

Lecture Notes for 1st Year Ph.D. Game Theory*

Navin Kartik[†]

Contents

1	Introduction	3
1.1	A (Very!) Brief History	3
1.2	Non-cooperative Game Theory	3
2	Strategic Settings	4
2.1	Extensive Form Representation	4
2.2	Strategies and Strategic Form of a Game	6
2.2.1	Strategies	6
2.2.2	Strategic (Normal) Form	7
2.2.3	Randomized Choices	8
2.3	An Economic Example	10
3	Simultaneous-Move Games	11
3.1	Dominance	11
3.1.1	Strictly Dominant Strategies	11
3.1.2	Strictly Dominated Strategies	12
3.1.3	Iterated Deletion of Strictly Dominated Strategies	13
3.1.4	Weakly Dominated Strategies	15
3.2	Rationalizability	17
3.3	Nash Equilibrium	19
3.3.1	Pure Strategy Nash Equilibrium	19
3.3.2	Mixed Strategy Nash Equilibrium	21
3.3.3	Finding Mixed Strategy Equilibria	24
3.3.4	Interpreting Mixed Strategy Equilibria	26
3.4	Normal Form Refinements of Nash Equilibrium	26
3.5	Correlated Equilibrium	30
3.6	Bayesian Nash Equilibrium and Incomplete Information	32
3.6.1	Examples	35
3.6.2	Purification Theorem	38
3.6.3	The Importance of Higher-Order Beliefs	39

*Last updated: March 11, 2020. These notes draw upon various published and unpublished sources, including notes by Doug Bernheim, Vince Crawford, George Mailath, David Miller, and Bill Sandholm. Thanks to David Miller and former and current students for comments.

[†]nkartik@gmail.com. Please feel free to send me suggestions or comments. I would particularly appreciate alerts to typos or other errors.

4	Dynamic Games and Extensive Form Refinements of Nash Equilibrium	41
4.1	The Problem of Credibility	41
4.2	Backward Induction and Subgame Perfection	42
4.2.1	Backward Induction	42
4.2.2	Subgame Perfect Nash Equilibrium	43
4.3	Systems of Beliefs and Sequential Rationality	46
4.3.1	Weak Perfect Bayesian Equilibrium	47
4.3.2	Sequential Equilibrium	51
5	Market Power	54
5.1	Monopoly	55
5.2	Basic Oligopoly Models	56
5.2.1	Bertrand oligopoly	56
5.2.2	Cournot oligopoly	57
5.3	Stackelberg Duopoly	58
5.4	Price Competition with Endogenous Capacity Constraints	59
6	Repeated Games	60
6.1	Description of a Repeated Game	61
6.2	The One-Shot Deviation Principle	63
6.3	A Basic Result	64
6.4	Finitely Repeated Games	65
6.5	Infinitely Repeated Games	67
6.5.1	Folk Theorems	70
7	Signaling and Cheap Talk	75
7.1	Costly Signaling	75
7.1.1	The Setting	75
7.1.2	Basic Properties	76
7.1.3	Separating Equilibria	77
7.1.4	Pooling Equilibria	80
7.1.5	Equilibrium Refinement	81
7.2	Cheap Talk	82
8	Mechanism Design	83
8.1	Revelation Principle	83
8.2	Vickrey-Clarke-Groves Mechanisms	84

1. Introduction

Game theory is a formal methodology and a set of techniques to study the interaction of *rational* agents in *strategic* settings.¹ ‘Rational’ here means the standard thing in economics: maximizing over well-defined objectives; ‘strategic’ means that what decisions one agent wants to take depends on what she thinks other agents will be doing, which in turn may depend on what those other agents think this agent will be doing. Note that *decision theory*—which you should have elements of earlier in the Micro core—is the study of how an individual makes decisions in non-strategic settings; hence game theory is sometimes also referred to as *multi-person* or *interactive decision theory*. The common terminology for the field comes from its putative applications to games such as poker, chess, etc.² However, the applications we are usually interested in have little directly to do with such games. In particular, these are what we call “zero-sum” games in the sense that one player’s loss is another player’s gain; they are games of pure conflict. In most interesting economic applications, there is typically a mixture of conflict and cooperation motives.

1.1. A (Very!) Brief History

Modern game theory as a field owes much to the work of John von Neumann. In 1928, he wrote an important paper on two-person zero-sum games that contained the famous Minimax Theorem, which we’ll see later on. In 1944, von Neumann and Oscar Morgenstern published their classic book, *Theory of Games and Strategic Behavior*, that extended the work on zero-sum games and also started cooperative game theory. In the early 1950s, John Nash made his seminal contributions to non-zero-sum games and started bargaining theory. In 1957, Robert Luce and Howard Raiffa published their book, *Games and Decisions: Introduction and Critical Survey*, popularizing game theory. In 1967–1968, John Harsanyi formalized methods to study games of incomplete information, which was crucial for widening the scope of applications. In the 1970s, there was an explosion of theoretical and applied work in game theory, and the methodology was well along its way to its current status as a preeminent tool in not only economics but other social sciences too.

1.2. Non-cooperative Game Theory

Throughout this course, we will focus on *noncooperative* game theory, as opposed to *cooperative* game theory. All of game theory describes strategic settings by starting with the set of *players*, i.e. the decision-makers. The difference between noncooperative and cooperative game theory is that the former takes each player’s individual actions as primitives, whereas the latter takes joint actions as primitives. That is, cooperative game theory assumes that binding agreements can be made by players within various groups and players can communicate freely in order to do so. We will take the noncooperative viewpoint that each player acts as an individual, and the possibilities for agreements and communication must be explicitly modeled. Except for brief discussions in Appendix A of Chapter 18 and parts of Chapter 22, [Mas-Colell, Whinston, and Green](#)

¹ Naturally, any attempt like this to provide a succinct definition of a rich subject is bound to be incomplete; nevertheless, this working definition is suitable for the focus in this course.

² Ironically, game theory actually has limited prescriptive advice to offer on how to play either chess or poker. For example, we know that chess is “solvable” in a sense to be made precise later, but nobody actually knows what the solution is! This stems from the fact that chess is simply too complicated to “solve” (at present); this is of course why the best players are said to rely significantly on their intuition or feel in addition to logical computation.

(1995)—hereafter, MWG—does not deal with cooperative game theory either. For an excellent introduction, see Chapters 13–15 in [Osborne and Rubinstein \(1994\)](#).

2. Strategic Settings

A game is a description of a strategic environment. Informally, the description must specify who is playing, what the rules are, what the outcomes are depending on any set of actions, and how players value the various outcomes.

Example 1. (Matching Pennies version A) Two players, *Anne* and *Bob*. Simultaneously, each picks *Heads* or *Tails*. If they pick the same, Bob pays Anne \$2; if they pick different, Anne pays Bob \$2.

Example 2. (Matching Pennies version B) Two players, *Anne* and *Bob*. First, Anne picks either *Heads* or *Tails*. Upon observing her choice, Bob then picks either *Heads* or *Tails*. If they pick the same, Bob pays Anne \$2; if they pick different, Anne pays Bob \$2.

Example 3. (Matching Pennies version N) Two players, *Anne* and *Bob*. Simultaneously, they pick either *Heads* or *Tails*. If they pick different, then they each receive \$0. If they pick the same, they wait for 15 minutes to see if it rains outside in that time. If it does, they each receive \$2 (from God); if it does not rain, they each receive \$0. Assume it rains with 50% chance.

In all the above examples, implicitly, players value money in the canonical way and are risk-neutral. Notice that [Example 1](#) and [Example 2](#) are zero-sum games — whatever Anne wins, Bob loses, and vice-versa. [Example 3](#) is not zero-sum, since they could both win \$2. In fact, it is a particular kind of coordination game. Moreover, it is also a game that involves an action taken by “nature” (who decides whether it rains or not).

2.1. Extensive Form Representation

- Work through extensive form representation of the examples first.

Let us now be more precise about the description of a game.

Definition 1. An extensive form game is defined by a tuple $\Gamma_E = \{\mathcal{X}, \mathcal{A}, I, p, \alpha, \mathcal{H}, H, \iota, \rho, u\}$ as follows:

1. A finite set of I players. Denote the set of players as $I = \{0, 1, \dots, I\}$. Players $1, \dots, I$ are the “real” players; player 0 is used as an “auxiliary” player, nature.
2. A set of nodes, \mathcal{X} .³
3. A function $p : \mathcal{X} \rightarrow \mathcal{X} \cup \{\emptyset\}$ specifying a unique immediate predecessor of each node x such that $p(x)$ is the empty-set for exactly one node, called the *root node*, x_0 .⁴

(a) The immediate successors of node x are defined as $s(x) = \{y \in \mathcal{X} : p(y) = x\}$.

³ Nodes are typically drawn as small solid circles, but note [fn. 4](#).

⁴ The root node is typically drawn as a small hollow circle.

- (b) By iterating the functions p and s , we can find all predecessors and successors of any node, x , which we denote $P(x)$ and $S(x)$ respectively. We require that that $P(x) \cap S(x) = \emptyset$, i.e. no node is both a predecessor and a successor to any other node.
 - (c) The set of *terminal nodes* is $T = \{x \in \mathcal{X} : s(x) = \emptyset\}$. Any non-terminal node $x \in \mathcal{X} \setminus T$ is a *decision node*.
4. A set of actions, \mathcal{A} , and a function $\alpha : \mathcal{X} \setminus \{x_0\} \rightarrow \mathcal{A}$ that specifies for each node $x \neq x_0$, the action which leads to x from $p(x)$. We require that α be such that if distinct $x', x'' \in s(x)$, then $\alpha(x') \neq \alpha(x'')$. That is, from any node, each action leads to a unique successor. The set of available actions at any node, x , is denoted $c(x) = \{\alpha(x')\}_{x' \in s(x)}$.
 5. A collection of information sets, \mathcal{H} , that forms a partition of $\mathcal{X} \setminus T$,⁵ and a function $H : \mathcal{X} \setminus T \rightarrow \mathcal{H}$ that assigns each decision node into an information set. We require that $c(x) = c(x')$ if $H(x) = H(x')$; that is, two nodes in the same information set have the same set of available actions. It is therefore meaningful to write $C(H) = \{a \in \mathcal{A} : a \in c(x) \ \forall x \in H\}$ for any information set $H \in \mathcal{H}$ as the set of choices available at H .
 6. A function $\iota : \mathcal{H} \rightarrow I$ assigning the player (possibly nature) to move at all the decision nodes in any information set. This defines a collection of information sets that any player i moves at, $\mathcal{H}_i \equiv \{H \in \mathcal{H} : i = \iota(H)\}$.
 7. For each $H \in \mathcal{H}_0$, a probability distribution $\rho(H)$ on the set $C(H)$.⁶ This dictates nature's moves at each of its information sets.
 8. $u = (u_1, \dots, u_I)$ is a vector of utility functions such that for each $i = 1, \dots, I$, $u_i : T \rightarrow \mathbb{R}$ is a von-Neumann Morgenstern payoff function that represents (expected utility) preferences for i over terminal nodes.

Keep in mind that when drawing game trees, we use dotted lines between nodes (Kreps) or ellipses around nodes (MWG) to indicate nodes that fall into the same information set.

- Work through examples of what the definition of an extensive form game rules out.

To avoid technical complications, we restrict attention in the formal definition above to *finite* games:

Assumption 1. *The set of nodes, \mathcal{X} , is finite.*

Remark 1. If \mathcal{X} is finite, then even if the set of actions, \mathcal{A} , is infinite, there are only a finite number of *relevant* actions; hence without loss of generality, we can take \mathcal{A} as finite if \mathcal{X} is finite.

At various points, we will study *infinite* games (where the number of nodes is infinite); the extension of the formal concept of a game to such cases will be intuitive and covered as needed.

It will often be convenient to talk about games without specifying payoffs for the players. Strictly speaking, this is called a *game form* rather than a game.

⁵ Recall that a partition is a set of mutually exclusive and exhaustive subsets.

⁶ One has to be a little careful in the definition if $C(H)$ is a continuum, which MWG ignore, and I will be casual about; cf. [Assumption 1](#) below.

Definition 2. A game form is an otherwise complete description of a game, only lacking payoff specification.

A property Y is said to be *mutual knowledge* if all players know Y (but don't necessarily know that others know it). A property Y is *common knowledge* if everyone knows Y , everyone knows that everyone knows Y , everyone knows that everyone knows that everyone knows Y , ..., ad infinitum. Clearly, common knowledge implies mutual knowledge but not vice-versa.⁷

Definition 3. A complete information game is one where all players' payoff functions (and all other aspects of the game) are common knowledge.

You might worry that restricting attention to complete information games is pretty limited: what about a version of Matching Pennies where Anne does not know whether Bob wants to match or not match? We'll see there is a beautiful trick to analyze such situations within the framework of complete information.

Remark 2. We will always assume in this course that the game form is common knowledge. So the only source of informational asymmetry across players at the outset of a game can be about payoffs. More generally, however, the term "incomplete information" can refer to any game where at the outset, one player knows something about the game that another does not.

Remark 3. We will restrict our attention to games with *perfect recall*, as is typical. Loosely, this means that a player never forgets (i) a decision she took in the past, and (ii) any information that she possessed when making a past decision. More precisely: (i) if $y \in P(x)$ then $H(x) \neq H(y)$ — no node and any predecessor can be in the same information set; and (ii) if $H(x'') = H(x')$, $x \in P(x')$, and $\iota(H(x)) = \iota(H(x'))$, then $\exists \hat{x} \in H(x)$ s.t. $\hat{x} \in P(x'')$ and $\alpha(s(x) \cap P(x')) = \alpha(s(\hat{x}) \cap P(x''))$. In words, (ii) says that if x'' and x' are in the same information set, x is a predecessor of x' , and the same player moves at both x and x' (and hence also x''), then there must be an action at x 's information set that can lead to both x' and x'' .

Another piece of terminology to be aware of, but we don't want to impose in general, is *perfect information*.

Definition 4. A game has perfect information if all information sets are singletons. Otherwise, it has imperfect information.

[Example 2](#) has perfect information, but [Example 1](#) and [Example 3](#) are of imperfect information. In terms of parlor games, chess has perfect information, whereas Mastermind has imperfect information.⁸

2.2. Strategies and Strategic Form of a Game

2.2.1. Strategies

A key concept in game theory is that of a player's *strategy*. A strategy, or a decision rule, is a *complete contingent plan* that specifies how a player will act at every information set that she is the decision-maker at, should it be reached during play of the game.

Definition 5. A [pure] strategy for player i is a function $s_i : \mathcal{H}_i \rightarrow \mathcal{A}$ such that $s_i(H) \in C(H)$ for all $H \in \mathcal{H}_i$.

⁷ Chapter 5 of [Osborne and Rubinstein \(1994\)](#) contains an accessible formal treatment of knowledge.

⁸ In case you don't know it, Mastermind is the game where player 1 chooses an ordered sequence of four colored pegs (unobservable to player 2) and player 2 has to arrive at it through a sequence of guesses.

It is very important to be clear about what a strategy is. Here is point of clarification. Consider a “game” where you are walking North on Amsterdam Ave and trying to get to the department’s entrance on 118th St.⁹ Every cross street on Amsterdam is a decision node. The set of actions at each node is {Turn right, Continue on Amsterdam}. Consider a strategy that specifies *Continue* at all streets South of 118th, and *Turn right* at the 118th Street node. For a full specification, the strategy has to specify what to do if you get to the 119th St node, the 120th St, and so on — even though you won’t actually get there if you follow the strategy! Remember: *complete contingent plan*. Moreover, do not confuse *actions* and *strategies*. An action is just a choice at a particular decision node. A strategy is a plan of action for every decision node that a player is the actor at. It may seem a little strange to define strategies in this way: why should a player have to plan for contingencies that his own actions ensure will never arise?! It turns out that what *would* happen at such “never-reached” nodes plays a crucial role in studying dynamic games, a topic we’ll spend a lot of time on later.

The set of available strategies for a player i is denoted S_i . In finite games, this is a $|\mathcal{H}_i|$ -dimensional space, where $|\mathcal{H}_i|$ is the number of information sets at which i acts. That is, $s_i \in S_i = \times_{H \in \mathcal{H}_i} C(H)$. Let $S = \prod_{i=1, \dots, I} S_i$ be the product space of all players’ strategy spaces, and $s = (s_1, \dots, s_I) \in S$ be a *strategy profile* where s_i is the i^{th} player’s strategy. We will sometimes write s_{-i} to refer to the $(I - 1)$ vector of strategies of all players excluding i , and therefore $s = (s_i, s_{-i})$.

In [Example 1](#), if we let Anne be player 1 and Bob be player 2, we can write $S_1 = S_2 = \{H, T\}$. Here, both players have 2 actions and also 2 strategies. In [Example 2](#), we have $S_1 = \{H, T\}$ whereas $S_2 = \{(H, H), (H, T), (T, H), (T, T)\}$ where any $s_2 = (x, y)$ means that player 2 plays x if player 1 plays H and y if player 1 plays T . Thus, even though both players continue to have 2 actions each, observe that player 1 has 2 strategies (as before), but now player 2 has 4 strategies. (Question: how many strategies does each player have in [Example 3](#)?)

2.2.2. Strategic (Normal) Form

Every [pure] strategy profile induces a sequence of moves that are actually played, and a probability distribution over terminal nodes. (Probability distribution because nature may be involved; if there is no randomness due to nature, then there will be a unique final node induced.) Since all a player cares about is his opponents’ actual play, we could instead just specify the game directly in terms of strategies and associated payoffs. This way of representing a game is known as the *Strategic* or *Normal* form of the game. To do this, first note that given a payoff function $u_i : T \rightarrow \mathbb{R}$, we can define an extended payoff function as the expected payoff for player i from a strategy profile s , where the expectation is taken with respect to the probability distribution induced on T by s . With some abuse of notation, I will denote this extended payoff function as $u_i : S \rightarrow \mathbb{R}$ again. Notice that the domain of u_i (S or T) makes it clear whether it is the primitive or extended payoff function we are talking about.

Definition 6. The normal form representation of a game, $\Gamma_N = \{I, S, u\}$, consists of the set of players, I , the strategy space, S , and the vector of extended payoff functions $u = (u_1, \dots, u_I)$.

Often, the set of players will be clear from the strategy space, so we won’t be explicit about the set I . For instance, the Normal Form for [Example 2](#) can be written as

⁹ This is really a decision-theory problem rather than a game, but it serves well to illustrate the point.

		Bob			
		(H, H)	(H, T)	(T, H)	(T, T)
Anne	H	2, -2	2, -2	-2, 2	-2, 2
	T	-2, 2	2, -2	-2, 2	2, -2

where we follow the convention of writing payoffs as ordered pairs (x, y) , with x being the payoff for the Row player (Anne) and y that of the Column player (Bob). This is a game where there is no role for Nature, so any strategy profile induces a unique terminal node. Consider on the other hand, [Example 3](#), where this is not the case. The Normal Form is

		Bob	
		H	T
Anne	H	1, 1	0, 0
	T	0, 0	1, 1

where the payoff of 1 if they match comes from the expected utility calculation with a 0.5 chance of rain (the expected payoff given the probability distribution induced over the terminal node).

***Normal Form Equivalence** There are various senses in which two strategic settings may be equivalent, even though they have different representations (in Normal or Extensive form). Indeed, as we already remarked, the same game can have different extensive form representations (think about MP-A and whether the game tree shows Bob moving first or Anne moving first). This is actually a deep question in Game Theory, but here is at least one simple case in which the equivalence should be obvious.

Definition 7 (Full Equivalence). Two normal form games, $\Gamma_N = \{I, S, u\}$ and $\tilde{\Gamma}_N = \{I, S, \tilde{u}\}$, are fully equivalent if for each $i = 1, \dots, I$, there exists $A_i > 0$ and B_i such that $\tilde{u}_i(s) = A_i u_i(s) + B_i$.

This definition is a consequence of the fact that the utility functions represent vNM expected-utility preferences, hence are only meaningful up to a linear transformation. For instance, this means that MP-A in [Example 1](#) is fully equivalent to another version of Matching Pennies where we just multiply players' payoffs by the constant 2. Makes sense, right?

2.2.3. Randomized Choices

Mixed Strategies Thus far, we have taken it that when a player acts at any information set, he deterministically picks an action from the set of available actions. But there is no fundamental reason why this has to be case. For instance, in MP-A, perhaps Bob wants to flip a coin and make his choice based on the outcome of the coin flip. This is a way of making a *randomized choice*. Indeed, as we'll see, allowing for randomization in choices plays a very important role in game theory.

Definition 8 (Mixed Strategy). A *mixed* strategy for player i is a function $\sigma_i : S_i \rightarrow [0, 1]$ which assigns a probability $\sigma_i(s_i) \geq 0$ to each pure strategy $s_i \in S_i$, satisfying $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

One way to think of this is that at the outset, i flips an $|S_i|$ -sided die (with the right probabilities for each side), and based on its outcome, decides which pure strategy to play. Clearly, a pure strategy is a

degenerate kind of mixed strategy, where $\sigma_i(s_i) = 1$ for some $s_i \in S_i$; in other words, mixed strategies subsume pure strategies. Sometimes, a mixed strategy that places positive probability on all pure strategies is called a *fully* or *totally mixed strategy*.¹⁰

Definition 9 (Fully Mixed Strategy). A strategy, σ_i , for player i is fully (or completely, or totally) mixed if $\sigma_i(s_i) > 0$ for all $s_i \in S_i$.

As a piece of notation, we denote the set of probability distribution on S_i as $\Delta(S_i)$, which is the simplex on S_i . The space of mixed strategies then is $\Delta(S_i)$, which I will often denote as Σ_i .

Notice now that even if there is no role for nature in a game, when players use (non-degenerate) mixed strategies, this induces a probability distribution over terminal nodes of the game. But we can easily extend payoffs again to define payoffs over a profile of mixed strategies as follows:

$$u_i(\sigma_1, \dots, \sigma_I) = \sum_{s \in S} [\sigma_1(s_1) \sigma_2(s_2) \cdots \sigma_I(s_I)] u_i(s).$$

Remark 4. For the above formula to make sense, it is critical that each player is randomizing *independently*. That is, each player is independently tossing her own die to decide on which pure strategy to play. This rules out scenarios such as two players jointly observing the roll of a “public” die, and then correlating their choice of individual pure strategies based on the die’s outcome. This independence assumption can be weakened in a more advanced treatment, but we maintain it throughout much of this course, except for brief remarks and the discussion in [Subsection 3.5](#).

Return to [Example 2](#). A mixed strategy for Anne can be specified by a single number $p_1 \in [0, 1]$ so that p_1 is the probability of playing the pure strategy H . This implicitly defines the probability of playing pure strategy T as $1 - p_1$. On the other hand, for Bob, a mixed strategy is a triple, $(q_1, q_2, q_3) \in [0, 1]^3$, where q_1 is the probability of playing (H, H) , q_2 is the probability of playing (H, T) , q_3 is the probability of (T, H) , and $1 - q_1 - q_2 - q_3$ is the probability of playing (T, T) .

Behavioral Strategies In the context of extensive form representations, there is an alternative way one can think about making randomized choices. Rather than randomizing over pure strategies, why not define a plan of action that specifies separately randomizing over the set of available actions at each information node? That is, in [Example 2](#), why can’t Bob simply specify how to randomize over Heads and Tails in each of two different scenarios: if Anne plays H , and if Anne plays T . Such a formulation is in fact feasible, and is called a *behavioral strategy*.

Definition 10 (Behavioral Strategy). A *behavioral* strategy for player i is a function $\lambda_i : \mathcal{A} \times \mathcal{H}_i \rightarrow [0, 1]$ which assigns a probability $\lambda_i(a, H) \geq 0$ to each action $a \in \mathcal{A}$ at information set $H \in \mathcal{H}_i$, satisfying $\forall H \in \mathcal{H}_i$, $\lambda_i(a, H) = 0$ if $a \notin C(H)$ and $\sum_{a \in C(H)} \lambda_i(a, H) = 1$.

To be clear, a behavioral strategy for Bob in [Example 2](#) would be a pair $(q_1, q_2) \in [0, 1]^2$ such that q_1 is the probability of playing H if Anne has played H and q_2 is the probability of playing H if Anne has played T . Implicitly then, $1 - q_1$ is the probability of playing T if Anne has played H and $1 - q_2$ is the probability of playing T if Anne has played T . Compare this to a mixed strategy for Bob described earlier.

¹⁰ There is no well established terminology for a strategy in $\Sigma_i \setminus S_i$; let’s use *non-pure strategy*.

As you probably guessed, in games of perfect recall, behavioral strategies and mixed strategies are equivalent. That is, for any player, for any behavioral strategy there exists a mixed strategy that yields exactly the same distribution over terminal nodes given the strategies (behavioral or mixed) of other players, and vice-versa.¹¹ The formal theorem is this, where an outcome means a probability distribution over terminal nodes:

Theorem 1 (Kuhn’s Theorem). *For finite games with perfect recall, every mixed strategy of a player has an outcome-equivalent behavioral strategy, and conversely, every behavioral strategy has an outcome-equivalent mixed strategy.*

I won’t prove this Theorem though the intuition is straightforward (you will work through a detailed example in a homework problem). Given the result, in this course, we will be a little casual and blur the distinction between mixed and behavioral strategies. Often, it is more convenient to use behavioral strategies in extensive form representations, and mixed strategies when a game is in strategic form. See [Osborne and Rubinstein \(1994, Section 11.4\)](#) for an excellent discussion.

2.3. An Economic Example

To illustrate a strategic setting with direct application to the study of markets, here is a classic model of imperfect competition. Conveniently, it also serves to introduce infinite action spaces. There are two firms, call them 1 and 2, producing an identical product. Market demand is given by $Q(P)$ with inverse demand $P(Q)$, both of which are decreasing functions mapping \mathbb{R}_+ to \mathbb{R}_+ . Firm i produces a non-negative quantity q_i at cost $c_i(q_i)$, with $c_i(0) = 0$. Notice that since quantities and price are in a continuum here, none of the following games is finite. Nonetheless, we will just adapt our definitions from earlier in straightforward ways.

Simultaneous quantity-setting (Cournot) Suppose that each firm must simultaneously pick a quantity, and the market price gets determined as $P(q_1 + q_2)$. In Normal form, this game has $S_i = \mathbb{R}_+$, $s_i = q_i$, and $u_i(s_i, s_{-i}) = s_i P(s_1 + s_2) - c_i(s_i)$. We can draw it in extensive form too, using various game tree notations to represent the infinite number of available actions.

Simultaneous price-setting (Bertrand) Suppose that each firm must simultaneously pick a non-negative price, and the market quantity gets determined by $Q(\min\{p_1, p_2\})$, with all sales going to the firm with lower price and a 50-50 split in case of equal prices. In Normal form, this game has $S_i = \mathbb{R}_+$, $s_i = p_i$, and

$$u_i(s_i, s_{-i}) = \begin{cases} Q(s_i) s_i - c_i(Q(s_i)) & \text{if } s_i < s_{-i} \\ \frac{1}{2} Q(s_i) s_i - c_i(\frac{1}{2} Q(s_i)) & \text{if } s_i = s_{-i} \\ 0 & \text{if } s_i > s_{-i} \end{cases}$$

¹¹ The equivalence also breaks down when the set of actions available to a player — and hence nodes in the game — is infinite, and in particular a continuum; see [Aumann \(1964\)](#). A third way of defining randomization that works for a very general class of games (including many infinite games) is the *distributional strategy* approach of [Milgrom and Weber \(1985\)](#), but it comes at the cost of being unnecessarily cumbersome for finite games, so we don’t use it typically.

Sequential quantity-setting (Stackelberg) Suppose now that firms sequentially pick quantities, where firm 2 observes firm 1's choice before acting. (Before reading further, see if you can represent the game in Normal form; it is an excellent check on whether you have fully grasped the difference between strategies and actions.) In Normal form, this game has $s_1 = q_1$ and $S_1 = \mathbb{R}_+$, $s_2 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ (i.e. s_2 is function) and S_2 is a function space defined by $S_2 = \{\text{functions from } \mathbb{R}_+ \text{ to } \mathbb{R}_+\}$, and

$$\begin{aligned} u_1(s_1, s_2) &= s_1 P(s_1 + s_2(s_1)) - c_1(s_1) \\ u_2(s_1, s_2) &= s_2(s_1) P(s_1 + s_2(s_1)) - c_2(s_2(s_1)) \end{aligned}$$

Note that even though both firms' strategy spaces are infinite, firm 1's strategy space lies in a finite-dimensional space (it is 1-dimensional); whereas the dimensionality of firm 2's strategy space is (uncountably) infinite.¹²

3. Simultaneous-Move Games

In this Section, we are going to study behavior in Normal form games. One can either view this as taking a possibly-complicated extensive form game and collapsing it into its Normal form (in which case we may be losing relevant "dynamic" considerations that we will come back to later), or as simply considering "static" games where players only move once and move simultaneously. In any case, the primitive in this section will be a game in its Normal form, $\Gamma_N = \{I, S, u\}$.

3.1. Dominance

3.1.1. Strictly Dominant Strategies

You've probably heard of the *Prisoner's Dilemma* game (MWG Figure 8.B.1). I'm going to reinterpret it as a game of trust.

Example 4. (Trust Game) The Trust Game has the following Normal form.

		Player 2	
		<i>Trust</i>	<i>Cheat</i>
Player 1	<i>Trust</i>	5, 5	0, 10
	<i>Cheat</i>	10, 0	2, 2

Observe that regardless of what her opponent does, player i is strictly better off playing *Cheat* rather than *Trust*. This is precisely what is meant by a strictly dominant strategy.

Definition 11 (Strictly Dominant strategy). A strategy $s_i \in S_i$ is a strictly dominant strategy for player i if for all $\tilde{s}_i \neq s_i$ and all $s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(\tilde{s}_i, s_{-i})$.

¹² Dealing with infinite dimensional spaces typically adds additional technical complications, but may be unavoidable as in this example.

That is, a strictly dominant strategy for i uniquely maximizes her payoff for any strategy profile of all other players. If such a strategy exists, it is highly reasonable to expect a player to play it. In a sense, this is a consequence of a player's "rationality".¹³

In the Trust Game, if both players play their strictly dominant strategies, the outcome of the game is *(Cheat, Cheat)*. But notice that this is a Pareto-dominated outcome. Another way to say this is that if the players could somehow write a binding contract that requires them to both play *Trust*, they would be better off doing that rather than playing this Trust Game. Lesson: self-interested behavior in games may not lead to socially optimal outcomes. This stems from the possibility that a player's actions can have a negative externality on another player's payoff. (Aside: think about the connection to the First Welfare Theorem.)

Exercise 1. *Prove that a player can have at most one strictly dominant strategy.*

Notice that we defined strictly dominant strategies by only considering alternative pure strategies for both player i and his opponents. Would it matter if we instead allowed mixed strategies for either i or his opponents? The answer is no.

Theorem 2. *If s_i is a strictly dominant strategy for player i , then for all $\sigma_i \in \Sigma_i \setminus \{s_i\}$ and $\sigma_{-i} \in \Sigma_{-i}$, $u_i(s_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$.*

Proof. For any σ_{-i} , σ_i , and s_i , we can write $u_i(s_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$ as

$$\sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \left[u_i(s_i, s_{-i}) - \sum_{\tilde{s}_i \in S_i} \sigma_i(\tilde{s}_i) u_i(\tilde{s}_i, s_{-i}) \right] > 0. \quad (1)$$

Since s_i is strictly dominant, $u_i(s_i, s_{-i}) - u_i(\tilde{s}_i, s_{-i}) > 0$ for all $\tilde{s}_i \neq s_i$ and all s_{-i} . Hence, $u_i(s_i, s_{-i}) - \sum_{\tilde{s}_i \in S_i} \sigma_i(\tilde{s}_i) u_i(\tilde{s}_i, s_{-i}) > 0$ for any $\sigma_i \in \Sigma_i \setminus \{s_i\}$. This implies inequality (1). \square

Exercise 2. *Prove that there can be no strategy $\sigma_i \in \Sigma_i$ such that for all $s_i \in S_i$ and $s_{-i} \in S_{-i}$, $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$.*

The preceding Theorem and Exercise show that there is absolutely no loss in restricting attention to pure strategies for all players when looking for strictly dominant strategies.

3.1.2. Strictly Dominated Strategies

What about if a strictly dominant strategy doesn't exist, such as in the following game?

Example 5. A game defined by the Normal form

		Player 2		
		a	b	c
Player 1	A	5, 5	0, 10	3, 4
	B	3, 0	2, 2	4, 5

¹³ There is a better way to say this, but it is too much work at this point to be formal about it. Roughly though, if there is a strict dominant strategy, then no other strategy is optimal (in the sense of maximizing payoffs) regardless of what a player believes his opponents are playing. "Rationality" means that the player is playing optimally for *some* belief he holds about his opponents' play. Thus, he must play his strictly dominant strategy.

