Strategic Dominance

Introduction

We return to the terrain of strategic-form (*aka* normal-form) games and ask: what predictions can we make about a game's outcome if we only assume that the players are rational? (To be definite… at this point we are *not* making assumptions about what the players believe about the other players' rationality or believe about the strategic choices the other players will make.) We typically won't be able to make a unique prediction; instead we'll usually find a multiplicity of outcomes which are consistent with the rationality of all the players. In fact, with this relatively weak assumption we often will be unable to *refine* the set of possible outcomes at all.[1] (In a later handout we'll ask the question: what predictions can we make about a game's outcome if we also assume that this rationality is "common knowledge?"[2] This will often yield more precise predictions than merely assuming that all players are rational.)

To begin our analysis we discuss the concept of strategic *dominance*. Then we will turn to the more precisely relevant concept of "never a best response." These notions will be the foundation for our study of *nonequilibrium* solution concepts. These concepts are nonequilibrium in the sense that they typically admit outcomes which are not Nash equilibria.[3] (Each player *will* be playing a best response to her

1 To refine a set of outcomes is to remove some outcomes from further consideration through the application of some acceptability criterion.

2 A loose definition of common knowledge will suffice for now: A fact is *common knowledge* if each player knows it, knows that all other players know it, knows that all other players know that all other players know it, *ad infinitum*.

3 Of course, we haven't encountered Nash equilibrium in this course yet. Many of you are nevertheless familiar with that concept and it is for those folk I include this comment.

beliefs about what strategies the other players choose. However, her beliefs may be false.) These nonequilibrium techniques are of interest for the following reason: If we can make a useful prediction using only nonequilibrium analysis, our conclusion can be much more compelling than if we had achieved the same result using the much stronger (and frequently more dubious) assumptions required by equilibrium analysis. Furthermore, by applying nonequilibrium techniques in our initial analysis of a game we will frequently greatly simplify our subsequent equilibrium analysis.

A strategy for a player is *dominated* if there exists another strategy for her which is better for her no matter what choices the opponents make. A rational player would never play a dominated strategy. We define what it means for a pure strategy to be dominated by another pure strategy, for a pure strategy to be dominated by a mixed strategy, and then for a mixed strategy to be dominated by a pure or mixed strategy.

Because a rational player would never play a dominated strategy, we can sometimes use a dominance analysis to rule out some outcomes as possibilities when the game is played by rational players. In some games a dominance analysis leads to a unique prediction of the outcome when players are rational; we say that these games are dominance solvable. In other games a dominance analysis results in no refinement of the set of possible outcomes. Other games lie between these two extremes: dominance analysis rejects some outcomes as impossible when the game is played by rational players but still leaves a multiplicity of outcomes.

Although we have established that a rational player would never play a dominated strategy, we have *not* established that any undominated strategy could be plausibly chosen by a rational player. I.e. we have not shown that a dominance analysis fully exhausts the implications of all players being rational.

Closely related to the concept of a strategy being dominated for a player is the idea that this strategy is “never a best response” for that player: No matter what beliefs she has about the actions of her opponents, she could not rationally choose to play that strategy. If a strategy is dominated, it can never be a best response. However, it is not obvious that the implication holds in the reverse direction. I.e. it’s not obvious that a strategy which is never a best response is also a dominated strategy. Therefore the set of strategies which are never best responses is weakly larger than the set of dominated strategies. An analysis based on whether strategies are possibly best responses *does* exhaust the implications of all players being rational: A strategy cannot be plausibly chosen by a rational player if and only if it is never a best response.

We will see—perhaps surprisingly—that in two-player games a strategy is never a best response if and only if it is dominated. For two-player games, then, a dominance analysis fully exploits the assumption that all players are rational. However, for games with three or more players, it is possible that an undominated strategy will yet never be a best response. Therefore we can sometimes rule out as a plausible choice a strategy even when it is undominated. For more-than-two-player games, then, a dominance argument need not fully exploit the assumption that all players are rational.

Recapitulation

Let's review the standard paradigm and notation. We have a finite set I of n players, $I=\{1,\ldots,n\}$. The i-th player has a nonempty finite set of pure strategies—her pure-strategy space S_i—available to her, from which she can choose one strategy $s_i\in S_i$. In general her choice of pure strategy will be probabilistically determined by a mixed strategy $\sigma_i\in\Delta(S_i)\equiv\Sigma_i$, which is a probability distribution over her pure-strategy space S_i; the probability that she chooses a particular pure strategy $s_i\in S_i$ is specified by $\sigma_i(s_i)$.

The players simultaneously choose their strategies. If they choose pure strategies, the resulting pure-strategy profile is some ordered n-tuple $\boldsymbol{s}=(s_1,\ldots,s_n)$, where $\boldsymbol{s}\in S=\mathsf{X}_{i\in I}S_i$. S is called the space of pure-strategy profiles. If they choose mixed strategies, the resulting mixed-strategy profile is some ordered n-tuple $\boldsymbol{\sigma}=(\sigma_1,\ldots,\sigma_n)$, where $\boldsymbol{\sigma}\in\Sigma\equiv\mathsf{X}_{i\in I}\Sigma_i$. Σ is called the space of mixed-strategy profiles.

When we want to focus on the strategy choices of all players except i, we write player i's deleted strategy profile $\boldsymbol{s}_{-i}=(s_1,\ldots,s_{i-1},s_{i+1},\ldots,s_n)$ for the pure-strategy case or $\boldsymbol{\sigma}_{-i}=(\sigma_1,\ldots,\sigma_{i-1},\sigma_{i+1},\ldots,\sigma_n)$ for the mixed-strategy case. We denote the space of player i's deleted pure- and mixed-strategy profiles by S_{-i} and Σ_{-i}, respectively. We write $\langle a_i,b_{-i}\rangle_i$ to indicate the strategy profile resulting when player i chooses strategy $a_i\in(S_i\cup\Sigma_i)$ and her opponents choose the deleted strategy profile $b_{-i}\in(S_{-i}\cup\Sigma_{-i})$.

After all the players simultaneously pick their individual strategies, a pure-strategy profile $\boldsymbol{s}\in S$ is realized.[4] Each player i receives an *ex post* payoff which depends on this $\boldsymbol{s}$, viz. $u_i(\boldsymbol{s})$.

The expected (*ex ante*) payoff to player i when the players all participate in the mixed-strategy profile $\boldsymbol{\sigma}\in\Sigma$ is

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

As a special case of (1) you can show that player i's expected payoff when she plays the pure strategy $s_i\in S_i$ against the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i}\in\Sigma_{-i}$ is

$$u_i(\langle s_i,\boldsymbol{\sigma}_{-i}\rangle_i)=\sum_{\boldsymbol{s}_{-i}\in S_{-i}}\left(\prod_{j\in I\setminus\{i\}}\sigma_j(s_j)\right)u_i(\langle s_i,\boldsymbol{s}_{-i}\rangle_i). \tag{2}$$

When player i believes that her $n-1$ opponents are playing the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i}\in\Sigma_{-i}$, the pure strategies which are best responses for i are given by i's best-response correspondence $\mathsf{BR}_i\colon\Sigma_{-i}\twoheadrightarrow S_i$,[5]

4 If all the players choose pure strategies, the realized pure-strategy profile $\boldsymbol{s}$ is simply the strategy profile the players choose. More generally, the players choose mixed strategies. Each player's mixed strategy results in a pure-strategy realization for her. The collection of these pure-strategy realizations is a realization of a pure-strategy profile.

$$\mathsf{BR}_i(\boldsymbol{\sigma}_{-i}) = \underset{s_i \in S_i}{\operatorname{arg\,max}}\, u_i(\langle s_i, \boldsymbol{\sigma}_{-i} \rangle_i), \tag{3}$$

which maps the space of deleted mixed-strategy profiles Σ_{-i} into subsets of the space of i's pure strategies S_i; i.e. $\mathsf{BR}_i\colon \Sigma_{-i} \twoheadrightarrow S_i$.

Fix some deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ by i's opponents. When player i has a unique pure-strategy best response—i.e. when $\mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$ is a singleton (alternatively, when $\#\mathsf{BR}_i(\boldsymbol{\sigma}_{-i}) = 1$)—she has no choice but to play that pure strategy. However, if there are two or more pure-strategy best responses, she can randomize over those pure-strategy best responses in any fashion, and she cannot put positive weight on any pure strategy which is not a best response. In other words, a mixed strategy $\sigma_i \in \Sigma_i$ is a best response for i against the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ if and only if it puts positive weight only on pure-strategy best responses; i.e. if and only if its support is contained in the best-response set:[6]

$$\operatorname{supp} \sigma_i \subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i}). \tag{4}$$

Strong dominance

Often a player's optimal strategy depends crucially on what she believes her opponents are doing. However, this is not always the case. We say that a strategy is *strongly* (or *strictly*) *dominant* for a player if it is strictly better than all of her other choices *regardless of the actions of her opponents.* When a rational player—someone who chooses her action so as to maximize her expected utility given her beliefs—has a dominant strategy, we can be quite confident in predicting that she would choose that dominant strategy. If every player had a dominant strategy, predicting the outcome of the game would be a trivial exercise: The predicted outcome would be the strategy profile in which each player chose her strongly dominant strategy. Fortunately or unfortunately, it is rarely the case that a player has a dominant strategy; however, notions of strategic dominance often remain useful. (In what follows, if I leave out the qualifier "strongly" or "strictly," interpret the statement as if the qualifier were explicit.)

