John Nachbar
Washington University in St. Louis
This Version: July 2, 2023

Game Theory Basics I: Strategic Form Games[1]

1 Introduction.

Game theory is an analytical framework for investigating conflict and cooperation. Early work was motivated by gambling and recreational games such as chess, hence the "game" in game theory. But it quickly became clear that the framework had much broader application. Today, game theory is used for mathematical modeling in a wide range of disciplines, including many of the social sciences, computer science, and evolutionary biology. Here, I draw examples mainly from economics.

These notes are an introduction to a mathematical formalism called a *strategic form* game (also called a *normal form* game). For the moment, think of a strategic form game as representing an atemporal interaction: each player (in the language of game theory) acts without knowing what the other players have done. An example is a single instance of the two-player game Rock-Paper-Scissors (probably already familiar to you, but discussed in the next section).

In companion notes, Game Theory Basics II: Extensive Form Games, I develop an alternative formalism called an *extensive form* game. Extensive form games explicitly capture temporal considerations, such as the fact that in standard chess, players move in sequence, and each player knows the prior moves in the game. As I discuss in the notes on extensive form games, there is a natural way to give any extensive form game a strategic form representation.

There is a third formalism called a game in *coalition form* (also called *characteristic function form*). The coalition form abstracts away from the details of what individual players do and focuses instead on what payoff allocations are physically possible, both for all players taken together and for every subset (coalition) of players. I do not (yet) have notes on games in coalition form.

A nice, short introduction to the study of strategic and extensive form games is Osborne (2008). A standard undergraduate text on game theory, one that I frequently use for my own course, is Gibbons (1992). Additional options are Osborne (2003), Watson (2013), and Tadelis (2013). A standard graduate game theory text is Fudenberg and Tirole (1991). I also recommend Myerson (1991), Osborne and Rubinstein (1994), and Mailath (2019). There are also good introductions to game theory in graduate microeconomic theory texts such as Kreps (1990), Mas-Colell et al. (1995), Jehle and Reny (2000), and Kreps (2023). There are useful short survey articles in Durlauf and Blume (2008), an updated version of which can often be accessed online through the library. Finally, Luce and Raiffa (1957) is a classic and still valuable text on game theory, especially for discussions of interpretation and motivation.

I start my discussion of strategic form games with Rock-Paper-Scissors.

2 An example: Rock-Paper-Scissors.

The game *Rock-Paper-Scissors (RPS)* is represented in Figure 1 in what is called a *game box*. There are two players, 1 and 2. Each player has three strategies in the

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

Figure 1: *A game box for Rock-Paper-Scissors (RPS).*

game: R (rock), P (paper), and S (scissors). Player 1 is represented by the rows while player 2 is represented by the columns. If player 1 chooses R and player 2 chooses P then this is represented as the pair, called a *strategy profile*, (R,P); for this profile, player 1 gets a *payoff* of -1 and player 2 gets a payoff of +1, represented as a *payoff profile* $(-1, 1)$.

For interpretation, think of payoffs as encoding preferences over winning, losing, or tying, with the understanding that S beats P (because scissors cut paper), P beats R (because paper can wrap a rock ...), and R beats S (because a rock can smash scissors). If both choose the same, then they tie. The interpretation of payoffs is actually quite delicate and I discuss this issue at length in Section 3.3.

This game is called *zero-sum* because, for any strategy profile, the sum of payoffs is zero. In any zero-sum game, there is a number V, called the *value* of the game, with the property that player 1 can guarantee that she gets at least V in expectation no matter what player 2 does and conversely player 2 can get $-V$ no matter what player 1 does. I provide a proof of this theorem in Section 7.2. In this particular game, $V = 0$; either player can guarantee that they get 0 in expectation by randomizing evenly over the three strategies.

Note that randomization is necessary to guarantee an expected payoff of at least 0. In Season 4 Episode 16 of the Simpsons, Bart persistently plays Rock against Lisa, and Lisa plays Paper, and wins. Bart here doesn't even seem to understand the game box, since he says, "Good old Rock. Nothing beats that." I discuss the interpretation of randomization in Section 3.4.

3 Strategic Form Games.

3.1 The Strategic Form.

I restrict attention in the formalism to finite games: finite numbers of players and finite numbers of strategies. Some of the examples, however, involve games with infinite numbers of strategies.

A *strategic form game* is a tuple $(I, (S_i)_i, (u_i)_i)$.

- I is the finite set, assumed not empty, of *players* with typical element i. The cardinality of I is N; I sometimes refer to N-player games. To avoid triviality, assume $N \geq 2$ unless explicitly stated otherwise.

- S_i is the set, assumed finite and not empty, of player i's *strategies*, often called *pure* strategies to distinguish from the mixed strategies described below.

  $S = \prod_i S_i$ is the set of *pure strategy profiles*, with typical element $s = (s_1, \ldots, s_N) \in S$.

- u_i is player i's *payoff function*: $u_i : S \to \mathbb{R}$. I discuss the interpretation of payoffs in Section 3.3.

Games with two players and small strategy sets can be represented via a game box, as in Figure 1 for Rock-Paper-Scissors. In that example, $I = \{1, 2\}$, $S_1 = S_2 = \{R, P, S\}$, and payoffs are as given in the game box.

As anticipated by the RPS example, we will be interested in randomization. For each i, let Σ_i be the set of probabilities over S_i, also denoted $\Delta(S_i)$. An element $\sigma_i \in \Sigma_i$ is called a *mixed strategy* for player i. Under σ_i, the probability that i plays pure strategy s_i is $\sigma_i[s_i]$.

A pure strategy s_i is equivalent to a degenerate mixed strategy, with $\sigma_i[s_i] = 1$ and $\sigma_i[\hat{s}_i] = 0$ for all $\hat{s}_i \neq s_i$. Abusing notation, I use the notation s_i for both the pure strategy s_i and for the equivalent mixed strategy.

Assuming that the cardinality of S_i is at least 2, the strategy σ_i is *fully (or completely) mixed* iff $\sigma_i[s_i] > 0$ for every $s_i \in S_i$. σ_i is *partly mixed* iff it is neither fully mixed nor pure (degenerate).

Since the strategy set is assumed finite, I can represent a σ_i as a vector in either $\mathbb{R}^{|S_i|}$ or $\mathbb{R}^{|S_i|-1}$, where $|S_i|$ is the number of elements in S_i. For example, if S_1 has two elements, then I can represent σ_1 as either $(p, 1-p)$, with $p \in [0, 1]$ (probability p on the first strategy, probability $1-p$ on the second), or just as $p \in [0, 1]$ (with the probability on the second strategy inferred to be $1-p$). And a similar construction works for any finite strategy set. Note that under either representation, Σ_i is compact and convex.

I discuss the interpretation of mixed strategies in Section 3.4. For the moment, however, suppose that players might actually randomize, perhaps by making use of coin flip or toss of a die. In this case the true set of strategies is actually Σ_i rather than S_i. One can think of S_i as a concise way to represent the true strategy set, which is Σ_i.

$\Sigma = \prod_i \Sigma_i$ is the set of *mixed strategy profiles*, with typical element $\sigma = (\sigma_1, \ldots, \sigma_N)$. A mixed strategy profile σ induces an independent probability distribution over S.

Example 1. Consider a two-player game in which $S_1 = \{T, B\}$ and $S_2 = \{L, R\}$. If $\sigma_1[T] = 1/4$ and $\sigma_2[L] = 1/3$ then the induced distribution over S can be represented in a game box as in Figure 2. □

	L	R
T	1/12	2/12
B	3/12	6/12

Figure 2: *An independent distribution over strategy profiles.*

Abusing notation, let $u_i(\sigma)$ be the *expected payoff* under this independent distribution; that is,

$$u_i(\sigma) = \mathbb{E}_\sigma[u_i(s)] = \sum_{s \in S} u_i(s)\sigma_1[s_1] \ldots \sigma_N[s_N].$$

where $\mathbb{E}_\sigma$ is the expectation with respect to the probability distribution over S induced by σ.

Example 2. In RPS (Figure 1), if Player 1 plays R and Player 2 randomizes $(1/6, 1/3, 1/2)$, then Player 1's expected payoff is

$$0 \times 1/6 + (-1) \times 1/3 + 1 \times 1/2 = 1/6.$$

□

Finally, I frequently use the following notation. A strategy profile $s = (s_1, \ldots, s_N)$ can also be represented as $s = (s_i, s_{-i})$, where $s_{-i} \in \prod_{j \neq i} S_j$ is a profile of pure strategies for players other than i. Similarly, $\sigma = (\sigma_i, \sigma_{-i})$ is alternative notation for $\sigma = (\sigma_1, \ldots, \sigma_N)$, where $\sigma_{-i} \in \prod_{j \neq i} \Sigma_j$ is a profile of mixed strategies for players other than i.

3.2 Correlation.

The notation so far builds in an assumption that any randomization is independent. To see what is at issue, consider the following game.

Example 3. The game box for one version of *Battle of the Sexes* is in Figure 3. The players would like to coordinate on either (A, A) or (B, B), but they disagree about which of these is better. □

	A	B
A	$8, 10$	$0, 0$
B	$0, 0$	$10, 8$

Figure 3: *A game box for Battle of the Sexes.*

If players in the Battle of the Sexes game depicted in Figure 3 can coordinate (say by pre-play communication) then it is possible that they would toss a coin prior to play and then execute (A, A) if the coin lands heads and (B, B) if the coin lands tails. This induces a *correlated* distribution over strategy profiles, which I represent in Figure 4.

	A	B
A	$1/2$	0
B	0	$1/2$

Figure 4: *A correlated distribution over strategy profiles for Battle of the Sexes.*

Note that under this correlated distribution, each player plays A half the time. If players were instead to play A half the time independently, then the distribution over strategy profiles would be as in Figure 5.

	A	B
A	$1/4$	$1/4$
B	$1/4$	$1/4$

Figure 5: *An independent distribution over strategy profiles for Battle of the Sexes.*

The space of all probability distributions over S is Σ^c, also denoted $\Delta(S)$, with generic element denoted σ^c. Abusing notation (again), I let $u_i(\sigma^c) = \mathbb{E}_{\sigma^c}[u_i(s)] = \sum_{s \in S} u_i(s)\sigma^c[s]$. For strategies for players other than i, the notation is $\sigma^c_{-i} \in \Sigma^c_{-i} = \Delta(S_{-i})$. The notation $(\sigma_i, \sigma^c_{-i})$ denotes the distribution over S for which the probability of $s = (s_i, s_{-i})$ is $\sigma_i[s_i]\sigma^c_{-i}[s_{-i}]$. Any element of Σ induces an element of Σ^c: an independent distribution over S is a special form of correlated distribution over S.

Remark 1. A mixed strategy profile $\sigma \in \Sigma$ is *not* an element of Σ^c, hence $\Sigma \not\subseteq \Sigma^c$. Rather, σ induces a distribution over S, and this induced distribution is an element of Σ^c.

In the Battle of the Sexes of Figure 3, to take one example, it takes three numbers to represent an arbitrary element $\sigma^c \in \Sigma^c$ (only three, rather than four, because the four numbers have to add up to 1, hence one number can always be inferred from the other three). In contrast, each mixed strategy σ_i can be described by a single number, as discussed earlier, hence σ can be represented as a pair of numbers, which induces a distribution over S. Thus, in this example, the set of independent strategy distributions over S is two dimensional while the set of all (correlated) strategy distributions is three dimensional. More generally, the set of independent strategy distributions over S is a lower dimensional subset of Σ^c. $\square$

3.3 Interpreting Payoffs.

In most applications in economics, payoffs are intended to encode choice by decision makers, and this brings with it subtleties that affect the interpretation and testability of the theory. The material in this section assumes some familiarity with standard decision theory, although I try to make the discussion as self-contained as possible.

To make the issues more transparent, I first introduce some additional structure. Let X be a finite set of possible *outcomes* for the game. Let $\gamma(s)$ give the outcome associated with strategy profile s; thus $\gamma : S \to X$.

- An outcome can encode (or represent) a profile of "prizes" for the players. In some games, this may take the form of a profile of numbers, which can be interpreted as monetary payments or penalties. This is frequently the case in economics applications of game theory. But the "prizes" can also be more abstract. For example, an outcome can specify which player (or players) "won" (without any more tangible prize), as in many recreational games. An outcome can be an assignment of partners (romantic, professional, organ donor). It can be multidimensional, as in market games where players trade bundles (vectors) of commodities. The possibilities are endless. Prizes can even be probabilistic; a prize might be a lottery ticket, for example.