You can easily convince yourself that there are no strictly dominant strategies here for either player. However, notice that regardless of whether Player 1 plays A or B , Player 2 does strictly better by playing b rather than a . That is, a is “strictly dominated” by b . This motivates the next definition, where we allow for mixed strategies explicitly.

Definition 12 (Strict Dominance). A (mixed) strategy $\sigma_i \in \Sigma_i$ is strictly dominated for player i if there exists a (mixed) strategy $\tilde{\sigma}_i \in \Sigma_i$ such that for all $s_{-i} \in S_{-i}$, $u_i(\tilde{\sigma}_i, s_{-i}) > u_i(\sigma_i, s_{-i})$. In this case, we say that $\tilde{\sigma}_i$ strictly dominates σ_i .

In words, $\tilde{\sigma}_i$ strictly dominates σ_i if it yields a strictly higher payoff regardless of what (pure) strategy rivals use. Note that the definition allows for both $\tilde{\sigma}_i$ and σ_i to be pure or non-pure. Using this terminology, we can restate [Definition 11](#): a strategy is strictly dominant for a player if it strictly dominates all other strategies for that player. Just as it is reasonable to expect a player to play a strictly dominant strategy if one exists; it is likewise reasonable that a player will *not* play a strictly dominated strategy — a consequence of rationality, again.

Why were we explicit about allowing for a strategy to be dominated by a mixed strategy in the definition? Here is a game where it does matter.

Example 6.

		Player 2		
		a	b	c
Player 1	A	1, 5	0, 6	2, 9
	B	1, 9	2, 6	0, 5

There is no pure strategy that strictly dominates any other pure strategy in this game. However, the mixed strategy σ_2 where $\sigma_2(a) = \sigma_2(c) = 0.5$ strictly dominates the strategy b for Player 2.

Remark 5. By the same argument as in [Theorem 2](#), there is no loss in only comparing against all the pure strategies for all other players when evaluating whether there is a strictly dominated strategy for a particular player.

Remark 6. Convince yourself that a mixed strategy will be strictly dominated if it puts positive probability on any pure strategy that is strictly dominated. (This implies what we already noted: a strictly dominant strategy must be a pure strategy.) However, a mixed strategy may be strictly dominated even if none of the pure strategies it puts positive probability on are strictly dominated. Check this in a variant of [Example 6](#) where b gives player 2 a payoff of 8 regardless of what player 1 does.

3.1.3. Iterated Deletion of Strictly Dominated Strategies

Return now to [Example 5](#). We argued that a is strictly dominated (by b) for Player 2; hence rationality of Player 2 dictates she won’t play it. But now, we can push the logic further: if Player 1 knows that Player 2 is rational, he should realize that Player 2 will not play strategy a . Notice that we are now moving from the rationality of each player to the *mutual knowledge* of each player’s rationality. That is, not only are Player 1 and Player 2 rational, but moreover, Player 1 knows that Player 2 is rational (and vice-versa). Assuming this is the case, once Player 1 realizes that 2 will not play a and “deletes” this strategy from the strategy space, then strategy A becomes strictly dominated by strategy B for Player 1. So now, if we iterate the knowledge of rationality once again (that is: there is mutual knowledge of rationality, and moreover, Player 2

knows that Player 1 knows that Player 2 is rational), then Player 2 realizes that 1 will not play A , and hence “deletes” A , whereafter b is strictly dominated by c for Player 2. Thus, Player 2 should play c . We have arrived at a “solution” to the game through the *iterated deletion of strictly dominated strategies* (IDSDS): (B, c) .

Here is a formal definition of IDSDS:

Definition 13. Given a game (I, S, u) , the set $S^\infty \subseteq S$ survives IDSDS if $S^\infty = \times_i S_i^\infty$ and there is a collection $(S_i^k)_{k \geq 0}$ for each player i satisfying:

1. $S_i^0 = S_i$ and $S_i^\infty = \bigcap_{k \geq 0} S_i^k$.
2. For all $k \geq 1$, $S_i^k \subseteq S_i^{k-1}$.
3. For all $k \geq 1$, each $s_i \in S_i^{k-1} \setminus S_i^k$ is strictly dominated in the game $(I, \times_j S_j^{k-1}, u)$.
4. No $s_i \in S_i^\infty$ is strictly dominated in the game (I, S^∞, u) .

Definition 14 (Strict-dominance solvable). A game is strict-dominance solvable if iterated deletion of strictly dominated strategies results in a unique strategy profile, i.e. there is a S^∞ surviving IDSDS with $|S^\infty| = 1$.

Since in principle we might have to iterate numerous times in order to solve a strict-dominance solvable game, the process can effectively only be justified by *common knowledge* of rationality.¹⁴ As with strictly dominant strategies, it is also true that most games are not strict-dominance solvable. Consider for example MP-A (Example 1): no strategy is strictly dominated.

You might worry whether the order in which we delete strategies iteratively matters. In other words, can different deletion processes produce different S^∞ in Definition 13? Insofar as we are working with *strictly* dominated strategies so far (and finite games), this is not the case.

Remark 7. In finite games, the order of deletion does not affect the set of strategies that survive the process of iterated deletion of strictly dominated strategies. In particular, if a game is strict-dominance solvable, the outcome is independent of the order in which strategies are iteratively deleted according to strict dominance.

You will be asked to prove the above claim as a homework exercise. The following example demonstrates the potential power of iteratively deleting strictly dominated strategies.

Example 7 (Linear Cournot). This is a specialized version of the Cournot competition game we introduced earlier. Suppose that inverse market demand is given by a linear function $p(Q) = a - bQ$, the cost functions for both firms are also linear, $c(q_i) = cq_i$, and the linear payoff functions are

$$u_i(q_i, q_{-i}) = q_i(a - b(q_i + q_{-i})) - cq_i,$$

which simplify to

$$u_i(q_i, q_{-i}) = (a - c)q_i - bq_i^2 - bq_iq_{-i}.$$

Assume that $a, b, c > 0$ and moreover $a > c$. To solve this game by iterated deletion of strictly dominated strategies, first define the “reaction” or “best response” functions $r_i : [0, \infty) \rightarrow [0, \infty)$ which specify firm i ’s

¹⁴ This isn’t quite right: in finite games, we would only need to iterate a finite number of times; hence in any given game, common knowledge of rationality isn’t quite necessary. But an arbitrarily large order of iteration of knowledge may be needed.

optimal output for any given level of its opponent's output. These are computed through the first order conditions for profit maximization (assuming an interior solution),¹⁵

$$a - c - 2br(q_{-i}) - bq_{-i} = 0.$$

Hence,

$$r(q_{-i}) = \frac{a - c}{2b} - \frac{q_{-i}}{2},$$

where I am dropping the subscript on r since firms are symmetric.

Clearly, r is a decreasing function; that is, the more the opponent produces, the less a firm wants to produce. The strategy space we start with for each firm is $S_i^0 = [0, \infty)$. For each firm, playing anything above $r(0)$ is strictly dominated, since the opponent plays at least 0. Hence, deleting strictly dominated strategies once yields $S_i^1 = [0, r(0)]$. Now, since the opponent plays *no more than* $r(0)$, it is (iteratively) strictly dominated for firm i to play *less than* $r(r(0))$, which I'll denote $r^2(0)$. Thus, the second round of iterated deletion yields the strategy space $S_i^2 = [r^2(0), r(0)]$. In the third round, since the opponent is playing *at least* $r^2(0)$, it is (iteratively) dominated for a firm to play *more than* $r(r^2(0))$, which of course I denote $r^3(0)$. So eliminating iteratively dominated strategies yields the space $S_i^3 = [r^2(0), r^3(0)]$, ... and so on, ad infinitum. The lower bounds of these intervals form a sequence $r^{2n}(0)$, and the upper bounds form a sequence $r^{2n+1}(0)$. Define $\alpha \equiv \frac{a-c}{b}$. You can check by expanding out some of the $r^n(0)$ formulae that for all $n = 1, 2, \dots$,

$$r^n(0) = -\alpha \sum_{k=1}^n \left(-\frac{1}{2}\right)^k.$$

This is a convergent series (by absolute convergence), and hence the intervals S_i^n converge to a single point. Thus the game is strict-dominance solvable. The solution can be found by evaluating the infinite series $-\alpha \sum_{k=1}^{\infty} \left(-\frac{1}{2}\right)^k$, which turns out to be $\frac{a-c}{3b}$.¹⁶ \square

3.1.4. Weakly Dominated Strategies

Consider the following game, which has no strictly dominated strategies (hence is not strict-dominance solvable).

Example 8. The normal form for a game is

		Player 2	
		a	b
Player 1	A	3, 4	4, 3
	B	5, 3	3, 5
	C	5, 3	4, 3

¹⁵ Check that the second order condition is satisfied.

¹⁶ To see this, observe that we can write

$$-\sum_{k=1}^{\infty} \left(-\frac{1}{2}\right)^k = -\sum_{k=1}^{\infty} \left[\left(-\frac{1}{2}\right)^{2k-1} + \left(-\frac{1}{2}\right)^{2k} \right] = \sum_{k=1}^{\infty} \left[\left(\frac{1}{2^{2k-1}}\right) - \left(\frac{1}{2^{2k}}\right) \right] = \sum_{k=1}^{\infty} \frac{1}{4^k} = \sum_{k=0}^{\infty} \frac{1}{4^k} - 1 = \frac{4}{3} - 1 = \frac{1}{3}.$$

Exercise 3. Prove that there are no strictly dominated pure strategies in this game.

Nonetheless, notice that unless Player 1 is *absolutely sure* that Player 2 is going to play a , he is strictly better off playing C rather than B . That is to say, for any (possibly mixed) strategy $\sigma_2 \neq a$, $u_1(C, \sigma_2) > u_1(B, \sigma_2)$. Moreover, $u_1(C, a) = u_1(B, a)$. Hence, C will do at least as well as B , and could do better. We say that B is weakly dominated by C . Generalizing,

Definition 15 (Weak dominance). A strategy $\sigma_i \in \Sigma_i$ is weakly dominated for player i if there exists a mixed strategy $\tilde{\sigma}_i \in \Sigma_i$ such that for all $s_{-i} \in S_{-i}$, $u_i(\tilde{\sigma}_i, s_{-i}) \geq u_i(\sigma_i, s_{-i})$ and for some $s_{-i} \in S_{-i}$, $u_i(\tilde{\sigma}_i, s_{-i}) > u_i(\sigma_i, s_{-i})$. In this case, we say that $\tilde{\sigma}_i$ weakly dominates σ_i .

We say that a strategy s_i is weakly dominant if it weakly dominates all other strategies, $\tilde{s}_i \neq s_i$. As with the case of strict dominance, it is important to allow for mixed strategies for player i in the definition of weak dominance, but not for the other players. Notice that we cannot appeal to “rationality” to justify the deletion of weakly dominated strategies, since a player might optimally play a weakly dominated strategy if he were certain that his opponents were going to play a particular strategy profile. That said, it has a lot of plausibility and can be useful in simplifying a complicated game. Just as with strict dominance, we can iteratively delete weakly dominated strategies. However, there is a subtlety here, because the order of deletion matters. To see this, continue the Example above. Having deleted B through weak dominance (by C), we can then delete for Player 2 strategy b since it is weakly dominated by strategy a , and finally iterate once more and delete strategy A for player 1. The iterative process has yielded the outcome (C, a) . On the other hand, starting from the outset all over, observe that A is also weakly dominated by C . If we delete A in the first round (rather than B as before), we can then delete a for Player 2 since it is weakly dominated by b ; and in third round, now delete B as it is weakly dominated by C . This process has led to the outcome (C, b) . This motivates two remarks.

Remark 8. The order of deletion can matter when iteratively deleting weakly dominated strategies. Arguably, this makes it less appealing as a solution concept. Oftentimes, in any round of deletion, we will delete *all* strategies that are weakly dominated for a player; this is referred to as (*iterated*) *admissibility*.¹⁷

Remark 9. There is no completely standard definition of what it means for a game to be weak-dominance solvable. Sometimes, it means that there is some order of iteratively deleting weakly dominated strategies that leads to a single strategy profile. Other times, it means that no matter what order in which we iteratively delete weakly dominated strategies, we end up at a unique (at least in terms of payoffs) strategy profile.¹⁸

One other terminological point: saying that a strategy dominates (or is dominated by) another is potentially ambiguous with regards to strict or weak dominance. Typically, dominance without a caveat means strict dominance, but the literature is not uniformly careful about this (I will try to remember to be!). Similarly, saying that a game is *dominance solvable* can mean either iterated deletion of strictly or weakly dominated strategies — typically, it means strict-dominance solvable.

¹⁷ If you are interested in an advanced “epistemics” approach to solution concepts, see [Brandenburger, Friedenberg, and Keisler \(2008\)](#) on foundations for iterated admissibility.

¹⁸ [Osborne and Rubinstein \(1994, p. 63\)](#) has a third (and substantively different) definition based on the idea of deleting all weakly dominated strategies in each round of deletion, i.e. iterated admissibility. Note that if we used this procedure on [Example 8](#), iterative deletion would not get to a unique strategy profile (or unique payoff profile): after deleting both A and B in the first round, we can proceed no further.

Exercise 4. In MP-B (Example 2), which strategies are strictly dominated for Player 2? Which are weakly dominated? Does the game have a unique prediction through iterated deletion of strictly-dominated strategies? What about through iterated deletion of weakly-dominated strategies?

Note that lots of games have no weakly dominated strategies, such as MP-A. On the other hand, there are some interesting and useful examples of games that do.

Example 9 (Second-Price Auction). A seller has one indivisible object. There are I bidders with respective valuations $0 \leq v_1 \leq \dots \leq v_I$ for the object; these valuations are common knowledge. The bidders simultaneously submit bids $s_i \in [0, \infty)$. The highest bidder wins the object and pays the *second highest* bid. Given a profile of bids, s , let $W(s) \equiv \{k : \forall j, s_k \geq s_j\}$ be the set of highest bidders. Bidder i gets utility

$$u_i(s_i, s_{-i}) = \begin{cases} v_i - \max_{j \neq i} s_j & \text{if } s_i > \max_{j \neq i} s_j \\ \frac{1}{|W(s)|}(v_i - s_i) & \text{if } s_i = \max_{j \neq i} s_j \\ 0 & \text{if } s_i < \max_{j \neq i} s_j. \end{cases}$$

In this game, it is weakly dominant for each player to bid his true valuation, that is to play $s_i = v_i$. To see this, define $m(s_{-i}) \equiv \max_{j \neq i} s_j$.

Suppose first $s_i > v_i$. Then for any strategy profile, s_{-i} , if $m(s_{-i}) > s_i$, $u_i(s_i, s_{-i}) = u_i(v_i, s_{-i}) = 0$. If $m(s_{-i}) \leq v_i$, then $u_i(s_i, s_{-i}) = u_i(v_i, s_{-i}) \geq 0$. Finally, if $m(s_{-i}) \in (v_i, s_i]$, then $u_i(s_i, s_{-i}) < 0 = u_i(v_i, s_{-i})$. Hence, $s_i = v_i$ weakly dominates all $s_i > v_i$.

Consider next $s_i < v_i$. If $m(s_{-i}) \geq v_i$, then $u_i(s_i, s_{-i}) = u_i(v_i, s_{-i}) = 0$. If $m(s_{-i}) < s_i$, then $u_i(s_i, s_{-i}) = u_i(v_i, s_{-i}) > 0$. Finally, if $m(s_{-i}) \in [s_i, v_i]$, then $0 = u_i(s_i, s_{-i}) < u_i(v_i, s_{-i})$. Hence, $s_i = v_i$ weakly dominates all $s_i < v_i$.

Therefore, in a second price auction, it seems reasonable that rational bidders should bid their true valuation. The bidder with the highest valuation wins, and pays the second highest valuation, v_{I-1} . Note that since bidding one's valuation is a weakly dominant strategy, it does not matter even if player i does not know the other players' valuations — even if valuations are only known to each player privately (rather than being common knowledge), it still remains a weakly dominant strategy to bid truthfully. We'll come back to this last point later in the course. \square

3.2. Rationalizability

Let's start by generalizing the notion of a best response we used in Example 7.

Definition 16 (Best Response). A strategy $\sigma_i \in \Sigma_i$ is a best response to the strategy profile $\sigma_{-i} \in \Sigma_{-i}$ if $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tilde{\sigma}_i, \sigma_{-i})$ for all $\tilde{\sigma}_i \in \Sigma_i$. A strategy $\sigma_i \in \Sigma_i$ is never a best response if there is no $\sigma_{-i} \in \Sigma_{-i}$ for which σ_i is a best response.

The idea is that a strategy, σ_i , is a best response if there is some strategy profile of the opponents for which σ_i does at least as well as any other strategy. Conversely, σ_i is never a best response if for every strategy profile of the opponents, there is some strategy that does strictly better than σ_i . Let me note here that we can also think of a best response in terms of optimal play against a *conjecture* about opponents' strategies. In particular, even if one knows (believes) that an opponent is using a pure strategy, one may be unsure *which* pure strategy, which is tantamount to the opponent mixing.

Clearly, in any game, a strategy that is strictly dominated is never a best response. In 2-player games, a strategy that is never a best response is strictly dominated. While this equivalence is true for all strategies in 2-player games, the exercise below asks you to prove it for pure strategies only (needless to say, because proving it for mixed strategies requires advanced tools).

Exercise 5. *Prove that in 2-player games, a pure strategy is never a best response if and only if it is strictly dominated.*¹⁹

In games with more than 2 players, there may be strategies that are *not* strictly dominated that are nonetheless never best responses.²⁰ As before, it is a consequence of “rationality” that a player should not play a strategy that is never a best response. That is, we can delete strategies that are never best responses. You can guess what comes next: by iterating on the knowledge of rationality, we iteratively delete strategies that are never best responses. The set of strategies for a player that survives this iterated deletion of never best responses is called her set of *rationalizable strategies* (Bernheim, 1984; Pearce, 1984).²¹ A constructive way to define this is as follows.

Definition 17 (Rationalizability).

1. $\sigma_i \in \Sigma_i$ is a 1-rationalizable strategy for player i if it is a best response to some strategy profile $\sigma_{-i} \in \Sigma_{-i}$.
2. $\sigma_i \in \Sigma_i$ is a k -rationalizable strategy ($k \geq 2$) for player i if it is a best response to some strategy profile $\sigma_{-i} \in \Sigma_{-i}$ such that each σ_j is in the convex hull of the set of $(k - 1)$ -rationalizable strategies for player $j \neq i$.
3. $\sigma_i \in \Sigma_i$ is rationalizable for player i if it is k -rationalizable for all $k \geq 1$.

Note a subtlety: at any round $k > 1$, we consider the *convex hull* of each opponent’s $(k - 1)$ -rationalizable strategies.²² The reason is that at any round, the set of “surviving strategies” for a player may not be convex (e.g., two pure strategies may each be best responses to some opponent strategy, but a mixture of them may not; recall Remark 6). Yet a player may be unsure which of her opponents’ surviving strategies will be used.

Remark 10. You should convince yourself that any strategy that does not survive iterated deletion of strictly dominated strategies (Subsection 3.1.3) is not rationalizable. (This follows from the earlier comment that a strictly dominated strategy is never a best response.) Thus, the set of rationalizable strategies is no larger than the set of strategies that survives iterated deletion of strictly dominated strategies. In this sense,

¹⁹ The “only if” part is a non-trivial problem, and I will give you a homework question that breaks down the steps. But as a hint now, you can use Kakutani’s fixed point theorem (see Lemma 1 on page 20).

²⁰ This stems from the fact that we are assuming that each of player i ’s opponents is choosing his strategy independently. If we were to allow for *correlated* strategies, then the notions of being strictly dominated and never a best response coincide regardless of the number of players. This points to *why* the notions coincide in 2-player games even in our way of doing things — each player only has one opponent, hence trivially, the opponents are choosing strategies independently.

²¹ Note that the order of deleting strategies that are never best responses doesn’t matter, since we are deleting strategies that are not even *weakly* optimal for some strategy profile of the opponents. This is analogous to the case with deleting strictly dominated strategies.

²² The convex hull of a set is the smallest convex set that contains that set. We interpret each element of the convex hull of a set of strategies as itself a mixed strategy.

rationalizability is (*weakly*) *more restrictive* than iterated deletion of strictly dominated strategies. It turns out that in 2-player games, the two concepts coincide. In n -player games ($n > 2$), they don't have to.²³

One way to think of rationalizable strategies is through an *infinite or circular chain of justification*. This is best illustrated through examples. Let us return to [Example 6](#). Is A rationalizable? Yes, by the following chain of justification: A is a best response to c , which is a best response to A . Here is another chain of justification that works: A is a best response to a , which is a best response to B , which is a best response to a . Is B rationalizable? Yes: B is a best response to a , which is a best response to B . Similarly, a and c are rationalizable, but b is not (and you know it is not rationalizable because we already saw that it is strictly dominated).

To see that how a chain of justification works when it involves a mixed strategy, consider a modification of [Example 6](#) so that the payoff cell for (B, c) is $(0, -5)$, everything else staying the same. Then, b is rationalizable by the following chain: b is a best response to the mixed strategy $(0.75)A + (0.25)B$,²⁴ this mixed strategy for Row is a best response to a , which is a best response to B , which is a best response to a .

Rationalizability is still a weak solution concept in the sense that the set of rationalizable strategies is typically large in any complicated game. For example, even in something as simple as MP-A ([Example 1](#)), every strategy is rationalizable. This is also true in the richer example of the 2nd price auction ([Example 9](#)).

Exercise 6. *Prove that every strategy is rationalizable in the 2nd price auction.*²⁵

I conclude this section by emphasizing that rationalizability is as far as we can go (in terms of refining our predictions for the outcome of a game) by using only common knowledge of rationality and the structure of the game. A little more precisely: *common knowledge of rationality and the structure of the game imply that players will play rationalizable strategies; conversely, any profile of rationalizable strategies is consistent with common knowledge of rationality and the structure of the game.*

3.3. Nash Equilibrium

Now we turn to the most well-known solution concept in game theory. We'll first discuss *pure strategy Nash equilibrium* (PSNE), and then later extend to mixed strategies.

3.3.1. Pure Strategy Nash Equilibrium

Definition 18 (PSNE). A strategy profile $s = (s_1, \dots, s_I) \in S$ is a pure strategy Nash equilibrium if for all i and $\tilde{s}_i \in S_i$, $u_i(s_i, s_{-i}) \geq u_i(\tilde{s}_i, s_{-i})$.

In a Nash equilibrium, each player's strategy must be a best response to *those strategies of his opponents that are components of the equilibrium*. Note that if s is a pure strategy Nash equilibrium, then it holds that for all i and $\sigma_i \in \Sigma_i$, $u_i(s_i, s_{-i}) \geq u_i(\sigma_i, s_{-i})$. In other words, s_i is a best response to s_{-i} if no pure strategy for i is strictly better against s_{-i} .

²³ Again, this is because we are not allowing for correlation in strategies across players. If we did, the concepts would coincide in general. See [fn. 20](#), and [Osborne and Rubinstein \(1994, Sections 4.1 and 4.2\)](#) for more details.

²⁴ Given that Row is playing $(0.75)A + (0.25)B$, Column gets expected payoff of 6 by playing either a or b , and 5.5 from playing c .

²⁵ The result does *not* directly follow from the fact that no strategy is strictly dominated, since the equivalence between iterated deletion of strictly dominated strategies and rationalizability is only for 2-player games.

Remark 11 (Nash equilibrium). There are various conceptual points to make about Nash equilibrium:

- Unlike with our earlier solution concepts (dominance and rationalizability), Nash equilibrium applies to a profile of strategies rather than any individual’s strategy. When people say “Nash equilibrium strategy”, what they mean is “a strategy that is part of a Nash equilibrium profile.”
- The term *equilibrium* is used because it connotes that *if a player knew* that his opponents were playing the prescribed strategies, then she is playing optimally by following her prescribed strategy. In a sense, this is like a “rational expectations” equilibrium, in that in a Nash equilibrium, a player’s beliefs about what his opponents will do get confirmed (where the beliefs are precisely the opponents’ prescribed strategies).
- Rationalizability only requires a player play optimally with respect to some “reasonable” conjecture about the opponents’ play, where “reasonable” means that the conjectured play of the rivals can also be justified in this way. On the other hand, Nash requires that a player play optimally with respect to what his opponents *are actually* playing. That is to say, the conjecture she holds about her opponents’ play is *correct*.
- The above point makes clear that Nash equilibrium is not simply a consequence of (common knowledge of) rationality and the structure of the game. Clearly, *each player’s strategy in a Nash equilibrium profile is rationalizable, but lots of rationalizable profiles are not Nash equilibria*.

Let’s look at some examples of how this works. In MP-B, the two PSNE are (H, TH) and (T, TH) . MP-A has no PSNE (Why?). In [Example 8](#), there are also two PSNE: (C, a) and (C, b) . Similarly, in MP-N ([Example 3](#)), there are two Nash equilibria: (H, H) and (T, T) . This last example really emphasizes the assumption of *correctly conjecturing* what your opponent is doing — even though it seems impossible to say which of these two Nash equilibria is “more reasonable”, any one is an equilibrium only if each player can correctly forecast that his opponent is playing in the prescribed way.

Exercise 7. Verify that in the Linear Cournot game ([Example 7](#)) and the 2nd price auction ([Example 9](#)), the solutions found via iterated dominance are pure strategy Nash equilibria. Prove that in the Linear Cournot game, there is a unique PSNE, whereas there are multiple PSNE in the 2nd price auction. (Why the difference?)

Remark 12. Every finite game of perfect information has a *pure strategy* Nash equilibrium.²⁶ This holds true for “dynamic” games as well, so we’ll prove it generally when we get there.

We next want to give some sufficient (but certainly not necessary) conditions for the existence of a PSNE, and prove it. To do so, we need the powerful fixed point theorem of Kakutani.

Lemma 1 (Kakutani’s FPT). *Suppose that $X \subset \mathbb{R}^N$ is a non-empty, compact, convex set, and that $f : X \rightrightarrows X$ is a non-empty and convex-valued correspondence with a closed graph.²⁷ Then there exists $x^* \in X$ such that $x^* \in f(x^*)$.*

²⁶ Recall that perfect information means all information sets are singletons.

²⁷ The correspondence f is convex-valued if $f(x)$ is a convex set for all $x \in X$; it has a closed graph if for all sequences $x^n \rightarrow x$ and $y^n \rightarrow y$ such that $y^n \in f(x^n)$ for all n , $y \in f(x)$. (Technical remark: in this context where X is a compact subset of \mathbb{R}^N , f having a closed graph is equivalent to it being upper hemi-continuous (uhc); but in more general settings, the closed graph property is necessary but not sufficient for uhc.)

(Question: why do we need the convex-valued assumption?)

It is also useful to define the idea of a best response correspondence (first recall the idea of a strategy being a best response, see [Definition 16](#)).

Definition 19 (BR Correspondence). The best response correspondence for player i , $b_i : S_{-i} \rightrightarrows S_i$, is defined by $b_i(s_{-i}) = \{s_i \in S_i \mid u_i(s_i, s_{-i}) \geq u_i(\tilde{s}_i, s_{-i}) \ \forall \tilde{s}_i \in S_i\}$.

By this definition, it follows that $s \in S$ is a pure strategy Nash equilibrium if and only if $s_i \in b_i(s_{-i})$ for all i . We apply this observation in proving the following existence theorem.

Theorem 3 (Existence of PSNE). Suppose each $S_i \subset \mathbb{R}^N$ is compact and convex (and non-empty); and each $u_i : S \rightarrow \mathbb{R}$ is continuous in s and quasi-concave in s_i .²⁸ Then there exists a PSNE.

Proof. Define the correspondence $b : S \rightrightarrows S$ by $b(s) = b_1(s_{-1}) \times \cdots \times b_I(s_{-I})$. A Nash equilibrium is a profile s^* such that $s^* \in b(s^*)$. Clearly, b is a correspondence from the non-empty, convex, and compact set S to itself.

Step 1: For all s , $b(s)$ is non-empty. This follows from the fact that each b_i is the set of maximizers of a continuous function u_i over a compact set S_i , which is non-empty by the Weierstrass Theorem of the Maximum.

Step 2: Each b_i (and hence b) is convex-valued. Pick any s_{-i} , and suppose that $s_i, \tilde{s}_i \in b_i(s_{-i})$. By definition of b_i , there is some \bar{u} such that $\bar{u} = u_i(s_i, s_{-i}) = u_i(\tilde{s}_i, s_{-i})$. Applying quasi-concavity of u_i , $u_i(\lambda s_i + (1 - \lambda)\tilde{s}_i, s_{-i}) \geq \bar{u}$ for all $\lambda \in [0, 1]$. But this implies that $\lambda s_i + (1 - \lambda)\tilde{s}_i \in b_i(s_{-i})$, proving the convexity of $b_i(s_{-i})$.

Step 3: Each b_i (and hence b) has a closed graph. Suppose $s_i^n \rightarrow s_i$ and $s_{-i}^n \rightarrow s_{-i}^n$ with $s_i^n \in b_i(s_{-i}^n)$ for all n . Then for all n , $u_i(s_i^n, s_{-i}^n) \geq u_i(\tilde{s}_i, s_{-i}^n)$ for all $\tilde{s}_i \in S_i$. Continuity of u_i implies that for all $\tilde{s}_i \in S_i$, $u_i(s_i, s_{-i}) \geq u_i(\tilde{s}_i, s_{-i})$; hence $s_i \in b_i(s_{-i})$.

Thus, all the requirements for Kakutani's FPT are satisfied. There exists a fixed point, $s^* \in b(s^*)$, which is a PSNE. \square

Remark 13. Note carefully that in Step 2 above, quasi-concavity a player's utility function plays a key role. It would be *wrong* (though tempting!) to say that $u_i(\lambda s_i + (1 - \lambda)\tilde{s}_i, s_{-i}) = \bar{u}$ by vNM. This is wrong because $\lambda s_i + (1 - \lambda)\tilde{s}_i$ is merely a point in the space S_i (by convexity of S_i), and *should not* be interpreted as a mixed strategy that places λ probability of s_i and $(1 - \lambda)$ on \tilde{s}_i . Right now you should think of $\lambda s_i + (1 - \lambda)\tilde{s}_i$ as just another pure strategy, which happens to be a convex combination of s_i of \tilde{s}_i . Indeed, this is the only point in the proof where quasi-concavity of u_i is used; come back and think about this point again after you see [Theorem 4](#).

Remark 14. A finite strategy profile space, S , cannot be convex (why?), so this existence Theorem is only useful for infinite games.

3.3.2. Mixed Strategy Nash Equilibrium

The previous remark motivates the introduction of mixed strategies. It is straightforward to extend our definition of Nash equilibrium to this case, and this subsumes the earlier definition of PSNE.

²⁸ Recall that a function $g : \mathbb{R}^N \rightarrow \mathbb{R}$ is quasi-concave if, for all $c \in \mathbb{R}$ and $x, y \in \mathbb{R}^N$ such that $g(x) \geq c$ and $g(y) \geq c$, $g(\lambda x + (1 - \lambda)y) \geq c \ \forall \lambda \in [0, 1]$.

Definition 20 (Nash Equilibrium). A strategy profile $\sigma = (\sigma_1, \dots, \sigma_I) \in \Sigma$ is a Nash equilibrium if for all i and $\tilde{\sigma}_i \in \Sigma_i$, $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tilde{\sigma}_i, \sigma_{-i})$.

To see why considering mixed strategies are important, observe that Matching Pennies version A (Example 1) has no PSNE, but does have a mixed strategy Nash equilibrium (MSNE): each player randomizes over H and T with equal probability. In fact, when player i behaves in this way, player $j \neq i$ is exactly indifferent between playing H or T ! That is, in the MSNE, *each player who is playing a mixed strategy is indifferent amongst the set of pure strategies he is mixing over*. This remarkable property is very general and is essential in helping us solve for MSNE in many situations. Before tackling that, let's first give an existence Theorem for Nash equilibria in finite games using mixed strategies.