First we will focus on when a pure strategy might be dominated. We will discuss what it means for one pure strategy for player i to dominate another of her pure strategies. We will see that we can assess dominance even if we restrict attention only to the opponents' pure strategies. After that we will see that sometimes a mixed strategy can dominate a pure strategy even though no pure strategy did. Lastly, we will how a mixed strategy can be dominated by another pure or mixed strategy.[7]

---

5 Recall that a correspondence yields a *set* for each value of its argument. "arg max" refers to the set of values of the argument which maximizes the maximand.

6 Recall that the support of a mixed strategy is the set of pure strategies to which it allocates positive probability; i.e. $\operatorname{supp} \sigma_i = \{s_i \in S_i\colon \sigma_i(s_i) > 0\}$.

7 A good reference concerning strong domination is Fudenberg and Tirole [1991].

Pure-strategy strong dominance

Consider two pure strategies for player i: $s_i, s_i' \in S_i$. We say that s_i' *strictly dominates* s_i if s_i' gives player i a strictly higher expected utility than does s_i for every possible deleted pure-strategy profile $\boldsymbol{s}_{-i} \in S_{-i}$ which her opponents could play, i.e. if

$$\forall \boldsymbol{s}_{-i} \in S_{-i}, \quad u_i(\langle s_i', \boldsymbol{s}_{-i} \rangle_i) > u_i(\langle s_i, \boldsymbol{s}_{-i} \rangle_i). \tag{5}$$

If s_i' strictly dominates s_i, a rational player would never choose s_i (if she believes her opponents will all choose pure strategies) because, regardless of her beliefs concerning her opponents' strategies $\boldsymbol{s}_{-i}$, she could increase her expected utility by choosing s_i' instead.

If there exists a strategy $s_i' \in S_i$ which strictly dominates $s_i \in S_i$ for player i, then we say that s_i is *strictly dominated* for player i. If there does not exist a pure strategy which dominates s_i, then we say that s_i is *undominated* (by pure strategies) for player i. If there exists an $s_i' \in S_i$ which dominates every other strategy $s_i \in S_i \backslash \{s_i'\}$,[8] then we say that s_i' is a *strictly dominant* strategy for player i. (We will see that a strategy can be dominated even if there does not exist a dominant strategy.) If a rational player has a dominant strategy, she must choose that strategy (if she believes that her opponents are all choosing pure strategies). (We will soon remove this "opponents are all choosing pure strategies" qualification.)

Note from (5) that in assessing whether one of her strategies is dominated, dominant, or undominated a player only needs to know her own payoff function $u_i: S \to \mathbb{R}$—i.e. how her payoff depends on the choices of all the players; she does not need to know anything about her opponents' payoffs or even that her opponents are rational.

Example: Analyzing pure-strategy dominance.

Pure-strategy domination is easy to analyze just by inspection of the payoff matrix. Consider Row's payoffs in the game in Figure 1.[9] We observe that c dominates b for Row, because $3>1$, $4>3$, and $0>-1$. Similarly, a dominates b. So, for either reason, b is a dominated strategy. However, neither a nor c dominates each other. (Against d, c is better than a, but against e, a is better than c.) Therefore a and c are both undominated by pure strategies. Further, Row doesn't have a dominant strategy, because she doesn't have a single strategy which dominates both of the remaining strategies. (A necessary, but not sufficient, condition for Row to have a dominant strategy is that every Row payoff in exactly one row appear in boldface. A sufficient condition is that no Row payoffs in any other row appears in boldface.)

8 Let A and B be sets. The *difference* (or *relative complement*) of A and B, denoted $A \backslash B$, is the set of elements which are in A but not in B, i.e. $A \backslash B = \{x \in A : x \notin B\}$. The difference of A and B is also sometimes written as $A - B$.

9 Row payoffs which are maximal within their columns and Column payoffs which are maximal within their row appear in boldface.

	d	e	f
a	2,1	**5,2**	**0**,−1
b	1,0	3,**3**	−1,2
c	**3**,7	4,**8**	**0**,6

Figure 1: Column has a dominant strategy; Row does not.

Now consider Column's payoffs. Strategy e dominates d, because $2>1$, $3>0$, and $8>7$. Similarly e dominates f. As a consequence both d and f are dominated strategies, and e is a dominant strategy for Column because it dominates all of his other strategies. (Note that every Column payoff in column e appears in boldface and that no other Column payoff appears in boldface.)

Restricting attention to opponents′ pure strategies

Look again at the definition of (5). If $s_i' \in S_i$ strictly dominates $s_i \in S_i$ in this sense, we are sure that s_i' will outperform s_i for player i as long as her opponents are choosing pure strategies (because the universal quantifier says for all $\boldsymbol{s}_{-i} \in S_{-i}$). But what if the opponents are playing mixed strategies? Can she still be sure that s_i' is a better choice than s_i? In general we want to allow for players choosing mixed strategies. It seems we would actually want the definition of (5) to be that s_i' strictly dominates s_i if the inequality holds for all possible mixed strategies by her opponents, i.e. if

$$\forall \boldsymbol{\sigma}_{-i} \in \Sigma_{-i}, \square\ u_i(\langle s_i', \boldsymbol{\sigma}_{-i}\rangle_i) > u_i(\langle s_i, \boldsymbol{\sigma}_{-i}\rangle_i). \tag{6}$$

Prima facie, the definition in (6) looks more difficult to satisfy than (5) because the inequality must hold in a larger set of cases. However, the two definitions are actually equivalent and the version in (5) is easier to verify because we only need to check a finite number of deleted pure-strategy profiles rather than all possible deleted mixed-strategy profiles. The two definitions are equivalent in the sense that, if s_i' strictly dominates s_i according to (5), then s_i' also strictly dominates s_i according to (6), and vice versa. Let's see why this is so. Clearly, satisfaction of the inequality in (6) implies satisfaction in (5) because the set of deleted pure-strategy profiles S_{-i} is included in the set of deleted mixed-strategy profiles Σ_{-i}.[10]

The other direction—that satisfaction of the inequality against all deleted pure-strategy profiles by the opponents implies satisfaction against all possible deleted mixed-strategy profiles—requires a little work to see. The idea is that player i's payoff to any pure strategy against a particular combination of mixed strategies by her opponents is a convex combination of her payoffs against deleted pure-strategy profiles by her opponents, where the coefficients of that convex combination depend only on her opponents' strategies. By playing the better (against pure strategies) pure strategy she will receive a higher convex combination than she would have by playing the inferior (against pure strategies) pure strategy.

---

10 This wording is so abusive as to be untrue. However it is so convenient an abuse that I (and the rest of the game theory community) permit it to stand. It is not true that $S_{-i} \subset \Sigma_{-i}$. The following, however, *is* true: For every $\boldsymbol{s}_{-i} \in S_{-i}$, form the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} = (\delta_1(s_1), \ldots, \delta_{i-1}(s_{i-1}), \delta_{i+1}(s_{i+1}), \ldots, \delta_n(s_n)) \in \Sigma_{-i}$. Then $\boldsymbol{s}_{-i}$ and $\boldsymbol{\sigma}_{-i}$ are equivalent in the sense that, for all $\sigma_i \in \Sigma_i$, $u_i(\langle \sigma_i, \boldsymbol{s}_{-i}\rangle_i) = u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i}\rangle_i)$.

In order to make this argument clearer consider an arbitrary deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$. We can use (2) to express $u_i(\langle s_i', \boldsymbol{\sigma}_{-i}\rangle_i)$ as a convex combination of $u_i(\langle s_i', \boldsymbol{s}_{-i}\rangle_i)$ terms, one for each $\boldsymbol{s}_{-i} \in S_{-i}$. Now assume that s_i' strictly dominates s_i in the sense of (5). Then we replace each $u_i(\langle s_i', \boldsymbol{s}_{-i}\rangle_i)$ term by something we know from (5) to be smaller, viz. $u_i(\langle s_i, \boldsymbol{s}_{-i}\rangle_i)$. The result is equal to $u_i(\langle s_i, \boldsymbol{\sigma}_{-i}\rangle_i)$, and we thereby establish the inequality (6) for an arbitrary $\boldsymbol{\sigma}_{-i}$ and therefore for all $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$. In other words symbols,

$$u_i(\langle s_i', \boldsymbol{\sigma}_{-i}\rangle_i) = \sum_{\boldsymbol{s}_{-i} \in S_{-i}} \left(\prod_{j \in I \setminus \{i\}} \sigma_j(s_j) \right) u_i(\langle s_i', \boldsymbol{s}_{-i}\rangle_i) > \sum_{\boldsymbol{s}_{-i} \in S_{-i}} \left(\prod_{j \in I \setminus \{i\}} \sigma_j(s_j) \right) u_i(\langle s_i, \boldsymbol{s}_{-i}\rangle_i) = u_i(\langle s_i, \boldsymbol{\sigma}_{-i}\rangle_i). \quad (7)$$

Therefore if s_i' strictly dominates s_i in the sense of (5), then it does in the sense of (6) as well.

