Lecture Notes

CGDE: Game Theory

September, 2019

Last updated: March 12, 2017

Zvika Neeman

1 Brief Introduction

- This is an introduction to *non-cooperative game theory*. There is also *cooperative game theory*, but the names are misleading. It is not that non-cooperative game theory is devoted to non-cooperation while cooperative game theory is devoted to cooperation. Rather, non-cooperative theory is strategic and analyzes strategic situations from the perspective of players, while cooperative game theory is *non-strategic* and analyzes situations from the perspective of possible coalitions of players.

- Strategic/Normal Form (one shot) games vs. Extensive Form (sequential) games:

  - "Matrices" vs. "Trees."

- Game theory is not really a ***theory*** but rather a ***language*** with specific vocabulary (e.g., game, player, strategy, beliefs, best response, equilibrium, etc.) and conventions (about how game should be described, equilibrium, etc.) that allow us to describe and analyze strategic situations. Almost any situation can be thought of as a strategic situation, hence the title of the course's textbook by Dixit & Nalebuff and "The Art of Strategy: A Game Theorist's Guide to Success in Business and Life."

- Show the "Battle of Wits" from *The Princess Bride* (Dir. Rob Reiner, 1987; based on the 1973 novel by William Goldman).

2 Strategic Form Games

2.1 Single-person decision making

One player facing uncertainty about Nature. Say the player chooses an element $a \in A$ and Nature makes a choice, $s \in S$.

Expected utility: Known probabilities $p \in \Delta(S)$ (where for any set X we denote the set of probability distributions over X by $\Delta(X)$) and a utility function $u : A \times S \to \mathbf{R}$. Player chooses $a \in A$ to solve $\max_{a \in A} \sum_{s \in S} p_s u(a, s)$. Write this for simplicity as $u(a, p)$

Subjective expected utility: p is not given, but subjective, in the mind of the decision maker.

$A \backslash S$	s	s'
a^1	5	1
a^2	1	5
a^3	4	0
a^4	2.2	2.2
a^5	2	4.5

If $p_s = 1$ choose a^1; $BR((1, 0)) = \{a^1\}$.

If $p_s = 0.3$, $p_{s'} = 0.7$ then the best reply is a^2; $BR((0.3, 0.7)) = a^2$.

If $p_s = 1/2$ choose a^5; $BR((0.5, 0.5)) = \{a^5\}$.

If $p_s = 1/3$ then choose either a^5 or a^2: both yield expected utility of $3\frac{2}{3}$; $BR((1/3, 2/3)) = \{a^2, a^5\}$.

No rational decision maker would choose a^3. It is never a best reply, regardless of beliefs. Formally: for any belief p, $a^3 \notin BR(p)$. That is, given any p there is a better $a \in A$ s.t. $u(a, p) > u(a^3, p)$. In particular for all p we can find a single strategy that is better, a^1. So we can also say that a^3 is dominated by a^1: $\forall s \in S$, $u(a^1, s) > u(a^3, s)$.

What about a^4? It is not dominated by any of the actions in A. But it is not a best reply. It is easiest to see this in a picture. On the horizontal axis we draw the probability of s', on the vertical the expected utility. (The decreasing line denotes the payoffs from a^1, the increasing is a^2, the flat is a^4, and the red is a^5).

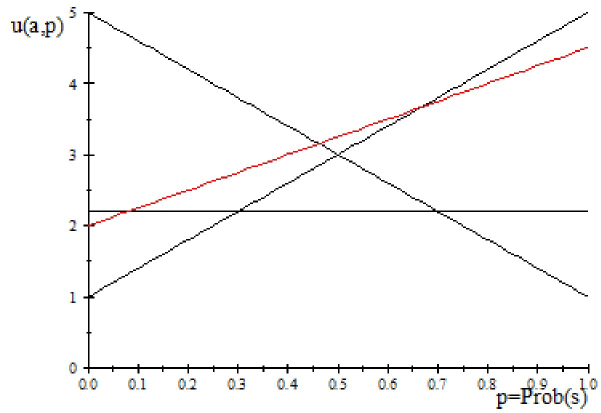

But a^4 is dominated–it is dominated by the mixed strategy of playing a^1 and a^2 with equal probabilities (*explain the sense in which being dominated is stronger than not being a best reply*). That mixed strategy gives expected payoffs of 3 regardless of what one believes about the opponent. Mixed strategies are written as $\alpha \in \Delta(A)$, where $\alpha(a)$ is the probability with which α chooses any $a \in A$, and the expected utility of a mixed strategy given state s is simply $\sum_a \alpha(a) u(a,s)$ and is also written as $u(\alpha, s)$. We also write $u(\alpha, p)$ for the expected utility of mixed strategy α when holding belief p over Nature.

Definition 1 *A strategy $\bar{\alpha} \in \Delta(A)$ is a best reply given belief $p \in \Delta(S)$, written $\bar{\alpha} \in \mathrm{BR}(p)$, if $u(\bar{\alpha}, p) \geq u(\hat{a}, p)$ for all $\hat{a} \in A$.*

The following definition follows immediately from Definition **??**.

Definition 2 *A strategy $\bar{\alpha}$ is never a best reply if for all beliefs $p \in \Delta(S)$ about Nature's (the opponent's) behavior there exists another pure strategy $\hat{a}$ (that may depend on the belief p) s.t. $u(\hat{a}, p) > u(\bar{\alpha}, p)$.*

Exercise 1 *(Not to be handed in.) Note that in Definition* **??** *the inequality is weak and in Definition* **??** *it is strict. Make sure you understand why this is the "right" way to define these notions.*

Exercise 2 *Explain why it makes no difference if $\hat{a} \in A$ is replaced by $\hat{\alpha} \in \Delta(A)$ in the two definitions above.*

Definition 3 *A strategy $\bar{\alpha} \in \Delta(A)$ is dominated if there is another (possibly mixed) strategy $\alpha \in \Delta(A)$ s.t. for all actions s of Nature (the opponent) $u(\alpha, s) > u(\bar{\alpha}, s)$. (We also say that $\bar{\alpha}$ is dominated by α.)*

The following theorem states that what we saw above in the example is general.

Theorem 1 *If the strategy set is finite (or if it is compact and the utility function is continuous), then a strategy is never a best reply if and only if it is dominated.*

Proof. That a strategy that is dominated is never a best reply is obvious. It remains to be shown that if a pure action is never a best reply then it is dominated. We prove this by arguing that if it is not dominated then it is a best reply to some belief.

So let a be a pure action that is not dominated. We want to show that it is a best response to some belief $p \in \Delta(S)$.

Notation: For any finite set X, the set $\Delta(X)$ is the set of probability measures over X, that is $\Delta(X) = \left\{p \in \mathbf{R}_+^{|X|} : \sum_{i=1}^{|X|} p_i = 1\right\}$, where $|X|$ is the number of elements (cardinality) of X. If X is an infinite subset of $\mathbf{R}$ then $\Delta(X)$ is the set of probability distributions over X in the usual sense. (For the technically minded we need to specify the σ–field over X, I will not delve into such issues.)

Let $|S| = K$. For any $\alpha \in \Delta(A)$, define the vector of payoffs

$$U_\alpha = (u(\alpha, s_1), u(\alpha, s_2) \ldots, u(\alpha, s_K))$$

Let V be the collection of these vectors,

$$V = \{(u(\alpha, s_1), u(\alpha, s_2) \ldots, u(\alpha, s_K)) : \alpha \in \Delta(A)\}.$$

Note that V is a convex set.

Define the collection of vectors U as follows,

$$U = \{U_a\} - V = \{u \in \mathbf{R}^K : u = U_a - v \text{ for some } v \in V\}.$$

Note that $0 = U_a - U_a$ belongs to U and that U is a convex set.

Define the strictly negative orthant of $\mathbf{R}^K$, denoted $\mathbf{R}_{--}^K$ and the weakly positive orthant, $\mathbf{R}_+^K$ as follows:

$$\begin{aligned}
\mathbf{R}_{--}^K &= \left\{v \in \mathbf{R}^K : v_k < 0 \text{ for all } k = 1 \ldots, K\right\} \\
\mathbf{R}_{+}^K &= \left\{v \in \mathbf{R}^K : v_k \geq 0 \text{ for all } k = 1 \ldots, K\right\}
\end{aligned}$$

Since a is not dominated, it must be that for every $\alpha \in \Delta(A)$ there exists some state $s_k \in S$ such that $u(a, s_k) - u(\alpha, s_k) \geq 0$. This means that every vector in U has at least one non-negative coordinate, i.e.

$$\mathbf{R}_{--}^K \cap U = \emptyset.$$

Since the negative orthant is also a convex set, we can apply the separating hyperplane theorem. Hence, there exists a vector $p \in \mathbf{R}^K$ s.t. $p \neq 0$ and a real number c such that

$$u \cdot p \geq c, \forall u \in U \text{ and } v \cdot p \leq c, \ \forall v \in \mathbf{R}^K_{--}.$$

In fact, it must be that $c = 0$. This is because $0 \in U$ so that $u \cdot p \geq c$ for all $u \in U \Rightarrow 0 \geq c$ and for every $\varepsilon > 0$, the vector $v^\varepsilon = (-\varepsilon, -\varepsilon, \ldots - \varepsilon) \in \mathbf{R}^K_{--}$ and $p \cdot v^\varepsilon = -\varepsilon \sum_k (p)_k$ so $v^\varepsilon \cdot p \leq c \Rightarrow -\varepsilon \sum_k (p)_k \leq c$ for all ε so $c \geq 0$.

Now since $v \cdot p \leq 0$ for all $v \in \mathbf{R}^K_{--}$, it must be that $p \in \mathbf{R}^K_+$, i.e., every coordinate of p is non-negative (otherwise we could find a vector $v \in \mathbf{R}^K_{--}$ such that $v \cdot p > 0$). And since $p \neq 0$, we can normalize p by

$$\hat{p} = \frac{1}{\sum_{k=1}^K p_k} p \in \mathbf{\Delta}(K)$$

since then $\sum_{k=1}^K \hat{p}_k = 1$. Note that the normalization preserves the inequality: $u \cdot \hat{p} \geq 0, \forall u \in U$. This inequality implies that

$$\hat{p} \cdot U_a \geq \hat{p} \cdot U_\alpha$$

for all $\alpha \in \Delta(S)$. And this is equivalent to

$$u(a, \hat{p}) \geq u(\alpha, \hat{p})$$

for all $\alpha \in \Delta(S)$. In other words a is a best-reply to $\hat{p}$. ∎

Remark 1 *The above proof is above the typical level of proofs of abstract results that I will expect for this course in terms of using formal results like the separating hyperplane theorem. On the other hand, the level of rigor and formal analysis required for specific games and examples that we* will *study will be similar.*

Exercise 3 *Recall the two definitions we gave, one for a strategy $\bar{\alpha}$ being dominated,* **??***, and one for a strategy $\bar{\alpha}$ never being a best reply,* **??***. In each definition consider making the following change: wherever it says $s \in S$ replace that with $p \in \Delta(S)$, and conversely. Would we get an equivalent, stronger, or weaker definition? Similarly, what would happen if wherever the definition says $a \in A$ we replaced that with $\alpha \in \Delta(A)$, and conversely. And what about both replacements?*

To be specific, one variant of the never-best-reply definition would be as follows: $\bar{\alpha} \in \Delta(A)$ is never best reply if for all $p \in \Delta(S)$ there exists $\hat{\alpha} \in A$ s.t. $u(\hat{\alpha}, p) \geq u(\bar{\alpha}, p)$. What is the relationship between this definition and the one in Definition **??***? (That is, if a strategy* *is a best reply according to Definition* **??***, does that imply it is a best reply according to the alternative definition in this exercise? What about the converse?)*

A variant of the dominance definition would be: A strategy $\bar{\alpha}$ *is dominated if there is another pure strategy* a^* *s.t. for all actions* s *of Nature (the opponent)* $u(a^*, s) > u(\bar{a}, s)$*.*

The question asks you to consider three variants on each definition, and discuss the relationships among them.

2.2 Strategic-form games: Definition

What is a game ? A (complete information) game in *normal* or *strategic form* is described as

$$G = (N, A, u)$$

where $N = \{1, ..., n\}$ describes the set of players; $A = A_1 \times \cdots \times A_N$ where for every $i \in N$ $A_i = \left\{a_1^i, ..., a_{m_i}^i\right\}$ describes the set of actions available to player i; and $u = (u_1, ..., u_n)$ where $u_i : A_1 \times \cdots \times A_n \to \mathbb{R}$ describes the payoff (utility) function for each player.

How is a Game Played ? It is assumed that players are expected utility maximizers, and that the players choose their strategies (strategies are mixtures, or lotteries, over actions; a degenerate lottery that is concentrated on one action is called a *pure strategy*) simultaneously, or at least without knowing the choices of the other players.

It is natural to describe strategic form games by game matrices. In 2 player games, the row player is usually referred to as Player 1 and the column player as Player 2; lines and columns correspond to the players' actions, respectively; payoffs are listed in pairs, with the payoff to Player 1 to the left.

Example.

	L	R
U	$3, 1$	$1, 2$
D	$3, 4$	$5, 7$

It is difficult to appreciate this at this stage, but strategic form games can represent a huge range of strategic situations (indeed, as we shall see later during the course, *all* possible strategic situations).

OK, but how *are* games played?

- Dominant strategies

- Iteratively deleting dominated strategies: thinking about the other thinking about ...
- Nash Equilibrium
- Refinements
- Other ? (other sensible possibilities ...)

Example 1 *Consider the following game of N players. $A_i = \{1, ..., M\}$. The player who guesses closest to 2/3 of the average of the other players gets a prize of* 100*. If more than one is closest the winner is chosen randomly.*

- *In this as well as in other games, a "good player" needs to be smart in two different senses: (1) computational ability, and (2) knowing what to expect or having the "correct beliefs".*

Mention the theory of k-level rationality.

2.3 Dominant Strategies

Definition 4 *A strategy is said to be dominant if it ensures a strictly higher payoff than any other strategy of the player, regardless of the strategies chosen by other players.*

Obviously, a player who has a dominant strategy would play it.

Example 2 ***The Prisoners' Dilemma***

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

The players choose their actions simultaneously (or independently). We expect that rational behavior (defined as optimization relative to some "reasonable beliefs") will produce a Pareto inefficient outcome. [Explain why this game is called the prisoners' dilemma]

[Show Dilbert video from Prisoners' Dilemma folder]

In laboratory experiments people do cooperate in the Prisoners' Dilemma and moreover rates of cooperation increases significantly if players are allowed to communicate with each other before playing the game (even though such communication should not make any difference strategically). The most likely explanation for this cooperation is that people's real payoffs in the game are different from the monetary payments induced by the experimenter, so that they don't really play the prisoners' dilemma game but some other game.

[Show Split-or-Steal videos. Explain the sense in which it is different from a standard PD. Show two examples, and the "weirdest ever" example]

Another important fact about the Prisoners' Dilemma is that it provides the quintessential counter-example to the First Welfare Theorem in economics that asserts that under certain conditions self interested behavior will produce Pareto efficient outcomes.

[Show Nash Equilibrium is inefficient from a Beautiful Mind]

Exercise 4 *What "real life scenarios" have the same structure as the Prisoners' Dilemma? (Contribution to a public good[1]; armament races; others?)*

[Show Dark Knight video from Prisoners' Dilemma folder]

Example 3 ***A variant on the Hotelling-Downs Model of Electoral Competition***[2]

Two political parties, a left wing and a right wing party, compete in electoral elections. Thus, $N = \{L, R\}$*. Each party chooses a policy, which is represented by a number between* -1 *and* 1*. The right wing party is constrained to choose a policy from the interval* $A_R = [0, 1]$*, and the left wing party is constrained to choose a policy from the interval* $A_L = [-1, 0]$*.*

The voter's "ideal points" are uniformly distributed over the interval $[-1, 1]$*. Given the two parties policies, each voter votes for the party that is closest to its ideal point. Voters who are indifferent randomize.*

Suppose that the objective of each party is to maximize the number (measure) of its voters. What policies will the parties choose? Show that the policy 0 *dominates all other policies for both parties. [Mention Black's Theorem (50s) about policy convergence to the median voter's preferred policy].*

2.4 Successive/Iterative Elimination of Strictly Dominated Strategies

Example 4 ***The Prisoners' Dilemma with a possibility to escape***

Suppose that in addition to cooperate and defect, the players also have the option to try and run away. Suppose that payoffs are given by the following matrix (R is for run):

[1] Players may either contribute to a public good or not (e.g., pollute or not); the quantity of the public good produced is given by $a\cdot$(sum of contributions), $\frac{1}{2} < a < 1$; payoffs are given by the quantity of the public good produced minus the player's own contribution.

[2] The example is a simplified version of Hotelling's model, which is a standard and much used model in "formal political science."

	C	*D*	*R*
C	*3,3*	*0,4*	*3,-1*
D	*4,0*	*1,1*	*1,-1*
R	*-1,3*	*-1,1*	*-1,-1*

We assume that escapees are killed during the attempt. If the other player keeps silent he gets off the hook, but if he snitches the deal is off and he is punished as if both had snitched.

Observe that in this new game, D no longer dominates C. But since R is dominated, it is obvious it will not be played by any rational player. And after R is eliminated, it is possible to also eliminate C.

This example demonstrates the process known as the **Iterative Elimination of Strictly Dominated Strategies.**

Definition 5 *A strategy $\bar{\alpha} \in \Delta(A)$ is dominated if there is another (possibly mixed) strategy $\alpha \in \Delta(A)$ s.t. for all actions s of the opponent (Nature) $u(\alpha, s) > u(\bar{\alpha}, s)$. (We also say that $\bar{\alpha}$ is dominated by α.)*

As implied by the next claim, the iterative elimination of strictly dominated strategies always gives rise to the same outcome.

Proposition 1 *The order of the elimination of strictly dominated strategies is of no consequence.*

Proof. This proof considers finite games, but the result is true also for compact continuous games as well. The proof relies on a basic but non-trivial result that if the strategy set is finite (or if it is compact and the utility function is continuous) then a strategy is undominated if and only if it is a best response. Assume two orders of deleting result in a different set of profiles, $A' = \prod_i A'_i$ and $A'' = \prod_i A''_i$. Let A^*_i be the set of strategies in A'_i and not in A''_i, and let a^*_1 be the first strategy in $A^* = \prod_i A^*_i$ that was deleted under the procedure that led to A''. Say, wlog, that a^*_i was deleted in the k^{th} round of the deletion procedure that led to A''. Since a^*_i is undominated, Theorem 1 implies it is a best reply to some belief on A'_{-i} (as it isn't deleted in the procedure that led to A'_i) it is also a best reply to a belief on opponents at the k^{th} round of the deletion procedure that led to A'' (since no strategy in A' is deleted in the procedure that leads to A'' before the k^{th} round). Hence it could not be deleted then. ■

Example 5 ***Iterative Elimination of Strictly Dominated Strategies***

Consider the following two player game.

	L	CL	CR	R
U	5,5	6,4	3,3	7,1
MU	4,1	0,3	0,4	5,12
MD	0,1	0,0	4,6	12,0
D	0,10	10,0	2,8	6,10

Observe that U dominates MU, L dominates CL, U dominates D, and L dominates R. We are left with the game:

	L	CR
U	5,5	3,3
MD	0,1	4,6

In this case, the set of strategies that survive iterative elimination of strictly dominated strategies does not yield a single outcome.

Exercise 5 *Describe an example of a two player 4×4 game in which a single action for each player survives successive elimination of dominated strategies, and that is such that each eliminated action except for the last one for each player is dominated by a strictly mixed strategy.*

Example 6 ***Iterative Elimination with an Infinite Number of Actions: Cournot Competition***

Suppose there are two firms $N = \{1, 2\}$. Each firm may produce a nonnegative quantity, $q_1, q_2 \in [0, \frac{1}{2})$, at zero cost.[3] *The Market price is $p = 1 - q_1 - q_2$ or 0, whichever is greater. Firm i's payoff is given by*

$$u_i = q_i \cdot p$$

Theorem 2 *The unique outcome that survives successive elimination of dominated strategies is $q_i = 1/3$.*

Proof. *The proof is via two claims.*

Claim 1 *If $q_j \geq x$ for some x, then $q_i = \dfrac{1-x}{2}$ strictly dominates any strategy $q_i > \dfrac{1-x}{2}$.*

Proof. *Plot $q_1(1 - q_1 - q_2)$ as a function of q_1 for some $q_2 \geq 0$ and notice where it peaks. Alternatively $du_i/dq_i = 1 - 2q_i - q_j \leq 1 - x - 2q_i < 0$ for $q_i > \frac{1-x}{2}$.* ■

[3] We assume that $q_1, q_2 < \frac{1}{2}$ because otherwise we can only eliminate weakly as opposed to strictly dominated strategies.

Claim 2 *If* $0 \leq q_j \leq x$ *for some* x*, then* $q_i = \dfrac{1-x}{2}$ *strictly dominates any strategy* $q_i < \dfrac{1-x}{2}$*.*[4]

Proof. *Plot* $q_1(1-q_1-q_2)$ *as a function of* q_1 *for some* $q_2 \geq 0$ *and notice where it peaks. Alternatively* $du_i/dq_i = 1 - 2q_i - q_j \geq 1 - x - 2q_i > 0$ *for* $q_i < \frac{1-x}{2}$*, so utility increases in* q_i *up to* $q_i = \frac{1-x}{2}$*.* ∎

Recall that $q_1, q_2 \in \left[0, \dfrac{1}{2}\right]$*. Claim 2 implies that after one round of elimination,* $q_1, q_2 \in \left[\dfrac{1}{4}, \dfrac{1}{2}\right]$*. After two rounds claim 1 implies that* $q_1, q_2 \in \left[\dfrac{1}{4}, \dfrac{3}{8}\right]$*, after three rounds Claim 2 again implies* $q_1, q_2 \in \left[\dfrac{5}{16}, \dfrac{3}{8}\right]$*, after four stages* $q_1, q_2 \in \left[\dfrac{5}{16}, \dfrac{11}{32}\right]$ *and so on. The limit is* $q_1, q_2 = \dfrac{1}{3}$*.*[5] ∎

Remark 2 *For an alternative proof see Bernheim (1984).*

Example 7 ***Iterative elimination of dominated strategies—Hotelling***

Consider the Hotelling game but where either firm / party can locate anywhere on a finite equally distributed subset of the interval $[0,1]$ *and each wants the largest number of customers / votes. Consumers / voters go to the closest firm / party. If the two parties are equally distant from the consumer/voter, then the consumer/voter randomizes with equal probability in choosing between the firms'/parties. So there are two players, with finite action sets* $A_i = \{0/m, 1/m, 2/m, \ldots (m-1)/m, 1\}$ *for some* $m \geq 5$*,*[6] *and*

$$u_i(a_i, a_j) = \begin{cases} a_i + (a_j - a_i)/2 & \text{if } \ a_j > a_i \\ 1/2 & \text{if } \ a_i = a_j \\ 1 - a_i + (a_i - a_j)/2 & \text{if } \ a_i > a_j \end{cases}$$

This describes the game. Now we solve by iteratively deleting dominated strategies.

Consider the strategy $a_i = 1$*. It is dominated by* $(m-1)/m$*. (It is also dominated by many other strategies, such as* $(m-2)/m$*, but that doesn't matter, the point is that it is dominated.) Note that* $(m-1)/m$ *is not dominated since if the opponent is located at 1*

[4] Notice that if $q_j > \frac{1}{2}$ then $q_i = \frac{1}{2}$ does not strictly dominate larger q_i's.

[5] **Explanation.** We have a sequence of intervals $\{[a_k, b_k]\}_{k \geq 1}$ that are such that $a_k < \dfrac{1}{3}$ and $b_k > \dfrac{1}{3}$ for every k, and such that the length of the kth interval is $\left(\dfrac{1}{2}\right)^k \searrow 0$ because by claim 2, $a_{k+1} = \dfrac{1-b_k}{2}$ and $b_{k+1} = \dfrac{1-a_k}{2}$ so the lenght of the $k+1$th interval is $b_{k+1} - a_{k+1} = \dfrac{b_k - a_k}{2}$.

[6] The result holds for smaller m but the notation is easier in what follows if we consider $m \geq 5$.

then $(m-1)/m$ *is the best one can do. A symmetric argument shows that* $1/m$ *dominates* 0 *and that* $1/m$ *is not dominated.*

After (and only after) 1 *(and similarly* 0*) are deleted, we see that* $(m-1)/m$ *is dominated by* $(m-2)/m$*.*

This process of deletion can continue until all that remains is $1/2$ *if* m *is even, and the two numbers closest to* $\frac{1}{2}$ *on either side if* m *is odd. [Why can't the process of deletion go past the point* $\frac{1}{2}$ *?]*

Exercise 6 *Consider a finite version of the Hotelling-Downs game with* $m=4$ *and* $m=8$ *and three parties / firms. What is the outcome of iteratively deleting dominated strategies?*

If we consider the case where the players can locate anywhere on the interval, so that the strategy space is $A_i=[0,1]$ *(so not restricting firms to locate at* k/m *for some fixed* m*) then the Hotelling-Downs game is not continuous. Therefore Theorem* **??** *does not apply. More disturbing is that the best-reply correspondence may be empty, that is, there are beliefs about the opponent against which there is no best reply. For example, if the opponent locates at* $1/3$ *then you would want to locate "just above* $1/3$*" but that is not a well defined point in* $[0,1]$*. It is possible to show that when* $A_i=(0,1)$ any *strategy in* $(0,1)$ *is a best reply.*

Exercise 7 *Show that if* $A_i=(0,1)$ *then any strategy in* $(0,1)$ *is a best reply.*

Iterative elimination of strictly dominated strategies has an appealing advantage, namely the analysis follows only from considerations of rationality. There is no need to assume that players know more about other players. It still assumes a lot: that all players know the *ordinal* payoffs in the game, are rational, and furthermore, that this is commonly known among the players to an arbitrarily large degree.

For example, is it reasonable that player 2 would play R in the game below?

	L	R
U	$3,4$	$2,5$
D	$2,1000000$	$0,-1000000$

Exercise 8 *Consider the following game of* N *players.* $A_i=\{1,...,M\}$*. The player who guesses closest to 2/3 of the average of the other players gets a prize of* 100*. If more than one is closest the winner is chosen randomly. Solve using iterated deletion of dominated strategies. (Hint: you can use the theorem mentioned above in the proof of Proposition 1 that states that a strategy is undominated if and only if it is a best response. Notice that this statement is equivalent to the statement that a strategy is dominated if and only if it is never a best response. Extra credit: once you identify that a certain strategy is dominated, can you identify a strategy that dominates it ?).*

2.4.1 Correlation / independence

How should one define beliefs over multiple opponents? One natural approach is that each player has beliefs over the possible strategy profiles (vectors) that the opponents may choose. For instance, if one of my opponents chooses among rows, U, D, and the other among columns, L, R, then I need to have beliefs over all pairs: $\{(U, L), (U, D), (D, L), (D, R)\}$. Thus a belief will be four non-negative numbers that sum to 1, say, (p_1, p_2, p_3, p_4).

However, one might think that players beliefs over their opponents are statistically (or stochastically) *independent.* In the example just considered I would have a belief $(q, 1-q)$ over which element of $\{U, D\}$ the row player chooses, and a belief $(r, 1-r)$ over which element of $\{L, R\}$ the column player chooses, and the joint belief over pairs in $\{(U, L), (U, D), (D, L), (D, R)\}$ is given by $(qr, q(1-r), (1-q)r, (1-q)(1-r))$.

This second alternative is a special case of the first, since the former allows for any four numbers (that are non-negative and sum to 1) while the latter only those allows for four such numbers that equal the product of the marginal probabilities. That is, the former allows for correlation, the latter does not.

The left hand table below is an example of beliefs that satisfy independence, and the left is an example that does not.

IND

	L	R
U	1/3	1/3
D	1/6	1/6

COR

	L	R
U	1/3	1/3
D	1/3	0

The left hand table below is the general form of beliefs that satisfy independence while the left is the general form of beliefs that need not satisfy independence. I specify the marginal probabilities on $\{U, D\}$ and $\{L, R\}$ as well in this table.

IND

		L	R
	marginals	r	$1-r$
U	q	qr	$q(1-r)$
D	$1-q$	$(1-q)r$	$(1-q)(1-r)$

COR

		L	R
	marginals	p_1+p_3	p_2+p_4
U	p_1+p_2	p_1	p_2
D	p_3+p_4	p_3	p_4

To discuss this more formally I introduce some notation. Given strategy sets A_i for $i = 1, ..., N$, let $A_{-i} = \prod_{j \neq i} A_j$. The set A_{-i} is the set of all strategy choices by all players other than i, so a typical element of A_{-i} is $(a_1, ..., a_{i-1}, a_{i+1}, ..., a_N)$. The set of all possible beliefs (allowing for correlation) of i over i's opponents is $\Delta(A_{-i}) \equiv \Delta\left(\prod_{j \neq i} A_j\right)$. Note that this is different from $\prod_{j \neq i} \Delta(A_j)$ which is the set of profiles (or vectors) of mixed strategies, or beliefs, over A_j for each player. When we specify a belief as an element $\alpha_{-i} \equiv (\alpha_j)_{j \neq i} \in \prod_{j \neq i} \Delta(A_j)$ we (implicitly) mean that beliefs satisfy independence so the probability of a pure strategy profile $(\bar{a}_j)_{j \neq i} \in A_{-i}$ is the product of the probability that each player $j \neq i$ will play $\bar{a}_j$, that is the probability of $(\bar{a}_j)_{j \neq i}$ is $\alpha_1(\bar{a}_1) \times \alpha_2(\bar{a}_2) \times \cdots$ $\times \alpha_{i-1}(\bar{a}_{i-1}) \times \alpha_{i+1}(\bar{a}_{i+1}) \times \cdots \ \times \alpha_N(\bar{a}_N) \equiv \prod_{j \neq i} \alpha_j(\bar{a}_j)$.

To clarify this notation further let me consider an example where player 1 chooses rows in $\{U, D\}$, 2 chooses columns in $\{L, R\}$, and 3 chooses matrices in $\{R, S, T\}$. In this case we have the following:

- The set of pairs of (pure) strategies that 1 and 3 can jointly take are $A_{-2} = \{U, D\} \times \{S, T, U\} = \{(U, R), (U, S), (U, T), (D, R), (D, S), (D, T)\}$.

- The set of beliefs over the pairs of (pure) strategies that 1 and 3 can jointly take is $\Delta(A_{-2}) = \{(p_1, p_2, ..., p_6) : p_i \geq 0, \sum p_i = 1\}$. That is, $\Delta(A_{-2})$ is the set of beliefs over A_{-2}, the set of all six non-negative numbers that sum to 1.

- The set of beliefs of 2 over 1 and 3 that are *independent* is the set $\{(qr_1, qr_2, qr_3, (1-q)r_1, (1-q)r_2, (1$ where q is the probability with which the row player chooses U and r_1 is the probability with which player 3 chooses R, r_2 is 3's probability of choosing S and r_3 is 3's probability of choosing T. This set is more easily and commonly described as $\Delta(A_1) \times \Delta(A_3)$ $= \{(q, 1-q) : q \in [0, 1]\} \times \{(r_1, r_2, r_3) : r_i \geq 0, \sum r_i = 1\}$, i.e., the set of pairs of mixed strategies for 1 and for 3 (where it is implicit that the joint distribution is independent).

I now use these concepts to reconsider our definitions of dominance, best replies and never best replies. The definition of dominance is unchanged because there are no beliefs in the definition. The definition of when a strategy is a best reply against a specific belief is unchanged because we specify the belief.

Definition 6 *A strategy $a_i \in A_i$ for player i is dominated if there is another (possibly mixed) strategy $\alpha_i \in \Delta(A_i)$ s.t. for all pure strategy profiles of the opponents α_i is strictly better: $\forall a_{-i} \in A_{-i}$ we have $u_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i})$.*

Definition 7 *A strategy $\bar{a}_i \in A_i$ is a best reply against $p \in \Delta(A_{-i})$ if $u_i(\bar{a}_i, p) \geq u_i(a_i, p)$ for all $a_i \in A_i$.*

In the above p can be stochastically independent or not. If we want to consider a p that is stochastically independent then – as in our notation above – we will write it as $\alpha_{-i} \in \prod_{j \neq i} \Delta(A_j)$. This is important when we define strategies that are never best replies.

Definition 8 *A strategy $\bar{a}_i \in A$ is never a best reply* ***against independent beliefs*** *if for all independent beliefs $\alpha_{-i} \in \prod_{j \neq i} \Delta(A_j)$ there exists another pure strategy $\hat{a}_i \in A_i$ (that may depend on $\bar{a}_i$) such that $u_i(\hat{a}_i, \alpha_{-i}) > u(\bar{a}_i, \alpha_{-i})$.*

Definition 9 *A strategy $\bar{a}_i \in A$ is never a best reply (against any beliefs – including those that allow for correlation) if for all beliefs $p \in \Delta(A_{-i})$ there exists another pure strategy $\hat{a}_i \in A_i$ (that may depend on $\bar{a}_i$) such that $u_i(\hat{a}_i, \alpha_{-i}) > u(\bar{a}_i, \alpha_{-i})$.*

To further clarify the difference between these last two definitions consider the following game. Players 1 and 2 choose rows and columns, and 3 matrices. Only player 3's payoffs are specified.

$$\begin{array}{ccc}
A & B & C \\
\begin{array}{c|c|c|}
 & L & R \\ \hline
U & 5 & 6 \\ \hline
D & 6 & 7 \\ \hline
\end{array}
&
\begin{array}{c|c|c|}
 & L & R \\ \hline
U & 7 & 6 \\ \hline
D & 6 & 5 \\ \hline
\end{array}
&
\begin{array}{c|c|c|}
 & L & R \\ \hline
U & 6 & 0 \\ \hline
D & 0 & 6 \\ \hline
\end{array}
\end{array}$$

Strategy C is a best reply against the following belief:

	L	R
U	$1/2$	0
D	0	$1/2$

This is easy to verify: all three strategies give the payoff 6.

However C is not a best reply against any independent belief of the form

	L	R
U	pq	$p(1-q)$
D	$(1-p)q$	$(1-p)(1-q)$

which is the independent product of the probability of the row player choosing U, for any $p \in [0,1]$, with any $q \in [0,1]$, which is the probability of the column player choosing L.

To see why C is not a best reply against any such belief note that if the probability of (U, L) is strictly greater than that of (D, R) then B yields more than C (because B then yields strictly more than 6 while C yields at most 6). Similarly if the probability of (D, R) is greater than that of (U, L) then A yields strictly more than C. So only when the probability of (U, L) and that of (D, R) are *equal* can C be a best reply. But when those two probabilities are equal C gets the same as A and B conditional on (U, L) or (D, R) and strictly less otherwise, and also when those two probabilities are equal the "otherwise" event must have strictly positive probability.

Our theorem from before continues to hold:

Theorem 3 *If the strategy space is finite (or if it is compact and the utility function is continuous) then a strategy is never a best reply against* ***any*** *belief (including those that are not necessarily independent) if and only if it is dominated,*

The proof is exactly as before. Why doesn't the theorem apply when we restrict attention to independent beliefs? The proof was for arbitrary beliefs. More specifically, when we constructed the beliefs using the separating hyperplane we have no way of knowing that it will be independent (in fact, in general, it will not).

Exercise 9 *Which, if either, of the following is true:*

- *If a strategy is never a best reply against any* ***independent*** *belief then it is dominated.*
- *If a strategy is dominated then it is never a best reply against any* ***independent*** *belief.*

Exercise 10 *Consider the game below where one player chooses rows (U or D), one chooses columns (L or R) and one chooses tables (A, B or C). Only the payoffs of the player who chooses tables are given. Show that C is not a best reply to any independent belief about the row and column players but that it is a best reply to a correlated belief. The latter is relatively easy; the former requires a bit more algebra, so here's a hint for the former: assume that the independent mixed strategies of the row and column player are such that A is at least as good as B (and then assume the converse).*

A

	L	R
U	5	5
D	0	5

B

	L	R
U	5	0
D	5	5

C

	L	R
U	0	2.6
D	2.6	0

2.4.2 Rationalizability

A strategy is rational if it is a best reply to some belief. A strategy of i is rationalizable if it is a best reply to some belief p_{-i} over strategies of the opponents, where each strategy of any opponent a_j that is assigned strictly positive probability is a best reply to some belief over j's opponents, and so on.