Theorem 4 (Existence of NE). *Every finite game has a Nash equilibrium (possibly in mixed strategies).*

Proof. For each i , given the finite space of pure strategies, S_i , the space of mixed strategies, Σ_i , is a (non-empty) compact and convex subset of $\mathbb{R}^{|S_i|}$. The utility functions $u_i : \Sigma \rightarrow \mathbb{R}$ defined by

$$u_i(\sigma_1, \dots, \sigma_I) = \sum_{s \in S} [\sigma_1(s_1) \sigma_2(s_2) \cdots \sigma_I(s_I)] u_i(s)$$

are continuous in σ and quasi-concave in σ_i (by linearity). Thus, Theorem 3 implies that there is a pure strategy Nash equilibrium of the infinite normal-form game $\langle I, \{\Sigma_i\}, \{u_i\} \rangle$; this profile is a (possibly degenerate) mixed strategy Nash equilibrium of the original finite game. \square

Remark 15. The critical need to allow for mixed strategies is that in finite games, the pure strategy space is not convex, but allowing players to mix over their pure strategies “convexifies” the space.

This does *not* mean that mixed strategies are not important in infinite games when the pure strategy space is convex. I illustrate through the following examples showing that a convex infinite game which does not have a pure strategy Nash equilibrium can nonetheless have a mixed strategy equilibrium. The setting is the price-competition analogue to the Linear Cournot we considered earlier (in Example 7).

Example 10 (Symmetric Linear Bertrand). Two firms compete in Bertrand price competition, each with identical linear costs given by $c(q_i) = cq_i$ ($c \geq 0$); market demand is given by a smooth decreasing function, $Q(p) > 0$. We will show that the unique PSNE is $s_1 = s_2 = c$. It is straightforward to verify that this is a PSNE, so let's argue that there cannot be any other PSNE. Wlog, suppose there is a PSNE with $s_1 \geq s_2$, with at least one $s_i \neq c$. There are three exhaustive cases to consider.

1. $s_2 > c$. Then firm 1 can do better by playing $\tilde{s}_1 = s_2 - \varepsilon$ (for some $\varepsilon > 0$), by the continuity of $Q(p)$.
2. $s_2 = c$. It must be that $s_1 > c$, but now firm 2 can do better by playing $\tilde{s}_2 = s_1$ (it makes positive rather than 0 profit).
3. $s_2 < c$. Then firm 2 is making losses, and can do better by playing $\tilde{s}_2 = c$. \square

Example 11 (Asymmetric Linear Bertrand). Continue with the setup of Example 10, but now suppose that costs are asymmetric, wlog $0 \leq c_1 < c_2$. Assume that there exists $\delta > 0$ such that $Q(p)[p - c_1]$ is strictly increasing on $p \in [c_1, c_2 + \delta]$.²⁹ You are asked to prove as an exercise that there is no PSNE in this game.

²⁹ The economics of this assumption is that it guarantees that the price firm 1 would charge if it were a monopolist in the market is strictly larger than c_2 .

However, we can construct a MSNE as follows: firm 1 plays $s_1 = c_2$, and firm 2 plays a mixed strategy, σ_2 which randomizes uniformly over all pure strategies in $[c_2, c_2 + \varepsilon]$, for some $\varepsilon > 0$. Denote the cdf of firm 2's price choice by $F(p; \varepsilon)$, with density $f(p; \varepsilon) = 1/\varepsilon$. Clearly, firm 2 is playing a best response to s_1 , since it earns 0 profits in equilibrium and cannot do any better. To show that firm 1 is playing optimally, we only need to show that it does not prefer to deviate to any $\tilde{s}_1 \in (c_2, c_2 + \varepsilon]$, since clearly any $\tilde{s}_1 < c_2$ or $\tilde{s}_1 > c_2 + \varepsilon$ does strictly worse than $s_1 = c_2$. Consider firm 1's profit function by picking any price in $[c_2, c_2 + \varepsilon]$, given firm 2's strategy:

$$\pi(p) = (1 - F(p; \varepsilon))(p - c_1)Q(p).$$

Differentiating gives

$$\pi'(p) = (1 - F(p; \varepsilon))[(p - c_1)Q'(p) + Q(p)] - f(p; \varepsilon)(p - c_1)Q(p).$$

By picking $\varepsilon > 0$ small enough we can make $\min_{p \in [c_2, c_2 + \varepsilon]} f(p; \varepsilon)$ arbitrarily large; it follows that for small enough $\varepsilon > 0$, $\pi'(p) < 0$ for all $p \in [c_2, c_2 + \varepsilon]$, which implies $s_1 = c_2$ is optimal for firm 1. \square

Exercise 8. *Prove that there is no PSNE in [Example 11](#). Why does [Theorem 3](#) not apply to this infinite game?*

This construction of a mixed strategy equilibrium to the asymmetric Bertrand game is due to [Blume \(2003\)](#). It obviates having to resort to “tricks” such as discretizing the price space, resolving price ties in favor of the lower cost firm, etc. — such assumptions would be needed to get existence in pure strategies.

***Infinite Pure Strategy Spaces** Let us record a mixed-strategy existence theorem for infinite games.

Theorem 5. *Suppose each S_i is a compact subset of a metric space and each $u_i : S \rightarrow \mathbb{R}$ is continuous. Then the game has a (possibly mixed) Nash equilibrium.*

The theorem relies on the following generalization of the Kakutani FPT to infinite-dimensional spaces; consult [Aliprantis and Border \(2006\)](#) for more on these definitions and results.

Lemma 2 (Fan-Glicksberg FPT). *Suppose that X is a non-empty, compact, convex subset of a locally convex Hausdorff (topological vector) space, and that $f : X \rightrightarrows X$ is a non-empty and convex-valued correspondence with a closed graph.³⁰ Then there exists $x^* \in X$ such that $x^* \in f(x^*)$.*

A proof of [Theorem 5](#) proceeds similarly to that of [Theorem 4](#) (and hence uses arguments similar to the proof of [Theorem 3](#)), but invoking the Fan-Glicksberg FPT instead of Kakutani. One invokes two facts about a non-empty compact metric space X : (i) the space of (Borel) probability measures on X is a non-empty, compact, and convex subset of a locally convex Hausdorff space;³¹ and (ii) if $g : X \rightarrow \mathbb{R}$ is a continuous function then $\int g d\mu$ is continuous in the probability measure μ . Point (i) implies that the set of mixed strategies (which, it bears emphasis, is infinite dimensional here) satisfies the domain hypotheses

³⁰ A topological vector space X is Hausdorff if for any $x \in X$ and $y \in X \setminus \{x\}$, there are neighborhoods N_x of x and N_y of y such that $N_x \cap N_y = \emptyset$. The space is locally convex if every neighborhood of zero includes a convex neighborhood of zero. In particular, every normed vector space is a locally convex Hausdorff space.

³¹ Actually, the set of these probability measures can be normed so that it is a (non-empty, compact, and convex) normed vector space.

of Fan-Glicksberg, while point (ii) gives continuity of payoffs in mixed strategies, which, along with the linearity of expected utility, implies that the best response correspondences satisfy the self-map's hypotheses in Fan-Glicksberg.

To sum up, if we allow for mixed strategies, we can always find Nash equilibria in finite games. In infinite games, Nash equilibria need not exist (construct an example!), but [Theorem 3](#) and [Theorem 5](#) gives sufficient conditions to assure existence. [Example 11](#) shows that these conditions are certainly not necessary.³²

3.3.3. Finding Mixed Strategy Equilibria

As noted earlier, the MSNE we computed for MP-A had the property that each player is indifferent, in equilibrium, between the pure strategies that he is randomizing over. The next result show that this is a general property.

Proposition 1 (MSNE Indifference Condition). *Fix a strategy profile, σ^* . Define $S_i^* \equiv \{s_i \in S_i \mid \sigma_i^*(s_i) > 0\}$ as the set of pure strategies that player i plays with positive probability according to σ^* . Then, σ^* is a Nash equilibrium if and only if*

1. $u_i(s_i, \sigma_{-i}^*) = u_i(s'_i, \sigma_{-i}^*)$ for all $s_i, s'_i \in S_i^*$;
2. $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*)$ for all $s_i \in S_i^*$ and $s'_i \in S_i$.

Proof. (Necessity.) If either condition fails, there are strategies $s'_i \in S_i$ and $s_i \in S_i^*$ such that $u_i(s'_i, \sigma_{-i}^*) > u_i(s_i, \sigma_{-i}^*)$. Construct the mixed strategy σ_i by setting $\sigma_i(\tilde{s}_i) = \sigma_i^*(\tilde{s}_i)$ for all $\tilde{s}_i \notin \{s'_i, s_i\}$ and $\sigma_i(s'_i) = \sigma_i^*(s_i) + \sigma_i^*(s'_i)$. Clearly, $u_i(\sigma_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*)$; hence σ_i^* is not a best response to σ_{-i}^* , and consequently, σ^* is not a Nash equilibrium.

(Sufficiency.) If both conditions hold, there is no strategy for player i that does strictly better against σ_{-i}^* than σ_i^* ; hence σ^* is Nash. \square

Corollary 1. *No strictly dominated strategy can be played with positive probability in a Nash equilibrium.*

Proof. Let σ^* be a Nash equilibrium in which $\sigma_i^*(s_i) > 0$ for some s_i that is strictly dominated. By strict domination of s_i , there exists a s'_i such that $u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$. (Why, even though the strict domination need not be by a pure strategy?) This contradicts condition (2) of [Proposition 1](#). \square

The necessity part of [Proposition 1](#) says that in any Nash equilibrium, a player must be indifferent over the pure strategies she is randomizing over.³³ This places a restriction on the mixed strategy of *her* opponents. That is, in general, player i will *not* be indifferent between playing s_i and $s'_i (\neq s_i)$, unless σ_{-i} is carefully chosen. Let's see how this operates.

Example 12 (Mixed Strategies in MP-A). Consider MP-A from [Example 1](#). We can solve for *all* mixed strategy equilibria as follows. At least one player must be non-degenerately randomizing over H and T .

³² See, for example, [Athey \(2001\)](#), [Jackson, Simon, Swinkels, and Zame \(2002\)](#), and [Reny \(1999\)](#) for even more on existence theorems.

³³ In fact, the logic establishes something stronger: if σ_i is a best response to some σ_{-i} (not necessarily part of a Nash equilibrium), then each $s_i \in \text{support}[\sigma_i]$ is a best response to σ_{-i} .

Wlog, suppose it is player 1 (Anne). For her to be indifferent between the two actions means that σ_2 (Bob's mixed strategy) must be such that $u_1(H, \sigma_2) = u_1(T, \sigma_2)$. Observe that

$$\begin{aligned} u_1(H, \sigma_2) &= 2\sigma_2(H), \\ u_1(T, \sigma_2) &= 2\sigma_2(T). \end{aligned}$$

Indifference thus requires that $\sigma_2(H) = \sigma_2(T) = \frac{1}{2}$. That is, Bob must randomize uniformly over his pure strategies for Anne to be willing to mix in equilibrium. By a symmetric argument, Anne must also be mixing uniformly (for Bob to be willing to mix at all, and in particular uniformly). Hence, the unique Nash equilibrium in MP-A is both players randomizing equally over H and T . \square

The next example applies the same idea to an important economic application, also showing that it works for infinite action spaces.³⁴

Example 13 (Common Value All-Pay Auction). There are $I > 1$ bidders for an object, each of whom values the object at $v > 0$. They all simultaneously submit bids, $s_i \geq 0$. The object goes to the highest bidder (randomly chosen amongst highest-bidders if there is a tie); everyone pays their bid to the auctioneer regardless of whether they win or not. Hence, payoffs for bidder i are $v - s_i$ if he wins, and $-s_i$ if he does not.

You can (and should!) verify that there are no PSNE in this game. To find a MSNE, we look for a *symmetric* MSNE, i.e. one where all players use the same mixed strategy. Let $F(x)$ denote the cdf over bids of a single player that is induced by the strategy; assume it has no atoms. For a player to be indifferent over all the bids that he is mixing over, it must be that for all $x \in \text{Supp}(F)$,

$$[F(x)]^{I-1}v - x = \bar{u}$$

for some constant utility level \bar{u} . Rearranging gives

$$F(x) = \left(\frac{\bar{u} + x}{v} \right)^{\frac{1}{I-1}}.$$

We know that $F(v) = 1$ (i.e., a player never bids more than his value, v) because any bid strictly above v is strictly dominated. Plugging in $F(v) = 1$ above yields $\bar{u} = 0$, and hence we get the solution

$$F(x) = \left(\frac{x}{v} \right)^{\frac{1}{I-1}}.$$

To complete the argument, we must show that no $x \notin \text{Supp}(F)$ yields a strictly higher utility than $\bar{u} = 0$, but this is immediate since $\text{supp}(F) = [0, v]$ and bidding any $x > v$ yields an expected payoff of $v - x < 0$ (because it wins for sure).

Notice that in the equilibrium, each player has an expected payoff of $\bar{u} = 0$ — competition amongst the buyers has left them with 0 surplus from the auction. You can check that the expected payment of any bidder is v/I , so that the auctioneer's revenue in expectation is v . \square

³⁴ Actually, if you were paying attention, the indifference property already came up in an infinite action space: the asymmetric Bertrand game of [Example 11](#).

Remark 16. All-pay auctions are useful models in various contexts. For example, in Industrial Organization and related fields, one can think of R&D as being an all-pay auction. That is, there are many firms competing against each other to develop a new product. Each firm independently and simultaneously decides how much money to sink into R&D. The “winner” is the one who invests the most money, but all players bear the R&D costs regardless of whether they win or not. All-pay auctions are also widely used in political economy to model lobbying or political campaigns.

3.3.4. Interpreting Mixed Strategy Equilibria

When we introduced mixed strategies in [Subsection 2.2.3](#), I suggested that the easiest way to think about them was as though players were rolling dice to determine which pure strategy to use. While this is pedagogically true, some people find it a little uncomfortable to think that agents are actually choosing their (pure) strategies through an act of explicit randomization. This may be particularly disconcerting given that we have already seen that in any Nash equilibrium, each player is *indifferent* over the set of pure strategies that he is mixing over! So why would a player then randomize, instead of just picking one with certainty? Of course, if he does pick one with certainty, this would in general destroy the indifference of the *other* players over the strategies they are randomizing over, and break the equilibrium altogether.

One response is to say that the player is indifferent, hence is happy to randomize. But this won’t sit well with you if you don’t like the idea that players randomize in practice. Fortunately, it turns out that we do not need players to be actually randomizing in a MSNE. All that matters is that as far as *other* players are concerned, player *i*’s choice *seem* like a randomized choice. That is, what matters is the uncertainty that other players have about *i*’s strategy. To give an example, consider an NBA basketball game where team *A* has possession, is down by 2 points, and there is only time for one more play. Team *A* is in a time-out, and has to decide whether to go for a 2-pointer to tie the game, or a 3-pointer to win.³⁵ Team *B* is obviously deciding whether to focus its defense against a 2-point shot or a 3-point shot. This is basically a generalized game of Matching Pennies, version A: team *A* wants to mismatch; team *B* wants to match. It may be that Team *A*’s coach has a deterministic way of deciding whether to go for the win or the tie — for example, he uses his star shooter’s morning practice 3-point accuracy as the critical factor. So long as Team *B* did not observe the shooter’s morning practice accuracy, it is a randomized choice as far as they are concerned. Hence, *B*’s belief about *A*’s play is a non-degenerate one, even though *A* may actually be playing a pure strategy based on some private information not available to *B*. That is, when we talk about *A*’s mixed strategy, we are really talking about *B*’s beliefs about *A*’s strategy.

This way of justifying mixed strategy equilibria is known as *purification* ([Harsanyi, 1973](#)). We’ll come back to it somewhat more precisely a bit later when we have studied incomplete information.

3.4. Normal Form Refinements of Nash Equilibrium

Many games have lots of Nash equilibria, and we’ve seen examples already. It’s natural therefore to ask whether there are systematic ways in which we can refine our predictions within the set of Nash equilibria. The idea we pursue here is related to weak dominance.

³⁵ Incidentally, NBA wisdom has it that the “road” team should go (more often) for a win, whereas a “home” team should go (more often) for the tie.

Example 14 (Voting Game). Suppose there are an odd number, $I > 2$, members of a committee, each of whom must simultaneously vote for one of two alternatives: Q (for status quo) or A (for alternative). The result of the vote is determined by majority rule. Every member strictly prefers the alternative passing over the status quo.

There are many Nash equilibria in this game. Probably the most implausible is this: every member plays $s_i = Q$; and this results in the status quo remaining. There is a more natural PSNE: every member plays $s_i = A$; and this results in the alternative passing. \square

Why is it Nash for everyone to vote Q in this game? Precisely because if all other players do so, then no individual player's vote can change the outcome. That is, no player is *pivotal*. However, it is reasonable to think that a player would vote conditioning on the event that he *is* pivotal. In such an event, he should vote A . One way to say this formally is that $s_i = Q$ is weakly dominated by $s_i = A$ (recall [Definition 15](#)). We suggested earlier that players should not play weakly dominated strategies if they believe that there is the slightest possibility that opponents will play a strategy profile for which the weakly dominated strategy is not a best response. One way to justify this belief is that a player assumes that even though his opponents might intend to play their Nash equilibrium strategies, they might make a mistake in executing them. This motivates the notion of *trembling-hand perfect Nash equilibrium*: Nash equilibria that are robust to a small possibility that players may make mistakes.

Given pure strategy spaces, $\{S_i\}_i$, and a function $\varepsilon : \bigcup S_i \rightarrow (0, 1)$, define

$$\Delta_\varepsilon(S_i) = \{\sigma_i \in \Delta(S_i) \mid \sigma_i(s_i) \geq \varepsilon(s_i) \text{ for all } s_i \in S_i\}$$

as the space of ε -constrained mixed strategies for player i . This is the set of mixed strategies for i that place at least $\varepsilon(s_i) > 0$ probability on each of pure strategy s_i . If you recall our initial discussion of randomization, every such strategy is a *fully mixed strategy* (see [Definition 9](#)). The idea here is that the non-zero probabilities on each pure strategy capture the notion of “unavoidable mistakes”. We can now define an ε -constrained equilibrium as a Nash equilibrium in which players play ε -constrained mixed strategies.

Definition 21. An ε -constrained equilibrium of a normal form game, $\Gamma_N \equiv \{I, \{S_i\}_{i=1}^I, \{u_i\}_{i=1}^I\}$, is a pure strategy Nash equilibrium of the perturbed game, $\Gamma_\varepsilon \equiv \{I, \{\Delta_\varepsilon(S_i)\}_{i=1}^I, \{u_i\}_{i=1}^I\}$.

Note that in the original game, $u_i : S \rightarrow \mathbb{R}$ (with mixed strategy profiles evaluated according to expected utility), while in the perturbed game we have extended to $u_i : \Delta_\varepsilon(S_1) \times \cdots \times \Delta_\varepsilon(S_I) \rightarrow \mathbb{R}$ using expected utility (with mixed strategy profiles—which are now compound lotteries over the original pure strategy profiles—again evaluated using expected utility).

A trembling-hand perfect equilibrium is any limit of a sequence of ε -constrained equilibria.

Definition 22 (Trembling-Hand Perfect Equilibrium). A Nash equilibrium, σ^* , is a trembling-hand perfect equilibrium (THPE) if there is a sequence, $\{\varepsilon^k\}_{k=1}^\infty$, with each $\varepsilon^k : \bigcup S_i \rightarrow (0, 1)$, such that $\varepsilon^k(s_i) \rightarrow 0$ for all s_i as $k \rightarrow \infty$, and an associated sequence of ε^k -constrained equilibria, $\{\sigma^k\}_{k=1}^\infty$, such that $\sigma^k \rightarrow \sigma^*$ as $k \rightarrow \infty$.

[Selten \(1975\)](#) proved the following existence theorem, paralleling that of Nash.

Theorem 6 (THPE Existence). *Every finite game has a THPE.*

Proof Sketch. Given a finite game, Γ_N , for any $\varepsilon : \bigcup S_i \rightarrow (0, 1)$, we can define the perturbed game Γ_ε as above. It is straightforward to verify that Γ_ε satisfies all the assumptions of [Theorem 3](#) when each $\varepsilon(s_i)$ is small enough,³⁶ and hence has a PSNE. Consider a sequence of perturbed games as $\varepsilon(s_i) \rightarrow 0$ for all s_i , and any sequence of equilibria. By the compactness of $\Delta(S)$, the equilibrium sequence has a convergent subsequence. Since $u_i : \Delta(S) \rightarrow \mathbb{R}$ is continuous, this subsequence converges to a Nash equilibrium of the original game. \square

Notice that the definition of THPE only requires that we be able to find *some* sequence of ε -constrained equilibria that converges to a THPE. A stronger requirement would be that *every* sequence of ε -constrained equilibria converge to it. Unfortunately, this requirement would be too strong in the sense that many (finite) games do not have equilibria that satisfy this property.³⁷

The definition of THPE given above using ε -constrained equilibria is conceptually elegant because it nicely captures the idea of mistakes that we started out with. However, it can be difficult to work with in practice. A useful result is the following.

Proposition 2. *A Nash equilibrium, σ^* , is a THPE if and only if there exists a sequence of fully mixed strategy profiles, $\sigma^k \rightarrow \sigma^*$, such that for all i and k , σ_i^k is a best response to σ_{-i}^k .*

Proof. (Only If.) Let σ^* be a THPE. Take the sequence of ε^k -constrained equilibria, $\sigma^k \rightarrow \sigma^*$, from the definition of THPE. We claim that for all i and s_i , if $\sigma_i^*(s_i) > 0$ then s_i is a best response to σ_{-i}^k for all k large enough (and hence σ_i^k is a best response to σ_{-i}^k for all k large enough). Suppose not. Then there is some player i and a sub-sequence of $\{\sigma^k\}$, call it $\{\sigma^n\}$, such that for all n , there is s_i^n with $u_i(s_i^n, \sigma_{-i}^n) > u_i(s_i, \sigma_{-i}^n)$. But then, as σ^n is an ε^n -constrained equilibrium, it must hold that $\sigma_i^n(s_i) \rightarrow 0 < \sigma_i^*(s_i)$, contradicting $\sigma_i^n \rightarrow \sigma_i^*$.

(If.) Fix $\sigma^k \rightarrow \sigma^*$ such that each σ_i^k is fully mixed and each σ_i^k is a best response to each σ_{-i}^k . Define $\varepsilon^k(s_i) \equiv \sigma_i^k(s_i)$ if $\sigma_i^*(s_i) = 0$ and $\varepsilon^k(s_i) \equiv 1/k$ if $\sigma_i^*(s_i) > 0$. Note that $\varepsilon^k(s_i) \rightarrow 0$ for all s_i as $k \rightarrow \infty$. We claim that σ^k is an ε^k -constrained equilibrium for all large enough k . This follows from the observation that by construction, if $\sigma_i^k(s_i) > \varepsilon^k(s_i)$ then $\sigma_i(s_i) > 0$, which, using the hypothesis that σ_i is a best response to σ_{-i}^k , implies that s_i is a best response to σ_{-i}^k (recall [Proposition 1](#)). In words: each σ_i^k only puts greater than minimum probability (minimum according to ε^k) on pure strategies that are best responses to σ_{-i}^k , and hence σ_i^k is a best response to σ_{-i}^k in the ε^k -perturbed game. \square

The Proposition says that in a THPE, σ^* , each σ_i^* is a best response to not only σ_{-i}^* (as required by Nash), but moreover, to each element of some sequence of fully mixed strategy profiles, $\sigma_{-i}^k \rightarrow \sigma_{-i}^*$. [Proposition 2](#) immediately implies:

Corollary 2. *If σ^* is a Nash equilibrium in which σ_i^* is fully mixed for each player i , then σ^* is a THPE.*

Furthermore, we can use the Proposition to show that no weakly dominated strategy can be part of a THPE:

³⁶ Without the qualification we could have $\Delta_\varepsilon(S_i) = \emptyset$.

³⁷ This is tied to the fact that there are games with multiple THPE, which cannot be justified by the same sequence of ε -constrained equilibria. One could alternatively consider a set-valued concept of equilibria, where rather than referring to a particular strategy profile as an equilibrium, one instead targets a collection of strategy profiles that jointly satisfy certain desiderata. For instance, can we find a (non-trivial) *minimal* collection of Nash equilibria with the property that every sequence of ε -constrained equilibria converges to some Nash equilibrium in the set? The set may not be a singleton, but perhaps it is still typically “small”. For a solution concept along these lines, see the notion of *stability* in [Kohlberg and Mertens \(1986\)](#).

Proposition 3 (THPE and Weak Dominance). *If σ^* is a THPE, then for all i : (1) σ_i^* is not weakly dominated; (2) for all $s_i \in S_i$ such that $\sigma_i^*(s_i) > 0$, s_i is not weakly dominated.*³⁸

Proof. The first claim is an immediate implication of [Proposition 2](#). For the second claim, take a profile σ^* such that for some i and $s_i \in S_i$ that is weakly dominated, $\sigma_i^*(s_i) > 0$. Then, by the definition of weak dominance, there exists $s'_{-i} \in S_{-i}$ and $\sigma_i \in \Sigma_i$ such that $u_i(\sigma_i, s'_{-i}) > u_i(s_i, s'_{-i})$, and for all $s_{-i} \in S_{-i}$, $u_i(\sigma_i, s_{-i}) \geq u_i(s_i, s_{-i})$. Now consider a strategy $\tilde{\sigma}_i$ constructed as follows: for any $\tilde{s}_i \in S_i \setminus \{s_i\}$, $\tilde{\sigma}_i(\tilde{s}_i) = \sigma_i^*(\tilde{s}_i) + \sigma_i^*(s_i)\sigma_i(\tilde{s}_i)$; and $\tilde{\sigma}_i(s_i) = \sigma_i^*(s_i)\sigma_i(s_i)$. (Exercise: verify that $\tilde{\sigma}_i$ defines a valid strategy.) It follows that for any σ_{-i} in which $\sigma_{-i}(s'_{-i}) > 0$, $u_i(\tilde{\sigma}_i, \sigma_{-i}) > u_i(\sigma_i^*, \sigma_{-i})$. Hence, σ_i^* is not a best response to any fully mixed strategy profile, σ_{-i} . By [Proposition 2](#), σ^* is not a THPE. \square

Combined with the existence result for THPE, the fact that THPE rules out weakly dominated strategies makes it an appealing refinement of Nash equilibrium. You might wonder if we can say more: is any Nash equilibrium in which players use strategies that are not weakly dominated a THPE? Unfortunately, this is not true in general.³⁹

Example 15. Consider the following three player game:

	L	R		L	R
T	1,1,1	1,0,1	T	1,1,0	0,0,0
B	1,1,1	0,0,1	B	0,1,0	1,0,0
	l			r	

You can verify that (B, L, l) is a Nash equilibrium where no strategy is weakly dominated. However, it is not a THPE, because for any fully mixed strategies for players 2 and 3 that assign sufficiently low probabilities to R and r respectively, player 1 strictly prefers to play T rather than B due to (L, r) being an order of magnitude more likely to occur than (R, r) . As an exercise, check that (T, L, l) is the *only* THPE in this game (hint: think first about the set of Nash equilibria). \square

The result is true in the case of two-player games, however.

Proposition 4. *In a two-player game, if σ^* is a Nash equilibrium where both σ_1^* and σ_2^* are not weakly dominated, then σ^* is a THPE.*

Proof Sketch. Assume that σ^* is a Nash equilibrium where both σ_1^* and σ_2^* are not weakly dominated. It can be shown that for each player, i , there exists some fully mixed strategy of the opponent, σ_j ($j \neq i$), such that σ_i^* is a best response to σ_j .⁴⁰ For any positive integer n , define for each player the mixed strategy $\sigma_i^n = \frac{1}{n}\sigma_i + (1 - \frac{1}{n})\sigma_i^*$. The sequence of fully mixed strategies profiles (σ_1^n, σ_2^n) converges to σ^* , and σ_i^* is a best response to each σ_j^n . By [Proposition 2](#), σ^* is a THPE. \square

³⁸ Strictly speaking, it is unnecessary to state the two parts separately, because the first implies the second. You can prove this along the lines of the proof below.

³⁹ As the example below and the logic of the proof of the following Proposition demonstrate, this is due to the independence we assume in the players' randomization when playing mixed strategies. If we allowed for correlation, it would indeed be true.

⁴⁰ See [Osborne and Rubinstein \(1994, p. 64, Exercise 64.2\)](#).

(Do you see where in the above proof we used the assumption of two players? Hint: think about where it fails to apply to [Example 15](#).)

Finally, we show that THPE is consistent with a particular class of reasonable Nash equilibria.

Definition 23 (Strict Nash Equilibrium). A strategy profile, s^* , is a strict Nash equilibrium if for all i and $s_i \neq s_i^*$, $u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_i^*)$.

Remark 17. There is no loss of generality in only considering PSNE in the definition of strict Nash equilibrium, since in any MSNE, a player is indifferent over the pure strategies he is mixing over.

In a strict Nash equilibrium, each player has a *unique* best response to his opponents' strategy profile. This implies that every strict Nash equilibrium is trembling-hand perfect.

Proposition 5 (Strict NE \implies THPE). *Every strict Nash equilibrium is a THPE.*

Proof. Exercise. □

Remark 18. The converse is certainly not true: not every THPE is a strict Nash equilibrium. (We've already seen multiple counter-examples.)

Sometimes a Nash equilibrium that is not strict is called a *weak* Nash equilibrium, but this is not very common terminology. Keep also in mind that there is a distinct concept from strict Nash called *strong* Nash equilibrium, which has to do with (lack of) “coalitional deviations”—outside our scope.

3.5. Correlated Equilibrium

As a final section on static games of complete information, we are going to briefly discuss the concept of *correlated equilibrium*. At a few different points earlier, I mentioned that requiring players to randomize independently when playing mixed strategies can be restrictive. Relaxing this is important for a fuller treatment for at least two reasons: (1) certain results depend on it (such as the equivalence of iterative deletion of strictly dominated strategies and rationalizability with more than 2 players); (2) it can be practically relevant in thinking about behavior in some strategic situations. To illustrate the latter point, we start with an example.

Example 16 (Battle of the Sexes). Spouses, Anne and Bob, must decide whether to go to the concert or the baseball game. Both want to coordinate with the other, but Anne prefers to coordinate on the concert whereas Bob prefers baseball. In normal form,

	b	c
B	2,6	0,0
C	0,0	6,2

You can verify that there are exactly three Nash equilibria to this game: (B, b) , (C, c) and a MSNE $(0.25B + 0.75C, 0.75b + 0.25c)$. The corresponding (expected) payoffs are: $(2, 6)$, $(6, 2)$, and $(1.5, 1.5)$. However, suppose that they can jointly observe a coin toss (or whether it is raining or sunny outside, or any other publicly observable random variable) before acting. Then they can attain a new payoff outcome, one that is more equitable than either PSNE, and Pareto-dominates the MSNE. For example, they toss a fair

coin, and if the outcome is Heads, they both play B ; if the outcome is Tails, they both play C .⁴¹ Clearly, given that the opponent is following the prescribed strategy, it is optimal for each player to follow it. In expectation, this achieves a convex combination of two PSNE payoffs, giving the payoff profile $(4, 4)$. More generally, by using an appropriately weighted randomization device, *any* convex combination of the the NE payoffs (or action-profiles) can be achieved. \square

The next example demonstrates that by using correlated but private signals, players may be able to do even better than by public randomization.

Example 17. Consider a modification of the Battles of the Sexes, as follows:⁴²

	b	c
B	5,1	0,0
C	4,4	1,5

Now, the three Nash equilibria are (B, b) , (C, c) and $(0.5B + 0.5C, 0.5b + 0.5c)$, with corresponding payoffs $(5, 1)$, $(1, 5)$, and $(2.5, 2.5)$. By using public randomization as before, any convex combination of these can be attained. But the players can do even more by using a device that sends each player correlated but privately observed signals. For example, suppose they hire an independent third party who tosses a three-sided fair die and acts as follows: she reveals whether the roll is 1 or in the set $\{2, 3\}$ to Anne; but to Bob, she reveals whether the roll is 3 or in the set $\{1, 2\}$. Consider a strategy for Anne where she plays B if told 1, and C if told $\{2, 3\}$; a strategy for Bob where he plays c if told 3, and b if told $\{1, 2\}$. Let us check that it is optimal for Anne to follow her strategy, given that Bob is following his: (i) when Anne is told 1, she knows that Bob will play b (since he will be told $\{1, 2\}$), hence it is optimal for her to play B ; (ii) when Anne is told $\{2, 3\}$, she knows that with probability 0.5, Bob was told $\{1, 2\}$ and will play b , and with probability 0.5, Bob was told 3 and will play c , hence she is indifferent between her two actions. A similar analysis shows that it is optimal for Bob to follow his strategy given that Anne is following hers. Hence, the prescribed behavior is self-enforcing, and attains a payoff of $(3\frac{1}{3}, 3\frac{1}{3})$, which is outside the convex hull of the original Nash equilibrium payoffs. \square

Generalizing from the examples, we now define a correlated equilibrium.