Earlier I said that, if a rational player believed that all of her opponents were choosing pure strategies, she would never play a dominated strategy, as defined by (5), and she would choose her dominant strategy if she had one. Now that we have seen that we can without loss of generality restrict our attention to opponents' pure strategies when assessing questions of dominance, we can dispense with this qualification that the player believes that all of her opponents are choosing pure strategies: We can categorically say that a rational player would never play a dominated strategy [as defined in (5)] and that she would choose her dominant strategy if she had one.

When all players have dominant strategies

Consider the case in which each player $i \in I$ 1 is rational, 2 knows her own payoffs u_i: $S \to \mathbb{R}$, and 3 has a dominant strategy s_i^*. Then each player must choose her dominant strategy s_i^*, and the result of the game must be the strategy profile $(s_1^*, \ldots, s_n^*)$. We say that such a game is *dominance solvable*.[11]

Example: The Prisoners' Dilemma is dominance solvable.

Consider the Prisoners' Dilemma game of Figure 2.[12] Each prisoner can either Fink or stay Mum. If they both stay Mum they are released for lack of evidence and get to keep their loot. Any prisoner who Finks is rewarded for squealing; he is also set free unless his partner also squeals. Any prisoner who is squealed upon is given a prison sentence. The reward to squealing does not compensate for the suffering of prison life. So if they both stay Mum, each is better off than if they both Finked. However, if one player stays Mum, the other gains by being the sole Fink.

	M	F
M	1, 1	−1, 2
F	2, −1	0, 0

Figure 2: A Prisoners' Dilemma

11 We will also apply this term later with regard to the iterated elimination of strictly dominated strategies.

12 The earliest written description of the Prisoners' Dilemma appears in Tucker [1950], which is reprinted in Straffin [1980].

Consider the Row suspect. Finking is better if Column stays Mum, because Row not only goes free (as she would if she stayed Mum too) but also pockets the reward for testifying. Finking is better if Column Finks, because jail is a foregone conclusion so Row might as well benefit from the fruits of testifying. So we see that, regardless of Column's action, Finking is a strictly better choice for Row. Therefore Finking is a dominant strategy for Row. Exactly the same analysis applies to Column's choice. Because each player has the dominant strategy Fink, the only outcome for the game with rational prisoners is (Fink, Fink).

Note that this analysis only requires that each player is rational and knows the deal *he* is offered by the authorities (i.e. his own payoffs). A prisoner does not need to know what deal the other prisoner is promised and does not even need to know that the other prisoner is rational.

However, game theory is not typically this easy. (Games where all players have dominant strategies are uninteresting game theoretically.) Consider the matching pennies game of Figure 3. Each player decides which side of a coin to show. Row prefers that the coins match; Column prefers that they be opposite. Here a player's best choice depends crucially on the choice of her opponent. (E.g. Column gains by playing T if Row plays H, but Column gains by playing H if Row plays T.) No strategy is dominant for either player.

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

Figure 3: Neither player has a dominant strategy in matching pennies.

Mixed-strategy dominance

We can use the above definition of pure-strategy dominance to rule out of consideration any pure strategy s_i which is dominated by some other pure strategy s_i'. Can we do more? In other words are there cases in which a pure strategy is dominated by some mixed strategy $\sigma_i' \in \Sigma_i$ of player i's but is not dominated by any pure strategy? The answer is yes.

So we define a more general notion of domination. Consider a pure strategy $s_i \in S_i$ and a mixed strategy $\sigma_i' \in \Sigma_i$ for player i. We say that σ_i' strictly dominates s_i if σ_i' gives player i a strictly higher expected utility than does s_i for every possible pure-strategy profile s_{-i} which her opponents could play, i.e. if

$$\forall \boldsymbol{s}_{-i} \in S_{-i}, \square\, u_i(\langle \sigma_i', \boldsymbol{s}_{-i} \rangle_i) > u_i(\langle s_i, \boldsymbol{s}_{-i} \rangle_i). \tag{8}$$

We use the obvious terminology, parallel to that we defined above in connection with pure-strategy dominance, concerning when a strategy is dominated, undominated, and dominant. (Exactly the same argument we used with respect to pure-strategy domination above works here to justify our restriction of attention to deleted pure-strategy profiles by the opponents.)

It's probably not obvious to you that allowing mixed strategies could increase the set of pure strategies which are dominated and therefore irrational. I.e. how do we know that it's not the case that any pure strategy s_i which is dominated by a mixed strategy σ_i' wouldn't be dominated more simply by some pure strategy s_i'? Let's look at a simple game to provide an example.

Example: A mixed strategy can dominate where no pure strategy can.

Consider the game in Figure 4. It's easy to verify that none of Row's pure strategies dominate each other. In particular neither Up nor Middle dominates Down. Even though Up is better than Down when Column plays left, Down is better than Up when Column plays right. Similar remarks can be made about the relationship between Middle and Down.

	l: $[q]$	r: $[1-q]$
U: $[p]$	6, 0	0, 6
M: $[1-p]$	0, 6	6, 0
D	2, 0	2, 0

Figure 4: A mixed strategy can dominate where no pure strategy can.

However, consider the mixed strategy for Row in which she plays Up one-half the time and Middle the remainder, i.e. $\sigma_R' = \frac{1}{2} \circ U \oplus \frac{1}{2} \circ M$. We will see that, even though neither Up nor Middle dominates Down, this mixture of Up and Middle *does* dominate Down. Let's compute Row's expected utility from this mixed strategy against both of Column's pure-strategy choices, left and right:

$$u_R(\sigma_R';\ l) = \tfrac{1}{2}(6) + \tfrac{1}{2}(0) = 3 > 2 = u_R(D;\ l), \tag{9}$$

$$u_R(\sigma_R';\ r) = \tfrac{1}{2}(0) + \tfrac{1}{2}(6) = 3 > 2 = u_R(D;\ r). \tag{10}$$

So we see that this mixture strictly outperforms Down regardless of which strategy Column chooses; therefore σ_i' dominates Down even though Down was undominated by pure strategies alone.

How did I know to choose the $(\frac{1}{2}, \frac{1}{2})$ mixing probabilities as I did? Let's write a mixture between Up and Middle more generally as

$$\sigma_R' = p \circ U \oplus (1-p) \circ M, \tag{11}$$

where $p \in [0, 1]$. We want to choose p such that Row's payoff to the mixed strategy exceeds her payoff to the pure strategy Down both when Column chooses left and when he chooses right. I.e. we want p to satisfy the two inequalities[13]

$$u_R(\sigma_R';\ l) = 6p + 0 \cdot (1-p) > u_R(D;\ l) = 2, \tag{12}$$

$$u_R(\sigma_R';\ r) = 0 \cdot p + 6(1-p) > u_R(D;\ r) = 2. \tag{13}$$

13 You can see how using the definition of dominance which only requires comparisons against a finite number of opponents' pure-strategy choices simplifies life tremendously!

Inequality (12) if satisfied whenever $p>\frac{1}{3}$; inequality (13) is satisfied whenever $p<\frac{2}{3}$. Therefore the mixture between Up and Middle σ_R' will dominate Down as long as we choose

$$p\in(\tfrac{1}{3},\tfrac{2}{3}). \tag{14}$$

My choice of $p=\frac{1}{2}$ was just one of many I could have made.