Using Theorem **??** and the preceding discussion on correlation one can prove that a strategy is rationalizable if and only if it survives iterated deletion of dominated strategies—so long as the beliefs over opponents allow for correlation. If one imposes independence on beliefs the solution concepts differ.[7]

[7] Observe that allowing for correlated beliefs implies that the concept of rationalizability is weaker, in the sense that it is easier to satisfy, than if we insist that players believe that other players play independently.

Exercise 11 *If players believe that other players play independently, which is larger, the set of rationalizable strategies, or the set of strategies that survive successive elimination of strictly dominated strategies ?*

2.5 Weakly Dominated Strategies

Definition 10 *A strategy a_i weakly dominates a strategy a_i' if for every profile a_{-i} of strategies of the other players*

$$u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i})$$

and there exists some profile $\hat{a}_{-i}$ such that $u_i(a_i, \hat{a}_{-i}) > u_i(a_i', \hat{a}_{-i})$.

Example 8 *The following game cannot be solved through the iterative elimination of strictly dominated strategies.*

$1,1$	$0,0$
$0,0$	$0,0$

But it can be solved through the (iterative) elimination of weakly dominated strategies.

Example 9 ***First-price auction with known values.***

A single object is for sale. The value to bidder 1 is 10 and to bidder 2 is 5. The highest bidder wins and pays his/her bid, in case of a tie they win with equal probability and pay if they win. There is a maximal bid of 15. (Consider only integer bids.) Bidding 15 is dominated for both, and then 14 is, ... and then 11 is, and then 10 is for bidder 2 (10 is not dominated for bidder 1 because generates a payoff of 0 and no other bid guarantees a strictly positive payoff). After 10 is eliminated for bidder 2, it can also be eliminated for bidder 1 because it is dominated by 9. Bidding 9 for bidder 2 would yield a negative amount, so zero, which may generate a positive payoff dominates 9 for 2. After that nothing is dominated, but there are weakly dominated strategies. After deleting 9 for 2, we can continue and observe that 2 will not bid more than 5 so 1 will not bid more than 6. Is this all? No. Bidding 0 is weakly dominated for 1 by bidding 1. Next, bidding zero and bidding 5 is weakly dominated for 2, since bidding 1 can get a positive payoff. But if 2 does not bid zero then bidding 1 for 1 is weakly dominated by bidding 2. Then for 2 bidding 1 is weakly dominated by bidding 2. Continuing in this way we can show that the only remaining bids are 5 for 1 and 4 for 2.

Exercise 12 *What is the result of iterated deletion of strongly dominated strategies if there is no limit on the bids? Which is a better model—the case where there is a (commonly known) limit, or no limit? (Remember this question when we study repeated games.)*

Example 10 ***The Outcome of Iterative Elimination of Weakly Dominated Strategies May Depend on the Order of Elimination.***

Consider the following game matrix.

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

One order of elimination is: M weakly dominates T, then R weakly dominates L which produces the outcomes (M,R) or (B,R) with payoffs $(2,1)$ or $(2,1/2)$. A second order of elimination is: M weakly dominates B, then L weakly dominates R which produces the outcomes (T,L) or (M,L) with payoffs $(1,1)$. A third order—simultaneous deletion—results in payoffs $(1,1)$ or $(2,1)$ and outcomes (M,L) or (M,R). (Notice that if there is a third player these different results can lead to different actions for that player, so the difference in the result of deleting according to different orders can be greater than they appear in this simple example.)

Exercise 13 ***Second-price auction with known values***

A single object is for sale. The value to bidder 1 is v_1 and to bidder 2 is v_2. The highest bidder wins and pays the second-highest bid, in case of a tie they win with equal probability and pay if they win. Prove that bidder i bidding v_i is weakly dominant (i.e., weakly dominates every other strategy). Show that if we solve using iterative deletion of weakly dominated strategies then the order of deletion matters to the auctioneer.

A natural question is whether there is an analog to Theorem **??**. The answer is yes: a strategy is weakly dominated if and only if it is never a cautious best reply, where a cautious best reply is a best reply to a belief that puts *strictly* positive probability on every pure strategy of the opponent.

Exercise 14 *Consider again exercise* **??** *but solve the case where $A_i = [0, M]$ using iterated deletion of weakly dominated strategies.*

2.6 Pure-Strategy Nash Equilibrium

Definition 11 *A profile of actions $(a_1^*, ..., a_n^*)$ is a pure-strategy Nash equilibrium if for every player i,*

$$u_i\left(a_i^*, a_{-i}^*\right) \geq u_i\left(a_i, a_{-i}^*\right)$$

for every $a_i \in A_i$. In other words, for every player i, $a_i^* \in \arg\max u_i\left(a_i, a_{-i}^*\right)$, or $a_i^* \in BR_i\left(a_{-i}^*\right)$.

Example 11 ***Battle of the Sexes***

	$O(pera)$	$B(oxing)$
O	$2,1$	$0,0$
B	$0,0$	$1,2$

This game has two pure-strategy Nash equilibria (O,O), (B,B).

Another famous coordination game is the following.

Example 12 ***Stag Hunt***

	$A(lone)$	$B(oth)$
A	$1,1$	$1,0$
B	$0,1$	$3,3$

The stag hunt game also has two pure strategy equilibria (A,A), (B,B). One might argue that A is "safer" which could make (A,A) more plausible in spite of the fact that it is less efficient. But note that since payoffs are in utils, there is no sense in which one is less risky than the other. Any argument in favor of (A,A) must involve some additional type of reasoning.

Example 13 ***Chicken***

	$C(hicken)$	$B(rave)$
C	$0,0$	$-1,1$
B	$1,-1$	$-2,-2$

[Show Chicken video from Cry-Baby 1990 Dir. John Waters]

Example 14 ***Matching pennies***

Two players each pick Heads (H) or Tails (T) simultaneously. The objective of the row player is to match the choice of the column player. The objective of the column player is to un-match the row player. This game is equivalent to the 'odd or even' game after *the row player announced 'even.'*

	$T\,(ails)$	$H\,(eads)$
T	$1,-1$	$-1,1$
H	$-1,1$	$1,-1$

There does not exist a pure strategy Nash equilibrium in this game.

Example 15 (Cournot Competition (Example ??) revisited) *The game:*

$$\begin{aligned} N &= 2 \\ A_i &= \mathbf{R}_+ \text{ (elements of } A_i \text{ are denoted by } q_i\text{)} \\ u_i\,(q_1, q_2) &= q_i \times p\,(q_1 + q_2) \end{aligned}$$

We assume $p\,(q_1+q_2) = 1 - (q_1+q_2)$ *if* $q_1+q_2 \leq 1$ *and* 0 *otherwise.*

Computing a Nash equilibrium. For any given q_2*, firm* 1*'s maximization problem is given by*

$$\begin{aligned} &\max_{q_1 \geq 0} q_1\,(1 - q_1 - q_2) \\ FOC &: \; q_1 = \frac{1-q_2}{2} \end{aligned}$$

Firm 2 solves a similar problem. The solution $q_1^* = q_2^* = \frac{1}{3}$*, which is called a Cournot equilibrium, is a Nash equilibrium of the game. (The vertical axis is* q_1*, the horizontal is* q_2*. The less steep line is* $q_1\,(q_2)$ *the steeper line is* $q_2\,(q1)$*.)*

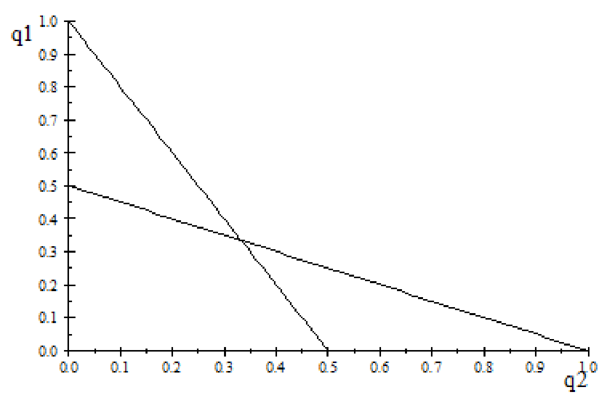

Remark 3 *In general, with* n *players, if we assume the products are homogenous, then we would have* $p\,(q) = 1 - \sum_{i=1}^{n} q_i$ *where* $q = (q_1, ..., q_n)$*. If the products are not homogenous (differentiated products) we would have more general functions and individualized prices,*

e.g. $p_i(q) = 1 - b_i q_i - b\sum_{j \neq i} q_j$*, if we think of all the other products as symmetric, or more generally* $p_i(q) = 1 - \sum b_j q_j$*. The case* $b_j > 0$ *corresponds to goods that are substitutes, and* $b_j < 0$ *complements. It is standard to assume* $b_i > |b_j|$*, the "own-price" effect is greater than the "cross-price" effect. (The case of equality is the homogenous case we study.)*

Exercise 15 *Find the Nash equilibria of the* $n-$*player homogenous good Cournot model.*

Observe that in contrast to the two-player case, with three players iterative deletion of dominated strategies does not yield the Nash equilibrium in Cournot competition.

Example 16 ***Bertrand competition (1883)***

Two firms operate in a market. The firms must decide simultaneously what price to charge. A single consumer buys one unit of the good from the cheaper firm. If both firms charge the same price, the consumer chooses the firm from which to buy randomly. The costs of production are zero for both firms.

The unique Nash equilibrium is $(0, 0)$*. Recall that a pure strategy Nash equilibrium is a profile of actions from which no player wishes to deviate. Note that at least one player benefits from deviating from any profile of prices* $(p_1, p_2) \neq (0, 0)$*.*

We cannot draw the best-reply correspondence for the case where firms can choose any price, since there is no best reply against the opponent choosing $p > 0$*: one wants to choose as close as possible to* p *but below.*

Exercise 16 *Optional: Solve for Bertrand price competition the case of non-homogenous price competition (defined similarly to the case of non-homogenous Cournot competition with prices instead of quantities, and quantities instead of prices) and draw the best-reply correspondence.*

Strategic substitutes and complements

When solving for comparative statics it turns out to be useful to consider two special cases of best replies: the case where they are increasing for all players and the case – as in Cournot quantity competition – where they are decreasing. This is useful because if we consider a policy that raises one person's BR then in the case of strategic complements all actions will move in the same direction while in the case of strategic substitutes they will move in the opposite direction. (Indeed if we consider "stable" equilibria – which we will discuss later – we can say more.)

Given $u(x, y)$ when is $BR(y)$ increasing? $BR(y)$ is defined implicitly by the first order condition (FOC) $u_1(BR(y), y) = 0$ (where the subscript of 1 indicates the derivitative with respect to the first argument). To find whether this function, $BR(y)$, is upward or downward sloping we take the derivative of this FOC: $u_{11}\left(BR\left(y\right),y\right)BR'\left(y\right)+u_{12}\left(BR\left(y\right),y\right)=0$. So $BR'\left(y\right)=-u_{12}\left(BR\left(y\right),y\right)/u_{11}\left(BR\left(y\right),y\right)$. If, as is common, we assume that u is concave in one's own decision variable (so that the FOC indeed characterizes the maximum) then $u_{11}<0$ and we have that the sign of BR' is the same as that of u_{12}.

Coaltional deviations

Nash equilibria are robust to deviations of single players, but not to "coalitional deviations" This implies that Nash equilibria can be Pareto inefficient, as in the stag-hunt game). There are "refinements" of Nash equilibrium that are concerned with this issue.

Obvious (?): allowing for coalitional deviations eliminates Pareto dominated Nash equilibria.

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

Less obvious (!): coalitional deviations can also eliminate other equilibria, and even some Pareto dominant equilibria.

Consider the simultaneous move game below, where player 3 chooses between the two matrices, where we denote the left matrix by l and the right by r.

	L	R
U	$3,3,0$	$0,0,0$
D	$0,0,0$	$2,2,2$

	L	R
U	$1,1,1$	$0,0,0$
D	$0,0,0$	$-1,-1,-1$

Here (D,R,l) is a Nash equilibrium but it is not robust to a coalitional deviation. Player 3 might be concerned that I and II will deviate to (U,L). Thus, the inferior Nash equilibrium (U,L,r) is more likely from this perspective, even though the coalition of 3 players would prefer (D,R,l) (this example touches upon the subject of coalitional stability and its relation to farsighted reasoning on part of the players).

Weakly dominated strategies

Nash equilibria may involve weakly dominated strategies; there are "refinements" to deal with this as well.

Example 17 *Consider again example* **??**. *This game has two Nash equilibria:* (U,L) *yielding* $(1,1)$ *and* (B,R) *yielding* $(0,0)$. *(Only* $(1,1)$ *does not use dominated strategies. We will discuss this further when considering refinements later.)*

2.6.1 Interpretation of Nash Equilibrium

What is required for pure strategy Nash equilibrium to be a plausible prediction for the way in which a game would be played?

1. (easy ?) Each player is rational, in the sense that he chooses the optimal response to other players' strategies, given his belief.
2. (hard ?) Each player needs to have a **correct** prediction of the behavior (or beliefs) of the other players.

Thus, compared to elimination of dominated strategies, Nash equilibrium involves less iterative/strategic thinking, and more knowledge about opponents. When might this knowledge be plausible?

1. When there is a social norm, which means that everyone knows how the game is usually played. For example, everybody knows what is the norm with respect to what side of the road to drive a car in different countries. Sometimes, it is enough that a game is played often enough and that play is stable over time to produce something that functions as a social norm in this context. Indeed, economists tend to think of social norms as Nash equilibria. But what's interesting about many social norms is that they are not Nash equilibria (at least not in any obvious way). Consider for example the norms of not spitting in public, picking after your dog, and tipping. Iy is interesting to think about what sustains these norms. It seems that the sanctions that support these norms are insufficient to support them on their own. What do you think ?
2. When the players can communicate and possibly reach an agreement on a self-enforcing plan of play.

   This may work when the game is a coordination game with a single pure strategy Nash equilibrium. However, consider the following example.

	L	R
T	$7,7$	$2,9$
B	$9,2$	$1,1$

   This game has two pure strategy NE, (B,L) and (T,R). If the row player says to the column player that column should play L because row is playing B, column could answer that row would want column to play L in any case. Communication is not necessarily informative or credible.

3. When the game is played repeatedly and the players follow some dynamic learning or evolutionary or mimicking process that converges to a Nash equilibrium.[8]

More exercises

Exercise 17 *Consider the three-player Hotelling – Downs game. Are there pure-strategy Nash equilibria where all choose the same location? Where two choose the same location? Where all choose different locations?*

Exercise 18 ***The Nash Bargaining game***

Each player chooses a number in $\mathbf{R}$*. If the sum is greater than 1 they both get 0. If the sum is equal to or less than 1 they each get what they asked for. What are the pure strategy Nash equilibria of this game?*

Exercise 19 *Find the pure strategy Nash equilibria in the second-price auction where players values are* 5 *and* 10 *for players 1 and 2 respectively.*

2.7 Mixed-Strategy Nash Equilibrium

A pure-strategy Nash equilibrium does not always exist. For example, it fails to exist even in the simple matching-pennies game above. We therefore extend the definition of Nash equilibrium to allow for mixed strategies. It can be shown that Nash equilibria that allow for mixed strategies exist under very general conditions.

Definition 12 *A profile of strategies* $(\alpha_1^*, ..., \alpha_N^*)$ *is a Nash equilibrium if for every player* i,

$$u_i\left(\alpha_i^*, \alpha_{-i}^*\right) \geq u_i\left(a_i, \alpha_{-i}^*\right)$$

for every $a_i \in A_i$.[9]

Remark 4 *Of course if the inequality holds for all* $a_i \in A_i$ *then it also holds for all* $\alpha_i \in \Delta(A_i)$.

[8]Mention Shapley's famous example of non-convergence of fictitious play in the following game, which is a limit case of generalized Rock, Paper, Scissors games. In this game, if the players start by choosing (a, B), the play will cycle indefinitely.

	A	B	C
a	0,0	1,0	0,1
b	0,1	0,0	1,0
c	1,0	0,1	0,0

[9]As discussed earlier, if α_i is a mixed strategy, then $\alpha_i(a_j) = \Pr(\text{player } i \text{ plays strategy } j)$. In this case $u_i(a_i, \alpha_j) = \sum_{a_j} u_i(a_i, a_j)\, \alpha_j(a_j)$ and $u_i(\alpha_i, \alpha_j) = \sum_{a_i, a_j} u_i(a_i, a_j)\, \alpha_i(a_i)\, \alpha_j(a_j)$.

Example 18 ***Battle of the Sexes (B-o-S) revisited***

	$O(pera)$	$B(oxing)$
O	$2,1$	$0,0$
B	$0,0$	$1,2$

When would player 2 randomize between O and B? Only in case of indifference, that is, only if playing O and B yields the same expected payoff, or:

$$\begin{aligned} p \cdot 1 + (1-p) \cdot 0 &= p \cdot 0 + (1-p) \cdot 2 \\ p &= \frac{2}{3} \end{aligned}$$

Similarly, player 1 will also randomize only in case of indifference, or:

$$\begin{aligned} q \cdot 2 + (1-q) \cdot 0 &= q \cdot 0 + (1-q) \cdot 1 \\ q &= \frac{1}{3} \end{aligned}$$

Therefore, $(p,q) = \left(\frac{2}{3}, \frac{1}{3}\right)$ *is a Nash equilibrium of the game.*

This is also true more generally: If player i plays a mixed strategy in equilibrium then i must be indifferent, so j's strategy must keep i indifferent. We demonstrate the equilibrium graphically: player 1's strategies are on the horizontal axis, and player 2's best reply correspondence $BR_2(p)$ is the red and dashed function that describes for each mixed strategy of player 1 what 2 would want to choose. Player 2's strategies are on the vertical axis and $BR_1(q)$ is the solid-black function (viewed as a function from the vertical axis onto the horizontal axis).

$$BR_1(q) = \begin{cases} p = 1 & q > \frac{1}{3} \\ 0 \le p \le 1 & q = \frac{1}{3} \\ p = 0 & q < \frac{1}{3} \end{cases} \qquad BR_2(p) = \begin{cases} q = 1 & p > \frac{2}{3} \\ 0 \le q \le 1 & p = \frac{2}{3} \\ q = 0 & p < \frac{2}{3} \end{cases}$$

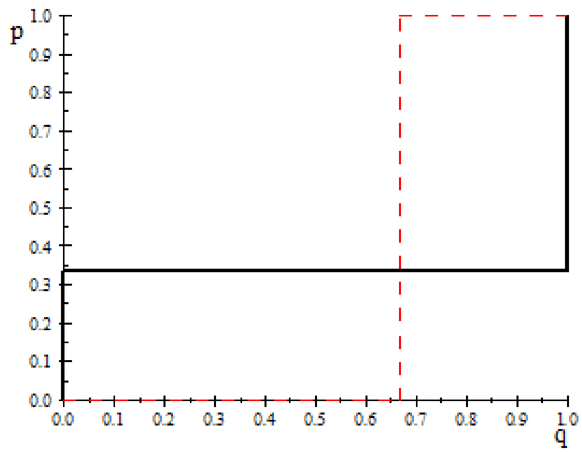

Expected payoffs to each player in the mixed strategy equilibrium are $\frac{2}{3}$, less than the payoffs to either player in either one of the pure strategy equilibria of the game.

Exercise 20 *Comparative statics in the B-o-S example.*

	L	R
T	$2+x, 1$	$0, 0$
B	$0, 0$	$1, 2$

As the payoff to the row player from the T, L equilibrium increases what happens to the mixed strategy equilibrium — specifically what happens to each player's mixed-strategy in the equilibrium — and what happens to the players' expected payoffs? Think of x being very large–is this what you would predict?

Example 19 ***Matching pennies revisited***

	$T(ails)$	$H(eads)$
T	$1, -1$	$-1, 1$
H	$-1, 1$	$1, -1$

In this game there exists a unique (mixed) Nash equilibrium given by $(p, q) = \left(\frac{1}{2}, \frac{1}{2}\right)$.

Example 20 ***Rock-Paper-Scissors***

	R	P	S
R	*0,0*	*-1,1*	*1,-1*
P	*1,-1*	*0,0*	*-1,1*
S	*-1,1*	*1,-1*	*0,0*

[Show three videos: Bud, Logic, how to win; there are many other references to the game in general culture; mention difficulty of randomization, and fact that people are bad at producing random sequences.]

Example 21 ***The Case of Kitty Genovese***

Based on a New York Time article from March 27th, 1964.[10] *In March 1964, a woman was murdered in Queens, New York. 38 neighbors witnessed the crime, yet no one called the police.*

A model:

n identical neighbors.

The (personal) cost of calling the police: 1.

The (personal) utility if someone calls the police: $x > 1$.

Nash equilibrium in pure strategies: exactly one neighbor calls the police. Such an equilibrium is implausible without some coordination.

We calculate a symmetric mixed-strategy Nash equilibrium: Every neighbor calls with probability p.

Each neighbor has to be indifferent between calling the police or not.

The payment if he chooses to call: $x - 1$.

The payment if he chooses not to call: $x \cdot (1 - (1 - p)^{n-1})$.

Indifference requires that: $x - 1 = x \cdot (1 - (1 - p)^{n-1})$.

And so we get: $p = 1 - (\frac{1}{x})^{\frac{1}{n-1}}$

The chance that someone will call the police:

$$\begin{aligned} 1 - (1-p)^n &= 1 - \left(1 - 1 + \left(\frac{1}{x}\right)^{\frac{1}{n-1}}\right)^n \\ &= 1 - \left(\frac{1}{x}\right)^{\frac{n}{n-1}} \end{aligned}$$

For $n = 1$ this expression equals 1.

The expression decreases as n increases and when $n \to \infty$ the probability converges to $1 - \frac{1}{x}$ as can be seen in the figure below that plots the function for $x = 3$ and for $n \geq 1$.

[10] See http://www.garysturt.free-online.co.uk/The%20case%20of%20Kitty%20Genovese.htm

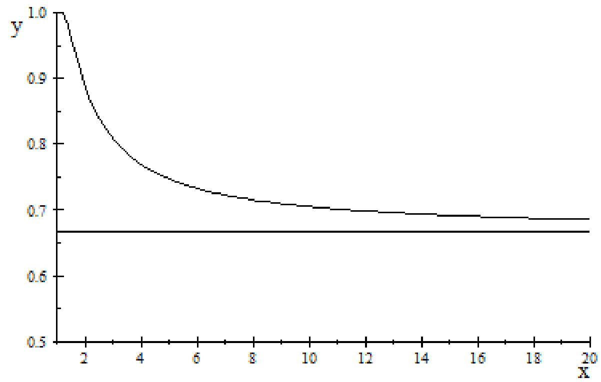

Exercise 21 *Find at least one mixed-strategy equilibrium in the Nash bargaining game of exercise* **??**.

Exercise 22 (Optional-not to be handed in-technically challenging) *Find a mixed-strategy equilibrium in the Nash bargaining game of exercise* **??** *that has an infinite support. Show that any pair of payoffs* (x, y) *such that* $x + y \leq 1$ *can be obtained in some equilibrium.*

2.7.1 Interpreting mixed-strategy equilibria

Mixed strategies can be thought of in three ways:

- a player actually mixing;
- uncertainty of one player about the other — where there is no actual mixing going on and the Nash equilibrium is a description of players' beliefs, not of their actions;
- a population mixing, where some proportion plays one way and others play differently.

Is a mixed strategies equilibrium plausible?

- Why should the player choose such probabilities which will keep the other player indifferent?
- Why should players know one another's actions or beliefs?
- If you think about a dynamic process of convergence to an equilibrium, would it be reasonable that the convergence will be to this equilibrium? (Depends on the process and the equilibrium. Show phase diagrams for Matching Pennies vs. Battle-of-Sexes)

- Zero-sum (and similarly constant-sum) games have various properties that make Nash equilibria more convincing. In all equilibria a player gets the same payoff, which is the maxmin, namely $\max_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in \Delta(A_2)} u_1(\alpha_1, \alpha_2)$, which equals the minmax $\min_{\alpha_2 \in \Delta(A_2)} \max_{\alpha_1 \in \Delta(A_1)} u_1(\alpha_1, \alpha_2)$, which in turn is equal to minus the other player's maxmin. Minmax and maxmin imply that players play so as to maximize their payoff under the assumption that the other players does the best it can to thwart their efforts. Intuitively, under the maxmin player 2 knows and best responds to player 1, which does the best it can in this situation; under the minmax it is player 1 who knows and best responds to 2, who does the best it can; so, intuitively, minmax $\geq$ maxmin. But remarkably, in zero sum games the minmax and maxmin are equal, and moreover, minmax and maxmin strategies are also a NashEquilibrium of the game. Moreover, the equilibria are also "exchangeable": if (α_1, α_2) and (α_1', α_2') are both equilibria, then so is (α_1, α_2'). Thus, the question of how one knows which equilibrium to play is not an issue. Unfortunately, true zero-sum games or situations seem to be rare in social sciences.

- There is another justification for mixed-strategy equilibria called purification. While we will not formally discuss the model, consider (loosely) the following enhancement of the B-o-S game: there is a large group of people playing the game, 2/3 of the population have payoffs (just) above 2 instead of 2 and 1/3 have (just) below 2. Then the 2/3-1/3 equilibrium has no one indifferent. Of course we chose the perturbation "conveniently". A remarkable result of Harsanyi is that this holds "almost" always for a "large" class of perturbations in the sense that for almost any game and any perturbation, almost any mixed strategy Nash equilibrium is a limit as perturbations decrease to zero, and any limit is a mixed strategy equilibrium of the game.

- In general the interpretation of mixed-strategy equilibria and their comparative statics is sometimes difficult. Nevertheless, in many games they are necessary for existence and often do have some interpretation.

Exercise 23

1. *Show that in general two player strategic form games*

$$\max_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in \Delta(A_2)} u_1(\alpha_1, \alpha_2) \leq \min_{\alpha_2 \in \Delta(A_2)} \max_{\alpha_1 \in \Delta(A_1)} u_1(\alpha_1, \alpha_2)$$

2. *Show that* $\max_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in \Delta(A_2)} u_1(\alpha_1, \alpha_2) = \min_{\alpha_2 \in \Delta(A_2)} \max_{\alpha_1 \in \Delta(A_1)} u_1(\alpha_1, \alpha_2)$ *in the zero-sum game below (the numbers in the table describe the payoffs of the row player)*

3	*-2*	*2*
-1	*0*	*4*
-4	*-3*	*1*

2.8 Existence of Nash Equilibrium

Denote the set of player i's best responses to the profile of other players' strategies α_{-i} by $BR_i(\alpha_{-i})$. To show that a game has a Nash equilibrium it suffices to show that there is a profile α^* of strategies such that $\alpha_i^* \in BR_i(\alpha_{-i}^*)$ for all $i \in N$. Define the set-valued function $BR : \Delta(A) \to \Delta(A)$ by $BR(\alpha) = \times_{i \in N} BR_i(\alpha_{-i})$. Then to prove the existence of Nash equilibrium we need to prove that there exists a vector of strategies α^* that is such that $\alpha^* \in BR(\alpha^*)$.

Fixed point theorems give conditions on functions under which the functions possess "fixed points" i.e., points that are mapped by the functions back into themselves. We shall rely on the following well known fixed point theorem.

Theorem 4 ***Kakutani's Fixed Point Theorem (1941).*** *Let X be a compact*[11] *convex*[12] *set of $\mathbb{R}^n$ and let $f : X \to X$ be a set-valued function for which for all $x \in X$ the set $f(x)$ is nonempty and convex,and the graph of f is closed. Then, there exists $x^* \in X$ such that $x^* \in f(x^*)$.*

Exercise 24 *For every one of the assumptions mentioned in Kakutani's Theorem find a counter-example that satisfies all the other assumptions yet fails to have a fixed point.*

Theorem 5 *A finite strategic form game $\langle N, (A_i), (u_i) \rangle$ has a Nash equilibrium.*

Proof. For every player $i \in N$ the set $BR_i(a_{-i})$ is nonempty (since it is a set of maximizers of a continuous function on a compact set) and convex (since expected payoffs are linear in probabilities, or $u_i(\lambda\alpha + (1-\lambda)\alpha', \alpha_{-i}) = \lambda u_i(\alpha, \alpha_{-i}) + (1-\lambda) u_i(\alpha', \alpha_{-i})$).

We show that the function BR has a closed graph. Let $\{\alpha^n\}$ and $\{\beta^n\}$ be two sequences that are such that $\beta^n \in BR(\alpha^n)$ for all n, $\alpha^n \to \alpha$ and $\beta^n \to \beta$, we need to show that $\beta \in BR(\alpha)$. The definition of BR implies that for every player $i \in N$, every strategy $\gamma_i \in \Delta(A_i)$, and every profile of other players' strategies α_{-i}^n,

$$u_i\left(\beta_i^n, \alpha_{-i}^n\right) \geq u_i\left(\gamma_i, \alpha_{-i}^n\right) \qquad \forall n \geq 1$$

[11] A set $X \subseteq \mathbb{R}^n$ is compact if it is closed and bounded. A set $X \subseteq \mathbb{R}^n$ is closed if for every convergent sequence of elements in the set $x_n \to x$ and $x_n \in X$ for every $n \geq 1$ then $x \in X$. A set $X \subseteq \mathbb{R}^n$ is bounded if there exists some number K such that $|x_i| \leq K$ for every $x \in X$.

[12] A set X is convex if for every $x, y \in X$ and $\lambda \in [0,1]$, $\lambda x + (1-\lambda) y \in X$.

Continuity of the expected payoff u_i on $\times_{i\in N}\Delta(A_i)$ then implies that for every $i \in N$, and strategy $\gamma_i \in \Delta(A_i)$

$$u_i(\beta_i, \alpha_{-i}) \geq u_i(\gamma_i, \alpha_{-i})$$

which implies that $\beta_i \in BR(\alpha_i)$.[13]

Therefore, by the Kakutani Fixed Point Theorem the mapping BR has a fixed point. ■

Remark 5 *The existence of Nash equilibrium in infinite games can be established along similar lines, provided that the set of players' actions A is compact and players' payoffs are continuous in their actions. If in addition the players payoffs are quasiconcave in their own actions then one can prove existence of a* pure-*strategy Nash equilibrium.*

Exercise 25 *A two player game is symmetric if $A^1 = A^2$ and $u^1(x,y) = u^2(y,x)$ for every $x, y \in A^1, A^2$. Prove that a finite symmetric two player game has a symmetric Nash equilibrium (where both players employ the same strategy). (Hint: use Kakutani's Fixed Point Theorem).*

2.8.1 Upper and lower hemicontinuity of the Nash equilibrium correspondence

A correspondence $\Gamma : A \to B$ is **upper hemicontinuous (uhc)** at a if for every open neighborhood of $\Gamma(a)$, V, there exists a neighborhood of a, U, such that $\Gamma(x) \subseteq V$ for every $x \in U$. Or, equivalently, if Γ has a closed graph and a compact image. A correspondence is uhc if it is uhc at every point. A correspondence $\Gamma : A \to B$ is **lower hemicontinuous (lhc)** at a if for every open set V that intersects $\Gamma(a)$ there exists a neighborhood of a, U, such that $\Gamma(x)$ intersects V for every $x \in U$. A correspondence is lhc if it is lhc at every point.

Example 22 *The correspondence*

$$\Gamma(x) = \begin{cases} 0 & 0 \leq x < 1 \\ [0,1] & x = 1 \end{cases}$$

is uhc but not lhc. The correspondence

$$\Gamma(x) = \begin{cases} [0,1] & 0 \leq x < 1 \\ 0 & x = 1 \end{cases}$$

is lhc but not uhc.

[13]Otherwise, suppose that some player i has a strategy γ_i that is such that $u_i(\beta_i, \alpha_{-i}) < u_i(\gamma_i, \alpha_{-i})$.Continuity of u_i implies that for every n that is large enough $u_i(\beta_i^n, \alpha_{-i}^n) < u_i(\gamma_i, \alpha_{-i}^n)$.A contradiction.

In the proof of existence of Nash equilibrium above we argued that the graph of the best-reply correspondence is closed, which means that it is uhc. This is a useful result more generally when studying equilibria of games. For example, sometimes it is easier to solve a finite approximation of an infinite game. For instance, we can consider a sequence of games with strategy spaces $\{0, 1/m, ...1\}$ converging to the infinite Hotelling game. (This also comes up in dynamic games, where longer- and longer finite versions of an infinite-horizon game are an approximation.) Since $\lim_{n\to\infty} BR_i\left((a_{-i})_n\right) \in BR\left(\lim_{n\to\infty}(a_{-i})_n\right)$ we can conclude that the limit of a sequence of Nash equilibria is a Nash equilibrium. So we can find an equilibrium of a limit game by taking limits of equilibria of approximating games.

However, because the best response correspondence is not lhc, not all Nash equilibria in the limit game are limits of equilibria in the approximating sequence. We saw that in the Hotelling-Downs game and will see other natural examples where this happens when we discuss repeated games.

Exercise 26 *(Hard) (Optional): Consider the following sequence of two-player games. In each game the strategy space may differ. Specifically it is* $0, 1/m, ..., 1$ *in game* m*. It is* $[0, 1]$ *in game "infinity". Assume that* (a_1^m, a_2^m) *is a Nash equilibrium in game* m*, where* a_i^m *might be a mixed or pure strategy and that* $a_i^m \to a_i$ *for* $i = 1, 2$.

1. *Assume that in all the games the payoffs are the same as in the "infinity" game, and that those payoffs are not continuous.*
   (a) *Explain why this does or does not imply that* (a_1, a_2) *is a Nash equilibrium using the discussion regarding upper-hemicontinuity from class.*
   (b) *Now assume that payoffs are as in the Hotelling location game, i.e., if in any game* $a_i < a_j$ *(for* $i, j = 1$ *or* 2*) then* $u_i(a_i, a_j) = \frac{a_1+a_2}{2}$ *and* $u_j(a_i, a_j) = 1 - \frac{a_1+a_2}{2}$, *and if* $a_i = a_j$ *then* $u_i(a_1, a_2) = u_j(a_1, a_2) = \frac{1}{2}$*. Assume also that you know that* $(a_1, a_2) = (a, a)$*, that is, it is a pure strategy symmetric strategy profile. Does this change your answer? Explain.*
2. *Assume that in all the games the payoffs are as in the Cournot competition game studied in class, i.e.,* $u_i(a_1, a_2) = a_1(1 - a_1 - a_2)$.
   (a) *Explain why this does or does not imply that* (a_1, a_2) *is a Nash equilibrium.*
   (b) *Now assume you know that* $(a_1, a_2) = (a, a)$*. Does this change your answer? Explain.*

2.9 Refinements and getting to Nash equilibria

There are two (related) issues regarding the solution concept of Nash equilibrium that have been studied extensively. One is refinements: there are Nash equilibria that seem "implausible" at some intuitive level, and theorists have tried to understand if there are systematic and "plausible" ways to rule out such equilibria. The second is how would we expect player to arrive at a Nash equilibrium—is there some process that can justify the concept. If so, perhaps it selects among equilibria, and as such also provides a refinement.

In general, refinements of Nash equilibria consider what would happen to the equilibrium if the environment were perturbed a little. Would the equilibrium survive? The answer depends on the perturbation. This is an enormous literature: equilibria can be perfect, proper, stable, explicable, hyperstable, divine, universally divine,...

In particular, we can imagine allowing a perturbation of the equilibrium strategy a little. Given such a perturbation, what should we require?

- We could ask that the best-reply functions move us closer (or not move us away). Then the equilibrium would be, in some sense, stable. More generally we could think of other "dynamic" precesses instead of just the best-reply dynamic. (E.g., think of a population, where in each period only a small set of players can change their strategy, or where they don't necessarily move to the best reply but to a better reply, or where they move taking into account that others are moving...). Within this class of questions we could ask whether an equilibrium is globally or locally stable, and whether it is locally stable against any or all possible perturbations.