Definition 24 (Correlated Equilibrium). A probability distribution p on the product pure strategy space, $S = S_1 \times \cdots \times S_I$, is a correlated equilibrium if for all i and all s_i chosen with positive probability under p ,

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i} | s_i) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} p(s_{-i} | s_i) u_i(s'_i, s_{-i}) \quad \forall s'_i \in S_i.$$

The way to think of this definition of correlated equilibrium is that everyone knows ex-ante that a “device” (or outside party) will choose the pure strategy profile s with probability $p(s)$, but each player only learns his component of the profile that was selected, s_i . We have a correlated equilibrium if all players want to follow their recommendation to play s_i , given that all other players are following their recommendations.

⁴¹ Indeed, your own personal experience in such situations may suggest that this is precisely how couples operate!

⁴² Unintendedly, I wrote this as Anne (row player) preferring baseball whereas Bob (column player) prefers the concert.

Observe that the key difference with the definition of Nash equilibrium is that the distribution p may have correlation between s_i and s_{-i} , and accordingly the optimality conditions account for the conditional probability of s_{-i} given s_i . This suggests that the set of correlated equilibria generalizes the set of Nash equilibria.

Proposition 6 (NE \implies Correlated Equilibrium). *Every Nash equilibrium is a correlated equilibrium.*

Proof. Given a Nash equilibrium, σ , simply define $p(s) = \prod_{i=1}^I \sigma_i(s_i)$. In this case, $p(s_{-i}|s_i) = p(s_{-i})$, and the optimality conditions for correlated equilibrium reduce to that of Nash. \square

Corollary 3. *Every finite game has a correlated equilibrium.*

Correlated equilibrium is a useful solution concept when thinking about pre-play communication (possibly through a mediator) and other contexts. However, we will stick to Nash equilibrium for the rest of this course.

3.6. Bayesian Nash Equilibrium and Incomplete Information

The assumption heretofore maintained that all players know each other's preferences over terminal nodes (or strategy profiles) is clearly a very restrictive one. In fact, it is reasonable to think that in many games, one doesn't really know what the opponents' payoffs are. For example, in the Bertrand or Cournot competition games, each firm may not know the other firm's cost function or cost parameter in the linear case. Alternatively, in the auction examples, players won't generally know their opponents' valuations for the object. Moreover, one may not even know one's own payoff from some strategy profiles, since the payoff could depend upon something that is only known to another player (Example 18 below illustrates).

Definition 25 (Incomplete Information Game). A game has (or is of) incomplete information when at least one player does not know the payoff that some player receives from some strategy profile (or terminal node).

Dealing with incomplete information would seem to require consideration of a player's beliefs about other players' preferences, her beliefs about their beliefs about her preferences, and so on, ad infinitum. This can quickly get very complicated. Fortunately, we have a now standard and beautiful way to approach this problem, due to Harsanyi. His idea is to *transform any game of incomplete information into a game of complete but imperfect information* as follows: we imagine that a player's payoffs are determined at the outset of the game by Nature, which chooses realizations of a random variable for each player. The vector of realizations of random variables determines the payoffs for all players from each strategy profile. A player observes the realization of his own random variable, but not necessarily that of others. As with other moves of Nature, the probability distribution that Nature uses for each random variable is common knowledge. This extended game is known as a *Bayesian game (of incomplete information)*, and we call a Nash equilibrium of this extended game a *Bayesian Nash equilibrium*. The realization of a player's random variable is often called his *type*.

Example 18 (For Love or Money). Suppose a wealthy Anne is considering whether to marry a pauper Bob, but is not completely sure whether he loves her (probability α), or just wants her money (probability $1 - \alpha$). This is a game where each player must choose to marry or not. If either one chooses not to marry, then both players get a payoff of 0. If both choose to marry, then Bob gets a payoff of 5 if he is a lover, and 3 if

he is a scoundrel; whereas Anne gets 5 if Bob is a lover and -3 if he is a scoundrel. Clearly, if it were known that Bob is a scoundrel, then every Nash equilibrium involves Anne choosing not to marry. Conversely, if Bob were known to be a lover, then there is a Nash equilibrium where both Anne and Bob choose marriage.

The question is, what happens if Bob privately knows his *type* (lover or scoundrel), but Anne does not? We can represent this game of incomplete information in extensive form using Nature’s move at the root node. In this extended game, a pure strategy for Anne is simply a choice of marry or not marry; so she has 2 pure strategies. However, for Bob, a pure strategy is a *contingent* plan that specifies whether to marry or not in each of two cases: if his type is lover, and if his type is scoundrel. Hence, Bob has 4 pure strategies in the extended game: (marry if lover, marry if scoundrel), (marry if lover, don’t if scoundrel), (don’t if lover, marry if scoundrel), and (don’t if lover, don’t if scoundrel). \square

We now define a Bayesian game formally.

Definition 26 (Bayesian Game). A Bayesian game is defined by a tuple $\langle I, \{A_i\}, \{u_i\}, \{\Theta_i\}, p \rangle$, where I is the set of players, A_i is the action space for player i , Θ_i is the set of types for player i , $u_i : A \times \Theta \rightarrow \mathbb{R}$ is the payoff function for player i , and $p : \Theta \rightarrow [0, 1]$ is the prior probability distribution over type profiles.⁴³ p is assumed to be common knowledge, but each player i only knows her own θ_i .

Note that in the Bayesian game, players privately learn their types and then choose their actions simultaneously. You should think of the action space for each player as possibly representing normal-form strategies of some underlying game. I use the terminology “action” rather than “strategy” because in the context of a Bayesian game, there is something else we will want to call a strategy.

Remark 19. We allow each utility function, u_i , to depend on the entire vector of types — not just player i ’s type. This is more general than MWG (p. 255), and is useful in many applications. We say that the game has *private values* when each u_i only depends on i ’s type. The more general situation has *interdependent values*, sometimes also referred to as *correlated values*.

If the prior distribution p is a product distribution, so that for each i the conditional distributional $p(\theta_{-i}|\theta_i)$ is independent of θ_i , we say the game has *independent types*.

Remark 20. The above formulation of a Bayesian game is *ex-ante*: there is a stage at which players have identical information/no private information, and they subsequently learn their private information. But in many applications—including [Example 18](#)—it is more natural to take an *interim* perspective: players simply begin with private information, and there is no literal ex-ante stage. From this perspective, we would define a Bayesian game as a tuple $\langle I, \{A_i\}, \{u_i\}, \{\Theta_i\}, \{q_i\} \rangle$, where the objects are as in [Definition 26](#), except that each $q_i : \Theta_i \rightarrow \Delta(\Theta_{-i})$. Here $q_i(\theta_i)$ captures type θ_i ’s (subjective) belief about the opponents’ type profile. Up to technical details, which don’t arise when Θ is finite, it is straightforward that for each i we can find a “prior” $p_i \in \Delta(\Theta)$ such that q_i obtains from p_i via Bayesian updating. For example, set $p_i(\theta) = (1/|\Theta_i|)q_i(\theta_{-i}; \theta_i)$; although there are infinitely many others, as we could use any full-support marginal over Θ_i instead of the uniform. But there is no guarantee that we can find single prior $p \in \Delta(\Theta)$ that is simultaneously valid for all players. That requires some degree of consistency in the subjective beliefs $\{q_i\}$, which is referred to as the *common prior assumption* (CPA). Under the CPA, the ex-ante interpretation is justified in an “as if” sense even if there is no real ex-ante stage. The CPA is typically assumed in applications and we will maintain it here—although I want to stress that there is no logical reason it must be required. One can study Bayesian

⁴³ Here, $A \equiv A_1 \times \cdots \times A_I$ and $\Theta \equiv \Theta_1 \times \cdots \times \Theta_I$.

games without the CPA, contrary to some misconceptions. Indeed the CPA can be controversial; see [Gul \(1998\)](#).

Remark 21. It is far from obvious that [Definition 26](#) of a Bayesian game is sufficient to represent all games of incomplete information. In particular, one key issue is, following the discussion after [Definition 25](#), whether Bayesian games (as defined above) are sufficiently rich to capture an infinite hierarchy of beliefs of every player about other players' beliefs. This was resolved by [Mertens and Zamir \(1985\)](#) and [Brandenburger and Dekel \(1993\)](#) with constructions of the so-called *universal type space*. This is material suited for a more advanced course, but see [Zamir \(2008\)](#) for a nice discussion. In a nutshell, they established that Bayesian games are indeed sufficient under some reasonable “coherence” conditions on belief hierarchies. Caveat: whether the CPA can be satisfied is another question.

A pure strategy for player i in a Bayesian game specifies what action to take for each of her possible type realizations. That is, a (pure) strategy is a mapping $s_i : \Theta_i \rightarrow A_i$. Given a profile of strategies for all players, we compute i 's expected payoffs as

$$\tilde{u}_i(s_1, \dots, s_I) = \mathbb{E}_\theta[u_i(s_1(\theta_1), \dots, s_I(\theta_I), \theta_i, \theta_{-i})].$$

With these preliminaries, we can define a (pure strategy) Bayesian Nash equilibrium in the natural way.

Definition 27 (Bayesian Nash Equilibrium). A (pure strategy) Bayesian Nash equilibrium of the Bayesian game, $\langle I, \{A_i\}, \{u_i\}, \{\Theta_i\}, p \rangle$, is a profile of (pure) strategies (s_1, \dots, s_I) such that for all i and all $s'_i \in S_i$, $\tilde{u}_i(s_i, s_{-i}(\cdot)) \geq \tilde{u}_i(s'_i, s_{-i})$.

In the natural way, we can allow for *mixed* strategies, which are just probability distributions over the pure strategies, and then define mixed strategy Bayesian Nash equilibrium. An important observation is that by the Nash existence theorem, a (possibly mixed) Bayesian Nash equilibrium (BNE) exists if for all i , A_i and Θ_i are finite.

The approach taken above is *ex ante* in the sense that players are choosing their strategies prior to knowing their types. On the other hand, one can imagine that a player picks an action once he knows his type, but not that of others; this is at the interim stage.⁴⁴ It is straightforward that a strategy can be part of a BNE if and only if it maximizes a player's expected utility conditional on each θ_i that occurs with positive probability.⁴⁵

Proposition 7. A profile of strategies, s , is a (pure strategy) BNE if and only if, for all i and all $\theta_i \in \Theta_i$ occurring with positive probability and all $a_i \in A_i$,

$$\mathbb{E}_{\theta_{-i}}[u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) \mid \theta_i] \geq \mathbb{E}_{\theta_{-i}}[u_i(a_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) \mid \theta_i].$$

Proof. Exercise (or see MWG p. 256). □

The Proposition implies that we can essentially think of each type of a player as maximizing its own payoffs independent of what other types of the player are doing. This is a very useful result since it permits

⁴⁴ *Ex post* would refer to a third stage where the types of all players become known, for example at the end of the game when payoffs are realized.

⁴⁵ This is for finite Θ_i ; in the infinite case, it applies to “almost all” θ_i .

us to “decompose”, type by type, the process of finding an optimal response for a player to any strategy profile of his opponents.

Remark 22. While we are focussing on BNE, one can also parallel other complete-information solution concepts. For example, a strategy s_i is *strictly dominant* if for all $s'_i \neq s_i$ and s_{-i} , $\tilde{u}_i(s_i, s_{-i}) > \tilde{u}_i(s'_i, s_{-i})$. In the private-value case, we can write this alternatively as for all θ_i (that occur with positive probability), $a_i \neq s_i(\theta_i)$, and a_{-i} , $u_i(s_i(\theta_i), a_{-i}, \theta_i) > u_i(a_i, a_{-i}, \theta_i)$.

In Bayesian games, a weakening of weak dominance is of interest: a strategy s_i is “very weakly dominant” if for all s'_i , and s_{-i} , $\tilde{u}_i(s_i, s_{-i}) \geq \tilde{u}_i(s'_i, s_{-i})$. (This terminology is not common. Another way to say it is that s_i is a best response to any strategy of the opponents. Do you see why it is weaker than weak dominance?) In the context of mechanism/market design, where the action space is frequently $A_i = \Theta_i$ and the strategy of interest is $s_i(\theta_i) = \theta_i$, this requirement is called *strategyproofness*.

Another concept of interest is that of *ex-post (Bayesian Nash) equilibrium*. A pure BNE s is an ex-post equilibrium if for all θ , i , and a_i , $u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta) \geq u_i(a_i, s_{-i}(\theta_{-i}), \theta)$. Note that we are using equilibrium information here; what we are dropping is information about other players’ types. Ex-post equilibrium is appealing because it means we can be agnostic about players’ beliefs about others’ information; it thus comes up in design problems.

3.6.1. Examples

Example 19 (BNE in For Love or Money). Continuing with [Example 18](#), let us now find the (pure strategy) BNE of this game. Denote Anne as player 1 and Bob as player 2; Bob’s types as l (lover) and c (scoundrel); and the actions are M (marry) and N (not marry). We could proceed in two ways, corresponding to the *ex-interim* and *ex-ante* definitions.

1. For the ex-interim procedure, we can treat a BNE as a triple $(s_1, s_2(l), s_2(c)) \in \{M, N\}^3$. Then (N, N, N) and (N, N, M) are always BNE. In addition, there is at least one other BNE, though what it is depends on the parameter α . If $\alpha \geq \frac{3}{8}$, (M, M, M) is a BNE. If $\alpha \leq \frac{3}{8}$, (N, M, M) is a BNE. Notice that (M, N, N) is never a BNE, and moreover, if $\alpha \in (0, 1)$, (M, M, N) , (M, N, M) , and (N, M, N) are not BNE.
2. For the ex-ante formulation, a BNE is a pair $(s_1, s_2) \in \{M, N\} \times \{MM, MN, NM, NN\}$, where for player 2, (MN) is the strategy of playing M if l and N if c , for example. You can write out the BNE in this formulation that are equivalent to the ex-interim formulation. \square

Example 20 (Incomplete Information 2nd Price Auction). Recall the 2nd price auction we studied in [Example 9](#). Now, each player’s valuation v_i is assumed to be *private* information. It is straightforward that it remains a BNE for each player to bid the truth, i.e. play $s_i(v_i) = v_i$. This follows from the fact that it was weakly dominant to bid the truth in the complete information case. \square

Example 21 (1st Price Auction). Now we consider a 1st price Auction, which is an auction where the highest bidder wins (bids are non-negative), the winner pays his own bid, and losers pay nothing. As before, if there is a tie, the winner is randomly selected from the highest bidders. Assume that each v_i is drawn independently from the distribution F that is continuous on \mathbb{R} and strictly increasing on the support $[\underline{v}, \bar{v}]$, where $\underline{v} \geq 0$. This is a setting of *independent private values*.

To find a BNE, we look for a symmetric pure-strategy equilibrium in which each player is playing the same strictly increasing and differentiable strategy, denoted $s^*(v)$. Suppose that all other bidders are playing $s^*(v)$. Player i 's expected payoff from a bid b when he has valuation v is⁴⁶

$$\pi(b, v) = [F((s^*)^{-1}(b))]^{I-1} [v - b].$$

The expected payoff in equilibrium for a bidder with valuation v is then

$$\Pi(v) = \pi(s^*(v), v) = [F(v)]^{I-1} [v - s^*(v)].$$

Note that $\Pi(\underline{v}) = 0$ because our hypothesis that $s^*(v)$ is strictly increasing (and that F is continuous) implies that a bidder with the lowest possible valuation wins with probability zero. Differentiating,

$$\Pi'(v) = \pi_1(s^*(v), v) \frac{ds^*}{dv}(v) + \pi_2(s^*(v), v) = \pi_2(s^*(v), v) = [F(v)]^{I-1},$$

where the second equality uses the envelope theorem, i.e., that $\pi_1(s^*(v), v) = 0$ because $s^*(v)$ is an optimal bid with valuation v .⁴⁷

Recall that by the Fundamental Theorem of Calculus,

$$\Pi(v) - \Pi(\underline{v}) = \int_{\underline{v}}^v \Pi'(x) dx.$$

Substituting in to the above from the previous derivations gives

$$[F(v)]^{I-1} [v - s^*(v)] = \int_{\underline{v}}^v [F(x)]^{I-1} dx,$$

which rearranges as

$$s^*(v) = v - \int_{\underline{v}}^v \left[\frac{F(x)}{F(v)} \right]^{I-1} dx.$$

It is clear that $s^*(v)$ is strictly increasing and differentiable, as required. Notice that $s^*(v) < v$ if $v > \underline{v}$, which makes sense: bidders should “shade down” their bids because winning is only beneficial if one pays less than one's valuation. (Indeed, verify that the strategy $s(v) = v$ is weakly dominated.)

To complete the argument, one must verify that $s^*(v)$ is in fact an optimal bid for type v —we only derived necessary conditions above—which I omit. \square

Remark 23. Above, we constructed a symmetric Bayes Nash equilibrium by assuming that bidders are using a differentiable pure strategy. In fact, one can prove that in *any* symmetric equilibrium, each bidder must use a continuous and strictly increasing (hence a.e. differentiable) pure strategy. This is proved by establishing that there cannot be a “gap” in the range of bids played in equilibrium (because then it would be suboptimal

⁴⁶ Below, let $(s^*)^{-1}(b) = \underline{v}$ for $b < s^*(\underline{v})$ and $(s^*)^{-1}(b) = \bar{v}$ for $b > \bar{v}$.

⁴⁷ You may worry that this is not assured for type \underline{v} , since $s^*(\underline{v})$ may be zero, and as such one is not assured the first order condition $\pi_1(s^*(\underline{v}), \underline{v}) = 0$. But this won't matter because it will be sufficient that the formula for $\Pi'(v)$ hold for $v > \underline{v}$.

to bid just above the gap; one should instead bid at or just above the bottom of the gap) and there cannot be any “atoms” in the equilibrium distribution of bids (because then no bidder would make an offer at or just below the atom as it is better to deviate to just above the atom, leading to a gap). Thus, the symmetric equilibrium is essentially unique.

Remark 24 (Revenue Equivalence). Let’s figure out the seller’s expected revenue in the 1st price auction if the symmetric equilibrium is played. It is simply the expected winning bid, or the expected bid of the bidder with the highest valuation. We may denote this $\mathbb{E}[s^*(V^{1:I})]$, where, for any integer $k \in [1, I]$, $V^{k:I}$ denotes the k -th highest valuation out of I bidders. It is common to refer to $V^{k:I}$ as the $(I + 1 - k)$ -th order statistic of F (with sample size I), so that $V^{1:I}$ is the I -th order statistic. To compute the expected revenue more explicitly, define $F^{1:I}(v) \equiv [F(v)]^I$, so that $F^{1:I}$ is the cdf of $V^{1:I}$. We have

$$\begin{aligned} s^*(v) &= v - \int_{\underline{v}}^v \left[\frac{F(x)}{F(v)} \right]^{I-1} dx \\ &= \frac{1}{F^{1:I-1}(v)} \left[F^{1:I-1}(v)v - \int_{\underline{v}}^v F^{1:I-1}(x) dx \right] \\ &= \frac{1}{F^{1:I-1}(v)} \int_{\underline{v}}^v x dF^{1:I-1}(x) \\ &= \mathbb{E}[V^{1:I-1} | V^{1:I-1} < v], \end{aligned} \tag{2}$$

where the third equality is derived through an integration by parts. This says that if a bidder has valuation v , he sets his bid equal to the expectation of the highest of the other $I - 1$ valuations, conditional on him having the highest overall valuation. So, the expected revenue for the seller is

$$\mathbb{E}[s^*(V^{1:I})] = \mathbb{E}[V^{1:I-1} | V^{1:I-1} < V^{1:I}] = \mathbb{E}[V^{2:I}], \tag{3}$$

where the first equality is from (2) and the second equality owes to iterated expectations.⁴⁸ But observe that $\mathbb{E}[V^{2:I}]$, which is the expectation of the 2nd highest valuation among I bidders, is also the expected revenue from a 2nd price auction! (Why?) We have thus proved that, given the relevant equilibria, the 1st price auction and the 2nd price auction generate exactly the same expected revenue for the seller, even though the payment rules and bidding behavior are very different.⁴⁹ This is a particular case of the *Revenue Equivalence Theorem*: the key points are that (i) the setting is one of independent private values (each bidder’s type

⁴⁸ More detail on the iterated expectations: for any vector of realized valuations, the second-highest valuation can also be viewed as the highest remaining valuation after excluding the highest valuation. Thus, $\mathbb{E}[V^{2:I}] = \mathbb{E}[\mathbb{E}[V^{1:I-1} | V^{1:I-1} < v]]$, where the outer expectation is with respect to $V^{1:I}$ having realization v . But this is just the middle expression of (3).

⁴⁹ Another way to see the revenue equivalence is to note that in the 2nd price auction, the expected payment for a bidder i with valuation v is

$$\begin{aligned} &\text{Prob}[i \text{ wins} | v_i = v] \times \mathbb{E}[\text{2nd highest bid} | v \text{ is the highest bid}] \\ &= \text{Prob}[i \text{ wins} | v_i = v] \times \mathbb{E}[\text{2nd highest value} | v \text{ is the highest value}] \\ &= F^{1:I}(v) \times \mathbb{E}[V^{1:I-1} | V^{1:I-1} < v], \end{aligned}$$

which, using our derivation (2), is the same as in the 1st price auction (because (2) gives the payment of bidder i with valuation v when he wins, which occurs with probability $F^{1:I}(v)$). If the expected payment is the same for every bidder (conditional on his type, and hence on expectation too), then the expected revenue is the same.

is his own valuation and types are drawn independently); (ii) both auction formats allocate the object to the bidder with the highest valuation; and (iii) a bidder with the lowest possible valuation (i.e., \underline{v}) gets 0 expected payoff in equilibrium. More generally, the issue of how to design an auction to maximize the seller's revenue lies in the field of mechanism design.

As a third example of Bayes Nash equilibrium, we look at Cournot competition game with incomplete information. It serves to illustrate a case where players' types are not independent.

Example 22 (Incomplete Information Cournot). The setting is the same as the Linear Cournot case considered in [Example 7](#), but modified so that each firm now has one of two potential cost parameters: c_H or c_L , where $c_H > c_L \geq 0$. Each firm's parameter is privately known to it alone, and the prior distribution is given by $p(c_H, c_H) = p(c_L, c_L) = \frac{1}{2}\alpha$ and $p(c_H, c_L) = p(c_L, c_H) = \frac{1}{2}(1 - \alpha)$, with $\alpha \in (0, 1)$ commonly known. Recall that inverse market demand is given by $p(Q) = a - bQ$; firm i 's cost function is $c_i q_i$, where $c_i \in \{c_H, c_L\}$.

Let's look for a symmetric pure-strategy BNE where each firm plays $s^*(c_H)$ and $s^*(c_L)$ for each of its two types. To solve for these, we observe that if a firm is of type c_H , then its maximization problem, taking as given that the other firm is using $s^*(\cdot)$, is:

$$\max_q [\alpha(a - b(s^*(c_H) + q) - c_H)q + (1 - \alpha)(a - b(s^*(c_L) + q) - c_H)q].$$

By hypothesis that $s^*(\cdot)$ is a symmetric equilibrium, a solution to the above problem must be $s^*(c_H)$. Taking the FOC (presuming an interior solution) thus gives an equilibrium condition:

$$\alpha(a - c_H - bs^*(c_H) - 2bs^*(c_H)) + (1 - \alpha)(a - c_H - bs^*(c_L) - 2bs^*(c_H)) = 0. \quad (4)$$

Similarly, the maximization problem when type is c_L is

$$\max_q [(1 - \alpha)(a - b(s^*(c_H) + q) - c_L)q + \alpha(a - b(s^*(c_L) + q) - c_L)q],$$

for which $s^*(c_L)$ must be a solution. This requires

$$(1 - \alpha)(a - c_L - bs^*(c_H) - 2bs^*(c_L)) + \alpha(a - c_L - bs^*(c_L) - 2bs^*(c_L)) = 0. \quad (5)$$

The equilibrium is found by solving Eqns. (4) and (5), which involves tedious but straightforward algebra. \square

3.6.2. Purification Theorem

In [Subsection 3.3.4](#), we introduced the notion of *purification* ([Harsanyi, 1973](#)) as a justification for MSNE. Using incomplete information, we can now state the idea more precisely. Start with a (finite) normal-form game, $\Gamma_N = \{I, \{S_i\}, \{u_i\}\}$. Let θ_i^s be a random variable with range $[-1, 1]$, and $\varepsilon > 0$ a constant. Player i 's perturbed payoff function depends on the collection $\theta_i \equiv \{\theta_i^s\}_{s \in S}$ and is defined as

$$\tilde{u}_i(s, \theta_i) = u_i(s) + \varepsilon \theta_i^s$$

We assume that the θ_i 's are independent across players, and each θ_i is drawn from a distribution P_i with density p_i .

Theorem 7 (Purification). *Fix a set of players, I , and strategy spaces $\{S_i\}$. For almost all vectors of payoffs, $u = (u_1, \dots, u_I)$,⁵⁰ for all independent and twice-differentiable densities, p_i on $[-1, 1]^{|S_i|}$, any MSNE of the payoffs $\{u_i\}$ is the limit as $\varepsilon \rightarrow 0$ of a sequence of pure-strategy BNE of the perturbed payoffs $\{\tilde{u}_i\}$. Moreover, the sequence of pure-strategy BNE are “essentially” strict.*

The Theorem, whose proof is quite complicated, is the precise justification for thinking of MSNE as PSNE of a “nearby” game with some added private information. Note that according to the Theorem, a single sequence of perturbed games can be used to purify all the MSNE of the underlying game. In his proof, Harsanyi showed that equilibria of the perturbed games exist (note that these are infinite games), and moreover, that in any equilibrium of a perturbed game, almost all $s_i(\theta_i)$ must be pure strategies of the underlying game, and that for almost all θ_i , $s_i(\theta_i)$ is the unique best response (this is the “essentially strict” portion of the result).

3.6.3. The Importance of Higher-Order Beliefs

In the examples/applications so far, higher order beliefs have not played a major role. But consider the following symmetric two-player game, where $x \in \{0, 2\}$:

	A	B
A	x, x	$-\varepsilon, 1$
B	$1, -\varepsilon$	$1, 1$

In what follows, $\varepsilon > 0$ is common knowledge; you can think of it as small. If it is commonly known that $x = 2$ then we have a coordination game: both (A, A) and (B, B) are strict Nash equilibria; (A, A) is Pareto dominant while (B, B) is risk dominant.⁵¹ If, on the other hand, $x = 0$ is common knowledge, then B is a strictly dominant strategy, so (B, B) is the unique Nash equilibrium.

Suppose there isn’t common knowledge about which game is being played, i.e., there is incomplete information about x . Assume the prior $\Pr(x = 2) > 1/2$ (and $\Pr(x = 0) = 1 - \Pr(x = 2)$). If players don’t receive any other information about x , then the corresponding Bayesian game has two Bayesian Nash equilibria: (A, A) and (B, B) ; both are strict and (A, A) remains Pareto dominant.

Let us add private information as follows. There is a “state” $\omega \in \Omega \equiv \{\omega_1, \omega_2, \dots, \omega_K\}$, where $K > 1$ is an odd integer. If $\omega < \omega_K$ then $x = 2$; if $\omega = \omega_K$ then $x = 0$. The distribution of ω is uniform. Player 1’s type is depicted by the partition

$$\{\{\omega_1\}, \{\omega_2, \omega_3\}, \dots, \{\omega_{K-1}, \omega_K\}\},$$

which represents what he learns about ω . Player’s 2 type is depicted by the partition

$$\{\{\omega_1, \omega_2\}, \dots, \{\omega_{K-2}, \omega_{K-1}\}, \{\omega_K\}\}.$$

(Exercise: write down the Bayesian game in our standard notation. Note that the distribution of types is correlated.)

⁵⁰ That is, for all but a set of payoffs of Lebesgue measure 0.

⁵¹ In a 2×2 two-player game, a strategy is risk dominant if it is the unique best response to the other player playing a 50-50 mixture.

The key property here is that while there can be arbitrarily high finite order of knowledge of x (as $K \rightarrow \infty$), there is never common knowledge of x .⁵² For example, when $\omega = \omega_{K-2}$, there is mutual knowledge that $x = 2$, but player 2 does not know that player 1 knows x ; rather player 2 puts probability 1/2 on player 1 thinking $x = 2$ and probability 1/2 on player holding a uniform distribution over $x \in \{0, 2\}$. When $\omega = \omega_{K-3}$, not only do both players know $x = 2$, but there is mutual knowledge of this (i.e., both players know that both players know $x = 2$); however, player 1 does not know that player 2 knows that player 1 knows x . Notice that for any integer M and any probability $p < 1$, as $K \rightarrow \infty$, there is M^{th} -order knowledge of x with ex-ante probability larger than p . It may appear, then, that for large K , it is “almost common knowledge” that $x = 2$.⁵³ But:

Claim 1. *In the above Bayesian game, IDSDS yields $s_i(t_i) = B$ for $i = 1, 2$ and all t_i .*

Proof. Deleting dominated strategies rules out any strategy that does not have $s_2(\{\omega_K\}) = B$. But then, $s_1(\{\omega_{K-1}, \omega_K\}) = A$ is iteratively dominated because B is risk dominant. The result follows by induction. \square

One take-away is that we have to be careful with what we mean by “almost common knowledge”: the topology with respect to which continuity with common knowledge is viewed is crucial (cf. [fn. 53](#).)

The argument in [Claim 1](#) is due to ideas in [Rubinstein \(1989\)](#); it starkly illustrates the logic by which outcomes in strategic settings can be sensitive to “tails” of higher-order beliefs. Notice that the argument is based on “contagion” — dominance for some types feeds into (iterative) dominance for other types. [Carlsson and Van Damme \(1993\)](#) introduced the terminology of “global games” to model situations in which each player observes the true payoffs of the game with a very small amount of noise; dominance plus contagion arguments prove powerful. I will give you a homework problem illustrating this. [Morris and Shin \(1998\)](#) popularized global games for applications.⁵⁴ More commonly, higher-order uncertainty is usually downplayed in applied settings by modeling them with simple type spaces.

⁵² It is common knowledge that player 2 knows x . (Why?) What is not common knowledge is whether player 1 knows x .

⁵³ However, for any probability p close enough to 1, no matter the value of $K > 1$, there is no state ω in which the following statement is true: both players ascribe probability (at least) p to $x = 2$, both players ascribe probability (at least) p to both players ascribing probability (at least) p to $x = 2$, ... ad infinitum. In other words, for all high p , there is a failure of “common p -belief” ([Monderer and Samet, 1989](#)), which is what drives the ensuing argument.

⁵⁴ While applied scholars often take the message of the global games literature to be that there is a “correct” unique equilibrium in settings that *prima facie* have multiple equilibria, this is not the right interpretation. See me for more discussion if you are interested, or have a look at [Weinstein and Yildiz \(2007\)](#).




Figure 1 – Extensive Form for Predation Game

4. Dynamic Games and Extensive Form Refinements of Nash Equilibrium

We now turn to a study of dynamic games, i.e. games where players are not just moving simultaneously. The underlying theme will be to refine the set of Nash equilibria in these games.

4.1. The Problem of Credibility

One view towards studying dynamic games is to simply write down their normal form, and then proceed as we did when studying simultaneous games. The problem however is that certain Nash equilibria in dynamic games can be very implausible predictions. Let's illustrate this through the following example.

Example 23 (Predation). Firm 1 (the entrant) can choose whether to enter a market against against a single incumbent, Firm 2, or exit. If Firm 1 enters, Firm 2 can either respond by fighting or accommodating. The extensive form and payoffs are drawn in Figure 1.

To find Nash equilibria of this game, we can write out the normal form as follows.

	f	a
e	-3,-1	2,1
x	0,2	0,2

Clearly, (x, f) is a Nash equilibrium, as is (e, a) .⁵⁵ However, (x, f) does not seem like a plausible prediction: conditional upon Firm 1 having entered, Firm 2 is strictly better off accommodating rather than fighting. Hence, if Firm 1 enters, Firm 2 should accommodate. But then, Firm 1 should foresee this and enter, since it prefers the outcome (e, a) to what it gets by playing x . \square

The problem in the Example is that the “threat” of playing f , that is fighting upon entry, is not credible. The outcome (x, f) is Nash because *if* Firm 2 would fight upon entry, then Firm 1 is better off exiting. However, in the dynamic game, Firm 1 should not believe such an “empty threat”. The crux of the matter is that the Nash equilibrium concept places no restrictions on players’ behavior at nodes that are

⁵⁵ There are also some MSNE involving x .

never reached on the equilibrium path. In this example, given that Firm 1 is playing x , any action for Firm 2 is a best response, since all its actions are at a node that is never reached when Firm 1 places x . Thus, by choosing an action (f) that it certainly wouldn't want to play if it were actually forced to act, it can ensure that Firm 1's [unique, in this case] best response is to play x , guaranteeing that it in fact won't have to act.