- In addition to encoding prizes, an outcome can also encode how the game was played. For example, as discussed in the companion notes (Game Theory Basics II: Extensive Form Games), chess can be represented as a game in strategic form. For chess, an outcome might encode not only who won (the prize) but also a record of the exact sequence of moves by the players.

Example 4. Recall RPS (Rock-Paper-Scissors) from Section 2. By assumption, the possible outcomes are

$$\{(\text{win}, \text{lose}), (\text{lose}, \text{win}), (\text{tie}, \text{tie})\},$$

where, for example, (win, lose) means that player 1 gets "win" (whatever that might mean) while player 2 gets "lose." In particular, I assume that the players have no particular interest in what strategies are played per se, only in who wins. γ is given by, for example, $\gamma(R, P) = (\text{lose}, \text{win})$. $\square$

I assume that player choice can be represented by a preference relation that is a complete preorder over X. That is, there is a complete preorder, which I call preferences, such that, given any subset of X, an outcome x is the most highly ranked outcome in that subset if, in fact, the player would have chosen x from that subset. This decision theoretic assumption is substantive (it builds in transitivity, in particular) but it is implicit in all of standard game theory. Moreover (and, because I have assumed that X is finite, this is a theorem rather than a new assumption), player i's preferences can in turn be represented by a *utility function*, say $v_i : X \to \mathbb{R}$, where $v_i(x) > v_i(\hat{x})$ iff player i strictly prefers x to $\hat{x}$. v_i will typically not be unique: two different utility functions represent the same preferences over X iff either is a strictly increasing function of the other.

Example 5. Returning to RPS, to simplify the discussion first note that the possible outcomes are exactly identified by which player wins or if they tie.

Suppose that either player strictly prefers win to tie and strictly prefers tie to lose. Any utility function that assigns a higher number to win than tie, and a higher number to tie than lose, will work. For example, I could have that, for player 1, win is -9, tie is -10, and lose is -100 and for player 2 that win is 11008, tie is 1, and lose is -90. In fact, and this is arbitrary at this point in the discussion, I take the utilities for both players to be win is 1, tie is 0, and lose is -1. □

Thus far, I have focused on pure strategy profiles. We are, however, also interested in mixed strategy profiles, which will induce probability distributions over X. The assumption in standard game theory is that players are *expected utility maximizers*. Formally, this means the following. As was the case for X, it is assumed that player choice over $\Delta(X)$ can be represented by a complete preorder over X and that this preorder can, in turn, by represented by a utility function (because $\Delta(X)$ is uncountably infinite, the second of these steps now requires an auxiliary technical assumption). As before, the utility function representing choice/preferences over $\Delta(X)$ is not unique. Expected utility maximization introduces the new assumption that out of the set of possible utility functions that represent i's preferences over $\Delta(X)$, there is one, say $\mathcal{V}_i : \Delta(X) \to \mathbb{R}$, for which there is a utility function over X, say $v_i : X \to \mathbb{R}$, such that, for any probability $\lambda \in \Delta(X)$, $\mathcal{V}_i(\lambda)$ is the expectation, with respect to λ, of $v_i(x)$:

$$\mathcal{V}_i(\lambda) = \mathbb{E}_\lambda[v_i(x)].$$

If λ_x is the degenerate distribution that assigns probability 1 to outcome x then $\mathcal{V}_i(\lambda_x) = v_i(x)$.

v_i is potentially much more tightly pinned down than before; it not only has to rank outcomes in X correctly but it has to have the additional property that expectations of v_i rank probabilities in $\Delta(X)$ correctly. It is possible for two different v_i to rank outcomes in X the same way but rank probability distributions in $\Delta(X)$ differently. I discuss this issue both immediately below and in Section 4.1, which provides an explicit example.

Finally, the payoff function is simply the composition of the v_i and γ functions: $u_i(s) = v_i(\gamma(s))$. Similarly (and abusing notation, as usual), $u_i(\sigma) = \mathbb{E}_\sigma[v_i(\gamma(s))]$.

Example 6. Returning yet again to RPS, if, for example, $v_1((\text{win}, \text{lose})) = 1$ then $u_1(R, S) = 1$, and so on. Again, note that there is an implicit assumption that v_i is compatible with a $\mathcal{V}_i$ that represents the player's choice over $\Delta(X)$.

As an illustrative calculation, suppose that player 2 randomizes $(1/6, 1/3, 1/2)$. Then player 1's expected payoffs are $1/6$ from R, $-1/3$ from P and $1/6$ from S. So, for this particular mixture by player 2, player 1 is indifferent between R and S but prefers either to P. □

With this foundation, I can now make a number of remarks about the interpretation of payoffs in decision theoretic terms.

1. Because the payoff function encodes choice, by construction, payoff maximization over X is built into payoffs; it is not a separate behavioral assumption. If you think players would make different choices, then you need to change the payoff representations to reflect this. And if you balk at the idea of using utility functions to represent choice, then we are going to have to move to some other formalism entirely; you can't both reject utility functions and continue to use payoff functions.

2. There is an old and still lively debate about whether the expected utility assumption is a good one, both for game theory in particular and for economics more generally. See, for example, the discussion of expected utility in Kreps (2012).

3. Since we care about choice over $\Delta(X)$, there is less flexibility in choosing v_i than there would be if we care only about representing choice over just X. For general expected utility, two v_i represent the same choice over $\Delta(X)$ iff either is an affine transformation of the other, with strictly positive slope: that is, $\hat{v}_i = av_i + b$, with $a > 0$. In fact, because the expected utility calculations that appear in game theory are special (they are constrained by the structure of the game), there is somewhat more flexibility in choosing v_i than this suggests. In a two-player game, for example, we can multiply all of player i's payoffs by a positive constant (as above) and we can add a constant to all of player i's payoffs corresponding to each pure strategy of the other player, with (potentially) a different constant for each of the other player's pure strategies (this is the new part), without changing the player's ranking over their mixed strategies. See Section 4.1 for more on this.

4. In many economic applications, an outcome is a profile of monetary values (profits, for example). It is common practice in such applications to assume that $u_i(s)$ equals the prize to i: if s gives i profits of \$1 billion, then $u_i(s) =$ \$1 billion. This assumption is substantive.

   (a) The assumption rules out phenomena such as altruism or envy. In contrast, the general formalism allows, for example, for $u_i(s)$ to be the sum of the individual prizes (a form of perfect altruism).

(b) The assumption rules out risk aversion. If players maximize expected payoffs, then to capture risk aversion, payoffs must be a concave function of profit. See also the example in Section 4.1.

(c) The assumption rules out the possibility that players care about more than the monetary prize, that they also care about how the game was played. The players might also care about pollution or worker welfare or national security, and so on. Again, these extra consideration are not ruled out by the game theory formalism itself but only by this particular implementation of that formalism.

5. It can be difficult to test game theory predictions such as Nash equilibrium in a lab. The experimenter can control γ, but the v_i, and hence the u_i, are in the heads of the subjects and not directly observable. In particular, it is not a violation of game theory to find that players are altruistic or spiteful. This flexibility of the game theory formalism is a feature, not a bug: the goal is to have a formalism that can be used to model essentially any strategic interaction.

6. A related point is that while it may be reasonable to assume that there is common knowledge of γ (everyone knows the correct γ, everyone knows that everyone knows the correct γ, and so on), there may not be even mutual knowledge of the v_i and hence of the u_i (players may not know the utility functions of the other players). In particular, in economics contexts, even if it makes sense to assume that players have common knowledge of rankings over outcomes (e.g., even if it is common knowledge that everyone cares only about their own income, and that everyone prefers more income to less), v_i may still not be common knowledge (because players may not know each other's attitudes toward risk, for example).

   Games in which there is not common knowledge of the u_i are called games of *incomplete information*. I discuss approaches to modeling such environments later in the course.

The interpretation of payoffs in terms of decision theory is not the only one possible. For example, in some applications of game theory to evolutionary biology, the strategies might be alleles (alternative versions of a gene) and payoffs might be the expected number of offspring.

Remark 2. $\mathcal{V}_i$ are v_i are both often called Von Neumann-Morgenstern (VN-M) expected utility functions; usage varies. v_i is also often called a Bernoulli utility function, especially for the special case in which the outcome is a profile of monetary payments (or penalties) and players care only about their own individual payments. I sometimes refer to v_i as a *subutility function*.

The basic idea of expected utility is usually credited to Daniel Bernoulli (1700-1782), who focused on the special case of probabilities over monetary payments. von

Neumann and Morgenstern (1947), as part of their work on providing a foundation for game theory, extended expected utility to abstract outcomes and also provided an axiomatic foundation. There were, subsequently, other axiomatizations, of which the best known is Herstein and Milnor (1953). Savage (1954) extended the expected utility idea to frameworks in which probabilities are subjective. Anscombe and Aumann (1963) provided an alternate, and in some ways simpler, formulation of Savage's theory for settings in which there is objective as well as subjective probability. (The leading example in Anscombe and Aumann (1963) is that the probabilities on a roulette wheel are objective while those for betting on horse races are subjective. A hardcore subjectivist would counter that all probabilities are subjective, even those for a roulette wheel or the toss of a coin or a die.) For more on decision theory in general and decision over probabilities in particular, see Kreps (1988), Mas-Colell et al. (1995), and Kreps (2012).

Note that the maintained assumption is that u_i is fixed as part of the description of the game but that σ is not. In terms of the notation introduced above, this means that v_i is fixed but λ is not. This is at variance with the standard construction of subjective expected utility as in Savage (1954), which takes v_i and λ to be jointly determind. See Perea (2020) for one take on this issue. $\square$

3.4 Interpreting Randomization.

There are three main interpretations of randomization in games. These interpretations are not mutually exclusive.

1. *Objective Randomization.* Each player has access to a randomization device. In this case, the true strategy set for player i is Σ_i. S_i is just a concise way to communicate Σ_i.

   Underlying objective randomization is the idea that in many games it can be important not to be predictable. This was the case in Rock-Paper-Scissors, for example (Section 2). A particular concern is that if the game is played repeatedly then one's behavior should not follow some easily detected, and exploited, pattern. An example along these lines is that military submarines occasionally implement hard turns in order to detect possible trailing submarines; such maneuvers are called "clearing the baffles." In order to be as effective as possible, it is important that these turns be unpredictable. Sontag and Drew (1998) reported that a captain of the USS Lapon used dice in order to randomize. Curiously, in Clancy (1984), a classic military techno-thriller, a critical plot point involves a CIA analyst correctly predicting when and how a (fictional) top Russian submarine commander would perform a "crazy Ivan" and clear the baffles of his submarine.

   Even if a pure strategy exhibits a pattern that can be detected and exploited in principle, it may be impossible to do so in practice if the pattern is sufficiently complicated. See, for example, Hu (2014).

2. *Empirical Randomization.* From the perspective of an observer (say an experimentalist), $\sigma_i[s_i]$ is the frequency with which s_i is played.

   The observer could, for example, be seeing data from a cross section of play by different players (think of a Rock-Paper-Scissors tournament with many simultaneous matchings). $\sigma_i[s_i]$ is the fraction of players in role i of the game who play s_i. Nash discussed this interpretation explicitly in his thesis, Nash (1950b). Alternatively, the observer could be seeing data from a time series: the same players play the same game over and over, and $\sigma_i[s_i]$ is the frequency with which s_i is played over time.

3. *Subjective Randomization.* From the perspective of player $j \neq i$, $\sigma_i[s_i]$ is the probability that j assigns to player i playing s_i.

   Consider again the cross sectional interpretation of randomization, in which many instances of the game are played by different players. If players are matched randomly and anonymously to play the game then, from the perspective of an individual player, the opponents are drawn randomly, and hence opposing play can be "as if" random even if the player knows that individual opponents are playing pure strategies.

   An important variant of the cross sectional interpretation of randomization is the following idea, due to Harsanyi (1973). As discussed in Section 3.3, players may know the γ of the game (giving prizes), but not the v_i of the other players (giving preferences over prizes), and hence may not know the u_i (giving payoffs) of the other players. Suppose that player j assigns a probability distribution over possible u_i, and for each u_i, player j forecasts play of some *pure* strategy s_i. Then, even though player j thinks that player i will play a pure strategy, player i's play is effectively random in the mind of player j because the distribution over u_i induces a distribution over s_i.