- Alternatively, we could ask that in the perturbed game there exist an equilibrium close to the given one. If there is such an equilibrium it would be, in a (different) sense, stable (or robust). Again we could ask for this robustness to be satisfied against one or any perturbation.

We will discuss some examples, but not explore this topic in any detail whatsoever.

2.9.1 Dynamic stability

Example 23 *B-o-S revisited*

Consider the mixed-strategy equilibrium in B-o-S. If we perturb the strategies a bit away from the mixed strategy equilibrium then the best replies will move us further away. The equilibrium is "unstable" in this sense. The pure strategy equilibria are locally stable in this sense: any perturbation away will move back.

In general mixed-strategy equilibria cannot be stable when we have simple best-reply dynamics; but one can consider other dynamics under which some mixed-strategy equilibria would, and others would not, be stable.

Example 24 *Cournot revisited*

The Nash equilibrium in the Cournot example is "stable" in this sense.

Exercise 27 *Consider an example like the Cournot example where the best replies are reversed: the horizontal axis is q_2 and the vertical is q_1, and the steeper function is 1's best reply, $q_1(q_2)$, and the function extends horizontally to the right along the axis (so that it describes the best reply for player 1 against any strategy of player 2). The shallower line is $q_2(q_1)$. Identify the Nash equilibria in the diagram, and identify which, if any, are stable.*

2.9.2 Trembling Hand Perfection

Alternatively we could ask that an equilibrium α^* be robust to perturbing the strategies in the sense that α^* is still an equilibrium in a perturbed game where strategies are perturbed in this way.

One motivation and model for this is that whenever a player plays a strategy he might tremble. That is, there is noise in the implementation of strategies. This assumption adds some "caution" into the game.

Definition 13 *A Nash equilibrium α is trembling-hand perfect (THP) (Selten (1975)) if there exists a sequence $\alpha^\varepsilon \to \alpha$, with $\alpha_i^\varepsilon(a_i) > 0$ for all i and $a_i \in A_i$, such that $\alpha_i \in BR\left(\alpha_{-i}^\varepsilon\right)$.*

That is, a Nash equilibrium is THP if each player's strategy is a best reply to a profile of opponents' strategies that are near their equilibrium strategies. That is, if it is a best reply to a small perturbation of their strategies. Notably, Selten also established the existence if THP equilibria.

Note that the perturbation is not correlated and that everyone has the same beliefs about the perturbation. Is this a "plausible" assumption? It is analogous to that of Nash equilibrium, but does it have the same "justifications"?

What are the equilibria when the trembles are small? In two-person games this is just like considering only Nash equilibrium that do not involve weakly dominated strategies (because such strategies cannot be best reply to totally mixed strategies). In games with more than two players there is a difference due to the issue of correlation.

Example 25 *Consider Example* **??**. *The Nash equilibrium* (U, L) *yielding* $(1,1)$ *is THP. To see this consider the sequence where the player plays the Nash equilibrium strategy with probability* $1 - 1/n$ *and the other actions with probability* $1/n$, *with* $n = 2, 3, \ldots$. *Then the NE strategy is a BR to any element in this sequence.*

Example 26 *Consider Example* **??**. *The NE* (M, L) *is thp. To see this consider the sequence where row plays* M *with probabilty* $1 - 2/n$ *and the other two startegies with* $1/n$ *and column plays* L *with* $1 - 1/n$ *and* R *with* $1/n$.

Exercise 28 *In Example* **??** *show that* (M, R) *is THP.*

In Example **??** *observe that* (D, R) *is not THP. (This is a one sentence observation.)*

Exercise 29 (Optional) *Using arguments related to those made (but not always proven) in these notes understand why: (i) a Nash equilibrium in two player games is thp iff it does not involve the use of weakly dominated strategies; (ii) a Nash equilibrium in games with more than two players is thp only if, but not necessarily if, it does not involve the use of weakly dominated strategies.*

One can go further and consider Nash equilibrium that remain after deleting (or even iteratively deleting) weakly dominated strategies, although the justification for that is even flimsier.

Exercise 30 *What are the pure strategy Nash equilibria of the Nash bargaining game that involve weakly dominated strategies?*

Exercise 31 *Consider the following three player game. Player 1 chooses an element of* $\{\alpha, \beta\}$, *2 chooses an element of* $\{A, B\}$ *and 3 chooses an element of* $\{a, b, c\}$. *Player 3 is completely indifferent. Player 1's payoffs and player 2's payoffs depend only on their own action and player 3's actions. So in particular 1 doesn't care what 2 does, and 2 doesn't care what 1 does. The tables below describe 1's and 2's payoffs.*

$1'$*s payoffs*

$1\backslash 3$	a	b	c
α	2	1	0
β	0	1	1

$2'$*s payoffs*

$2\backslash 3$	a	b	c
A	0	1	2
B	1	1	0

1. *Is* (β, B, b) *a Nash equilibrium? Does it involve a weakly dominated strategy for any player?*

2. *Is (β, B, b) a thp equilibrium? Try doing this without reading the following hint, but if you need it then: Consider a sequence converging to 3's equilibrium strategy of b and verify if it can be that B is a best response against that sequence and that β is a best response against that sequence.*

3. *What does this say about the relationship between thp and Nash equilibria that do not use weakly dominated strategies in $n > 2$ player games. Specifically, is it possibly true that if a strategy profile is a Nash equilibrium that doesn't use weakly dominated strategies then is must be a thp equilibrium? How about the converse?*

4. *Really optional (i.e., above the level required for the class): For the one that you only know is "possibly" true, do you think it is true? If not can you come up with a counter example, and if yes can you come up with a proof?*

2.9.3 Evolutionary Stability

Game theory is also applied to the study of evolution. In this context, the players are not assumed to be rational, or even capable of choosing a strategy, rather, they are genetically programmed to play specific strategies. This gives related, but different, notions of dynamic processes and of stable equilibria.

Utility is then interpreted as the player's ability to survive through the expected number of offspringm, which is called *fitness* by biologists. A strategy which results in a larger number of offspring – even by a little will take over the population after enough generations. There are cases where one strategy is simply superior to the other (two eyes), and in other cases it depends on the strategies of the other individuals in the population (aggression).

The simplest model is of a population of one type: there is a very large number of individuals who are randomly matched, in pairs, to play the game. Each player has the same set of possible strategies, and the payment depends only on the strategies, and not on the identity of the specific individual. This kind of model can be represented by a symmetric game.

What characterizes an outcome which would be stable over time?

- Nash equilibrium (mutations are unsuccessful)
- Symmetric (since there is only one population)
- Robust against mutations.

Example 27 *An individual "picks a gender" upon birth (or conception). Denote a choice of "Female" by X and a choice of "male" by Y.*

	X	Y
X	$0,0$	$1,1$
Y	$1,1$	$0,0$

There are no symmetric equilibria in pure strategies. There is an equilibrium in mixed strategies: Every individual is a male with probability of $\frac{1}{2}$ and a female with a probability of $\frac{1}{2}$. Therefore $\frac{1}{2}$ of the population are male and $\frac{1}{2}$ females, and no mutant can gain an advantage. (The study of why genders are 1/2-1/2 and more generally why there is sex is of course much more complex...)

Example 28 *Stag-Hunt or hunting skill together-alone: Suppose that it is possible to develop strong muscles (X) or communication skills (Y). And again, this is decided at birth or conception.*

	X	Y
X	$3,3$	$3,0$
Y	$0,3$	$10,10$

Nash equilibria with pure strategies: X,X or Y,Y.

Note that the Nash equilibrium in mixed strategies is unstable: If 30% of the population have communication skills and a small set of mutants who cooperate enter then they will do better: the mutants will do as well against the rest of the population (who are split 70/30) and strictly better against themselves.

Formally: ESS – Evolutionary Stable Strategies

Definition 14 *α is a ESS if:*

1. *For any $\alpha', u(\alpha,\alpha) \geq u(\alpha',\alpha)$. (Symmetric Nash equilibrium - a mutant has no advantage).*

2. *If for some $\alpha', u(\alpha,\alpha) = u(\alpha',\alpha)$, then $u(\alpha,\alpha') > u(\alpha',\alpha')$. (If a mutant has the same survival probability against the existing population as do other members of the populations, then it has a strictly lower survival probability against another mutant, thus they would become extinct).*

The following provides an equivalent definition.

Proposition 2 *α is an ESS iff there exists $\bar{\varepsilon}$, such that for every $\varepsilon \in (0,\bar{\varepsilon})$, $u(\alpha, \varepsilon\alpha' + (1-\varepsilon)\alpha) > u(\alpha', \varepsilon\alpha' + (1-\varepsilon)\alpha)$.*

In words the proposition says that α is better than α' when a small group of any population α' invades.

One can also consider a dynamic process, where in each period there is a population distribution, and strategies grow according to their success against the distribution and study the stable limits of such a process. This "Darwinian" process, called replicator dynamics, is closely related, but not quite equivalent, to ESS.

A drawback to ESS as a solution concept is that they may not exist.

Exercise 32 *Consider the symmetric version of B-o-S:*

	X	Y
X	$0,0$	$1,3$
Y	$3,1$	$0,0$

Is the mixed-strategy equilibrium an ESS? Replace the $(0,0)$ *in the* (X,X) *cell by* $(2,2)$ *and answer the same question. Replace the* (Y,Y) *cell with* $(2,2)$ *and answer the same question.*

Exercise 33 (Optional) *Find the set of Nash equilibria and ESS in the following two games. Is there a difference in the sense in which the Nash equilibria are stable or unstable in these games? In the matrix I specify only the row player's payoff as the games are symmetric. The left-hand game is a version of rock-papers-scissors.*

	A	B	C
T	1	2	0
M	0	1	2
B	2	0	1

	A	B	C
A	1	1	0
B	0	1	1
C	1	0	1

3 Extensive-Form Games

3.1 Definition of Perfect-Information Extensive-Form Games

Another natural way to describe games is through their "extensive form" or "tree structure." We will begin with a special subclass of such games—games with perfect information in which players move in sequence and know everything that was done before. In general the formal notation for extensive-form games is cumbersome and I will not present it in all detail in class, leaving you to see the details in one of the texts.

Definition 15 *An extensive form game with perfect information has the following components:*

- *A (game)* tree*:*
  - *a set of* nodes*, and links between them*
    - *all but one with a unique* predecessor
    - *and with no* cycles*.*
    - *Some more terminology: The node without a predecessor is called the* root*. Observe that each node may have more than one* successor*. Starting from the root and following an arbitrary successor defines a* path*, the path until some node is the* history *of that node. A node with no successor is a* terminal *node.*
- *A set N of players.*
- *A map from the set of non-terminal nodes to the set of players determining who moves at that node.*
- *A map from the set of paths to $\mathbf{R}^{\mathbf{N}}$, the payoff from that path. (In a finite length tree we can think of this as a map from terminal nodes into $\mathbf{R}^{\mathbf{N}}$).*

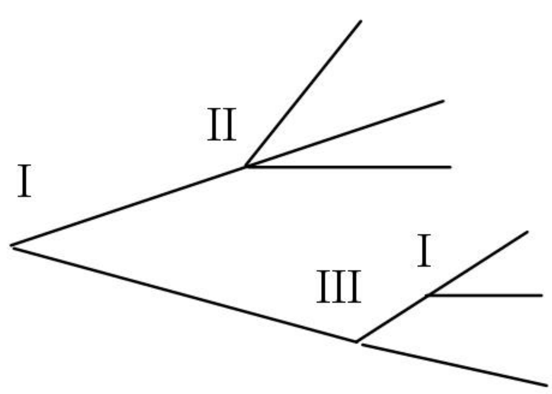

Every perfect information extensive-form game corresponds to a normal form game.[14] A pure strategy of a player is a specification of an action at each node where that player moves. The payoffs of a profile of strategies is the payoff from the path that the profile induces. (A strategy defined in this way involves for each player a specification of actions even at nodes that a player's own earlier action precludes.) We will discuss this transformation in detail later, as well the translation of mixed-strategies, and focus on the extensive form and pure strategies for now. (Later we will see that every normal-form game corresponds to an extensive-form game, but not one of perfect information.)

Why study perfect information games if they can be translated into a normal-form game? (The same reasons will apply to extensive-form games that do not have perfect information.)

- They are often much easier to describe. Consider the two person game where in odd periods player 1 chooses one of two actions and in even periods player 2 chooses one of two actions. In a K period version (where K is even) each player has $2^{K/2}$ strategies; if the game is infinite horizon then while there are only countably many periods there are uncountably many strategies. Describing the game, the strategies and so on is often much easier using extensive-form games.

- The additional information of the tree structure may be useful in selecting among multiple equilibria.

3.2 Backwards induction

The intuitively "obvious" solution for finite-length perfect-information games is backwards induction: start at the end and work backwards. (If length is not finite we have nowhere to start this process.) After understanding this concept we will go back to connect it to the concepts we have studied.

Note that this approach does not say what to do if a player is indifferent. When a player is indifferent we will adopt a tie-breaking rule, for example in a bargaining game we will assume that players accept offers when they are indifferent between rejecting them and accepting them.

Example 29 *Conflict in the Balkans*

[14]This "mapping" is many-to-one. That is, different perfect-information extensive-form games may have the same normal form; we will see this later.

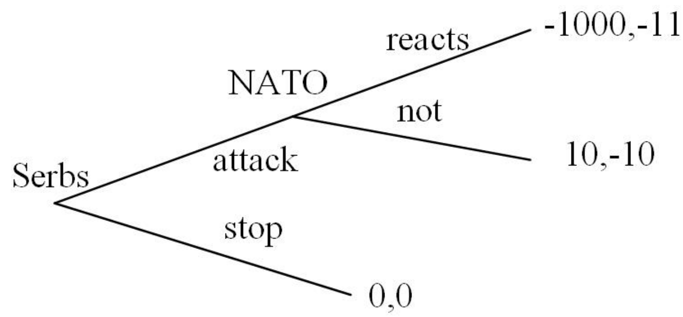

If the Serbs attack, then "obviously" NATO will not react. (NATO threatening to react is said to be "non-credible.")

What can NATO do? For example, it can change its own payoff. If Clinton declares that NATO will react, then if NATO does not react, Clinton will be called a liar and Clinton's payoff could drop from -10 to -12. This would make Clinton's commitment to react credible and would cause the Serbs to stop their attack.

[Show Video Teilspielperfekt from Batman]

Example 30 ***Inflation, or the value of [commitment]***

The game:

Stage 1: Workers request a nominal wage of w.

Stage 2: The central bank sets the inflation rate π.

The real wage is $w_r = \frac{w}{1+\pi}$. The demand for labor is $Q(w_r) = 1 - w_r = 1 - \frac{w}{1+\pi}$. The utility of the workers is their total wage - $u_w = w_r \cdot Q(w_r) = \frac{w}{1+\pi}(1 - \frac{w}{1+\pi})$. The central bank maximizes $U_b = \lambda Q - \frac{\pi^2}{2}$.

Solving by backwards induction:

Given w, the bank solves $\max u_b = \lambda(1 - \frac{w}{1+\pi}) - \frac{\pi^2}{2}$. So $\frac{\lambda w}{(1+\pi)^2} - \pi = 0 => \lambda w = \pi(1+\pi)^2$ $=> \pi = ...$ (an increasing function of λw).

The workers deduce $\pi(\lambda w)$. Their maximum is obtained when $w_r = \frac{1}{2}$, therefore they demand $w = \frac{1}{2}(1 + \pi(\lambda w))$.

This gives both w and π. Note that even without calculating we see $w_r = \frac{1}{2}$, $Q = \frac{1}{2}$, and $\pi > 0$ regardless of λ. So the bank would like to commit to zero inflation because it implies that workers would request a lower nominal wage which would raise demand for labor, but once a wage has been set, it has an incentive to set $\pi > 0$ to increase demand for labor or output. A commitment to zero or low inflation may be achieved by having the bank's objective function described by law as targeting inflation, not output.

Example 31 ***Stackelberg Competition***

Suppose that two firms need to choose their quantities and that their costs of production are zero. Demand is given by $P(q_1, q_2) = 1 - q_1 - q_2$. *Firm* 1 *is the "leader" – it moves first. Firm* 2 *is the "follower" – it moves after firm 1. The game-tree looks as follows:*

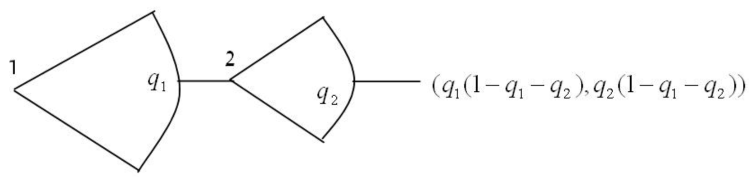

Given that firm 1 *chose* q_1, *what would be optimal for firm* 2 *to do? Because firm 2's payoff function is given by:* $u_2(q_1, q_2) = q_2(1 - q_1 - q_2)$ *it would set its quantity depending on firm 1's quantity as follows:*

$$q_2(q_1) = \frac{1 - q_1}{2}$$

if $q_1 \leq 1$ *and zero otherwise. Firm 1 can anticipate firm 2's behavior and therefore chooses its own quantity to maximize the following objective function*

$$\max_{q_1 \geq 0} u_1(q_1, q_2(q_1)) = q_1 \left(1 - q_1 - \frac{1 - q_1}{2}\right)$$

It would therefore set $q_1^* = \frac{1}{2}$, *which would lead firm 2 to set* $q_2^* = \frac{1}{4}$. *(Recall that under Cournot competition both firms set their quantities equal to* $\frac{1}{3}$.

Example 32 ***Take-it-or-leave-it bargaining (A variation on the famous Ultimatum Game)***

Two players have to divide \$4 between them. Player 1 proposes a division of the money in integer units $(1, 3)$ *or* $(2, 2)$ *or* $(3, 1)$, *where the first amount is the amount for player 1 and the second for 2; assume she is not allowed to propose 0 or 4 for herself. If player 2 agrees, then this division is executed. Otherwise, the money is lost. The game-tree of this game can be described as follows:*

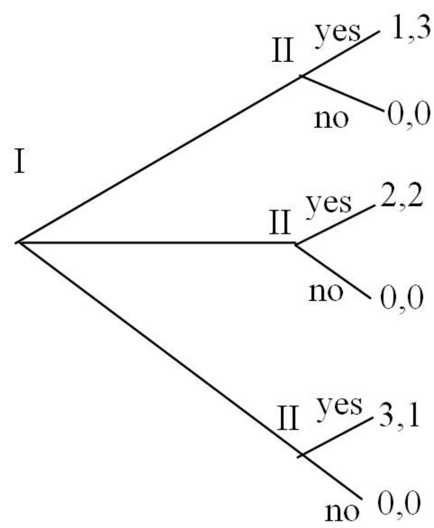

The "backwards induction" solution can be seen as follows. Player 2's threat to refuse a division is [not credible], so that player will 1 get the best possible division for herself.

If we include the endpoints in the above game then backwards induction is no longer pinned down uniquely. Player 1 *could offer 0 expecting to get 4 or could offer* 1 *expecting that an offer of zero results in rejection. Both outcomes survive backwards induction.*

If we consider the continuum game then 2 will accept any positive offer and be indifferent after an offer of zero. In this case it is common to identify the offer of 0 with 2 accepting as the solution since the solution in finite approximations of the game has player 1 getting closer and closer to this outcome. Indeed, when we solve for subgame perfect (SGP) equilibria we will see that the only SGP equilibrium is 1 offering 0 and 2 accepting.

Show video from Kubrick's "Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb" (1964)]

Example 33 ***Iterated (finite) bargaining.***

Consider the game where players alternate in making offers: after an offer the opponent can accept or reject the offer; if they accept the game is over, and if they reject they can make the next offer. Assume that each period (which is made up of an offer and a reply) is discounted by δ. *(Money in the future is worth* δ *less; the item itself is depreciating (ice-cream is melting); the opportunity may disappear with probability* $1-\delta$ *(e.g., the opponent will find an alternative person with whom to interact, or may die...). Assume that individuals accept offers when indifferent (for the same reason that we view an offer of zero in the preceding one-shot game as the backwards-induction outcome).*

Two rounds: In the second round the person who makes the offer gets the whole pie. That pie is worth 1 *then, but* δ *in the preceding round. So in the preceding round that person* *will accept any offer that is at least δ. Thus agreement is reached immediately with an offer of $(1-\delta,\delta)$.*

Exercise 34 *Solve the three and four period versions, and find the solution for general n.*

Exercise 35 *Consider a bargaining game where instead of each period being discounted by δ there is a cost of making an offer to the person making the offer: each time you make an offer you pay some $c < 1$. Solve by backwards induction the one-shot and the two- and three-period games, and find the solution for general n.*

Example 34 ***Chess*** *(Zermello's Theorem)*

The game of chess can be described by a finite tree. Chess is a perfect-information zero-sum game. By applying backward induction we can see that either White has a winning strategy, or Black has a winning strategy, or both White and Black have strategies that ensure a tie.

Exercise 36 *Prove Zermello's Theorem. Hint: use induction on the length of the game.*

This result, which is probably the first result in Game Theory, implies that chess is a "trivial" game in a certain sense, but is it?

Exercise 37 *item The game of bi-chess has the same rules as ordinary chess, except that each player can move twice when it is its turn to play. That is, White begins and plays twice, then Black plays twice, and so on. Show that White can at least ensure a draw in the game of bi-chess (namely, Black does not have a winning strategy).*

Exercise 38 *Consider the following two player board game. There is a rectangular board with $n \times m$ squares. The two players alternate their turns. In its turn, each player chooses a square from among those that have not been eliminated yet, and all the squares above and to the rigut of this chosen square are eliminated. The player who chooses and eliminates the last remaining square loses the game.*

1. *Find the winning strategy of the first player when the board is square (i.e., $n \times n$).*
2. *Prove that the first player can ensure it wins also when the board is rectangular (hint: it is possible to prove that the first player has a winning strategy without explicitly describing its strategy).*

Example 35 ***The Centipede Game*** *(Rosenthal, 1986)*

The two-legged centipede (duopod):

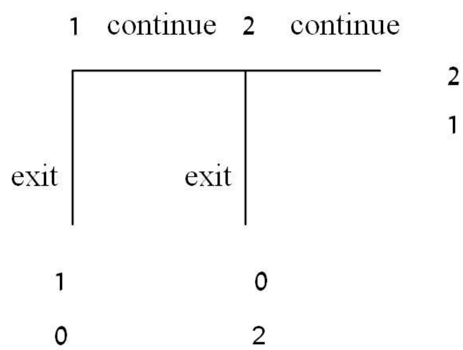

The four-legged centipede (quadripod)

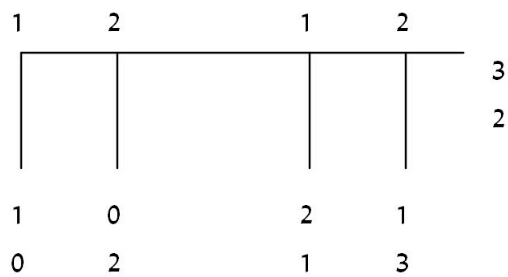

The milion-legged centipede (mega-pod?)

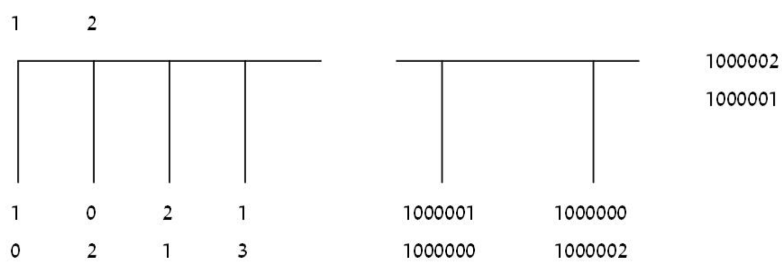

Backwards induction implies that the first to move will quit and payoffs will be $(1,0)$. Is this a good prediction? If units are in ten-thousand dollars, perhaps, if in one-million dollar units probably, if in one-billion definitely. In any case this shows (again) the limitations of a certain solution concept. However, in contrast to criticisms of rationalizability it is worth pointing out that backwards induction does *not* formally correspond to common knowledge of rationality. It is consistent with common knowledge of rationality for players not to exit immediately. As we will see this is related to deleting weakly dominated strategies about which we already observed that iterative deletion does not correspond to common knowledge of rationality (which only involves iterative deletion of [strongly] dominated strategies).

Exercise 39 (Double Marginalization) *Demand for your tea is determined by the following function: $Q = 1 - p$, where p is the price you charge. The cost of making a bottle of tea is $c = \$.10$.*

1. *If you could sell your tea directly to the public what price would you charge?*

2. *However, thats not how the world works. You have to sell your tea via a retailer. The retailer has a monopoly and faces the same demand curve that you do. Thus the retailer chooses a price p_r to maximize its profits:*

   Retailer profits $= (p_r - p_w) \cdot (1 - p_r)$.

   where p_w is the wholesale price that the retailer pays you for its tea.

   Note that we are assuming here that the retailer doesnt have any costs other than the wholesale cost of buying from you.

   Figure out what the retailer will charge as a function of the wholesale price and then using that formula figure out what is the profit-maximizing wholesale price you should charge. Note that your profits are $(p_w - 0.10) \cdot (1 - p_r)$. Who makes more, you or the retailer ?

Exercise 40 (The Rotten Kid Theorem)

1. *Suppose a parent and child play the following game, first analyzed by Becker (1974). First, the child takes an action, A, that produces income for the child, $I_C(A)$, and income for the parent, $I_P(A)$. (Think of $I_C(A)$ as the child's income net of any costs of the action A.) Second, the parent observes the incomes I_C and I_P and then chooses a bequest, B, to leave to the child. The child's payoff is $U(I_C + B)$; the parent's is $V(I_P - B) + kU(I_C + B)$, where $k > 0$ reflects the parent's concern for the child's well-being. Assume that: the action is a nonnegative number, $A \geq 0$; the income functions $I_C(A)$ and $I_P(A)$ are strictly concave and are maximized at $A_C > 0$ and $A_P > 0$, respectively; the bequest B can be positive or negative; and the utility functions U and V are increasing and strictly concave. Prove the "Rotten Kid" Theorem: in the backwards- induction outcome, the child chooses the action that maximizes the family's aggregate income, $I_C(A) + I_P(A)$, even though only the parent's payoff exhibits altruism.*

2. *Now suppose the parent and child play a different game, first analyzed by Buchanan (1975). Let the incomes I_C and I_P be fixed exogenously. First, the child decides how much of the income I_C to save (S) for the future, consuming the rest ($I_C - S$) today. Second, the parent observes the child's choice of S and chooses a bequest, B. The child's payoff is the sum of current and future utilities: $U_1(I_C - S) + U_2(S + B)$. The parent's*

payoff is $V(l_C - B) + k(U_1(I_C - S) + U_2(S + B))$. Assume that the utility functions U_1, U_2, and V are increasing and strictly concave. Show that there is a "Samaritan's Dilemma": in the backwards-induction outcome, the child saves too little, so as to induce the parent to leave a larger bequest (i.e., both the parent's and child's payoffs could be increased if S were suitably larger and B suitably smaller).

3.3 The Relationship between Perfect-Information Extensive and Strategic Form Games and between Nash equilibria and BI outcomes

As noted every extensive form game can be represented as a strategic-form game. The set of players is identical. The set of pure strategies is identical. For any vector of strategies, payoffs in the normal form game matrix are the expected payoffs obtained in the extensive form game when the vector of strategies is played.

Example 36 ***The Balkan game revisited*** *(example* **??***).*

	R	N
A	$-1000, -11$	$10, -10$
S	$0, 0$	$0, 0$

Note that the BI outcome corresponds to deleting the weakly dominated strategy R and then deleting the strongly dominated strategy S.

Exercise 41 *Consider the game*

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$0, 0$

Find two different perfect-information games with this (common) strategic form.

Exercise 42 *Specify the strategic-form game for the two-legged centipede of example* **??***.*

What is the size of the matrix for the strategic-form description of the integer version of the T-o-L game (of example **??***).*

A Nash equilibrium in a perfect-information extensive form game is defined exactly as is a Nash equilibrium in a strategic form game: a profile of strategies such that no player has a strictly positive gain from deviating. Of course, the set of Nash equilibria of a perfect-information extensive form game is the set of Nash equilibria of the strategic-form game corresponding to that extensive-form game.

Example 37 ***Balkan revisited*** *(examples* **??** *and* **??***).*

In the Balkan game there is a second Nash equilibrium: (S, R)*. This Nash equilibrium is not a backwards induction outcome, and it involves use of a weakly dominated strategy.*

Proposition 3 *The BI strategies constitute a Nash equilibrium of the game. (In generic perfect-information games—including perfect-information games where all the payoffs are distinct—the BI outcome is the same as that obtained by iteratively deleting weakly dominated strategies in a suitable order.)*

Example 38 ***T-o-L (Ultimatum) revisited*** *(example* **??***)*

What is the set of Nash equilibria for the T-o-L game (integer and continuum)? Any initial offer $(x, 1-x)$ *with the strategy of the replier being to reject any offer less than* $1-x$ *and either accept or reject any offer above* $1-x$*.*

Remark 6 *A similar game has been played in experiments in the US, Israel, Slovenia and Russia. The results were such that player 1 offered a division around 70% - 30%. Most offers were around 70% and most responders rejected offers lower than 30%. Observe that this is a Nash equilibrium outcome of the game but not a backward induction outcome of the game.*

Exercise 43 *What* outcomes *are possible in a Nash equilibrium of the various centipede games of example* **??***? Are the* strategies *in all the Nash equilibria the same as under BI ?*

3.4 Subgame Perfect Equilibrium

In infinite-length perfect-information games we cannot apply backwards induction. A subgame perfect equilibrium (SPE or SGP equilibrium) is a refinement of Nash equilibrium in extensive form games that generalizes the idea of backward induction. It generalizes BI to infinite games of perfect information as we'll see here, and to a larger class of extensive-form games as we'll see next.

A **subgame** of a perfect-information game is obtained by considering the specification of nodes, players and payoffs after any specific node in the tree. That is, we define a node in the tree to be the root, ignore anything before that node, and thereby have a new game that is a *subgame* of the original game.

Definition 16 *A subgame perfect equilibrium is a profile of strategies that is a Nash equilibrium in all subgames of the game (including the original game).*

Of course, when we consider whether the strategy profile is a Nash equilibrium in the subgame, we ignore the specification of the strategy preceding the subgame.

Example 39 *Show SPE and Nash in a simple "threat game." Explain that SPE is a refinement of Nash equilibrium that rules out the use of "non-credible threat". Show kidnapping video.*

Example 40 *In the four-legged centipede the strategy profile ((exit,continue),(exit,exit)) gives the Backwards induction outcome and is a Nash equilibrium but is not a subgame perfect equilibrium. In the first and last subgames the profile specifies a Nash equilibrium. But in the second and third subgames player the at least one of the players is not playing a best reply to the other. Note that if we considered player* 1*'s overall strategy, then since* 1 *is exiting at the start, there is no benefit to changing the action at his second decision node. But focusing on the specification of the strategies conditional on being in the subgame we see that ((continue),(exit)) is not a Nash equilibrium.*

Exercise 44 *Prove* carefully *that in any finite or infinite horizon centipede game there is no pure strategy subgame perfect equilibrium other than "always exit". (In an infinite horizon version the first player specifies a subset of the odd numbers which are the nodes at which she exits, the second specifies a subset of the even numbers which are the nodes at which he exits, and we need to specify the payoffs if both never exit. In that case assume they both get 0.)*

Exercise 45 ***Strategic exit in a declining market***

Consider the following duopoly environment. The market is decreasing in size over time, so that in period t monopoly profits for firm 1 *are* $510 - 25t$ *and for firm* 2 *they are* $51 - 2t$, *while duopoly profits are* $105 - 10t$ *and* $10.5 - t$ *respectively. Each period has two steps: first one firm chooses whether to exit and then (after observing the choice of the first) the other firm chooses whether to exit. Exit is irrevocable: after exiting a firm cannot reenter. After both decisions are made the payoffs for that period are obtained: 0 in that period for any firm that exits, monopoly profits for period t if a firm is alone in that period, and duopoly profits in that period if both firms have not yet exited. Find the subgame perfect equilibria for the case where firm 1 goes first in each period and for the case where firm 2 goes first in each period.*

Remark 7 ***Timing games****: The preceding game is an example of a general class of games called timing games in which players decide when to exit / enter. Such games can be used to model R&D, wars of attrition, dynamic auctions, entry and (as above) exit, and more.*

3.5 The one-stage deviation principle

A strategy profile $s = (s_1, ..., s_n)$ is resistant against a one stage deviation, if there is no subgame in which a player can gain by a single deviation.

Proposition 4 *In every finite-length game, and in every infinite-length game with bounded and discounted payoffs, profile s is a SPE if and only if it is resistant against a one stage deviation.*

That is, if something is not a subgame perfect equilibrium then there must be a subgame and a player who can gain by changing her strategy at *one* node in that subgame.

Proof. (Sketch) Assume first the game is of finite length. If s is not a subgame perfect equilibrium then there is a subgame G and a player i who has a deviation s_i' in G that is profitable. If s_i' differs from s_i in only one period then we are done. If not, consider the subgame G' starting from the last period in which the deviation differs from s_i. If that deviation alone is not profitable in G', then it is redundant in the sense that we can change that action back to the action specified by s_i and still have a profitable deviation in G. If it is, then we have found a subgame with a profitable one stage deviation. Continuing in this manner we can change s_i' back to s_i backwards from the last place where it differs.

If the game is discounted and infinite the same idea works as follows. If there is a profitable deviation, the gain is at least some ε. The total possible present value of the difference in payoffs from some T onwards is bounded by $\sum_{t=T}^{\infty} \delta^t 2M$ (where M is the maximal absolute value of payoffs in any period), so for some T this is less than $\varepsilon/2$. So consider the game until T where there is a gain of at least $\varepsilon/2$. This shows that if there exists a profitable deviation, then there exists a profitable deviation in the game up to time T. To finish the proof, apply the preceding argument for finite games. ■

Example 41 ***The importance of discounting*** *Consider an infinite single person game, where in each period t the person has to choose u_t equal to 1 or 0, and the payoffs are given by the limit of the average payoffs: $\lim_{T\to\infty} \sum_{t=1}^{T} u_t/T$. So if the person chooses in even period 1 and in odd periods 0 the person's payoff is $1/2$. Clearly there is a profitable deviation: always choose 1. But in any subgame a single change of 0 to 1 will not effect the limit. Now assume that the payoffs are given by discounting $\sum u_t \delta^t$, where $\delta < 1$. Consider any subgame, and observe that changing from 0 to 1 in any one period will yield a profit.*

3.6 Rubinstein's Model of Alternating Offer Bargaining [15]

Suppose that two players are bargaining over a "pie" of size one. Bargaining is an infinite version of the bargaining game discussed in example **??**.

The game tree can be drawn as follows:

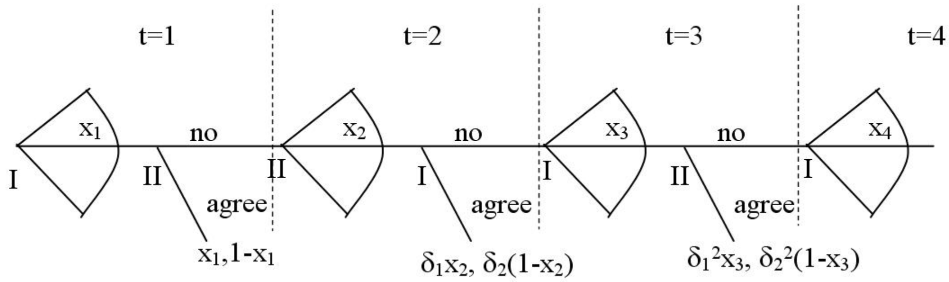

For simplicity, we denote both Player 1's and Player 2's offers in terms of Player 1's share of the pie. We assume that Players 1 and 2 discount their future payoffs according to the discount factors δ_1 and δ_2 respectively. We also assume that if the players keep on rejecting each other's offer forever, then they both obtain a payoff of zero.