4.2. Backward Induction and Subgame Perfection

The natural way to solve the problem above is to require that a player's strategy specify optimal actions at *every node of the game tree*. That is, when contemplating an action, a player takes as given that the relevant node has been reached, and thus should playing something that is optimal here on out (given her opponents' strategies). This is the principle of *sequential rationality* (which we define formally later). In [Example 23](#), action f is not optimal conditional on the relevant node being reached; only a is. Thus, sequential rationality requires that 2's strategy be a , to which the unique best response for 1 to play e , resulting in the more reasonable outcome (e, a) .

4.2.1. Backward Induction

In general, we can try to apply this logic to any extensive game in the following way: start at the “end” of the game tree, and work “back” up the tree by solving for optimal behavior at each node, determining optimal behavior earlier in the game by anticipating the later optimal behavior. This procedure is known as *backward induction*. In the class of finite games with perfect information (i.e. finite number of nodes and singleton information sets), this is a powerful procedure.

Theorem 8 (Zermelo). *Every finite game of perfect information has a pure strategy Nash equilibrium that can be derived through backward induction. Moreover, if no player has the same payoffs at any two terminal nodes, then backward induction results in a unique Nash equilibrium.*

Proof. The uniqueness part is straightforward. The rest is in MWG pp. 272–273. □

Zermelo's Theorem says that backward induction can be powerful in various *finite* games. For example it implies that even a game as complicated as chess is solvable through backward induction; in chess, one of the following mutually exclusive statements is true:

- White has a strategy that results in a win for him regardless of what Black does;
- Black has a strategy that results in a win for him regardless of what White does;
- Each player has a strategy that results in either a draw or a win for him regardless of what the other player does.

In this sense, chess is “solvable”; alas, no-one knows what the solution is!^{[56](#),[57](#)}

⁵⁶ Note well what the application to chess is: it doesn't just say that every chess game must end in one of the players winning or in a draw (this is trivial, modulo the game ending—see the next footnote); it says that one of the players has a strategy that *guarantees* a win for the player (regardless of what strategy the other player uses), or both player have strategies that *guarantee* at least a draw (regardless of what the strategy the other player uses).

⁵⁷ There is actually a subtle caveat here with regards to real chess. Although game theorists often think of a chess as a finite game, and this is needed to apply the Theorem above directly, in fact the official rules of chess make it an infinite game. Nevertheless, the result above concerning chess (that either one player has a winning strategy, or both players have strategies that individually ensure at least a draw) is true. See [Ewerhart \(2002\)](#).

Corollary 4. *Every finite game of perfect information has a PSNE.*

As in illustration of how backward induction operates, here is a famous game.

Example 24 (Centipede Game). Two players, 1 and 2, take turns choosing one of two actions each time, *continue* or *stop*. They both start with \$1 in their respective piles, and each time i says *continue*, \$1 is taken away from his pile, and \$2 are added to the other player's pile. The game automatically stops when both players have \$1000 in their respective piles. Backward induction implies that a player should say *stop* whenever it is his turn to move. In particular, Player 1 should say *stop* at the very first node, and both players leave with just the \$1 they start out with.⁵⁸ \square

4.2.2. Subgame Perfect Nash Equilibrium

We are now going to define a refinement of Nash equilibrium that captures the notion of backward induction. To do so, we need some preliminaries. Recall from [Definition 1](#) that an extensive form game, Γ_E , specifies a host of objects, including a set of nodes, \mathcal{X} , an immediate predecessor mapping $p(x)$ that induces a successor nodes mapping $S(x)$, and a mapping $H(x)$ from nodes to information sets.

Definition 28 (Subgame). A subgame of an extensive form game, Γ_E , is a subset of the game such that

1. (a) there is a unique node in the subgame, x^* , such that $p(x^*)$ is not in the subgame. Moreover, (b) x^* is not a terminal node, and (c) $H(x^*) = \{x^*\}$;
2. a node, x , is in the subgame if and only if $x \in \{x^*\} \cup S(x^*)$;
3. if node x is in the subgame, then so is any $\tilde{x} \in H(x)$.⁵⁹

Remark 25. Notice that any extensive form game as whole is always a subgame (of itself). Thus, we use the term *proper subgame* to refer to a subgame where $x^* \neq x_0$ (recall that x_0 is the root of the game).

Exercise 9. Draw an extensive form game and indicate three different parts of it that respectively each fail the three components of the definition of a subgame.

The key feature of a subgame is that it is a game in its own right, and hence, we can apply the concept of Nash equilibrium to it. We say that a strategy profile, σ , in the game Γ_E *induces* a Nash equilibrium in a particular subgame of Γ_E if the [probability distribution over] moves specified by σ for information sets in the subgame constitute a Nash equilibrium when the subgame is considered as a game by itself.

Definition 29 (Subgame Perfect NE). A Nash equilibrium, σ^* , in the extensive form game, Γ_E , is a subgame perfect Nash equilibrium (SPNE) if it induces a Nash equilibrium in *every* subgame of Γ_E .

In finite extensive form games with possibly imperfect information, we conduct *generalized backward induction* as follows:

1. Consider the maximal subgames (that is, subgames that have no further proper subgames) and pick a Nash equilibrium in each maximal subgame (one exists! why?).

⁵⁸ Experiments with this game show that most players tend to continue until there is a substantial sum of money in both their piles, and then one will say *stop*, almost always before the game automatically stops.

⁵⁹ Actually, this condition together with condition 1(a) implies 1(c). Can you prove it?




Figure 2 – Extensive Form for Predation with Niches Game

2. Replace each maximal subgame with a “terminal” node that has the payoffs of the Nash equilibrium we picked in the subgame. After replacement, call this a “reduced” game.
3. Iterate the process on a successive sequence of “reduced” games until the whole tree has been replaced with a single terminal node.

It should be intuitive that the process of generalized backward induction yields a SPNE of a game of imperfect information, and conversely every SPNE survives generalized backward induction.⁶⁰ The following example illustrates the idea.

Example 25 (Predation with Niches). This is an extended version of the Predation game from [Example 23](#). Now, Firm 1 (the entrant) first chooses to enter or not. If it enters, then the two firms simultaneously choose a niche of the market to compete in: a (A) or b (B). Niche b is the “larger” niche. The extensive form and payoffs are drawn in [Figure 2](#).

To find the pure strategy SPNE equilibria of this game, we employ generalized backward induction as follows. Notice that there is only one proper subgame here. We can write out the normal form of this proper subgame as follows.

	A	B
a	-6,-6	-1,1
b	1,-1	-3,-3

There are two PSNE in the subgame: (a, B) and (b, A) . If we replace the subgame with a terminal node corresponding to (a, B) payoffs, then it follows that Firm 1 prefers to play x at its first move. If on the other hand we replace the subgame with a terminal node corresponding to (b, A) payoffs, then it follows that Firm 1 prefers to play e at its first move. Therefore, the two pure strategy SPNE of this game are (xa, B) and (eb, A) . \square

⁶⁰ See MWG Proposition 9.B.3 (p. 277) for a precise statement.

Since every subgame of a finite game has a Nash equilibrium, and a SPNE of the whole game is derived by “folding” together Nash equilibria of each subgame, we have the following existence result.

Theorem 9 (SPNE Existence). *Every finite game has a SPNE.*

Remark 26. MWG — and other textbooks — tend not to mention the above result because existence of SPNE can also be derived as a corollary to the existence of *Sequential Equilibrium*, which we discuss later on.

Obviously, in games of perfect information, every SPNE can be derived through the basic backward induction process, since generalized backward induction reduces to this with perfect information. Using Zermello’s Theorem, it follows that

Corollary 5 (Pure Strategy SPNE). *Every finite game of perfect information has a pure strategy SPNE. Moreover, if no player has the same payoffs at any two terminal nodes, there is a unique SPNE.*

Exercise 10. Recall the centipede game from [Example 24](#). Prove that the unique SPNE is where each player plays “stop” at every node he is the actor at (this is simply an application of the results above!). Show that there are a plethora of Nash equilibria. However, prove that in every Nash equilibrium, Player 1 must play “stop” at the first node with probability 1 (hint: you have to consider mixed strategies).

Example 26 (Finite Horizon Bilateral Bargaining). Players 1 and 2 are bargaining over the split of $v > 0$ dollars. The game lasts a finite odd number of $T \in \{1, 3, \dots\}$ periods. In period 1, player 1 offers player 2 some amount $b_1 \in [0, v]$, which player 2 can either accept or reject. If player 2 accepts, the game ends, with the proposed split enforced. If player 2 rejects, then we move onto period 2, where player 2 offers player 1 some amount $b_2 \in [0, v]$, and player 1 can either accept or reject. The game proceeds in this way for up to T periods; in the $T + 1$ period, the game necessarily ends and both players get nothing if agreement has not been reached. Players discount the future by a factor $\delta \in (0, 1)$, so that a dollar received in period t gives a payoff of δ^{t-1} . What is the set of SPNE?

Remarkably, there is a unique SPNE. (This is not a consequence of Zermello’s Theorem.) We find the SPNE through backward induction. Start at the end: in the final period, T , player 2 weakly prefers to accept any offer $b_T \geq 0$, and strictly so if $b_T > 0$. Given this, the *only* offer for player 1 in period T that can be part of a SPNE is $b_T = 0$, which player 2 must respond to by accepting in a SPNE.⁶¹ The corresponding payoffs in the subgame starting at period T are $(\delta^{T-1}v, 0)$. Now consider the subgame starting in period $T - 1$, where player 2 is the proposer. Player 1 weakly prefers to accept any offer that gives him a payoff of at least $\delta^{T-1}v$, and strictly so if it provides a payoff of strictly more than $\delta^{T-1}v$. Thus, the only offer for player 2 that can be part of a SPNE is $b_{T-1} = \delta v$, which player 1 must respond to by accepting in a SPNE. The corresponding payoffs in the subgame starting at period $T - 1$ are thus $(\delta^{T-1}v, \delta^{T-2}(1 - \delta)v)$. Now consider period $T - 2$. The above logic implies that the only offer that is part of a SPNE is such that $\delta^{T-3}b_{T-2} = \delta^{T-2}(1 - \delta)v$, i.e. $b_{T-2} = \delta(1 - \delta)v$, which must be accepted in a SPNE.

⁶¹ Any $b_T > 0$ is strictly worse for player 1 than some $b_T - \varepsilon > 0$, which player 2 necessarily accepts in a SPNE. If player 2 rejects $b_T = 0$ with positive probability, then player 1 would strictly prefer to offer some $b_T = \varepsilon > 0$. Note that player 2 is indifferent when offered 0, but it is critical that we resolve his indifference in favor of accepting in order to have a SPNE.




Figure 3 – Extensive Form for Predation version 2 Game

Continuing this process all the way back to period 1, by induction we see that there is a unique SPNE that involves the sequence of offers, $b_T^* = 0$ and for all $t \in \{1, \dots, T-1\}$,

$$b_{T-t}^* = -v \sum_{\tau=1}^t (-\delta)^\tau = \frac{\delta(1 - (-\delta)^t)}{1 + \delta} v.$$

In any period t , the responder's accepts an offer if and only if it is at least b_t^* . Player 1's equilibrium payoff is therefore

$$\pi_1^* = v - b_1 = v \left[1 - \frac{\delta(1 - (-\delta)^{T-1})}{1 + \delta} \right] = v \frac{1 + \delta^T}{1 + \delta},$$

and player 2's equilibrium payoff is $\pi_2^* = b_1 = v \frac{\delta - \delta^T}{1 + \delta}$.

Observe that player 1's equilibrium payoff is larger than player 2's; this because player 1 is both the first and last proposer—each aspect confers some benefit.⁶² As $T \rightarrow \infty$, the equilibrium payoff vector $\rightarrow (v/(1 + \delta), v\delta/(1 + \delta))$, which in turn is $\approx (v/2, v/2)$ when $\delta \approx 1$. In other words, when the bargaining horizon is long and both players are patient, each gets close to half the pie. Impatience pushes in favor of the first proposer. \square

Exercise 11. Construct a Nash equilibrium in the finite horizon bargaining game in which player 1 gets a payoff of 0.

4.3. Systems of Beliefs and Sequential Rationality

A limitation of the preceding analysis is subgame perfection is powerless in dynamic games where there are no proper subgames.

Example 27 (Predation version 2). Modify the Predation game in [Example 23](#) so that Firm 1 now has two ways in which it can enter the market. The extensive form is in [Figure 3](#). Now, the set of SPNE is exactly

⁶² Notice that in a 2-period version of this problem, for large δ , the second proposer gets most of the pie in equilibrium; on the other hand, for small δ , the first proposer gets most of the pie in equilibrium.

the same as the set of NE, because there are no proper subgames to this game. Hence, the NE (x, f) is a SPNE. But (x, f) is no more plausible here than it was in the original Predation example; regardless of whether Firm 1 entered using e_1 or e_2 , given that it has actually entered, Firm 2 is strictly better off playing a rather than f . \square

4.3.1. Weak Perfect Bayesian Equilibrium

Accordingly, we need a theory of “reasonable” choices by players at all nodes, and not just at those nodes that are parts of proper subgames. One way to approach this problem in the above example is to ask: could f be optimal for Firm 2 when it must actually act for *any* belief that it holds about whether Firm 1 played e_1 or e_2 ? Clearly, no. Regardless of what Firm 2 thinks about the likelihood of e_1 versus e_2 , it is optimal for it to play a . This motivates a formal development of beliefs in extensive form games. Recall the notation we use in such games: the set of nodes is \mathcal{X} , the set of successor nodes to any x is denoted $S(x)$, \mathcal{H} is the set of information sets, $H(x)$ is the information set that a node x belongs to, and $\iota(H)$ is the player who acts at an information set H .

Definition 30 (System of Beliefs). A system of beliefs is a mapping $\mu : \mathcal{X} \rightarrow [0, 1]$ such that for all $H \in \mathcal{H}$, $\sum_{x \in H} \mu(x) = 1$.

In words, a system of beliefs, μ , specifies the relative probabilities of being at each node of an information set, for every information set in the game. Obviously, $\mu(x) = 1$ for all x such that $H(x) = \{x\}$. That is, at singleton information sets, beliefs are degenerate.

Using this notion of beliefs, we can state formally what it means for strategies to be sequentially rational given some beliefs. Let $\mathbb{E}[u_i \mid H, \mu, \sigma_i, \sigma_{-i}]$ denote player i 's expected utility starting at her information set H if her beliefs regarding the relative probabilities of being at any node, $x \in H$ is given by $\mu(x)$, and she follows strategy σ_i while the other plays follow the profile of strategies σ_{-i} . Think about this via behavioral strategies: conditional on any node in the information set H , the behavioral strategy σ induces a probability distribution over terminal nodes, and so some expected utility for player i ; we integrate these expected utilities over all the nodes in H using weights given by μ .

Definition 31 (Sequential Rationality). A strategy profile, σ , is sequentially rational at information set H given a system of beliefs, μ , if

$$\mathbb{E}[u_{\iota(H)} \mid H, \mu, \sigma_{\iota(H)}, \sigma_{-\iota(H)}] \geq \mathbb{E}[u_{\iota(H)} \mid H, \mu, \tilde{\sigma}_{\iota(H)}, \sigma_{-\iota(H)}]$$

for all $\tilde{\sigma}_{\iota(H)} \in \Sigma_{\iota(H)}$.

A strategy profile is sequentially rational given a system of beliefs if it is sequentially rational at all information sets given that system of beliefs.

In words, a strategy profile is sequentially rational given a system of beliefs if there is no information set such that once it is reached, the actor would strictly prefer to deviate from his prescribed play, given his beliefs about the relative probabilities of nodes in the information set and opponents' strategies.

With these concepts in hand, we now define a *weak perfect Bayesian equilibrium*. The idea is straightforward: strategies must be sequentially rational, and beliefs must be derived from strategies whenever possible

via Bayes rule. Recall that the statistical version of Bayes Rule (discrete case) says that given any events A, B, C , where B occurs with positive probability given that C does,

$$\text{Prob}(A|B, C) = \frac{\text{Prob}(B|A, C)\text{Prob}(A|C)}{\text{Prob}(B|C)}. \quad (6)$$

Definition 32 (Weak PBE). A profile of strategies, σ , and a system of beliefs, μ , is a weak perfect Bayesian equilibrium (weak PBE), (σ, μ) , if:

1. σ is sequentially rational given μ ; and
2. μ is derived from σ through Bayes rule whenever possible. That is, for any information set H such that $\text{Prob}(H|\sigma) > 0$, and any $x \in H$,

$$\mu(x) = \frac{\text{Prob}(x|\sigma)}{\text{Prob}(H|\sigma)}. \quad (7)$$

The first part of the definition is the requirement of sequential rationality, and the 2nd part is the consistency of beliefs as embodied by Bayes Rule.⁶³ Keep in mind that strictly speaking, a weak PBE is a strategy profile-belief pair. However, we will sometimes be casual and refer to just a strategy profile as a weak PBE. This implicitly means that there is at least one system of beliefs such the pair forms a weak PBE.

Remark 27. The moniker “weak” in weak PBE is because absolutely no restrictions are being placed on beliefs at information sets that do not occur with positive probability in equilibrium, i.e. on *out-of-equilibrium* information sets. To be more precise, no consistency restriction is being placed; we do require that they be well-defined in the sense that beliefs are probability distributions. As we will see, in many games, there are natural consistency restrictions one would want to impose on out of equilibrium information sets as well.

To see the power of weak PBE, we return to the motivating example of Predation version 2.

Example 28 (Weak PBE in Predation version 2). Recall that [Example 27](#) had (x, f) as a SPNE. I claim that it is not a weak PBE equilibrium (or more precisely, there is no weak PBE involving the strategy profile (x, f)). This is proved by showing something stronger: that there is a unique weak PBE, and it does not involve (x, f) . To see this, observe that any system of beliefs in this game, μ , can be described by a single number, $\lambda \in [0, 1]$, which is the probability that μ places on the node following e_1 . But for any λ , the uniquely optimal strategy for Firm 2 is a . Hence, only a is sequentially rational for any beliefs. It follows that *only* (e_1, a) can be part of a weak PBE. Given this, we see that beliefs are also pinned down as $\lambda = 1$ by Bayes Rule. There is thus a unique strategy-belief pair that forms a weak PBE in this game. \square

Proposition 8 (Nash and weak PBE). *A strategy profile, σ , is a Nash equilibrium of an extensive form game if and only if there exists a system of beliefs, μ , such that*

1. σ is sequentially rational given μ at all H such that $\text{Prob}(H|\sigma) > 0$ (not necessarily at those H such that $\text{Prob}(H|\sigma) = 0$); and

⁶³ To see that the 2nd part of the definition is just Bayes Rule, think of the the left-hand side of (7) as $\text{Prob}(x|H, \sigma)$ and the right-hand side as $\frac{\text{Prob}(H|x, \sigma)\text{Prob}(x|\sigma)}{\text{Prob}(H|\sigma)}$ and then compare to (6). We are able to simplify because $\text{Prob}(H|x, \sigma) = 1$ by the hypothesis that $x \in H$.




Figure 4 – Extensive Form for Predation version 3 Game

2. μ is derived from σ through Bayes rule at every information set H such that $\text{Prob}(H|\sigma) > 0$.

Proof. Exercise. □

Since we have already seen that not every NE is a weak PBE, and since the only difference in the above Proposition with the definition of a weak PBE is that sequential rationality is only required at a subset of information sets (those that occur with positive probability, rather than all of them), we immediately have the following corollary.

Corollary 6. *Every weak PBE is a Nash equilibrium; but not every Nash equilibrium is a weak PBE.*

(You might be wondering what we can say about SPNE and weak PBE equilibrium: [Example 28](#) showed that a SPNE need not be a weak PBE ... for the converse, hold off for just a little bit!)

The next example shows how to solve for weak PBE in a more complicated case than before.

Example 29 (Predation version 3). This is yet another variant of a predation game. The extensive form is drawn in [Figure 4](#). The key difference with before is that Firm 2's optimal action if it must act depends on whether Firm 1 played e_1 or e_2 . To solve for weak PBE, let μ_1 be the probability the system of beliefs assigns to the node following e_1 , let σ_f be Firm 2's probability of playing f , and σ_x , σ_{e_1} and σ_{e_2} denote the respective probabilities in Firm 1's strategy.

First, observe that x can never part of a weak PBE because e_2 strictly dominates x for Firm 1, hence it is never sequentially rational for Firm 1 to play x . Next, observe that it is sequentially rational for Firm 2 to put positive probability on f if and only if $-1(\mu_1) + -1(1 - \mu_1) \geq -2\mu_1 + (1 - \mu_1)$, i.e. if and only if $\mu_1 \geq \frac{2}{3}$. Now we consider two cases.

1. Suppose that $\mu_1 > \frac{2}{3}$ in a weak PBE. Then Firm 2 must be playing f (this is uniquely sequentially rational given the beliefs). But then, Firm 1 must be playing e_2 (since $\gamma > 0 > -1$), and Bayes rule requires $\mu_1 = 0$, a contradiction.
2. Suppose that $\mu_1 < \frac{2}{3}$ in a weak PBE. Then Firm 2 must be playing a (this is uniquely sequentially rational given the beliefs). But then, Firm 1 must be playing e_1 (since $3 > 2 > 0$), and Bayes rule beliefs requires $\mu_1 = 1$, a contradiction.




Figure 5 – Extensive Form for [Example 30](#)

Therefore, any weak PBE has $\mu_1 = \frac{2}{3}$. Given the earlier observation that $\sigma_x = 0$ in any weak PBE, Bayes rule holds if and only if $\sigma_{e_1} = \frac{2}{3}$ and $\sigma_{e_2} = \frac{1}{3}$. For this to be sequentially rational for Firm 1 requires it to be indifferent between e_1 and e_2 , which is the case if and only if $-1\sigma_f + 3(1 - \sigma_f) = \gamma\sigma_f + 2(1 - \sigma_f)$, i.e. if and only if $\sigma_f = \frac{1}{\gamma+2}$.

We conclude that there is a unique weak PBE in this game. \square

Exercise 12. Solve for the weak PBE in the above game when $\gamma \in (-1, 0)$.

Now we return to the issue of relating weak PBE to Nash equilibria, in particular, to SPNE. As we saw earlier in [Example 28](#), not every SPNE is a weak PBE. The following example demonstrates that not every weak PBE is a SPNE.

Example 30 (wPBE $\not\subseteq$ SPNE). In the game given by [Figure 5](#), (b, u, l) is a weak PBE (with what beliefs?) but the unique SPNE is (t, u, r) . [Question: what can you say about player 2's strategy in any wPBE?]

The intuition behind the example is simple: weak PBE places minimal restrictions on behavior at a subgame that is off the equilibrium path (i.e., not reached under the equilibrium strategies), while SPNE requires more. Loosely speaking, within an off-path subgame, weak PBE just rules out dominated actions, while SPNE requires mutual best responses. We have thus proved the following.

Proposition 9. A weak PBE need not be a SPNE; a SPNE need not be a weak PBE.

Exercise 13. Prove that in any extensive form game of perfect information, the set of weak PBE is the same as the set of SPNE.

The fact that a weak PBE may not be subgame perfect (in imperfect information games) motivates a strengthening of the solution concept. One strengthening is that of a *perfect Bayesian equilibrium* (PBE, without the “weak” moniker). This requires that the strategy profile-belief pair be not only a weak PBE,




Figure 6 – Extensive Form for [Example 31](#)

but moreover induce a weak PBE in every proper subgame.⁶⁴ Clearly, this would assure subgame perfection. However, it is still too weak in its restrictions on beliefs off the equilibrium path, as the following example demonstrates.

Example 31 (Implausible Beliefs in PBE). [Figure 6](#) shows a game form.⁶⁵ Suppose payoffs were specified so that it were strictly dominant for player 1 to choose t . Then, there are a continuum of PBE, all with the same strategy profile, (t, a) , but supported by different beliefs: any $\mu(x_h) \in [0, 1]$ and $\mu(x_l) = 1 - \mu(x_h)$. The reason we can do this is that there are no proper subgames, hence the set of PBE is the same as set of weak PBE, and no restrictions are placed on beliefs at information sets that are off the equilibrium path. However, it should be clear that the only reasonable belief for player 2 is $\mu(x_h) = \mu(x_l) = \frac{1}{2}$, since player 1's move cannot depend on N 's choice, and N picks h and l with equal probabilities.⁶⁶ \square

4.3.2. Sequential Equilibrium

A stronger equilibrium concept that we now introduce is that of *sequential equilibrium*, due to [Kreps and Wilson \(1982\)](#).

Definition 33 (Sequential Equilibrium). A strategy profile, σ , and a system of beliefs, μ , is a sequential equilibrium (SE), (σ, μ) , if:

1. σ is sequentially rational given μ ;
2. There exists a sequence of *fully mixed* strategy profiles $\{\sigma^k\}_{k=1}^{\infty} \rightarrow \sigma$, such that $\{\mu^k\}_{k=1}^{\infty} \rightarrow \mu$, where μ^k denotes the beliefs derived from strategy profile σ^k using Bayes rule.

⁶⁴ In textbook treatments, PBE is often defined as something that is a bit stronger than the definition I have given; that is, not only is it a weak PBE that is subgame perfect, but other subtle conditions are typically imposed. See for example [Fudenberg and Tirole \(1991a, pp. 331-333\)](#) or MWG (p. 452).

⁶⁵ Recall that this is called a game form rather than a game because payoffs are not specified.

⁶⁶ To keep it simple, this example is such that even the implausible PBE are outcome-equivalent (i.e. have the same strategy profile) to plausible PBE. But it is easy to construct variants where this need not be the case: see MWG Example 9.C.5 (p. 289) for one.

Part 1 of the definition is the same as that of weak PBE. The 2nd part is a consistency of beliefs notion that is more demanding than that of weak PBE.⁶⁷ To interpret it, first note that given a fully mixed strategy profile, every information set is reached with positive probability, and hence Bayes rule completely determines a belief system. Thus, the definition roughly requires that the equilibrium beliefs be “close to” beliefs that are fully determined via Bayes rule from a fully mixed strategy profile that is “near by” the equilibrium strategy profile. There is obviously a connection here with trembling hand perfect equilibrium,⁶⁸ since these “near by” fully mixed strategy profiles can be thought of as arising from mistakes in playing the equilibrium strategies. Since sequential equilibrium places more restrictions on belief consistency than weak PBE, it follows that *every sequential equilibrium is a weak PBE*. The converse is not true, as we now show by example.

Example 32. Recall [Example 30](#) that had (b, u, l) as weak PBE strategies. Let us argue that the unique SE strategy profile is the unique SPNE: (t, u, r) . Let σ_u and σ_d denote the probabilities used by player 2; H_3 denote the only non-singleton information set; and let the four decision nodes be denoted as x_1, x_2, x_{3u} and x_{3d} respectively. In any fully mixed strategy profile, σ , Bayes rule implies

$$\mu_\sigma(x_{3u}) = \frac{\text{Prob}(x_{3u}|\sigma)}{\text{Prob}(H_3|\sigma)} = \frac{\text{Prob}(x_{3u}|x_2, \sigma)\text{Prob}(x_2|\sigma)}{\text{Prob}(H_3|x_2, \sigma)\text{Prob}(x_2|\sigma)}.$$

[Read $\text{Prob}(x_{3u}|x_2, \sigma)$ as the probability that x_{3u} is reached given that x_2 has been reached and σ is the profile being played, and so forth.]

Canceling terms and noting that $\text{Prob}(H_3|x_2, \sigma) = 1$, we have

$$\mu_\sigma(x_{3u}) = \text{Prob}(x_{3u}|x_2, \sigma) = \sigma_u.$$

Thus, for any sequence of fully mixed profiles $\{\sigma^k\}_{k=1}^\infty$ that converges to a profile σ^* , the limit of the sequence of beliefs derived from Bayes rule necessarily has $\mu_{\sigma^*}(x_{3u}) = \sigma_u^*$. Since player 2 being sequentially rational imposes that $\sigma_u^* = 1$ in any SE, σ^* , it follows that $\mu_{\sigma^*}(x_{3u}) = 1$ and hence player 3 must play r in any sequential equilibrium. It is then uniquely sequentially rational for player 1 to play t in any SE. \square

Remark 28 (Lower hemicontinuity of Beliefs in a SE). To get a better intuition for how sequential equilibrium works, it is useful to think about the mapping from mixed strategies to beliefs for the relevant player(s) in a game. Consider [Example 30](#). Holding fixed a σ_u , we can think of the allowable beliefs in a weak PBE by defining the maximal belief correspondence $B(\sigma_t)$ which satisfies the property that $B(\sigma_t)$ is derived from Bayes Rule whenever possible, i.e. so long as $\sigma_t > 0$. Graphically, this is illustrated in [Figure 7](#). The key

⁶⁷ It is possible to show that for any strategy profile σ , there is a system of beliefs satisfying this consistency requirement.

⁶⁸ Note, however, that we only defined THPE for normal form games. It turns out that a normal form THPE need not even be a SPNE. There is a definition of THPE for extensive form games that I won't pursue in detail (cf. MWG pp. 299-300). Briefly, an extensive form THPE is defined using normal form THPE but applied to the “agent normal form” of the extensive form, in which each information set of a player is treated as belonging to a separate copy of that player—so that trembles are independent across a player's information sets. The upshot is that extensive form THPE is slightly more demanding than sequential equilibrium (i.e., every extensive form THPE is a sequential equilibrium strategy profile, but not vice-versa); nonetheless, generically they coincide (i.e., if we fix a finite extensive game form, the two concepts produce different sets of equilibria for a set of payoffs of Lebesgue measure 0).




Figure 7 – SE and weak PBE Beliefs

point of course is that when $\sigma_t = 0$, Bayes Rule does not apply, and any $B(0) \in [0, 1]$ is permitted in a weak PBE. However, SE requires that $B(0)$ be the limit of $B(\sigma_t)$ as $\sigma_t \rightarrow 0$. In other words, SE requires that we take a selection from correspondence B that satisfies *lower hemicontinuity*. Indeed, this was one of motivations behind how [Kreps and Wilson \(1982\)](#) arrived at their definition of SE.

Exercise 14. Verify that in any SE in [Example 31](#), $\mu(x_h) = \mu(x_l) = \frac{1}{2}$, and work through the lower hemicontinuity of beliefs idea for this case.

Earlier, we noted that it is clear that every SE is a weak PBE. In fact, one can show that any SE is subgame perfect (MWG Proposition 9.C.2), hence it follows that every SE is a PBE (not just a weak PBE). We can summarize with the following partial ordering.

Theorem 10 (Solution Concepts Ordering). *In any extensive form game, the following partial ordering holds amongst solution concepts:*

$$\{SE\} \subseteq \{PBE\} \subseteq \begin{matrix} \{weak\ PBE\} \\ \cap \\ \{SPNE\} \end{matrix} \subseteq \begin{matrix} \{weak\ PBE\} \\ \cup \\ \{SPNE\} \end{matrix} \subseteq \{NE\}.$$

It is fair to say that in extensive form games, sequential equilibrium is a fairly uncontroversial solution concept to apply, at least amongst those who are willing to accept any type of “equilibrium” concept (rather than sticking to just rationalizability, for example). However, it turns out that even sequential equilibrium may not be completely satisfying in some games.

Example 33 (Predation version 4). This is another version of the Predation game, whose extensive form is drawn in [Figure 8](#). I leave it as an exercise to show that (x, f) and (e_1, a) are both SE strategy profiles (the latter is straightforward; showing the former requires you to find a sequence of fully mixed strategies and implied beliefs that make the required equilibrium beliefs consistent). However, the equilibrium play (x, f) can be argued to be implausible. Observe that by playing e_2 , player 1 *will get a payoff of no more than -1*,




Figure 8 – Extensive Form for Predation version 4

with certainty. Hence, that strategy is strictly dominated by the strategy x . Accordingly, if player 2 must actually move in the game, he should realize that player 1 must have played e_1 (with the anticipation that player 2 will play a), to which his best response is indeed a . Given this logic, player 1 should play e_1 rather than x ; hence (x, f) is not plausible. \square

This is an example of *forward induction*, because we are requiring player 2 to reason about what he should play at his information set not just on the basis of what will happen *after* his move (that is backward induction, which is degenerate in this game for player 2 since he is the last mover), but moreover on the basis of what player 1 could have rationally done *prior* to 2's move.

Exercise 15. Prove that in the [Example 33](#), there is a SE with strategy profile (x, f) .