   A subtlety with the subjective interpretation is that if there are three or more players, then two players might have different subjective beliefs about what a third player might do. This is assumed away in our notation, where σ_i does not depend on anything having to do with the other players.

3.5 The Reduced Strategic Form.

The strategic form may contain pure strategies that are "redundant."

Consider Figure 6, which gives the game box for a variant of Battle of the Sexes (Figure 3) that has an added third strategy, B', for player 1. B' is redundant in the sense that it generates the same payoffs as B for either player, regardless of what player 2 does. From the standpoint of payoff maximization, neither player cares whether player 1 chooses B' rather than B, or vice versa.

Now consider Figure 7, another Battle of the Sexes variant. Pure strategy C is also redundant, in the sense that player 1 can generate the same payoffs, for either

	A	B
A	$8, 10$	$0, 0$
B	$0, 0$	$10, 8$
B'	$0, 0$	$10, 8$

Figure 6: *An alternate game box for Battle of the Sexes.*

	A	B
A	$8, 10$	$0, 0$
B	$0, 0$	$10, 8$
C	$4, 5$	$5, 4$

Figure 7: *An alternate game box for Battle of the Sexes.*

player, as C by randomizing 50:50 between A and B. C is, in effect, simply the name of a particular mixed strategy.

Given a game in strategic form, the *reduced strategic form* is the game in which strategies that are redundant in the above senses have been removed. Arguably, the redundant strategies just add notational clutter, and can be excluded without affecting the analysis of the game in any meaningful way. There are, however, two caveats.

First, the claim isn't that redundant strategies would never be played but rather that the reduced form can be analyzed first and then any redundant strategies can be brought back into consideration, if for some reason this is desired. For example, in Battle of the Sexes, if the analysis of the reduced form (Figure 6) predicts that player 1 would play strategy B, then the same analysis predicts that in the variant game of Figure 6, player 1 might play either B or B', or might even randomize between them.

Second, when we get to extensive from games, there will be solution concepts, such as subgame perfect Nash equilibrium, that, strictly speaking, require using the strategy set generated by the extensive form, even if some of those strategies turn out to be redundant in the above sense.

4 Nash Equilibrium.

4.1 The Best Response Correspondence.

Given a profile of opposing (mixed) strategies $\sigma_{-i} \in \Sigma_{-i}$, let $\mathrm{BR}_i(\sigma_{-i})$ be the set of mixed strategies for player i that maximize player i's expected payoff; formally,

$$\mathrm{BR}_i(\sigma_{-i}) = \{\sigma_i \in \Sigma_i : \forall \hat{\sigma}_i \in \Sigma_i, u_i(\sigma_i, \sigma_{-i}) \geq u_i(\hat{\sigma}_i, \sigma_{-i})\}.$$

An element of $\mathrm{BR}_i(\sigma_{-i})$ is called a *best response* to σ_{-i}.

Given a profile of strategies $\sigma \in \Sigma$, let $\mathrm{BR}(\sigma)$ be the set of mixed strategy profiles $\hat{\sigma}$ such that, for each i, $\hat{\sigma}_i$ is a best response to σ_{-i}. Formally,

$$\mathrm{BR}(\sigma) = \{\hat{\sigma} \in \Sigma : \forall i \; \hat{\sigma}_i \in \mathrm{BR}_i(\sigma_{-i})\} .$$

BR is a correspondence on Σ. Since Σ_i is compact for each i, Σ is compact. For each i, expected payoffs are continuous, which implies that for any $\sigma \in \Sigma$ and any i, $\mathrm{BR}_i(\sigma_{-i})$ is not empty. Thus, BR is a non-empty-valued correspondence on Σ.

The definition of best response builds in an assumption that players are expected payoff maximizers; see also the discussion in Section 3.3. One consequence of expected payoff maximization is that payoffs can be rescaled in certain ways without affecting the best response correspondence.

1. For any s_{-i}, one can add any constant to $u_i(s_i, s_{-i})$, the same constant for every s_i but possibly differing across s_{-i}, without affecting player i's best response correspondence.

2. One can multiply u_i by any positive constant, the same constant for every s_i and every s_{-i}, without affecting player i's best response correspondence.

Example 7. The game in which player 1's payoffs are

	L	R
T	5	11
B	2	14

generates the same best response correspondence for player 1 as the game with payoffs

	L	R
T	1	0
B	0	1

To see this, subtract 2 from the first column, 11 from the second, and then multiply all payoffs by 1/3. □

On the other hand, a non-linear transformation of payoffs will typically change the best response correspondence. In particular, if the transformation is concave, then the best response correspondence will embody risk aversion, as in the next example.

Example 8. Consider the game in which player 1's payoffs are

	L	R
T	10	0
M	0	10
B	4	4

If player 2 randomizes equally over L or R then the set of player 1's best responses is any mixture over T and M, for an expected payoff of 5.

If we take the square root, a concave transformation, then payoffs are (approximately),

	L	R
T	3.2	0
M	0	3.2
B	2	2

In this case, if player 2 randomizes equally over L or R, then player 1's unique best response is B.

For more on this issue, see Weinstein (2016). □

4.2 Nash Equilibrium.

The single most important solution concept for games in strategic form is *Nash equilibrium* (NE). Nash himself called it an "equilibrium point." I discuss some aspects of the intellectual history of Nash equilibrium in Section 8.

A NE is a (mixed) strategy profile σ^* such that for each i, σ_i^* is a best response to σ_{-i}^*, hence $\sigma^* \in \mathrm{BR}(\sigma^*)$: a NE is a fixed point of BR.

Definition 1. *Fix a game. A strategy profile $\sigma^* \in \Sigma$ is a* Nash equilibrium *(NE) iff σ^* is a fixed point of* BR. *If σ^* is a NE then the associated vector of expected payoffs is a Nash equilibrium payoff profile.*

I will go through a number of examples of NE shortly.

As a point of terminology, it is important to remember that the Nash equilibrium is a strategy profile, not a payoff profile.

A NE is *pure* if all the strategies are pure. It is *fully mixed* if every σ_i is fully mixed. If a NE is neither pure nor fully mixed, then it is *partly mixed*. Strictly speaking, a pure NE is a special case of a mixed NE; in practice, however, I may sometimes (sloppily) write mixed NE when what I really mean is fully or partly mixed.

In separate notes (Game Theory Basics III), I survey some motivating stories for NE. It is easy to be led astray by motivating stories, however, and confuse what NE "ought" to be with what NE, as a formalism, actually is. For the moment, therefore, I focus narrowly on the formalism.

Remark 3. From Section 4.1, we know that certain types of payoff transformations do not affect the best response correspondence. It follows that these transformations also do not affect the set of Nash equilibria. □

The following result was first established in Nash (1950a).

Theorem 1. *Every (finite) strategic form game has at least one NE.*

Proof. For each i, Σ_i is compact and convex and hence Σ is compact and convex. BR is a non-empty valued correspondence on Σ and it is easy to verify that it has a closed, convex-valued graph. By the Kakutani fixed point theorem (see my notes on fixed point theorems), BR has a fixed point $\sigma \in \mathrm{BR}(\sigma)$. ■

Theorem 1 is not true if we restrict attention to pure strategies. Rock-Paper-Scissors (RPS; Section 2) is a canonical example of a game that has no pure strategy NE. Instead, RPS has a unique NE that is fully mixed and in it both players randomize equally across their three strategies; see also Example 11.

Remark 4. Non-finite games may not have NE. A trivial example is the game in which you name a number $\alpha \in [0, 1)$, and I pay you a billion dollars with probability α (and nothing with probability $1 - \alpha$). Assuming that you prefer more (expected) money to less, you don't have a best response and hence there is no NE.

The difficulty in the above example was that the strategy set was not compact. A related example has $\alpha \in [0, 1]$, which is now compact, but with payoff of zero if you choose $\alpha = 1$. Here the problem is that payoff function is not continuous.

The most straightforward extension of Theorem 1 is to games in which each player's strategy set is a compact subset of some measurable topological space and the payoff functions are continuous. In this case, the set of mixed strategies is weak-* compact. One can take finite approximations to the strategy sets, apply Theorem 1 to get a NE in each finite approximation, appeal to compactness to get a convergent subsequence of these mixed strategy profiles, and then argue, via continuity of payoffs, that the limit must be an equilibrium in the original game.

This approach will not work if the utility function is not continuous and, unfortunately, games with discontinuous utility functions are fairly common in economics. A well known example is the Bertrand duopoly game, discussed in Example 16 in Section 5. For a survey of the literature on existence of NE in general and also on existence of NE with important properties (e.g., pure strategy NE or NE with monotonicity properties), see Reny (2008). □

Remark 5. A related point is the following. The Brouwer fixed point theorem states that if $D \subseteq \mathbb{R}^N$ is compact and convex and $f : D \to D$ is continuous then f has a fixed point: there is an $x \in D$ such that $f(x) = x$. Brouwer is the underlying basis for the Kakutani fixed point theorem, which was used in the proof Theorem 1, and one can also use Brouwer more directly to prove NE existence (as, indeed, Nash himself subsequently did for finite games in Nash (1951)). Thus, Brouwer implies NE.

The converse is also true: given a compact convex set $D \subseteq \mathbb{R}^N$ and a continuous function $f : D \times D$, one can construct a game in which the strategy set for player 1 is D and, if a NE exists, it must be that $s_1 = f(s_1)$. One such construction can be found here. □

4.3 NE and Randomization.

The following result says that a best response gives positive probability to a pure. There is no standard name for this result; I refer to it as the **Randomization Theorem**.

Theorem 2 (Randomization Theorem)**.** *For any finite game, $\sigma_i \in \mathrm{BR}_i(\sigma_{-i})$ iff for every s_i such that $\sigma_i[s_i] > 0$, $s_i \in \mathrm{BR}_i(\sigma_{-i})$.*

Proof. $\Rightarrow$. Since the set of pure strategies is assumed finite, one of these strategies, call it s_i^*, has highest expected payoff (against σ_{-i}) among all pure strategies. Let the expected payoff of s_i^* be c_i^*. For any mixed strategy σ_i, the expected payoff is the convex sum of the expected payoffs to i's pure strategies. Therefore, c_i^* is the highest possible payoff to any mixed strategy for i, which implies that, in particular, s_i^* is a best response, as is any other pure strategy that has an expected payoff of c_i^*.

I now argue by contraposition. Suppose $s_i \notin \mathrm{BR}_i(\sigma_{-i})$, hence $u_i(s_i, \sigma_{-i}) < c_i^*$. If $\sigma_i[s_i] > 0$, then $u_i(\sigma_i, \sigma_{-i}) < c_i^*$, hence $\sigma_i \notin \mathrm{BR}_i(\sigma_{-i})$, as was to be shown.

$\Leftarrow$. All pure strategies that are best responses must have the same expected payoff, say c_i^*. Therefore, any mixture over pure best responses will have expected payoff c_i^*. And any payoff to any other mixed strategy is a convex sum of expected payoffs to pure strategies, and the maximum such payoff is c_i^*. Therefore, any mixed strategy can have expected payoff at most c_i^*. ■

Theorem 2, the Randomization Theorem, implies that a NE mixed σ_i will give positive probability to two different pure strategies only if those pure strategies each earn the same expected payoff as the mixed strategy. That is, in a NE, a player is *indifferent* between all of the pure strategies that he plays with positive probability. This provides a way to compute mixed strategy NE, at least in principle.

Example 9 (Finding mixed NE in 2 x 2 games). Consider a general 2 x 2 game given by the game box in Figure 8. Let $q = \sigma_2(L)$. Then player 1 is indifferent between

	L	R
T	a_1, a_2	b_1, b_2
B	c_1, c_2	d_1, d_2

Figure 8: *A general 2 x 2 game.*

T and B iff

$$qa_1 + (1-q)b_1 = qc_1 + (1-q)d_1,$$

or

$$q = \frac{d_1 - b_1}{(a_1 - c_1) + (d_1 - b_1)},$$

provided, of course, that $(a_1 - c_1) + (d_1 - b_1) \neq 0$ and that this fraction is in $[0, 1]$.