The intuition for the successful domination of Down by a mixture of Up and Middle can be more clearly explained when we consider Column's choice between left and right as a mixed strategy, $\sigma_C=q\circ l\oplus(1-q)\circ r$. In Figure 5 we plot Row's expected payoff to Up, Middle, Down, and to the $(\frac{1}{2},\frac{1}{2})$ mixture of Up and Middle, all as a function of Column's mixing parameter q. It is easy to calculate that

$$u_R(U;\,q)=6q+0\cdot(1-q)=6q, \tag{15}$$

$$u_R(M;\,q)=0\cdot q+6(1-q)=6-6q. \tag{16}$$

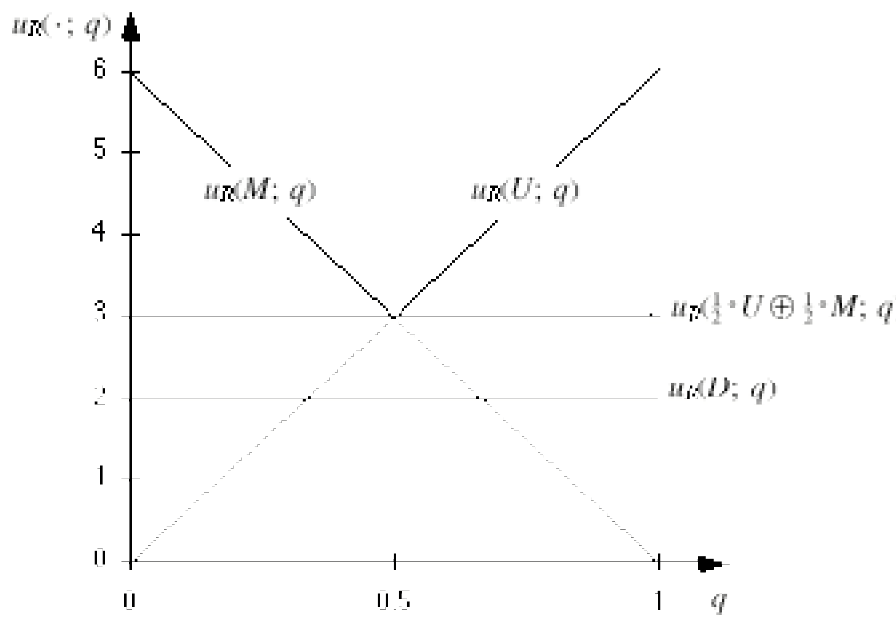

Figure 5: Row's pure- and mixed-strategy payoffs as a function of Column's mixed strategy.

Now we can ask the question: why was the mixture successful at dominating Down when neither Up nor Middle was successful on its own? We can see that Up is very good for Row against high q, i.e. when Column puts a lot of weight on left; but Up is very bad against low q. Conversely, Middle is bad against high q but good against low q. So we see that Up and Middle are complementary: one's regime of strength is the other's regime of weakness. By mixing between these two highly variable performers, we obtain a mixed strategy which is uniformly mediocre. In this example uniformity was a virtue when finding dominating strategies—it doesn't matter how good a strategy is at its best if it is too bad at its worst. Mediocrity is tolerable as long as it is good enough. Even if Row's payoffs to Down were somewhat higher, we would still have been able to find a dominating mixture. You should be able to show that this would be possible as long as Row's common payoff to Down were less than 3.

Dominated mixed strategies

Let's summarize our discussion of strict dominance so far. First we defined what it meant for one pure strategy to dominate another pure strategy. We justified restricting attention to pure strategies of the opponents when assessing the performance of a conjectured dominating strategy. We then defined strict dominance of a pure strategy by a mixed strategy. We saw that a pure strategy could be dominated by a mixed strategy even if it is not dominated by any pure strategy.

Note that thus far we have only discussed what it meant for a pure strategy to be dominated. But that raises a further question: If we know what pure strategies are dominated and therefore not rationally chosen, what can we say about the mixed strategies which are dominated and therefore not rationally chosen by a rational player?

Part of the answer is simple: Any mixed strategy which puts positive probability on a dominated strategy is itself dominated. I leave it as an exercise to generalize (6) in the natural way in order to define what if would mean for a mixed strategy to be strictly dominated (by another mixed strategy). Then show that, if some mixed strategy σ_i has a dominated pure strategy in its support, you could construct another mixed strategy σ_i' which strictly dominates σ_i.

However, this does not mean that any mixed strategy which puts positive probability only upon undominated pure strategies is necessarily undominated itself. We saw above that a pure strategy could be dominated by a mixed strategy even though it was undominated by any pure strategy. Very similarly, a nondegenerate mixed strategy σ_i can be dominated by another mixed strategy σ_i' (even by a pure strategy) even though σ_i puts no weight on dominated pure strategies.

Example: A mixed strategy over undominated pure strategies can be dominated.

To see that a mixed strategy can be dominated even when it puts no weight on dominated pure strategies we can modify the game in Figure 4 slightly, by changing Row's payoff to Down from $2 \to 4$. See Figure 6.

	l: $[q]$	r: $[1-q]$
U: $[p]$	6,0	0,6
M: $[1-p]$	0,6	6,0
D	4,0	4,0

Figure 6: Down is now better than it was.

Now we again consider the mixture $\sigma_R' = \frac{1}{2} \circ U \oplus \frac{1}{2} \circ M$, and plot Row's payoff to this mixed strategy and to her pure strategies U, M, and D as functions of Column's mixed strategy q. As we can see in Figure 7, D now dominates the mixed strategy even though neither pure strategy U or M was dominated.

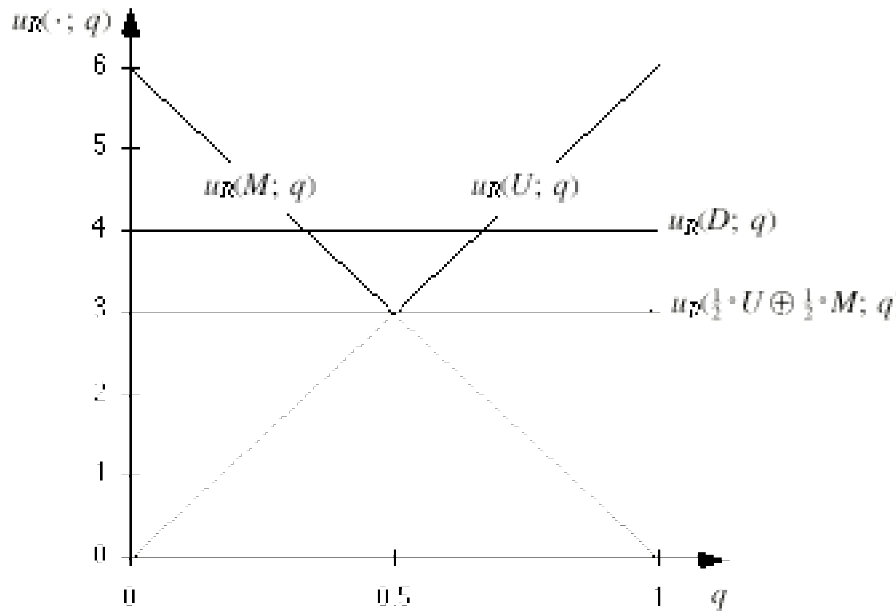

Figure 7: An equal mixture of U and M is dominated by D even though neither U nor M is dominated.

In fact we can show that *every* nondegenerate mixture of Up and Middle is dominated. It will be useful to plot the expected payoff to Row from a mixture of Up and Middle as a function of Column's mixed strategy q. Let's compute Row's expected payoff to the mixed strategy $p\circ U \oplus (1-p)\circ M$:

$$u_R(p\circ U \oplus (1-p)\circ M;\,q) = 6[pq + (1-p)(1-q)] = 3 + 6(2p-1)(q-\tfrac{1}{2}). \tag{17}$$

I wrote the payoff in this form to show that the graph of Row's expected payoff, as a function of q, of every such mixed strategy is simply a rotation of her pure-strategy payoff to U (or M) about the intersection ($q=\frac{1}{2}, u=3$) of the two pure-strategy payoff graphs.[14] To see this note, that when $q=\frac{1}{2}$, $u_R(p\circ U \oplus (1-p)\circ M) = 3$, for all p. Note also that the slope of graph of $u_R(p\circ U \oplus (1-p)\circ M)$ varies smoothly with p, coinciding with the slopes of the graphs of $u_R(M;\,q)$ and $u_R(U;\,q)$ for $p=0$ and $p=1$, respectively. See Figure 8.

It is obvious from Figure 8 that Row's pure strategy Down dominates every mixture of Up and Middle such that $\frac{1}{3} < p < \frac{2}{3}$. However, for $p \leq \frac{1}{3}$, the mixture is at least as good as Down for small q; for $p \geq \frac{2}{3}$, the mixture is least as good as Down for q near one. So we still need to show that these small and large p mixtures of Up and Middle which are undominated by a pure strategy *are* dominated by mixed strategies.

Let's consider the case where $p = {}^5/_6$. (See Figure 9.) I will show that the payoff to this mixture of Up and Middle is dominated by a mixture of Up and Down. Exactly as we saw before concerning the graph

14 The payoff to the mixed strategy is a convex combination of the pure-strategy payoffs; therefore the graph of the mixed-strategy payoff is a vertical convex combination of the pure-strategy graphs. Therefore the mixed-strategy graph must pass through the intersection of the pure-strategy graphs.

of the payoff to a mixture of Up and Down, you can show that the graph of the payoff to a mixture of Up and Down is a rotation of the payoff to Up about the point of intersection of the graphs of the pure-strategy payoffs to Up and Down. The mixture of Up and Down I propose as a dominating mixed strategy chooses Up with probability $\frac{2}{3}$. You can show that the graph of Row's expected payoff to this mixture has the same slope as her payoff to the mixture of Up and Middle, but is strictly higher, as shown in Figure 9. Therefore this mixture of Up and Down dominates the mixture of Up and Middle.