Applying forward induction to refine predictions in a game can be somewhat controversial.⁶⁹ However, as you will see when studying signaling games, without it, the set of (sequential) equilibria can be unbearably large; whereas applying it yields a narrow set that accords with intuition.⁷⁰

5. Market Power

In this section, we are going to look at a few static models of markets with a small number of firms. We've already seen some examples before: two-firm Bertrand and Cournot competition. The idea now is to

⁶⁹ This is partly because there is some tension between the motivating idea behind forward induction and the idea of trembles, which is used to intuitively justify SE. In particular, forward induction operates on the notion that an out-of-equilibrium action should be interpreted as a “rational” decision by a player, and subsequent players should ask themselves what to play based on this interpretation (think of the logic we used in [Example 33](#)). On the other hand, trembles are based on the idea that out-of-equilibrium actions result from “mistakes” rather than intended deviations. This is a subtle issue that is more suitable for an advanced course, but do come to talk to me if you are interested in discussing it more.

⁷⁰ For those of you more cynical: this is of course the usual rhetoric that means “accords with the dominant thinking of many people who have thought about this.”

generalize those examples into some broad ideas of oligopoly markets. Throughout, we take a partial-equilibrium approach, focussing on only one market.

5.1. Monopoly

You've seen competitive market analysis earlier in the year; that is a benchmark against which we can view the effect of market power. We begin this section by looking at the other extreme where a single firm produces the good, known as *a monopoly*.

Market demand for the good at price $p \geq 0$ is given by a function $x(p)$, where $x : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{+\infty\}$. Implicit in this formulation is the assumption that consumers are not strategic players, and they are anonymous to the firm in the sense that the firm cannot charge separate prices to different consumers. We assume that $x(p) = 0$ for all $p \geq \bar{p} \in (0, \infty)$, and that $x(\cdot)$ is strictly decreasing on $[0, \bar{p}]$ (thus $x(q) = +\infty \implies q = 0$). The monopolist knows the demand function, and can produce quantity $q \in \mathbb{R}_+$ at a cost $c(q)$, where $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Define the *inverse demand function*, $p : [0, x(0)] \rightarrow \mathbb{R}_+$ by $p(q) = \min\{p : x(p) = q\}$. Observe that $p(0) = \bar{p}$, and for all $q \in (0, x(0)]$, $p(q) = x^{-1}(q)$, so that $p(\cdot)$ is strictly decreasing on $[0, x(0)]$.

Rather than writing the monopolist's objective as choosing price, it is convenient to take the equivalent approach of maximizing quantity, so that the objective is

$$\max_{q \in \mathbb{R}_+} p(q)q - c(q). \quad (8)$$

Under some assumptions, this problem will have a unique solution that can be found by analyzing the first order condition. The assumptions are:

- i) $p(\cdot)$ and $c(\cdot)$ are twice differentiable on their domains. [So we can take 1st and 2nd derivatives.]
- ii) $p(0) > c'(0)$. [So choosing $q = 0$ will not be optimal for the monopolist.]
- iii) There is a unique $q^e \in (0, \infty)$ such that $p(q^e) = c'(q^e)$; and $p'(q^e) < c''(q^e)$. [This will be the socially optimal quantity, and also ensure a solution to the monopolist's problem.]
- iv) $p''(q)q + 2p'(q) - c''(q) < 0$ for all $q \in (0, q^e]$. [Will ensure that FOC is sufficient.]

Assumption (iv) seems convoluted, but is satisfied under various more primitive assumptions. For example, it holds under linear demand ($p(q) = a - bq$) and convex costs.

Recall that in a competitive market, price = marginal cost, so that q^e defined above is the unique competitive market quantity, and also that the socially optimal quantity.⁷¹ In contrast:

Proposition 10. *Under the stated assumptions, the monopolist's problem has a unique solution, $q^m \in (0, q^e)$, given by*

$$p'(q^m)q^m + p(q^m) = c'(q^m). \quad (9)$$

The monopolist produces less than the socially optimal quantity. Moreover, $p(q^m) > c(q^m)$, so that price under monopoly exceeds marginal cost.

⁷¹ The latter part of Assumption (iii) is needed to guarantee that q^e is indeed a social optimum. Can you see why?

Proof. If a solution exists to (8), call it q^m , Assumption (i) implies that the objective is differentiable and must satisfy the FOC:

$$p'(q^m)q^m + p(q^m) \leq c'(q^m), \quad \text{with equality if } q^m > 0.$$

The LHS above is marginal revenue while the RHS is marginal cost. By Assumption (iii), $p(q) < c'(q)$ for all $q > q^e$, and since $p'(\cdot) < 0$, it follows that $q^m \in [0, q^e]$. Since a continuous function attains a maximum on a compact set, there is a solution to (8).

Assumption (ii) implies that $q^m > 0$ and the FOC must hold with equality. Assumption (iv) implies that the objective is strictly concave on $(0, q^e]$, hence there is a unique solution to the FOC. \square

The key observation here is that a monopolist recognizes that by reducing quantity, it increases revenue on *all the units sold* because the price goes up on all units, captured by the term $p'(q^m)q^m$ in (9). On the other hand, the direct effect of reducing quantity only decreases revenue on the marginal unit, captured by the term $p(q^m)$ in (9). When the quantity is q^e , the marginal reduction in revenue is compensated equally by the cost savings, since $p(q^e) = c'(q^e)$. Thus, the *inframarginal* effect of raising revenue on all other units makes it optimal for the monopolist to produce quantity below q^e . Note that the inframarginal effect is absent for firms in a competitive market.

The *deadweight welfare loss from a monopoly* can be quantified as

$$\int_{q^m}^{q^e} [p(q) - c'(q)] dq.$$

Graphically, this would be the region between the inverse demand curve and the marginal cost curve that is foregone under a monopoly relative to a competitive market.

Remark 29. As is suggested by the above discussion of inframarginal effects, the social inefficiency arising from a monopoly is crucially linked to the (often plausible) assumption that the monopolist must charge the same price to all consumers. If the monopolist could instead perfectly *price discriminate* in the sense of charges a distinct price to each consumer (knowing individual demand functions), then the inefficiency would disappear — although all the surplus would be extracted by the monopolist. See MWG (p. 387) for a formal treatment.

5.2. Basic Oligopoly Models

Let us now turn to oligopoly markets with at least 2 firms. We already studied price competition (Bertrand) and quantity competition (Cournot) with exactly two firms (a duopoly). The first task is to generalize to an arbitrary number of firms.

5.2.1. Bertrand oligopoly

The general Bertrand case is straightforward extension of the duopoly analysis and left as an exercise.

Exercise 16. Consider the Bertrand model of [Example 10](#), except that there are now an arbitrary number of $n \geq 2$ firms, each with symmetric linear costs. Assume that if multiple firms all charge the lowest price, they each get an equal share of the market demand at that price. Show that when $n > 2$, (i) there are pure

strategy Nash equilibria where not all firms charge marginal cost; (ii) but in any Nash equilibrium, all sales take place at price = marginal cost (do it first for PSNE; then extend it to MSNE).⁷²

Thus, under Bertrand competition, we see that just two firms is sufficient to make the market perfectly competitive. Although this is striking, it doesn't seem realistic in some applications. As we will see next, the Cournot model is more satisfying in this regard (although arguably not as appealing in the sense that we often think about firms as choosing prices rather than quantities—though in some cases, quantity choice is not unreasonable).

5.2.2. Cournot oligopoly

We generalize the linear Cournot setting of [Example 7](#) as follows. There are $n \geq 2$ identical firms. Each simultaneously chooses quantity $q_i \in \mathbb{R}_+$. Each has a twice differentiable, strictly increasing, and weakly convex cost function $c(q)$, where $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. The inverse market demand given the *total quantity*, $Q = \sum_i q_i$, is given by $p(Q)$, where $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is twice differentiable, strictly decreasing, and weakly concave for all $Q \geq 0$.⁷³ Assume that $p(0) > c'(0)$. Also assume that there is a unique socially efficient total quantity $Q^e \in (0, \infty)$. Social efficiency requires that the aggregate output be distributed efficiently across firms; since costs are weakly convex, one way to do this is to have each firm produce $\frac{1}{n}Q^e$ (this will be the unique way if costs are strictly convex). Thus, we have $p(Q^e) = c'(\frac{1}{n}Q^e)$, i.e. the marginal cost for each firm is the price at the socially efficient aggregate quantity.

Now let us look for a symmetric pure strategy Nash equilibrium for the profit-maximizing firms in this market. Taking as given other firms' quantity choices, firm i 's objective is to choose q_i to maximize

$$\pi_i(q_i, q_{-i}) = p\left(\sum_{j \neq i} q_j + q_i\right)q_i - c(q_i).$$

In a symmetric equilibrium, we cannot have $q_i = 0$, since $p(0) > c'(0)$. So a symmetric equilibrium must have $q_i > 0$ for all i . This means that the FOC must be satisfied with equality for each firm:

$$p'\left(\sum_{j \neq i} q_j + q_i\right)q_i + p\left(\sum_{j \neq i} q_j + q_i\right) - c'(q_i) = 0. \quad (10)$$

The assumptions we made guarantee that the problem is strictly concave when $q_i > 0$ (for any q_{-i}), so the FOC is also sufficient for a maximum. In a symmetric equilibrium, each firm produces the same quantity, $q^* = \frac{1}{n}Q^*$, where Q^* solves

$$p'(Q^*)\frac{1}{n}Q^* + p(Q^*) = c'\left(\frac{1}{n}Q^*\right). \quad (11)$$

Our assumptions guarantee that there is a unique solution $Q^* > 0$ to the above equation. We therefore conclude that:

Proposition 11. $Q^e > Q^* > q^m$, where q^m is the monopoly quantity. Moreover, the market price is strictly above each firm's marginal cost.

⁷² For the MSNE argument, you can assume that there is a well-defined monopoly price (i.e., $(p - c)Q(p)$ has a unique, finite maximum).

⁷³ The assumption that $p(\cdot)$ is weakly concave is not a conventional property of demand functions, although it is satisfied for linear demand. It is assumed to ensure that the payoff function for each firm is quasiconcave in its own quantity, which recall is an assumption to guarantee existence of PSNE in infinite games ([Theorem 3](#)).

Proof. If $Q^* = q^m$, the LHS of (11) is strictly greater than the LHS of (9); while the RHS is weakly smaller than that of (9). Thus equation (11) is not satisfied when $Q^* = q^m$. By concavity of LHS and convexity of RHS of (11), we must have $Q^* > q^m$.

If $Q^* = Q^e$, the RHS of (11) is equal to second term of the LHS, so is strictly smaller than the LHS, since $p' < 0$. Again, by concavity/convexity, we must have $Q^* < Q^e$.

That price is above marginal cost follows immediately from (11) and that $p' < 0$. \square

Thus, the total quantity under Cournot oligopoly with 2 or more firms lies strictly between the socially optimal quantity and the monopoly quantity. The intuition is as follows: following the same logic as in the monopoly case, each firm has an incentive to reduce its individual quantity when the aggregate output is the social optimum, because of the inframarginal effect. However, as compared to the monopoly case, reducing quantity slightly only has an individual inframarginal benefit of $p'(Q)\frac{1}{n}Q$, since each firm has $\frac{1}{n}$ market share. This leads to a smaller reduction in quantity than would occur under a monopolist. To put it differently: when reducing quantity, each oligopolist has a positive externality on all the other firms (since it leads to an increase in price on all the units they sell) that it does not internalize.

This line of reasoning suggests that $n \rightarrow \infty$, the incentive to reduce quantity starting at the social optimum would vanish, since each individual firm would only be producing a very small quantity and thus the individual inframarginal gain is small. On the other hand, one might worry that because there are more and more firms, even if any individual firm's quantity reduction (relative to the social optimum) vanishes, the total quantity reduction does not. This turns out not to be the case under so long as the socially optimal quantity remains bounded.⁷⁴ In the proposition below, we introduce a subscript index for the number of firms.

Proposition 12. *Assume that $\{Q_n^e\}_{n=1}^\infty$ is bounded. Then $\lim_{n \rightarrow \infty} Q_n^* = \lim_{n \rightarrow \infty} Q_n^e$.*

Proof. Rewrite equation (11) as

$$-p'(Q_n^*)\frac{1}{n}Q_n^* = p(Q_n^*) - c'(\frac{1}{n}Q_n^*).$$

Let a finite upper bound on Q_n^e be \bar{Q}^e . Proposition 11 implies that for any n , $Q_n^* \in [q^m, \bar{Q}^e]$, hence Q_n^* is uniformly bounded. Since p is twice differentiable, $p'(Q_n^*)$ is also uniformly bounded. Thus the LHS above converges to 0 as $n \rightarrow \infty$. So the RHS must, as well. This implies the result. \square

Exercise 17. *Using the linear Cournot model of Example 7, but now allowing for $n \geq 2$ identical firms, solve explicitly for the symmetric PSNE quantities and market price. Verify directly the statements of Proposition 11 and Proposition 12 for this example.*

5.3. Stackelberg Duopoly

In a Stackelberg duopoly, both firms choose quantities, but one firm moves first and the second observes the first mover's choice before making its choice. The important point is that the first mover is able to make a commitment to its quantity that the second mover must incorporate into its choice. The standard

⁷⁴ The socially optimal quantity could, in principle, go to ∞ even though demand does not change, because by adding firms we are changing the aggregate production technology. For instance, when $c(\cdot)$ is strictly convex with $c(0) = c'(0) = 0$, then with a very large number of firms, society can produce large quantities at close to zero total cost by making each firm produce a very small quantity.

application is to a market where there is an incumbent who is facing a potential entrant, with the incumbent being able to pre-commit before the entrant arrives.

To illustrate the main idea, consider a linear market competition model, where inverse demand is given by $p(Q) = \max\{0, a - bQ\}$, and two identical firms have linear costs $c(q_i) = cq_i$, where $a, b, c > 0$ and $a > c$. Firms choose non-negative quantities, but firm 1 chooses its quantity before firm 2 (the game is one of perfect information). We call this the linear Stackelberg Duopoly. Note that a strategy for firm 2 is a function from \mathbb{R}_+ to \mathbb{R}_+ .

We begin by observing that there are plenty of pure strategy Nash equilibria.

Exercise 18. Show that in the linear Stackelberg Duopoly, for any $(q_1, q_2) \in \mathbb{R}_+^2$ such that $q_1 \leq \frac{a-c}{b}$ and $q_2 = \frac{a-c-bq_1}{2b}$, there is a pure strategy Nash equilibrium where firm 1 chooses q_1 and firm 2 chooses q_2 on the equilibrium path.

However, most of these Nash equilibria rely upon incredible threats by firm 2: they are not subgame perfect.

Proposition 13. The linear Stackelberg Duopoly has a unique SPNE: firm 2 plays $s_2(q_1) = \max\{0, \frac{a-c-bq_1}{2b}\}$ and firm 1 plays $s_1 = \frac{a-c}{2b}$. The equilibrium path quantity choices are therefore $q_1 = \frac{a-c}{2b}$ and $q_2 = \frac{a-c}{4b}$, so that total equilibrium quantity is $\frac{3(a-c)}{4b}$.

Proof. Exercise. □

Thus, in the linear Stackelberg duopoly, we see that firm 1 is more aggressive (produces more) and firm 2 is less aggressive (produces less) than in the Cournot duopoly. Total market quantity is between the total of a Cournot duopoly ($\frac{2(a-c)}{3b}$) and the efficient quantity ($\frac{a-c}{b}$). You can verify that firm 1's profit is higher here than in the Cournot game, and also higher than firm 2's. On the other hand, firm 2 does worse here than in the Cournot game.

The themes apply to general Stackelberg duopoly settings (not just this linear case). The main point is that there is a *first mover advantage* to firm 1. Since quantity choices are *strategic substitutes* (each firm's best response is a decreasing function of the other firm's quantity), the first mover is able to exploit its ability to commit by increasing its quantity relative to the Cournot individual quantity, ensuring that the second mover will reduce its quantity in response.

Remark 30. Notice that we have restricted attention to quantity competition above in treating the Stackelberg setting. This is the standard practice. But, one could study sequential moves by the two firms with price competition. I will assign a homework problem on this.

5.4. Price Competition with Endogenous Capacity Constraints

In studying Bertrand competition to this point, we have assumed that a firm can produce any quantity that is demanded at the price it charges. From a short run perspective, this is unrealistic: given that various factors of production are fixed in the short run (after all, that is the definition of short run), there may be an upper bound on how much a firm can supply, i.e. there may be a *capacity constraint*. In this section, I mention how capacity constraints can be used to “bridge” the Bertrand and Cournot models: intuitively, the quantity choices in Cournot can be viewed as long-run choices of capacity, so that firms choose prices in the short-run a la Bertrand given these capacity choices.

The following example illustrates how capacity choices can fundamentally change the nature of Bertrand competition.

Example 34. Consider the linear Bertrand duopoly model of [Example 10](#). Suppose now that each firm is exogenously capacity constrained so that it cannot produce more than $\bar{q} > 0$ units. Without specifying all the details of how the market works, let us assume only that if $p_i > p_j$, then firm j will sell at price p_j a quantity that does not exceed its capacity, and if this is not sufficient to cover the market demand at price p_j , then firm i will sell to some strictly positive mass of consumers at price p_i .

The main observation is that if $\bar{q} \geq Q(c)$, then there is a PSNE where both firms price at marginal cost. But if $\bar{q} < Q(c)$, then it is no longer a Nash equilibrium for both firms to price at marginal cost. Why?

As you will see in a homework problem, for an appropriate specification of how the market works and a wide range of capacity constraints, \bar{q}_1 and \bar{q}_2 (this allows the constraints to differ across firms), the unique equilibrium in the pricing game involves both firms setting an identical price equal to $p(\bar{q}_1 + \bar{q}_2)$, where $p(\cdot)$ is the inverse demand function. In other words, the outcome of Bertrand with exogenous capacity constraints is just like what happens in a Cournot model if firms happen to choose quantities equal to these capacity constraints! But now, consider a prior stage where firms get to endogenously choose their capacity constraints, anticipating that after these are mutually observed, they will play a subsequent Bertrand-with-capacity-constraints pricing game. Intuitively, SPNE of this two-stage game have to yield capacity choices equal to the Nash quantity choices of the Cournot model, since in any subgame following capacity choices (q_1, q_2) , the resulting payoffs are the same as in the Cournot model with quantity choices (q_1, q_2) . See [Kreps and Scheinkman \(1983\)](#) for a thorough analysis. This formally justifies the interpretation of Cournot quantity competition as capturing long-run capacity choices followed by short-run Bertrand price competition.

Remark 31. In the above two-stage model, it is essential that the capacity choices are mutually observed before the price competition. If firms cannot observe each other's capacities when choosing prices, then it is as though they choose capacities and prices simultaneously. Convince yourself that in such a variation, Nash equilibria must yield zero profit (marginal cost pricing), just like unconstrained Bertrand competition.

6. Repeated Games

An important class of dynamic games are so-called *repeated games*. They are used to study strategic interactions that are ongoing over time. For instance, a typical application of the Prisoner's Dilemma or Trust Game ([Example 4](#)) is to represent a situation where two parties have to exert some individually costly effort in order to achieve a socially beneficial outcome, but there is an incentive to free ride on the other party. Often, such interactions do not occur once and for all; rather they occur repeatedly, and each player is able to observe something about the past history of play and condition her own behavior on that information.⁷⁵ This implies that actions are strategically linked over time. Going back to the Trust Game, one intuition is that under repeated interaction, a player may not Cheat because she fears that her partner will retaliate by in turn Cheating in the future. (We will see under what conditions this reasoning can be justified.) As a general principle, the inter-temporal strategic linkage can result in a much richer set of possible behavior than mere repetition of the static game prediction.

⁷⁵ For example, you can think about your own experience exerting effort in a study group.

6.1. Description of a Repeated Game

Stage Game. Formally, a repeated game consists of repetitions of a *stage game*. Although the stage game itself could be a dynamic game, we will focus on the normal-form representation of the stage game, so that a stage game is a one-period simultaneous move game of complete information given by $\langle I, \{A_i\}, \{\pi_i\} \rangle$, where I is the set of players, each A_i is the set of actions (or pure strategies of the stage game) for player i , and $\pi_i : A \rightarrow \mathbb{R}$ is the stage-game von Neumann- Morgenstern utility function for player i (where $A := A_1 \times \cdots \times A_I$). Assume the stage game is finite in the sense that I and A are finite. As usual, this can be relaxed at the cost of technical complications.⁷⁶ In the standard way, we can extend each $\pi_i : \Delta(A) \rightarrow \mathbb{R}$. I will use the notation α_i to denote an element of $\Delta(A_i)$ and refer to this as a *mixed action*.

Repeated Game. A repeated game (sometimes also referred to as a *supergame*) is formed by a repetition of the stage game for $T \in \{1, \dots, \infty\}$ periods. If T is finite, we call it a *finitely repeated game*; if $T = \infty$, we call it an *infinitely repeated game*. We will only consider repeated games of *perfect monitoring*: this means that at the end of every period, every player observes the actions chosen by all players in that period.⁷⁷

Strategies. It is convenient to represent strategies in the repeated game as behavioral strategies. Denote the history of actions at period t as $h^t = \{a^1, \dots, a^{t-1}\}$, where for any t , $a^t = (a_1^t, \dots, a_I^t)$. Thus, a_i^t refers to player i 's action at time t . Let H^t denote the space of all possible histories at time t , and let $H := \bigcup_{t=1}^T H^t$ denote the space of all possible histories in the repeated game. A (pure) strategy for player i in the repeated game can be represented as a function $s_i : H \rightarrow A_i$. In words, it specifies what action to take in any period following any history at that period. A (behavioral) mixed strategy can be represented as $\sigma_i : H \rightarrow \Delta(A_i)$.

Payoffs. If T is finite, then we could represent payoffs as usual via terminal nodes. But in an infinitely repeated game, there are no terminal nodes. So we need to take a different approach. Notice that any pure strategy profile in the repeated game maps into a unique profile of actions in each period (i.e., a unique path through the game tree), and a mixed strategy maps into a distribution over actions in each period. Thus, any (possibly mixed) strategy profile generates a sequence of expected payoffs for each player in each period. It suffices therefore to specify how a player aggregates a sequence of period-payoffs into an overall utility index. The standard approach is to assume exponential discounting, so that a sequence of expected payoffs (v_i^1, v_i^2, \dots) yields player i an aggregate utility of

$$\tilde{u}_i(\{v_i^t\}) = \sum_{t=1}^T \delta^{t-1} v_i^t,$$

⁷⁶ For example, if A is not finite, one needs to assume that a Nash equilibrium exists in the stage game, and also make a boundedness assumption on payoffs.

⁷⁷ Repeated games of *imperfect monitoring*, or *unobservable actions*, are important but beyond the current scope. An excellent textbook reference is [Mailath and Samuelson \(2006\)](#).

where $\delta \in [0, 1)$ is the discount factor, assumed to be the same across players for simplicity.⁷⁸ When dealing with infinitely repeated games, it is useful to normalize this into *average discounted payoffs*,

$$u_i(\{v_i^t\}) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} v_i^t. \quad (12)$$

To understand the normalization, let U_i be the value of (12) for some sequence of (possibly time-varying) payoffs in each period. Now if we insert the constant U_i in place of v_i^t for all t into (12), we would get $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} U_i$, which evaluates precisely as U_i . Thus, in the infinite horizon, we are justified in interpreting (12) as the average discounted payoff: it gives a utility level such that if the player received that utility level in every period, he would have the same average discounted utility as the discounted average of the original stream of period-utilities. Notable then is that by using the average discounted payoff we put the entire repeated game payoff on the “same scale” as any one period.

Remark 32. Above, I have assumed that $\delta < 1$. This means that periods are not treated symmetrically, since a player has a preference for the same period-payoff sooner than later. For finitely repeated games there is no difficulty whatsoever with allowing $\delta = 1$. For infinitely repeated games, the difficulty is that when $\delta = 1$, total payoffs can be unbounded (even though stage-game payoffs are bounded), which can create problems. Nevertheless, if one wants to treat periods symmetrically in an infinitely repeated, there are criteria that can be used instead of exponential discounting, but I will not pursue those here.⁷⁹

Remark 33. The previous remark also points to why the normalization from \tilde{u}_i to u_i (by multiplying by $(1 - \delta)$) is useful for infinitely repeated games: it allows us to take limits as $\delta \rightarrow 1$ while keeping discounted payoffs bounded.

Remark 34. The discount factor can be interpreted in the literal way of time preference. But there is another interpretation: it can capture uncertainty about when the game will end. That is, an infinitely repeated game can be interpreted as situation where players know that the game will end in finite time with probability one, but are unsure of exactly when it will end. To be specific, suppose that players have a discount factor of $\rho \in [0, 1]$ and in addition, also think that conditional on the game having reached period t , it will continue to period $t + 1$ with probability $\lambda \in [0, 1]$ (if the game does not continue, we normalize all payoffs thereafter to zero). Under the condition that either λ or ρ is strictly less than one 1, one can show that this induces the same expected payoff as a setting where the game continues with probability 1 in each period, but players have a discount factor $\delta = \rho\lambda$.

Given this setup for repeated games, we will look at Nash and SPNE equilibria as usual. Observe that the beginning of every period $t > 1$ marks a proper subgame, and these are the only proper subgames because since we have assumed simultaneous moves in the stage game.

⁷⁸ To write payoffs directly as a function of mixed strategies, we would therefore write

$$\tilde{u}_i(\sigma) = \sum_{t=1}^T \delta^{t-1} \sum_{h^t \in H^t} \Pr(h^t | \sigma) \pi_i(\sigma(h^t)).$$

⁷⁹ Probably the most common alternative is known as *limit of means criterion*, which evaluates a sequence of period-payoffs by the formula $\lim_{T \rightarrow \infty} \inf \frac{1}{T} \sum_{t=1}^T v_i^t$ (note: $\lim \inf$ because limit may not exist); another alternative is known as the *overtaking criterion* which just uses $\lim_{T \rightarrow \infty} \inf \sum_{t=1}^T v_i^t$.

6.2. The One-Shot Deviation Principle

An essential result in the study of repeated games—indeed, dynamic games in general—is the *one-shot deviation principle*. It says that when considering profitable deviations for any player from a SPNE, it suffices to consider strategies where he plays as he is supposed to at all information sets except one, i.e. he only behaves differently at a single history. To state the idea precisely, let $\sigma_i|_{h^t}$ be the restriction of strategy σ_i to the subgame following history h^t .

Definition 34 (Profitable One-Shot Deviation). Fix a strategy profile σ . A profitable one-shot deviation for player i is a strategy $\sigma'_i \neq \sigma_i$ s.t.

1. there is a unique history $h^{t'}$ such that for all $\tilde{h}^t \neq h^{t'}$, $\sigma'_i(\tilde{h}^t) = \sigma_i(\tilde{h}^t)$.
2. $u_i(\sigma'_i|_{h^{t'}}, \sigma_{-i}|_{h^{t'}}) > u_i(\sigma|_{h^{t'}})$.

In the Definition above, the first part is what it means to be a one-shot deviation: there is exactly one history at which the strategies differ. Note, however, that even if a deviation is one-shot, it can have a significant effect on the path of play, since the behavior at histories subsequent to $h^{t'}$ can depend significantly on exactly what was played at $h^{t'}$ (we'll see examples later). The second part of the definition is what it means to for the deviation to be profitable. Note here that the profitability is defined *conditional* on the history $h^{t'}$ being reached, even though it may not actually be reached given the profile σ . This means that a Nash equilibrium can have profitable one-shot deviations off the equilibrium path. Yet:

Proposition 14 (One-Shot Deviation Principle). *A strategy σ is a SPNE if and only if there are no profitable one-shot deviations.*

Proof. The “only if” is immediate: if there is a profitable one-shot deviation from σ , it cannot be a SPNE. What needs to be shown is the “if” part. We prove the contrapositive: if σ is not a SPNE then there is a profitable one-shot deviation.

So pick any strategy profile σ and suppose it is not a SPNE. Then there exists a deviation $\tilde{\sigma}_i$ that is profitable for some player i in a subgame following some history \hat{h}^t . Without loss of generality, assume that $\hat{h}^t = \emptyset$, so that the deviation is profitable in the entire game (this is wlog because we can just “zoom in” on the subgame and treat it as an entire game restricting everything that follows to the subgame). Let $\varepsilon > 0$ be the discounted utility gain from the deviation in absolute (rather than average) terms, i.e. $\varepsilon := \tilde{u}_i(\tilde{\sigma}_i, \sigma_{-i}) - \tilde{u}_i(\sigma)$. Since the stage game is finite,⁸⁰ there exists some (potentially large) $T^* < \infty$ s.t.

$$\frac{\delta^{T^*}}{1 - \delta} \left(\max_{a \in A} \pi_i(a) - \min_{a \in A} \pi_i(a) \right) < \frac{\varepsilon}{2}.$$

This implies that at least half the discounted utility gain from deviating to $\tilde{\sigma}_i$ must accrue to player i by period T^* — for no matter how big the gains in every period after T^* , the discounted value is less than $\frac{\varepsilon}{2}$.⁸¹ Hence, there must exist a profitable deviation $\hat{\sigma}_i$ that differs from σ_i at only a finite number of histories.

⁸⁰ or, more generally, by an assumption that payoffs in the stage game are bounded

⁸¹ There is nothing important about choosing a half. The point is that the gains from the deviation in terms of period-payoffs cannot accrue only too far out in the future, because of discounting. Note that this holds if the game is finitely repeated even if there is no discounting.

Now we use induction to complete the proof. Look at any history, $h^{t'}$, such that there is no possible later history at which $\hat{\sigma}_i$ differs from σ_i . (This is well-defined because $\hat{\sigma}_i$ and σ_i only differ at a finite number of histories.) If we modify $\hat{\sigma}_i$ by only replacing $\hat{\sigma}_i(h^{t'})$ with $\sigma_i(h^{t'})$, is this new strategy still a profitable deviation from σ_i ? There are only two possibilities:

- If the answer is NO, we are done, since we could take σ_i , switch it at only history $h^{t'}$ with $\hat{\sigma}_i(h^{t'})$, and we have created a profitable one-shot deviation.
- If the answer is YES, redefine $\hat{\sigma}_i$ by just replacing the behavior at $h^{t'}$ with $\sigma_i(h^{t'})$, and go back to the inductive step. Eventually, we must hit the previous bullet and would have found a profitable one-shot deviation. \square

Remark 35. That the game is a repeated game is not important for the one-shot deviation principle. The argument applies to various classes of extensive form games that are either finite, or more generally, satisfy a “*continuity-at-infinity*” condition that is ensured by discounting.⁸² The canonical class is referred to as multi-stage games with observable actions: roughly, the only nonsingleton information sets capture simultaneous moves. See [Fudenberg and Tirole \(1991a\)](#).

The one-shot deviation principle is very useful because it means that we can always focus on “simple” deviations when thinking about SPNE. Let me stress that the one-shot deviation principle does *not* apply to Nash equilibria; we’ll see an example later.

6.3. A Basic Result

The one-shot deviation principle allows us to make a simple but useful observation: repetition of stage-game Nash equilibria is a SPNE of a repeated game. To state this precisely, say that a strategy profile σ is *history-independent* if for all h^t and \tilde{h}^t (i.e., any two histories at the same time period), $\sigma(h^t) = \sigma(\tilde{h}^t)$.

Proposition 15. *A history-independent profile, σ , is a SPNE if and only if for each t , $\sigma(h^t)$ (for any h^t) is a Nash equilibrium of the stage game.*

Proof. For the if part, observe that because each $\sigma(h^t)$ is a Nash equilibrium of the stage game, there is no profitable one-shot deviation, and hence no profitable deviation by the one-shot deviation principle. For the only if part, suppose that some $\sigma(h^t)$ is not a Nash equilibrium of the stage game. Then some player i has a profitable deviation where he deviates at h^t but otherwise plays just as σ_i ; since all others are playing history-independent strategies, this is indeed a profitable deviation in the subgame starting at h^t . \square

It is important to emphasize two things about the Proposition: first, it applies to both finitely- and infinitely-repeated games; second, it does not require that the *same* Nash equilibrium be played at every history; only that in any given period, the Nash equilibrium being played not vary across possible histories at that period. So, for example, if we consider a finitely- or infinitely-repeated version of Battle of the Sexes ([Example 16](#)), the following strategies form a SPNE: in every odd period, regardless of history, row plays B and column plays b ; in every even period, regardless of history, they play C and c respectively.

⁸² Here is a (somewhat contrived, but perfectly valid) counter-example illustrating why some condition is needed in infinite games: suppose a single player makes an infinite sequence of decisions of Left or Right. If he chooses Left an infinite number of times, he gets a payoff of 1, otherwise he gets 0. Then there is a profitable deviation from the strategy of playing Right always (e.g., deviate to playing Left always), but there is no profitable one-shot deviation.