Similarly, if $p = \sigma_1(T)$, then player 2 is indifferent between L and R if

$$pa_2 + (1-p)c_2 = pb_2 + (1-p)d_2,$$

or

$$p = \frac{d_2 - c_2}{(a_2 - b_2) + (d_2 - c_2)},$$

again, provided that $(a_1 - c_1) + (d_1 - b_1) \neq 0$ and that this fraction is in $[0, 1]$.

Note that the probabilities for player 1 are found by making the *other* player, player 2, indifferent. A common error is to do the calculations correctly, but then to write the NE incorrectly, by flipping the player roles. □

Example 10 (NE in Battle of the Sexes). Battle of the Sexes (Example 3) has two pure strategy NE: (A, B) and (B, A). There is also a mixed strategy NE with

$$\sigma_1(A) = \frac{8 - 0}{(10 - 0) + (8 - 0)} = \frac{4}{9}$$

and

$$\sigma_2(A) = \frac{10 - 0}{(10 - 0) + (8 - 0)} = \frac{5}{9}$$

which I can write as $((\frac{4}{9}, \frac{5}{9}), (\frac{5}{9}, \frac{4}{9}))$; where $(\frac{4}{9}, \frac{5}{9})$, for example, means the mixed strategy in which player 1 plays A with probability $4/9$ and B with probability $5/9$.

The mixed NE here is generally viewed as implausible. It has lower expected payoff, for either player, than either of the two pure NE, and it is unstable under typical stories of how players might adjust their behavior when out of equilibrium. □

Example 11 (NE in Rock-Paper-Scissors). In Rock-Scissor-Paper (Section 2), the NE is unique and in it each player randomizes $1/3$ each across all three strategies. That is, the unique mixed strategy NE profile is

$$\left(\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)\right).$$

It is easy to verify that if either player randomizes $1/3$ across all strategies, then the other payer is indifferent across all three strategies. □

Example 12. Although every pure strategy that gets positive probability in a NE is a best response, the converse is not true: in a NE, there may be best responses that are not given positive probability. A trivial example where this occurs is a game where the payoffs are constant, independent of the strategy profile. Then players are always indifferent and every strategy profile is a NE.

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

Figure 9: *Matching Pennies.*

5 More NE examples.

Example 13. Figure 9 provides the game box for a game called Matching Pennies, which is simpler than Rock-Paper-Scissors but in the same spirit. This game has one equilibrium and in this equilibrium both players randomize 50:50 between the two pure strategies. □

Example 14. Figure 10 provides the game box for one of the best known games in game theory, the Prisoner's Dilemma (PD) (the PD is actually a set of similar games; this is one example). PD is a stylized depiction of friction between joint incentives (the sum of payoffs is maximized by (C, C)) and individual incentives (each player has incentive to play D).

	C	D
C	$4, 4$	$0, 6$
D	$6, 0$	$1, 1$

Figure 10: *A Prisoner's Dilemma.*

It is easy to verify that D is a best response no matter what the opponent does. The unique NE is thus (D, D).

A PD-like game that shows up fairly frequently in popular culture goes something like this. A police officer tells a suspect, "We have your partner, and he's already confessed. It will go easier for you if you confess as well." Although there is usually not enough detail to be sure, the implication is that it is (individually) payoff maximizing not to confess if the partner also does not confess; it is only if the partner confesses that it is payoff maximizing to confess. If that is the case, the game is not a PD but something like the game in Figure 11, where I have written payoffs in a way that makes them easier to interpret as prison sentences (which are bad).

	Silent	Confess
Silent	$0, 0$	$-4, -2$
Confess	$-2, -4$	$-3, -3$

Figure 11: *A non-PD game.*

This game has three NE: (Silent, Silent), (Confess, Confess), and a mixed NE in which the suspect stays silent with probability $1/3$. In contrast, in the (true) PD, each player has incentive to play D regardless of what he thinks the other player will do. □

Example 15 (The Cournot Duopoly). The earliest use of game theoretic reasoning in an economic problem that I know of is the duopoly (i.e., two firm) model in Cournot (1838).[2]

There are two players (firms). For each firm $S_i = [0, 1]$. Note that this is *not* a finite game; one can work with a finite version of the game but analysis of the continuum game is easier in many respects.

For the pure strategy profile $(q_1, q_2) \in [0, 1]$ (q_i for "quantity"), the payoff to player i is,

$$(1 - q_1 - q_2)q_i - cq_i,$$

where $c \in (0, 1)$, provided $q_1 + q_2 \leq 1$. If $q_1 + q_2 > 1$, then the payoff to firm i is $-cq_i$, which is strictly negative if $q_i > 0$.

The idea here is that firm i produces a quantity of a good, best thought of as perishable, and then the market clears somehow, so that each receives a price $1 - q_1 - q_2$, provided the firms don't flood the market ($q_1 + q_2 > 1$); if they flood the market, then the price is zero. I have assumed that the firm's payoff is its profit. There is a constant marginal cost c but no fixed cost.

I will focus on pure strategies.[3] It is easy to verify that for $q_2 \leq 1 - c$, the best response for firm 1 is,

$$\frac{1 - c - q_2}{2}.$$

If $q_2 > 1 - c$, the best response is to produce nothing.

There is a NE in which both firms produce $(1-c)/3$. One can find this by noting that, in pure strategy equilibrium, $q_1^* = (1 - c - q_2^*)/2$ and $q_2^* = (1 - c - q_1^*)/2$; this system of equations has a unique solution with $q_1^* = q_2^* = q^* = (1-c-q^*)/2$, which implies $q^* = (1 - c)/3$. I will show in Example 25 that this NE is the unique NE (even if mixed strategies are considered).

Total output in NE is $2(1-c)/3$ and price is $(1+2c)/3$. The NE outcome of the Cournot game is thus intermediate between monopoly (where quantity is $(1 - c)/2$ and price is $(1 + c)/2$)) and perfect competition (where quantity is $1 - c$ and price is c).

If $c = 0$, then it is also a NE for each firm to produce an output of 1, since the price will be zero no matter what the other firm does, and because output is costless. The presence of this NE, which relies on the unrealistic assumption of zero production costs, illustrates the fact that sometimes a model can be made simpler to

[2]It is not clear what model Cournot actually had in mind; see, for example, Magnan de Bornier (1992). What I'm describing here is the model commonly ascribed to Cournot.

[3]Because of the kink in the payoff function when $q_1 + q_2 = 1$, analyzing mixed strategies is messy.

analyze, in some respects, by adding making the model a tiny bit more complicated (and realistic). □

Example 16 (Bertrand Duopoly). Writing in the 1880s, about 50 years after Cournot, Bertrand objected to Cournot's use of quantity setting on the grounds that quantity competition is unrealistic. In Bertrand's model, firms set prices rather than quantities.

Assume (to make things simple) that there are no production costs, either marginal or fixed. Let p_i denote the price charged by firm i (in particular, p_i is not a probability here). I allow p_i to be any number in $\mathbb{R}_+$, so, once again, this is *not* a finite game.

- $p_i \geq 1$ then firm i gets a payoff of 0, regardless of what the other firm does. Assume henceforth that $p_i \leq 1$ for both firms.
- If $p_1 \in (0, p_2)$ then firm 1 receives a payoff of $(1 - p_1)p_1$. In words, firm 1 gets the entire demand at p_1, which is $1 - p_1$. Revenue, and hence profit, is $(1 - p_1)p_1$. Implicit here is the idea that the good of the two firms are homogeneous and hence perfect substitutes: consumers buy which ever good is cheaper; they are not interested in the name of the seller.
- If $p_1 = p_2 = p$, both firms receive a payoff of $(1 - p)p/2$. In words, the firms split the market at price p.
- If $p_1 > p_2$, then firm 1 gets a payoff of 0.

Payoffs for firm 2 are defined symmetrically. It is not hard to see that firm 1's best response is as follows.

- $p_2 > 1/2$ then firm 1's best response is $p_1 = 1/2$.
- If $p_2 \in (0, 1/2]$ then firm 1's *best response set is empty*. Intuitively, firm 1 wants to undercut p_2 by some small amount, say $\varepsilon > 0$, but $\varepsilon > 0$ can be made arbitrarily small. Technically, the existence failure for best response arises because of a discontinuity in the payoff function. One can address the existence failure by requiring that prices lie on a finite grid (be expressed in pennies, for example), but this introduces other issues.
- If $p_2 = 0$, then any $p_1 \in \mathbb{R}_+$ is a best response.

Notwithstanding the fact that best response isn't defined over much of the domain, there is a pure NE, namely the price profile $(0, 0)$.

The Bertrand model thus predicts the competitive result with only two firms. At first this may seem paradoxical, but remember that I am assuming that the goods are homogeneous. The bleak prediction of the Bertrand game (bleak from the viewpoint of the firms; consumers do well and the outcome is efficient) suggests that firms have strong incentive to differentiate their products, and also strong incentive to collude (which may be possible if the game is repeated, as I discuss later in the course).

Example 17. (Bertrand-Edgeworth Duopoly.) Francis Edgeworth, writing in the 1890s, about ten years after Bertrand, argued that firms often face capacity constraints, and these constraints can affect equilibrium analysis. Continuing with the Bertrand duopoly of Example 16, suppose that firms can produce up to 1/2 at zero cost but then face infinite cost to produce any additional output. This would model a situation in which, for example, each firm had stockpiled a perishable good; this period, a firm can sell up to the amount stockpiled, but no more, at zero opportunity cost (zero since any amount not sold melts away in the night).

Suppose that sales of firm 1 (taking into consideration the capacity constraint and ignoring the trivial case with $p_1 > 1$) are given by,

$$q_1(p_1, p_2) = \begin{cases} 0 & \text{if } p_1 > p_2 \text{ and } p_1 > 1/2, \\ 1/2 - p_1 & \text{if } p_1 > p_2 \text{ and } p_1 \leq 1/2, \\ (1 - p_1)/2 & \text{if } p_1 = p_2 \leq 1, \\ 1 - p_1 & \text{if } p_1 < p_2 \text{ and } p_1 > 1/2, \\ 1/2 & \text{if } p_1 < p_2 \text{ and } p_1 \leq 1/2. \end{cases}$$

The sales function for firm 2 is analogous.

For interpretation, suppose that there are a continuum of consumers, with a consumer whose "name" is $\ell \in [0, 1]$ wanting to buy at most one unit of the good at a price of at most $v_\ell = 1 - \ell$. Then, integrating across consumers, market demand facing a monopolist would be $1 - p$, as specified.

With this interpretation, consider the case $p_1 > p_2$ and $p_1 \leq 1/2$. Then $p_2 < 1/2$ and demand for firm 2 is $1 - p_2 > 1/2$. But because of firm 2's capacity constraint, it can only sell 1/2, leaving $(1 - p_2) - 1/2 > 0$ potentially for firm 1, provided firm 1's price is not too high. Suppose that the customers who succeed at buying from firm 2 are the customers with the highest valuations (they are the customers willing to wait on line, for example). Then the left over demand facing firm 1, called *residual* demand, is $1/2 - p_1$.

Note that this specification of residual demand is a substantive assumption about how demand to the low priced firm gets rationed. We could instead have assumed that customers are assigned to firm 2 randomly. That would have resulted in a different model.

In contrast to the Bertrand model, it is no longer an equilibrium for both firms to charge $p_1 = p_2 = 0$. In particular, if $p_2 = 0$, then firm 1 is better off setting $p_1 = 1/4$, which is the monopoly price for residual demand. In fact, there is no pure strategy equilibrium here. There is, however, a symmetric mixed strategy equilibrium, which can be found as follows.

Note first that the maximum profit from residual demand is 1/16 ($= 1/4 \times 1/4$). A firm can also get a profit of 1/16 by being the low priced firm at a price of 1/8 (and selling 1/2). The best response for firm 1 to pure strategies of firm 2 is then,

$$\begin{aligned} 1/2 &\text{ if } p_2 > 1/2, \\ \emptyset &\text{ if } p_2 \in (1/8, 1/2], \\ 1/4 &\text{ if } p_2 \leq 1/8. \end{aligned}$$

The best response for firm 2 to pure strategies of player 1 is analogous.

Intuitively, the mixed strategy equilibrium will involve a continuous distribution, and not put positive probability on any particular price, because of the price undercutting argument that is central to the Bertrand game. So, I look for a symmetric equilibrium in which firms randomize over an interval of prices. The obvious interval to look at is $[1/8, 1/4]$.