This game has many Nash equilibria. For example, for Player 1 to always demand x and to refuse any offer of less than x, and for Player 2 to always demand $1-x$ and refuse any offer less than $1-x$ is a Nash equilibrium of the game for any $x \in [0,1]$. But are these Nash equilibria also subgame perfect equilibria? No. Fix such an equilibrium where Player 1 offers x. Consider the subgame that begins after Player 1 offered some $x' > x$ in the first period. If Player 2 refuses Player 1's offer, then, given Player 1's strategy, Player 2's payoff in the game cannot be higher than $\delta_2(1-x)$. Therefore if Player 1 demands more than x but less than

$$\begin{aligned} 1-\delta_2(1-x) &= 1-\delta_2+\delta_2 x \\ &> (1-\delta_2)\,x+\delta_2 x = x \end{aligned}$$

in the first period, then Player 2's best reply is to accept because it would give her a payoff strictly larger than $\delta_2(1-x)$.

Proposition 5 ***(Rubinstein, 1982)*** *The alternating-offer game above has a unique subgame-perfect equilibrium. In this equilibrium, Player 1 always offers $\bar{x}$, and accepts any offer that*

[15]This game is also known as the Rubinstein-Stähl alternating offer bargaining game. Stähl (1972) described this game in the early 70s, but did not identify or establish the existence of a unique subgame perfect equilibrium for the game.

gives him at least $\bar{y}$ and Player 2 always offers $\bar{y}$ and accepts any offer that gives her at least $\bar{x}$ where

$$\bar{y} = \delta_1 \bar{x} \qquad \text{and} \qquad 1 - \bar{x} = \delta_2 (1 - \bar{y})$$

Thus, in equilibrium, the game ends with agreement on $\bar{x}$ in round 1.[16]

Proof. First, notice that the game is stationary: the sub-game which starts in period $3, 5...$ is a (discounted) replica of the entire game. The sub-game which start in periods $2, 4...$ is a replica of the entire game, only with the roles of the players reversed.

Assume that there is at least one SPE (we will later prove that). We will see that there is only one possibility of payments in equilibrium.

Let us look at the sub-game which starts in period 3. Let $\bar{v}_1^3$ be the highest (supremum) payment that player 1 can obtain in a SPE (in terms of period 3, not discounted to a previous period). Therefore, in period 2 player 1 will accept any offer above $\delta_1 \bar{v}_1^3$ so player 2 will get at least $\underline{v}_2^2 \geq 1 - \delta_1 \bar{v}_1^3$. In period 1, the highest payment that player 1 can demand is $1 - \delta_2 \underline{v}_2^2$, otherwise player 2 will refuse, and so $\bar{v}_1^1 \leq 1 - \delta_2 \underline{v}_2^2 \leq 1 - \delta_2 (1 - \delta_1 \bar{v}_1^3)$. However, the sub-game which starts in period 3 (with payments in terms of period 3) is identical to the subgame starting in period 1. Therefore $\bar{v}_1^1 \equiv \bar{v}_1^3 = \bar{v}_1$. Plugging that in the preceding inequality will yield the following: $\bar{v}_1 \leq \frac{1-\delta_2}{1-\delta_1\delta_2}$.

Similarly, let $\underline{v}_1^3$ denote the lowest payment player 1 can obtain in period 3 in a SPE (in terms of period 3). In period 2, player 2 will not demand more than $\bar{v}_2^2 = 1 - \delta_1 \underline{v}_1^3$ (in terms of period 2) - otherwise player 1 will refuse. Therefore, in period 1, player 1 will demand no less than $\underline{v}_1^1 \geq 1 - \delta_2 \bar{v}_2^2 \geq 1 - \delta_2 (1 - \delta_1 \underline{v}_1^3)$. However, $\underline{v}_1^1 \equiv \underline{v}_1^3 = \underline{v}_1$, and so $\underline{v}_1 \geq \frac{1-\delta_2}{1-\delta_1\delta_2}$.

Hence, we reach the conclusion that there is only one payment which player 1 can obtain in a SPE: $v_1 = \underline{v}_1 = \bar{v}_1 = \frac{1-\delta_2}{1-\delta_1\delta_2}$.

Similarly, we can see that $v_2 = \frac{1-\delta_1}{1-\delta_1\delta_2}$ - this is the payment that player 2 will get in terms of period 2.

So, whenever it is 1's turn to propose, it must offer a partition $(v_1, 1 - v_1) = \left(\frac{1-\delta_2}{1-\delta_1\delta_2}, \delta_2 \frac{1-\delta_1}{1-\delta_1\delta_2} \right)$ and whenever it is its turn to respond 1 must accept any offer larger than or equal to $(1 - v_2, v_2)$.

If player 2 refuses in period 1 he gets v_2 in terms of period 2, which is equivalent to $\delta_2 \frac{1-\delta_1}{1-\delta_1\delta_2}$ in period 1. In period 1 player 1 can therefore get at least $1 - \delta_2 \frac{1-\delta_1}{1-\delta_1\delta_2} = \frac{1-\delta_2}{1-\delta_1\delta_2}$ as 2 would accept such an offer, and 1 cannot get more, as 2 would reject anything worse and 1 would then end up with less.

Similarly, 1 cannot do better than accept 2's offer whenever possible.

[16] Solving for $\bar{x}$ and $\bar{y}$, we see that $\bar{x} = \frac{1-\delta_2}{1-\delta_1\delta_2}$ and $\bar{y} = \frac{\delta_1 - \delta_1\delta_2}{1-\delta_1\delta_2}$.

What about existence? Notice that the following strategies are a SPE: player 1 always offers a partition v_1, $1 - v_1$, player 2 always offers $1 - v_2$, v_2, and player 1 agrees to offers of at least $1 - v_2$, and player 2 agrees to offers of at least $1 - v_1$, where $1 - v_1 = \delta_2 v_2$ and $1 - v_2 = \delta_1 v_1$.

It can be easily verified that no player, in any subgame, can make a one-stage profitable deviation from that strategy. Therefore, according to the one-stage-deviation principal, this is a SPE. ■

Conclusions:

- Show symmetry between 1 and 2. They both make the same partition when it is their turn to propose.
- If $\delta_1 \to 1$ (1 is very patient) we get a partition of $(1, 0)$.
- If $\delta_2 \to 0$ (2 is very impatient) we get a partition of $(1, 0)$.
- If $\delta_1 = \delta_2 = \delta$, we get a partition of $\frac{1}{1+\delta}$, $\frac{\delta}{1+\delta}$ => Player 1 has an advantage.
- This advantage goes to 0 as $\delta \to 1$ as then the partition converges to $(1/2, 1/2)$.
- Because the players divide the pie immediately, there is no waste (an efficient solution). This led to much research on what might cause delays in bargaining; the answer is that in certain environments if there is incomplete information about the players' preferences then delay can arise. But there must be some change in the environment from the simple one here. Various other extensions have been investigated, including the presence of outside options, strikes, identity of proposer/responder, more than two parties, etc.

What happens if offers are given at a faster rate?

Instead of a period of length 1, let us look now at a period of length $1/k$. What are the discounting factors of a period of length $1/k$? They are $\delta_2^{1/k}$, $\delta_1^{1/k}$ (because after k periods of length $1/k$ we will get $(\delta_i^{1/k})^k = \delta_i$. Therefore, the partition is: $\frac{1-\delta_2^{1/k}}{1-\delta_1^{1/k}\delta_2^{1/k}}$, $1 - \frac{\delta_2^{1/k}}{1-\delta_1^{1/k}\delta_2^{1/k}}$, which, as $k \to \infty$ converges to $\frac{(-\log \delta_2)}{(-\log \delta_1)+(-\log \delta_2)}$, $\frac{(-\log \delta_1)}{(-\log \delta_1)+(-\log \delta_2)}$ (by L'Hopital's Rule).

- If $\delta_1 = \delta_2$ the partition will be $(\frac{1}{2}, \frac{1}{2})$: The first player no longer has an advantage.
- A more patient player will get a larger share of the pie: $1 > \delta_1 > \delta_2 > 0 \Rightarrow -\log \delta_1 < -\log \delta_2$

Exercise 46 *Consider the three-person bargaining game where in periods* $1, 4, 7, \ldots$ *player 1 makes offers, in periods* $2, 5, 8, \ldots$ *player 2 makes offers, and in periods* $3, 6, 9, \ldots$ *player 3 makes offers. An offer is a three-way split. After an offer in period* t *is made the other players say in turn (say with the one whose number is lower going first) whether they accept or reject. If they both accept the game is over with that split, otherwise the next period,* $t+1$*, begins.*

Show that the split $\left(\frac{1}{1+\delta+\delta^2}, \frac{\delta}{1+\delta+\delta^2}, \frac{\delta^2}{1+\delta+\delta^2}\right)$ *is a subgame-perfect-equilibrium payoff. (Not very difficult.) Hint: use strategies similar to those proposed in the two-player game.*

Optional and difficult: Show that any split is a subgame-perfect-equilibrium payoff.

Exercise 47 *Consider an alternating offers bargaining game like Rubinstein's bargaining game but with* $n = 1000$ *periods. In odd periods, player 1 makes an offer and player 2 can either accept or reject; and in even periods, player 2 makes an offer and player 1 can either accept or reject. In the last period, if player 1 refuses Player 2's offer, then they both get a payoff of zero. Assume that* $\delta_1 < 1$ *and* $\delta_2 = (1 + \delta_1)/2$*. What does the subgame perfect equilibrium outcome converge to as* $\delta_1 \nearrow 1$*?*

3.7 Mixed strategies in extensive-form games

One way to define mixed strategies is exactly as in the strategic form games. The player mixes over strategies where a strategy defines an action at each node. But in the extensive-form one can think of an alternative method of defining mixed strategies: the player can mix at each decision node. This is called a behavior strategy.

Definition 17 *A pure strategy of player* i *specifies an action at each decision node of* i*. A mixed strategy is a probability distribution over pure strategies. A behavioral strategy is a probability distribution over actions for each decision node of* i*.*

Proposition 6 *(Kuhn) Every mixed strategy is equivalent to a behavior strategy in the sense that whatever pure (or mixed or behavior) strategy the opponents choose, the behavior strategy and the mixed strategy yield the same distribution over branches in the tree. This mapping from mixed to behavior strategies is many to one and onto, so every behavior strategy is also equivalent to at least one mixed strategy and possibly more.*

Example 42 *Consider the extensive-form game where player I chooses* U *or* D *and after* U *player II chooses* L *or* R *and after* D *player II chooses* A *or* B*. Thus player II has four pure strategies, LA, LB, RA, RB, and the mixed strategies are probabilities distributions over those four,* $q = (q_{LA}, q_{LB}, q_{RA}, q_{RB})$*. A behavior strategy is a pair of probabilities,* *one over* $\{L, R\}$ *and one over* $\{A, B\}$, $(p_L, p_R), (r_A, r_B)$. *To find the behavior strategy corresponding to a mixed strategy* q *we just find the probabilities of* L *vs.* R *and of* A *vs.* B: $p_L = q_{LA} + q_{LB}$, $p_R = q_{RA} + q_{RB}$, $r_A = q_{LA} + q_{RA}$ *and* $r_B = q_{LB} + q_{RB}$. *For example consider* $q = (1/15, 4/15, 1/3, 1/3)$; *this gives* $p = (1/3, 2/3)$ *and* $r = (2/5, 3/5)$. *Another mixed strategy* q *with the same behavioral strategy* q, r *can be seen by simply multiplying out* p *and* r; $\hat{q} = (2/15, 3/15, 4/15, 2/5)$ *(which when we add up gives the same* p *and* r*).*

What is the difference between two distinct mixed strategies that have the same behavioral strategies? They allow for "redundant" correlation: you can choose a different probability between A and B depending on whether you choose L or R.

3.8 Definition of Imperfect-Information Extensive-Form Games

A perfect information game describes strategic environment where players move in sequence and when a player chooses an action at a node she knows all actions chosen by players at all nodes that precede her decision node. Naturally there are strategic situations some players have information about choices by other players, but not as much information as is contained in the description of a perfect-information game. How can we describe such games? Using information sets: a player cannot distinguish between different nodes in a given information set.

In describing a game of imperfect information we add to our previous description of an extensive-form game a specification of information sets of a player. A set of nodes may belong to the same information set of a player if the number of successor nodes is the same (otherwise the player can tell the nodes apart by the choices available). The game has perfect recall if the player cannot distinguish the nodes based on information earlier in the tree. If she can, then the fact that the nodes belong to the same information set implies that the player cannot distinguish between them, which means she must have imperfect recall. Because if recall was perfect, these nodes would not have belonged to the same information set. (For a formal definition see one of the texts.) (Mention the absent-minded driver's paradox.)

The following is an example of an imperfect-information game with perfect recall.

Example 43 *If we changed player III to II that would have imperfect recall as II knew before whether she moved after I moved to the left or to either the middle or right node. If we changed III to I then that would have imperfect recall as I knows whether he moved right or middle.*

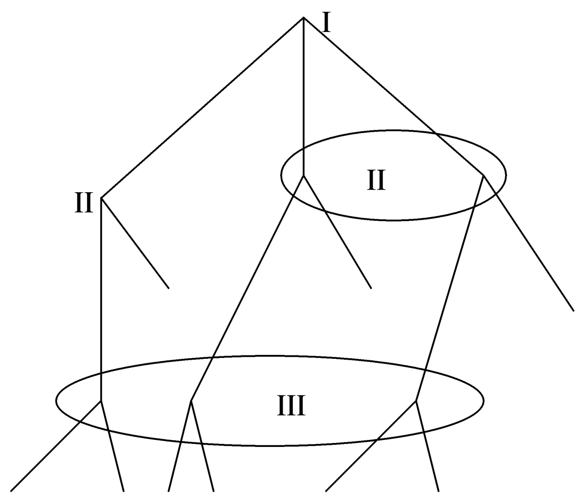

We saw before that every perfect-information game can be described as a strategic-form game. Now we have a converse: every strategic-form game can be described as an extensive-from game with perfect recall. In general there may be many imperfect-information games that describe a strategic-form game.

Example 44 *The extensive-form game corresponding to matching pennies or the prisoners' dilemma or the battle of the sexes has one player at a root node with two actions, and the second player has an information set consisting of the two successor nodes with two actions at each node.*

Exercise 48 *Optional: Make sure you see how to write down (including payoffs) the imperfect-information extensive-form game corresponding to each of the three games mentioned in the preceding example.*

Make sure you understand how to write down the strategic form of the tree in example **??** *(add payoffs as symbols: $a, b, c...$).*

Strategies (pure and mixed and behavioral) in an imperfect-information game are defined as in perfect-information games; the only difference is that an action is now a decision at an information set, not at a node. The relationship between mixed and behavioral strategies is unchanged.

3.9 Nash and SPE

A *Nash equilibrium* of an imperfect-information game is defined exactly as it was in the perfect-information and strategic form games: a profile of strategies such that no player can gain by changing strategies.

A *subgame-perfect* equilibrium is defined exactly as it was in the case of perfect-information games: a profile of strategies that specifies a Nash equilibrium in every subgame. (Note that it is easier to see what any profile strategies define as strategies in the subgame using behavioral rather than mixed strategies.)

However, we need to define a *subgame* with more care now. Intuitively, the first condition is that a subgame can start at any node that is a singleton information set, so you cannot "split" an information set when starting the game. Moreover, once you have found a node at which to start a subgame you cannot split an information set that follows. That is, consider any successor node to the new root node you have chosen, and the information set that contains that successor node. All the nodes in that information set must be successors to the root you have chosen.

Example 45 ***Example ?? continued.*** *In this tree you cannot start a subgame at the two nodes that comprise II's right information set since that would split II's information set. You cannot start a subgame at the singleton information set of player II on the left since that would split III's information set.*

A problem with subgame perfection is that it can depend on "irrelevant" changes in the game tree.

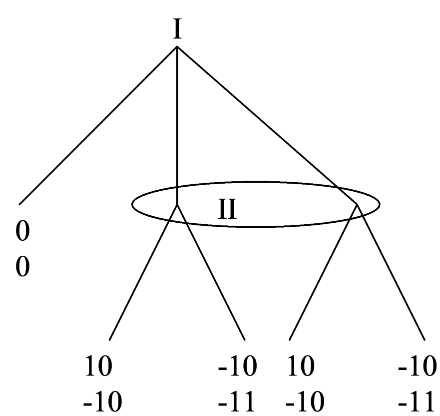

In the tree above there are indeed no subgames, so the Nash equilibrium where I moves L and II moves r is also subgame perfect. Contrast this with the game where II has only one node. Even though we would presumably not want our predictions to change between the two games, once there is one node SGP rules out (L, r). This is also strategically the same game as the game below in which the SGP equilibrium has I going either M or R and II playing l.

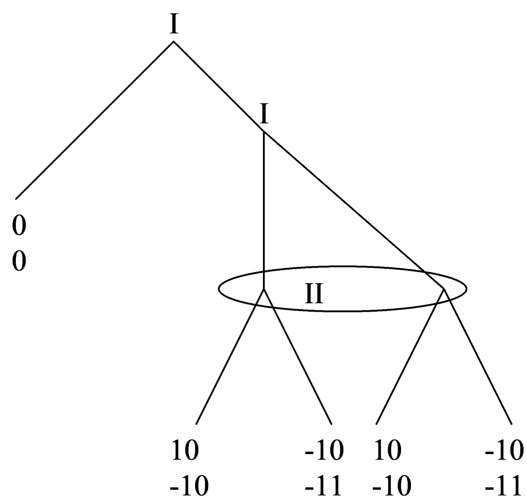

There are solution concepts for games that deal with this, but we do not have time to discuss them in this course. We will continue to work with SPE since it is strong enough for our current purposes.

Definition 18 ***(Terminology, really)*** *Given a profile of strategies in a dynamic game we can distinguish between the information sets and subgames that are* ***on path*** *— i.e., that occur with positive probability according to the profile — and* ***off path****. If the profile is an equilibrium then we distinguish between information sets* ***on the equilibrium path*** *and* ***off the equilibrium path****.*

Example 46 *In the above tree in a profile in which player I plays L then everything other than the left most branch is off path.*

In generic games with perfect information we saw that SGP equilibria coincide with the backwards induction solution and with the solution of iteratively deleting weakly dominated strategies. In games with imperfect information there is no simple relationship between SGP equilibria and iteratively deleting weakly dominated strategies (or backwards induction).

Exercise 49 *Draw the game tree (including information sets) for the following situation. Player 1 chooses whether to invest a lot or a little. After 1 chooses, 2 observes 1's choice and chooses between investing a lot or a little. After player 2 chooses it is player 3's turn to chose between one of two actions (say invest or quit), where player 3 observes 2's action, but player 3 does not observe player 1's actions. (No need to specify payoffs.)*

Exercise 50 *For games (a) through (d) on the following pages, find the normal form and the set of (pure) Nash equilibria. See if the reduced normal form is dominance solvable and, if so, give the solution. Also, for (a), solve by backward induction.*

a)

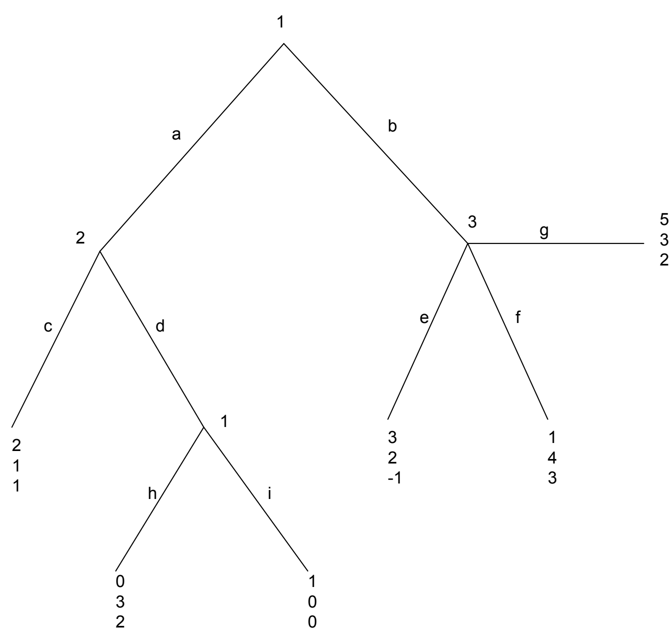

b)

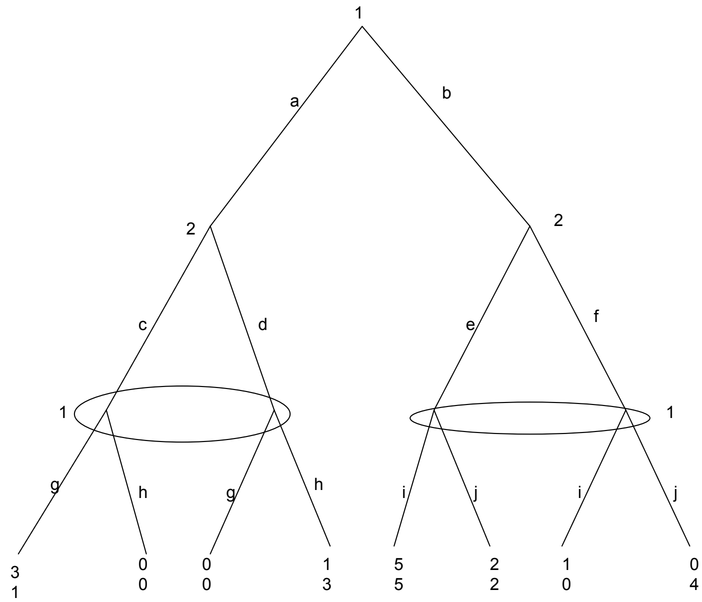
1
a
b
2
2
c
d
e
f
1
1
g
h
g
h
i
j
i
j
3 1
0 0
0 0
1 3
5 5
2 2
1 0
0 4

c)

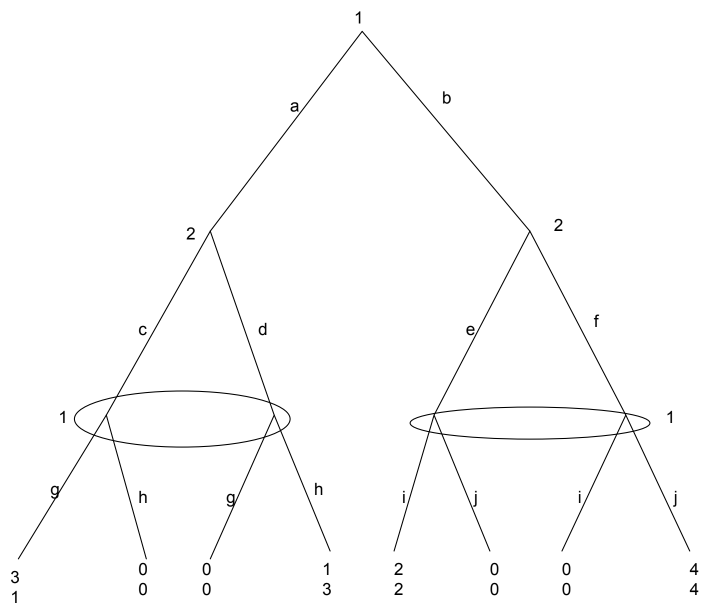
1
a
b
2
2
c
d
e
f
1
1
g
h
g
h
i
j
i
j
3 1
0 0
0 0
1 3
2 2
0 0
0 0
4 4

d)

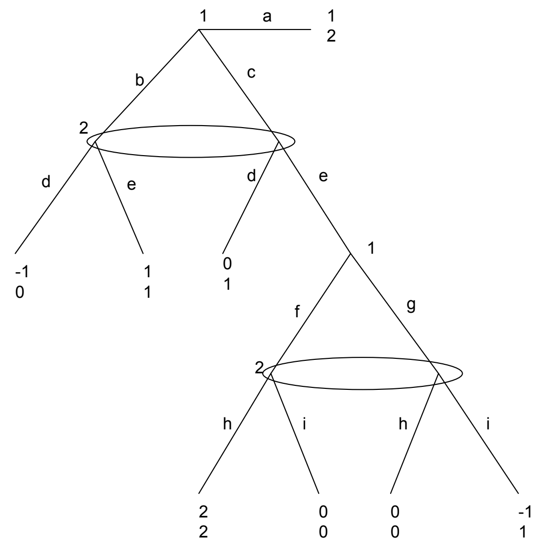

Exercise 51 ***A unique SGP equilibrium with no weakly dominated strategies.*** *Consider the game where player 1 can play* out *and get* $1/2$ *or play* in *which leads to a game of matching pennies. Subgame perfection implies that the payoff of the subgame after* in *is zero, so in the unique SGP equilibrium player 1 goes out. In the strategic form of this game there are no weakly dominated strategies (check !).*

Example 47 ***A unique solution given by iteratively deleting weakly dominated strategies but multiple SGP equilibria: burning money***

Consider the game where player 1 first chooses between the following matrices, and then, after 1's choice is observed they play the corresponding matrix. Note that the right-hand matrix is equivalent to the left-hand one except that player 1 (the row player) has "burned" some money/utility.

	A	B
U	$9,6$	$0,0$
D	$7,7$	$6,9$

	A'	B'
U'	$7.5,6$	$-1.5,0$
D'	$5.5,7$	$4.5,9$

Note that there are several SGP equilibrium paths in this game: 1 plays R followed by (U',A') *(backed by play of* (D,B) *after L); 1 plays L followed by* (D,B) *(backed by play of* (D',B') *after R) and 1 plays L followed by* (U,A) *(backed by play of* (U',A') *after R).*

The following analysis corresponds exactly to iteratively deleting weakly and strongly dominated strategies in the strategic form. (Make sure you see this.) The strategies R, U, D' *and* R, D, D' *give at most* 5.5 *while* L, D, U' *and* L, D, D' *both give at least* 6 *so strongly dominate the former. Having thus deleted* D' *in the* R *game, the column player will choose* A' *after she observes* 1 *selecting* R. *(This corresponds to* A, B' *and* B, B' *being weakly dominated in the game that remains after the first deletions by* A, A' *and* B, A' *respectively.) Note that in the game that remains* 1 *can guarantee a payoff of* 7.5 *by playing* R, U, U' *or* R, D, U'. *So at this stage* L, D, U' *is strongly dominated (by* R, U, U'*). Thus, if* 2 *observes* 1 *playing the left game, since the only strategy remaining for* 1 *after* L *is to play* U, 2 *does not play* B. *(That is,* B, A' *and* B, B' *are weakly dominated.) Thus* 1 *obtains* 9 *by playing* L, U, U' *while by playing* R, U, U' *player* 1 *will obtain only* 7.5. *Thus the option of burning money is never used and enables* 1 *to get his preferred outcome! Does this deletion process make sense? At first sight perhaps yes. But note that playing* R, U' *justified deleting* L, D *but then we deleted* R, U'*... This is an example where undertaking a costly action can signal one's intent. Later we will see how it can signal one's information.*

Explain how this example illustrates the idea of ***forward induction****. Describe the "i'm cold, i.e., please close the window" example. The idea is that if 2 reaches a certain information set, he should look to the beginning of the tree in order to try to explain why he reached this particular information set. With backward induction, we say that the choice at any information set depends only on what comes after that information set and hence that we can work backward. Here we say that decisions depend on where you think you are in the tree and you have to look backward to see this. (Terminology is confusing: backward induction means you look forward, so you work backward. Forward induction means you look backward, so you work forward.)*

I will not offer a precise definition of forward induction because there is no widespread agreement on precisely what the idea means. The "burning–the–dollar" example above is a good illustration of how forward induction can pin things down quite a lot. In this example, dominance solvability pinned us down to a unique solution, while SPE would leave a lot of possibilities. On the other hand, it also seemed to pin things down rather unintuitively.

Exercise 52 *In example* **??***:*

1. *For one of the paths that can arise in a pure-strategy SGP equilibrium (described in the example) find* two *SGP equilibria that have that path. (Note: a SGP equilibrium path is the path of play that the SGP equilibrium strategies would generate; the strategies specify also off-path behavior.)*

2. *Find a SGP equilibrium path of play not described in the example and specify the SGP equilibrium startegies that result in this path of play. Note: It will involve mixing on the path of play.*

3. *True or false: In this example there are no SGP that involve pure strategies along the equlibrium path for which mixed startegies are required off the path.*

Exercise 53 ***Battle of the Sexes with burning a dollar.*** *Suppose that two players play the battle of the sexes game below:*

	B	O
B	$3, 1$	$0, 0$
O	$0, 0$	$1, 3$

Suppose that before the game is played, Player 1 (the row player) can either burn a dollar (that is worth one unit of payoff) or not. Player 1's action is observed by Player 2. The players are both risk-neutral.

1. *Describe the game in extensive and strategic forms.*

2. *"Solve" the game through successive elimination of weakly dominated strategies. Given your answer, what is the rational, if any, for burning the dollar?*

3.10 Sequential Rationality and Refinements of Nash Equilibrium

An example due to Kreps and Wilson (1982) (Figure 4).

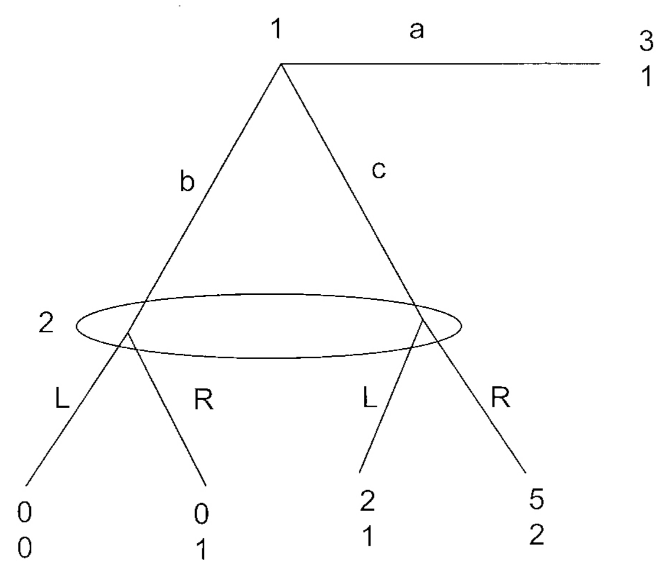

Note that the only subgame is the game itself so the Nash equilibria and subgame perfect equilibria are the same. The normal form is

	L	R
a	$3,1$	$3,1$
b	$0,0$	$0,1$
c	$2,1$	$5,2$

There are two pure strategy Nash equilibria: (a, L) and (c, R) (show at both the strategic and extensive form games). Let's look at the former more closely. This is an equilibrium because if 1 chooses c, 2 will choose L and 1 is made worse off. But what would 2 do if 1 actually chose c? At this point, 2 knows he is at one of two nodes. No matter what probabilities he assesses to the two, he is better off choosing R as it strictly dominates L. If we had done this without the information set, SPE would have said this was an incredible threat and ruled it out. (Show.) Why should putting this information set here make this threat any more credible?

In other words, in this example, no matter what beliefs 2 has at this information set, this strategy is not sequentially rational. Hence subgame perfection did not do what we wanted it to do.

The idea of weak perfect Bayesian equilibrium (weak PBE for short) is supposed to deal with the problem more directly. What we do is to require sequential rationality at *all* information sets, reached and unreached. That is, a weak PBE is a profile of strategies, π, and a system of beliefs μ such that

1. the strategies π are sequentially rational at all information sets given the beliefs μ
2. the beliefs μ are given by Bayes' rule for all information sets that are reached with a positive probability given the strategies π.

Because we're requiring sequential rationality everywhere, we see that in games of perfect information, this picks up the backward induction solution.

Notice also that this notion does better on our example above. The problem was that for any beliefs 2 might have, R is the optimal strategy and subgame perfection did not require this. Weak PBE does, so the only weak PBE (in pure or mixed strategies) is (c, R).

Weak PBE (as the word "weak" in the name should suggest) is not an ideal notion either. It moves in the right direction by attacking the issue of sequential rationality directly, but it doesn't pin down beliefs sufficiently. As a result, a weak PBE may not even be subgame perfect. To see this, consider the following example (Figure 5).

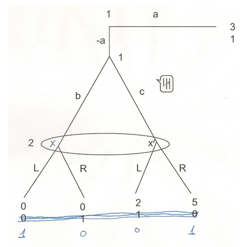

This game has a unique SPE. In the one proper subgame, 1 has a strictly dominant strategy of c, so the subgame equilibrium must be (c, R). Hence the unique SPE is $(\bar{a}c, R)$.

However, (ac, L) is a weak PBE. The beliefs for 2 just have to put probability greater than 1/2 on node x. Given such beliefs, L is sequentially rational for 2. Given 2's strategy, c is sequentially rational for 1 at his second information set. Clearly, given this specification of the subgame, a is sequentially rational for 1 at his initial information set. The problem is that beliefs of 2 are completely unconstrained because his information set has zero probability. It is not reasonable for 2 to put probability greater than 1/2 on x when 1 has a dominant strategy in the subgame of c ! But these beliefs are never "tested" in equilibrium and so are never shown to not make sense.

There are two main approaches in the literature to deal with this problem. The more prominent is called *perfect Bayesian equilibrium* or PBE. This approach simply adds certain constraints which, unfortunately, are very specific to the nature of the game being analyzed. If you look at the Fudenberg and Tirole paper which introduced this idea, you find one definition for one kind of game and a different definition for another kind. The reason this is widely used, though, is that it is much easier to work with than the second approach, which is called *sequential equilibrium.* PBE will be heavily used in the last third of the course.

Sequential equilibrium (due to Kreps and Wilson) can be viewed as adding a particular condition on the beliefs. Fortunately, this condition is *not* game–specific and does give us a lot of nice properties. Unfortunately, it's a pain in the neck to check and it is not obvious what it does. I should mention that Kreps and Wilson were the first to define the notion of sequentially rational, so their contribution is more major than my order of presentation might make it seem.

More formally, say that a behavior strategy is *totally mixed* if it gives every action strictly positive probability at every information set. Notice that if we have a totally mixed behavior strategy, there is a unique belief defined from it by Bayes' rule since every information set will have positive probability. If π is totally mixed, I'll write $\mu(\pi)$ to denote this unique belief. Given (π, μ), we say that beliefs are *consistent* if there exists a sequence $\pi^k \to \pi$ such that each π^k is *totally mixed* and $\mu(\pi^k) \to \mu$. A *sequential equilibrium* is an assessment (μ, π) that is consistent and sequentially rational.

In other words, take a sequentially rational assessment (μ, π). If you want to show it is a sequential equilibrium, you simply construct a sequence of totally mixed strategies converging to π and define the beliefs from Bayes' Rule implied by each strategy in the sequence. If the sequence of beliefs converges to μ, you're done.

3.11 Examples of Sequential Equilibria and Discussion

Let's look at some examples. First, it shouldn't be surprising that sequential equilibrium and weak PBE coincide in the game in Figure 4. After all, the key there was that sequential rationality pinned everything down. In particular, (c, R) is sequentially rational — more precisely, for any μ, the assessment consisting of these behavior strategies and the beliefs μ is sequentially rational.

Consistency is also easy to verify. To be consistent with Bayes' Rule, μ will have to assign probability 1 to the node where 1 chose c. So let 1 choose c with probability $1-(1/n)$ and b and a each with probability $(1/2)(1/n)$. By Bayes' Rule, the μ_n we get from this says that 1 chose c with probability

$$\frac{1-\frac{1}{n}}{1-\frac{1}{n}+\frac{1}{2n}} = \frac{2n-2}{2n-1}$$

which goes to 1 as $n \to \infty$.

What about the game in figure 5? First, let's verify that the subgame perfect equilibrium is sequential. Recall that it is $(\bar{a}c, R)$. With these strategies, every information set is reached. Hence consistency requires beliefs of probability 1 on x'. With these beliefs, this is sequentially rational. Hence it is sequential. Next, let's see if the other weak PBE (ac, L) (the one that isn't subgame perfect) is sequential. We already saw that it is sequentially rational with beliefs giving x probability at least 1/2. So the only thing left to check is consistency. Let p_n be the probability on $\bar{a}$ and q_n the probability on b. These must both go to zero. Given this, the probability on x is

$$\frac{p_n q_n}{p_n q_n + p_n(1-q_n)} = q_n.$$

Hence this probability must go to zero. So no belief with probability at least 1/2 on x is consistent. Hence this is *not* a sequential equilibrium.

