prefer L regardless of his type; i.e. we need receiver's response to R to be d (or else the sender would deviate). Thus $q \leq \frac{2}{3}$. Thus $[(L, L), (u, d), p = .5, q = .5]$ is a pooling equilibrium.

4. If the sender plays $(m(t_1), m(t_2)) = (R, L)$, then the receiver must have beliefs $p = 0$ and $q = 1$. Hence the receiver's best response is $(a(L), a(R)) = (u, u)$, and in both cases the sender gets a payoff of 2. If a type 1 sender plays L, the receiver plays u, so sender's payoff is 1, so there is no incentive to deviate. If a type 2 sender plays R, the receiver plays u, and the sender's payoff would be 1, so there is no incentive to deviate. Thus $[(R, L), (u, u), p = 0, q = 1]$ is a separating perfect Bayesian equilibrium.

2. $a^*(R) = d$ and $a^*(L) = u$ regardless of p. So either type of sender would deviate. Thus there is no perfect Bayesian equilibrium with pooling on R.

3. $(m^*(t_1), m^*(t_2)) = (L, R)$, $p = 1, q = 0$. Then $a^*(L) = u$ and $a^*(R) = d$. Thus a t_1 sender would not deviate, but a t_2 sender would. Thus there is no perfect Bayesian equilibrium with this type.

4.2 Intuitive Criterion (Cho, Kreps)

Consider the Beer, Quiche game.

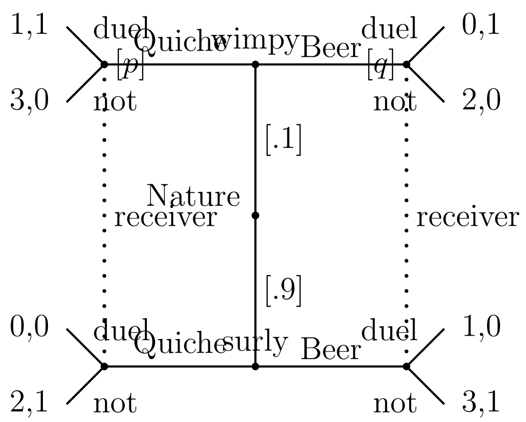

The pooling equilibria are

1. [(Quiche,Quiche), (not,duel), $p = .1, q \geq .5$].

2. [(Beer,Beer), (duel,not), $q = .1, p \geq .5$].

We argue that the first equilibrium is not sensible. The intuitive criterion says that a wimpy type can never gain by sending message Beer, but a surly type can. So if I see Beer, I must put probability 1 on type surly.

Definition. A message $m_j \neq m^*(t_i) \in M$ is **equilibrium-dominated** for type $t_i \in T$ if type t_i's equilibrium payoff $u_S(t_i, m^*(t_i), a^*(m_j))$ is greater than t_i's highest possible payoff from m_j:

$$u_S(t_i, m^*(t_i), a^*(m_j)) > \max_{a_k \in A} u_S(t_i, m_j, a_k).$$

Definition. The (Weak) Intuitive Criterion. If the information set following m_j is off the equilibrium path and m_j is equilibrium-dominated for type t_i, then (if possible) the receiver's belief $\mu(t_i \mid m_j)$ must be 0. This is possible so long as m_j is not equilibrium-dominated for all types $t_i \in T$.

4.3 Moral Hazard

One party may undertake actions that

1. affect the other party's value of the transaction,

2. the second party cannot perfectly monitor or enforce the action.

We will focus on a moral hazard relationship between an owner (principal) and a manager (agent), which is called the Principal-Agent problem.

Assume that the agent has an exogenously given reservation utility and that he accepts any contract that yields at least this utility. The optimal contract is the one that achieves the desired objective as a subgame perfect equilibrium of the following game.