For it to be an equilibrium for firm 1 to randomize over the interval $[1/8, 1/4]$, firm 1 has to get the same profit from any price in this interval.[4] The profit at either $p_1 = 1/8$ (where the firm gets capacity constrained sales of 1/2) or $p_1 = 1/4$ (where the firm sells 1/2-1/4=1/4 to residual demand) is 1/16. Thus for any other $p_1 \in (1/8, 1/4)$ it must be that, letting $F_2(p_1)$ be the probability that player 2 charges a price less than or equal to p_1,

$$F_2(p_1)p_1(1/2 - p_1) + (1 - F_2(p_1))p_1(1/2) = 1/16,$$

where the first term is profit when $p_2 < p_1$ and the second term is profit when $p_2 > p_1$; because the distribution is continuous, the probability that $p_2 = p_1$ is zero and so I can ignore this case.

Solving for F_2 yields

$$F_2(p_1) = \frac{1}{2p_1} - \frac{1}{16p_1^2}.$$

The symmetric equilibrium then has $F_1 = F_2$, where F_2 is given by the above expression. The associated density is

$$f_2(p_1) = \frac{1}{8p_1^3} - \frac{1}{2p_1^2}$$

and is graphed in Figure 17. This makes some intuitive sense. If we imagine this equilibrium being repeated, then most of the time prices will be close to 1/8, reflecting the undercutting logic of the Bertrand game, but from time to time one of the prices will be close to 1/4.

Edgeworth argued that in a situation such as this, with no pure strategy equilibrium, prices would just bounce around. Nash equilibrium makes a more specific prediction: the distribution of prices is generated by the mixed strategy equilibrium.

There is an intuition, first explored formally in Kreps and Scheinkman (1983), that the qualitative predictions of Cournot's quantity model can be recovered in a

[4]In a continuum game like this, this requirement can be relaxed somewhat, for measure theoretic reasons, but this is a technicality as far as this example is concerned.

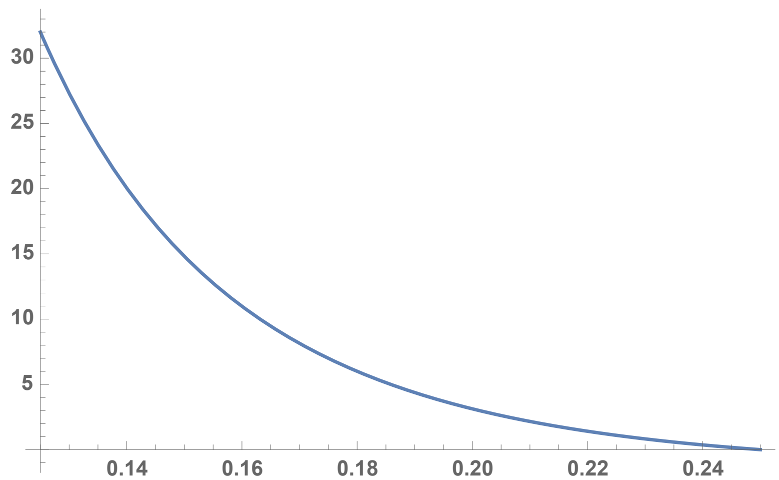

Figure 12: Density for the NE in the Bertrand-Edgeworth Game example.

more realistic, multi-period setting in which firms make output choices in period 1 and then sell that output in period 2. The Edgeworth example that I just worked through is an example of what the period 2 analysis can look like. If the capacity constraints had been, instead, 1/3 each (the equilibrium quantities in the Cournot game), then you can verify that the equilibrium of the Bertrand-Edgeworth game would, in fact, have been for both firms to charge a price of 1/3. □

6 Dominance and Rationalizability.

6.1 Strict Dominance.

Informally, σ_i strictly dominates $\hat{\sigma}_i$ iff σ_i always generates strictly higher expected payoff than $\hat{\sigma}_i$, regardless of the strategies of the other players.

Definition 2.

- σ_i strictly dominates $\hat{\sigma}_i$ *iff for any profile of opposing pure strategies* s_{-i}*,*

$$u_i(\sigma_i, s_{-i}) > u_i(\hat{\sigma}_i, s_{-i}).$$

- $\hat{\sigma}_i$ *is* strictly dominated *iff there exists a strategy that strictly dominates it.*

σ_i strictly dominates $\hat{\sigma}_i$ iff σ_i has strictly higher expected payoff against any opposing mixed strategy profile.

Theorem 3. *σ_i strictly dominates $\hat{\sigma}_i$ iff for any mixed strategy profile $\sigma_{-i} \in \Sigma_{-i}$,*

$$u_i(\sigma_i, \sigma_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}).$$

Proof. Almost immediate because of the linear nature of expected payoffs:

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})\sigma_{-i}[s_{-i}].$$

and similarly for $u_i(\hat{\sigma}_i, \sigma_{-i})$. ■

Remark 6. If a pure strategy s_i is strictly dominated, then it will, in particular be strictly dominated by a strategy that gives zero probability weight to s_i itself. This can somewhat expedite the search for a strictly dominating (mixed) strategy. □

Definition 3. *A strategy is* strictly dominant *iff it strictly dominates every other strategy.*

Since two strategies cannot strictly dominate each other, there can be at most one strictly dominant strategy. For a similar reason, a strictly dominant strategy, if it exists, must be pure.

A strategy that is strictly dominant is unambiguously good: no matter what you think the other players might do, you want to play the strictly dominant strategy, provided one exists. Conversely, a strategy that is strictly dominated is unambiguously inferior: there is another strategy that always does better, regardless of what the other players do. In particular, a strictly dominated strategy can never be given positive probability in a Nash equilibrium.

Example 18. Recall the Prisoner's Dilemma of Example 14 and Figure 10. Then, for either player, C is strictly dominated by D; D is strictly dominant. □

Example 19. Consider the game in Figure 13 (I've written payoffs only for player 1); I also used this game in Example 8 in Section 4.1.

	L	R
T	10	0
M	0	10
B	4	4

Figure 13: *A Dominance Example.*

B is strictly dominated for player 1 by the mixed strategy $(1/2, 1/2, 0)$ (which gets a payoff of 5 no matter what player 2 does, whereas B always gets 4). There is no strictly dominant strategy.

Note that B is not strictly dominated by any pure strategy. So it is important in this example that I allow for dominance by a mixed strategy. □

Example 20. Consider the game in Figure 14.

B is not strictly dominated for player 1. In particular, focusing on mixed strategies that give zero weight to B (See Remark 6), any mixture over T and M that gives payoff at least 6 to L gives a payoff less than 4 against R, and vice versa. □

	L	R
T	10	0
M	0	10
B	6	6

Figure 14: *Another Dominance Example.*

6.2 Strict Dominance and Never a Best Response.

Definition 4. *A mixed strategy $\hat{\sigma}_i$ is* never a best response (NBR) *iff there does not exist a profile of opposing mixed strategies $\sigma_{-i} \in \Sigma_{-i}$ such that $\hat{\sigma}_i \in \mathrm{BR}_i(\sigma_{-i})$.*

Note that NBR is distinct from the condition recorded in Theorem 3 because the definition of NBR does not immediately imply that if σ_i is NBR then there is a strategy that strictly dominates it. The conjecture that this is in fact true is valid when $N = 2$ (there are only two players) but not when $N \geq 3$. I first consider the easier case, when $N = 2$.

Theorem 4. *Suppose that $N = 2$. A strategy $\hat{\sigma}_i$ is strictly dominated iff it is NBR.*

Proof. This is a special case of Theorem 5 below, but note in passing that the $\Rightarrow$ direction is implied by Theorem 3. It is the $\Leftarrow$ direction that is difficult. ■

Example 21. Consider again the game in Figure 13 in Example 19. As already noted in Example 19, B is strictly dominated for player 1 by the mixed strategy $(1/2, 1/2, 0)$ (which gets a payoff of 5 no matter what player 2 does, whereas B always gets 4).

One can similarly check that B is NBR by considering all mixtures over L and R. Let q be the probability on L. If $q > 1/2$ then T is strictly optimal. If $q < 1/2$ then M is strictly optimal. If $q = 1/2$ then any mixture over T and M is optimal. But B is never optimal. □

Example 22. Consider the game in Figure 14. As already noted in Example 20, B is not strictly dominated.

Let q be the probability on L. Then B is a best response if $q \in [2/5, 3/5]$; it is the unique best response for $q \in (2/5, 3/5)$.

Note that B is not a best response to any pure strategy for player 2. So it is important in this example that I allow for mixed strategies by player 2. □

As these examples illustrate, the checks for NBR and for being strictly dominated are, loosely speaking, mirror images of each other.

- For NBR, one has to consider all possible mixtures over opponent strategies, but it suffices to focus on pure strategies for oneself (by Theorem 2, the Randomization Theorem, if a strategy is not a best response then, in particular, there is a pure strategy that has higher expected payoff).

- For strictly dominated, it suffices to consider pure strategies for the opponent, but to find a dominating strategy one may have to consider mixtures for oneself.

When $N \geq 2$, we can encounter the issue illustrated in the next example.

Example 23. Figure 15 gives payoffs for player 1 in a three-player game in which player chooses between L (Left), M, and R, not explicitly labeled in the diagram, player 2 chooses between $2A$ and $2B$, and player 3 chooses between $3A$ and $3B$.

	$3A$	$3B$
$2A$	5	0
$2B$	0	5

	$3A$	$3B$
$2A$	8	8
$2B$	8	0

	$3A$	$3B$
$2A$	0	8
$2B$	8	8

Figure 15: *NBR and Strict Dominance When $N > 2$.*

One can verify that in this game, L is NBR. But L is *not* strictly dominated: any mixture over M and R that yields more than 5 against $(2A, 3A)$ yields less than 5 against $(2B, 3B)$. □

In Example 23, even though L is NBR, L is nevertheless a best response to the *correlated distribution* in which $(2A, 3A)$ and $(2B, 3B)$ are played with equal probability. This leads to the correct conjecture that a version of Theorem 4 holds for any N provided we weaken NBR to allow correlation by the opposing players. I discuss below what this correlation might represent.

Somewhat abusing notation, let $\mathrm{BR}_i(\sigma^c_{-i})$ denote the set of strategies for player i that are best responses to the (possibly correlated) distribution $\sigma^c_{-i} \in \Sigma^c_{-i} = \Delta(S_{-i})$.

Definition 5. *Fix a game. A mixed strategy $\hat{\sigma}_i$ is* never a best response to any (possibly correlated) distribution (CNBR) *iff there does not exist a distribution over opposing strategies $\sigma^c_{-i} \in \Sigma^c_{-i}$ such that $\hat{\sigma}_i \in \mathrm{BR}_i(\sigma^c_{-i})$.*

The following result reduces to Theorem 4 when $N = 2$.

Theorem 5. *In any finite game, a strategy σ_i is strictly dominated iff it is CNBR.*

Proof. $\Rightarrow$. This follows from a minor modification of the proof of the $\Rightarrow$ direction of Theorem 3.

$\Leftarrow$. Suppose that $\hat{\sigma}_i$ is CNBR. I need to show the existence of a mixed strategy σ_i that strictly dominates $\hat{\sigma}_i$.

For each $\sigma^c_{-i} \in \Sigma^c_{-i}$, construct the vector $w(\sigma^c_{-i}) \in \mathbb{R}^{|S_i|}$ where the coordinate corresponding to pure strategy s_i is given by

$$u_i(s_i, \sigma^c_{-i}) - u_i(\hat{\sigma}_i, \sigma^c_{-i}).$$

Let W be the set of $w(\sigma^c_{-i})$ for all $\sigma^c_{-i} \in \Sigma^c_{-i}$. W is equal to the convex hull of the finite set $w(s_{-i})$, for all $s_{-i} \in S_{-i}$, and is therefore compact and convex.