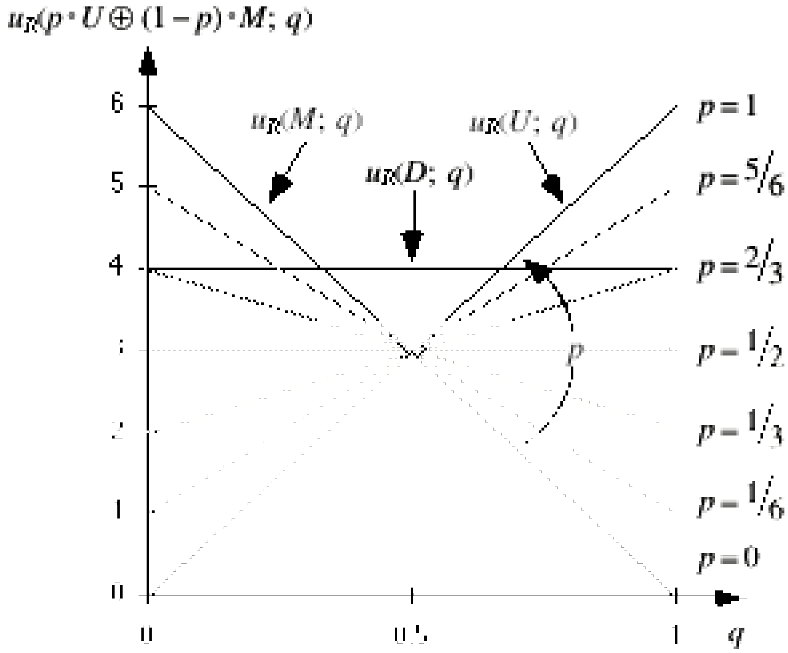

Figure 8: The payoff to a mixture of Up and Middle, as a function of q, is a rotation of a pure-strategy payoff function about the intersection of the two pure-strategy payoffs.

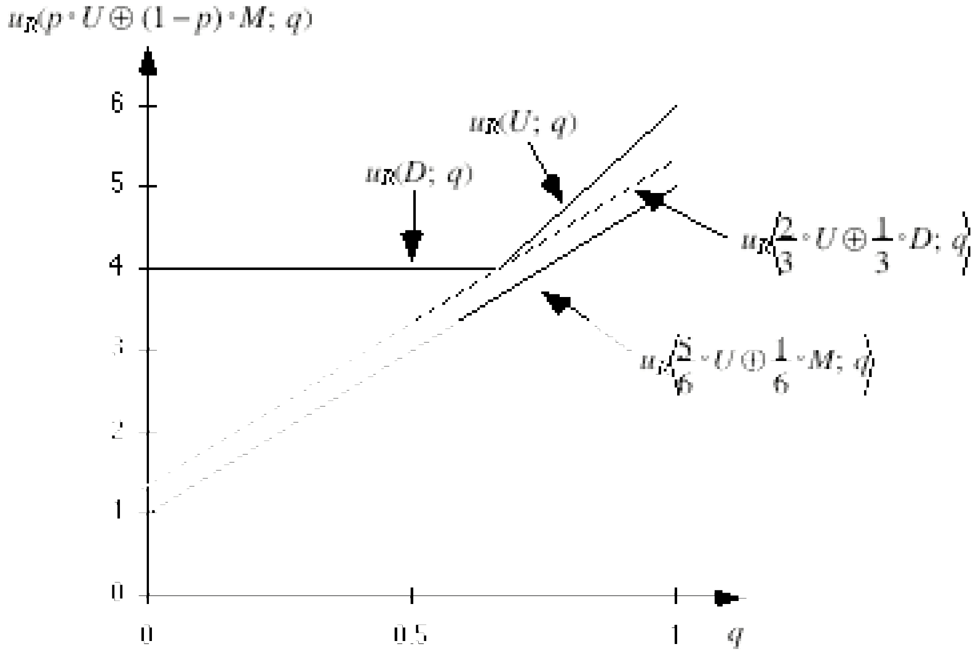

Figure 9: A mixed strategy, undominated by pure strategies, can be dominated by a mixed strategy.

Of course this example doesn't show that *every* nondegenerate mixture of Up and Middle is dominated as I claimed earlier. How do I know that I can *always* find a mixture which dominates any given mixture of Up and Middle? Now I'll give a more general argument.

In Figure 10 I've indicated the upper envelope of Row's pure-strategy payoff functions.[15] Consider any line tangent to this upper envelope. Such a tangent might coincide with one of the three pure-strategy segments; if not, it must be tangent at the intersection between either a the Middle and Down pure-strategy payoff functions or b the Down and Up payoff functions. Therefore any line tangent to the upper envelope corresponds to either 1 a pure strategy, 2 a mixture between Middle and Down, or 3 a mixture of Down and Up.

---

15 Let $f_1, \ldots, f_n$ be functions from some common domain X into the reals, i.e. $f_i\colon X \to \mathbb{R}$. Then the *upper envelope* of these functions is itself a function $\bar{f}\colon X \to \mathbb{R}$ defined by: $\bar{f}(x) \equiv \max\{f_1(x), \ldots, f_n(x)\}$.

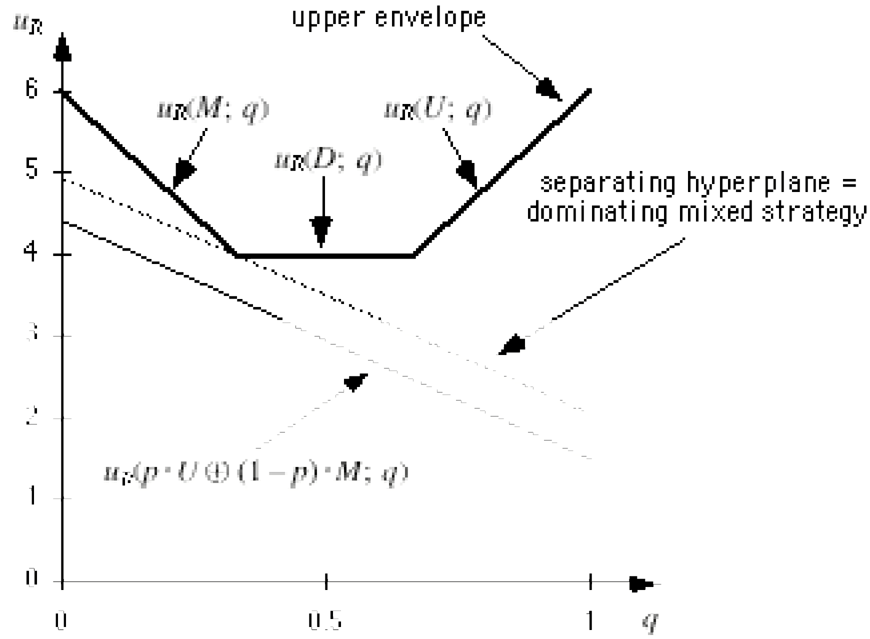

Figure 10: The separating hyperplane theorem guarantees the existence of a mixed strategy which dominates any nondegenerate mixture of Up and Middle.

Note from Figure 8 that any nondegenerate mixture of Up and Middle lies strictly below the upper envelope displayed in Figure 10. What I need to show then is that, for any such nondegenerate mixture, there exists a line tangent to the upper envelope which lies strictly above the mixture of Up and Middle. Then, because every such line corresponds to some mixed strategy, I will have shown that there exists a mixed strategy which dominates the mixture of Up and Middle.

To show that there indeed exists a line tangent to the upper envelope and above the mixture of Up and Middle, I pull out the separating hyperplane theorem.[16] First note that the graph above the upper envelope is a convex set.[17] The set of points below the mixture of Up and Middle is also a convex set.[18] These two convex sets are disjoint, therefore there exists a hyperplane which separates the two sets and is tangent to the first.[19]

---

16 There are several versions of the separation theorem; see Debreu [1959] or Debreu [1959]. The theorem I would want would say: "Let A and B be disjoint closed convex sets. Then there exists a hyperplane which supports A (i.e. includes a point on the boundary of A) and strictly separates B (i.e. so that B is contained within the interior of one of the closed half planes determined by the hyperplane)." I haven't found a reference for this exact theorem; typically you find a supporting hyperplane theorem which addresses the issue of a point on the boundary of A and you can find separating theorems which don't guarantee that the hyperplane intersects the boundary of one of the sets. Nevertheless I assert the theorem in quotes with confidence.

17 The set of points above the graph of a convex function is convex. Each of the pure-strategy bounding lines is a convex function. The points above the upper envelope is the intersection of the points above each bounding line; therefore the points above the upper envelope are the intersection of convex sets and, hence, are a convex set.

18 The set of points below the graph of a concave function is convex. A hyperplane is both convex and concave.

19 We will see later that this type of construction relies crucially on there only being two players.

Domination and never-a-best-response

Consider a strategy $\sigma_i \in \Sigma_i$ for player $i \in I$ and beliefs $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ which player i holds about the actions of the other players. If σ_i is not a best response for player i to these beliefs $\boldsymbol{\sigma}_{-i}$, then she must have available a better strategy $\sigma_i' \in \Sigma_i$; i.e. $\exists \sigma_i' \in \Sigma_i$ such that $u_i(\langle \sigma_i', \boldsymbol{\sigma}_{-i} \rangle_i) > u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i)$. It's conceivable that there are no beliefs to which σ_i is a best response; we say that σ_i is *never a best response* for i if

$$\forall \boldsymbol{\sigma}_{-i} \in \Sigma_{-i},\ \exists \sigma_i' \in \Sigma_i,\ u_i(\langle \sigma_i', \boldsymbol{\sigma}_{-i} \rangle_i) > u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i). \tag{18}$$

If σ_i is a dominated strategy for player i, then there exists a strategy $\sigma_i' \in \Sigma_i$ which is better-for-i than σ_i regardless of the actions $\boldsymbol{\sigma}_{-i}$ of the other players; i.e.[20]

$$\exists \sigma_i' \in \Sigma_i,\ \forall \boldsymbol{\sigma}_{-i} \in \Sigma_{-i},\ u_i(\langle \sigma_i', \boldsymbol{\sigma}_{-i} \rangle_i) > u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i). \tag{19}$$

Note that these two conditions for being never a best response and being dominated, respectively, though similar, are not identical. The order of appearance of the universal and existential quantifiers is reversed. From (19) you can easily deduce (18); i.e. a dominated strategy is never a best response. However, (18) does not simply imply (19); we leave open the possibility that a strategy may never be a best response but is yet undominated by any other strategy.

Despite the nonexistence of a simple, logical, syntactical deduction of (19) from (18), we will see that in two-player games it is indeed the case that any strategy which is never a best response is also dominated; i.e. when $n = 2$, $(19) \Leftrightarrow (18)$.[21] Then we will see an example which shows that this relationship need not hold in general when there are more than two players.

In two-player games: never-a-best-response ⇔ dominated

Without loss of generality consider player 1. Here's a preview sketch of our demonstration that, for two-player games, a strategy is never a best response if and only if it is dominated: We have already observed that if a strategy is dominated it is never a best response. Now we show that if a strategy is never a best response then it is dominated. The mathematical space we will consider is player 1's expected utility plotted against player 2's mixed strategy. The milestones of the demonstration are: 1 The graph of player 1's payoff to any player-1 mixed strategy as a function of player 2's mixed strategy will be a hyperplane in this space.[22] 2 The hyperplane corresponding to a dominated strategy lies everywhere strictly below the hyperplane corresponding to a dominating strategy. 3 The upper envelope

---

20 Of course, we know that this set of inequalities could be equivalently written with the substitutions $\boldsymbol{s}_{-j} \to \boldsymbol{\sigma}_{-j}$ and $S_{-j} \to \Sigma_{-j}$. I write these this way to emphasize the relationship to the previous set of inequalities.

21 Pearce [1984: Appendix B, Lemma 3] provides an elegant and short proof of this result by constructing a new, zero-sum game from the original, not-necessarily–zero-sum game and exploiting the existence of a Nash equilibrium. This proof however does not deliver the same graphical intuition as the approach taken here. Myerson [1991] proves a theorem, which implies this result in two-player games, using a linear-programming formulation. Fudenberg and Tirole [1991] sketch the proof of a related theorem which has the more explicitly separating-hyperplane flavor of our approach.

22 Let $f\colon X \to Y$. The *graph* of f is $\{(x, f(x))\colon x \in X\} \subset X \times Y$. Therefore the graph of a function lies in the Cartesian product of its domain and target set.

of the hyperplanes corresponding to player 1's pure strategies describes player 1's expected utility when she chooses her strategy optimally given player 2's strategy. The hyperplane corresponding to any strategy which is never a best response must lie everywhere strictly below this upper envelope. 4 Any hyperplane which is tangent to this upper envelope corresponds to some mixed strategy for player 1. 5 If a strategy is never a best response, there exists a hyperplane which is tangent to the upper envelope and which lies everywhere strictly above the hyperplane for the never-a-best-response strategy. Therefore there exists a strategy which dominates the never-a-best-response strategy, and therefore the never-a-best-response strategy is dominated. (Figure 11 illustrates this theorem in a case where player 2 has only two pure strategies. Therefore his mixed-strategy space is the one-dimensional unit simplex, which we parameterize by $q\in[0,1]$.)

1□ The graph of player 1's payoff to any mixed strategy will be a hyperplane in this space.

We have previously seen that the payoff to a player from a mixed-strategy profile is a linear function of the mixing probabilities of any one player:[23] For any $i,k\in I$,

$$u_i(\boldsymbol{\sigma})=\sum_{\boldsymbol{s}_k\in S_k}\sigma_k(s_k)\left[\sum_{\boldsymbol{s}_{-k}\in S_{-k}}\left(\prod_{j\in I\setminus\{k\}}\sigma_j(s_j)\right)u_i(\boldsymbol{s})\right].\tag{20}$$

In particular, player 1's payoff to the mixed-strategy profile (σ_1,σ_2) is the linear (and therefore convex and concave) function of the $\{\sigma_2(s_2)\}_{s_2\in S_2}$

$$u_1((\sigma_1,\sigma_2))=\sum_{s_2\in S_2}c(s_2;\sigma_1)\sigma_2(s_2),\tag{21}$$

where the coefficients $c(s_2;\sigma_1)$ are given by c: $S_2\times\Sigma_1$ defined by

$$c(s_2;\sigma_1)=\sum_{s_1\in S_1}\sigma_1(s_1)u_1((s_1,s_2)).\tag{22}$$

Therefore the graph of $u_1(\sigma_1,\cdot)$: $\Sigma_2\to\mathbb{R}$ (i.e. for fixed $\sigma_1\in\Sigma_1$) is a hyperplane in the space $\Sigma_2\times\mathbb{R}$, which we denote by $\mathscr{H}(\sigma_1)=\{(\sigma_2,u((\sigma_1,\sigma_2))): \sigma_2\in\Sigma_2\}$.

[23] See the "Strategic-Form Games" handout.

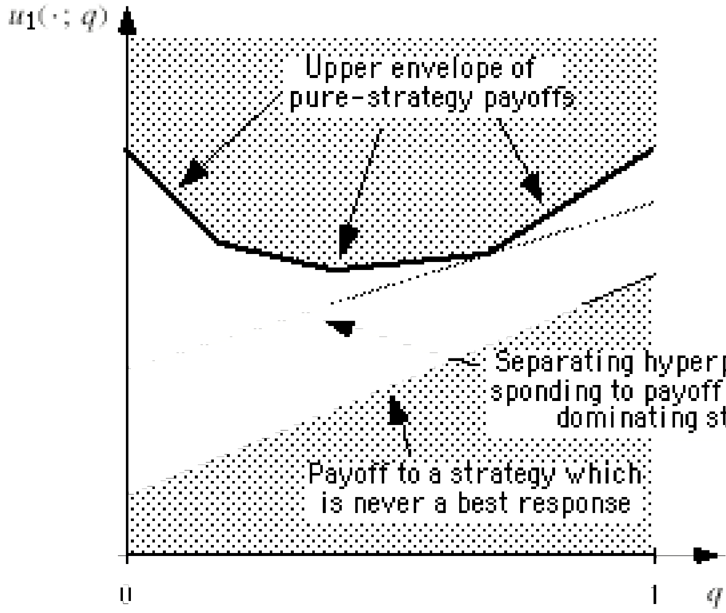

Figure 11: A strategy which is never a best response must be dominated in a two-player game.

2□ The hyperplane corresponding to a dominated strategy lies everywhere strictly below the hyperplane corresponding to a dominating strategy.

If $\sigma_1{}'$ dominates σ_1, then, for all $\sigma_2 \in \Sigma_2$, $u_1((\sigma_1{}', \sigma_2)) > u_2((\sigma_1, \sigma_2))$ and therefore the hyperplane corresponding to $\sigma_1{}'$ lies everywhere above σ_1; i.e. $\mathscr{H}(\sigma_1{}')$ lies everywhere above $\mathscr{H}(\sigma_1)$.

3□ The upper envelope of the hyperplanes corresponding to player 1's pure strategies describes player 1's expected utility when she chooses her strategy optimally given player 2's strategy. The hyperplane corresponding to any strategy which is never a best response must lie everywhere strictly below this upper envelope.

Let $f_1, \ldots, f_n$ be functions from some common domain X into the reals, i.e. f_i: $X \to \mathbb{R}$. Then the *upper envelope* of these functions is itself a function $\bar{f}$: $X \to \mathbb{R}$ defined by: $\bar{f}(x) \equiv \max\{f_1(x), \ldots, f_n(x)\}$. Consider the finite set of functions $\{u_1(s_1, \cdot)\}_{s_1 \in S_1}$, where each $u_1(s_1, \cdot)$: $\Sigma_2 \to \mathbb{R}$. (These are player 1's expected-payoff functions to her pure strategies.) Let $\bar{u}_1$ be the upper envelope of the $\{u_1(s_1, \cdot)\}_{s_1 \in S_1}$, i.e. $\bar{u}_1$: $\Sigma_2 \to \mathbb{R}$ and $\forall \sigma_2 \in \Sigma_2$, $\bar{u}_1(\sigma_2) = \max\,\{u_1(s_1, \sigma_2)$: $s_1 \in S_1\}$. For a given $\sigma_2 \in \Sigma_2$, if player 1 chooses her strategy optimally she receives the expected payoff $\max_{s_1 \in S_1} u_1(s_1, \sigma_2) = \bar{u}_1(\sigma_2)$. Therefore the upper envelope function $\bar{u}$ describes player 1's expected payoff as a function of player 2's strategy when player 1 chooses a best response to player 2's strategy.

Denote the graph of this upper envelope function by $\mathscr{U}$; i.e. $\mathscr{U} = \{(\sigma_2, \bar{u}_1(\sigma_2))$: $\sigma_2 \in \Sigma_2\}$. This upper-envelope surface is composed of a finite number of linear segments (i.e. pieces of hyperplanes).[24] Each

[24] Consider any pure strategy $s_1 \in S_1$ for player 1. Let $F(s_1)$ be the set of player-2 mixed strategies for which s_1 is a best response for player

of these faces corresponds to some player-1 pure strategy s_1 in the sense that for all σ_2 directly below this face (i.e. in the projection of this face upon Σ_2) s_1 is a best response by player 1 to player 2's choice of σ_2. The intersections of two or more faces are linear manifolds (i.e. intersections of hyperplanes). Such a linear manifold might be a single point, a line, or a set of higher dimension. Consider the projection of such a linear manifold onto Σ_2. For each σ_2 belonging to this projection, any of the pure strategies corresponding to the intersecting faces is a best response by player 1 to player 2's choice of σ_2.

Note that, for every player-1 strategy $\sigma_1 \in \Sigma_1$, the hyperplane corresponding to σ_1, viz. $\mathscr{H}(\sigma_1)$, must lie everywhere weakly below the upper envelope $\mathscr{U}$. I.e. $\forall \sigma_1 \in \Sigma_1$, $\forall \sigma_2 \in \Sigma_2$, $u_1((\sigma_1, \sigma_2)) \leq \bar{u}_1(\sigma_2)$. If, to the contrary, for some $\sigma_2 \in \Sigma_2$, it were true that $\mathscr{H}(\sigma_1)$ lay strictly above $\mathscr{U}$, then we would have the contradiction that $u_1((\sigma_1, \sigma_2)) > \bar{u}_1(\sigma_2) = \max\{u_1((s_1, \sigma_2)) \colon s_1 \in S_1\} = \max\{u_1((\sigma_1, \sigma_2)) \colon \sigma_1 \in \Sigma_1\}$.[25]

If $\sigma_1 \in \Sigma_1$ were a best response by player 1 to some beliefs $\sigma_2 \in \Sigma_2$, then $u_1((\sigma_1, \sigma_2)) = \bar{u}_1(\sigma_2)$.[26] Therefore, if σ_1 were never a best response, then $\forall \sigma_2 \in \Sigma_2$, $u_1((\sigma_1, \sigma_2)) < \bar{u}_1(\sigma_2)$, and therefore the hyperplane $\mathscr{H}(\sigma_1)$ corresponding to this never-a-best-response strategy must lie everywhere strictly below the upper envelope $\mathscr{U}$.

4☐ Any hyperplane which is tangent to this upper envelope corresponds to some mixed strategy for player 1.

Consider a hyperplane which is tangent to the upper envelope $\mathscr{U}$. If this hyperplane is coincident with the upper envelope along an entire face of the upper envelope then it corresponds to (i.e. is the graph of) the expected payoff to player 1, as a function of player 2's mixed strategy, from the pure strategy which corresponds to that face. If this hyperplane is tangent at a linear manifold which is the intersection of two faces, the tangent hyperplane corresponds to the expected payoff to player 1 when she plays some mixture over the two pure strategies corresponding to the intersecting faces. More generally, if the hyperplane is tangent to the upper envelope at some point $(\sigma_2, \bar{u}_1(\sigma_2))$, then this tangent hyperplane corresponds to the payoff to player 1 from some mixed strategy σ_1 which puts weight only on those pure strategies which correspond to the faces which intersect at the point of tangency. I.e. $\exists \sigma_1 \in \Sigma_1$ such that $\operatorname{supp} \sigma_1 \subset \mathrm{BR}_1(\sigma_2)$ and such that the tangent hyperplane is the graph of $u_1(\sigma_1, \cdot)$ over Σ_2.[27]

5☐ If a strategy is never a best response, there exists a hyperplane which is tangent to the upper envelope and which lies everywhere strictly above the hyperplane for the never-a-best-response strategy.

1; i.e. $F(s_1) = \{\sigma_2 \in \Sigma_2 \colon s_1 \in \mathrm{BR}_1(\sigma_2)\}$. You can show that for all $s_1 \in S_1$, $F(s_1)$ is a convex set. Therefore $\{F(s_1)\}_{s_1 \in S_1}$ is a finite partition of Σ_2, because the cells of the collection are pairwise disjoint and exhaust Σ_2. Therefore, for any $s_1 \in S_1$, the graph of the upper envelope $\mathscr{U}$ over $F(s_1)$ is a hyperplane.

25 Make sure you understand why $\max\{u_1(s_1, \sigma_2) \colon s_1 \in S_1\} = \max\{u_1(\sigma_1, \sigma_2) \colon \sigma_1 \in \Sigma_1\}$!

26 Recall that a mixed strategy $\sigma_1 \in \Sigma_1$ is a best response by player 1 to the mixed strategy $\sigma_2 \in \Sigma_2$ if and only if $\operatorname{supp} \sigma_1 \subset \mathrm{BR}_1(\sigma_2)$; i.e. it puts positive weight only upon pure-strategy best responses.

27 I'm not offering an accessible proof for this claim at the present time. See me this time next year to see whether I've worked out an elegant argument!

Therefore there exists a strategy which dominates the never-a-best-response strategy, and therefore the never-a-best-response strategy is dominated.

Now I pull out the separating hyperplane theorem.[28] First note that the graph above the upper envelope is a convex set.[29] Consider the graph of the payoff to player 1 from a strategy which is never a best response. It is a hyperplane and therefore concave. The set of points below this hyperplane is also a convex set.[30] These two convex sets are disjoint; therefore there exists a hyperplane which is tangent to the higher set and strictly separates the lower set. Because this separating hyperplane is tangent to the upper envelope $\mathcal{U}$, it corresponds to a mixed strategy for player 1. Because the separating hyperplane strictly separates the lower convex set, it lies everywhere strictly above the hyperplane corresponding to the never-a-best-response strategy. Therefore there exists a mixed strategy for player 1 which dominates the never-a-best-response strategy. ☺

Three or more players: never-a-best-response $\not\Rightarrow$ dominated

We have seen that in two-player games any strategy which is never a best response must also be dominated. Now we show that this equivalence between never being a best response and being dominated does not extend generally to games with more than two players. We show this by exhibiting a three-player game in which player 3 will have a strategy which is never a best response to any pair of mixed strategies by the two opponents yet this strategy will not be dominated by any other strategy of player 3's.

Consider the game in Figure 12.[31] There are three players; player 1 chooses a row (U or D); player 2 chooses a column (l or r); and player 3 chooses a matrix (A, B, C, or D). We will only be concerned with the payoffs to player 3, so we only include her payoffs in the matrices of Figure 12. Each of players 1 and 2 have two pure strategies, therefore a mixed strategy for each of them can be written in terms of a single parameter. We indicate the probability that player 1 chooses U by p and the probability that player 2 chooses l by q.

---

28 There are several versions of the separation theorem; see Takayama [1985: 39–45] or Debreu [1959]. The theorem I would want would say: "Let A be a closed convex set and B be a convex set, where A and B are disjoint. Then there exists a hyperplane which supports A (i.e. includes a point on the boundary of A) and strictly separates B (i.e. so that B is contained within the interior of one of the closed half planes determined by the hyperplane)." I haven't found a reference for this exact theorem; typically you find a supporting hyperplane theorem which addresses the issue of a point on the boundary of A and you can find separating theorems which don't guarantee that the hyperplane intersects the boundary of one of the sets. Nevertheless I assert the theorem in quotes with confidence.

29 The set of points above a convex function is convex. Each of the pure-strategy bounding lines is a convex function. The points above the upper envelope is the intersection of the points above each bounding line; therefore the points above the upper envelope are the intersection of convex sets and, hence, are a convex set.

30 The set of points below a concave function is convex. (A hyperplane is both convex and concave.)

31 This game is given in problem 2.7 of Fudenberg and Tirole [1991].

		l [q]	r [$1-q$]		l [q]	r [$1-q$]		l [q]	r [$1-q$]		l [q]	r [$1-q$]
[p]	U	9	0		0	9		0	0		6	0
[$1-p$]	D	0	0		9	0		0	9		0	6
		A			B			C			D	

Figure 12: Player 3′s payoffs in a three-player game as a function of Row′s and Column′s mixed strategies p and q, respectively. Player 3 chooses the matrix.

We first compute player 3's expected payoff to each of her four pure strategies as functions of the mixed strategies of her two opponents.

$$u_3(A;p,q)=9pq, \tag{23a}$$

$$u_3(B;p,q)=9[p(1-q)+(1-p)q]=9(p+q-2pq), \tag{23b}$$

$$u_3(C;p,q)=9(1-p)(1-q), \tag{23c}$$

$$u_3(D;p,q)=6[pq+(1-p)(1-q)]=6(1+2pq-p-q). \tag{23d}$$

Player 3's expected payoff to each of her four pure strategies is plotted as a function of p and q in Figure 13.

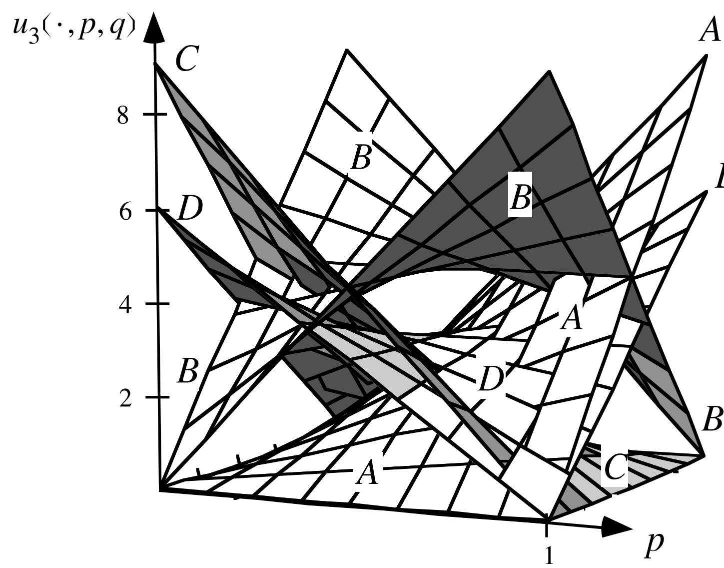

Figure 13: Player 3′s pure-strategy payoffs as functions of player 1 and 2′s deleted mixed-strategy profile (p,q) in the game from Figure 12.

We notice immediately that, unlike a two-player game, the graphs of the selected player's payoffs when plotted against the deleted strategy profile of her opponents are not hyperplanes.[32] Although we previously observed that a player's payoff is linear in a specified other player's mixing probabilities, her payoff is not linear in the mixing probabilities of more than one players.

I want to show that D is not dominated by any of the pure strategies A, B, or C or in fact by any mixture of A, B, and C. Yet we will see that D is not a best response to any deleted mixed-strategy profile $\sigma_{-3}=(p,q)$ of the opponents.

It is clear from Figure 13 or from equations (23) that D is not dominated by any of the other three pure strategies. When $p=q=1$, D gives a higher payoff than B or C. When $p=q=0$, D gives a higher payoff than A (or B, but we've already shown that B doesn't dominate D).

Now we want to show that there is no mixture of A, B, and C which dominates D.[33] Call our candidate mixed strategy $\sigma_3=r\circ A\oplus(1-r-s)\circ B\oplus s\circ C$, where $r,s\geq 0$ and $r+s\leq 1$. In order that σ_3 dominate D, it must yield a higher payoff against all deleted strategy profiles $(p,q)\in[0,1]^2$. In particular this must be the case when $p=q=1$. Against this deleted mixed-strategy profile player 3's payoff to the mixed strategy σ_3 is $9r$, because B and C yield a zero payoff there, and D pays off 6. In order that $9r>6$, we must choose $r>@/3$. Similarly, when $p=q=0$, C provides the only nonzero payoff to σ_3. Therefore, in order that the payoff from σ_3 exceed the payoff of 6 from D, the weight on C must be such that $s>@/3$. Therefore the two requirements, in order that σ_3 dominate D, viz. $r,s>@/3$ are incompatible with the requirement that $r+s\leq 1$. Therefore there does not exist any strategy which dominates D.

You can show algebraically that D is never a best response. However, for our purposes it will suffice to clearly show graphically that the payoff from D lies everywhere below the upper envelope of the pure-strategy payoffs. (See Figure 14.) This example shows that the equivalence between never being a best response and being dominated does not extend to more than two-player games.

Summary and preview of coming attractions

We began our preparation for the study of nonequilibrium solution concepts by introducing the notion of strategic dominance. We defined what it means for one pure strategy for a player to dominate another of her pure strategies. We showed that to evaluate questions of strategic dominance for a player we can with total generality assume that her opponents are playing pure strategies. We generalized the notion of strategic dominance so as to define what it means for a mixed strategy to dominate another mixed strategy. This generalization is useful because it is possible for a strategy to be dominated by a mixed strategy even when it is undominated by pure strategies.

A rational player would never play a dominated strategy. Therefore we can rule out as observed outcomes any which involves any player playing a dominated strategy. (When we later make a stronger

32 In a two-player game, a deleted strategy profile is simply a strategy, because there is only one opponent.

33 Thanks to Sam Dinkin for supplying this part of the demonstration!

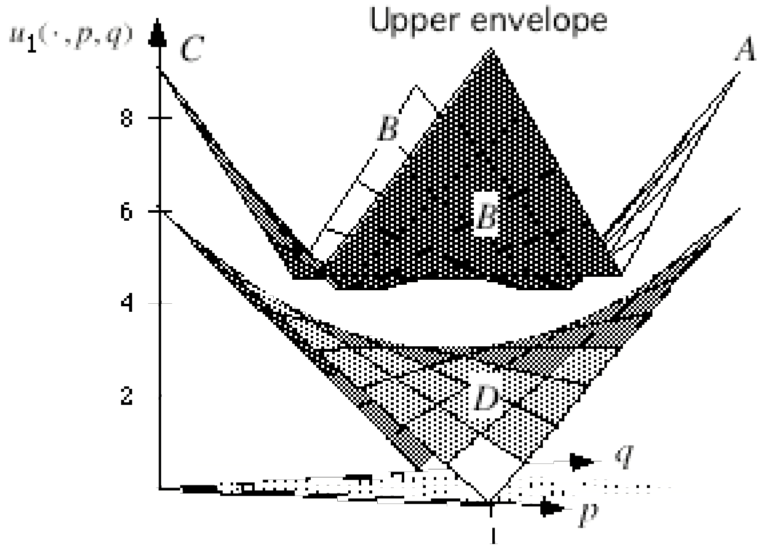

Figure 14: The payoff to D lies everywhere below the upper envelope of player 3′s pure-strategy payoffs.

assumption that the rationality of all players is *common knowledge*, we will apply a stronger technique called the *iterated elimination of dominated strategies* to further refine the set of possible observed outcomes.)

Although we established that a rational player would never play a dominated strategy, it need not be the case that any undominated strategy could be plausibly chosen by a rational player. I.e. a dominance analysis need not fully exhaust the implications of all players being rational. We then introduced what it means for a strategy to be *never a best response*. We argued that a strategy cannot be plausibly chosen by a rational player if and only if it is never a best response.

We saw that in two-player games a strategy is never a best response if and only if it is dominated. For two-player games, then, a dominance analysis does fully exploit the assumption that all players are rational. (And the iterated elimination of dominated strategies will fully exhaust the implications of common knowledge of rationality.)

However, for games with three or more players, it is possible that an undominated strategy will yet never be a best response. Therefore we can sometimes rule out as a plausible choice a strategy even when it is undominated. For more-than-two-player games, then, a dominance argument will not fully exploit the assumption that all players are rational. (And therefore the iterated elimination of dominated strategies will not fully exhaust the implications of common knowledge of rationality. We will instead iteratively eliminate strategies which are never best responses. This process will yield the nonequilibrium solution concept of *rationalizability*.)

References

Debreu, Gerard [1959] "Separation Theorems for Convex Sets," *SIAM Review* **1** 2 (July): 95–98.

Fudenberg, Drew and Jean Tirole [1991] *Game Theory*, MIT Press.

Myerson, Roger B. [1991] *Game Theory: Analysis of Conflict*, Harvard University Press.

Straffin, Philip [1980] "The Prisoner's Dilemma," *UMAP Journal* **1**: 101–103.

Takayama, Akira [1985] *Mathematical Economics*, second edition, Cambridge University Press.

Tucker, A.W. [1950] "A Two-Person Dilemma," mimeo.