Here is the model. Let $A = \{a_1, \ldots, a_N\}$ be the finite set of actions for the agent and $S = \{s_1, \ldots, s_M\}$ the finite set of signals. Each action produces a probability distribution over signals. Let $\Pi_{mn} = \text{Prob}[s_m \mid a_n]$. A **contract** is a vector $\{W(s_1), \ldots, W(s_M)\}$. Contracts are written in terms of signals, not actions. $u(w, a_n) = u(w) - d(a_n)$ is the utility of the agent, and μ_0 is the reservation utility of the agent. Both principal and agent are expected utility maximizers.

Assume

1. Every signal is possible following every action: $\Pi_{mn} > 0$ for all m, n.

2. Utility of wages, $u(\cdot)$, is strictly increasing and strictly concave: $u' > 0$, $u'' < 0$.

We will consider the case where signals are the principal's realized gross profit, so

$$B(a_n) = \sum_{m=1}^{M} \Pi_{mn} s_m$$

is the principal's gross expected profit. Let $x_m = u(W(s_m))$. Let $v = u^{-1}$, so $v(x_m) = W(s_m)$. The principal's problem is to solve

$$\max_{a_n, x_1, \ldots, x_m} B(a_n) = \sum_{m=1}^{M} \Pi_{mn} v(x_m)$$

such that

1. $$\sum_{m=1}^{M} \Pi_{mn} x_m - d(a_n) \geq \mu_0,$$

   the participation constraint,

2. $$\sum_{m=1}^{M} \Pi_{mn} x_m - d(a_n) \geq \sum_{m=1}^{M} \Pi_{mn} x_m - d(a_{n'})$$

   for all $a_{n'} \in A$, the incentive compatibility constraint.

We solve the problem in two steps.

1. For each $a_n \in A$, we find the cheapest way to induce the agent to take action a_n, i.e. for each $a_n \in A$, we solve

   $$C(a_n) = \min_{x_1, \ldots, x_m} \sum_{m=1}^{M} \Pi_{mn} v(x_m)$$

   such that (1) and (2) above are satisfied.

2. We solve for the action which maximizes the principal's expected net profit:

$$\max_{a_n \in A} B(a_n) - C(a_n).$$

Example. (Varian 25.4) Two actions and two signals. $A = \{a_1, a_2\}$, where a_1 and a_2 are effort levels, with $a_2 > a_1$, and $S = \{\text{success,failure}\}$. Let $\Pi_{sn} = \text{Prob}[\text{success}|\ a = a_n]$, $\Pi_{fn}[\text{failure}|\ a = a_n] = 1 - \Pi_{sn}$, $n = 1, 2$. Assume that $\Pi_{s2} > \Pi_{s1}$, and hence $\Pi_{f1} > \Pi_{f2}$. The implication is that $\frac{\Pi_{s2}}{\Pi_{s1}} > \frac{\Pi_{f2}}{\Pi_{f1}}$.

We solve

$$\min_{x_s, x_f} \Pi_{sn} v(x_s) + \Pi_{fn} v(x_f)$$

such that $\Pi_{sn} x_s + \Pi_{fn} x_f - a_n \geq \mu_0$ and $\Pi_{sn} x_s + \Pi_{fn} x_f - a_n \geq \Pi_{sn'} x_s + \Pi_{fn'} x_f - a_{n'}$ whenever $a_n \neq a_{n'}$. We may sketch the agent's indifference curves over *contracts.* All of the contracts which give the agent expected utility $\bar{u}$ when they take action a_1 are

$$\Pi_{s1} x_s + \Pi_{f1} x_f - a_1 = \bar{u},$$

$$x_f = \frac{\bar{u} + a_1}{\Pi_{f1}} - \frac{\Pi_{s1}}{\Pi_{f1}} x_s. \tag{1}$$

Similarly, given the action a_2, the indifference curve is

$$x_f = \frac{\bar{u} + a_2}{\Pi_{f2}} - \frac{\Pi_{s2}}{\Pi_{f2}} x_s. \tag{2}$$

We must also worry about the incentive constraint. We need

$$\Pi_{s1} x_s + \Pi_{f1} x_f - a_1 = \Pi_{s2} x_s + \Pi_{f2} x_f - a_2,$$

or

$$x_f = \frac{\Pi_{s2} - \Pi_{s1}}{\Pi_{f2} - \Pi_{f1}} x_s \cdot \frac{a_2 - a_1}{\Pi_{f1} - \Pi_{f2}}.$$

But

$$\Pi_{s2} - \Pi_{s1} = (1 - \Pi_{f2}) - (1 - \Pi_{f1}) = \Pi_{f1} - \Pi_{f2},$$

so

$$x_f = x_s - \frac{a_2 - a_1}{\Pi_{f1} - \Pi_{f2}}.$$

We now work out the isocost line for the principal. If the agent takes action a_1, then contracts that cost $\bar{c}$ are denoted by $\Pi_{s1}v(x_s) + \Pi_{f1}v(x_f) = \bar{c}$. Differentiate totally to get $\Pi_{s1}v'(x_s)\ dx_s + \Pi_{f1}v'(x_f)\ dx_f = 0$, or

$$\frac{dx_f}{dx_s} = -\frac{\Pi_{s1}v'(x_s)}{\Pi_{f1}v'(x_f)}.$$

Note that if $x_s = x_f$, then $\frac{dx_f}{dx_s} = -\frac{\Pi_{s1}}{\Pi_{f1}}$. This means action a_1 isocost curves are tangent to action a_1 isocost curves along the 45° line. The optimal contract for implementing x_f and action a_1 has $x_s = x_f$.

Example. Sam is employed by Clampett Oil Company. He gives recommendations on where to drill. There are two effort levels (where h denotes the number of hours worked). If $h = 400$, then the company makes \$250 000 with probability .7 and \$ -50 000 with probability .3. If $h = 0$, then the company makes \$250 000 with probability .5 and \$ -50 000 with probability .5. Sam's utility over wages and ours worked is $u(w, h) = \sqrt{w} - \frac{h}{200}$ (in thousands). Sam's alternative is to be a professional dart thrower, which gives utility 8.

1. Suppose Clampett Oil cannot monitor Sam's hours. Find the minimum cost of inducing $h = 400$ or $h = 0$.

   (a) Let us solve the problem of inducing $h = 400$. We must solve $\min_{x_s,x_f} .7x_s^2 + .3x_f^2$ such that
   $$.7x_s + .3x_f - 2 \geq 8 \tag{1}$$
   and
   $$.7x_s + .3x_f - 2 \geq .5x_s + .5x_f. \tag{2}$$
   From the graph, we know that the solution occurs where (1) and (2) intersect. We therefore have $.7x_s + .3x_f = 10$ and $.2x_s - .2x_f = 2$, so $x_2 = 13$ and $x_s = 13$ and $x_f = 3$, i.e. the minimum cost of implementing $h = 400$ is \$121 000.

   (b) Let us solve the problem of implementing $h = 0$. We have $.5x_s + .5x_f = 8$ and $x_s = x_f$, so $x_s = x_f = 8$. So the minimum cost is \$64 000.

2. Which contract should Clampett Oil give Sam? The profit for inducing $h = 400$ is $.7(250000) + .3(-50000) - 121000 = 39000$, whereas the profit for inducing $h = 0$ is $.5(250000) + .5(-50000) - 64000 = 36000$. So the contract would be to offer Sam w_s =169 000 in case of success and w_f =9000 in case of failure.

3. What contract would Clampett offer if they could observe h?

(a) The first-best solution with $h = 400$ would be obtained by solving $.7x_s + .3x_f - 2 = 8$ and $x_s = x_f$. Hence we have $x_s = x_f = 10$. Thus the cost is \$100 000. The profit is $.7(250000) + .3(-50000) - 100000 = 60000$.

(b) We solved the other case before. The profit is only \$36 000. Hence Clampett Oil should induce $h = 400$ by offering Sam \$100 000 in any case.

Chapter 5

Problems

1. There are two baseball teams that are preparing for a three-game series. Each team has three pitchers. Team 1 has an Ace, a mediocre pitcher, and a scrub. Team 2 has two mediocre pitchers and a scrub. A pitcher can only be used once in the series. The probabilities of winning the game depending on the pitcher match-ups are as follows:

Match up	Outcome
Ace versus mediocre player	Ace wins with probability .8
Ace versus scrub	Ace wins with probability .9
Mediocre pitcher versus scrub	Mediocre pitcher wins with probability .6
Same versus same	Each wins with probability .5

   Assume a win is worth 1 and a loss is worth 0. Suppose the three games are played sequentially. The pitchers for each game are chosen simultaneously, but after a game is played it is common knowledge what pitchers were used. Remember each team can only use each pitcher once. Write out the extensive form of the game. (The extensive form has 18 terminal nodes.)

2. One day, long ago, two women and a baby were brought before King Salmon. Each claimed to be the true mother of the baby. King Salmon pondered dividing the baby in half and giving an equal share to each mother, but decided that in the interest of the baby he would use game theory to resolve the dispute instead. Here is what the king knew. One of the women is the true mother but he does not know which one. It is common knowledge among the two women who the true mother is. The true mother places a value of 10 on having the baby and 0 on not having it. The woman who is not the true mother places a value of 5 on having the baby and 0 on not having it. The king proposed the following game

to determine who would get to keep the baby. One woman is chosen to go first. She is player 1, and the other woman is player 2. She must announce "mine" if she wants to claim to be the real mother and "hers" if she wants to claim that the other woman (player 2) is the true mother. If she says "hers" the baby goes to player 2 and the game ends. If she says "mine," player two gets to move. Player 2 announces "MINE" or "HERS." If player 2 says "MINE," she gets the baby and pays King Salmon an amount v. Meanwhile, player 1 also has to pay King Salmon 1. If player 2 says "HERS," Player 1 gets the baby at no cost and player 2 makes no payments to King Salmon.

(a) Draw the extensive form for each game:

i. The true mother is player 1.

ii. The true mother is player 2.

(b) Assume that the mothers play the (unique) subgame perfect equilibrium in each game. Find a value v such that the true mother receives the baby regardless of which mother is player 1.

(c) Draw the normal form representation of each of the extensive form (using your value of v from part (b).) Are there any Nash equilibria in which the true mother does not receive the baby? If so, why does this happen under Nash equilibria but not under subgame perfect equilibria?

3. Two firms, A and B, produce a homogeneous good that they sell in two markets, 1 and 2. Each firm has the cost function

$$C(q_1^j + q_2^j) = \frac{(q_1^j + q_2^j)^2}{2},$$

where q_1^j and q_2^j are the quantities firm $j = A, B$ sells in market 1 and market 2, respectively. The inverse demand function is the same in each market and is given by $P(Q_i) = 20 - Q_i$, where $Q_i = q_i^A + q_i^B$ for $i = 1, 2$.

(a) Write firm A's profit (from sales in both markets) as a function of $q_1^A, q_2^A, q_1^B, q_2^B$.

(b) Consider the case where the firms sell the good in both markets simultaneously. Write down the strategy set for firm j. Find the Nash equilibrium.

(c) Consider the case where the two markets are open sequentially. Both firms sell output in market 1 and then, after seeing the sales in market 1, both firms sell output in market 2. Define a strategy for firm j. Find the unique subgame perfect equilibrium. Provide intuition as to why in subgame perfect equilibrium both firms do better in this case than in the previous case.

item In professional baseball there is an interesting arbitration process. If the team and a player cannot agree on the terms of the player's contract then each side submits a salary to the arbitrator. The arbitrator chooses one salary or the other (not a compromise), and that salary becomes binding. Suppose it is common knowledge that the arbitrator has in mind a salary of \$$V$. Draw the reaction correspondences for the team and the player and identify the Nash equilibrium.

4. Consider an industry in which there are two firms, a manufacturer and a retailer. The manufacturer moves first, choosing a wholesale price p_W for its product. The (monopolistic) retailer observes p_W and then decides what quantity q to buy from the manufacturer to then resell to the public. If the retailer sells q units of output, then the retail price is $p_R = A - q$, where $A > 0$. Assume for simplicity that manufacturing costs are zero and the retailer's only costs are the payments to the manufacturer. Hence, if the wholesale price is p_W and the retailer chooses to buy and then sell q units, then the profits of the manufacturer and retailer are, respectively,

$$\Pi_M(p_W, q) = p_W q$$

and

$$\Pi_R(p_W, q) = (A - q)q - p_W q.$$

   (a) Find the subgame perfect equilibrium of this game.

   (b) In the subgame perfect equilibrium what is each firm's profit?

   (c) Suppose that the two firms merge to become a single firm that maximizes joint profits: $\Pi_M(p_W, q) + \Pi_R(p_W, q)$. What happens to the retail price as a result of the merger? How does the profit of the merged firm compare to the sum of the profits in part (b)?

5. Consider the following strategic game.

	Rock	Paper	Scissors
Rock	0,0	$-1, 1$	$1, -1$
Paper	$1, -1$	0,0	$-1, 1$
Scissors	$-1, 1$	$1, -1$	0,0

   (a) Can the game be solved by iterated elimination of strictly dominated strategies? If so, what is the solution?

   (b) Is there a (non-degenerate) mixed strategy Nash equilibrium for the same? If so, what is it?

6. Players 1 and 2 are bargaining over how to split one dollar. Both players simultaneously name shares they would like to have, s_1 and s_2, where $0 \leq s_1, s_2 \leq 1$. If $s_1 + s_2 \leq 1$, then the dollar is divided with each player receiving the share named. If $s_1 + s_2 > 1$, then each player receives \$0. What are the pure-strategy Nash equilibria of this game? Is the dollar divided in every Nash equilibrium?

7. For the extensive form game below (which has only one player), consider the mixed strategy $\sigma_1 = \left(\frac{1}{3}\ell_1\ell_2\ell_3, \frac{1}{18}r_1\ell_2\ell_3, \frac{1}{6}r_1r_2\ell_3, \frac{4}{9}r_1r_2r_3\right)$, where $r_1\ell_2\ell_3$, for example, denotes the pure strategy in which player 1 chooses r_1 at his first decision node, ℓ_2 at his second decision node, and ℓ_3 at his third decision node.

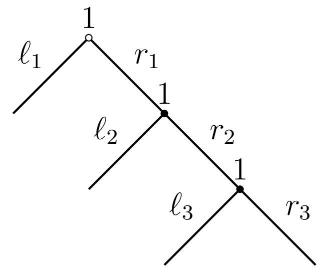

What is player 1's equivalent behavioral strategy?

8. Consider the following game in extensive form.

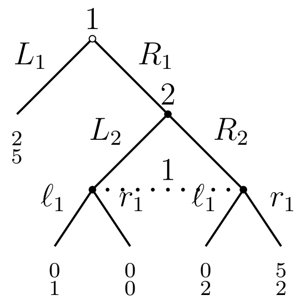

(a) Compute a subgame perfect equilibrium.

(b) If possible, find a Nash equilibrium of the game that is not a subgame perfect equilibrium.

9. Consider the following two-player game, Γ. Each player can either work hard (h) or shirk (s). If both work hard they each get a payoff of 2. If both shirk they each get a payoff of 0. If one shirks and the other works hard, the person who works hard gets a payoff of -1 and the person who shirks gets a payoff of 0. The game Γ is preceded by another simultaneous game in which the two players announce whether they wish to "Play" or "Not Play" the game Γ. If they both announce "Play" then they play the game Γ. Otherwise, each gets a payoff of 1, i.e., strategy profiles (Play, Not Play), (Not Play, Play), and (Not Play, Not Play) give each player a payoff of 1.

   (a) Compute all the Nash equilibria of the subgame Γ.

   (b) Is there a subgame-perfect equilibrium in which both players choose shirk with positive probability in the subgame Γ?

10. Consider the following static Bayesian Game:

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

   (2) Player 1 (the row player) learns whether nature has drawn Game 1 or Game 2, but player 2 does not.

   (3) Player 1 chooses either T or B and Player 2 simultaneously chooses either L or R.

	L	R
T	10,10	0,0
B	0,0	4,4

	L	R
T	4,4	0,0
B	0,0	8,8

   (a) Write out the extensive-form version of this game.

   (b) Write down the normal-form version of the extensive form in part (a). Find all the pure strategy Bayesian equilibria.

11. Consider the two-player Bayesian game in which $S_1 = \{T, B\}$ and $S_2 = \{L, R\}$, each player has two types $\Theta_1 = \Theta_2 = \{0, 1\}$, and each type profile is equally likely, i.e. $\rho(0,0) = \rho(0,1) = \rho(1,0) = \rho(1,1) = 1/4$. Payoffs are given below:

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

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

	L	R
T	2,4	0,2
B	0,1	1,0

	L	R
T	2,2	0,0
B	4,0	1,2

The type profile determines which of the four matrices maps actions to payoffs. If, for example, $t_1 = 0$ and $t_2 = 1$, then the upper right matrix shows how player payoffs depend on actions. In that case, if player 1 chooses T and player 2 chooses R then their payoffs are 1 and 0, respectively.

(a) Show that $(s_2(0), s_2(1)) = (R, L)$ is **NOT** part of a Bayesian equilibrium.

(b) Find a pure strategy Bayesian equilibrium.

12. Two individuals each have \$1 deposited in a bank and must simultaneously decide whether to withdraw their \$1 or leave it deposited in the bank. The bank has invested all money in a risky, illiquid asset. If one or both players attempt to withdraw their \$1, the bank will have to liquidate the asset for one-half the initial value and the proceeds will be divided among the claimants. Since there are two players that means \$2 would be liquidated for \$1, which would be divided 1 way (if one claimant) or 2 ways (if two claimants). If neither player withdraws their \$1, the money will remain invested and each player will either gain 20 cents or lose 20 cents depending on the success or failure of the bank's risky investment. Success occurs with probability p and failure occurs with probability $1 - p$. The game is described below. Player 1 chooses either withdraw or not withdraw (w or n, respectively) and player 2 simultaneously chooses either withdraw or not withdraw (W or N, respectively.)

	W	N
w	.5,.5	1,0
n	0,1	1.2,1.2

	W	N
w	.5,.5	1,0
n	0,1	.8,.8

(a) Describe the set of pure-strategy Nash equilibria for each value of p in the interval $[0, 1]$. Identify which equilibria are efficient.

(b) Suppose now that player 1 is informed about the outcome of the bank's investment before she decides whether to withdraw or not withdraw. Player 2 still only knows the probability of success or failure. These facts are common knowledge. Suppose the game is still played simultaneously. Describe the set of pure-strategy Bayesian equilibria for each value of p in ther interval $[0, 1]$. Briefly comment on how private information affects the outcome of this game.

13. Consider a first-price, sealed-bid auction in which the bidder's valuations are independently and uniformly distributed on $[0,1]$. Show that if there are n bidders, then the strategy of bidding $(n-1)/n$ times one's valuation is a symmetric Bayesian equilibrium of the auction.

14. Consider the following static Bayesian game:

    (i) Nature determines whether the payoffs are as in game 1 or as in game 2, each game being equally likely.

    (ii) Player 1 (the row player) learns whether nature has drawn game 1 or game 2, but player 2 does not.

    (iii) Player 1 chooses either T or B, player 2 observes player 1's choice, and then chooses either L or R.

	L	R
T	0,0	10,4
B	6,6	0,2

	L	R
T	2,4	0,4
B	6,0	2,6

Find all the pure-strategy pooling and separating perfect Bayesian equilibria in the resulting signaling game. Be sure to specify the equilibrium strategies and beliefs in each case.

15. An evil dictator might have weapons of mass destruction (WMD). A neighboring country regards each possibility as being equally likely and must decide whether or not to invade. You must analyze the signaling game in which the evil dictator makes an announcement about his weapons status and the neighboring country decides whether or not to invade. Find all the pure-strategy Bayesian equilibria. Are there any separating equilibria? Are there any pooling equilibria?

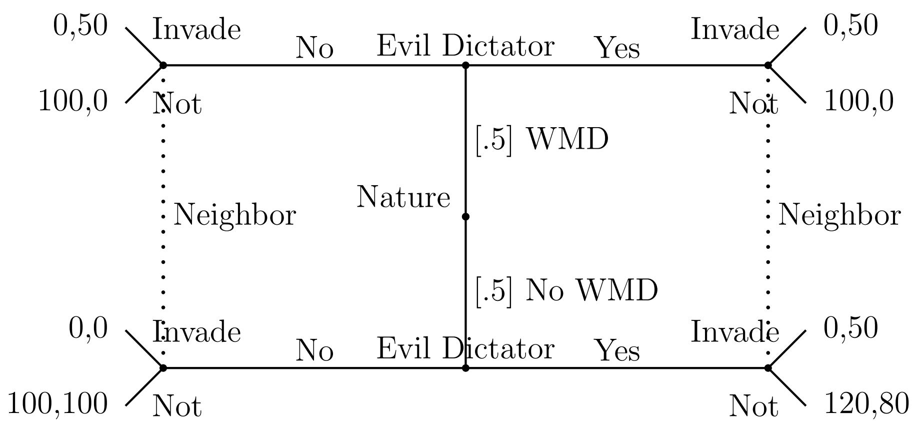

16. Once there was an evil lord that visited a poor serf each fall and demanded half of his harvest. Both the lord and the serf knew that a good harvest produced 10 bushels of potatoes and that a bad harvest produced only 6. In addition, it was common knowledge that the probability of a good harvest was 0.6. The serf knew which type of harvest had occurred. The lord was told the type of harvest by the serf and had to decide whether or not to believe him. If he believed the serf he took the potatoes he was offered and left. If he didn't believe the serf, he took what the serf offered and punished him for lying by burning down his house.

    The serf kept potatoes in the cellar of his house. Hence, the truth would be revealed once the house had burnt down and his potato holdings were exposed. Punishing a dishonest serf gave the lord an added boost of pleasure equal to $c > 0$. On the other hand, if the fire revealed that the serf had been telling the truth, the lord's happiness was diminished by c. While the lord felt bad when he was wrong, he would not compensate the serf by giving the serf some of his crop back. Also, neither the serf nor the lord liked baked potatoes, so any potatoes that were hidden in the cellar were worthless if the house was burned. The payoffs of this signaling game are shown below.

    (a) Find all the pure-strategy perfect Bayesian equilibria. Are there any separating equilibria? Are there any pooling equilibria?

    (b) Do any of the pooling equilibria fail the (weak) intuitive criterion?

    (c) Suppose the lord can commit to a strategy in advance. Namely, the lord can announce in advance what he will do in response to each message by the serf, and the announcement is credible. What strategy for the lord maximizes his expected payoff? Set $c = 1$.

17. Two identical firms each simultaneously choose a nonnegative quantity of output. There are no costs. The payoff to firm i as a function of the outputs is $\Pi_i(q_i, q_{-i}) = (100 - q_i - q_{-i})q_{-i}$. Show that for each firm, any level of output greater than 50 is strictly dominated.

18. On the hit TV show Survivor Thailand two tribes competed in a game they called Thai 21. The game starts with 21 flags in a circle. The teams take turns removing flags from the circle until they are all gone. The team that removes the last flag wins immunity. At each turn teams must remove 1, 2, or 3 flags. Team 1 moves first. How many flags should Team 1 take on its first move?

19. Consider the two-player Bayesian game in which $S_1 = \{T, B\}$ and $S_2 = \{L, R\}$, each player has two types $\Theta_1 = \Theta_2 = \{0, 1\}$, and each type profile is equally likely, i.e. $p(0, 0) = p(0, 1) = p(1, 0) = p(1, 1) = \frac{1}{4}$. Payoffs are given below:

		$t_2 = 0$			$t_2 = 1$	
		L	R		L	R
$t_1 = 0$	T	0,0	1,2	T	0,1	1,0
	B	2,1	0,0	B	2,0	0,2
		L	R		L	R
$t_1 = 1$	T	2,0	0,2	T	2,1	0,0
	B	0,1	1,0	B	0,0	1,2

The type profile determines which of the four matrices maps actions to payoffs. If, for example, $t_1 = 0$ and $t_2 = 1$, then the upper right matrix shows how player payoffs depend on actions. In that case, if player 1 chooses T and player 2 chooses R then their payoffs are 1 and 0, respectively.

(a) Is $(s_1(0), s_1(1)) = (T, B)$ part of a Bayesian equilibrium?

(b) Is $(s_2(0), s_2(1)) = (L, L)$ part of a Bayesian equilibrium?

20. Find all the pure-strategy perfect Bayesian equilibria of the following game in extensive form.

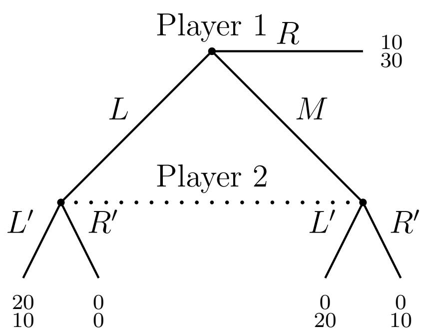

21. Once upon a time a signaling game was played between a princess and a frog. The frog was the "Sender." He could either say he was a "prince" or a "frog." The princess was the receiver. She could either kiss the frog, in which case he might turn into a prince. Or, she could eat him. It was well known that 10% of the frogs in the kingdom would turn into princes when kissed by a princess. For reasons that arise from commonly known facts about frogs and princesses, the payoffs of the signaling game are as presented below.

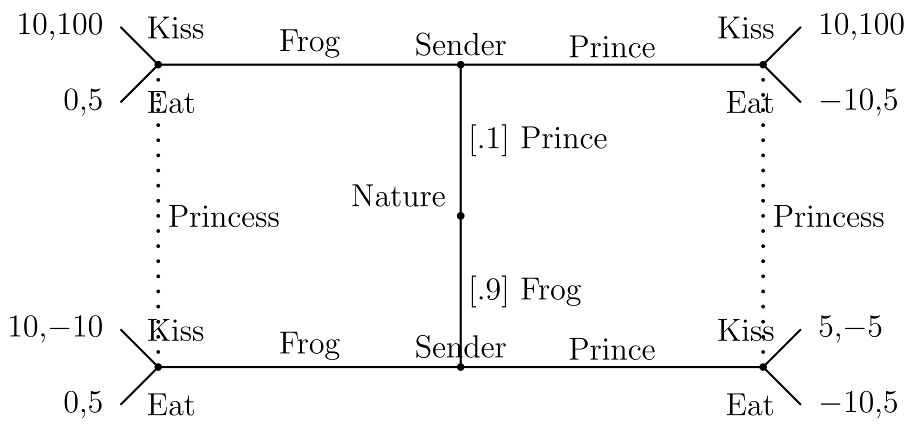

(a) Find all the pure-strategy perfect Bayesian equilibria. Are there any separating equilibria? Are there any pooling equilibria?

(b) Do any of the pooling equilibria fail the (weak) intuitive criterion?

22. An agent has effort levels $a_1 =$ and $a_2 = 3$. The probability distribution over sales for each level of effort is given below.

	Sales	Sales
Effort	s_1 =\$100	s_2 =\$150
$a_1 = 1$	.8	.2
$a_2 = 3$	.4	.6

The expected profit for the principal when the agent takes action a_i is expected sales conditional on a_i minus expected costs conditional on a_i. The agent has utility over wages and actions given by $U(w, a_i) = w^{1/2} - a_i$ and a reservation utility of 1. Contracts can be written in terms of sales, but not actions.

(a) Find the minimum cost of implementing each action.

(b) Which action does the principal implement to maximize profit?