Consequently, there are an infinite number of pure strategy SPNE in the infinitely repeated Battle of Sexes, or a large number in the finitely repeated case when there are many periods.

Although [Proposition 15](#) is not terribly exciting, one reason to state it is that it implies existence of SPNE in infinitely repeated games (under our maintained assumptions).

Corollary 7. *If the stage game has a Nash equilibrium,⁸³ then the repeated game has a SPNE.*

6.4. Finitely Repeated Games

[Proposition 15](#) doesn't say anything about strategies that are not history independent. Moreover, in games with a unique stage game Nash equilibrium (such as Matching Pennies or the Trust Game), it doesn't offer existence beyond repetition of the same (possibly mixed) action profile. As it turns out, there are no other SPNE in such cases for finitely repeated games!

Proposition 16. *Suppose $T < \infty$ and the stage game has a unique (possibly mixed) Nash equilibrium, α^* . Then the unique SPNE of the repeated game is the history-independent strategy profile σ^* s.t. for all t and h^t , $\sigma^*(h^t) = \alpha^*$.*

Proof. The argument is by generalized backward induction. In any SPNE, we must have $\sigma^*(h^T) = \alpha^*$ for all h^T . Now consider period $T-1$ with some history h^{T-1} . There are no dynamic incentive considerations, since no matter what players do at this period, they will each get $\pi_i(\alpha^*)$ in the last period. Thus each player must be playing a stage-game best response to others' play in this period. Since α^* is the unique stage-game Nash equilibrium, we must have $\sigma^*(h^{T-1}) = \alpha^*$; otherwise there is a profitable deviation for some player. Induction yields the same conclusion for all periods. \square

The Proposition implies, for instance, that no matter how many periods it is repeated, so long as there are only a finite number of repetitions, the Trust Game has a unique SPNE where each player plays Cheat in every period.

On the other hand, we can have interesting history-*dependent* play in finitely repeated games if there are multiple Nash equilibria of the stage game. In such cases, there can be SPNE in the repeated game where some periods involve action profiles being played that are not Nash equilibria of the stage game. The following example illustrates some general principles about sustaining "good" outcomes in repeated games.

Example 35 (Modified Prisoner's Dilemma). Let $\delta = 1$ (no discounting) and $T = 2$ and consider the following stage game:

	L	C	R
T	10,10	2,8	-5,13
M	8,2	5,5	0,0
B	13,-5	0,0	1,1

There are two PSNE in the stage game: MC and BR. This has a Prisoner's Dilemma flavor because in the stage game, both players would be better off if they could somehow manage to play TL, but the problem is that each has a unilateral deviation from that profile that is profitable (B and R respectively). So the stage game PSNE are pareto-inefficient. From [Proposition 15](#) we know there are at least four history-independent

⁸³ To be clear: this is automatic if the stage game is finite, but the statement is intended to emphasize that the result also applies to infinite stage games that have Nash equilibria.

SPNE of the repeated game, where the outcomes are playing MC in each period, BR in each period, or MC in one period and BR in the other.

But we can actually do better: there is a SPNE of the repeated game where TL is played in the first period. Consider the following strategy profile: TL is played in the first period, and in the 2nd period, $s^2(TL) = MC$ and $s^2(h^2) = BR$ for all $h^2 \neq TL$. Clearly, there are no profitable deviations in the 2nd period alone, since MC and BR are both Nash equilibria of the stage game (why would it not suffice to only note that MC is?). In the first period, a player's expected payoff from following his prescribed strategy (given that the opponent is playing his prescribed strategy) is $10 + 5 = 15$. By deviating from T or L respectively, either player can get at most $13 + 1 = 14$. Hence, there is no profitable deviation.

Now suppose we relax the assumption that $\delta = 1$. Plainly, if δ is close enough to 1, the same logic would apply. How high must δ be in order to sustain a pure-strategy SPNE (PSSPNE) where TL is played in the first period? As already noted, in period 2, following any history, either MC or BR must be played. Intuitively, the best hope of sustaining TL in period 1 is to play the above-specified s^2 in period 2 — this provides the “highest reward” for playing TL in period 1 and the “worst punishment” for playing anything else, subject to the requirement of subgame perfection. So we can sustain TL in the first period of a PSSPNE if and only if $10 + 5\delta \geq 13 + \delta$, or $\delta \geq \frac{3}{4}$.

A related point is that if we added periods, then the requirement on the discount factor becomes less demanding. As an illustration, consider now $T = 3$. If $\delta \geq \frac{3}{4}$, we can sustain outcomes of TL in the first two periods. If $\delta < \frac{3}{4}$, we know from the above analysis that TL cannot be sustained in the 2nd period of a PSSPNE. But how about having it played in just the first period? This can be done in a PSSPNE if and only if $10 + 5\delta + 5\delta^2 \geq 13 + \delta + \delta^2$, or $\delta \geq \frac{1}{2}$.⁸⁴

As a final variation, suppose again $T = 2$, but we now modify the game so that each player has an additional action in the stage game, D (for Destruction) such that the payoffs from DD are $(-x, -x)$, the payoffs from aD for any $a \neq D$ is $(0, -x)$ and symmetrically the payoffs from Da for any $a \neq D$ are $(-x, 0)$. Let $x > 0$. Since D is strictly dominated for both players, it does not change any of the prior analysis for PSSPNE of the repeated game. However, if we turn to Nash equilibria of the repeated game, things are quite different.⁸⁵ TL can be sustained in the first period in a Nash equilibrium so long as $10 + 5\delta \geq 13$, or $\delta \geq \frac{3}{5}$. A NE strategy profile would be as follows: play TL in the first period, and $s^2(TL) = MC$ and $s^2(h^2) = DD$ for all $h^2 \neq TL$. This is a Nash equilibrium when $\delta \geq 3/5$ because given column's play, a best response for row is either to follow the prescribed behavior or to play B in both periods (but not to play B followed by D), and the former is at least as good as the latter if and only if $\delta \geq 3/5$. Notice that when $\delta < 3/5$, the row player, say, would deviate with a *two-stage deviation* — this reiterates that the one-shot deviation principle does not apply to Nash equilibria. Of course, even when $\delta \geq 3/5$, the prescribed strategy profile is not a SPNE because at any $h^2 \neq TL$ (which are off the equilibrium path), players are not playing optimally. Thus there are profitable one-shot deviations according to Definition 34; they just occur at histories that are never reached. \square

Remark 36. Note well that actions which are dominated in the stage-game *can* be played in a SPNE of the repeated game, even on the equilibrium path. They just cannot be played in the last period of the (finitely) repeated game. Indeed, we have demonstrated this point in the example above: why so?

⁸⁴ Showing the “only if” part requires considering various pure strategy profiles and ruling them out as SPNE.

⁸⁵ You can check that in the original game without D , one cannot sustain TL in the first period in a pure strategy Nash equilibrium unless $\delta \geq 3/4$, just as with PSSPNE.

General principles illustrated by the example:

1. What is needed to sustain stage-game outcomes that are not sustainable in a one-shot game is that the ability to “reward” current-period behavior with future behavior, or the flip side of the coin, to “punish” deviations by switching to less-desirable future behavior.
2. The importance of the future — either via high discount factor or the length of the future horizon — is key: the future losses from deviations must outweigh current gains.
3. Looking only at Nash equilibria instead of SPNE allows a greater scope to deter current deviations, since more threats about future play are possible. Although this can be helpful in sustaining cooperative outcomes, it is not satisfactory when these threats are incredible (just as NE that are not subgame perfect are not compelling in simple dynamic games).

Exercise 19. Does [Proposition 16](#) apply to Nash equilibria? That is, does a finitely repeated game have a unique Nash equilibrium if the stage-game has a unique Nash equilibrium? Prove or give a counter-example.

6.5. Infinitely Repeated Games

Now we turn to infinitely repeated games where $T = \infty$. (Keep in mind that $\delta < 1$ throughout now, and when we refer to “payoff” of the infinitely repeated game, we mean the normalized or average discounted payoff.) Our goal is to develop simple versions of a classic result known as the *Folk Theorem*. Let us begin with an example.

Example 36 (Infinitely Repeated PD). Consider an infinitely repeated game where the stage game is the following Prisoner’s Dilemma.⁸⁶

		Player 2	
		C	D
Player 1	C	5, 5	0, 6
	D	6, 0	2, 2

As we already noted, [Proposition 16](#) implies that with only a finite number of repetitions, there is a unique SPNE (no matter the discount factor), with the outcome of DD in every period. In the current infinitely repeated case, we also know that history-independent repetition of DD is a SPNE. But are there other SPNE?

Grim-Trigger: Consider the following pure strategy for each player: in the first period, play C ; in any subsequent period, play C if the history is such that neither player has ever played D before, and play D otherwise. This strategy is known as *grim-trigger*. Under what conditions, if any, is it a SPNE for both players to play grim-trigger? By the one-shot deviation principle, we only need to check that there is no profitable one-shot deviation. In principle, this could be very complicated to check, since there are an infinite number of histories. But given the simple structure of the strategy profile, it turns out to be straightforward. First, observe that there is no profitable one-shot deviation at any history where D has already been played by some player in the past, since such a one-shot deviation would only lower the current period payoff and

⁸⁶ The payoffs are slightly modified from the original Trust Game, but preserving the same qualitative structure — the reason will become clear.

not affect future payoffs.⁸⁷ Thus, we can focus on one-shot deviations at histories where CC has always been played in the past. Wlog, we can now consider a one-shot deviation for player 1 of playing D at period 1 (why?). In so doing, player 1 triggers a switch from perpetual CC to (DC, DD, DD, \dots) . So the deviation is not profitable if and only if:

$$5 \geq (1 - \delta) \left[6 + 2 \frac{\delta}{1 - \delta} \right],$$

or $\delta \geq \frac{1}{4}$. Remarkably, so long as the discount factor is not too low, we have a SPNE where on the equilibrium path there is mutual cooperation in all periods! *The reason we get this stark difference compared to the finitely repeated game is that now there is no specter of the last period hanging over the players.*^{88,89} Mathematically, we have a failure of continuity—specifically, lower hemi-continuity—of the SPNE payoff set at $T = \infty$.

Tit-for-Tat SPNE: Now consider another strategy: play C in the first period; in any subsequent period, play whatever the opponent played in the previous period. This is known as *tit-for-tat*. For mutual play of tit-for-tat to be a SPNE, we must again show that there is no profitable one-shot deviation. Wlog, focus on player 1's deviations. Given the structure of tit-for-tat, whether a one-shot deviation is profitable at some history h^t only depends upon the action profile played in the previous period, a^{t-1} , since this is what determines how the opponent plays in the current period (i.e., what happened in periods $t-2, t-3, \dots$ is irrelevant). So we consider the four possibilities for a^{t-1} . First, suppose $a^{t-1} = CC$ or $t = 1$. In the subgame starting at h^t , not deviating leads to CC in every period on the path of play; a (one-shot) deviation to D leads to DC, CD, DC, CD, \dots . So the constraint is

$$\begin{aligned} 5 &\geq (1 - \delta) [6 + \delta(0) + \delta^2(6) + \delta^3(0) + \delta^4(6) + \dots] \\ &= (1 - \delta) \frac{6}{1 - \delta^2}, \end{aligned}$$

or $\delta \geq \frac{1}{5}$. Next, suppose $a^{t-1} = CD$. Not deviating at h^t leads to DC, CD, DC, \dots whereas deviating to C leads to CC, CC, \dots . So the constraint is

$$(1 - \delta) \frac{6}{1 - \delta^2} \geq 5,$$

or $\delta \leq \frac{1}{5}$. Third, suppose $a^{t-1} = DC$. Not deviating at h^t leads to CD, DC, CD, \dots whereas deviating leads to DD, DD, \dots . So the constraint is

$$(1 - \delta) \frac{6\delta}{1 - \delta^2} \geq 2,$$

⁸⁷ You might wonder why we even had to consider such a deviation, since if players follow the prescribed strategies, there never would be such a history. But it is crucial that we rule out such deviations because we are looking at SPNE (and the one-shot deviation principle relies on subgame perfection).

⁸⁸ It is worth emphasizing at this point that what is important is not that the game literally be infinitely repeated, but rather that there always be a (non-vanishing) possibility that there will be another period, i.e. that “today is not the end”.

⁸⁹ The difference between finitely- and infinitely-repeated games is less stark when there are multiple Nash equilibria of the stage game. See [Benoit and Krishna \(1985\)](#).

or $\delta \geq \frac{1}{2}$. Finally, suppose $a^{t-1} = DD$. Not deviating at h^t leads to DD, DD, ... whereas deviating leads to CD, DC, CD, ... So the constraint is

$$2 \geq (1 - \delta) \frac{6\delta}{1 - \delta^2}$$

or $\delta \leq \frac{1}{2}$. Plainly, there is no δ that can satisfy all the four requirements. Thus, mutual tit-for-tat is not a SPNE in this game, no matter the discount factor. (An exercise below will clarify how generally this point holds.)

Tit-for-Tat NE: On the other hand, is mutual tit-for-tat a Nash Equilibrium? Answering this requires some care, because we cannot appeal to the one-shot deviation principle any longer, so we have to consider all possible deviations. You asked to show in an exercise below that if the opponent is playing tit-for-tat, then one of the three following strategies is a (not necessarily unique) best response for a player: either (i) play tit-for-tat; (ii) play D in every period; or (iii) alternate between D and C , beginning with D in period 1. From this, it follows that tit-for-tat is a best response if and only if

$$5 \geq (1 - \delta) \max \left\{ \frac{6}{1 - \delta^2}, 6 + 2 \frac{\delta}{1 - \delta} \right\},$$

or $\delta \geq \frac{1}{4}$. Thus, mutual tit-for-tat is a NE for all sufficiently high discount factors. (This is not a general property of Prisoner Dilemma games; the exact specification of the payoff matrix matters, and for some specifications mutual tit-for-tat is not a NE for any δ .⁹⁰)

Other SPNE: So far we have seen that for sufficiently high discount factors, we can achieve a payoff profile in SPNE of the repeated game that is equal to the efficient payoff profile (5, 5) of the stage game, and also one equal to the stage game Nash equilibrium payoff profile (2, 2). But we can also achieve various other payoff profiles. For example, consider a strategy that modifies grim-trigger as follows: play C in the first period; in any even period, play D ; in any odd period > 1 , play D if either player ever played D in a prior odd period, otherwise play C . One can show that mutual play of this strategy is a SPNE if δ is sufficiently large. In this SPNE, each player gets a payoff of $(1 - \delta) (5 + 2\delta + 5\delta^2 + 2\delta^3 + \dots) = (1 - \delta) \left(\frac{5}{1 - \delta^2} + \frac{2\delta}{1 - \delta^2} \right) = \frac{5 + 2\delta}{1 + \delta}$. As you would expect, this converges to 3.5 as $\delta \rightarrow 1$, which is the simple average of the stage-game payoffs from CC and DD . \square

Exercise 20. Suppose the stage game prisoner's dilemma has different payoffs: $u(C, C) = (3, 3)$, $u(D, D) = (1, 1)$, $u(C, D) = (0, 4)$ and $u(D, C) = (4, 0)$. For what discount factors (if any) is tit-for-tat a SPNE in the infinitely repeated game?

Exercise 21. In [Example 36](#), show that the modified grim-trigger profile described at the end is a SPNE for all discount factors high enough, and identify the minimum discount factor needed.

Exercise 22. Show that in [Example 36](#), if player 2 plays tit-for-tat, a best response for player 1 is either (i) play tit-for-tat; (ii) play D in every period; or (iii) alternate between D and C , beginning with D in period 1. [Hint: first consider the payoff to player 1 from using any strategy such that the path of play is CC in every period. Then argue that if 1's best response is a strategy such that he takes action D in some period, then either strategy (ii) or (iii) is a best response.]

⁹⁰ For example, if you replace the DC and CD payoffs by 10,0 and 0,10 respectively, then you can check that the incentive constraint we just derived cannot be satisfied for any δ ; hence mutual tit-for-tat would not be a NE. (This is the reason I changed payoffs from the original Trust game.)

6.5.1. Folk Theorems

The natural question raised by the discussion in [Example 36](#) is: what are all the payoff profiles that can be achieved by a SPNE of an infinitely repeated game? It is this question that various *folk theorems* provide answers to. We will state two such theorems, starting with the more straightforward one. We need two definitions.

Definition 35 (Feasible Payoffs). Let $\tilde{V} := \{(\pi_1(a), \dots, \pi_I(a))\}_{a \in A}$ be the set of all payoff vectors that are attained by some action profile of the stage game. The set of *feasible payoffs*, V , is defined as $V := \text{co}(\tilde{V})$, i.e., the convex hull of \tilde{V} .

To interpret this, draw the set of feasible payoffs for [Example 36](#). The reason we are interested in the convex hull of the payoffs from action profiles is that as $\delta \rightarrow 1$, any payoff that is in the convex hull can be obtained (ignoring any equilibrium or incentive issues for now, just as a matter of “technology”) by having players play an appropriate sequence of (possibly time-varying) action profiles. Indeed, recall the last construction in [Example 36](#). The general point was made by [Sorin \(1986\)](#).

It is intuitive that no payoff vector *outside* the set of feasible payoffs can be achieved as the average discounted payoffs in a SPNE (or Nash Equilibrium, for that matter) of an infinitely repeated game. But is any feasible payoff vector supportable in SPNE? It is not hard to see that the answer is no (can you provide an example?). But it turns out that “almost everything of interest” in V — in the sense of economic interest, not mathematical — with sufficiently patient players.

Definition 36 (Nash-threat Payoff). Any player i ’s Nash-threat payoff is

$$\underline{v}_i := \inf\{v_i : \exists \text{ stage-game (possibly mixed) Nash equilibrium } \alpha \text{ s.t. } \pi_i(\alpha) = v_i.\}$$

Theorem 11 (Nash-threats Folk Theorem). *Fix a stage game and pick any $v \in \mathbb{R}^I$ such that $v \in V$ and for all i , $v_i > \underline{v}_i$. There is $\underline{\delta} \in [0, 1)$ such that for all $\delta > \underline{\delta}$, there is a SPNE of the infinitely repeated game with average discounted payoff profile v .*

Proof. To simplify the proof, we will make two assumptions that can be dispensed with:

- at the start of each period, players observe the realization of a (sufficiently rich) public randomization device that allows them to correlate their strategies.⁹¹ This means that in any period, they can play any mixed action profile in $\Delta(A)$, as opposed to just mixed action profiles in $\Delta(A_1) \times \dots \times \Delta(A_I)$.
- at the end of each period, players observe not just the realization of the mixed action profile being played that period, but the actual mixture itself. This means that if a player is supposed to mix over actions in a particular way (potentially degenerately) at some history but deviates to some other mixture, this will be observed by everyone.

Since $v \in V$ is a feasible payoff vector, there is some $\alpha^* \in \Delta(A)$ such that $\pi(\alpha^*) = v$. Now consider the following strategy profile:

⁹¹ Formally, let $\{\omega^1, \dots, \omega^t, \dots\}$ be a sequence of independent draws from a uniform distribution on $[0, 1]$. We expand the notion of a history as follows: for each $t \geq 1$, $h^t = (a^1, \dots, a^{t-1}, \omega^1, \dots, \omega^t)$. Thus, players are able to condition their behavior at t on not only the history of actions but also the history of (and current) realizations of the randomization device.

1. In the first period, play α^* .
2. At any history where α^* was played in every prior period, play α^* .
3. At any history h^t where some $\alpha \neq \alpha^*$ was played in some prior period, let $t' := \min\{\tilde{t} : \alpha^{\tilde{t}} \neq \alpha^*\}$ and let $j := \min\{i : \alpha_i^{t'} \neq \alpha_i^*\}$. That is, j is a “first deviator”. Let α be the stage-game Nash equilibrium (possibly mixed) such that $\pi_j(\alpha) = \underline{v}_j$. (If there are multiple such stage-game Nash equilibria, any can be used, but pick the same one every time.) At h^t , players play α .

Let’s argue that this strategy profile is a SPNE. On the path of play, α^* is played every period. Observe that once there is a deviation in some period, call it t , players are just repeating the same stage-game Nash equilibrium profile regardless of what happens in any period following t . Thus, by the logic of [Proposition 15](#), there are no profitable deviations in any subgame following the first deviation. So it suffices to argue that a unilateral first deviation is not profitable. The first deviator, call him j , can gain at most some finite amount of period-utility in the period he deviates. But in all future periods, he foregoes $v_j - \underline{v}_j > 0$. Let $d := \min_i \{v_i - \underline{v}_i\} > 0$. Since $\lim_{\delta \rightarrow 1} \frac{\delta}{1-\delta} d = \infty$, no deviation is profitable for large enough δ . \square

The strategy profile used in the proof is known as *Nash Reversion*, since the key idea is to punish a deviator by just playing a stage-game Nash equilibrium in every period thereafter that gives him the lowest payoff amongst all stage-game Nash equilibria. Note the parallel with grim-trigger in the repeated Prisoner’s Dilemma. In a homework problem, you will apply this idea to a simple model of repeated market power.

Nash reversion is a very intuitive way to punish a player for deviating from the desired action profile. It turns out, however, that there may be more “severe” punishment schemes than Nash reversion. The following example illustrates.

Example 37 (Minmax Punishments). Consider the following stage-game:

		Player 2	
		C	D
Player 1	C	3, 3	0, 4
	D	4, 1	1, 0

Observe that player 1 has a strictly dominant strategy of D , whereas player 2’s best response is to “mismatch” with player 1’s action. So the unique Nash equilibrium of the stage game is DC .

In the infinitely repeated game, can we find a SPNE where CC is played every period with high enough discount factor? Nash reversion is of no help, since player 1 prefers the the unique Nash equilibrium of the stage game to CC . Nevertheless, one can do it (and you are asked to, below!). Note that it is easy to sustain play of CC in a Nash Equilibrium (with patient players); what makes the question interesting is the requirement of subgame perfection. \square

Exercise 23. For the above example, for high enough discount factors, construct a SPNE where CC is played in every period on the equilibrium path.

To generalize the example, make the following definition.

Definition 37 (Individually Rational Payoffs). A player i ’s *minmax* payoff of the stage game is given by:

$$\underline{v}_i := \min_{\alpha_{-i}} \max_{\alpha_i} \pi_i(\alpha_i, \alpha_{-i}).$$

A vector of payoffs, $v = (v_1, \dots, v_I)$ is *strictly individually rational* if for all i , $v_i > \underline{\underline{v}}_i$.

You’ve seen the idea of minmax in a homework problem a while back. A player’s minmax payoff is the lowest payoff his opponents can “force” him down to in the stage game so long as he plays an optimal response to his opponents’ play. Thus, a player will obtain at least his minmax payoff in a Nash equilibrium of the stage game.⁹² It is in this sense that no payoff below the minmax payoff is “individually rational.” But then, in any Nash equilibrium—indeed, any rationalizable strategy profile—of the infinitely repeated game, a player’s average discounted payoff cannot be lower than his minmax payoff either, for he can assure at least his minmax payoff by just playing a strategy where at each history he myopically best responds to what his opponents are doing at that history. Our final result of this section is that essentially any vector of feasible and strictly individually rational payoffs can be obtained as a SPNE of an infinitely repeated game.

Theorem 12 (Minmax Folk Theorem). *Fix a (finite) stage game and let $V^* \subseteq V$ be the set of feasible and strictly individually rational payoffs. Assume the interior of V^* is nonempty.⁹³ For any $v \in V^*$, there is $\underline{\delta} \in [0, 1)$ such that for all $\delta > \underline{\delta}$ there is a SPNE of the infinitely repeated game with average discounted payoff profile v .*

A full proof is somewhat involved, but we can illustrate the key idea how minimax rather than Nash reversion can be used as a credible threat.

Proof (Partial). Let us assume there are only two players, that the payoff vector we seek, v , is the payoff vector from some pure action profile, a^* (i.e., $v = \pi(a^*)$), and that $v_i > \underline{\underline{v}}_i^P := \min_{a_{-i}} \max_{a_i} \pi_i(a_i, a_{-i})$. The last assumption means that each player is getting strictly more than his *pure strategy minmax*, not just his individually rational payoff (cf. fn. 92). Let a_i^m be some action of i that holds his opponent j to j ’s pure strategy minmax payoff, and consider the following *stick-and-carrot strategy*:

1. Play a_i^* initially or if a^* was played in previous period.
2. If there was a deviation from (1), play a_i^m K times and then restart (1).
3. If there is a deviation from (2), begin (2) again.

The length K of phase 2 (the “punishment phase”) will be determined momentarily. Let $\bar{v}_i := \max_a \pi_i(a)$ be the maximum stage payoff for i , and let $v_i^m := \pi_i(a_i^m, a_{-i}^m)$ be i ’s stage payoff in phase (2). Note that

$$v_i^m \leq \underline{\underline{v}}_i^P < v_i \leq \bar{v}_i.$$

Pick any integer $K > 0$ such that for both players i , $K(v_i - v_i^m) > \bar{v}_i - v_i$.

All we need to check is that there is no profitable one-shot deviation from phase 1 or phase 2.

⁹² Note that the definition above explicitly considers mixed strategies for both player i and his opponents. This is not innocuous. For example, if you consider simultaneous Matching Pennies where the utility from matching is 1 and mismatching is 0 for player 1, say, then the minmax payoff for player 1 is 0.5 according to our definition; whereas if we only consider pure strategies, the “pure strategy minmax” payoff would be 1.

⁹³ The interior is relative to I -dimensional Euclidean space.

There is no profitable one-shot deviation from phase 1 of the strategy profile if

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t v_i \geq (1 - \delta) \left(\bar{v}_i + \sum_{t=1}^K \delta^t v_i^m + \sum_{t=K+1}^{\infty} \delta^t v_i \right)$$

$$\iff \sum_{t=1}^K (v_i - v_i^m) \geq \bar{v}_i - v_i.$$

Since the LHS $\rightarrow K(v_i - v_i^m)$ as $\delta \rightarrow 1$, this inequality holds for all $\delta < 1$ large enough because of our choice of K above. Intuitively, a deviation from phase 1 triggers phase 2, in which a player will get a per-period payoff no higher than his pure strategy minmax, which is strictly less than his payoff in phase 1; if phase 2 is long enough then a patient player will not find this profitable.

In phase 2, deviation is most tempting in the initial period, as that is when there are the most remaining punishment periods. A one-shot deviation in the initial period is not profitable if

$$(1 - \delta) \left(\sum_{t=0}^{K-1} \delta^t v_i^m + \sum_{t=K}^{\infty} \delta^t v_i \right) \geq (1 - \delta) \left(\underline{v}_i^P + \sum_{t=1}^K \delta^t v_i^m + \sum_{t=K+1}^{\infty} \delta^t v_i \right)$$

$$\iff v_i^m + \delta^K v_i \geq \underline{v}_i^P + \delta^K v_i^m$$

$$\iff \delta^K v_i + (1 - \delta^K) v_i^m \geq \underline{v}_i^P.$$

The last inequality holds when δ is large enough because $v_i > \underline{v}_i^P$. Intuitively, given the specification of phase 3, the effect of a deviation in the initial period of phase 2 is to swap a future period's payoff of v_i with a payoff of \underline{v}_i^P today; when a player is patient, this is not profitable.⁹⁴

The above logic shows why phase 2 is the “stick” that enforces phase 1, while at the same time phase 1 acts as a “carrot” for carrying out phase 2. This carrot is used to ensure that a player carrying out the punishment of his opponent even when it is not in her myopic interest to do so. \square

See [Mailath and Samuelson \(2006, p. 101\)](#) for a full proof. The general idea remains that a player who deviates from the equilibrium path is minmaxed by his opponents. To induce the opponents to carry out the minmaxing, they must be given a small reward once the punishment phase is completed. The assumption that the interior of V^* is nonempty ensures that rewards/punishment incentives can be provided to each player independently of the others. In particular, the condition rule out any pair of players having identical payoff functions.⁹⁵

We can apply the Theorem to [Example 37](#) by deducing what the set V^* is there. (Draw a picture.) We see immediately that the payoff profile from $CC(3, 3)$ is supportable in SPNE as $\delta \rightarrow 1$, which you were asked to prove earlier with an explicit equilibrium construction.

Remark 37 (Comments on Folk Theorems). [Theorem 11](#) and [Theorem 12](#) show that standard equilibrium requirements do very little to narrow down predictions in infinitely repeated games:⁹⁶ in terms of payoffs,

⁹⁴ Note that if we fix $\delta < 1$, then the last inequality places an upper bound on K . But that upper bound diverges to ∞ as $\delta \rightarrow 1$. We have fixed an arbitrary large enough K , and we are then taking δ large.

⁹⁵ In fact, instead of assuming V^* has a nonempty interior, the following weaker “nonequivalent utilities” condition suffices ([Abreu, Dutta, and Smith, 1994](#)): either there are two players, or no two players have identical preferences in the sense that there are no $i, j \in I$ ($i \neq j$) and $a > 0$, $b \in \mathbb{R}$ with $\pi_i(\cdot) = a\pi_j(\cdot) + b$.

⁹⁶ There are also folk theorems for finitely repeated games when there are multiple Nash equilibria of the stage

more or less “anything can happen” as $\delta \rightarrow 1$ (subject to being feasible and strictly individually rational). Some comments:

1. This can make comparative statics difficult.
2. With repeated interactions, players are able to get around inabilities to write binding contracts (at least in the perfect monitoring environment we have been considering). Anything achievable through a binding contract can also be attained as a non-binding equilibrium.
3. Which equilibrium gets played could be thought of determined by some kind of pre-play negotiation amongst players. It is natural to then think that the payoff profile will be efficient, and which efficient profile is selected may depend on bargaining powers. (There is not much formal work on this yet.)
4. Although the Theorems only discuss SPNE payoffs, one can say a fair bit about what strategies can be used to attain these payoffs. An important technique is known as *self-generation*; see [Abreu, Pearce, and Stacchetti \(1990\)](#).
5. Although they yield multiplicity of equilibria, folk theorems are important because of what they tell us about *how* payoffs can be achieved, what the qualitative structure of reward and punishment is (more relevant under imperfect monitoring, which we haven’t studied here), and what the limitations are.
6. In applications, people usually use various refinements to narrow down payoffs/equilibria, such as efficiency, stationarity, Markov strategies, symmetric strategies, etc.
7. It is also important to understand what payoffs are supportable for a given δ (not close to 1). There is some work on this.

game, see for example [Benoit and Krishna \(1985\)](#).

7. Signaling and Cheap Talk

Dynamic games of incomplete information are particularly rich in application. In this section, we study a class of such games that involve *asymmetric information*, viz. situations where some players hold private information that is valuable to other players.

7.1. Costly Signaling

We first consider the classic idea of [Spence \(1973\)](#). We'll frame it in the context of labor markets, as he did, but it is important to stress that the same principle has been applied in a host of other contexts.

7.1.1. The Setting

There is a worker and two firms.⁹⁷ For simplicity, we study here the canonical case where there are only two types of workers, θ_H and θ_L , with $\theta_H > \theta_L > 0$. Let $\lambda = \text{Prob}(\theta = \theta_H) \in (0, 1)$ be the fraction of high type workers. A worker's type is his private information. Prior to entering the labor market, workers can obtain education at some cost. In particular, a worker of type $\theta \in \{\theta_L, \theta_H\}$ can obtain education level $e \geq 0$ at a cost $c(e, \theta)$. We assume that $c_1(e, \theta) > 0$ and $c_{12}(e, \theta) < 0$. In words, the first condition says that acquiring more education is always costly on the margin; the second condition says that the marginal cost is lower for the higher type worker. It is critical that firms observe a worker's education level after he has acquired it (but they do not observe his productivity type). We set the reservation wage for both types to be 0; the important assumption here is that it is the same for both types, but otherwise 0 is a normalization. The game that the worker then plays with the firms is as follows:⁹⁸

1. Nature chooses worker's type, $\theta \in \{\theta_L, \theta_H\}$, according to $\lambda = \text{Prob}(\theta = \theta_H)$.
2. Having privately observed θ , worker chooses e .
3. Having observed e but not θ , each firm $i \in \{1, 2\}$ simultaneously offers a wage, w_i .
4. Worker accepts one or neither job.

Payoffs are as follows: the worker gets $u(w, e, \theta) = w - c(e, \theta)$ if she accepts an offer and $-c(e, \theta)$ if she does not; the firm that employs the worker gets $\theta - w$; the other firm gets 0. Note that as specified, education is absolutely worthless in terms of increasing productivity — it is solely an instrument to potentially signal some private information.⁹⁹ This is known as *purely dissipative* signaling.

⁹⁷ We could also do this more workers and more firms. Two firms is sufficient because when firms compete for workers through their wage offerings, Bertrand competition between any two firms is sufficient to drive wages up to the marginal product of labor.