In fact, one can prove that every sequential equilibrium is subgame perfect.

Does consistency exhaust the properties we would want beliefs to satisfy? Unfortunately, no. (fig 7)

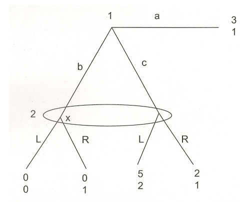

In this game, (a, R) is a sequential equilibrium, supported by having 2 put large enough probability on x. But note that b is dominated by both a and c, so it seems quite unreasonable for 2 to think 1 would play b. Surely it is more sensible that if 2 finds himself unexpectedly called on to play, he concludes that 1 must have played c.

A similar example: Figure 8.

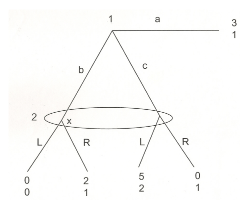

(a, R) is still sequential — again put enough probability on node x. Now the argument that 1 would have preferred c to b does not follow so clearly: c doesn't dominate b anymore. However, a still dominates b, suggesting that 2 should know that 1 would never play b.

Kohlberg and Mertens use this example to motivate what they call *forward induction*. The idea is that if 2 reaches his information set, he should look to the *beginning* of the tree in order to try to explain why he reached the information set. With backward induction, we say that the choice at any information set depends only on what comes after that information set and hence that we can work backward. Here we say that decisions depend on where you think you are in the tree and you have to look backward to see this. (Terminology is confusing: backward induction means you look forward, so you work backward. Forward induction means you look backward, so you work forward.)

Unlike backward induction type arguments where we look from where we are toward the end of the tree, it is the fact that 1 could have had a payoff of 3 from choosing a that determines the beliefs here. Forward induction argues that 2 should believe that 1 chose c and should accordingly choose R. Hence forward induction rules out the (a, R) equilibrium.

I will not offer a precise definition of forward induction because there is no widespread agreement on precisely what the idea means. The basic idea is what I've sketched: when you see a deviation from equilibrium, you look to the beginning of the game in search of an explanation, not just to the end. However, pinning this down more exactly is not straightforward.

I will say, though, that Kohlberg and Mertens have offered a refinement of equilibrium, called stable equilibrium, which is intended to capture the idea. Also, it turns out that if a game is dominance solvable (i.e., the normal form is dominance solvable), then this solution is the unique stable equilibrium. Thus dominance solvability captures at least some of what is meant by forward induction.

The "burning–the–dollar" example above is a good illustration of how forward induction can pin things down quite a lot. As we saw there, dominance solvability pinned us down to a unique solution, while SPE (and sequential eq) would leave a lot of possibilities. On the other hand, it also seemed to pin things down rather unintuitively.

Exercise 54

Find all the pure strategy Nash equilibria in the game depicted in Figure 6 below (Hint: there are three). Do these equilibria satisfy sequential rationality? Are they sequential equilibria?

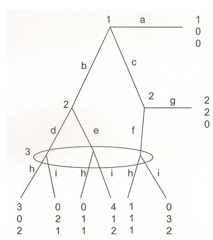

4 Repeated Games

In principle, repeated games are specific examples of extensive form games, and so need no special analysis. But, we have previously mentioned that if a game were repeated several times, then players' incentives would change. They may perceive incentives to establish a reputation, to cooperate, etc. As we shall see below, the study of repeated games provides a useful framework for discussing cooperation and reputation.

4.1 Finitely Repeated Games

Consider for example the repeated prisoners' dilemma game, where the one-shot game is in the table below.

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

Observe that in the twice repeated PD each player has 32 pure strategies (not 4): what to do in the first period, and what to do in the second period after each possible first period outcome (of which there are four).

Exercise 55 *How many strategies are there for each player in the thrice repeated PD ?*

The subgames in a repeated game are easy to identify: they are also repeated games. In a game that is repeated finitely many times, say N, the subgames are various shorter repetitions; in a game that is infinitely repeated there are infinitely many subgames, each of which is the same infinitely repeated game.

Suppose that the Prisoners' Dilemma game above is repeated N times. This has a unique SGP equilibrium in which in every subgame (D, D) is played. In all subgames in the last round there is a unique Nash equilibrium, (D, D). Given that in the last round (D, D) is played play in the preceding round has no impact on the last round, so in the preceding round (D, D) must be played, and so on.

It may be helpful to use the one-step-deviation principle (henceforth OSDP) also to see this. Consider a profile where in some subgame (D, D) is not played. Consider the last period following some history, say period n following h, in which (D, D) is not played, and consider the subgame starting from that point. By switching to D the player who was not playing D in period n gains in that period, and nothing changes afterwards (since afterwards in all subgames (D, D) is played regardless of that player's action in n). That always playing D is SGP follows immediately from the OSDP.

Proposition 7 *In a finite repetition of any game with a unique Nash equilibrium, the SGP equilibrium involves playing that Nash equilibrium in all periods and after every history.*

The finitely repeated PD has an additional feature. In general a finite repetition of a game with a unique Nash equilibrium can have multiple Nash equilibria, but in the PD the unique Nash equilibrium *outcome* is that (D, D) is always played.

Exercise 56 *Prove the claim in the preceding paragraph, namely that in all Nash equilibria of the PD equlibrium path of play always involves (D, D). How then does the set of Nash equilibrium strategty profiles differ from that of SGP equilibrium profiles ?*

Example 48 *Consider the following game played twice. In this game we have added another dominated strategy.*

	C	D	E
C	$3, 3$	$0, 4$	$-1, -1$
D	$4, 0$	$1, 1$	$-1, -1$
E	$-1, -1$	$-1, -1$	$-2, -2$

The following is a symmetric Nash equilibrium: each player plays C and if anyone plays something other than C in the first period then they play E in the second period, otherwise they play D in the second period. This is not a SGP equilibrium.

The next example shows that multiple Nash equilibria in the stage game permit more cooperation in SPE of the repeated game.

Example 49 *Consider the following game played twice.*

	C	D	E
C	$3,3$	$0,4$	$0,4$
D	$4,0$	$1,1$	$1/2,1/2$
E	$4,0$	$1/2,1/2$	$2,2$

Then the following is a symmetric SGP equilibrium. Play C, and afterwards, if both player played C in the first period then play E, and otherwise play D in the second period.

Note that we can now think of a new issue that comes up: *renegotiation.* Consider the players conversing after the first round in which a player chose action D, and imagine that player apologizing and saying to the opponent as follows: "At this point we can either play (D,D) or (E,E). We *both* lose from (D,D), hence we should play (E,E)." Of course, players who anticipate renegotiation would then feel free to play D in the first round, which destroys the possibility of moving beyond the stage game Nash equilibrium outcome. While we will not study renegotiation in repeated games, it is worth noting that it provides another – different in spirit – refinement of Nash and subgame perfect equilibrium.

4.2 Infinitely repeated games

4.2.1 Payoffs and strategies

We will only discuss discounted payoffs (not limit-of-the-averages, and not other criteria, such as overtaking).[17] We will also only consider equal discounting. Thus payoffs given

[17] Two other criteria that appear in the literature are the **Limit of means**, or

$$U_i = \lim\inf_{T\to\infty} \frac{1}{T+1}\sum_{t=0}^{T} u_t$$

and the **Overtaking** criterion, or

$$(v_i^t) \succ (w_i^t) \Leftrightarrow \lim\inf_{T\to\infty}\sum_{t=1}^{T}(v_i^t - w_i^t) > 0.$$

According to the former criterion players ignore histories of finite lengths. For example, $(1,1,1...) \sim (0,1,1...)$. The latter criterion gives a large weight to the infinite future.

actions a_t in period t are $\sum \delta^t u_i(a_t)$.

It is easier to work with normalized payoffs. Since multiplying payoffs by a constant has no effect we multiply by $(1-\delta)$ so the payoffs from $(a_t)_{t=0}^{\infty}$ are $\sum \delta^t u_i(a_t)(1-\delta)$. This is convenient because payoffs then lie in the space of payoffs for one play of the game. For example, if players repeatedly play the same profile a then payoffs are $\sum \delta^t u_i(a_t)(1-\delta) = \sum \delta^t u_i(a)(1-\delta) = u_i(a)\sum \delta^t(1-\delta) = u_i(a)$ (recall that $\sum \delta^t = \frac{1}{(1-\delta)}$).

For δ close to 1, by playing different profiles over time players can obviously get (close to) any convex combination of payoffs $u(a)$ for all a. That is, they can potentially get payoffs in the (rational) convex hull of the payoffs of the one-shot game. In the case of the PD this is the interior of the diamond in the following diagram.

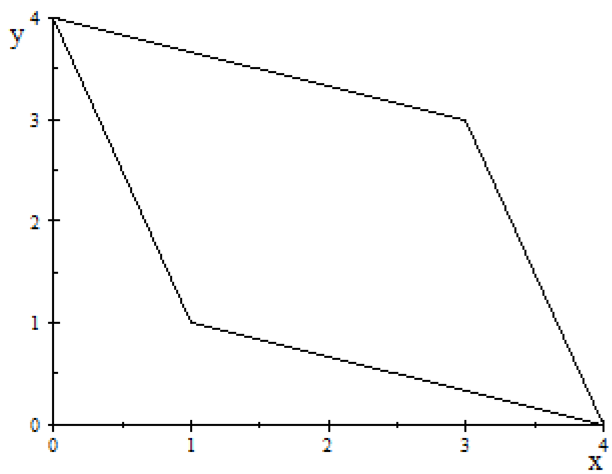

More generally, each profile generates a payoff $u(a)$ and the possible payoffs in the repeated game are within the convex hull of these points, and they are (close to) any point in the hull if δ is close to 1.

Strategies in repeated games specify an action in each period as a function of the *history* until that period. A history in period t is the set of profiles played in the periods until t, an element of $A^t = A \times A \times \cdots \times A$, where A is the set of all action profiles in the one-shot game. So a pure strategy for i specifies for any history up to any period an action for i in the following period. It is a function from $\cup_{t=1}^{\infty} A^t$ into A_i. A behavior strategy is a probability over actions in each period following any history, that is a function from $\cup_{t=1}^{\infty} A^t$ into $\Delta(A_i)$ (where $\Delta(A_i)$ is the set of randomizations over A_i).

Suppose that the Prisoners' Dilemma game above is repeated infinitely many times. What would be an equilibrium of this repeated game? Note that because there is no last period, the game cannot be analyzed using backward induction. Although the world is finite, an infinitely repeated game may be a more compelling modeling device than a finitely repeated game because of the following two reasons:

- In practice, people often ignore the "last period" and so an infinite model is a more appropriate model than a finite game where the last period looms large and has a strong effect on the analysis.

- In some cases, there is no well defined end. In any given period, with some probability, there is another period, and with the complimentary probability the game ends. For example, suppose that in every period the probability that the game continues to the next period is $\frac{1}{2}$. The probability that the game will last for more than 20 periods is

$$1 - \sum_{k=1}^{20} \left(\frac{1}{2}\right)^k \simeq \frac{1}{1,000,000}$$

  Therefore we are used to thinking about such a game as a finite game, although given that we have reached the period n, there is a high probability for the game to continue to period $n+1$.

4.2.2 Equilibria

We proceed to describe different equilibria of the game.

Example 1. "Grim-Trigger" Strategies

Suppose that the players employ the following strategies: play C in every period. Once the other player played D, switch to playing D forever. Observe that:

1. If the players are "sufficiently patient" then for both players to play grim-trigger strategies is a Nash equilibrium of the game. Players need to be sufficiently patient so that the one period gain they get from playing D against the other player's C is smaller than the discounted loss they incur from playing (D, D) instead of (C, C) forever.

2. For both players to play grim-trigger strategies is not a subgame perfect equilibrium of the game. To see this, consider a subgame that begins after one of the players deviated and played D. In this subgame, the deviating player plays C in the first period after the deviation while the other player plays D, which is not a Nash equilibrium. However, grim-trigger strategies can be easily modified so that they are subgame perfect.

Specifically, change the grim-trigger strategy to the following: play C in every period. Once any player played D, switch to playing D forever.[18]

Example 2. Tit-for-tat

Suppose that players employ the following strategies: start playing C. In every period play what the other player has played in the previous period (i.e., reward play of C in the previous period with cooperation in the current period, and penalize defection in the previous period with defection in the current period). Tit-for-tat became famous as the strategy that won a large competition that was organized by the political scientist Robert Axelrod in the 80s.[19]

1. If the players are sufficiently patient, then for both to play tit-for-tat is a Nash equilibrium of the game. This is because along the equilibrium path, players play (C, C) in every period. Playing D instead of C produces no more than $4 + 1 + 3 + 3 + \cdots$ compared to $3 + 3 + 3 + \cdots$.

2. Note however that for both to play tit-for-tat is not a subgame perfect equilibrium. If the other player deviated and played D, then it is better to continue playing C and continue getting C forever, than playing D to which the other player would respond by D in the next period, so that, at best, the sequence of payoffs is given by $(4, 1, 3, 3, \ldots)$. But, of course, to not respond to a defection with some retaliation is not an equilibrium either. Unfortunately, there is no easy fix that would make tit-for-tat subgame perfect.

4.2.3 A Nash Folk Theorem for the Discounting Criterion

We showed that by using "trigger" and other strategies it is possible to obtain Nash equilibrium outcomes of infinitely repeated games that are different from the outcomes obtained by repetitions of Nash equilibria of the stage game. We now ask more generally what outcomes or payoffs can be sustained in a Nash equilibrium of a repeated game.

Suppose that the repeated game consists of an infinite number of repetition of the stage game $G = \langle N, (A_i), (u_i) \rangle$. A payoff profile of the stage game G is *feasible* if it is a convex combination

$$\sum_{a \in A} \alpha_a u(a)$$

[18] In general, any profile of strategies where players play a Nash equilibrium of the stage game in every stage of the repeated game is also a Nash equilibrium of the repeated game.

[19] The article that describes the competition "The Evolution of Cooperation" was published in the journal *Science* in 1981. It is (currently) one of the most cited articels ever published there.

with rational α coefficients.

Example. We saw above a figure that depicts the set of feasible payoffs in the Prisoner's Dilemma game.

If $w = \sum_{a \in A} \alpha_a u(a)$ is a feasbile profile and $\alpha_a = \dfrac{\beta_a}{\gamma}$ for each $a \in A$, where every β_a is an integer and $\gamma = \sum_{a \in A} \beta_a$, then a sequence of outcomes in the repeated game that consists of an indefinite repetition of a cycle of length γ in which each $a \in A$ is played for β_a periods yields an average payoff over the cycle, and hence in the entire repeated game, of w. If the players' discount factor δ is close to 1, then the players' discounted payoff from the game is close to w.

Define player i's *minmax payoff in* G, denoted v_i, to be the lowest payoff that the other players can force upon player i :[20]

$$v_i = \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i})$$

A payoff profile w in G for which $w_i > v_i$ for every $i \in N$ is called *strictly enforceable.*

Example. Feasible and strictly enforceable payoffs in the Prisoner's Dilemma game.

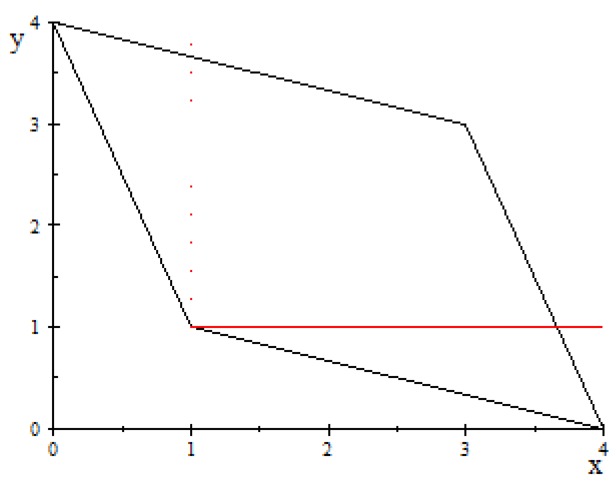

[20]Observe that

$$\min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i}) \geq \max_{a_i \in A_i} \min_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})$$

Or in words, the lowest payoffs that the other player can force upon player i is larger than or equal to the highest payoff that player i can ensure regardless of other players' strategies. It can be shown however that for two player games, the two are equal.

Denote by $p_{-i} \in A_{-i}$ one of the solutions of the minimization problem in the definition of the minmax payoff v_i above. The collection of actions p_{-i} is the most severe per period "punishment" that the other players can inflict upon player i in G.

Proposition (Nash Folk Theorem for the Discounting Criterion). Let w be a strictly enforceable feasible payoff profile of the finite game $G = \langle N, (A_i), (u_i) \rangle$. For every $\varepsilon > 0$ there exists a $\delta \in (0,1)$ large enough and a payoff profile w' of G for which $|w - w'| < \varepsilon$ such that w' is a Nash equilibrium payoff profile of the δ-discounted infinitely repeated game of G.

Proof. Fix $\epsilon > 0$. Let $w = \sum_{a \in A} \frac{\beta_a}{\gamma} u(a)$ be a feasible strictly enforceable payoff profile, where each β_a is an integer and $\gamma = \sum_{a \in A} \beta_a$, and let (a^t) be the cycling sequence of action profiles of which the cycle (of length γ) contains β_a repetitions of a for each $a \in A$. Suppose that δ is sufficiently large so that the discounted payoff to the players from following the cycle (a^t), denoted w', satisfies $|w - w'| < \varepsilon$.

Let s_i be the strategy of player i that chooses a_i^t in each period t unless there was a previous period t' in which a single player deviated from $a^{t'}$, in which case it chooses $(p_{-j})_i$, where j is the deviant in the first such period t'. The strategy profile s is a Nash equilibrium of the repeated game since a player j who deviates recieves at most his minmax payoff v_j in every subsequent period and so the gain from deviation in any given period, which because G is finite is bounded from above, is smaller than the loss, which is given by

$$\sum_{t=1}^{\infty} \delta^t \left(u_j \left(a_j^t, a_{-j}^t \right) - v_j \right) = \frac{w_j'' - v_j}{1 - \delta}$$

and that becomes arbitrarily large as δ approaches 1 (w'' is different from w' because it "starts at a different point in the cycle"; as δ approaches 1, both w'' and w' converge to w and so are strictly larger than v_j) The discount factor δ has to be chosen in such a way that w' is sufficiently close to w and that deviation is not worthwhile. Finally, the payoff profile that is generated by s is w' as required. ■

- Explain that in the SPE folk theorem the problem is how to induce players to penalize other players when penalizing other players may be more costly for the players than being penalized themselves ...

Remark. The proof of the Folk Theorem above raises the following issues:

1. Since penalizing another player can be costly for the players, is it in the interest of players to penalize other players who deviated from equilibrium play? In other words,

is the "punishment phase" of the equilibrium that is described in the proof of the theorem above itself an equilibrium? Investigating the extent to which it is possible to have a subgame perfect folk theorem answers this question. The result is that, in general, it is possible to replace Nash with subgame perfect equilibrium in the Folk Theorem above.

2. Another concern is whether or not it is efficient to penalize players. If not, then penalties can be renegotiated away, which raises the question of what can be sustained in a renegotiation proof equilibrium. See Farrell and Maskin (*Games and Economic Behavior,* 1989) and Benoît and Krishna (*Econometrica*, 1993) for details.

3. Yet another question is what can be done when the players' actions are not fully observable. The work on this subject distinguishes between the relatively easier case where all the players have access to some public signal about past play, and the harder case where each player observes a private signal about other players' past play.

Exercise 57 *Consider the following game. Draw the feasible set. Describe strategies that are a SGP equilibrium and that give approximately any payoff above* $(1,1)$ *for* δ *high enough.*

	A	B
A	$0,0$	$3,1$
B	$1,3$	$0,0$

4.2.4 The Nash-threat folk theorem

Consider the infinitely repeated PD, with discount factor δ. One SGP equilibrium is for the players to always play D, with payoffs $(1,1)$. However, if they are sufficiently patient they may be willing to cooperate and get more now with the threat that if they deviate they will revert to the Nash equilibrium. That is, consider the *trigger* strategies where a player plays C so long as everyone played C, and once someone plays D, they play D forever.

Let's check if this is a SGP equilibrium. Using the OSDP we need to check deviating at nodes after everyone has always cooperated and at nodes after which someone has deviated. The latter nodes result in playing (D,D) forever which we have already seen is a SGP equilibrium. In the former the gain is $4-3=1$ and the cost is that in the future instead of getting 3 repeatedly you get 1 repeatedly. Thus the deviation is worthwhile iff $(1-\delta) > \delta(3-1)$, so if $\delta < 1/3$. For $\delta \geq 1/3$ the strategies constitute a SGP equilibrium.

In this way if δ is high enough players can get any payoff strictly above any Nash equilibrium of the one shot game. In the repeated PD this means any payoff above $(1,1)$. So they can get anything within the feasible region that is above and to the right of the horizontal and vertical red lines.

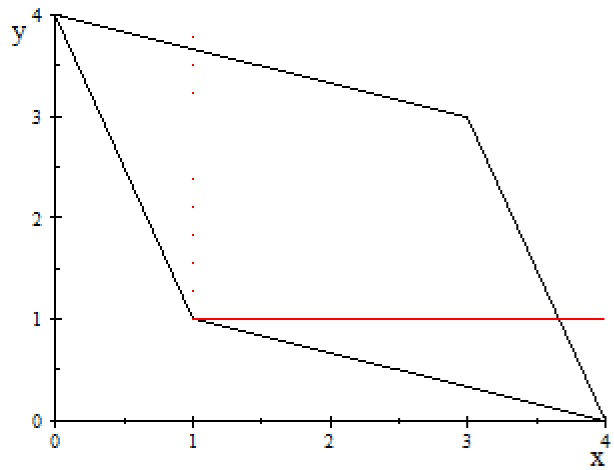

Indeed this approach works in any repeated game using any Nash equilibrium.

Exercise 58 *Consider the following game. Draw the feasible set. Describe strategies that are a SGP equilibrium and that give approximately any payoff above* $(1,1)$ *for* δ *high enough.*

	A	B
A	$0,0$	$3,1$
B	$1,3$	$0,0$

4.2.5 The minmax folk theorem(s)

In the Nash-threat folk theorem the SGP equilibrium which is the repeated one-shot Nash equilibrium serves as a punishment for deviations from many paths. We can also use it as a reward to get payoffs below the nash equilibrium. For example in the game below, there is a unique Nash equilibrium in the one-shot game with payoffs $(1,1)$.

	D	E
D	$1,1$	$1/5,1/5$
E	$1/5,1/5$	$0,0$

It is obvious we cannot have an equilibrium in which a player gets less than $1/5$, since a player can guarantee herself at least $1/5$ by playing D. But interestingly we can get payoffs below 1. Consider the following strategies. Play (E,E) for several, say k, periods, and then play D forever. If anyone fails to play E in one of these first k periods then play E for *another* k periods. We use the OSDP to check if this is a SGP. There are $k+1$ types of subgames: the subgames where there are k more periods of E left, and subgames where D is supposed to be played. In the latter there is clearly no profitable deviation. In the former one can deviate in any period, which will result in a one period gain of $1/5$ and a delay in shifting to the (D, D) path. If a player switches from E to D in one of the k periods then there will be k more periods before getting to the (D, D) path. So if it is worthwhile to switch in one of those periods it is worthwhile to switch in the first. Deviating in the first of the k periods leads to a gain of $1/5$ right away and a delay in the receipt of 1 instead of 0 in k periods. That is, switching to D results in the sequence of payoffs $(1/5, 0, ...0, 1, 1, ...)$ while playing according to the equilibrium yields $(0, ...0, 1, 1, ...)$ where in both the 0 is repeated k times. The difference is $1/5 - \delta^k$, so this deviation is not worthwhile if $1/5 - \delta^k \leq 0$, or $\delta^k \geq 1/5$. So how low can we make the payoffs in a SGP equilibrium of this sort? If (E, E) is played for k periods followed by (D, D) forever then the payoffs are δ^k, which as we saw must be at least $1/5$. So we can get payoffs down to the amount that a player can guarantee herself – namely $1/5$ – for any $\delta \geq 1/5$.

It is of course "odd" to work so hard to get low payoffs, but this illustrates a general idea. Also remember that once we get low payoffs as SGP equilibria we can use those equilibria as punishments to help sustain cooperation. The cooperation that can be obtained is higher the greater is the worst punishment; so the worst SGP equilibrium helps obtain the most cooperation.

In the game below on the left we saw that if $\delta < 1/3$ we cannot sustain the highest payoff (at least not using trigger strategies). Consider adding the strategy of E as in the game on the right. This would not help expand the SGP equilibria in a finitely repeated game since there is a unique SGP equilibrium. But it does enable getting the $(3, 3)$ payoff with δ even lower than $1/3$. We just saw that in the game with just D and E we can get payoffs arbitrarily close to $1/5$. Adding C does not change that calculation. But using the SGP equilibrium that gives payoffs of $1/5$ we can more easily get 3. Consider the following strategies: play C so long as everyone has played C; if anyone ever fails to play C play the SGP equilibrium that yields payoffs $(1/5, 1/5)$. Once again using the OSDP (and our previous analysis that gave that payoff of $1/5$) we see that this is a SGP equilibrium if the stream of payoffs $(4, 1/5, 1/5, ...)$ is not preferred to $(3, 3, ...)$ i.e., if $(1-\delta)\, 4 + \delta/5 \leq 3$, so if $1 \leq 3.8\delta$, or $\delta \geq 1/3.8$. The range of δ for which we can obtain 3 has increased and indeed in this method we can obtain anything in the feasible set that is above $1/5$ — the amount that a player can guarantee herself — for both players, if δ is close enough to 1.

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

	C	D	E
C	$3,3$	$0,4$	$0,4$
D	$4,0$	$1,1$	$1/5,1/5$
E	$4,0$	$1/5,1/5$	$0,0$

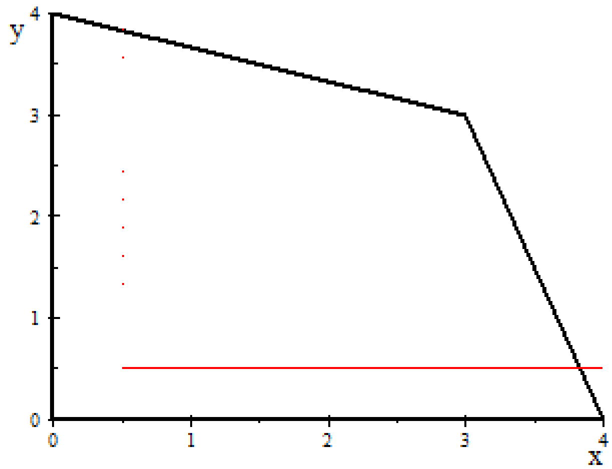

We say that a utility profile is feasible if it lies in the convex hull of the payoffs of the one-shot game; we say that it is enforceable if it is greater than the minmax payoff $v_{*i} \equiv \min_{\alpha_{-i}} \max_{\alpha_i} u_i(\alpha_i, \alpha_{-i})$ for all i.

Proposition 8 ***The Folk Theorem*** *Let G be a game, such that the payment set is n-dimensional (as the number of players), let v be a feasible and enforceable payment. For every $\varepsilon > 0$ there is a SGP equilibrium s and $\bar{\delta} < 1$, such that for every $\bar{\delta} < \delta$, $\|u(s) - v\| < \varepsilon$.*

Thus, as the players become increasingly patient, the set of SGP-equilibrium payoffs converges to the feasible and enforceable set. (Technical detail: We do not need limits to get the Pareto frontier; the limits are needed in order to approximate elements of the convex hull of one-shot payoffs through a sequence of profiles over time and in order to include the lower boundary.)

Proof. (sketch) Assume, for simplicity, that v is obtained when the profile a is being played.

Take a point v' in which every player gets lower payments than v, yet more than his minmax payment, v_{*i}: for every i, $v_{*i} < v'_i < v_i$.

Denote $v'(i) = (v'_1 + \varepsilon, ..., v'_{i-1} + \varepsilon, v'_i, v'_{i+1} + \varepsilon, v'_n + \varepsilon)$ (Everyone, except i, receives an award ε). (Assume, for simplicity, that this point is reached when profile $a(i)$ is being played.)

Let M_i denote the largest gap between i's payments in G: $M_i = \max_{a,a' \in A} u_i(a) - u_i(a')$.

Let m_j be the profile giving j's minmax payoff, m_j solves $\min_{\alpha_{-i}} \max_{\alpha_i} u_i(\alpha_i, \alpha_{-i})$.

Let w_i^j denote i's payment when j is being minmaxed, i.e., $w_i^j = u_i(m_j)$.

Let N satisfy $M_i < N(v_i' - v_{*i})$ for all i, i.e., $N > \max_i M_i / (v_i' - v_{*i})$ hence for δ close enough to 1, a punishment of length N cancels out any possible gain from a one stage deviation.

(Throughout this proof we ignore, for simplicity, the difficulty to observe a deviation from a mixed strategy. This can be dealt with; see Fudenberg and Tirole).

v is being played.
A deviation is bad news for everyone:
the benchmark now becomes v'.
v' is being played +
a prize to players who punish: $v(i)$.[21]

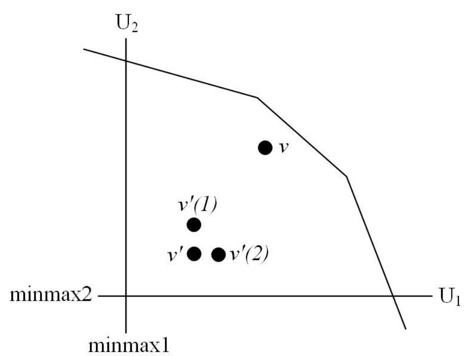

The Structure of the Equilibrium:

Phase I: (The equilibrium path): a is played persistently. If a single player, j, deviates, we move to phase IIj.

Phase IIj: (j is being punished): m_j is played N times, and then move to phase IIIj. If a single player, k (including $k = j$) deviates, we move to phase IIk.

Phase IIIj: (Whoever has punished player j is being compensated): $a(j)$ is played persistently. If a single player, i, deviates, we move to phase IIi.[22] ■

Exercise 59 *Check that this is a SGP equilibrium for $\bar{\delta} < \delta$, by using the OSDP.*

[22] After players have been compensated, why not revert back to v ? This is not necessary in order to sustain the SGP, but is it possible ?

Exercise 60 *Consider the following game. Describe strategies that are a SGP equilibrium and that give approximately any payoff above* $(1,1)$ *for* δ *high enough.*

	A	B
A	$0,0$	$3,1$
B	$1,3$	$4,4$

4.2.6 Example: War of attrition

I present this as a simultaneous-move (discrete-time) timing game. (In the exercises you saw a sequential-move timing game.) Two animals fight for a prize whose value at time $t = 0, 1, ...$ is $v > 1$. Fighting (continuing to stay in) costs 1 each period. As soon as one player quits the other gets the prize; if both quit in the same period neither get the prize. Payoffs in each period t are discounted by δ. This game has many interpretations, as an all-pay dynamic auction, as R&D competition, exit from a market,....

A strategy here is a specification of whether to leave at each information set that can be reached. The information set at time t is reached only if both players continued until time t. So the strategy involves specifying a time to quit.

Two pure-strategy SGP equilibria of the game are given by one player continuing forever and the other quitting immediately. Let's find a symmetric equilibrium, in which in each period t each player quits with probability p_t.

Remark 8 *If in any* t *player* 1 *quits with probability* 1 *then the opponent does not quit with probability one in that period and if in any* t *player 1 quits with probability zero then 2 does not quit with probability zero, so in a symmetric equilibrium* $p_t \in (0,1)$.

Exercise 61 *Prove the above claim.*

Exercise 62 *Consider an SGP equilibrium in which in* t *player 1 quits with probability* 1. *Can you say with what probabilty does player 2 quit in* t*? Consider an SGP equilibrium in which in* t *player 1 quits with probability* 0. *Can you say with what probability does player 2 quit in* t*?*

Remark 9 *Describing strategies by specifying* p_t *at the information set at time* t *is a behavior-strategy description; the equivalent mixed strategy is to assign probability* $\Pi_{\tau=0}^{t}(1-p_\tau)\,p_t$ *to the strategy* "quit at t if the other player hasn't quit yet."

For this to be an equilibrium the player must be indifferent between quitting and continuing in period t. Quitting yields a payoff of 0 (in addition to the costs of having reached

t), continuing yields v if the opponent quits in t, which happens with probability p_t. and -1 if the opponent does not quit (with probability $1-p_t$). (It is important to understand where the -1 comes from. If the opponent does not quit then the next period is reached, and one again is playing a behavior strategy with positive probability on quitting, so the payoff must be exactly what one gets from quitting in the next period.) Indifference implies $0 = p_t v + (1-p_t)(-1)$, so $p_t = 1/(1+v)$; as one would expect the greater the prize the higher is the probability of continuing.

Exercise 63 *Assume that if the players both quit in the same period they each get the prize with probability 1/2. Solve for the symmetric equilibrium.*

Now consider the continuous time version of this game: in each instant each player can quit. A strategy is a specification of a time in $[0,\infty)$, a mixed strategy is a CDF G. Looking for a symmetric mixed-strategy equilibrium, we assume the opponent plays G and seek the payoff for a player exiting at t. It is $G(t)\,v-(1-G(t))\,t-\int_0^t g(\tau)\,\tau \mathrm{d}\tau$. If the player mixed at t then the player must be indifferent between t and small increase in t. Taking the derivative of the payoff w.r.t. t gives $g(t)\,v-1+g(t)\,t+G(t)-g(t)\,t=0$, so $g(t)/(1-G(t))=1/v$. (Make sure you see the "economic" intuition for this condition.) If $G(t)=1-e^{-\frac{t}{v}}$ this condition is satisfied.

We confirm that this is the limit as time gets small of the outcome of the discrete-time game. If we assume that the interval is of length Δ then the cost per stage of continuing is $-\Delta$, and the same analysis as in the discrete-time game before yields that $pv+(1-p)\,\Delta=0$ so $p=\Delta/(\Delta+v)$. Then $1-G(t)=(1-p)^{t/\Delta}$ because t/Δ is the number of periods until t, and $(1-p)^{t/\Delta}=\left(\frac{v}{v+\Delta}\right)^{t/\Delta}$. Finally taking Δ to zero gives $\lim_{\Delta\to 0}\left(\frac{v}{v+\Delta}\right)^{t/\Delta} = 1/\lim_{\Delta\to 0}\left(1+\frac{\Delta}{v}\right)^{t/\Delta} = 1/\lim_{x\to 0}(1+x)^{xt/v} = e^{-t/v}$ (where the equality before last define $x=\Delta/v$).

5 Games with Incomplete Information

5.1 Introduction

Uncertainty about the state of the world (payoffs) can be captured by identifying payoffs with expected payoffs. But how can we describe a situation in which different players know that the state of the world is determined randomly according to some stochastic process and that they will receive different information about the state of the world? For example, demand for a good will be high or low next week, and an incumbent firm will know the demand, but a firm considering entry today will not have the market research reports. We already have the tools for this—we add a player, Nature, and give players different information sets about Nature's actions. The players know this environment—that "Nature" will determine if demand is high or low, that one player will receive the information and the other will not. So the game itself continues to be known, and in fact is common knowledge among the players, as before. These are called games with moves by Nature or Bayesian games, and are solved and studied exactly as are games of imperfect information. Before we study such games in detail, let's consider just a possible information structure—a move by Nature and the information sets.

Example 50 *Nature chooses H or L with probabilities 1/5, 4/5, player 1 will be informed of this, and 2 will not. If the state is H then* 1 *assigns probabilities* $(1,0)$ *to* H, L *while* 2 *assigns probabilities* $(1/5, 4/5)$*. If the state is L then* 1*'s beliefs are* $(0,1)$ *but* 2*'s are unchanged. Thus in each state we know the players' belief. Note that this generates "higher-order" beliefs (beliefs about beliefs): we know that in both states* 1 *assigns probability* 1 *to* 2*'s beliefs being* $(1/5, 4/5)$ *and that* 2 *assign probability* $1/5$ *to* 1*'s belief being* $(1,0)$ *and* $4/5$ *to it being* $(0,1)$*.*

A different example. It is known that 1 *will be informed whether the state is L or H, but* 2 *might or might not be informed and in any case not perfectly; with probability* $1/3$ 2*'s market research will give a signal about the state, where the signal is* l *or* h*, and it gives the corresponding lower-case letter signal with probability* $3/4$*; with probability* $2/3$ 2*'s market research will give no information to* 2*. Now* 2*'s beliefs are either* $(1/5, 4/5)$ *or* $(\Pr(H|h), \Pr(L|h))$ *or* $(\Pr(H|l), \Pr(L|l))$*. Note that* $\Pr(H|h) = \frac{(1/5)(3/4)}{(1/5)(3/4)+(4/5)(1/4)} = \frac{3}{7}$*,* $\Pr(L|l) = \frac{(4/5)(3/4)}{(4/5)(3/4)+(1/5)(1/4)} = \frac{12}{13}$ *and so on. Now Nature determines both the demand (H or L) and the information that 2 gets. In the state* (L,l) *2's beliefs about H and L are* $(1/13, 12/13)$*, 1's beliefs are still* $(0,1)$*;* 2*'s beliefs about* 1*'s beliefs are that they are* $(0,1)$ *with probability* $(1/13, 12/13)$*, and* $(1,0)$ *otherwise;* 1*'s beliefs about* 2 *depend on* 1*'s "type" L or H; with probability* $2/3$ *they are* $(1/5, 4/5)$ *regardless of* 1*'s type; if* 1 *is of type H then it believes that* 2*'s beliefs are given by* $(3/7, 4/7)$ *with probability* $(1/3)(3/4) = 1/4$ *and by* $(1/13, 12/13)$ *with probability* $(1/3)(1/4) = 1/12$*, and so on.*

[Show "no common knowledge ?" video]

Consider now a situation where a player knows his own preferences but not that of others. In contrast to the preceding situation where both firms know that demand will be high or low and are studying the future here there is no "first stage" before information is obtained, every player is "born" knowing the information. These are called Bayesian games. Harsanyi argued that such situations should be modeled exactly as the previous example—with a move by Nature and suitable information structures. This is because given any set of "rational" players each player has beliefs about the other players' preferences, about their beliefs about everyone's preferences, and so on. We saw above that any commonly known model of Nature with information sets gives rise to players having beliefs, beliefs about beliefs, etc.; Harsanyi's insight was that any hierarchy of beliefs about beliefs comes from some commonly known model of this form. We will not delve into this technically and conceptually challenging topic of how to model what are called games of incomplete information. We will simply apply Harsanyi's insight and model both situations using Bayesian games–game of imperfect information with moves by Nature. Because of Harsanyi's insight that games of incomplete information can be modeled as games of imperfect information with a move by nature, in the literature the latter are often (imprecisely) called games of incomplete information.

The examples we will study include a firm that knows its own cost, but not a rival's; more generally each player knowing only his or her own preferences, sellers having better information than buyers about quality of goods; bidders in an auction know their own willingness to pay but not that of others, and so on.

We will begin by studying what are called static games of incomplete information: Nature moves, players get information, and then players play a strategic form game. We will then turn to dynamic games of incomplete information.

Example 51 ***Second price auction.*** *The information structure is that each person learns his own value for the object being auctioned, and nothing about the other. Then, the players participate in a second price auction. A strategy for player i in this game is a specification of a bid at each information set that i might have. So it is a function from possible values into bids, $s_i : V_i \to \mathbf{R}$, where V_i is the set of possible values for player i.*

Exercise 64

1. *Prove that bidding the true value is a weakly dominant strategy for each bidder.*

2. *Find an asymmetric equilibrium for the second price auction with two bidders whose valuations lie between zero and one.*

Example 52 *Two firms engage in Cournot competition. Demand is given by $p(Q) = 3 - Q$ where Q is the total quantity produced in the market. The cost of production is uncertain. Each firm has cost 1 or 2 with equal probability, independently of the cost of the other firm. Each firm knows its own cost realization but not the cost of the other firm. Both firms are interested in maximizing their expected profits.*

We search for a symmetric equilibrium. Denote the quantity produced by a firm whose cost is 1 and 2 by q_1 and q_2, respectively. In equilibrium, it must be that each type of the firm *(i.e., both firm with cost 1 and firm with cost 2) chooses its quantity q so as to maximize its expected profits. A firm with cost 1 chooses q to maximize*

$$\frac{1}{2} \cdot q\,(3 - q - q_1) + \frac{1}{2} \cdot q\,(3 - q - q_2) - q.$$

(The cost of production is 1 and the other firm has cost 1 and 2 with equal probabilities.) A firm with cost 2 chooses q to maximize

$$\frac{1}{2} \cdot q\,(3 - q - q_1) + \frac{1}{2} \cdot q\,(3 - q - q_2) - 2q.$$

(The cost of production is 2 and the other firm has cost 1 and 2 with equal probabilities.)

In a symmetric equilibrium where i chooses q depending on its cost and the opponent chooses q_i depending on its cost, it follows that:

$$q_1 = \arg\max_{q \geq 0} \left[\frac{1}{2} \cdot q\,(3 - q - q_1) + \frac{1}{2} \cdot q\,(3 - q - q_2) - q\right] \tag{1}$$

and

$$q_2 = \arg\max_{q \geq 0} \left[\frac{1}{2} \cdot q\,(3 - q - q_1) + \frac{1}{2} \cdot q\,(3 - q - q_2) - 2q\right] \tag{2}$$

The first order condition for the maximization of (1) is:

$$4 - 4q - q_1 - q_2 = 0.$$

The fact that q_1 is the maximizer therefore implies that it must satisfy the following equation:

$$4 - 5q^1 - q^2 = 0.$$

Similarly, the first order condition for the maximization of (2) is:

$$1 - 2q - \frac{q_1}{2} - \frac{q_2}{2} = 0$$

and the fact that q_2 is the maximizer therefore implies that it must satisfy the following equation:

$$2 - q_1 - 5q_2 = 0$$

The solution is:

$$q_1 = \frac{3}{4}; q_2 = \frac{1}{4}.$$

This is also a Nash equilibrium of the game.

5.2 Bayesian Games: The General Formulation

A Bayesian game is a quintuple

$$G = \langle N, A, T, p, u \rangle$$

where

$N = \{1, ..., n\}$ – is a set of players.

$A = A_1 \times \cdots \times A_n$ where A_i is a set of actions for player i.

$T = T_1 \times \cdots \times T_n$ where T_i is a set of types for player i; each player is informed of his own type. A player's type describes the player's "private information."

$p(t) \geq 0$ is a common prior probability, $\sum_{t \in T} p(t) = 1$. The players' beliefs can be derived from p through Bayesian updating.

$u = (u_1, ..., u_n)$ where $u_i : A \times T \to \mathbb{R}$ is player i's payoff (utility) function. (Players are assumed to be expected utility maximizers.)

Remark 10 *The payoff to player i may depend on other players' types and other players' actions as well as on his own type and action.*

The notion of type allows us to model both different "preference" types (e.g., two types of player 1 may be distinguished by their different utility function), and different "belief" types (e.g., two types of player 1 may be distinguished by their different beliefs about player 2's utility function).

A common prior is a standard assumption in economic applications, it is quite restrictive and thereby helps narrow down predictions.

As mentioned above, it can be shown that this formulation of a Bayesian game can be used to describe any situation that involves incomplete or asymmetric information. Showing that this is indeed the case is beyond the scope of this course.[23]

A strategy for player i is given by a function

$$s_i : T_i \to \Delta(A_i).$$

This formulation captures the fact that players know their own types, but not other players' types. The beliefs of type t_i of player i about other players' types is given by Bayes rule

$$p(t_{-i}|t_i) = \frac{p(t)}{p(t_i)}$$

[23] The standard reference is Mertens and Zamir (IJGT, 1985); Myerson's game theory textbook contains an accessible treatment of this difficult subject.

Definition 19 *A profile of strategies* $(s_1^*, ..., s_N^*)$ *is a Bayesian-Nash equilibrium if for every type* t_i *of every player* i

$$\sum_{t_{-i}} p(t_{-i} | t_i) \, u\left(s_i^*(t_i), s_{-i}^*(t_{-i}); t_i, t_{-i}\right) \geq \sum_{t_{-i}} p(t_{-i} | t_i) \, u\left(s_i, s_{-i}^*(t_{-i}); t_i, t_{-i}\right)$$

for every $s_i \in \Delta(A_i)$. *In other words, every action* $a_i^*(t_i)$ *in the support of* $s_i^*(t_i)$ *satisfies*

$$a_i^*(t_i) \in \arg\max_{a_i \in A_i} \sum p(t_{-i} | t_i) \cdot u_i\left(a_i, s_{-i}^*(t_{-i}); t_i, t_{-i}\right).$$

5.3 Examples

Example 53 *Two firms engage in Cournot competition. Demand is given by* $p(Q) = a_1 - Q$, $p(Q) = a_2 - Q$, *and* $p(Q) = a_3 - Q$, *with probabilities* $\frac{1}{6}$, $\frac{1}{3}$, *and* $\frac{1}{2}$, *respectively. Firm* 1 *knows whether the* y *intercept, denoted* a, *is at* a_1 *or at one of* $\{a_2, a_3\}$, *and firm* 2 *knows whether* a *is in* $\{a_1, a_2\}$ *or equals* a_3.[24] *The cost of production is zero for both firms. Both firms are interested in maximizing their expected profits.*

Each firm has two types: firm A has types $\{\{1\}, \{2, 3\}\}$ *and firm B has types* $\{\{1, 2\}, \{3\}\}$, *which we will write as* $\{1, 23\}$ *and* $\{12, 3\}$.

The state of the world is described by the pair of types and the state of demand, so there are 12 possibilities with the prior probabilities as follows:

	a_1		a_2		a_3	
	12	3	12	3	12	3
1	1/6	0	0	0	0	0
23	0	0	1/3	0	0	1/2

A observes the row, generating the following conditional probabilities:

[24]For example, one firm's research department focuses on answering the question "is demand low?" and the other on "is demand high?" Or, firm A observes whether demand in area A is high or low, and firm B observes whether demand in area B is high or low, and it cannot be that demand in area A is high if demand in area B is low. For example, if the good in question is a luxury good, then if demand for it is low in upper-middle class neighborhoods it must also be low in working class neighborhood.

$A \backslash B$	low	high
low	$a_3, \frac{1}{2}$	$a_2, \frac{1}{3}$
high	0	$a_1, \frac{1}{6}$

	$a_1, 12$	$a_1, 3$	$a_2, 12$	$a_2, 3$	$a_3, 12$	$a_3, 3$
1	1	0	0	0	0	0
23	0	0	$\frac{1/3}{1/3+1/2} = \frac{2}{5}$	0	0	$\frac{1/2}{1/2+1/3} = \frac{3}{5}$

B observes the column generating the following conditional probabilities

	$a_1, 1$	$a_1, 23$	$a_2, 1$	$a_2, 23$	$a_3, 1$	$a_3, 23$
12	$\frac{1/6}{1/6+1/3} = \frac{1}{3}$	0	0	$\frac{2}{3}$	0	0
3	0	0	0	0	0	1

We search for a symmetric equilibrium. Denote the quantity produced by type j of firm i by q_j^i. In equilibrium, it must be that

$$
\begin{aligned}
q_1^A &= \arg\max_{q \geq 0} q\left(a_1 - q - q_{12}^B\right) \\
q_{23}^A &= \arg\max_{q \geq 0} \left[\frac{2}{5} \cdot q\left(a_2 - q - q_{12}^B\right) + \frac{3}{5} \cdot q\left(a_3 - q - q_3^B\right)\right]
\end{aligned}
$$

and

$$
\begin{aligned}
q_{12}^B &= \arg\max_{q \geq 0} \left[\frac{1}{3} \cdot q\left(a_1 - q - q_1^A\right) + \frac{2}{3} \cdot q\left(a_2 - q - q_{23}^A\right)\right] \\
q_3^B &= \arg\max_{q \geq 0} q\left(a_3 - q - q_{23}^A\right)
\end{aligned}
$$

We thus get four linear equations with four unknowns:

$$
\begin{aligned}
q_1^A &= \frac{a_1 - q_{12}^B}{2}; \qquad 10 q_{23}^A = 2\left(a_2 - q_{12}^B\right) + 3\left(a_3 - q_3^B\right) \\
6 q_{12}^B &= a_1 - q_1^A + 2\left(a_2 - q_{23}^A\right); \qquad q_3^B = \frac{a_3 - q_{23}^A}{2}
\end{aligned}
$$

The solution is given by:

$$
\begin{aligned}
q_1^A &= \frac{77}{171} a_1 - \frac{26}{171} a_2 + \frac{2}{57} a_3; \qquad q_{23}^A = -\frac{4}{171} a_1 + \frac{28}{171} a_2 + \frac{11}{57} a_3 \\
q_{12}^B &= \frac{17}{171} a_1 + \frac{52}{171} a_2 - \frac{4}{57} a_3; \qquad q_3^B = \frac{2}{171} a_1 - \frac{14}{171} a_2 + \frac{23}{57} a_3
\end{aligned}
$$

Example 54 *Two firms engage in Cournot competition. Demand is given by $p(Q) = 3 - Q$ where Q denotes the total quantity produced. The cost of firm B is zero. The cost of firm A is equally likely to be either zero or one. If the cost is zero, then firm B knows it with probability $\frac{1}{2}$, and doesn't know it with probability $\frac{1}{2}$. If the cost is one, then firm B doesn't know it. Firm A knows its own cost but doesn't know whether firm B knows it, and this is commonly known between the two firms. Both firms are interested in maximizing their* *expected profits. It is as if B spies after A's cost. Either it discovers it has the special machine that's needed to produce at zero cost, or it doesn't, and in this case doesn't know if A has the machine or not.*

Each firm has two types: firm A has types $\{0, 1\}$ and firm B has types $\{k, d\}$. (There are also cost possibilities $\{0, 1\}$, but since it is commonly known that A's type and the costs are perfectly correlated we can combine these to get a situation that is easier to describe than the preceding example.)[25] *The left table specifies the prior, the middle the conditional probabilities of where each row conditions on a different event, B's conditionals are in the third table and each column is a different conditioning event.*

	k	d
0	1/4	1/4
1	0	1/2

	k	d
0	1/2	1/2
1	0	1

	k	d
0	1	1/3
1	0	2/3

We search for an equilibrium. Denote the quantity produced by type j of firm i by q^i_j. In equilibrium, it must be that

$$\begin{aligned} q_0^A &= \arg\max_{q \geq 0} \left[\frac{1}{2} \cdot q \left(3 - q - q_k^B\right) + \frac{1}{2} \cdot q \left(3 - q - q_d^B\right) \right] \\ q_1^A &= \arg\max_{q \geq 0} \left[q \left(3 - q - q_d^B - 1\right) \right] \end{aligned}$$

and

$$\begin{aligned} q_k^B &= \arg\max_{q \geq 0} \left[q \left(3 - q - q_0^A\right) \right] \\ q_d^B &= \arg\max_{q \geq 0} \left[\frac{1}{3} \cdot q \left(3 - q - q_0^A\right) + \frac{2}{3} \cdot q \left(3 - q - q_1^A\right) \right] \end{aligned}$$

We thus get four linear equations with four unknowns:

$$\begin{aligned} 4q_0^A &= 6 - q_k^B - q_d^B; \qquad 2q_1^A = 2 - q_d^B \\ 2q_k^B &= 3 - q_0^A; \qquad 6q_d^B = 9 - q_0^A - 2q_1^A \end{aligned}$$

[25] If it weren't commonly known that firm A doesn't know what firm B knows, then we would have had to introduce additional types of firm A that describe what A believes about B. If it wasn't then commonly known that B doesn't know A's beliefs about what B knows, we would have had to introduce additional types for firm B that describe its beliefs and so on.

In this sense, the assumption of common knowledge helps to "close the model."

The solution is given by:

$$\begin{aligned} q_0^A &= \frac{31}{33}; \qquad (>) \qquad q_1^A = \frac{13}{33} \\ q_k^B &= \frac{34}{33}; \qquad (<) \qquad q_d^B = \frac{40}{33} \end{aligned}$$

Notice that firm B is better off when the spy does not discover anything. This might suggest that firm B would have liked to commit, if it could, not to know firm A's cost. But that ignores the effect on A's behavior that B might know the cost. We can study whether B wants to commit only by studying the game where the probability with which B succeeds in learning A's type is a parameter p.

Example 55 *As before two firms engage in Cournot competition. Demand is given by* $p(Q) = 3 - Q$ *where Q denotes the total quantity produced. The cost of firm B is zero. The cost of firm A is equally likely to be either zero or one. If the cost is zero, then firm B is informed of it with probability p, and is not informed with probability* $1-p$*. If the cost is one, then firm B doesn't know it. Firm A knows its own cost but doesn't know whether firm B knows it, and this is commonly known between the two firms. Both firms are interested in maximizing their expected profits.*

Each firm has two types: firm A has types $\{0,1\}$ *and firm B has types* $\{k,d\}$*. (There are also cost possibilities*$\{0,1\}$*, but since it is commonly known that A's type and the costs are perfectly correlated we can combine these to get a situation that is easier to describe than the preceding example.) The left table specifies the prior, the middle the conditional probabilities of where each row conditions on a different event, B's conditionals are in the third table and each column is a different conditioning event.*

	k	d
0	$p/2$	$(1-p)/2$
1	0	$1/2$

	k	d
0	p	$1-p$
1	0	1

	k	d
0	1	$\frac{1-p}{2-p}$
1	0	$\frac{1}{2-p}$

We search for a symmetric equilibrium. Denote the quantity produced by type j of firm i by q_j^i*. In equilibrium, it must be that*

$$\begin{aligned} q_0^A &= \arg\max_{q\geq 0} \left[p \cdot q\left(3 - q - q_k^B\right) + (1-p)\cdot q\left(3 - q - q_d^B\right)\right] \\ q_1^A &= \arg\max_{q\geq 0} \left[q\left(3 - q - q_d^B - 1\right)\right] \end{aligned}$$

and

$$\begin{aligned} q_k^B &= \arg\max_{q\geq 0}\left[q\left(3-q-q_0^A\right)\right] \\ q_d^B &= \arg\max_{q\geq 0}\left[\frac{1-p}{2-p}\cdot q\left(3-q-q_0^A\right)+\frac{1}{2-p}\cdot q\left(3-q-q_1^A\right)\right] \end{aligned}$$

We thus get four linear equations with four unknowns in two parameters:

$$\begin{aligned} 2q_0^A &= p\left(3-q_k^B\right)+(1-p)\left(3-q_d^B\right) \\ 2q_1^A &= 2-q_d^B \\ 2q_k^B &= 3-q_0^A \\ 2(2-p)\,q_d^B &= \left(3-q_1^A\right)+(1-p)\left(3-q_0^A\right) \end{aligned}$$

The solution is: $\left[q_0^A=\frac{13p-22}{15p-24}, q_1^A=\frac{7p-10}{15p-24}, q_d^B=\frac{16p-28}{15p-24}, q_k^B=\frac{16p-25}{15p-24}\right]$

$u_d^B=\left[\frac{1-p}{2-p}\cdot q\left(3-q-q_0^A\right)+\frac{1}{2-p}\cdot q\left(3-q-q_1^A\right)\right]$

$=\left(\frac{1-p}{2-p}\cdot\frac{16p-28}{15p-24}\left(3-\frac{16p-28}{15p-24}-\frac{13p-22}{15p-24}\right)+\frac{1}{2-p}\cdot\frac{16p-28}{15p-24}\left(3-\frac{16p-28}{15p-24}-\frac{7p-10}{15p-24}\right)\right)=$

$=\frac{1416p-1056p^2+256p^3-616+1160p-352p^2-952}{2016p-1170p^2+225p^3-1152}$ *is drawn in black.*

Similarly, u_k^B *is drawn in red, and* u^B *is drawn in blue.*

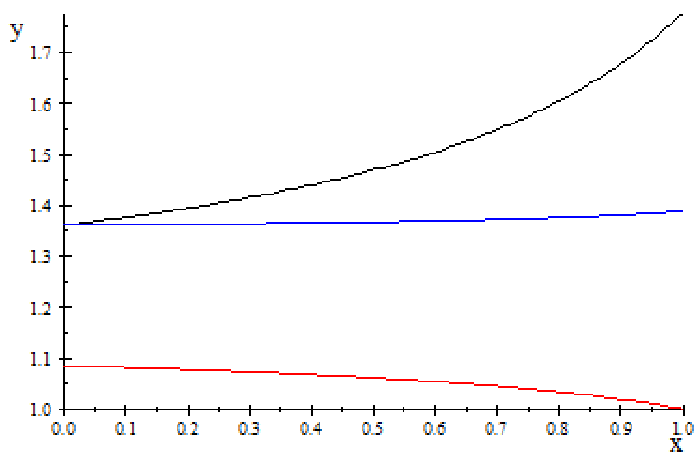

We see that for any p, B's profits are higher as type d than as type k, but B's overall profits are increasing in p.

Example 56 *Read examples 8.E.1 and 8.E.2 from MWG*

Exercise 65 *Demand can be High, denoted H, or Low, denoted L, with equal probabilty. Player $i=1,2$ observes the signal h_i with probability $2/3$ if H and with probabilty $1/3$ if L,* *and otherwise observes signal l_i. The actions of the players are invest or not. The payoffs of the players are symmetric and are as follows. If they do not invest thier payoff is zero. If both invest their payoff is 3 if demand is H and $-B$ if demand is L. If one invests then the payoff to that one is 4 if demand is H and 1 if demand is L.*

Find the pure strategy symmetric equilibria for all values of B.

Find the pure asymmetric equilibria for all values of B.

Exercise 66 *Consider a Bayesian game in which player 1 may be either type a or type b, where type a has probability .9 and type b has probably .1. Player 2 has no private information. Depending on player 1's types, the payoffs to the two players depend on their actions in $A_1 = \{U, D\}$ and $A_2 = \{L, R\}$ as shown in the following table (payoffs if 1 has types a and b are on the left and right matrices, respectively).*

	L	R
U	$2, 2$	$-2, 0$
D	$0, -2$	$0, 0$

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

Compute all Bayesian Nash equilibria of this game.

Exercise 67 (An exchange game) *Each of the two players receives a ticket on which there is a number in the finite set $\{x_1, x_2, \ldots, x_m\}$, where $x_1 < x_2 < \cdots < x_m$. The number on a players' ticket is the size of a prize that she may receive. The two prizes are identically and independently distributed, with distribution function F. Each player knows only the huber on her own ticket. Each player is asked independently and simultaneously whether she wants to exchange her prize for the other player's prize. f both players agree then the prizes are exchanged; otherwise each player receives her own prize. The players' objective is to maximize their expected payoffs. Model this situation as a Bayesian game and show that in any Bayesian Nash equilibrium the highest prize that either player is willing to exchange is the smallest possible prize x_1.*

Exercise 68 (A parlor game) *Formulate the following parlor game as an extensive-form game with imperfect information. First player 1 receives a card that is either H or L with equal probabilities. Player 2 does not see the card. Player 1 may announce that her card is L, in which case she must pay \$1 to player 2, or may claim that her card is H, in which case player 2 may choose to concede or to insist on seeing player 1's card. If player 2 concedes then he must pay \$1 to player 1 . If he insists on seeing player 1's card then player 1 must pay him \$4 if her card is L and he must pay her \$4 if her card is H. Define the set of strategies of each player and find the Nash equilibria of this game.*

5.4 Auctions

We model auctions as Bayesian games. We first discuss private and then common value auctions of a single object.

5.4.1 Private Value Auctions

– Suppose there are $n \geq 1$ bidders (potential buyers). The number n is *exogenously given.*

– Each bidder $i \in \{1, ...n\}$ knows his own valuation (hence the term private values), which we denote by $v_i \geq 0$. The bidders do not know the other bidders' valuations, but believe that they are each independently distributed according to some distribution function F.

– The payoff to bidder i if it wins the object and pays a price p_W is $v_i - p_W$; its payoff if it just pays p_L without winning the object is $-p_L$ (bidders may pay even if they don't win the object; for example, if the auction requires participation fees). Since we assume that bidders are (risk-neutral) expected utility maximizers, it follows that the payoff to bidder i if it wins the object with probability q and pays an expected price $p = qp_W + (1-q)\,p_L$ is[26]

$$q \cdot v_i - p.$$

First-Price Auction

Description. Bidders submit their bids simultaneously.[27] The highest bidder wins the object and pays its bid. Other bidders pay nothing. In case of a tie, the winning bidder is chosen randomly from among those who submitted the highest bid.

[Show "Rules of the Game" video from Butch and Cassidy]

We denote bidder i's bid by b_i. The rules of the auction define a Bayesian game as follows:

– Players: $N = \{1, ...n\}$;

– Actions (possible bids): $B_i = \{b_i \geq 0\}$;

[26] More generally, the expected utility of a risk averse bidder with a utility function from money $u(\cdot)$ is

$$qu\,(v_i - p_W) + (1-q)\,u\,(-p_L)\,.$$

[27] Bidding need not be literally simultaneous. What's important is that bidders don't know other bidders' bids at the time they make their own bids. So, for example, writing bids into envelopes qualifies as simultaneous bidding even if it's not all done at the exact same time.

– Types: $V_i = \{v_i \geq 0\}$;

– Prior: We assume that bidders' valuations are independently and identically distributed (i.i.d.). Because of independence,

$$\Pr(v_1 \leq \overline{v}_1, v_2 \leq \overline{v}_2, ..., v_n \leq \overline{v}_n) = \prod_{i=1}^{n} F(\overline{v}_i).$$

– Payoffs:

$$\begin{aligned} u_i(v,b) &= (v_i - b_i)\Pr(i \text{ wins}) \\ &= \begin{cases} v_i - b_i & b_i > \max\{b_j : j \neq i\} \\ \frac{v_i - b_i}{\#\left\{j : b_j = \max\limits_{k\in\{1,\dots,n\}}\{b_k\}\right\}} & b_i = \max\{b_j : j \neq i\} \\ 0 & b_i < \max\{b_j : j \neq i\} \end{cases} \end{aligned}$$

A **strategy** for bidder i is given by a bidding function $b_i(v_i) : V_i \longmapsto B_i$ that specifies a bid for whatever valuation bidder i may have for the object.

A **Bayesian-Nash equilibrium** is given by a collection of bidding functions $\{b_i(\cdot)\}_{i\in\{1,\dots,n\}}$ such that for every type v_i of every bidder i

$$b_i(v_i) \in \arg\max_{b_i}(v_i - b_i)\Pr(i \text{ wins the object}).$$

Example. Suppose that $n = 2$, and that $v_1, v_2 \sim U[0,1]$ independently of each other. We search for a Bayesian-Nash equilibrium in increasing linear strategies, that is, such that $b_i(v_i) = a_i + c_i v_i$, $a_i \geq 0, c_i > 0$. Note that we do not restrict the bidders' strategy space. We show that an equilibrium in linear strategies exists. Further analysis, which is beyond the scope of this course, reveals that such an equilibrium not only exists but is in fact unique in this example.

Note that because by assumption $c_j > 0$ the probability that the two bidders' bids tie is zero. It follows that in equilibrium, it must be that

$$b_i(v_i) \in \arg\max_{b_i}(v_i - b_i)\Pr(b_i \geq a_j + c_j v_j).$$

Now,

$$\begin{aligned}\Pr\left(b_i \geq a_j + c_j v_j\right) &= \Pr\left(v_j \leq \frac{b_i - a_j}{c_j}\right) \\ &= \frac{b_i - a_j}{c_j}\end{aligned}$$

provided $0 \leq \frac{b_i - a_j}{c_j} \leq 1$. This is indeed the case because optimality requires that $a_j \leq b_i \leq a_j + c_j$. It therefore follows that

$$b_i\left(v_i\right) \in \arg\max_{b_i}\left(v_i - b_i\right)\left(\frac{b_i - a_j}{c_j}\right).$$

Working through the necessary and sufficient first-order condition we get

$$b_i\left(v_i\right) = \frac{a_j}{2} + \frac{v_i}{2},$$

or, because by assumption, $b_i\left(v_i\right) = a_i + c_i v_i$,

$$a_i = \frac{a_j}{2}, c_i = \frac{1}{2}.$$

By symmetry,

$$a_j = \frac{a_i}{2}, c_j = \frac{1}{2},$$

from which it follows that $a_1 = a_2 = 0$ and the equilibrium is given by

$$b_i\left(v_i\right) = \frac{v_i}{2}, \qquad i \in \{1, 2\}.$$

Note the the bidders always bid less than their valuations. ■

Before solving for the equilibrium in a more general model let's see why the strategies must be monotonic (non-decreasing) in equilibrium. Assume some type v optimally wants to bid b. Consider a type who values the good more, $\hat{v} > v$. That type cannot optimally bid less. By bidding less the probability of winning goes down and so does the payment. But a type with a higher value cares more about winning so that the decrease in probability must be at least as costly to this type as to v, and the savings is the same. This property is called the single-crossing property–the marginal willingness to pay is increasing in the player's type.

Formally, assume the opponent adopts some strategy s_j. This generates a probability distribution over bids, and in particular determines for each bid b the probability that i

wins if he bids b, denoted by $G(b)$.[28] Of course G is non-decreasing. Assume that type v_i optimally bids b_i^* against s_j and that $\hat{b}_i^*$ is the optimal bid for $\hat{v}_i > v_i$.

Claim 3 *$\hat{b}_i^* \geq b_i^*$ except if $G(b_i^*) = G\left(\hat{b}_i^*\right) = 0$, i.e., except if both types lose with probability 1 when they bid optimally.*

Proof. Since b_i^* is optimal for v_i is it better than bidding anything else, in particular better than bidding $\hat{b}_i^*$, and conversely.

$$\begin{aligned} G(b_i^*)(v_i - b_i^*) &\geq G\left(\hat{b}_i^*\right)\left(v_i - \hat{b}_i^*\right) && (1) \\ G\left(\hat{b}_i^*\right)\left(\hat{v}_i - \hat{b}_i^*\right) &\geq G(b_i^*)(\hat{v}_i - b_i^*) && (2) \end{aligned}$$

Adding and cancelling gives

$$\begin{aligned} G(b_i^*)(v_i - b_i^*) + G\left(\hat{b}_i^*\right)\left(\hat{v}_i - \hat{b}_i^*\right) &\geq G(b_i^*)(\hat{v}_i - b_i^*) + G\left(\hat{b}_i^*\right)\left(v_i - \hat{b}_i^*\right) \\ \left(G\left(\hat{b}_i^*\right) - G(b_i^*)\right)(\hat{v}_i - v_i) &\geq 0. \end{aligned}$$

Since $\hat{v}_i > v_i$ we have $G\left(\hat{b}_i^*\right) - G(b_i^*) \geq 0$.

If $G\left(\hat{b}_i^*\right) - G(b_i^*) > 0$ then, since G is non-decreasing we must have $\hat{b}_i^* > b_i^*$.

If $G\left(\hat{b}_i^*\right) = G(b_i^*) > 0$ then we can divide equations (**??**) and (**??**) by $G(b_i^*) = G\left(\hat{b}_i^*\right)$ and get $\hat{b}_i^* \geq b_i^*$ and $\hat{b}_i^* \leq b_i^*$ so $\hat{b}_i^* = b_i^*$. (The intuition for this is simple: if $G\left(\hat{b}_i^*\right) = G(b_i^*)$ then both b_i^* and $\hat{b}_i^*$ win with the same probability, so [if this probability is not zero] regardless of value i would choose the lower bid so they cannot differ.)

If $G\left(\hat{b}_i^*\right) = G(b_i^*) = 0$ then both v_i and $\hat{v}_i$ are optimally bidding so low that they always lose. Then, in principle, they could bid differently. For exaple if they value the object at 5 and 10, but the opponent always bids 20 then they can bid anything below 20. ■

Remark 11 *The above type of argument is very common. In the example of the auction one can actually prove that the strategy must be strictly increasing but we will not. Note that since s is monotonic s it is differentiable almost everywhere.*

We now compute a Bayesian-Nash equilibrium of the first price auction. Suppose that for each bidder to bid $\beta(x_i)$ where $\beta : [0, \omega] \to \mathbb{R}$ is increasing and differentiable is a Bayesian-Nash equilibrium of the first price auction.

[28] For example, if v_j is uniform on $[0,1]$ (i.e. F is uniform on $[0,1]$) and j bids $s_j(v_j) = v_j/2$ then G is uniform on $[0, 1/2]$. If instead $s_j(v_j) = v_j^2$ then $G(x) = \Pr\left\{v_j : v_j^2 \leq x\right\} = \Pr\left\{v_j : v_j \leq \sqrt{x}\right\} = F(\sqrt{x}) = \sqrt{x}$.

A heuristic computation of β. The expected payoff to bidder 1 with valuation x who bids b when other bidders bid according to β is given by

$$\Pr(1 \text{ wins with } b) \cdot (x-b) = G\left(\beta^{-1}(b)\right) \cdot (x-b)$$

where $G = F^{N-1}$ denotes the distribution function of the random variable Y_1, which is the maximum of $N-1$ independently drawn valuations that are drawn according to F.[29]

Maximizing this expression with respect to b yields the first-order condition:

$$\frac{g\left(\beta^{-1}(b)\right)}{\beta'\left(\beta^{-1}(b)\right)}(x-b) - G\left(\beta^{-1}(b)\right) = 0$$

where $g = G'$ denotes the density of Y_1 (recall that for any function $f : X \to Y$, $\frac{df^{-1}(y)}{dy} = \frac{1}{f'(x)}$). Because bidding according to β is an equilibrium, $b = \beta(x)$ (or $\beta^{-1}(b) = x$) the previous equation yields the following differential equation:

$$\begin{aligned} \frac{g(x)}{\beta'(x)}(x-\beta(x)) - G(x) &= 0 \\ G(x)\beta'(x) + g(x)\beta(x) &= xg(x) \end{aligned}$$

or

$$\frac{d}{dx}\left(G(x)\beta(x)\right) = xg(x).$$

Integrating both sides according to x yields:

$$G(x)\beta(x) = \int_0^x yg(y)\,dy + C$$

[29]Observe that

$$\begin{aligned} \Pr(1 \text{ wins with } b) &= \Pr(b > \beta(x_2), ..., \beta(x_n)) \\ &= \Pr(b > \beta(x_2)) \cdots \Pr(b > \beta(x_n)) \\ &= \Pr\left(x_2 < \beta^{-1}(b)\right) \cdots \Pr\left(x_n < \beta^{-1}(b)\right) \\ &= \Pr\left(x < \beta^{-1}(b)\right)^{N-1} \\ &= F\left(\beta^{-1}(b)\right)^{N-1} \\ &= G\left(\beta^{-1}(b)\right). \end{aligned}$$

(observe that the fact that $G(0) = 0$ implies that $C = 0$) or[30]

$$\begin{aligned} \beta(x) &= \frac{1}{G(x)} \int_0^x y g(y)\, dy \\ &= E[Y_1 | Y_1 \leq x]. \end{aligned}$$

Remark. This derivation is heuristic because the differential equation is only a necessary condition for equilibrium.

Remark. Observe that the fact that F is continuous and increasing implies that $E[Y_1 | Y_1 \leq x] < x$ for $x > 0$. This formula also shows that the bid is increasing with N because the first order statistic Y_1 increases with N.

We verify that $\beta^{\mathrm{I}}(x) = E[Y_1 | Y_1 \leq x]$ is indeed a Bayesian-Nash equilibrium of the first-price auction. Suppose that all $N - 1$ bidders bid according to β^{I}. We show that to bid according to β^{I} is a best response. It is not optimal to bid $b > \beta^{\mathrm{I}}(\omega)$. The expected payoff to a bidder who has valuation x if she bids b is calculated as follows. Denote $\beta^{\mathrm{I}}(z) = b$ or $z = \left(\beta^{\mathrm{I}}\right)^{-1}(b)$.

$$\begin{aligned} \Pi(b, x) &= G(z)\left(x - \beta^{\mathrm{I}}(z)\right) \\ &= G(z)\, x - G(z)\, E[Y_1 | Y_1 \leq z] \\ &= G(z)\, x - \int_0^z y g(y)\, dy \\ &= G(z)\, x - G(z)\, z + \int_0^z G(y)\, dy \\ &= G(z)(x - z) + \int_0^z G(y)\, dy \end{aligned}$$

where the 4th equality follows from integration by parts.

Sufficiency follows from the fact that

$$\begin{aligned} \Pi\left(\beta^{\mathrm{I}}(x), x\right) - \Pi\left(\beta^{\mathrm{I}}(z), x\right) &= \int_0^x G(y)\, dy - \left(G(z)(x - z) + \int_0^z G(y)\, dy\right) \\ &= G(z)(z - x) - \int_x^z G(y)\, dy \\ &\geq 0 \end{aligned}$$

[30] Note that by L'Hopital's rule:

$$\lim_{x \searrow 0} \beta(x) \overset{\mathcal{L}}{=} \frac{0 \cdot g(0)}{g(0)} = 0.$$

regardless of whether $z > x$ or $z < x$ (demonstrate this on a figure with a plot of G).

Remark 12 Notice that all but the lowest type have strictly positive gains—*this is their "information rent", the profits they get because the seller does not know their type.* The lowest type earns zero.

Reserve Price. If the seller sets a reserve price $r > 0$, then bidders with valuations below r cannot possibly win. A bidder with valuation r bids $\beta^{\mathrm{I}}(r) = r$ in equilibrium (because by bidding r it wins if every other bidder has a valuation below r). The analysis above can be repeated to show that in this case $\beta^{\mathrm{I}}(x) = E[\max\{Y_1, r\} | Y_1 \le x]$ for $x \ge r$ and zero otherwise is a Bayesian-Nash equilibrium of the first-price auction. The fact that $E[\max\{Y_1, r\} | Y_1 \le x] > E[Y_1 | Y_1 \le x]$ for $x \ge r$ suggests that the seller may be able to increase its expected revenue by setting a positive reserve price.

Uniqueness of Equilibrium. See Lebrun (IER, 1999) for a proof that an equilibrium exists for the first price auction in private value environments and for sufficient conditions it is unique. The method of proof used by Lebrun is based on the mathematical theory of existence and uniqueness of solutions to systems of partial differential equations.

Second-Price Auction (Vickrey Auction)

Description. Bidders submit their bids simultaneously. The highest bidder wins the object and pays the second highest bid. Other bidders pay nothing. In case of a tie, the winning bidder is chosen randomly from among those submitted the highest bid.

Bidder i's bid is denoted b_i. The rules of the auction define a Bayesian game as follows:

- Players: $N = \{1, ...n\}$;
- Actions (possible bids): $B_i = \{b_i \ge 0\}$;
- Types: $V_i = \{v_i \ge 0\}$;
- Prior: We assume that bidders' valuations are i.i.d. Because of independence,

$$\Pr(v_1 \le \overline{v}_1, v_2 \le \overline{v}_2, ..., v_n \le \overline{v}_n) = \prod_{i=1}^{n} F(\overline{v}_i).$$

– Payoffs:

$$u_i(v,b)=\begin{cases} v_i-\max\{b_j:j\neq i\} & b_i>\max\{b_j:j\neq i\} \\ \frac{v_i-b_i}{\#\left\{j:b_j=\max\limits_{k\in\{1,\ldots,n\}}\{b_k\}\right\}} & b_i=\max\{b_j:j\neq i\} \\ 0 & b_i<\max\{b_j:j\neq i\} \end{cases}$$

A strategy for bidder i is given by a bidding function $b_i(v_i): V_i \longmapsto B_i$ that specifies a bid for whatever valuation bidder i may have for the object.

We saw that in a first price auction, bidders bid below their valuations. How should they bid in a second price auction?

Claim. Bidding the true valuation, $b_i(v_i)=v_i$, is a (weakly) dominant strategy for each bidder.

Proof. Distinguish among the following three cases. (1) $\max\{b_j:j\neq i\}>v_i$. Bidder i is better off not winning the object because the only way it can win is by paying more than the object is worth to her. Hence bidding v_i is optimal. (2) $\max\{b_j:j\neq i\}<v_i$. Bidder i wins with the truthful bid v_i and pays only $\max\{b_j:j\neq i\}$ which is independent of his bid. Hence, bidder i cannot do better by changing his bid. (3) $\max\{b_j:j\neq i\}=v_i$. Bidder i's payoff is zero in this case regardless of his bid, hence any bid, and in particular v_i, is optimal. To complete the proof we need to show that truthful bidding is strictly better than any other alternative for some profile of other bidders' bids. ■

The great advantage of the second price auction is that it considerably simplifies bidders' decision problems about how much to bid because bidding the true valuation is a dominant strategy regardless of the bidders' beliefs. In spite of this advantage, the second price auction is seldom used in practice. This is perhaps due to:

1. bidders are reluctant to reveal their true valuations because they fear this will make them vulnerable in the future,

2. auctions in which it is clear that the object is sold for a low price although someone is willing to pay a lot may be embarrassing for the seller, and

3. it is difficult for the winning bidder to verify that the seller tells the truth about the second highest bid.

Remark. The second price auction also has another equilibrium. Suppose for example that $n=2$, and $v_1,v_2\sim U[0,1]$ independently. Observe that $b_1(v_1)\equiv 0$ and $b_2(v_2)\equiv 1$ is also a

BNE of the second price auction. This BNE is less plausible than the BNE described above because it is supported by weakly dominated strategies.

Reserve Price. The setting of a positive reserve price by the seller has no effect on the bidders' incentives to bid truthfully. Bidding the true willingness to pay is still a dominant strategy for the bidders. A positive reserve price may nevertheless increase the expected revenue to the seller in the event that the second highest bid falls below the reserve price.

Revenue Equivalence

The four most widely used auction mechanisms are the first-price, second-price, English, and Dutch auctions.

Dutch Auction. The price decreases continuously starting from some high prespecified level following a "price clock." The first bidder who stops the price clock wins the object and pays the price at which she stopped the clock. Other bidders pay nothing. This auction form has the advantage of being very quick. It was used to sell tulip bulbs in Holland (hence its name) and is used to sell fish in Japan.

English Auction. An auctioneer keeps raising the price as long as bidders indicate they are willing to pay more. The auction ends when bidders are unwilling to bid up the price any further. The highest bidder wins, and pays the amount it bid. Other bidders pay nothing. This auction format has several variants that are distinguished by the information that is conveyed to the bidders with respect to other bidders' previous bids. One version is the one described here which is often used in art auctions. Another version is the one used in E-Bay where the auctioneer is replaced by an automated system.

Remark. The First-price and Dutch auctions are equivalent. And, the second-price and English auction are "strategically" equivalent in private value settings.

Moreover, recall the second price auction. In the symmetric equilibrium with two bidders and a uniform distribution over the interval $[0, V]$ type v bids v and wins with probability $\frac{v}{V}$. What is that types expected payment? It is the expected value of the opponent's value conditional on that value being lower than v. This is $E(v_j | v_j < v) = v/2$. So the expected payment of type v conditional on winning is $v/2$ and the expected payment (unconditional) is $s(v) F(v) = \frac{v}{2} \times \frac{v}{V}$, which is identical to bidder v's expected payment under the first price auction. This identical result is not a coincidence, but an example of the Revenue-Equivalence Theorem.

It turns out that the equivalence between different types of auction is a much more general phenomenon.

Revenue Equivalence Theorem (Myerson, 1981). In independent private-value settings with continuous and increasing distribution functions, any two auction mechanisms A and B that (1) induce identical probabilities of winning as a function of the bidders' types, $p_i^A(v_1, ..., v_n) = p_i^B(v_1, ..., v_n)$ for every bidder i and bidders' valuations $v_1, ..., v_n$, and (2) give the lowest type of each of the bidders the same expected payoffs, generate the same expected revenue for the seller.

Corollary. All the four auction mechanisms mentioned above generate the same expected revenue to the seller.

Two other important remarks are:

1. All the four auction mechanisms mentioned above are efficient (they all assign the object to the bidder with the highest valuation). However, while the second price and English auctions continue to be efficient also in asymmetric environments (where bidders draw their valuations from different distributions), the first price auction is not. Can you see why?

2. It can be shown (see Myerson, 1981, for details) that, if in addition, bidders valuations are i.i.d., then all the four auction mechanisms mentioned above are optimal provided the seller sets an appropriately chosen reservation price.

5.4.2 Common Value Auctions and the Winner's Curse

Example. Consider an English auction, say, of an old phonebook or another essentially valueless book with many pages. Suppose that I pay 1 cent per page of the book. How much are you willing to bid?

In many cases, bidders' valuations have common elements (e.g., an oil well has the same value for all bidders; a painting has identical resale value to all bidders, etc.). This common element gives rise to the winner's curse. Intuitively, the winner of the auction is the bidder who has the *most optimistic* signal/opinion about the value of the object. But a more precise estimate of the value of the object is obtained by aggregating the information contained in *all* the bidders' signals, which are lower.

Assume that $v_i = v + \varepsilon_i$ with $\varepsilon_i \sim F$ and $E(\varepsilon_i) = 0$; each bidder has the same common value but observes a noisy (mean zero) signal. (Also called a mineral-rights auction.) We begin with studying the new feature that arises and solve a two-value problem.

In the first price auction it is *not* an equilibrium to use the equilibrium strategies as a function of the expected value based on one's own signal. The expected value is v_i. If bidders' strategies are monotone in their estimates, and i wins that means $v_i > v_j$. The expected value conditional on the information that i wins is lower! Specifically $E(v|v_i, v_j) = (v_i + v_j)/2$. Namely, conditional on v_i and $v_j < v_i$ the expectation of v is lower than v_i. The first-price symmetric optimal strategy $b^*(v)$ for the independent-private values case calculated the optimal trade-off between the probability of winning on the one hand, and the price paid when winning on the other hand, *assuming the value was* v. But now the value conditional on winning is lower, so the strategy cannot be optimal.

Example: Second Price Auction. Suppose that $n = 2$, and that each bidder observes an independent signal $x_1, x_2 \sim U[0,1]$ about the value of the object. The value to both bidders is given by $v_1 = v_2 = x_1 + x_2$. The winner's curse is the name given to describe the difference between a bidder's "naive" expected valuation of the object

$$x_i + E[x_j] = x_i + \frac{1}{2}$$

and the bidder's valuation of the object conditional on (the information deduced from) winning the object

$$\begin{aligned} x_i + E[x_j | x_j \leq x_i] &= x_i + \frac{x_i}{2} \\ &< x_i + \frac{1}{2}. \end{aligned}$$

(notice that we assume that bidders employ symmetric bidding strategies here). For small values of x_i this difference can be very large because winning the auction with a low value of x_i conveys very negative information about the value of x_j.

However, in equilibrium, rational bidders (are supposed to) anticipate the winner's curse and therefore should not be subject to it.

Suppose that the bidders participate in a second-price auction. We search for a symmetric Bayesian-Nash equilibrium in increasing linear strategies, $b_i(x_i) = a + cx_i$, $a \geq 0, c > 0$.[31]

Because by assumption $c > 0$ the probability that the two bidders' bids tie is zero. It follows that in equilibrium, it must be that

$$b(x_i) \in \arg\max_{b_i} E[x_i + x_j - b_j | b_i \geq b_j] \Pr(b_i \geq b_j).$$

[31] The example also has a continuum of asymmetric equilibria given by $b_1(x_1) = c_1 x_1$ and $b_2(x_2) = c_2 x_2$ where c_1 and c_2 are such that

$$c_1 + c_2 = c_1 c_2.$$

Now,

$$\begin{aligned}
\Pr\left(b_i \geq b_j\right) &= \Pr\left(b_i \geq a + cx_j\right) \\
&= \Pr\left(x_j \leq \frac{b_i - a}{c}\right) \\
&= \frac{b_i - a}{c}
\end{aligned}$$

provided $\frac{b_i - a}{c}$ is between 0 and 1. As before, this follows from the fact that a rational bidder would not bid less than a or more than $a + c$. Because i knows x_i and b_i, for i

$$\begin{aligned}
E\left[x_i + x_j - b_j \left| b_i \geq b_j\right.\right] &= E\left[x_i + x_j - a - cx_j \left| b_i \geq a + cx_j\right.\right] \\
&= x_i - a + (1 - c)\, E\left[x_j \left| x_j \leq \frac{b_i - a}{c}\right.\right] \\
&= x_i - a + (1 - c)\frac{b_i - a}{2c}.
\end{aligned}$$

It follows that

$$b\left(x_i\right) \in \arg\max_{b_i}\left(x_i - a + (1 - c)\frac{b_i - a}{2c}\right)\left(\frac{b_i - a}{c}\right),$$

or

$$\left(x_i - a + (1 - c)\frac{b_i - a}{2c}\right)\frac{1}{c} + \left(\frac{b_i - a}{c}\right)\left(\frac{1 - c}{2c}\right) = 0.$$

And upon rearranging,

$$b\left(x_i\right) = \frac{a}{1 - c} - \frac{cx_i}{1 - c}.$$

Our assumption that $b_i\left(x_i\right) = a + cx_i$ implies that

$$c = -\frac{c}{1 - c} \Longrightarrow c = 2$$

and

$$a = \frac{a}{1 - c} = -a \Longrightarrow a = 0.$$

In equilibrium,

$$b\left(x_i\right) = 2x_i.$$

It is instructive to compare the equilibrium bidding strategy with the "naive" estimate of the value of the object. For small values of x_i bidder i submits very low bids because winning the auction with such a low signal conveys very bad news regarding the value of the object; on the other hand, for large values of x_i bidder i submits very high bids relative to the bidder's "naive" estimate because losing the auction with such a high signal conveys very good news regarding the value of the object (or bad news about the value of not winning).

Remark. Bidding the true expected valuation is *not* a dominant strategy in the second price auction with common values. The proof that showed that bidding the true valuation is a dominant strategy in a second price auction in a private values setting relied on the fact that a change in a bidder's bid does not change the value of winning (because it does not affect the price paid upon winning). However, in a common values setting changing one's bid changes what one learns from winning, and therefore changes the expected value of winning the object. This implies that the proof that worked in the private values case does not work in the case of common values. To convince yourself of this point, show that $b_1(x_1) = 2x_1$ is not a best response for bidder 1 in the example above if bidder 2 employs a different bidding strategy than $b_2(x_2) = 2x_2$.

Example 57 *Consider a second-price auction and assume v is either* 1 *or* 2 *with probability* 1/2*, and that if it is* 1 *then everyone sees a low signal, while if it is 2 each bidder observes an independent signal that is the high signal with probability* 1/2 *and the low signal with probability* 1/2. *Then conditional on the high signal the probability that the value is 2 is* 1. *Conditional on one low signal* $\Pr\{v = 1|low\} = \frac{1/2}{1/2+1/4} = 2/3$ *so the expected value is* $2/3 + 2 \times (1/3) = 4/3$. *Conditional on two low signals* $\Pr\{v = 1|low,low\} = \frac{1/2}{1/2+1/8} = \frac{4}{5}$ *so the expected value is* $4/5 \times 1 + 1/5 \times 2 = 6/5 < 4/3$.

First we will see that bidding 4/3 *when the signal is low is not consistent with equilibrium. Conditional on a low signal and winning you know the opponent observed a low signal. But then* $\Pr\{v = 1|low,low,win\} = \Pr\{v = 1|low,low\}$ *and the expected value is lower.*

Now consider the case where both bid 6/5. *Then conditional on winning this is exactly the expected value. Bidding less obviously implies no gain, bidding more implies no gain as well (you win with higher probability but pay the expected value so continue to gain zero in expectation).*

Exercise 69 *Consider the same environment, but with a change in signals as follows. When* $v = 2$ *everyone sees a high signal; when* $v = 1$ *each bidder sees a low signal with probability 1/2 and a high signal with probability 1/2 (the signals are drawn independently). Is there a winner's curse?*

In the second-price auction the equilibrium is to bid $b(x) = E(v|\text{both signals are } x)$. (See Milgrom and Weber (*Econometrica*, 1982, A Theory of Auctions and Competitive Bidding, Theorem 6). For now we will see this in examples.

Example 58 *In the mineral-rights model above $E(v|\text{one signal is } x) = E(v|\text{both signals are } x)$, so there is no change. The intuition is that while conditional on winning the value goes down, so does the payment. But in general this is not the case.*

Example 59 ***Another common values example with a First Price Auction***

Assume x_i is iid according to F on $[0, X]$ and $v_i = x_1 + x_2$.

Consider a symmetric strictly increasing differentiable equilibrium.

$$s_i(x_i) \in \arg\max_b \Pr\{s_j(x_j) < b\} E(x_i + x_j - b|b > s_j(x_j))$$

where

$$\Pr\{s_j(x_j) < b\} = F\left(s_j^{-1}(b)\right)$$

and

$$\begin{aligned} E(x_i + x_j - b|b > s_j(x_j)) &= x_i - b + E(x_j|s_j(x_j) < b) \\ &= x_i - b + E\left(x_j|x_j < s_j^{-1}(b)\right) \end{aligned}$$

where the second equality follows if we assume s is strictly increasing. Assume also that F is uniform on $[0,1]$ (so $f = 1$ and $F(z) = z$) and conjecture a linear strategy for j: $s_j(x_j) = a_j x_j + c_j$, so $x_j = s_j^{-1}(b) = (b - c_j)/a_j$. Then the optimality condition becomes

$$\begin{aligned} s_i(x_i) &\in \arg\max_b \Pr\{s_j(x_j) < b\} E(x_i + x_j - b|b > s_j(x_j)) \\ &= \arg\max_b \left(s_j^{-1}(b)\right)\left(x_i - b + E\left(x_j|x_j < s_j^{-1}(b)\right)\right) \\ &= \arg\max_b \left(s_j^{-1}(b)\right)\left(x_i - b + s_j^{-1}(b)/2\right) \\ &= \arg\max_b \left(\frac{b - c_j}{a_j}\right)\left(x_i - b + \frac{b - c_j}{2a_j}\right) \end{aligned}$$

Taking the derivative and setting it equal to zero to find the optimal b for i given x_i yields $b = \frac{a_j x}{2a_j - 1} + \frac{a_j c_j - c_j}{2a_j - 1}$. So we obtain a linear solution. Note that if the outcome of this process weren't linear it would mean that when the opponent chooses a linear strategy the best reply need not be linear so the asumption of an equilibrium in linear strategies would be violated.

In a symmetric equilibrium $b = s(x_i) = ax_i + c = \frac{ax}{2a-1} + \frac{ac-c}{2a-1}$ which implies $a = 1$ and $c = 0$ and thus $s(x_i) = x_i$.

Contrast this with an "analogous" private values case where $v_i = x_i + E(x_j) = x_i + 1/2$. Then $s^{IPV}(x_i) = 1/2 + x_i/2$. (By analogy to what we have already solved.) So $s^{IPV}(x_i) >$

$s^{CV}(x_i)$ *(since* $1/2 > x_i/2$ *except for* $x_i = 1$*).*

Note that in the pure common-values case there is no question of efficiency. (There is still the question of whether the price converges to the true value in a large economy. It does.)

A major question about auctions and other mechanisms of allocation is whether they aggregate information efficiently. That is, assume that individuals have private information about the value to the other player. Will the outcome "reveal" all that information, in the sense that the price will reflect it and the winner will be the same as if it were public. There are several positive answers to this, when there are many players. This is a more advanced literature that we will not study

5.4.3 Public goods with incomplete information

There are two players. Each player knows their own cost of providing a public good, $c \in [\underline{c}, \bar{c}]$ distributed according to F. The benefit to each player from the public good is commonly known to be 1. The good is provided if either player contributes to it. It is convenient to depict the game as

	Contribute	Don't
C	$1 - c_1, 1 - c_2$	$1 - c_1, 1$
D	$1, 1 - c_2$	$0, 0$

The strategies are a function specifying for each type whether they contribute (1) or not (0), i.e., $s_i : [\underline{c}, \bar{c}] \to \{0, 1\}$, and

$$u_i(s_i, s_j, c_i, c_j) = \max\{s_1(c_1), s_2(c_2)\} - c_i s_i$$

Finding the equilibrium

In games in economic environments where there are finitely many actions and continuum of types the equilibrium strategies are often given by thresholds: types below a certain cutoff do one action, and above that cutoff they do another, and so on. Intuitively this follows from a monotonicity in payoffs. In fact, this monotonicity implies more generally that strategies are monotonic.

One way to think about equilibrium is in terms of whether a type would like to mimic–behave like–a different type. In this example, if for some beliefs about the opponent a type contributes then any type with a lower cost will want to contribute given those same beliefs, and if a type doesn't contribute given some beliefs then any higher type will not want to contribute with those same beliefs. Since beliefs in this game are independent, we have argued that equilibrium strategies are of the threshold type.

So assume the opponent chooses threshold c_j^*. Thus j's strategy is to contribute with lower types and not with higher types, and since we have a CDF what any particular type does is irrelevant as it occurs with probability 0. Formally $s_j^*(c_j) = 1$ iff $c_j \leq c_j^*$. Which type of i is indifferent? $u_i\left(\text{Contribute}, s_j^*, c_i, F\right) = 1 - c_i$; $u_i\left(\text{Don't}, s_j^*, c_i, F\right) = F\left(c_j^*\right)$, so $1 - c_i = F\left(c_j^*\right)$. In a symmetric equilibrium $1 - c^* = F(c^*)$. For example, if $F(c) = c/2$ (uniform on $[0, 2]$) then $c^* = 2/3$.

Exercise 70 *Understand that for F with support on $[a, b]$ with $0 \leq a < b \leq 1$ there is a symetric equilibrium and (two) asymmetric equilibria.*

For F uniform with support on $[a, b]$ with $0 \leq a < b \leq 1$ compare the expected utilities of type $v \in [a, b]$ for each player in the asymmetric equilibria with those in the symmetric equilibirum. Now compare the ex ante expected utility (integrate what you get times the density over v in $[a, b]$). Can you say anything about general F?

For F with support on $[0, k]$ for $k > 1$ argue that there is no equilibrium in which one player always behaves the same way regardless of his type.

Exercise 71 *Solve the uniform case on $[0, 2]$ with $n > 2$ players. What happens as $n \to \infty$?*

Exercise 72 *Solve this with $n = 2$ but private information about values; costs are $1/2$ and values are distributed according to G.*

Exercise 73 *Harder (but not optional). Solve this with $n = 2$ but private information about costs and values; costs are distributed according to F and values according to G.*

5.4.4 Auctions

First price independent private value auction (IPV)

Assume that each player's *private* value is *independently* distributed according to some F with density f. The first price auction is an auction where the highest bidders wins the object and pays its bid. Losers pay nothing. A strategy is a bid for each possible value. We will focus on equilibria in continuous and strictly increasing strategies. (One can argue that any equilibrium of this game must have continuous strictly increasing strategies.) Since they are strictly increasing and F has a density the probability of a tie is zero so we don't need to worry about those.

Before solving for the equilibrium let's see why the strategies must be monotonic (non-decreasing). Assume some type v optimally wants to bid b. Consider a type who values the good more, $\hat{v} > v$. That type cannot optimally bid less. By bidding less the probability of winning goes down and so does the payment. But a type with a higher value cares more about winning so that the decrease in probability must be at least as costly to this type as to v, and the savings is the same. This property is called the single-crossing property–the marginal willingness to pay is increasing in the player's type.

Formally, assume the opponent adopts some strategy s_j. This generates a probability distribution over bids, and in particular determines for each bid b the probability that i wins if he bids b, denoted by $G(b)$.[32] Of course G is non-decreasing. Assume that type v_i optimally bids b_i^* against s_j and that $\hat{b}_i^*$ is the optimal bid for $\hat{v}_i > v_i$.

Claim 4 *$\hat{b}_i^* \geq b_i^*$ except if $G(b_i^*) = G\left(\hat{b}_i^*\right) = 0$, i.e., except if both types lose with probability 1 when they bid optimally.*

Proof. Since b_i^* is optimal for v_i is it better than bidding anything else, in particular better than bidding $\hat{b}_i^*$, and conversely.

$$G(b_i^*)(v_i - b_i^*) \geq G\left(\hat{b}_i^*\right)\left(v_i - \hat{b}_i^*\right) \tag{3}$$
$$G\left(\hat{b}_i^*\right)\left(\hat{v}_i - \hat{b}_i^*\right) \geq G(b_i^*)(\hat{v}_i - b_i^*) \tag{4}$$

adding and cancelling gives

$$\begin{aligned} G(b_i^*)(v_i - b_i^*) + G\left(\hat{b}_i^*\right)\left(\hat{v}_i - \hat{b}_i^*\right) &\geq G(b_i^*)(\hat{v}_i - b_i^*) + G\left(\hat{b}_i^*\right)\left(v_i - \hat{b}_i^*\right) \\ \left(G\left(\hat{b}_i^*\right) - G(b_i^*)\right)(\hat{v}_i - v_i) &\geq 0. \end{aligned}$$

Since $\hat{v}_i > v_i$ we have $G\left(\hat{b}_i^*\right) - G(b_i^*) \geq 0$.

If $G\left(\hat{b}_i^*\right) - G(b_i^*) > 0$ then, since G is non-decreasing we must have $\hat{b}_i^* > b_i^*$.

If $G\left(\hat{b}_i^*\right) = G(b_i^*) > 0$ then we can divide equations (**??**) and (**??**) by $G(b_i^*) = G\left(\hat{b}_i^*\right)$ and get $\hat{b}_i^* \geq b_i^*$ and $\hat{b}_i^* \leq b_i^*$ so $\hat{b}_i^* = b_i^*$. (The intuition for this is simple: if $G\left(\hat{b}_i^*\right) = G(b_i^*)$ then both b_i^* and $\hat{b}_i^*$ win with the same probability, so [if this probability is not zero] regardless of value i would choose the lower bid so they cannot differ.)

If $G\left(\hat{b}_i^*\right) = G(b_i^*) = 0$ then both v_i and $\hat{v}_i$ are optimally bidding so low that they always lose. Then, in principle, they could bid differently. For exaple if they value the object at 5 and 10, but the opponent always bids 20 then they can bid anything below 20. ■

[32] For example, if v_j is uniform on $[0,1]$ (i.e. F is uniform on $[0,1]$) and j bids $s_j(v_j) = v_j/2$ then G is uniform on $[0,1/2]$. If instead $s_j(v_j) = v_j^2$ then $G(x) = \Pr\left\{v_j : v_j^2 \leq x\right\} = \Pr\{v_j : v_j \leq \sqrt{x}\} = F(\sqrt{x}) = \sqrt{x}$.

Remark 13 *The above type of argument is very common.* I have posted notes about the general connection between single crossing and monotonic best replies. In the example of the auction one can actually prove that the strategy *must* be *strictly* increasing but we will not. Note that since it is monotonic s is differentiable almost everywhere.

The expected utility of i bidding b_i when i's value is v and the opponent bids according s_j and has distribution F, is $u_i(b_i, s_j, v_i, F) = F\left(s_j^{-1}(b_i)\right)(v_i - b_i)$. Assume for simplicity that F is uniform on $[0, V]$. Then the payoff is $\left(s_j^{-1}(b_i)/V\right)(v_i - b_i)$. The derivative w.r.t b_i is $\frac{v_i - b_i}{V s_j'\left(s_j^{-1}(b_i)\right)} - \frac{s_j^{-1}(b_i)}{V}$. At a symmetric equilibrium $s_j(v_i) = s_i(v_i) = s(v_i) = b_i$ so we have $\frac{v - s(v)}{V s'(s^{-1}(s(v)))} - \frac{s^{-1}(s(v))}{V} = 0$ i.e., $\frac{v - s(v)}{V s'(v)} - \frac{v}{V} = 0$, so $v - s(v) - v s'(v) = 0$. The solution is $s(v) = v/2$. So the expected payment of type v conditonal on winning is $v/2$ and the expected payment (unconditional) is $s(v)F(v) = \frac{v}{2} \times \frac{v}{V}$.

Remark 14 *Note that strictly increasing and symmetric guarantees efficiency—the person with the higher value wins.*

For some more intuition let's drop the assumption that F is uniform. The first-order condition is $\frac{f\left(s_j^{-1}(b_i)\right)}{s_j'\left(s_j^{-1}(b_i)\right)}(v_i - b_i) - F\left(s_j^{-1}(b_i)\right)$, and at a symmetric equilibrium $F(v)s'(v) + f(v)s(v) = vf(v)$, or $\frac{d}{dv}(F(v)s(v)) = vf(v)$. Observe that $s(0) = 0$. [Proof: If $s(0) = \varepsilon > 0$ then $s(v) \geq \varepsilon$ for $v \in (0, \varepsilon)$ and then $s(\varepsilon)$ wins with strictly positive probability at a bid greater than ε, which is not optimal.] Apply $\int_0^x$ to both sides to get $F(x)s(x) = \int_0^x vf(v)\,\mathrm{d}v$ or $s(x) = \int_0^x vf(v)\,\mathrm{d}v/F(x) = \mathrm{E}(v|v < x)$. [Note for future reference that this is related to integrating by parts.] The intuition is that in an equilibrium with strictly increasing strategies you know you lose against a bidder with values above yours. So when you win it must be against a bidder with lower values, hence you condition on $v < x$. Moreover, *very roughly* speaking, if you bid below the conditional expected value (e.g., by bidding the conditional expected bid) then an opponent with lower value would gain by trying to bid more than you.

The above is only suggestive since we used the necessary FOC, which we didn't show are sufficient. The proof that this is an equilibrium is as follows. In this proof we use again the "trick" where instead of thinking of a person with value v deviating from $b = s(v)$ to some other bid b', we think of the person behaving like another type v' such that $s(v') = b'$. Note that in general it is not enough to check for such deviations since there might be bids that no type makes, but the equilibrium we are considering is continuous and atomless so the only bids not made are greater than the maximal bid and lower than the minimal bid and those are not optimal. [Bidding the maximal bid wins with probability 1 because the strategy is atomless; bidding the lowest bid loses with probability 1.] So bidding $s(v')$ with value v

yields $F(v')(v - s(v')) = F(v')v - F(v')\,\mathrm{E}(v|v < v')$ [since the opponent plays this strategy by assumption] $= F(v')v - \int_0^v xf(x)\,\mathrm{d}x = F(v')v - F(v')v' + \int_0^{v'} F(x)\,\mathrm{d}x$ [integrating by parts — see above] $= F(v')(v - v') + \int_0^{v'} F(x)\,\mathrm{d}x$. The benefit from bidding v is just the second term, $\int_0^v F(x)\,\mathrm{d}x$, since the first is zero. So the change in payoff from bidding v instead of v' is $\int_0^v F(x)\,\mathrm{d}x - F(v')(v - v') - \int_0^{v'} F(x)\,\mathrm{d}x = F(v')(v' - v) - \int_v^{v'} F(x)\,\mathrm{d}x$ which is non-negative (note that F is increasing). [Draw a picture.]

Remark 15 Recall the second price auction. In the symmetric equilibrium type v bids v and wins with probability $\frac{v}{V}$. What is that types expected payment? It is the expected value of the opponent's value conditional on that value being lower than v. This is $E(v_j|v_j < v) = v/2$. So the expected payment of type v conditional on winning is $v/2$ and the expected payment (unconditional) is $s(v)F(v) = \frac{v}{2} \times \frac{v}{V}$. This identical result is not a coincidence, but an example of the revenue-equivalence theorem which we will return to shortly.

Remark 16 Notice that all but the lowest type have strictly positive gains—*this is their "information rent", the profits they get because the seller does not know their type.* The lowest type earns zero.

Exercise 74 *Solve the above for the n player auction. As n goes to infinity what do bids converge to?*

5.4.5 Wars of attrition and all-pay auctions

$u_i(s_i, s_j, v_i, v_j) = -b_i$ if $b_j \geq b_i$ and $u_i(s_i, s_j, v_i, v_j) = v_i - b_j$ if $b_j < b_i$

Again we look for a strictly increasing continuous symmetric equilibrium. That any equilibrium is non-decreasing you can prove–be sure you see how.

$u_i(b, s_j, v_i, F) = \int_0^{s_j^{-1}(b)} (v_i - s_j(v_j)) f(v_j)\,dv_j - b\left(1 - F\left(s_j^{-1}(b)\right)\right)$

We will not solve this further for now.

Exercise 75 *Very optional / quite involved. Find the FOC. Substitute to find a candidate equation for a symmetric equilibrium: $s(v) = \int_0^v \frac{xf(x)}{1-F(x)}dx$. (Remark: If F is exponential ($F(v) = 1 - e^{-v}$) the bid $s(v) = v^2/2$ solves the equation.)*

Exercise 76 *Describe the payoffs from actions given values — $u_i(s_i, s_j, v_i, v_j)$ — and the overall payoffs given a distribution of the opponent and the opponents" strategy — $u_i(b, s_j, v_i, F)$ — for the following game: each player submits a bid and pays their bid, and wins iff they bid more than anyone else.*

Exercise 77 *Both games above (in the exercise and before these two exercises) correspond to a situation where both players make payments that depend on the bids, and in which the higher bidder wins. What is the critical difference in the environment being modeled that would determine which game is a better model?*

5.4.6 Revenue equivalence

Consider any "auction-like" game in which each player i chooses an action $b_i \in \mathbf{R}$ as a function i's type v_i and in which the game specifies for any profile of actions b who gets the object and how much everyone pays.

Let $\tau_i(b) : \mathbf{R}^n \to \mathbf{R}$ and $\chi_i(b) : \mathbf{R}^n \to [0,1]$ denote the expected payment and probability of getting the good for i, when everyone bids $b_1, ..., b_n$. (Of course $\Sigma_i \chi_i(b) = 1$.)

In equilibrium each player chooses a (possibly random) bid given their value $s_i : \mathbf{R} \to \Delta(\mathbf{R})$.

Fix henceforth some equilibrium of the game. Given this equilibrium we can find the induced functions $t_i(v) : \mathbf{R}^n \to \mathbf{R}$ and $q(v) : \mathbf{R}^n \to [0,1]$ which specify the expected payment for i and probability of i receiving the object given a vector of values (types) and given the equilibrium strategies. That is, $t_i(v) = \mathrm{E}\tau(s_1(v_1), ..., s_n(v_n))$ (where the expectation is over the bids in case s_i is mixed for some i) and similarly for b_i.

Finally, let $T_i(v_i)$ and $Q_i(v_i)$ be the expected payment and probability of winning for i: $T_i(v_i) = \mathrm{E}(t_i(v)|v_i) = \int_{v_{-i}} t_i(v_{-i}, v_i) F(v_{-i}|v_i) = \int_{v_{-i}} t_i(v_{-i}, v_i) F(v_{-i}) = \int \cdots \int t_i(v_{-i}, v_i) \left[\prod_{j \neq i} d\mathrm{F}(v_j)\right]$, where the equality before last uses independence, and similarly, $Q_i(v_i) = \mathrm{E}(q_i(v)|v_i) = \int_{v_{-i}} q_i(v_{-i}, v_i) F(v_{-i})$.

What are v_i's payoffs from this equilibrium? $U_i(v_i) = v_i Q_i(v_i) - T_i(v_i)$.

In the game any player i of type v_i could deviate to any other action b. In equilibrium of course such a deviation is not profitable. In particular i could deviate to the action chosen by a different type of i, $\tilde{v}_i$. This is exactly the same type of analysis we adopted above–we consider v_i deviating to behave like $\tilde{v}_i$. What are v_i's payoffs from such a deviation to $\tilde{v}_i$? They are

$$U(\tilde{v}_i|v_i) = v_i Q_i(\tilde{v}_i) - T_i(\tilde{v}_i).$$

So in equilibrium the following *incentive compatibility (IC)* conditions must hold: for every i and $\tilde{v}_i$:

$$v_i Q_i(\tilde{v}_i) - T_i(\tilde{v}_i) \leq v_i Q_i(v_i) - T_i(v_i).$$

Therefore

$$U_i(v_i|v_i) = U_i(v_i) = \max_{\tilde{v}_i} U_i(\tilde{v}_i|v_i) = \max_{\tilde{v}_i} v_i Q(\tilde{v}_i) - T_i(\tilde{v}_i).$$

So $U_i(v_i)$ is a maximum among linear functions which means it is the upper envelope of a collection of linear functions which means it is convex. [To see geometrically why consider the figure. The three lines represent $U_i(\tilde{v}_i|v_i)$ functions for different values of $\tilde{v}_i$. These functions are linear in v_i. For each v_i you chose the $\tilde{v}_i$ that gives the maximal value. The maximal one is the upper envelope, i.e. starts with the green line, then is the red one and finally is the black one — a convex function.]

put figure here

Convex functions are absolutely continuous which means that they are almost everywhere differentiable and the integral of their derivative equals the function $U_i(v_i) = U_i(0) + \int_0^{v_i} U_i'(t)\,\mathrm{d}t$.

By the envelope theorem $U_i'(t) = \frac{d}{dt}\max_{\tilde{v}_i} U_i(\tilde{v}_i|t) = \frac{d}{dt}\max_{\tilde{v}_i} tQ(\tilde{v}_i) - T_i(\tilde{v}_i) = Q(t)$ where the last equality follows since in equilibrium it is $\tilde{v}_i = t$ that maximizes $U_i(\tilde{v}_i|t)$. So in *any* equilibrium the utility of type v_i of i depends only on the payoff that the lowest value type gets, $U_i(0)$ and the probability with which the different types win.

$$U_i(v_i) = U_i(0) + \int_0^{v_i} Q_i(t)\,\mathrm{d}t$$

So, for instance, in any efficient equilibrium, in which the highest value wins, so that in all efficient equilibrium $Q_i(\cdot)$ is the same function, the utility of any type of player i is the same up to a constant, $U_i(0)$. Since $U_i(0)$ gets utility 0 from the auction itself, if we cannot force participation (this is called the participation constraint) and if don't pay for participation, we have $U_i(0)$ equals the "reservation utility" (the utility they would get from not participating in the auction). (In our case zero.) Thus, in all efficient mechanisms in which the lowest value type gets their reservation value, all types get the same expected utility.

Since the utility of type v_i is the probability of winning times the value, $Q(v_i)v_i$, less the expected payment, $T_i(v_i)$, the expected payment of each type must be the same.

Hence also the expected payments received by the auctioneer must be the same.

"Application" Consider the war of attrition game of section **??**. Assume there exists an increasing symmetric equilibrium. Show that the proposed strategy, $s(v) = \int_0^v \frac{xf(x)}{1-F(x)}\mathrm{d}x$, in exercise **??** is this equilibrium.

By revenue equivalence the expected payment in such an equilibrium must be $\int_0^v xf(x)\,\mathrm{d}x$ (the expected payment in the symmetric equilibrium of the second-price auction in which everyone bids their value).

The expected payment when everyone adopts bidding strategy $s(v)$ in this war of attrition is $\int_0^v s(t)\,\mathrm{d}F(t) + (1 - F(v))\,s(v)$. (In the game you pay the opponents bid when you bid more, i.e., when your value is higher in an increasing equilibrium, and you pay your bid when the opponent bids more.)

By the revenue equivalence theorem we must have $\int_0^v s(t)\,\mathrm{d}F(t) + (1 - F(v))\,s(v) = \int_0^v xf(x)\,\mathrm{d}x$. Taking derivatives we have $s(v)\,f(v) - f(v)\,s(v) + (1 - F(v))\,s'(v) = vf(v)$. Therefore $s'(v) = vf(v)/(1 - F(v))$ and hence $s(v) = \int_0^v \frac{xf(x)}{1-F(x)}\mathrm{d}x$.

Exercise 78 *Use this method to find the symmetric equilibrium in the first price auction.*

5.4.7 Common values and auctions—the winner's curse

Assume that $v_i = v + \varepsilon_i$ with $\varepsilon_i \sim F$ and $E(\varepsilon_i) = 0$; each bidder has the same common value but observes a noisy (mean zero) signal. (Also called a mineral-rights auction.) We will not solve the general auction, but understand the new feature that arises and solve a two-value problem.

In the first price auction it is *not* an equilibrium to use the equilibrium strategies as a function of the expected value based on ones own signal. The expected value is v_i. If you bid v_i and win that means $v_i > v_j$. The expected value conditional on the information that you win is lower! Specifically $E(v|v_i, v_j) = (v_i + v_j)/2$, conditional on v_i and $v_j < v_i$ the expectation of v is lower than v_i. The first-price symmetric optimal strategy $b^*(v)$ for the independent-private values case calculated the optimal trade-off between the probability of winning on the one hand, and the price paid when winning on the other hand, *assuming the value was* v. But now the value conditional on winning is lower, so the strategy cannot be optimal.

In the second-price auction the equilibrium is to bid $b(x) = E(v|\text{both signals are } x)$. (See Milgrom and Weber (*Econometrica*, 1982, A Theory of Auctions and Competitive Bidding, Theorem 6). For now we will see this in examples.

Example 60 *In the mineral-rights model above* $E(v|\textit{one signal is } x) = E(v|\textit{both signals are } x)$, *so there is no change. The intuition is that while conditional on winning the value goes down, so does the payment. But in general this is not the case.*

Example 61 *Consider then a second-price auction and assume* v *is either* 1 *or* 2 *with probability* 1/2, *and that if it is* 1 *then everyone sees a low signal, while if it is* 2 *each* *bidder observes an independent signal that is the high signal with probability* $1/2$ *and the low signal with probability* $1/2$. *Then conditional on the high signal the probability that the value is 2 is* 1. *Conditional on one low signal* $\Pr\{v=1|low\} = \frac{1/2}{1/2+1/4} = 2/3$ *so the expected value is* $2/3 + 2\times(1/3) = 4/3$. *Conditional on two low signals* $\Pr\{v=1|low,low\} = \frac{1/2}{1/2+1/8} = \frac{4}{5}$ *so the expected value is* $4/5 \times 1 + 1/5 \times 2 = 6/5 < 4/3$.

First we will see that bidding $4/3$ *when the signal is low is not consistent with equilibrium. Conditional on a low signal and winning you know the opponent observed a low signal. But then* $\Pr\{v=1|low,low,win\} = \Pr\{v=1|low,low\}$ *and the expected value is lower.*

Now consider the case where both bid $6/5$. *Then conditional on winning this is exactly the expected value. Bidding less obviously implies no gain, bidding more implies no gain as well (you win with higher probability but pay the expected value so continue to gain zero in expectation).*

Exercise 79 *Consider the same environment, but with a change in signals as follows. When* $v=2$ *everyone sees a high signal; when* $v=1$ *each bidder sees a low signal with probability 1/2 and a high signal with probability 1/2 (the signals are drawn independently). Is there a winner's curse?*

Example 62 ***Another common values example***

Assume x_i *is iid according to* F *on* $[0,X]$ *and* $v_i = x_1 + x_2$.

Consider a symmetric strictly increasing differentiable equilibrium.

$$s_i(x_i) \in \arg\max_b \Pr\{s_j(x_j) < b\} E(x_i + x_j - b | b > s_j(x_j))$$

where

$$\Pr\{s_j(x_j) < b\} = F\left(s_j^{-1}(b)\right)$$

and

$$\begin{aligned} E(x_i + x_j - b | b > s_j(x_j)) &= x_i - b + E(x_j | s_j(x_j) < b) \\ &= x_i - b + E\left(x_j | x_j < s_j^{-1}(b)\right) \end{aligned}$$

where the second equality follows if we assume s *is strictly increasing. Assume also that* F *is uniform on* $[0,1]$ *(so* $f=1$ *and* $F(z)=z$*) and conjecture a linear strategy for* j*:*

$s_j(x_j) = a_j x_j + c_j$, so $x_j = s_j^{-1}(b) = (b - c_j)/a_j$. Then the optimality condition becomes

$$
\begin{aligned}
s_i(x_i) &\in \arg\max_b \Pr\{s_j(x_j) < b\} E(x_i + x_j - b | b > s_j(x_j)) \\
&= \arg\max_b \left(s_j^{-1}(b)\right)\left(x_i - b + E\left(x_j | x_j < s_j^{-1}(b)\right)\right) \\
&= \arg\max_b \left(s_j^{-1}(b)\right)\left(x_i - b + s_j^{-1}(b)/2\right) \\
&= \arg\max_b \left(\frac{b - c_j}{a_j}\right)\left(x_i - b + \frac{b - c_j}{2a_j}\right)
\end{aligned}
$$

Taking the derivative and setting it equal to zero to find the optimal b for i given x_i yields $b = \frac{a_j x}{2a_j - 1} + \frac{a_j c_j - c_j}{2a_j - 1}$. So we obtain a linear solution. Note that if the outcome of this process weren't linear it would mean that when the opponent chooses a linear strategy the best reply need not be linear so the asumption of an equilibrium in linear strategies would be violated.

In a symmetric equilibrium $b = s(x_i) = ax_i + c = \frac{ax}{2a-1} + \frac{ac-c}{2a-1}$ which implies $a = 1$ and $c = 0$ and thus $s(x_i) = x_i$.

Contrast this with an "analogous" private values case where $v_i = x_i + E(x_j) = x_i + 1/2$. Then $s^{IPV}(x_i) = 1/2 + x_i/2$. (By analogy to what we have already solved.) So $s^{IPV}(x_i) > s^{CV}(x_i)$ (since $1/2 > x_i/2$ except for $x_i = 1$).

Note that in the pure common-values case there is no question of efficiency. (There is still the question of whether the price converges to the true value in a large economy. It does.)

A major question about auctions and other mechanisms of allocation is whether they aggregate information efficiently. That is, assume that individuals have private information about the value to the other player. Will the outcome "reveal" all that information, in the sense that the price will reflect it and the winner will be the same as if it were public. There are several positive answers to this, when there are many players. This is a more advanced literature that we will not study, but will comment upon.

5.5 Bilateral Trade or the Buyer-Seller problem

5.5.1 The double-auction mechanism

Assume a buyer and seller each with independent values, $v \sim F^v$, and costs, $c \sim F^c$, where both are distributions on some interval $[\underline{v}, \bar{v}]$ and $[\underline{c}, \bar{c}]$. If the intervals do not overlap then we don't need any information in order to determine that there should always (or never) be trade. Moreover, when there should always be trade we can find a price $p \in (\bar{c}, \underline{v})$ such that if we determine that as the price then the agents will transact and we have efficiency.

But if the intervals overlap we don't know what price to set. Consider then the following mechanism. Each has to offer a price, and if the selling price is no higher than the buying price, they transact at the average price . Otherwise they do not transact. In the efficient outcome they have to tell the truth, otherwise they will forgo beneficial opportunities of trade. But, intuitively, this is not an equilibrium because by exaggerating one gives up small probabilities of trade and gains price. (If the other tells the truth, and you say $v-\varepsilon$ instead of v then you gain $\varepsilon/2$ in all cases of trade, and lose trade if the other person has cost $c\in[v,v+\varepsilon]$, but in these cases your gain was in the order of ε, so your loss is of order ε^2 and gain is of order ε.)

Exercise 80 *Show that the strategies are weakly monotonic (either increasing or decreasing).*

We will assume the distribution is uniform on $[0,1]$. Then the optimal prices satisfy

$$\begin{aligned}
s_b(v) &\in \arg\max_p \Pr\{s_s(c)\le p\} E\left(v-\frac{p+s_s(c)}{2}\middle| s_s(c)\le p\right)\\
s_s(c) &\in \arg\max_p \Pr\{s_b(v)\ge p\} E\left(\frac{p+s_b(v)}{2}-c\middle| s_b(v)\ge p\right)
\end{aligned}$$

There are many equilibria of this game.

Exercise 81 *Find an equilibrium where there is never trade. Show that there is an equilibrium in which there is trade if and only if $v\ge\bar{p}$ and $c\le\bar{p}$, for any $\bar{p}\in[0,1]$. Do these equilibria involve weakly dominated strategies?*

We will find an equilibrium in linear strategies: $s_b(v)=a_bv+d_b$, $s_s(c)=a_sc+d_s$. (Note that this includes the case of telling the truth, even though we already "know" that is not an equilibrium strategy.) Consider the buyer bidding p and the seller playing $s_s(c)=a_sc+d_s$.

$$\begin{aligned}
&\Pr\{s_s(c)\le p\} E\left(v-\frac{p+s_s(c)}{2}\middle| s_s(c)\le p\right)\\
=\ &\Pr\{a_sc+d_s\le p\} E\left(v-\frac{p+a_sc+d_s}{2}\middle| a_sc+d_s\le p\right)\\
=\ &\Pr\left\{c\le\frac{p-d_s}{a_s}\right\}\left(v-\frac{p}{2}-\frac{1}{2}E\left(a_sc+d_s\mid a_sc+d_s\le p\right)\right)\\
=\ &\frac{p-d_s}{a_s}\left(v-\frac{p}{2}-\frac{p+d_s}{4}\right)
\end{aligned}$$

So (taking the derivative, setting it equal to zero, and solving): $s_b(v)=2v/3+d_s/3\Rightarrow a_b=2/3$ and $d_s/3=d_b$

Similarly,

$$\begin{aligned}
& \Pr\{s_b(v) \geq p\} E\left(\frac{p + s_b(v)}{2} - c \middle| s_b(v) \geq p\right) \\
= \; & \Pr\{a_b v + d_b \geq p\} E\left(\frac{p + a_b v + d_b}{2} - c \middle| a_b v + d_b \geq p\right) \\
= \; & \left(1 - \frac{p - d_b}{a_b}\right)\left(-c + \frac{p}{2} + \frac{1}{2}(p + a_b + d_b)\right)
\end{aligned}$$

So (taking derivatives, setting equal to zero and solving): $s_s(c) = 2c/3 + (a_b + d_b)/3 \Rightarrow a_s = 2/3$ and $\left(\frac{2}{3} + \frac{d_s}{3}\right)/3 = d_s$, so $d_s = 1/4$ and hence $d_b = 1/12$.

So the equilibrium is: $s_b(v) = 2v/3 + 1/12$ and $s_s(c) = 2c/3 + 1/4$.

In this equilibrium there is trade iff $2v/3 + 1/12 > 2c/3 + 1/4 \iff v > c + 1/4$. Trade fails to occur even when it is beneficial (e.g. if $b + 1/4 > v > c$).

Myerson and Satterthwaite proved that *no bargaining or trading mechanism of any sort obtains efficiency* in the example above. It turns out that the double auction we solved for is the best one can get. To be precise about what it means that this is the best , recall for each value of (c, v) we can compute the optimal surplus (gain) from trade. It is zero if $c > v$ and it is $v - c$ if $v > c$. So the maximal possible surplus is $\int_v \left(\int_{c<v} (v - c) f_c \mathrm{d}c\right) f_v \mathrm{d}v$. For any mechanism we can calculate the actual surplus. The mechanism determines a probability of trade for any pair of types (c, v); denote this probability by $q(c, v)$. Then when those are the actual types they have expected surplus $q(c, v) \times (v - c)$. So the overall expected surplus of a mechanism with probability of trade $q(c, v)$ for each pair (c, v) is $\int_v \left(\int_{c<v} q(c, v)(v - c) f_c \mathrm{d}c\right) f_v \mathrm{d}v$. The double auction maximizes this value, but we see that it is strictly less than the maximal possible surplus that could be obtained if agents did not have private information about their costs and values.

Recall the Coase theorem: by assigning property rights get Pareto efficiency *if bargaining is efficient*. Here we see that in a "natural" mechanism bargaining is not efficient, which shows that the Coase theorem might be irrelevant in many environments.

5.5.2 On the existence of an ex post efficient budget balanced individually rational mechanism for trade

We will show in a discrete example a version of the Myerson-Satterthwaite result. In contrast to the overlapping-intervals case studied above, we will see that there may be inefficiency, but it need not occur – it depends on the parameters.

We will not be specific about the nature of the inefficiency.

Some mechanism design problems – as we saw in studying auctions – admit the existence of ex-post efficient, ex-post budget balanced, and interim individually rational mechanisms, but others – such as the bilateral trade or buyer–seller problem we are studying here – do not. We now clarify these terms.

Ex post efficiency means that given the realized values of everyone the outcome is efficient. Given quasi-linear preferences that we have assumed this means that the allocation of the goods maximizes the sum of utilities.

Ex post budget balance means that while there are transfers of money among players (e.g., between buyers and sellers) there is no money transferred outside the set of players or received from outside the set of players.

Interim Individually rational means that the players can choose whether to participate in the mechanism, and they choose this based on knowing their type.

The most famous result that establishes the *impossibility* of ex-post efficient, budget balanced, and individually rational mechanisms is due to Myerson and Satterthwaite (1983). Here, we consider a simpler 2×2 version of their model that is due to Matsuo (1989).

So consider the following mechanism design problem. There is a buyer and a seller. The buyer is interested in buying an object which is owned by the seller. The buyer's value (type) for the object is either v_1 or v_2. The seller's reservation value (cost, type) is either c_1 or c_2. Suppose that each profile of types is equally likely (i.e., the buyer's and seller's types are independent, and both the buyer and the seller are equally likely to be of either type). Suppose that the buyer's and seller's types are "symmetric" in the following sense:

$$\overset{c_1}{\cdot} \longleftarrow A \longrightarrow \overset{v_1}{\cdot} \longleftarrow -D- \longrightarrow \overset{c_2}{\cdot} \longleftarrow A \longrightarrow \overset{v_2}{\cdot}$$

A mechanism for this bargaining environment is any game form that specifies a message set for each player and a mapping from message profiles to outcomes, where an outcome is (i) the probability with which the buyer obtains the object, as a function of the messages sent by the buyer and seller and (ii) the price the buyers pays as a function of the messages sent by the buyer and seller.

The revelation principle implies that if we are interested in studying the range of equilibrium outcomes, then no loss of generality is entailed by restricting attention to truth-telling equilibria in direct revelation mechanisms.

A direct revelation mechanism is composed of two functions: $t\,(c, v)$, which described the expected payment from the buyer to the seller when their types are given by (c, v), and $q\,(c, v)$, which described the probability with which the buyer obtains the object when types are given by (c, v). A direct revelation mechanism is thus characterized by the following eight-tuple $\langle q_1, q_2, q_3, q_4, t_1, t_2, t_3, t_4 \rangle$:

$$\begin{array}{c|c|c|} q(c,v) & c_1 & c_2 \\ \hline v_1 & q_1 & q_2 \\ \hline v_2 & q_3 & q_4 \\ \hline \end{array} \qquad \begin{array}{c|c|c|} t(c,v) & c_1 & c_2 \\ \hline v_1 & t_1 & t_2 \\ \hline v_2 & t_3 & t_4 \\ \hline \end{array}$$

Let

$$\begin{aligned} u(v', v) &\equiv E_c [vq(c, v') - t(c, v')] \\ h(c', c) &\equiv E_v [t(c', v) - cq(c', v)] \end{aligned}$$

Definition. A direct revelation mechanism $\langle t, q \rangle$ is incentive compatible and individually rational if and only if

$$U(v) \equiv u(v,v) \geq u(v', v) \qquad \forall v, v' \in \{v_1, v_2\} \tag{5}$$

$$H(c) \equiv h(c,c) \geq h(c', c) \forall c, c' \in \{c_1, c_2\} \tag{6}$$

$$U(v) \geq 0 \qquad \forall v \in \{v_1, v_2\} \tag{7}$$

$$H(c) \geq 0 \qquad \forall c \in \{c_1, c_2\} \tag{8}$$

Remark 17 Alternatively, instead of focusing our attention on incentive compatible direct revelation mechanisms we could consider the probability of trade and expected payment in a Nash equilibrium and denote those by $q(c,v)$ and $t(c,v)$. In this case, IC and IR would follow from the fact that what we consider is a truth-telling equilibrium in the game where the players can choose whether to play (and if not they get 0) and if they play they send a message specifying their type. The fact that these two approaches are identical illustrates the revelation principle in this context.

Definition 20 *A direct revelation mechanism $\langle t, q \rangle$ is ex-post efficient if $q(c,v) = 1$ whenever $v > c$, or in matrix form:*

$$\begin{array}{c|c|c|} q(c,v) & c_1 & c_2 \\ \hline v_1 & 1 & 0 \\ \hline v_2 & 1 & 1 \\ \hline \end{array}$$

Proof. There exists an ex-post efficient incentive compatible and individually rational direct revelation mechanism $\langle t, q \rangle$ if and only if

$$c_2 - v_1 \leq (v_2 - c_2) + (v_1 - c_1) \qquad \text{(i.e., iff } D \leq 2A\text{)}$$

We first show that if $D > 2A$, then no ex-post efficient mechanism exists. Suppose to the contrary that such a mechanism exists. The IC constraint for v_2 implies

$$U(v_2) \geq v_2 E_c[q(c, v_1)] - E_c[t(c, v_1)] + v_1 E_c[q(c, v_1)] - v_1 E_c[q(c, v_1)]$$

or

$$U(v_2) \geq U(v_1) + (v_2 - v_1) E_c[q(c, v_1)]. \tag{1}$$

Similarly, the IC constraint for c_1 implies

$$H(c_1) \geq H(c_2) + (c_2 - c_1) E_v[q(c_2, v)]. \tag{2}$$

Ex-post efficiency implies that $E_c[q(c, v_1)] = E_v[q(c_2, v)] = \frac{1}{2}$. Plug this into (**??**) and use IR – B to get

$$U(v_2) \geq \frac{v_2 - v_1}{2}.$$

Similarly, (**??**) and IR – S imply

$$H(c_1) \geq \frac{c_2 - c_1}{2}.$$

Since v_2 and c_1 each occurs with probability $\frac{1}{2}$, the sum of the ex-ante expected payoffs is at least

$$\frac{1}{2} \cdot (U(v_2) + H(c_1)) \geq \frac{v_2 - v_1}{4} + \frac{c_2 - c_1}{4} = \frac{A + D}{2}.$$

But the maximum ex-ante surplus is

$$\frac{1}{4} \cdot ((v_2 - c_2) + (v_2 - c_1) + (v_1 - c_1) + 0) = \frac{4A + D}{4},$$

which is smaller than $\frac{A+D}{2}$ if $D > 2A$. A contradiction.

Next, we show that if $D \leq 2A$ then the following direct revelation mechanism is incentive compatible, individually rational, and ex-post efficient: $q(c, v) = 1$ if and only if $c < v$, and $t(c, v)$ is given by the following matrix:

$t(c,v)$	c_1	c_2
v_1	v_1	0
v_2	$\frac{v_2+c_1}{2}$	c_2

Intuition: The problem is to get the "high value" types v_2 and c_1 to reveal their types. To accomplish this goal, these types are given the most favorable prices possible if they reveal their identity.

Ex-post efficiency, IR, and IC for v_1 and for c_2 are immediate. IC for v_2 follows from

$$\frac{1}{2}\cdot\underbrace{\left(v_2-\frac{v_2+c_1}{2}\right)}_{A+\frac{D}{2}}+\frac{1}{2}\cdot\underbrace{(v_2-c_2)}_{A}\geq\frac{1}{2}\cdot\underbrace{(v_2-v_1)}_{A+D}.$$

$$\begin{aligned} 2A+\frac{D}{2} &\geq A+D \\ &\iff 2A\geq D. \end{aligned}$$

IC for c_1 follows similarly. ■

Remark 18 The ex-post efficient mechanism for the case where $D \leq 2A$ is not incentive compatible when $D > 2A$ because in this case v_2 and c_1 would rather report v_1 and c_2, respectively. If $D \leq 2A$, the benefit that v_2 and c_1 would obtain from reporting v_1 and c_2, respectively, is too small relative to the lower probability they would get to trade, and so ex-post efficiency is possible.

5.6 Voting: the swing voters' curse, sequential voting, and unanimity in juries

Consider the following example (taken from Feddersen and Pesendorfer, AER, 1996). "There are two candidates, the status quo (candidate 0) and the alternative (candidate 1). Voters are uncertain about the cost of implementing the alternative. This cost is either high (state 0) or low (state 1). All voters prefer the status quo if the cost is high and the alternative if the cost is low. At least one of the voters is [perfectly] informed and knows the costs with certainty. However, voters do not know the exact number of informed voters in the electorate. All of the uninformed voters share a common knowledge prior that with .9 probability the cost is high and the status quo is the best candidate. Suppose that all voters (informed and uninformed alike) vote only on the basis of their updated prior. All of the informed voters vote for the status quo if the cost is high and the alternative if the cost is low while all of the uninformed voters vote for the status quo in both states. The informed voters are behaving rationally while the uninformed are not. An uninformed voter is only pivotal if some voters have voted for the alternative. But this can only occur if the cost is low and the informed voters vote for the alternative. Therefore, an uninformed voter can affect the election outcome only if the cost is low. Consequently, an uninformed voter should vote for the alternative. On the other hand, it cannot be rational for all uninformed voters to vote for the alternative. In this case each uninformed voter would prefer to vote for the status quo. Thus it is not optimal for uninformed voters to vote only on the basis of their prior information. In this example there is an easy solution for the uninformed voters: abstention. ...Given the behavior of the other voters uninformed voters suffer the swing voter's curse: they are strictly better off abstaining than by voting for either candidate. This is true even though uninformed voters believe that the status quo is almost certainly the best candidate."

Exercise 82 *There are 2 states, 0 and 1 and two proposals, the status quo "zero" and the alternative, "one". Players have utility 1.1 if the status quo is chosen in state 0 and utility 1 if the alternative is chosen in state 1, and zero otherwise. (They enjoy matching the state, but it's slightly more important to keep the status quo in state 0 than to switch in state 1.) Each state is a priori equally likely. In each state each player observe a signal that is correct with probability 0.6. The alternative is selected if at least 2/3 vote for it.*

Assume there are 2 voters. When is voter 1 pivotal? Assuming all but voter 1 vote according to their signal, then conditional on voter 1 being pivotal what is the probability that the alternative is the right choice? What is 1's optimal vote if she observes the signal that the state is 0? What if the signal is that the state is 1? Is it an equilibrium to vote informatively?

Assume there are 100 voters. When is voter 1 pivotal? Assuming all but voter 1 vote according to their signal, then conditional on voter 1 being pivotal what is approximately the probability that the alternative is the right choice? What is 1's optimal vote if she observes the signal that the state is 0? What if the signal is that the state is 1? Is it an equilibrium to vote informatively?

The preceding exercise shows that it cannot be that everyone is voting informatively in a large electorate. One can expand the example to allow for different preferences plus a common component as above. (That is, there is a private value component to the election (e.g. tastes) and a common value component (e.g., competence).) Then we will get "monotonic" equilibria where a segment who most prefers the status quo votes for the status quo, a segment who most prefers the alternative votes for the alternatives, and a middle segment votes informatively. The percent that votes informatively converges to zero as the population grows, but is still large enough for all the private information about the common value "aggregates": the outcome is the same as if the private information were public.

Exercise 83 ***Voting continued: sequential (symmetric binary) voting*** *Assume you are given a symmetric equilibrium in the simultaneous voting game with* 21 *voters who have preferences and information as in the preceding exercise and using the same 2/3 rule. In the* *given equilibrium each voter votes* 1 *with probability* p *if she observes a signal that the state is 1 and she votes 0 with probability* q *if she observes the state is 0. Answer the following without calculations.*

Can it be that both p *and* q *are strictly between zero and one?*

If p *is strictly between zero and one, what is* q*?*

If q *is strictly between zero and one, what is* p*?*

Which of these, if either, is inconsistent with equilibrium?

Now imagine they move sequentially one after the other, and observe the votes. Assume that all voters except for the 14th to move play the symmetric equilibrium of the simultaneous move game—that is they ignore the votes of everyone else. Now it is the turn of the 14th, and she has seen 13 votes all for i *preceding her move, and observes the signal that the state is* $j \neq i$*. Is it a best reply to ignore the preceding votes and vote according to the equilibrium?*

Voting continued: Juries

We may discuss this in class; for more detail see the paper by Feddersen and Pesendorfer in the American Political Science Review.

5.7 No trade

Assume two individuals with private information considering a bet. That is, there is a state of Nature, $\omega \in \Omega$, information partitions P_i, for $i = 1, 2$, and a possible "bet", which is transfer from 1 to 2 conditional on the state of nature: $\tau(\omega)$ is the payment of 1 to 2 in state ω. Thus $\tau(\omega)$ is the amount that 1 loses and 2 gains from the bet. The two players must agree for the bet to take place. (We will focus on bets with expected ex ante return zero and no price for the bet, but as we'll see the argument applies to having positive or negative transfer prices. Moreover, this is equivalent to a bet with a positive ex ante expected value as well. For example, purchasing a stock is like such a bet.) Assume an arbitrarily small transaction cost, denoted by ε, to agreeing to a bet. Then no one ever agrees to the bet. The proof is immediate. A type t_i of player i will agree to a bet only if the expected return to i conditional on the type t_i is strictly positive, since it has to cover the transaction cost. Then the ex ante payoff is strictly positive for both players. But the ex ante payoff to one is the negative of the ex ante payoff to the other.

Assume (wlog) that for each pair of types there is a unique state. Then formally we have types $t_i \in T_i$ and actions Y, N and payoffs in state (t_1, t_2) if both say yes in that state are equal to $-\tau(t_1, t_2) - \varepsilon, \tau(t_1, t_2) - \varepsilon$. Therefore if t_2 says Y then $\sum_{t_1} \tau(t_1, t_2)\, p(t_1|t_2) > \varepsilon$. So, adding up for all t_2 we have $\varepsilon < \sum_{t_2} p(t_2) \sum_{t_1} \tau(t_1, t_2)\, p(t_1|t_2) = \sum_{t_1,t_2} \tau(t_1, t_2)\, p(t_1, t_2)$. But the same argument for 1 shows $\sum_{t_1,t_2} -\tau(t_1, t_2)\, p(t_1, t_2) > \varepsilon$. Obviously this cannot happen so it cannot be that both say Y for some types, but if one always says N then the other only pays the transaction cost, so must say N as well.

We will also consider a numeric example in class to make this more concrete.

Note that a critical assumption in the above is the common prior.

Remark 19 *This example is closely related to Aumann's "Agreeing to disagree" paper. There he shows that the conditional probabilities that two individuals assign to an event cannot differ if the value is common knowledge. Formalizing this requires defining common knowledge which we will not do at this stage. But we will discuss it n class via an example like the one we'll present on the above.*

Exercise 84 *Assume that individuals are risk averse—can they ever agree to bet (i.e., both say yes in some state of the world) in equilibrium, and if so what is their expected gain in the state that they agree? Assume they are risk neutral: can they agree to bet and if so what is their expected gain?*

Exercise 85 *Consider the following information structure. Player 1 information partition is* $\{\{1\},\{2,3\},\{4\}\}$*, 2's is* $\{\{1,2\},\{3,4\}\}$*. The bet has the following transfers from 1 to 2* $\tau(1)=-1, \tau(2)=3, \tau(3)=-7, \tau(4)=15$*. Assume that each player says yes by mistake with probability* μ*. That is, if they try to say no, then the actually say yes with probability* μ *and no with probability* $1-\mu$*. Thus their actions are the choice of attempting to say yes (which succeeds) and attempting to say no (which succeeds with probability* μ*). Find the set of equilibria when* $\mu=1/2$*. Find the set of equilibria when* $\mu=1/10$*.*

5.8 Correlated equilibrium

Consider a game with asymmetric information where *both players payoffs do not depend on either player's type.* That is, $u(a_i,a_j,t_i,t_j)=u(a_i,a_j)$. This is equivalent to playing a given game after observing some correlating device whose outcome is payoff irrelevant. This can nonetheless impact payoffs. Consider the following game.

	L	R
U	$7,7$	$2,9$
D	$9,2$	$0,0$

This game has 3 Nash equilibria, with payoffs $(2,9),(9,2)$ and $(4.5,4.5)$. A public correlating device can obtain convex combinations of these. Consider now the type space with common prior in the table below obtained for example by someone choosing randomly among $\{1,2,3\}$ and informing the column player whether or not the outcome is 3 and the row player whether or not it is 1. If players choose according to the label corresponding to their type then we have an equilibrium and the ex ante payoff is 6.

	t_2^L	t_2^R
t_1^U	1/3	1/3
t_1^D	1/3	0

It should be intuitive that no matter how complicated a payoff irrelevant type space is, what matters is only the actions the types play. So any two types that play the same action in equilibrium can be combined into one joined type that cannot tell the original two apart. By the sure-thing principle if both uncombined types found it optimal to play some action a then the joined type will as well. So in studying the effect of adding payoff irrelevant information we can wlog restrict attention to type spaces with the number of types of a player equal to the number of actions of that player. So in the above example, the set of equilibrium distributions we can find is determined as follows.

Consider the following type space (where of course all probabilities non negative and sum to 1).

	t_2^L	t_2^R
t_1^U	p	q
t_1^D	r	s

How does t_1^U compare U to D? U gives expected payoff of $\frac{p}{p+q}7 + \frac{q}{p+q}2$ while D gives $\frac{p}{p+q}9$. There are three other similar equations. If all are satisfied we have an equilibrium where each type plays its designated strategy. Moreover, as argued before, any equilibrium of a the game preceded by any type space is equivalent to some such equilibrium. So we can characterize the set of distributions over outcomes that can occur in this way—four linear equations determine the set of equilibrium distributions over outcomes.

5.9 Bad Reputation

In a game with incomplete information, "reputation" is defined as the updated probability with which a player is believed to be of a certain type, usually the "good"/"strong"/"aggressive" type. This belief affects play and so players try to manipulate it to their advantage.

"Good reputation" vs. "bad reputation" models hinge on whether the activities taken to manipulate beliefs enhance or hurt welfare. For example, production of high quality products in order to support the belief that one is a "good" producer enhances welfare. Indeed, people are used to the idea of working hard in order to build a good reputation. The next example illustrates that sometimes there is no choice but to develop a bad reputation.

The example is due to Ely and Valimaki (QJE, 2003).

- Suppose that when a car has a problem it either needs a tune up or engine replacement, each with probability $\frac{1}{2}$.
- Suppose that the payoff to the motorist if she gets what she need is 10. The payoff if she gets the wrong treatment is -20. The payoff if she doesn't get any treatment is 0. (All payoffs include the relevant price which is fixed through the analysis.)
- The mechanic is either "honest" or not (this is the mechanic's type). The payoff of an honest mechanic is the same as that of the motorist. So if the game is that the mechanic is honest and everyone knows this, he does the right thing always and everyone takes their cars to him. Everyone's payoff is 10.
- Suppose however that there's a small probability that the mechanic is dishonest and replaces every engine, regardless of whether it is needed or not.

Claim. The uncertainty about the car mechanic's type can cause the car repair market to completely break down.

Reason: Suppose that people continue to go to the car mechanic and that the first M people who go all get their engines replaced. If M is really big, people are going to figure this must be the dishonest mechanic and they'll all stop going to him. Of course, there is a chance that the mechanic is honest and it just so happened that this is what the first M people really needed. Still, if M is very large, this seems very unlikely relative to the hypothesis that the car mechanic is dishonest.

So suppose that people will no longer go to the car mechanic if he replaces the first M engines. Suppose that the car mechanic has replaced the first $M-1$ engines and one more motorist shows up.

If the car mechanic is dishonest, then he'll replace the engine whether the car needs it or not.

If the car mechanic is honest, then if the car only needs a tune up, he'll happily do that. But what if the car really needs an engine replacement? An honest car mechanic knows that if he does an engine replacement, no one will ever come to his shop again. But if he does a tune up, everyone will realize that he's not dishonest and motorists will come to him all the time. So the payoff of an honest car mechanic from doing a tune up is -20 today plus 10 every period from tomorrow onward. The payoff from doing an engine replacement is 10 today but 0 forever after. If the car mechanic is patient enough, the former will be better. Let's suppose that the car mechanic is indeed patient enough for this to be true.

So we have seen that an honest car mechanic will do a tune up whether it's needed or not, and a dishonest car mechanic will do an engine replacement whether it's needed or not.

Consider now the payoff to the motorist. If she is facing an honest mechanic, then her expected payoff is

$$\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot (-20) = -5.$$

If she is facing a dishonest car mechanic, she gets the same expected payoff. This implies that the motorist is better off not going to the car mechanic.

Hence, we have shown that if the car mechanic does engine replacements on the first $M - 1$ cars, no one will show up again.

But now the argument works backward. The conclusion is that no one will ever go to the car mechanic at all.