Because $\hat{\sigma}_i$ is CNBR, every $w \in W$ has at least one strictly positive coordinate. Therefore W has an empty intersection with the negative orthant $\mathbb{R}^K_-$, which is a closed and convex set. By the standard Separating Hyperplane Theorem (see my notes on convex sets), there is a $\lambda \in \mathbb{R}^K$ and an r such that for any $b \in \mathbb{R}^K_-$ and any $w \in W$,

$$\lambda \cdot b < r < \lambda \cdot w.$$

Note that this implies $\lambda \neq 0$. Since the origin is contained in $\mathbb{R}^K_-$, the inequalities imply $r > 0$. Since the negative unit vectors are contained in $\mathbb{R}^K_-$, the inequalities imply $\lambda > 0$ (i.e., weakly positive in all coordinates and strictly positive in at least one coordinate), since otherwise one could take b to be a positive multiple of the negative unit vector corresponding to a coordinate for which λ is negative, violating the inequalities. Finally, define

$$\sigma_i[s_i] = \lambda[s_i] / \sum_{t_i \in S_i} \lambda[t_i].$$

By construction, $\sum_{s_i \in S_i} \sigma_i[s_i] w(\sigma^c_{-i}) > 0$ for every σ^c_{-i}, which implies that σ_i strictly dominates $\hat{\sigma}_i$. ■

Does correlation make sense? Suppose that one adopts the methodological viewpoint, which is standard in classical game theory, that players act independently. If they appear to an outsider observer to correlate, then that is because they have access to a correlation device, such as a publicly observable coin toss. And if that is the case, then the description of the game was incomplete: the correlation device should have been modeled as part of the game. The notes on extensive form games provide a framework for doing this modeling explicitly. In this modified game, we are back to players acting independently after observing the output of the correlation device.

Arguing against this viewpoint is the idea that player i, in choosing which strategy to play, is optimizing with respect to beliefs about the strategies chosen by the other players. That is, in discussing player i's behavior, we should interpret σ_{-i} or σ^c_{-i} as player i's beliefs about the behavior of the other players rather than a factual statement as to what the other players are actually doing. Player i's beliefs could exhibit correlation, even if player i is certain that the other player's are acting independently. A standard story along these lines is that player i believes that some feature of the other players' background, such as going through the same school system, might serve as a de facto correlation device. And even if, in principle, you feel that such de facto correlation devices should be modeled explicitly, in practice this is infeasible.

If you agree that a player's beliefs about other players may exhibit correlation, then it make sense *not* to dismiss a strategy as unambiguously bad if it is NBR but not CNBR. And by Theorem 5, this matches with strict dominance arguments. In Example 23, for instance, L is *not* strictly dominated even though it is NBR.

6.3 Rationalizability.

If a pure strategy s_i is NBR then it cannot be part of any NE. Therefore, when searching for NE, one can simplify the search by first deleting NBR strategies. This motivates the following procedure.

Let $S_i^1 \subseteq S_i$ denote the strategy set for player i after deleting all pure strategies for i that are NBR. S_i^1 is not empty since, as noted in Section 4.1, the best response correspondence is not empty-valued. The S_i^1 form a new game, and for that game we can ask whether any strategies are NBR. Deleting those strategies, we get $S_i^2 \subseteq S_i^1 \subseteq S_i$, which again is not empty. And so on.

Since the game is finite, this process terminates after a finite number of rounds in the sense that there is a t such that for all i, $S_i^{t+1} = S_i^t$. Let S_i^R denote this terminal S_i^t and let S^R be the product of the S_i^R. The strategies S_i^R are called the *rationalizable* strategies for player i; S^R is the set of *rationalizable (pure) strategy profiles*. Rationalizability was first introduced into game theory in Bernheim (1984) and Pearce (1984).

Theorem 6. *Fix a game. If σ is a NE and $\sigma_i[s_i] > 0$ then $s_i \in S_i^R$.*

Proof. Almost immediate by induction and Theorem 2, the Randomization Theorem. ■

Remark 7. The construction of S^R uses maximal deletion of NBR strategies at each round of deletion. This does not matter: one can show that if, whenever possible, one deletes at least one strategy for at least one player at each round, then eventually the deletion process terminates, and the terminal S^t is S^R. □

Example 24. In the Prisoner's Dilemma (Example 14, Figure 10), C is NBR. D is the unique rationalizable strategy. □

Example 25. Consider the Cournot duopoly (Example 15). One can verify that $q_i = (1-c)/2$ strictly dominates $q_i > (1-c)/2$ and hence any $q_i > (1-c)/2$ is NBR. Deleting $q_i > (1-c)/2$ for both firms/players, one can verify that $q_i = (1-c)/4$ strictly dominates $q_i < (1-c)/4$. Continuing in this way eventually eliminates every strategy other than $q_i = (1-c)/3$. The pure strategy profile in which both firms produce $(1-c)/3$ is thus the unique NE.

This game has an infinite strategy set and this iterative deletion procedure takes an infinite number of steps. Hence, strictly speaking, this example is not covered by the above formal discussion. It is immediate, however, that any strategy that is deleted by iterated deletion of strictly dominated strategies, even in this infinite setting, cannot be part of a Nash equilibrium. □

A natural, but wrong, intuition is that any strategy in S_i^R gets positive probability in at least one NE.

Example 26. Consider the game in Figure 16. $S_i^R = \{A, B, C\}$ for either player. But the unique NE can be shown to be (A, A). □

	A	B	C
A	$4,4$	$1,1$	$1,1$
B	$1,1$	$2,-2$	$-2,2$
C	$1,1$	$-2,2$	$2,-2$

Figure 16: *Rationalizability and NE.*

A related point can be illustrated by Rock-Paper-Scissors (RPS; Section 2). All three actions in RPS are rationalizable, and each is in fact played with positive probability in the NE. But the NE provides a more precise prediction for RPS than merely that all three strategies will be played: the NE predicts that the three strategies will each be played with equal probability.

6.4 Iterated Strict Dominance Deletion, Rationalizability, and Correlated Rationalizability.

The results in Section 6.2 imply that, when $N = 2$, we can construct S^R by either iterated deletion of NBR strategies or iterated deletion of strictly dominated strategies. When $N > 3$, however, the set of strategies that are NBR may be larger than the set of strategies that are strictly dominated, and this implies that S^R may be strictly smaller than the set obtained by iterated deletion of strictly dominated strategies. I refer to the latter set, which corresponds to the set obtained by iterated deletion of CNBR strategies, as the *correlated rationalizable set*, denoted S^{CR}.

If the goal of working with rationalizability is purely to narrow the search for the NE, then S^R is the right version of rationalizability to use. But S^{CR} is arguably the right version of rationalizability to use if the goal is instead to understand what players might "reasonably" do based on introspective reasoning of the form, "doing *this* is optimal for me provided I think my opponents will do *that*, and doing *that* is optimal for my opponents provided that they think . . . ." But these are deep waters. See Brandenburger and Dekel (1987).

Another way to say this is that if you find iterated deletion of strictly dominated strategies compelling, then know that this procedure leaves you with S^{CR}, which may be strictly larger than S^R if $N > 2$.

6.5 Weak Dominance and Admissibility.

Weak dominance is defined analogously to strict dominance.

Definition 6.

- σ_i weakly dominates $\hat{\sigma}_i$ *iff for any profile of opposing pure strategies* s_{-i}*,*

$$u_i(\sigma_i, s_{-i}) \geq u_i(\hat{\sigma}_i, s_{-i}).$$

with strict inequality for at least one s_{-i}*.*

- $\hat{\sigma}_i$ *is* weakly dominated *iff there is a strategy that weakly dominates it.*

Thus, if σ_i strictly dominates $\hat{\sigma}_i$, then it weakly dominates it, but not necessarily conversely. The analog of Theorem 3, with an almost identical proof, is the following.

Theorem 7. *σ_i weakly dominates $\hat{\sigma}_i$ iff for any* fully mixed *σ_{-i},*

$$u_i(\sigma_i, \sigma_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}).$$

Informally, a weakly dominated strategy is imprudent: while it might be a best response to some *pure* s_{-i}, it will not be a best response if i thinks that every s_{-i} is possible, even if some receive only very low probability.

Definition 7. *A strategy is* weakly dominant *iff it weakly dominates every other strategy.*

Since two strategies cannot weakly dominate each other, there can be at most one weakly dominant strategy. For a similar reason, a weakly dominant strategy, if it exists, must be pure.

Definition 8. *A NE is* admissible *iff it does not put positive weight on a weakly dominated strategy.*

Admissibility is an example of *refinement* of Nash equilibrium. A refinement takes the full set of NE and then eliminates equilibria that are, in some sense, unreasonable. I discuss admissibility further in separate notes on refinements.

Example 27. (An Entry Deterrence Game.) Consider the game of Figure 17.

	A	F
I	$15, 15$	$-1, -1$
O	$0, 35$	$0, 35$

Figure 17: Admissibility.

This game can be interpreted as a simple model of entry deterrence in a market with an incumbent monopolist, player two. A potential entrant, player one, has to decide whether to enter (I or O). If the entrant stays out, the incumbent gets the monopoly profit. But if the entrant goes in, the incumbent can either acquiesce (A) or fight (F).

This game has infinitely many NE. One NE is (I, O). There is also an infinite set of NE of the form (O, q), where q is the probability that player 2 plays A, satisfies $q \leq 1/16$. In the latter equilibria, the entrant is deterred from choosing I by the incumbent threatening that, should that happen, the incumbent will choose F with probability at least $15/16$, leaving the entrant with a non-positive expected payoff.

All of the (O, q) are inadmissible because F is weakly dominated for the incumbent: F only makes sense for the incumbent if the incumbent is certain the entrant won't come in. The standard interpretation is that F represents a threat that is not credible for the payoffs as given. The only admissible NE is (I, O).

It is worth underscoring that the point here is not that the (O, q) profiles are not NE, because they are, but rather all of these NE are arguably implausible. Put differently, even if you buy into the idea that non-NE are unreasonable, this example shows that many NE profiles may also be unreasonable. □

Recall (Theorem 4) that when $N = 2$, a strategy is strictly dominated iff it is never a best response (NBR). Recall also (Theorem 5) that this equivalence extends to $N > 2$ provided one weakens NBR to CNBR, in effect allowing other players to correlate. The analog of this result for weak dominance is the following.

Theorem 8. *In any finite game, a strategy is weakly dominated iff it is not a best response to any* correlated *strategy profile that is* fully mixed*.*

Proof. There is a proof in Myerson (1991) ■

Finally game theorists sometimes use the term "dominance" in an ambiguous manner. Because there are differences between weak and strong dominance (for example, order can matter for the *iterated* deletion of weakly dominated strategies), I try to be explicit about which concept I am using.

7 Other Topics.

7.1 How many NE are there?

Fix players and strategy sets and consider all possible assignments of payoffs. With N players and $|S|$ pure strategy profiles, payoffs can be represented as a point in $\mathbb{R}^{N|S|}$. Wilson (1971) proved that there is a precise sense in which the set of Nash equilibria is finite and odd (which implies existence, since 0 is not an odd number) for most payoffs in $\mathbb{R}^{N|S|}$. As we have already seen, Rock-Paper-Scissors has one equilibrium, while Battle of the Sexes has three. A practical implication of the finite and odd result, and the main reason why I am emphasizing this topic in the first place, is that if you are trying to find all the equilibria of the game, and so far you have found four, *then you are probably missing at least one.*

Although the set of NE in finite games is typically finite and odd, there are exceptions.

Example 28 (A game with two NE). A somewhat pathological example is the game in Figure 18. This game has exactly two Nash equilibria, (A, A) and (B, B). In particular, there are no mixed strategy NE: if either player puts any weight on A than the other player wants to play A for sure.

Note that this game box exhibits ties in the payoffs to the players. For example, the two cells corresponding to player 1 choosing B both have the same payoffs for

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

Figure 18: *A game with two Nash equilibria.*

both players. Because of these ties, this game violates the condition in Wilson (1971). □

The previous example seems somewhat artificial, and in fact it is extremely unusual in applications to encounter a game where the set of NE is finite and even. But it is quite common to encounter strategic form games that have an infinite number of NE; such games are common when the strategic form represents a game with a non-trivial temporal structure.

Example 29. Recall the entry deterrence game of Example 27. As discussed there, this game has an infinite number of NE, although only one of these NE, (I, A), is admissible.

The usual interpretation of this game is that it is the strategic form representation of an interaction in which player 1, the potential entrant into a market dominated by player 2, moves first. If player 1 chooses not to enter (chooses O), then the planned/threatened response by player 2 becomes irrelevant. This leads to a payoff tie in the O row of the game box. Because of this payoff tie, this game violates the condition in Wilson (1971). □

Assuming that the number of equilibria is in fact finite and odd, how many equilibria are typical? In 2×2 games, the answer is 1 or 3. What about in general? A lower bound on how large the set of NE can possibly be in finite games is provided by $L \times L$ games (two players, each with L strategies) in which, in the game box representation, payoffs are (1, 1) along the diagonal and (0, 0) elsewhere. In such games, there are L pure strategy Nash equilibria, corresponding to play along the diagonal, and an additional $2^L - (L+1)$ fully or partly mixed NE. The total number of NE is thus $2^L - 1$, which is the number of non-empty strategy subsets. This example is robust; payoffs can be perturbed slightly and there will still be $2^L - 1$ NE.

This is extremely bad news. First, it means that the maximum number of NE is growing exponentially in the size of the game. This establishes that the general problem of calculating all of the equilibria is computationally intractable. Second, it suggests that the problem of finding even one NE may, in general, be computationally intractable; for work on this, see Daskalakis et al. (2009). Note that the issue is not whether algorithms exist for finding one equilibrium, or even all equilibria. For finite games, there exist many such algorithms. The problem is that the time taken by these algorithms to reach a solution can grow explosively in the size of the game. For results on the number of equilibria in finite games, see

McLennan (2005).

7.2 Zero-Sum Games.

Recall that a two-player game is *zero-sum* iff $u_1 + u_2 = 0$, or $u_2 = -u_1$. It is uncommon for games in economic applications to be zero sum (or, more generally, strictly competitive, as defined in Remark 15 below). Nevertheless, the theory for zero-sum games is often useful. One reason is that results from the theory of zero sum games can sometimes be applied to other contexts by constructing an artificial zero-sum game. I give examples of this is in Remark 10 and Remark 13 below.

For any two-player game, not necessarily zero-sum, σ_i is a *security strategy* for player i iff σ_i solves

$$\max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}).$$

Informally, it as if the game has been changed to one with a temporal structure: player i chooses a mixed strategy, and then the other player moves after seeing i's choice of mixed strategy. Player i's security strategy maximizes i's payoff in this modified game under the assumption that the other player always acts to minimize i's payoff (which is consistent with the other player's payoffs in a zero-sum game but not in general). Player i's *security payoff* is

$$V_i = \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}).$$

Remark 8. Without loss of generality, one can compute i's security strategy by considering only pure strategies for the other player. The reason is that for any σ_i, σ_{-i} solves the minimization problem

$$\min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$$

iff, for every s_{-i} such that $\sigma_{-i}[s_{-i}] > 0$, s_{-i} is also a solution. □

Example 30. In Rock-Paper-Scissors (RPS), the security strategy for either player is to randomize 1/3 on each pure strategy; the security payoff is 0. Note that the security strategies are the Nash equilibrium strategies and the security values are the NE expected payoffs. Theorem 9 below will document that this is a general feature of zero-sum games.

As noted in Remark 8, one can compute player 1's security strategy using only pure strategies for player 2. But it is essential that player 1 be able to randomize and that player 2 not be able to observe the pure strategy that player 1 eventually chooses. If these assumptions are violated, then player 2 can drop player 1's payoff in RPS all the way down to -1. □

Example 31. Consider the Battle of the Sexes game of Example 3. In this game, the security strategy for player 1 is $(\frac{5}{9}, \frac{4}{9})$ and the security strategy for player 2 is $(\frac{4}{9}, \frac{5}{9})$. This is the mirror image of this game's mixed strategy NE (see Example 10).

Note that if player 1 plays the security strategy $(\frac{5}{9}, \frac{4}{9})$, then it is optimal for player 2 to play A: player 2's security strategy is not a best response to player 1's security strategy. For this reason, the profile of security strategies is, arguably, not a plausible prediction for this game. In fact, a case can be made that the mixed strategy NE is not plausible either. The plausible predictions, arguably, are the two two pure NE, (A, A) and (B, B), and distributions over those (e.g., play either pure strategy NE with probability 1/2); see the discussion in Example 3 and Example 10.

Note that Battle of the Sexes is not zero-sum, nor can it be derived from a zero-sum game via the operations discussed in Section 4.1. The problem with using the security strategy in games that are not zero-sum is that it fails to take into consideration the actual incentives of the other player. □

Theorem 9. *Consider any finite zero-sum game.*

1. *If σ^* is a NE, then for each i*

$$u_i(\sigma^*) = V_i. \tag{1}$$

In particular, $V_1 = -V_2$ and all NE have the same payoffs.

2. *σ^* is a NE iff it is a profile of security strategies.*

Proof. Let $\sigma^* = (\sigma_1^*, \sigma_2^*)$ be a Nash equilibrium. By the definition of NE,

$$u_1(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*).$$

Moreover,

$$\max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \geq \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2).$$

The analogous inequalities for player 2 yield

$$u_2(\sigma_1^*, \sigma_2^*) = \max_{\sigma_2} u_2(\sigma_1^*, \sigma_2) \geq \min_{\sigma_1} \max_{\sigma_2} u_2(\sigma_1, \sigma_2).$$

Since $u_2 = -u_1$, this implies,

$$u_1(\sigma_1^*, \sigma_2^*) = \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \leq \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2).$$

Putting this all together:

$$\begin{aligned} \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2) &\geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \\ &= u_1(\sigma_1^*, \sigma_2^*) \\ &= \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \\ &\geq \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2). \end{aligned} \tag{2}$$

On the other hand, for any $\hat{\sigma}_1$ and $\hat{\sigma}_2$,

$$u_1(\hat{\sigma}_1, \hat{\sigma}_2) \geq \min_{\sigma_2} u_1(\hat{\sigma}_1, \sigma_2)$$

which implies

$$\max_{\sigma_1} u_1(\sigma_1, \hat{\sigma}_2) \geq \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2).$$

Applying the analogous argument to u_2, and again using the fact that $u_2 = -u_1$, yields,

$$\min_{\sigma_2} u_1(\hat{\sigma}_1, \sigma_2) \leq \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2).$$

Setting $\hat{\sigma}_i = \sigma_i^*$ and combining yields

$$\begin{aligned} \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2) &\geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \\ &= u_1(\sigma_1^*, \sigma_2^*) \\ &= \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \\ &\geq \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2). \end{aligned} \tag{3}$$

Combining Inequality 2 and Inequality 3,

$$\begin{aligned} \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2) &= \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \\ &= u_1(\sigma_1^*, \sigma_2^*) \\ &= \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \\ &= \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2). \end{aligned} \tag{4}$$

Equality 4 establishes Equality 1 for player 1. The claim for player 2 follows by using $u_2 = -u_1$.

Still assuming that σ^* is a NE, the above conclusion that

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

implies that σ_1^* is a security strategy for player 1. Again, the claim for player 2 follows by using $u_2 = -u_1$.

Finally, for any $\hat{\sigma}_1$, $\hat{\sigma}_2$,

$$u_1(\hat{\sigma}_1, \hat{\sigma}_2) \geq \min_{\sigma_2} u_1(\hat{\sigma}_1, \sigma_2)$$

$$\max_{\sigma_1} u_1(\sigma_1, \hat{\sigma}_2) \geq u_1(\hat{\sigma}_1, \hat{\sigma}_2)$$

If $\hat{\sigma}_1$ and $\hat{\sigma}_2$ are security strategies then

$$\min_{\sigma_2} u_1(\hat{\sigma}_1, \sigma_2) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

$$\min_{\sigma_1} u_2(\sigma_1, \hat{\sigma}_2) = \max_{\sigma_2} \min_{\sigma_1} u_2(\sigma_1, \sigma_2)$$

Since $u_2 = -u_1$, the latter implies

$$\max_{\sigma_1} u_1(\sigma_1, \hat{\sigma}_2) = \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

By Theorem 1, a NE exists (this is a critical step in the proof; see Remark 12) which means that Equation 4 holds and in particular that

$$\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2). \tag{5}$$

Putting this all together, we get

$$u_1(\hat{\sigma}_1, \hat{\sigma}_2) \geq \min_{\sigma_2} u_1(\hat{\sigma}_1, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \hat{\sigma}_2) \geq u_1(\hat{\sigma}_1, \hat{\sigma}_2)$$

which implies

$$u_1(\hat{\sigma}_1, \hat{\sigma}_2) = \max_{\sigma_1} u_1(\sigma_1, \hat{\sigma}_2),$$

which says that σ_1 is a best response to $\hat{\sigma}_2$. Analogously, σ_2 is a best response to $\hat{\sigma}_1$, and so $(\hat{\sigma}_1, \hat{\sigma}_2)$ is a NE. ■

Remark 9. Depending on the author, the term "MinMax Theorem" (or, more frequently, MiniMax Theorem) refers to some combination of Theorem 9 and

$$\min_{\sigma_2} \max_{\sigma_1} u_2(\sigma_1, \sigma_2) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2),$$

which was Equation 5 in the proof of Theorem 9.

Theorem 9 and Equation 5 were first established in von Neumann (1928) and were among the first general results in game theory.[5] □

Remark 10. A version of Equation 5 is true for general finite games, rather than just zero-sum games. To see this, consider a general finite game and construct a new zero-sum game with two players, i and $-i$. Player $-i$ has the pure strategy set S_{-i} and payoff function $\hat{u}_{-i} = -u_i$. A mixed strategy for player $-i$ is an element of $\Delta(S_{-i})$; if there are three or more players in the original game, a mixed strategy for player $-i$ in the new game corresponds to a correlated distribution over the pure strategies of the other players in the original game. Then Equation 5 implies that

$$\min_{\sigma^c_{-i}} \max_{\sigma_i} u_i(\sigma_i, \sigma^c_{-i}) = \max_{\sigma_i} \min_{\sigma^c_{-i}} u_i(\sigma_i, \sigma^c_{-i}).$$

Remark 11. This, in turn, raises the question of when in general it is true that

$$\min_{y} \max_{x} f(x, y) = \max_{x} \min_{y} f(x, y)$$

The answer, due to Nikaido (1954) and Sion (1958), is: whenever the domain of f is a topological vector space, f is quasiconcave and upper semi-continuous in x and quasiconvex and lower semi-continuous in y. All of these conditions are satisfied trivially in the case of finite games. □

[5]Yes, von Neumann (1928) is in German; no, I haven't read it in the original.

Remark 12. The proof given for Theorem 9 relied on Theorem 1, ensuring the existence of NE. Theorem 9, however, predates Nash Equilibrium and can be proved using tools that are in sense more elementary than the fixed point argument that underlies Theorem 1. The proof based on existence of NE is, however, arguably easier to grasp. □

Remark 13. The difficult direction of Theorem 5, which characterizes strict dominance in terms of best response, can be proved by constructing an artificial zero-sum game much as in Remark 10. The zero-sum proof, rather than the Separating Hyperplane Theorem proof, is the one that appears in Pearce (1984).

Remark 14. Theorem 9 and Equation 5 are closely related to the Duality Theorem from linear programming. See Myerson (1991). □

The next result, a corollary of Theorem 9, says that the set of NE of a zero-sum game have a rectangular structure.

Theorem 10 (Interchangeability)**.** *Consider a finite zero-sum game and let σ^* and $\hat{\sigma}$ be NE. Then $(\sigma_1^*, \hat{\sigma}_2)$ and $(\hat{\sigma}_1, \sigma_2^*)$ are also NE.*

Proof. Immediate from the second half of Theorem 9, which establishes the equivalence between NE and profiles of security strategies. ■

Remark 15. Say that a game is *strictly competitive* if it can be transformed into a zero-sum game by the operations discussed in Section 4.1; see Adler et al. (2009). Because these transformations leave the best response correspondence unchanged, most of the results for zero-sum games extend to, or have close analogs for, strictly-competitive games. A trivial example of a strictly competitive game is a *constant-sum* game: there is a $c \in \mathbb{R}$ such that $u_1 + u_2 = c$. □

7.3 The structure of NE.

As in Section 7.1, if players and strategy sets are fixed, then payoff functions can be represented as a vector in $\mathbb{R}^{N|S|}$, giving payoffs for each player and each strategy profile. Let

$$\mathcal{N} : \mathbb{R}^{N|S|} \to \mathbb{P}(\Sigma)$$

be the NE correspondence: for any specification of payoffs $u \in \mathbb{R}^{N|S|}$, $\mathcal{N}(u)$ is the set of NE for the game defined by u. By Theorem 1, $\mathcal{N}$ is non-empty-valued.

The following result states that the limit of a sequence of NE is a NE (in particular, the set of NE for a fixed u is closed) but that for some games there are NE for which there are no nearby NE in some nearby games.

Theorem 11. *$\mathcal{N}$ is upper hemicontinuous but (for $|S| \geq 2$) may not be not lower hemicontinuous.*

Proof.

1. *Upper Hemicontinuity.* Since Σ is compact, it suffices to prove that $\mathcal{N}$ has a closed graph, for which it suffices to prove that every point (u, σ) in the complement of graph($\mathcal{N}$) is interior. Take any (u, σ) in the complement of graph($\mathcal{N}$). Then there is an i and a pure strategy s_i for which

$$u_i(s_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) > 0.$$

By continuity of expected utility in both u_i and σ, this inequality holds for all points within a sufficiently small open ball around (u, σ), which completes the argument.

2. *Lower Hemicontinuity.* Consider the game in figure 19. For any $\varepsilon > 0$, the

	L	R
A	$1, \varepsilon$	$1, 0$

Figure 19: *The NE correspondence is not lower hemicontinuous.*

unique NE is (A, L). But if $\varepsilon = 0$ then any mixed strategy profile is a NE, and in particular (A, R) is a NE. Therefore, for any sequence of $\varepsilon > 0$ converging to zero, there is no sequence of NE converging to (A, R). This example implies a failure of lower hemicontinuity in any game with $|S| \geq 2$.

■

The set of NE need not be connected, let alone convex. But one can prove that for any finite game, the set of NE can be decomposed into a finite number of connected components; see Kohlberg and Mertens (1986). In the Battle of the Sexes (Example 10 in Section 4.3), there are three components, namely the three NE. In the entry deterrence game of Example 27, there are two components: the pure strategy (In, A) and the component of $(O, (q, 1-q))$ for $q \in [0, 1/16]$.

7.4 Correlated Equilibrium.

I discuss correlated equilibrium by means of examples. For the general definition see a text such as Fudenberg and Tirole (1991) or Myerson (1991). The original paper is Aumann (1974).

Consider first the game of Battle of the Sexes (Figure 3). As noted in Section 3.2, it is reasonable in this game to imagine that the players toss a coin prior to play so that they play the two pure strategy NE with equal probability. This induces a correlated distribution over strategy profiles that can be interpreted as a a new kind of equilibrium, a correlated equilibrium: given the recommendation of the coin, and the assumption that the other player will follow that recommendation, a player has no incentive to deviate from the recommendation.

In the Battle of the Sexes example, the correlated equilibrium distribution was a convex combination of NE distributions. But it also possible for a correlated equilibrium distribution to lie outside of the convex hull of the NE distributions. Consider the following game, called Chicken, in Figure 20.

	J	D
J	$6,6$	$2,7$
D	$7,2$	$0,0$

Figure 20: *The Chicken.*

The motivation is that two drivers are racing against other. The one who jumps last wins, but if neither jumps, they crash (into each other or into a wall/cliff) and die. In the formalism, D (for "drive") is intended to be interpreted as, "let the other driver jump first." The canonical reference is the 1955 James Dean movie, Rebel Without a Cause. (Strictly speaking, the linked clip shows an outcome not considered by the game box formalism.)

There is a correlated equilibrium distribution of the form given in Figure 21, which is not a convex combination of NE distributions. In this correlated equi-

	J	D
J	$1/3$	$1/3$
D	$1/3$	0

Figure 21: *An independent distribution over strategy profiles for Battle of the Sexes.*

librium, a correlation device recommends either J or D to each player with a 1/3 chance that the joint recommendation will be (J, D), (D, J), or (D, D). A player does not see the recommendation made to the other player.

In particular, if the device recommends J, then the player infers that there is a 50:50 chance that the other player has been told to play J. Therefore J has an expected payoff of 4 while D has an expected payoff of 3.5, so that J is indeed optimal. On the other hand, if the device recommends D, then the player infers that the other player has been told J, in which case D is indeed optimal.

In a finite game, the set of correlated equilibrium distributions is a convex polytope defined by a set of linear inequalities. It contains the convex hull of the set of NE distributions, sometimes (as just illustrated) strictly. This implies, in particular, that correlated equilibria exist. But one can establish existence of correlated equilibria directly by relatively elementary means; in particular, one can do so without use of a fixed point theorem.

8 Some Game Theory History.

The concept of Nash equilibrium seems quite natural, and indeed it had been anticipated by Cournot over a hundred years earlier. This raises the question of why Nash equilibrium is named in honor of Nash rather than, say, Cournot. As already noted, Nash himself just referred to it as an "equilibrium point." The following comments are based on Nachbar and Weinstein (2016), written for a special section of the Notices of the American Mathematical Society honoring Nash after his death.

In proposing his solution concept, Nash made two methodological contributions. First, there was widespread confusion in the economics literature prior to Nash about what was a model and what was a solution to that model. It was common, for example, to see references to *Cournot equilibrium* and *Bertrand equilibrium*, as though Cournot and Bertrand were using different equilibrium concepts. There is even a vexed literature called *conjectural variations* that tries to formalize this. Nash's methodological position was that Cournot equilibrium and Betrand equilibrium refer to the same solution concept, what we now call Nash equilibrium, applied to different models, rather than to different solution concepts applied to the same model. See also Mayberry et al. (1953).

Second, prior to Nash, the state of the art in dealing with games was the von Neumman-Morgenstern (VN-M) solution, developed in von Neumann and Morgenstern (1947) as an outgrown of von Neumman's earlier work on zero-sum games (see Section 7.2). From a modern (post-Nash) perspective, the VN-M solution is a mash-up of two different approaches to game theory, one for games in strategic form and another for games in coalition form. The coalition form abstracts away from the details about what individual players do and focuses instead on what payoff allocations are physically possible, both for all players taken together and for subsets of players. Because of its hybrid nature, the VN-M solution sometimes obscures strategic issues. For example, in the Prisoner's Dilemma, the payoff profiles predicted by the VN-M solution are all Pareto efficient; in particular, the VN-M solution rules out the payoff profile corresponding to the Nash equilibrium (in strictly dominant strategies, no less), namely (D, D). In effect, the VN-M solution assumes away the tension between efficiency and individual opportunism that the Prisoner's Dilemma is designed to illustrate.

More broadly, it is nearly impossible to use the VN-M solution to study incentives, one of the central topics of modern economics research. Nash argued that game theory should maintain a division between the study of strategic form and coalition form games, and Nash offered what we now call Nash equilibrium as a solution concept for games in strategic form. The *Nash program*, see Serrano (2008), bridges the two approaches by investigating whether, or under what circumstances, the analysis of strategic form games gives the same qualitative answers as the analysis of coalition form games.

Remark 16. Nash called the analysis of games in strategic form "non-cooperative game theory" and the analysis of games in coalition form "cooperative game theory." This terminology has stuck, which is unfortunate, I think, because it sets up a false expectation that the strategic form, being "non-cooperative," rules out cooperation. This is not correct, as the Nash program, alluded to above, underscores. □

References

Adler, I., C. Daskalakis, and C. H. Papdimitriou (2009): "A Note on Strictly Competitive Games," in *WINE*.

Anscombe, F. and R. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34.

Aumann, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

Bernheim, B. D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028.

Brandenburger, A. and E. Dekel (1987): "Rationalizability and Correlated Equilibrium," *Econometrica*, 55, 1391–1402.

Clancy, T. (1984): *The Hunt for Red October*, U.S. Naval Institute.

Cournot, A. (1838): *Researches into the Mathematical Principles of the Theory of Wealth*, New York: Kelley, translation from the French by Nathaniel T. Bacon. Translation publication date: 1960.

Daskalakis, C., P. Goldberg, and C. Papdimitriou (2009): "The Complexity of Computing a Nash Equilibrium," *SIAM Journal on Computing*, 39, 195–259.

Durlauf, S. N. and L. E. Blume, eds. (2008): *The New Palgrave Dictionary of Economics*, Palgrave.

Fudenberg, D. and J. Tirole (1991): *Game Theory*, Cambridge, MA: MIT Press.

Gibbons, R. (1992): *Game Theory for Applied Economists*, Princeton, NJ: Princeton University Press.

Harsanyi, J. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Eequilibrium Points," *International Journal of Game Theory*, 2, 1–23.

Herstein, I. N. and J. Milnor (1953): "An Axiomatic Approach to Measurable Utility," *Econometrica*, 21, 291–297.

Hu, T.-W. (2014): "Unpredictability of Complex (Pure) Strategies," *Games and Economic Behavior*, 88, 1–15.

Jehle, G. A. and P. J. Reny (2000): *Advanced Microeconomic Theory*, Addison Wesley, second ed.

Kohlberg, E. and J.-F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003–1037.

Kreps, D. (1988): *Notes on the Theory of Choice*, Underground Classics, Westview Press.

——— (1990): *A Course in Microeconomic Theory*, Princeton, NJ: Princeton University Press.

——— (2012): *Microeconomic Foundations I: Choice and Competitive Markets*, Princeton, NJ: Princeton University Press.

——— (2023): *Microeconomic Foundations II: Imperfect Competition, Information, and Strategic Interaction*, Princeton, NJ: Princeton University Press.

Kreps, D. and J. Scheinkman (1983): "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14, 326–337.

Luce, R. D. and H. Raiffa (1957): *Games and Decisions*, Wiley.

Magnan de Bornier, J. (1992): "The "Cournot-Bertrand Debate": A Historical Perspective," *History of Political Economy*, 24, 623–656.

Mailath, G. J. (2019): *Modeling Strategic Behavior*, World Scientific.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995): *Microeconomic Theory*, New York, NY: Oxford University Press.

Mayberry, J. P., J. F. Nash, and M. Shubik (1953): "A Comparison of Treatments of a Duopoly Situation," *Econometrica*, 21, 141–154.

McLennan, A. (2005): "The Expected Number of Nash Equilibria of a Normal Form Game," *Econometrica*, 75, 141–174.

Myerson, R. (1991): *Game Theory*, Cambridge, MA: Harvard University Press.

Nachbar, J. and J. Weinstein (2016): "Nash Equilibrium," *Notices of the American Mathematical Society*, 65, 494–496.

Nash, J. (1950a): "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences*, 36, 48–49.

——— (1950b): "Non-Cooperative Games," Ph.D. thesis, Princeton.

——— (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54, 286–295.

Nikaido, H. (1954): "On von Neumann's minimax principle," *Pacific Journal of Mathematics*, 4, 65–72.

Osborne, M. and A. Rubinstein (1994): *A Course in Game Theory*, Cambridge, MA: MIT Press.

Osborne, M. J. (2003): *An Introduction to Game Theory*, Oxford University Press.

——— (2008): "Strategic and Extensive Form Games," in *The New Palgrave Dictionary of Economics, 2nd Edition*, Palgrave MacMillan Ltd.

Pearce, D. G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1049.

Perea, A. (2020): "A Foundation for Expected Utility in Decision Problems in Games," EPICENTER Working Paper No. 22 (2020).

Reny, P. J. (2008): "Non-Cooperative Games (Equilibrium Existence)," in *The New Palgrave Dictionary of Economics*, ed. by L. E. B. Steven N. Durlauf, New York: Mcmillan, second ed.

Savage, L. (1954): *Foundations of Statistics*, New York: John Wiley.

Serrano, R. (2008): "Nash Program," in *The New Palgrave Dictionary of Economics*, ed. by S. Durlauf and L. Blume, Palgrave, 2nd ed.

Sion, M. (1958): "On general minimax theorems," *Pacific Journal of Mathematics*, 8, 171–176.

Sontag, S. and C. Drew (1998): *Blind Man's Bluff*, PublicAffairs.

Tadelis, S. (2013): *Game Theory: An Introduction*, Princeton University Press.

von Neumann, J. (1928): "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, 100, 295–320.

von Neumann, J. and O. Morgenstern (1947): *Theory of Games and Economic Behavior*, Princeton University Press, second ed.

Watson, J. (2013): *Strategy: An Introduction to Game Theory*, W. W. Norton, third ed.

Weinstein, J. (2016): "The Effect of Changes in Risk Attitude on Strategic Behavior," *Econometrica*, 84, 1881–1902.

Wilson, R. (1971): "Computing Equilibria of N-Person Games," *SIAM Journal of Applied Mathematics*, 21, 80–87.