⁹⁸ You can think of the multiple workers case as simply each worker playing this game simultaneously with the firms, with types drawn independently across workers.

⁹⁹ Of course, in practice, education also has a productivity-enhancing purpose ... we hope. Even in that case, it can serve as a signaling instrument. As you'll see, what is important is the difference in marginal costs of acquiring education for the different worker types.

7.1.2. Basic Properties

The goal is to study whether and how education can be used by workers to signal their type to the firms. In what follows, we are going to analyze *pure strategy weak PBE* that satisfy the following additional property: for any level of education chosen by the worker, e , both firms have the same belief over the worker's type. That is, if $\mu_i(e)$ represents the belief of firm $i \in \{1, 2\}$ that the worker is type θ_H given that he has chosen education e , then we require that $\mu_1(e) = \mu_2(e)$ for all e .¹⁰⁰ It turns out that the set of such weak PBE is identical to the set of sequential equilibria in this model (strictly speaking, in a discretized version of the model, since sequential equilibrium is only defined for finite games); see [Fudenberg and Tirole \(1991b\)](#) for a more general equivalence of (a version of) PBE and sequential equilibria in (a class of) signaling games. For short, I will just refer to any of these weak PBE as an equilibrium in the ensuing discussion.

Using generalized backward induction, we start analyzing play at the end of the game.

Stage 4: In the last stage, sequential rationality requires that a worker accepts a job offer from the firm that offers the higher wage, so long as it is non-negative.¹⁰¹ If they both offer the same wage, he randomly accepts one of them.

Stage 3: Given Stage 4 behavior of the worker, we can see that for any belief $\mu(e)$, both firms must offer the wage

$$w(e) = \mathbb{E}[\theta|\mu(e)] = \mu(e)\theta_H + (1 - \mu(e))\theta_L. \quad (13)$$

To see this, note that a firm expects strictly negative payoffs if it hires the worker at a wage larger than $\mathbb{E}[\theta|\mu(e)]$, and strictly positive payoffs if it hires him at less than $\mathbb{E}[\theta|\mu(e)]$. Since firms are competing in Bertrand price competition, the unique sequentially rational best responses are $w(e)$ as defined above.

What more can be said about $w(e)$? At this stage, not much, except that in any equilibrium, for all e , $w(e) \in [\theta_L, \theta_H]$ because $\mu(e) \in [0, 1]$. Note that in particular, we cannot say that $w(\cdot)$ even need be increasing. Another point to note is that there is a 1-to-1 mapping from $\mu(e)$ to $w(e)$. Remember that equilibrium requires $\mu(e)$ to be derived in accordance with Bayes rule applied to the worker's on-path education choices.

Stage 2: To study the worker's choice of education, we must consider her preferences over wage-education pairs. To this end, consider the utility from acquiring education e and then receiving a wage w , for type θ : $u(w, e, \theta) = w - c(e, \theta)$. To find indifference curves, we set $u(w, e, \theta) = \bar{u}$ (for any constant \bar{u}), and implicitly differentiate, obtaining

$$\left. \frac{dw}{de} \right|_{u=\bar{u}} = c_1(e, \theta) > 0.$$

Thus indifference curves are upward sloping in e - w space, and moreover, at any particular (e, w) , they are steeper for θ_L than for θ_H by the assumption that $c_{12}(e, \theta) < 0$. Thus, an indifference curve for type θ_H crosses an indifference curve for type θ_L only once (so long as they cross at all). This is known as the (*Spence-Mirlees*) *single crossing property*, and it plays a key role in the analysis of many signaling models. [Figure 9](#) shows a graphical representation.¹⁰²

¹⁰⁰ MWG call this a PBE, and indeed, in accordance with our terminology, any weak PBE equilibrium with commonality of beliefs will be subgame perfect. However, subgame perfection does not require the commonality of beliefs; this is an added restriction, albeit a natural one.

¹⁰¹ Strictly speaking, he can reject if the wage is exactly 0, but we resolve indifference in favor of acceptance, for simplicity. This is not important.

¹⁰² Here is another way to state the requisite single-crossing property in a manner that is portable to many types:




Figure 9 – The Single-Crossing Property

Obviously, the choice of e for a worker of either type will depend on the wage function $w(e)$ from Stage 3. But in turn, the function $w(e)$ (or equivalently, $\mu(e)$) must be derived from Bayes rule for any e that is chosen by the worker (of either type). For any e not chosen by either type, any $w(e) \in [\theta_L, \theta_H]$ is permissible since Bayes rule doesn't apply. This flexibility in specifying “off-the-equilibrium-path” wages yields a multiplicity of equilibria. Equilibria can be divided into two classes:

1. Separating equilibria. Here, the two types of worker choose different education levels, thus “separating” themselves.
2. Pooling equilibria. Here, the two types choose the same education level, thereby “pooling” together.

We study each of them in turn.¹⁰³

7.1.3. Separating Equilibria

Let $e^*(\theta)$ denote a worker's equilibrium education choice, and $w^*(e)$ denote an equilibrium wage offer given the equilibrium beliefs $\mu^*(e)$.

Claim 2. *In any separating equilibrium, $w^*(e^*(\theta_i)) = \theta_i$ for $i \in \{L, H\}$, i.e., a worker is paid her marginal product.*

Proof. By definition, in a separating equilibrium, the two types choose different education levels, call them $e^*(\theta_L) \neq e^*(\theta_H)$. Bayes rule applies on the equilibrium path, and implies that $\mu(e^*(\theta_L)) = 0$ and $\mu(e^*(\theta_H)) = 1$. The resulting wages by substituting into equation (13) are therefore θ_L and θ_H respectively. \square

$\forall (w, e) > (w', e')$, where $>$ is in the vector sense, and $\forall \theta > \theta'$, $u(w, e, \theta') \geq u(w', e', \theta') \implies u(w, e, \theta) > u(w', e', \theta)$. Given the separability $u(w, e, \theta) = w - c(e, \theta)$, it is sufficient that c satisfy decreasing differences, which under smoothness of c can be stated as $c_{e\theta}(e, \theta) < 0$.

¹⁰³ The two classes are exhaustive given the restriction to pure strategies. If we were to consider mixed strategies, there would be a third class of *partial pooling* or *hybrid* equilibria where types could be separating with some probability and pooling with some probability.




Figure 10 – Separating Equilibrium Low Type’s Allocation

Claim 3. *In any separating equilibrium, $e^*(\theta_L) = 0$, i.e., the low type worker chooses 0 education.*

Proof. Suppose towards contradiction that $e^*(\theta_L) > 0$. By Claim 2, type θ_L ’s equilibrium utility is $\theta_L - c(e^*(\theta_L), \theta_L)$. If she instead chose education level 0, she would receive a utility of at least $\theta_L - c(0, \theta_L)$, because $w^*(0) \geq \theta_L$. Since $c_1(\cdot, \theta_L) > 0$, it follows that the worker gets a strictly higher utility by deviating to an education level of 0, a contradiction with equilibrium play. \square

Claim 2 and Claim 3 combined imply that the equilibrium utility for a low type is $u(\theta_L, 0, \theta_L)$. This puts the low type on the indifference curve passing through the point $(0, \theta_L)$ in e - w space, as drawn in Figure 10.

We can use this picture to construct a separating equilibrium. By Claim 2, the θ_H worker must receive a wage of θ_H , hence an allocation somewhere on the horizontal dotted line at θ_H . If the allocation were to the left of where $IC(\theta_L)$ crosses that dotted line, then type θ_L would prefer to mimic θ_H worker rather than separate (i.e. it would prefer to choose the education level that θ_H is supposed to, rather than 0); this follows from the fact that allocations to the “left” of a given indifference curve are more desirable. So, a candidate allocation for the high type is education level \hat{e} with wage θ_H in Figure 10. That is, we set $e^*(\theta_H) = \hat{e}$ and $w^*(\hat{e}) = \theta_H$, where \hat{e} is formally the solution to

$$u(\theta_H, \hat{e}, \theta_L) = u(\theta_L, 0, \theta_L).$$

It remains only to specify the wage schedule $w^*(e)$ at all points $e \notin \{0, \hat{e}\}$. Since we are free to specify any $w^*(e) \in [\theta_L, \theta_H]$, consider the one that is drawn in Figure 11.

Given this wage schedule, it is clear that both types are playing optimally by choosing 0 and \hat{e} respectively, i.e., neither type strictly prefers choosing any other education level and receiving the associated wage over its prescribed education and associated wage. Beliefs (or wages) are correct on the equilibrium path, and thus firms are playing optimally. Thus, e^* and w^* as defined is in fact a separating equilibrium. It is obvious that there are various wage schedules that can support the same equilibrium education choices: an alternate schedule that works, for example, is $w(e) = \theta_L$ for all $e \in [0, \hat{e})$ and $w(e) = \theta_H$ for all $e \geq \hat{e}$.




Figure 11 – A Separating Equilibrium

The more interesting question is whether there are other education levels that can be sustained in a separating equilibrium. [Claim 3](#) says that the low type must always play $e(\theta_L) = 0$, but could we vary $e^*(\theta_H)$? Yes. Let \bar{e} be the education level that solves

$$u(\theta_H, \bar{e}, \theta_H) = u(\theta_L, 0, \theta_H).$$

In words, \bar{e} is the education level that makes type θ_H indifferent between acquiring \bar{e} with pay θ_H and acquiring 0 with pay θ_L . The single-crossing property stemming from $c_{12}(\cdot, \cdot) < 0$ ensures that $\bar{e} > \hat{e}$. A separating equilibrium where $e^*(\theta_H) = \bar{e}$ is illustrated in [Figure 12](#).

It follows from the construction's logic that for every $e \in [\hat{e}, \bar{e}]$, there is a separating equilibrium with $e^*(\theta_H) = e$; and there is no separating equilibrium with $e^*(\theta_H) \notin [\hat{e}, \bar{e}]$. It is easy to see that we can Pareto-rank these separating equilibrium allocations.

Proposition 17. *A separating equilibrium with $e^*(\theta_H) = e_1$ Pareto-dominates a separating equilibrium with $e^*(\theta_H) = e_2$ if and only if $e_1 < e_2$.*

Proof. Straightforward, since firms are making 0 profit in any separating equilibrium, the low type of worker receives the same allocation in any separating equilibrium, and the high type of worker prefers acquiring less education to more at the same wage. \square

So the first separating equilibrium we considered with $e^*(\theta_H) = \hat{e}$ Pareto-dominates all others with different high-type allocations; it has the least inefficiency via the costly signaling. This outcome—note that because there the flexibility with off-path beliefs means that there are many equilibria leading to the same outcome or on-path behavior—is referred to as the *least-cost separating equilibrium outcome* or the *Riley outcome* after [Riley \(1979\)](#).




Figure 12 – A Separating Equilibrium with $e^*(\theta_H) = \bar{e}$

7.1.4. Pooling Equilibria

Now we turn to the other class of equilibria, where both types of worker choose the same education level, thus pooling together. That is, we are looking for equilibria where $e^*(\theta_H) = e^*(\theta_L) = e^P$ for some e^P . In any such equilibrium, Bayes rule implies that $\mu^*(e^P) = \lambda$, hence from equation (13), $w^*(e^P) = \mathbb{E}(\theta)$. To determine which e^P are feasible in a pooling equilibrium, define \underline{e} as the education level that makes type θ_L indifferent between acquiring education \underline{e} with wage $\mathbb{E}(\theta)$ and acquiring education 0 with wage θ_L . Formally, \underline{e} is the solution to

$$u(\mathbb{E}(\theta), \underline{e}, \theta_L) = u(\theta_L, 0, \theta_L)$$

Since $\mathbb{E}(\theta) \in (\theta_L, \theta_H)$, it follows that $\underline{e} \in (0, \hat{e})$. Figure 13 shows the construction of a pooling equilibrium with $e^P = \underline{e}$. Of course, there are multiple wage schedules that can support this pooling choice of $e^P = \underline{e}$.

Moreover, the logic of this construction implies that there is a pooling equilibrium for any $e^P \in [0, \underline{e}]$; but not for any $e^P > \underline{e}$. The reason for the latter is that a worker of type θ_L would strictly prefer to choose education 0 and get $w^*(0) \geq \theta_L$ rather than choose $e^P > \underline{e}$ and get wage $\mathbb{E}(\theta)$. We can also Pareto-rank the pooling equilibria.

Proposition 18. *A pooling equilibrium with education level e^P Pareto-dominates a pooling equilibrium with education level \tilde{e}^P if and only if $e^P < \tilde{e}^P$.*

Proof. Straightforward, since firms are making 0 profit (in expectation) in any pooling equilibrium, both types of worker receive the same wage in any pooling equilibrium, and both type of workers strictly prefer acquiring lower education levels for a given wage. \square

Note also that any pooling equilibrium is completely wasteful in the sense both types of worker would be better off if the ability to signal had been absent altogether, and the market just functioned with no education acquisition and a wage rate of $\mathbb{E}(\theta)$. Contrast this with separating equilibria, where at least the high type is able to reap some benefit from signaling in terms of a higher wage (though it may not compensate him for the cost of signaling relative to the absence of the signaling altogether).




Figure 13 – A Pooling Equilibrium

7.1.5. Equilibrium Refinement

The fact that we have been able to construct a continuum of equilibrium allocations in both separating and pooling classes is somewhat troublesome. To put it another way, is there a reason why some of these equilibria should be thought of as more “reasonable” than others? The latitude in selecting $w(e)$ for all e that are not chosen in equilibrium is what leads to the multiplicity of equilibria. Based on the culminating discussion in [Section 4](#), you might guess that we have to turn to forward induction arguments. No time like the present, so let’s do so without further ado!

The refinement we analyze is called *equilibrium dominance* or the *intuitive criterion* and the ideas are due to [Cho and Kreps \(1987\)](#) and [Banks and Sobel \(1987\)](#).¹⁰⁴

Definition 38 (Equilibrium Dominance). A signaling equilibrium, $(e^*(\theta), w^*(e))$ and beliefs $\mu^*(e)$, satisfies the equilibrium dominance condition if $\mu^*(\tilde{e}) = 1$ for any \tilde{e} such that

1. \tilde{e} is not chosen by either type in the equilibrium;
2. for all $w \in [\theta_L, \theta_H]$, $u(w^*(e^*(\theta_L)), e^*(\theta_L), \theta_L) > u(w, \tilde{e}, \theta_L)$;
3. for some $w \in [\theta_L, \theta_H]$, $u(w^*(e^*(\theta_H)), e^*(\theta_H), \theta_H) < u(w, \tilde{e}, \theta_H)$.

What does the equilibrium dominance condition require? Condition 2 is the key. It says that type θ_L gets strictly higher utility in the equilibrium than any $w \in [\theta_L, \theta_H]$ it could get in return for choosing the out-of-equilibrium education \tilde{e} . Condition 3 says that there is some $w \in [\theta_L, \theta_H]$ that would make type θ_H prefer acquiring \tilde{e} if it received w in return, relative to what it gets in equilibrium. Note that it is sufficient to check this condition using the most attractive wage, i.e., $w = \theta_H$.

¹⁰⁴ Strictly speaking, there is a difference between the equilibrium dominance criterion and the intuitive criterion ([Cho and Kreps, 1987](#), Section 3). In general, the former is stronger (i.e., more restrictive) because, vaguely speaking, they turn on the same principle but the former entails iterated applications. The criteria are, however, equivalent in our environment with two types. I use the term “equilibrium dominance” because it is more informative than “intuitive criterion”.

The intuition behind this condition is quite simple: since type θ_L can only do worse by playing \tilde{e} relative to what it gets in equilibrium (so long as it gets in response some $w \in [\theta_L, \theta_H]$ — sequential rationality is important), if we do observe the choice of \tilde{e} , we ought to rule out the possibility that it was made by θ_L . Moreover, there is some response $w \in [\theta_L, \theta_H]$ that justifies θ_H choosing \tilde{e} . Thus, we must infer that the deviation to \tilde{e} was made by θ_H , hence we must put $\mu^*(\tilde{e}) = 1$, or equivalently $w^*(\tilde{e}) = \theta_H$. (This is a good time to go back and look at [Example 33](#) and the logic we discussed there, it's very similar.)

It turns out that this equilibrium dominance condition is very powerful in signaling games with 2 types of the privately informed player. Let's apply it to the current model.

Proposition 19. *The only signaling equilibria (amongst both pooling and separating) that satisfy the equilibrium dominance condition are the separating equilibria with $e^*(\theta_H) = \hat{e}$.*

Proof. For any separating equilibrium with $e^*(\theta_H) > \hat{e}$, consider any out-of-equilibrium $\tilde{e} \in (\hat{e}, e^*(\theta_H))$. It can be seen (graphical argument) that Conditions 2 and 3 of [Definition 38](#) are met. So we must have $\mu^*(\tilde{e}) = 1$, or equivalently, $w^*(\tilde{e}) = \theta_H$. But then type θ_H has a profitable deviation to \tilde{e} .

For any pooling equilibrium with education level e^P (recall $e^P \leq \underline{e} < \hat{e}$), define \check{e} by $u(\mathbb{E}(\theta), e^P, \theta_L) = u(\theta_H, \check{e}, \theta_L)$. Note that $\check{e} \in (e^P, \hat{e})$, and because of the single-crossing property, $u(\mathbb{E}(\theta), e^P, \theta_H) < u(\theta_H, \check{e}, \theta_H)$. Consider the out-of-equilibrium $\tilde{e} = \check{e} + \varepsilon$ for a small $\varepsilon > 0$ such that $u(\mathbb{E}(\theta), e^P, \theta_H) < u(\theta_H, \tilde{e}, \theta_H)$. It can be seen (graphical argument) that Conditions 2 and 3 of [Definition 38](#) are met. So we must have $\mu^*(\tilde{e}) = 1$, or equivalently, $w^*(\tilde{e}) = \theta_H$. But then type θ_H has a profitable deviation to \tilde{e} . \square

Thus, application of the equilibrium dominance condition yields a unique equilibrium outcome (i.e., the equilibria that survive can only differ in off-the-equilibrium path wages), which is the least-cost separating outcome. In fact, hybrid or partially-pooling equilibria (see [fn. 103](#)) are also eliminated by the equilibrium dominance condition.

7.2. Cheap Talk

We'll also cover the classic cheap-talk model of [Crawford and Sobel \(1982\)](#). **TBAdded.**

8. Mechanism Design

We only have time to peek into mechanism design.

We consider a framework with $I \equiv \{1, \dots, I\}$ agents. Each agent has a type $\theta_i \in \Theta_i$, which is the agent's private information. There is prior distribution of the vector $\theta \in \Theta \equiv \Theta_1 \times \dots \times \Theta_I$.¹⁰⁵ There is a set Z of allocations/alternatives/outcomes. Each agent has a vNM utility function $u_i(z, \theta)$. The environment has private values if we can reduce to $u_i(z, \theta_i)$; otherwise there is interdependent values. (Recall the terminology from our discussion of Bayesian games.)

A canonical special case is to take $Z \equiv X \times T$, where $T \equiv \mathbb{R}^I$, with the interpretation that $t \in \mathbb{R}^I$ is the vector of transfers paid by the agents. In this case, we refer to just the element of X as the allocation. We will assume *quasi-linearity* here: utility for each i is given simply by $u_i(x, \theta) - t_i$. Note that this imposes risk neutrality with respect to the transfer.

We will be interested in achieving or implementing (in a sense to be made precise below) an objective given by a social choice function, SCF, or allocation rule $\psi : \Theta \rightarrow Z$. In the quasi-linear environment, typical objectives of interest include efficiency (maximize $\sum_I u_i(x, \theta)$),¹⁰⁶ revenue maximization (maximize $\sum_I t_i$), or profit maximization (maximize $u_0(x, \theta) + \sum_I t_i$, with the interpretation that $u_0(x, \theta)$ is the “cost” for some additional player, usually called the principal).

A *mechanism* is a game form, given by $(A_1 \dots A_I, g)$, or just (A, g) for short, where $A \equiv A_1 \times \dots \times A_I$ and $g : A \rightarrow Z$. The interpretation is that for each i , A_i is the pure action set in some game set up by the designer; note that the game need not be a simultaneous-move game, just like in our discussion for Bayesian games. A mechanism, together with the other primitives (utilities and type distribution) defines a Bayesian game; a (pure) strategy for an agent in this game is $s_i : \Theta_i \rightarrow A_i$.

We say that a mechanism *implements SCF* ψ in a given solution concept if the Bayesian game has an “equilibrium” of that concept, (s_1^*, \dots, s_I^*) , such that for all θ , $\psi(\theta) = g(s_1^*(\theta_1), \dots, s_I^*(\theta_I))$. “Equilibrium” should be understood broadly; it is just some solution concept. Note that we are restricting to pure strategy equilibria. This can be relaxed. More importantly, we are focussing on what is called *weak/partial* implementation — i.e., we are not requiring that every equilibrium must lead to the desired outcomes.¹⁰⁷ Three common solutions concepts are Bayesian Nash equilibrium (BNE), ex post BNE, and dominant strategies. It is important to note that dominant strategies here refers to the “very weakly dominant” notion we mentioned in Remark 22, i.e., for each agent i , s_i^* should be a best response to any strategy profile of the opponents in the Bayesian game.

8.1. Revelation Principle

A *direct mechanism* has $A_i = \Theta_i$ for all i . A direct mechanism is said to be *incentive compatible* if it has a truth-telling equilibrium, i.e, all agents playing $s_i^*(\theta_i) = \theta_i$ is an equilibrium. An incentive compatible direct

¹⁰⁵ As usual, for general results below we treat Θ_i as finite. Furthermore, to simplify, I'll assume every θ_i has positive prior probability. These can be relaxed.

¹⁰⁶ The interpretation of this as (Pareto) efficiency requires some clarification. The best way to think of it is efficiency subject to holding constant the sum of transfers. That is, given any θ , if we don't choose x to maximize $\sum_I u_i(x, \theta)$, we can always find a redistribution of transfers that will lead to a Pareto improvement while not changing the sum of transfers.

¹⁰⁷ That requirement is called *full implementation*. Classic references on this include Maskin (1999) (circulated since 1977) and Jackson (1991). There is a nice survey by Jackson (2001).

mechanism implements a SCF if the implementation is achieved in the truth-telling equilibrium.

The following result is the *revelation principle*.

Proposition 20. *If some mechanism implements a SCF in dominant strategies (BNE), then there is an incentive compatible direct mechanism that also implements the SCF in dominant strategies (BNE).*

Proof. We prove it for dominant strategies and the case of private values. Suppose (A, g) implements ψ . Write $s^*(\theta)$ as shorthand for $(s_i^*(\theta_i))_{i=1}^I$ and similarly $s_{-i}(\theta_{-i})$. Consider the direct mechanism (Θ, γ) , where $\gamma(\theta) \equiv g(s^*(\theta))$.

For any i , θ_i , $\hat{\theta}_i$, and θ_{-i} , the dominance of s_i^* implies that

$$u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i) \geq u_i(g(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})), \theta_i),$$

and hence that

$$u_i(\gamma(\theta_i, \theta_{-i}), \theta_i) \geq u_i(\gamma(\hat{\theta}_i, \theta_{-i}), \theta_i),$$

But this means that in the direct mechanism (Θ, γ) , truth-telling is a [very weakly!] dominant strategy for each player. Moreover, it is clear that the truth-telling equilibrium of the direct mechanism implements the SCF.

The argument for BNE and private values is analogous, and a good exercise. The argument is also similar absent private values. \square

The revelation principle is a crucial result because it affords a tremendous simplification when we search for what SCFs can be implemented: not only can we restrict attention to a relatively small space of mechanisms, but moreover, the constraints we need consider are just those of incentive compatibility.

8.2. Vickrey-Clarke-Groves Mechanisms

Now consider an environment with private values and quasi-linearity. Let $x^* : \Theta \rightarrow X$ be some (ex post) efficient allocation rule: $\forall \theta, x^*(\theta) \in \arg \max_{x \in X} \sum_i u_i(x, \theta_i)$. (Assume maximizers exist.)

Consider the following transfer rule:

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

where $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ is arbitrary. A direct mechanism with allocation and transfer functions $(x^*, (t_i^G)_{i=1}^I)$ is called a *Groves mechanism*. (Hence the superscript G on t_i .)

Proposition 21. *Any Groves mechanism is dominant strategy incentive compatible.*

Proof. Type θ_i of agent i 's payoff from reporting $\hat{\theta}_i$ when others report $\hat{\theta}_{-i}$ is

$$\begin{aligned} & u_i(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i) - t_i^G(\hat{\theta}_i, \hat{\theta}_{-i}) \\ &= u_i(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i) + \sum_{j \neq i} u_j(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j) - h_i(\hat{\theta}_{-i}). \end{aligned} \tag{14}$$

Since x^* is an efficient allocation rule, it follows that the function above is maximized over $\hat{\theta}_i$ when $\hat{\theta}_i = \theta_i$. Dominant strategy incentive compatibility follows. \square

The key intuition here is that the specification of the transfer rule, along with an efficient allocation rule, ensures that each agent is maximizing efficiency given his type and that his opponents' types are their reports. To make that clear, consider the first two terms in (14); we can ignore the third term as it is independent of $\hat{\theta}_i$. You see that by announcing $\hat{\theta}_i = \theta_i$, agent i maximizes the sum of those two terms, producing an efficient allocation for the type profile $(\theta_i, \hat{\theta}_{-i})$. To say this yet another way, ignoring the h_i term, a Groves mechanism pays each agent the total utility of the other agents; hence, when including his own utility, each agent ends up maximizing the total utility of all agents — he internalizes any externalities caused by his report.

Let $x_{-i}^* : \Theta_{-i} \rightarrow X$ denote an efficient allocation rule if we ignore agent i 's payoff: $\forall \theta, x_{-i}^*(\theta_{-i}) \in \arg \max_{x \in X} \sum_{j \neq i} u_j(x, \theta_j)$. A special case of Groves mechanisms, called the *Vickrey-Clarke-Groves* (VCG) mechanism, obtains when we set

$$h_i(\theta_{-i}) = \sum_{j \neq i} u_j(x_{-i}^*(\theta_{-i}), \theta_j),$$

so that the transfer rule becomes

$$t_i^V(\theta) = \sum_{j \neq i} u_j(x_{-i}^*(\theta_{-i}), \theta_j) - \sum_{j \neq i} u_j(x^*(\theta), \theta_j). \quad (15)$$

The direct mechanism here is defined so as to make each agent precisely pay for the externality his presence (or his announcement) creates. To see this, note that the first term in t_i^V above is the total utility for all agents but i from an efficient allocation that ignores i 's presence, while the second term is their total utility from an efficient allocation when i is present. In particular, $t_i^V(\theta) = 0$ if i is not “pivotal”, i.e., if $x_{-i}^*(\theta_{-i}) = x^*(\theta)$.

The VCG mechanism (or Groves mechanisms, more broadly) is important because it says that we can achieve efficiency despite private information in various problems. Here is one example.

Example 38. Consider allocating an indivisible good. Let $X = I$, denoting which agent gets the good. Let $u_i(i, \theta_i) = \theta_i$ and $u_i(j, \theta_i) = 0$ for $j \neq i$. So the agent who gets the good gets payoff θ_i , and all other agents get 0. Consequently, ignoring the issue of ties (i.e., assuming $\theta_j \neq \theta_i$ for all $i, j \in I$), there is a unique efficient allocation rule:

$$x^*(\theta) = \arg \max_{j \in I} \theta_j, \text{ and } x_{-i}^*(\theta_{-i}) = \arg \max_{j \neq i} \theta_j.$$

Plugging into (15) yields

$$t_i^V(\theta) = \begin{cases} \max_{j \neq i} \theta_j & \text{if } x^*(\theta) = i \\ 0 & \text{otherwise.} \end{cases}$$

That is, the VCG mechanism in this context is simply a second-price auction!

References

- ABREU, D., P. K. DUTTA, AND L. SMITH (1994): “The Folk Theorem for Repeated Games: A New Condition,” *Econometrica*, 62(4), 939–948.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58(5), 1041–1063.
- ALIPRANTIS, C. D., AND K. C. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer.
- ATHEY, S. (2001): “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” *Econometrica*, 69(4), 861–889.
- AUMANN, R. J. (1964): “Mixed and Extensive Strategies in Infinite Extensive Games,” in *Advances in Game Theory*, ed. by M. Dresher, L. Shapley, and A. Tucker, vol. 52 of *Annals of Mathematics Studies*, pp. 627–650. Princeton University Press.
- BANKS, J. S., AND J. SOBEL (1987): “Equilibrium Selection in Signaling Games,” *Econometrica*, 55(3), 647–661.
- BENOIT, J.-P., AND V. KRISHNA (1985): “Finitely Repeated Games,” *Econometrica*, 53(4), 905–922.
- BERNHEIM, B. D. (1984): “Rationalizable Strategic Behavior,” *Econometrica*, 52(4), 1007–28.
- BLUME, A. (2003): “Bertrand Without Fudge,” *Economics Letters*, 78(2), 167–168.
- BRANDENBURGER, A., AND E. DEKEL (1993): “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory*, 59(1), 189–198.
- BRANDENBURGER, A., A. FRIEDENBERG, AND H. J. KEISLER (2008): “Admissibility in Games,” *Econometrica*, 76(2), 307–352.
- CARLSSON, H., AND E. VAN DAMME (1993): “Global games and equilibrium selection,” *Econometrica: Journal of the Econometric Society*, pp. 989–1018.
- CHO, I.-K., AND D. KREPS (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102(2), 179–221.
- CRAWFORD, V., AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50(6), 1431–1451.
- EWERHART, C. (2002): “Backward Induction and the Game-Theoretic Analysis of Chess,” *Games and Economic Behavior*, 39, 206–214.
- FUDENBERG, D., AND J. TIROLE (1991a): *Game Theory*. MIT Press, Cambridge, MA.
- (1991b): “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53(2), 236–260.

- GUL, F. (1998): “A Comment on Aumann’s Bayesian View,” *Econometrica*, 66(4), 923–927.
- HARSANYI, J. C. (1973): “Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points,” *International Journal of Game Theory*, 2, 1–23.
- JACKSON, M., L. SIMON, J. SWINKELS, AND W. ZAME (2002): “Communication and Equilibrium in Discontinuous Games of Incomplete Information,” *Econometrica*, 70(5), 1711–1740.
- JACKSON, M. O. (1991): “Bayesian Implementation,” *Econometrica*, 59(2), 461–77.
- JACKSON, M. O. (2001): “A Crash Course in Implementation Theory,” *Social Choice and Welfare*, 18(4), 655–708.
- KOHLBERG, E., AND J.-F. MERTENS (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54(5), 1003–1037.
- KREPS, D., AND R. WILSON (1982): “Sequential Equilibria,” *Econometrica*, 50(4), 863–894.
- KREPS, D. M., AND J. A. SCHEINKMAN (1983): “Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes,” *Bell Journal of Economics*, 14(2), 326–337.
- MAILATH, G. J., AND L. SAMUELSON (2006): *Repeated Games and Reputations*. Oxford University Press.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*. Oxford University Press.
- MASKIN, E. (1999): “Nash Equilibrium and Welfare Optimality,” *Review of Economic Studies*, 66(1), 23–38.
- MERTENS, J.-F., AND S. ZAMIR (1985): “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*, 14(1), 1–29.
- MILGROM, P. R., AND R. J. WEBER (1985): “Distributional Strategies for Games of Incomplete Information,” *Mathematics of Operations Research*, 10(4), 619–632.
- MONDERER, D., AND D. SAMET (1989): “Approximating common knowledge with common beliefs,” *Games and Economic Behavior*, 1(2), 170–190.
- MORRIS, S., AND H. S. SHIN (1998): “Unique equilibrium in a model of self-fulfilling currency attacks,” *American Economic Review*, pp. 587–597.
- OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT Press.
- PEARCE, D. G. (1984): “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 52(4), 1029–50.
- RENY, P. J. (1999): “On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games,” *Econometrica*, 67, 1029–1056.
- RILEY, J. G. (1979): “Informational Equilibrium,” *Econometrica*, 47(2), 331–359.

- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior under “Almost Common Knowledge.”,” *American Economic Review*, 79(3), 385–91.
- SELTEN, R. (1975): “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 4, 25–55.
- SORIN, S. (1986): “On Repeated Games with Complete Information,” *Mathematics of Operations Research*, 11(1), 147–160.
- SPENCE, M. (1973): “Job Market Signaling,” *Quarterly Journal of Economics*, 87(3), 355–374.
- WEINSTEIN, J., AND M. YILDIZ (2007): “A structure theorem for rationalizability with application to robust predictions of refinements,” *Econometrica*, 75(2), 365–400.
- ZAMIR, S. (2008): “Bayesian games: Games with incomplete information,” Discussion Paper Series dp486, Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem.