Strategic-Form Games

Introduction

We'll look at noncooperative games which are played only once, which involve only a finite number of players, and which give each player only a finite number of actions to choose from. We'll consider what is called the *strategic* (or *normal*) form of a game. Although our formal treatment will be more general, our exemplary paradigm will be a two-person, simultaneous-move matrix game.

The strategic (or "normal") form of a game is a natural and adequate description of a simultaneous-move game. It is also a useful platform on which to perform at least some of our analysis of games which have a more complicated temporal and information structure than a simultaneous-move game has. (In order to perform the remaining analysis of these games, however, we'll later introduce and use the "extensive form.")

We will define a strategic-form game in terms of its constituent parts: players, actions, and preferences. We will introduce the notion of mixed strategies, which are randomizations over actions. Our first step in the analysis of these games will be to solve the Easy Part of Game Theory, viz. the problem of what choice a rational player would make given her beliefs about the choices of her opponents. Later we will turn to the Hard Part of Game Theory: what beliefs the players can rationally hold concerning the choices of their opponents.

Individual strategies

We have a nonempty, finite set I of $n \in \mathbb{N} \equiv \{1, 2, \ldots\}$ players

$$I = \{1, \ldots, n\}. \tag{1}$$

The i-th player, $i \in I$, has a nonempty set of *strategies*—her *strategy space* S_i—available to her, from which she can choose one strategy $s_i \in S_i$.[1] Note—as indicated by the "*i*" subscript—that each player has her own strategy space S_i. Therefore each player has access to her own possibly unique set of strategies. We will assume that each player's strategy space is finite.

When necessary we will refer to these as *pure strategies* in order to distinguish them from *mixed strategies*, which are randomizations over pure strategies.

Example:

Consider a two-player game between Robin and Cleever. Suppose Robin has two actions available to her: Up and Down. Then her strategy space S_{R} would be

$$S_{\mathrm{R}} = \{\text{Up}, \text{ Down}\}.$$

When she plays the game she can choose only one of these actions. So her strategy s_{R} would be either $s_{\mathrm{R}} = \text{Up}$ or $s_{\mathrm{R}} = \text{Down}$. Likewise, suppose that Cleever can move left, middle, or right. Then his strategy space is

$$S_{\mathrm{C}} = \{\text{left}, \text{ middle}, \text{ right}\}.$$

Strategy profiles

For the time being it will be useful to imagine that all players pick their strategies at the same time: player 1 picks some $s_1 \in S_1$, player 2 picks some $s_2 \in S_2$, etc. We can describe the set of strategies chosen by the n players as the ordered n-tuple:[2]

---

1 A strategy need not refer to a single, simple, elemental action; in a game with temporal structure a strategy can be a very complex sequence of actions which depend on the histories of simple actions taken by all other players. We will see this clearly when we learn to transform an extensive-form description of a game into its strategic form. The name "strategic form" derives precisely because the present formalism ignores all this potential complexity and considers the strategies as primitives of the theory (i.e. as units which cannot be decomposed into simpler constituents).

2 In this introduction I'm using boldface notation to represent multicomponent entities in hopes that this will help you keep straight the

$$\boldsymbol{s} = (s_1, \ldots, s_n). \tag{2}$$

This n-dimensional vector of individual strategies is called a *strategy profile* (or sometimes a *strategy combination*). For every different combination of individual choices of strategies we would get a different strategy profile $\boldsymbol{s}$. The set of all such strategy profiles is called the *space of strategy profiles* S. It is simply the Cartesian product of the strategy spaces S_i for each player.[3] We write it as[4]

$$S \equiv S_1 \times \cdots \times S_n = \mathop{\mathsf{X}}_{i=1}^{n} S_i = \mathop{\mathsf{X}}_{i \in I} S_i. \tag{3}$$

Example (continued):

Considering Robin as player 1, if she chose $s_{\mathrm{R}} = \text{Down}$ and Cleever chose $s_{\mathrm{C}} = \text{middle}$, then the resulting strategy profile would be:

$$\boldsymbol{s} = (\text{Down, middle}).$$

The space of all strategy profiles for this example is

$$S = S_{\mathrm{R}} \times S_{\mathrm{C}} = \{(\text{Up, left}), (\text{Up, middle}), (\text{Up, right}), (\text{Down, left}), (\text{Down, middle}), (\text{Down, right})\}.$$

Player i is often interested in what strategies the other $n-1$ players choose. We can represent such an $(n-1)$-tuple of strategies, known as a *deleted strategy profile*, by[5]

$$\boldsymbol{s}_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n). \tag{4}$$

To each player i there corresponds a *deleted strategy profile space* S_{-i}, which is the space of all possible strategy choices $\boldsymbol{s}_{-i}$ by her opponents of the form in (4), i.e.[6]

---

distinction between strategies, for example, and vectors of strategies. However, don't get spoiled: most papers and texts in game theory don't do this. And I'll stop doing it soon.

3 The Cartesian product (or *direct product*) of n sets is the collection of all ordered n-tuples such that the first elements of the n-tuples are chosen from the first set, the second elements from the second set, etc. E.g., the set of Cartesian coordinates $(x, y) \in \mathbb{R}^2$ of the plane is just the Cartesian product of the real numbers $\mathbb{R}$ with itself, i.e. $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$. For another example, let $A = \{1, 2\}$ and $B = \{\alpha, \beta, \delta\}$. Then $A \times B = \{(1, \alpha), (1, \beta), (1, \delta), (2, \alpha), (2, \beta), (2, \delta)\}$. More formally, $A_1 \times \cdots \times A_m = \{(a_1, \ldots, a_m){:}\ \forall i \in \{1, \ldots, m\},\ a_i \in A_i\}$. When we form the Cartesian product of m copies of the same set S, we simply write $S^m = S \times \cdots \times S$.

4 Note that the player set I is ordered. We avoid ambiguity concerning the order in which the Cartesian product is formed when the notation "$\mathsf{X}_{i \in I}$" is used by adopting the obvious convention, which is expressed in this case by $\mathsf{X}_{i=1}^{n}$.

5 In other words this is a strategy profile with one strategy (that of player i) deleted from it. The formal definition obviously does not work quite right if $i = 1$ or $i = n$, but the necessary modifications for these cases should be obvious.

6 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 simply $A - B$. The set $I \backslash \{i\} = \{1, \ldots, i-1, i+1, \ldots, n\}$, when $1 < i < n$. Note that "$I \backslash i$" and "$\{I\} \backslash \{i\}$" are incorrect attempts at expressing this set.

$$S_{-i} \equiv S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n = \mathop{\mathsf{X}}_{\substack{j=1 \\ j\neq i}}^{n} S_j = \mathop{\mathsf{X}}_{j\in I\setminus\{i\}} S_j. \tag{5}$$

If we want to single out the strategy-decision problem for the particular player i, it is useful to write a strategy profile $\boldsymbol{s}\in S$ as a combination of her strategy $s_i\in S_i$ and the deleted strategy profile $\boldsymbol{s}_{-i}\in S_{-i}$ of the strategies of her opponents. So you will typically see

$$\boldsymbol{s}\equiv(s_i,\boldsymbol{s}_{-i}). \tag{6}$$

As another example of this notation: $(\tilde{s}_i,\boldsymbol{s}_{-i})=(s_1,\ldots,s_{i-1},\tilde{s}_i,s_{i+1},\ldots,s_n)$. The usage in (6) is notationally abusive.[7] To avoid this we can define for each player $i\in I$ a function $\langle\cdot,\cdot\rangle_i$: $S_i\times S_{-i}\to S$ defined by $\langle a,(b_1,\ldots,b_{i-1},b_{i+1},\ldots,b_n)\rangle_i=(b_1,\ldots,b_{i-1},a,b_{i+1},\ldots,b_n)$.[8] Therefore $\boldsymbol{s}=\langle s_i,\boldsymbol{s}_{-i}\rangle_i$.

Example (continued):

Because we're using the letters R and C rather than the numbers 1 and 2 to identify the players in this example and because there are only two players, the correspondence I'm about to draw might seem artificial, but the idea should be clear. Robin's deleted strategy profile is the profile of the strategies chosen by Robin's opponents, viz. by Cleever. Therefore

$$S_{-\mathsf{R}}=S_{\mathsf{C}}=\{\text{left},\ \text{middle},\ \text{right}\},$$

$$S_{-\mathsf{C}}=S_{\mathsf{R}}=\{\text{Up},\ \text{Down}\}.$$

If $\boldsymbol{s}=(\text{Up},\text{middle})$, then $s_{-\mathsf{R}}=\text{middle}$ and $s_{-\mathsf{C}}=\text{Up}$.

Payoffs

The game is played by having all the players simultaneously pick their individual strategies. This set of choices results in some strategy profile $\boldsymbol{s}\in S$, which we call the *outcome* of the game. Each player has a set of preferences over these outcomes $\boldsymbol{s}\in S$. We assume that each player's preferences over lotteries over S can be represented by some von Neumann-Morgenstern utility function u_i: $S\to\mathbb{R}$.[9]

At the conclusion of the game, then, each player $i\in I$ receives a *payoff* $u_i(\boldsymbol{s})=u_i(\langle s_i,\boldsymbol{s}_{-i}\rangle_i)$.[10] Note that

---

7 The strategy-profile notation $(s_i,\boldsymbol{s}_{-i})$ is abusive because the position into which the first argument is meant to be inserted depends on the "i" subscripts. But we can't depend on those subscripts being visible. For example, in a three-player game where $S_1=S_2=S_3=\{U,M,D\}$, what strategy profile does $(U,(M,D))$ represent? It could mean any of (U,M,D), (M,U,D), or (M,D,U) for $i=1,2,3$, respectively.

8 However, you will not see this $\langle\cdot,\cdot\rangle_i$ notation anywhere else… yet!

9 Therefore we can extend the domain of each u_i to be the set of lotteries over outcomes in S. We perform the extension such that u_i has the expected-utility property: For any lottery over S of the form $p\circ\boldsymbol{s}\oplus(1-p)\circ\boldsymbol{s}'$, where $\boldsymbol{s},\boldsymbol{s}'\in S$ and $p\in[0,1]$, $u_i(p\circ\boldsymbol{s}\oplus(1-p)\circ\boldsymbol{s}')=pu_i(\boldsymbol{s})+(1-p)u_i(\boldsymbol{s}')$.

10 I use the soon-to-be-the-universal-standard notation "(…)" to enclose argument lists of functions, operators, etc., rather than the more

the payoff each player receives depends not only on the strategy she picked but also on the strategies which all the other players picked. In other words the payoff to any one player depends on the entire strategy profile played. The individual payoffs for all the n players for a particular strategy profile $\boldsymbol{s}$ define a payoff vector $\boldsymbol{u}(\boldsymbol{s})$ for that strategy profile $\boldsymbol{u}(\boldsymbol{s}) \equiv (u_1(\boldsymbol{s}), \ldots, u_n(\boldsymbol{s}))$; i.e. $\boldsymbol{u}\colon S \to \mathbb{R}^n$.

We can fully describe our game then by the triple $(I, S, \boldsymbol{u})$, i.e. by a player set I, a space of strategy profiles S, and a vector $\boldsymbol{u}$ of von Neumann-Morgenstern utility functions defined over S.

Denote by $\#S_i$ the *cardinality* of i's strategy space S_i (i.e. the number of strategies in S_i). The set of all payoffs to all players can be represented by a $\#S_1 \times \cdots \times \#S_n$ matrix.[11] Each cell in the payoff matrix corresponds to a particular strategy profile and contains the n-tuple payoff vector which specifies the payoff to each player when the strategy profile corresponding to that cell is played.[12] When there are only two players, this matrix is easily represented in two dimensions. The first player chooses a row and the second player chooses a column. (See the next example.) The payoffs for three-player games can be represented as $\#S_3$ matrices each of which is of dimension $\#S_1 \times \#S_2$. In this case players 1 and 2 choose a row and a column, respectively, and player 3 chooses the particular $\#S_1 \times \#S_2$ matrix.

Example (continued):

We can cook up payoffs to Robin and Cleever for each possible strategy profile (i.e., for each combination of individual strategy choices) to arrive at the following possible payoff matrix:

		C[leever/column]		
		l	*m*	*r*
R[obin/ow]	*U*	2,8	7,7	0,3
	D	1,5	7,4	9,6

Figure 1: A typical payoff matrix.

where the abbreviations for the strategies which label the rows and columns have the obvious meanings. The first payoff of each pair is by convention the one that Row (Robin) receives. The second is that received by Column (Cleever).

To make explicit the connection between the payoff matrix and the payoff function $u_i(\boldsymbol{s})$ formalism from above, we note two examples:

$$u_{\mathsf{R}}((U, l)) = 2,$$

common standard parentheses "(…)". I believe my alternative convention helps to quickly distinguish between instances of multiplication and instances of *operation upon*. For example, does "$P(x+a)$" mean a function P evaluated at $(x+a)$? Or does it mean a number P multiplied by $(x+a)$? This distinction *should* be clear from context alone. However, economists are notably prone to the notational abuses of 1 suppressing arguments and 2 identifying a *function* with its *values*.

11 The multidimensional matrix will have a side for each player, and the number of boxes along each side will be equal to the number of strategies the corresponding player has available to her.

12 In the two-player case the payoff matrix is often called a *bimatrix*, because each element is an ordered pair.

$u_C((D,r))=6.$

Best responses to pure strategies

We typically assume in game theory that all players are *rational*. This means that each player will choose an action which *maximizes her expected utility given her beliefs* about what actions the other players will choose. The Easy Part of Game Theory is figuring out what a player will do given her beliefs. This is the problem we focus on now. We ask the question: if player i knows (read "believes with certainty") which strategy each of her opponents will pick, what strategy should she pick? Obviously, she should pick a *best response* to the plays of her opponents.

Definition

We say that a strategy $s_i^* \in S_i$ for player i is a *best response by i to the deleted strategy profile* $\mathbf{s}_{-i} \in S_{-i}$ iff[13]

$$(\forall s_i \in S_i)\ \ u_i(\langle s_i^*, \mathbf{s}_{-i}\rangle_i) \geq u_i(\langle s_i, \mathbf{s}_{-i}\rangle_i), \tag{7}$$

or, equivalently,[14]

$$s_i^* \in \operatorname*{arg\,max}_{s_i \in S_i} u_i(\langle s_i, \mathbf{s}_{-i}\rangle_i). \tag{8}$$

What is this definition saying? With a few semesters of Varian behind you, this notation should not be too obfuscating, but it's worthwhile to go through it carefully. We're fixing our attention on one particular player i and we assume we know the strategy choice $s_j \in S_j$ made by each player $j, j \in I \backslash \{i\}$. If $s_i^* \in S_i$ is a best response by player i to our assumed set of strategies $\mathbf{s}_{-i} \in S_{-i}$ played by the other $n-1$ players, then it must give player i a payoff—this is the left-hand side of the inequality (7)—at least as large as she would get if she played any other strategy $s_i \in S_i$ from her set of allowed strategies S_i. (This is where the universal quantifier "$\forall s_i \in S_i$" comes in.)

Note that the inequality in definition (7) is weak. The best response may not give player i strictly more than any other choice of strategy.[15] But it is at least tied with any other strategies. In other words, you cannot always talk about *the* best response for a player to some set of plays by everyone else, but you can always talk about *a* best response.[16] Therefore we will not necessarily have a best-response *function* which specifies player i's unique best response to some deleted strategy profile $\mathbf{s}_{-i} \in S_{-i}$, but we

13 "iff" ≡ "if and only if."

14 Note that arg max refers to all values of the argument which maximize the objective function; therefore the arg max yields a set.

15 The weakness of the inequality also makes it possible to write the universal quantifier expression as "$\forall s_i \in S_i$" rather than as "$\forall s_i \in S_i \backslash \{s_i^*\}$".

16 For each $\mathbf{s}_{-i} \in S_{-i}$, define $g: S_i \to \mathbb{R}$ by $g(s_i)=u_i(\langle s_i, \mathbf{s}_{-i}\rangle_i)$. Player i's strategy space S_i is a nonempty, finite set. Therefore its image under g, viz. $g(S_i) \equiv \{g(s_i):\ s_i \in S_i\}$, is a nonempty, finite, ordered set and thus has a maximum and therefore has a maximizer in S_i. Therefore a best response exists.

will have a *best-response correspondence for player i,* $\mathsf{BR}_i\colon S_{-i} \twoheadrightarrow S_i$, where $\mathsf{BR}_i(\boldsymbol{s}_{-i}) \subset S_i$ specifies the *set* of best responses for player i to an arbitrary deleted strategy profile $\boldsymbol{s}_{-i} \in S_{-i}$.[17,18] We can write

$$\mathsf{BR}_i(\boldsymbol{s}_{-i}) = \{s_i^* \in S_i\colon\ s_i^* \text{ is a best response by } i \text{ to } \boldsymbol{s}_{-i}\}, \tag{9}$$

or, more formally,

$$\begin{aligned} \mathsf{BR}_i(\boldsymbol{s}_{-i}) &= \{s_i^* \in S_i\colon \forall s_i \in S_i, u_i(\langle s_i^*, \boldsymbol{s}_{-i}\rangle_i) \geq u_i(\langle s_i, \boldsymbol{s}_{-i}\rangle_i)\}, \\ &= \operatorname*{arg\,max}_{s_i \in S_i} u_i(\langle s_i, \boldsymbol{s}_{-i}\rangle_i). \end{aligned} \tag{10}$$

(We will soon replace this definition of the best-response correspondence BR_i with a more general one.)

Example (continued):

Let's first find a best response for Robin given that Cleever plays "right." She only has two strategies to choose from: Up and Down. We first compute her payoffs for each alternative given that Cleever plays right. We read these from the first element of the payoff pairs in the rightmost column of the matrix in Figure 1:

$$u_{\mathsf{R}}(U, r) = 0,$$

$$u_{\mathsf{R}}(D, r) = 9.$$

Clearly playing Down gives Robin the strictly higher utility given that Cleever plays right. Therefore her best response to Cleever's $s_{\mathsf{C}} = \text{right}$ is $s_{\mathsf{R}}^* = \text{Down}$; we write $\mathsf{BR}_{\mathsf{R}}(r) = \{D\}$.[19]

Note that, against middle by Cleever, Robin is indifferent between Up and Down, because

$$u_{\mathsf{R}}(U, m) = u_{\mathsf{R}}(D, m) = 7.$$

Therefore Robin has two best responses to middle; we write $\mathsf{BR}_{\mathsf{R}}(m) = \{U, D\}$.

Now let's find a best response for Cleever given that Robin will play Down. Cleever has three alternatives, so we compute his utility for each one, given that Robin plays Down, by picking the *second* elements (because Cleever is the second player) from the payoff pairs in the bottom ("Down") row. We get

---

17 Recall that a *correspondence* $\varphi\colon X \twoheadrightarrow Y$ is a "set-valued function" or, more properly, a mapping which associates to every element $x \in X$ in the domain a *subset* of the target set Y. In other words, $\forall x \in X$, $\varphi(x) \subset Y$. The "$\twoheadrightarrow$" symbol is used to distinguish between mapping into subsets of the target set and mapping into the target set itself as the notation "$\rightarrow$" for a function would imply.

18 Let A and B be sets. A is a *subset* of B if every element of A is an element of B. If A is a subset of B, we write $A \subset B$. In this case, we also say that B *contains* A, which we write $B \supset A$. Note from the definition that every set is a subset of itself; i.e. $A \subset A$. Warning: Some authors use the symbol "□" for this meaning and reserve "$\subset$" for "proper subset;" viz. "$A \subset B \Rightarrow (A$ □ B and $A \neq B)$."

19 It might seem somewhat pedantic to include the braces in these cases where the best response correspondence has but a single value; I do so because a correspondence is *set valued.*

$u_C(D,l)=5,$

$u_C(D,m)=4,$

$u_C(D,r)=6.$

We see that Cleever gets his maximum payoff, given that Robin is playing Down, when he plays right. Therefore Cleever's best response to $s_R=$ Down is $s_C{}^*=$ right; i.e. $\mathrm{BR}_C(D)=\{r\}$.

Similarly you can find the best response(s) for either player for any choice of strategy by the other player. It is often a useful prelude to the analysis of a game to determine all the best responses to all possible pure-strategy choices of the other players and to indicate them on the game matrix. For example, in Figure 2, I have made boldface the payoffs which are maximal given the opponent's pure strategy. (I.e. Each of Robin's bolded payoffs is maximal in its column; each of Cleever's bolded payoffs is maximal in its row.)

		C[leever/column]		
		l	*m*	*r*
R[obin/ow]	*U*	**2**,**8**	**7**,7	0,3
	D	1,5	**7**,4	**9**,**6**

Figure 2: Robin vs. Cleever with best responses to others' pure strategies indicated.

Mixed strategies

So far we have restricted our attention to *pure strategies*; these typically have direct economic interpretations, and these are the actions which define the consequent payoffs. If every player plays a pure strategy, then the payoffs to all players are *deterministic*—there is no uncertainty concerning the payoffs resulting from a specified pure-strategy profile.

However, we will find it useful to expand each player's possible choices to include *mixed strategies*—randomizations over her pure strategies. A player will ultimately execute exactly one of her pure-strategy choices. However, which particular pure strategy she executes is determined by a randomization which is specified by her mixed strategy. Furthermore, we stipulate that each player's randomization is independent of the other players' randomizations.[20] When a player $i\in I$ chooses a mixed strategy, every other player $j\in I\backslash\{i\}$ might be uncertain about which pure strategy $s_i\in S_i$ the i-th player will choose.[21,22]

---

20 This assumption is sometimes dropped. This leads to the notion of *correlated strategies*. See for example Aumann [1987].

21 Mixed strategies will be useful when we study equilibrium concepts for games. Equilibrium may fail to exist if we restrict ourselves to pure strategies; existence is guaranteed if we have access to mixed strategies.

22 You might object to mixed strategies because you don't believe that people would actually pick actions randomly. I advise that you just hold on to your objections for a while and accept for the time being that the concept will prove useful. We will later return to a discussion of the various interpretations which can be given to mixed strategies. This discussion will be more meaningful once you

Formally, we say that a mixed strategy σ_i for player i is a *probability distribution* over player i's pure-strategy choices S_i; we write $\sigma_i \in \Delta(S_i)$.[23] If player i is playing the mixed strategy σ_i and if s_i is one of the pure strategies available to her, i.e. $s_i \in S_i$, then the probability that she will choose the pure strategy s_i is denoted by $\sigma_i(s_i)$. (So we could also view the mixed strategy σ_i as a function from player i's pure-strategy space S_i into the unit interval $[0, 1]$, i.e. $\sigma_i\colon S_i \to [0, 1]$.) We denote player i's mixed-strategy space by Σ_i. (We see that $\Sigma_i = \Delta(S_i)$.)

A mixed strategy specifies a value on $[0, 1]$ for each $s_i \in S_i$. Each player chooses one and only one pure strategy $s_i \in S_i$ in a single play of the game. Therefore any mixed strategy $\sigma_i \in \Delta(S_i)$ must be such that the sum of the probabilities associated with the pure strategies is unity, i.e.

$$\sum_{s_i \in S_i} \sigma_i(s_i) = 1. \tag{11}$$

This property is indeed satisfied as a result of σ_i being a probability distribution over S_i. This is the justification for using a probability distribution to represent a mixed strategy.

Analogously to our treatment of pure strategies: A *mixed-strategy profile* $\boldsymbol{\sigma}$ is an n-tuple of individual mixed strategies, one for each player; e.g. $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$, where $\forall i \in I$, $\sigma_i \in \Sigma_i$. The *space of mixed-strategy profiles* Σ is just the Cartesian product of the individuals' mixed-strategy spaces, viz. $\Sigma \equiv \mathsf{X}_{i \in I}\, \Sigma_i$. Player i's deleted mixed-strategy space is $\Sigma_{-i} = \mathsf{X}_{j \in I \setminus \{i\}}\, \Sigma_j$. We can also extend the domain of the $\langle \cdot, \cdot \rangle_i$ function in the natural way so that $\boldsymbol{\sigma} = \langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i$.[24]

We say that the *support* of the mixed strategy $\sigma_i \in \Sigma_i$ is the set of pure strategies to which σ_i assigns positive probability,[25] i.e.

$$\operatorname{supp}(\sigma_i) \equiv \{s_i \in S_i\colon \sigma_i(s_i) > 0\}. \tag{12}$$

I.e. the support of the mixed strategy σ_i consists of those pure strategies which player i could conceivably play if she chose the mixed strategy σ_i. For all $\sigma_i \in \Sigma_i$, $\operatorname{supp}(\sigma_i) \subset S_i$.

The unit simplex

At this point we pause to define the $(k-1)$-dimensional *unit simplex*. The $(k-1)$-dimensional unit simplex is the set of k-vectors whose components 1 are all nonnegative and 2 sum to one.[26,27] This

---

have seen how the theory develops; therefore this is a case where the motivation and interpretation are better left to follow rather than precede the exposition.

23 When A is a finite set, $\Delta(A)$ is the set of all probability distributions over A.

24 I.e. we extend the domain of $\langle \cdot, \cdot \rangle_i$ to include $\Sigma_i \times \Sigma_{-i}$ so that $\langle \cdot, \cdot \rangle_i\colon [(S_i \times S_{-i}) \cup (\Sigma_i \times \Sigma_{-i})] \to (S \cup \Sigma)$, where the restriction of $\langle \cdot, \cdot \rangle_i$ to $S_i \times S_{-i}$ is $\langle \cdot, \cdot \rangle_i\colon (S_i \times S_{-i}) \to S$ and the restriction to $\Sigma_i \times \Sigma_{-i}$ is $\langle \cdot, \cdot \rangle_i\colon (\Sigma_i \times \Sigma_{-i}) \to \Sigma$. Both restrictions have the same symbolic definition: $\langle a_i, (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n) \rangle_i = (b_1, \ldots, b_{i-1}, a_i, b_{i+1}, \ldots, b_n)$.

25 This is analogous to the support of a random variable, which is the closure of the set of values which are assigned positive probability.

26 Note that these two conditions taken together guarantee that each component is weakly less than one.

27 This simplex is composed of vectors with k components. It gets its name, viz. "$k-1$," because it is a $(k-1)$-dimensional subspace embedded in a k-dimensional world. $(k-1)$ is also the number of "degrees of freedom" each vector has: Once $k-1$ components have

concept is useful here because any probability distribution over a finite set must belong to a unit simplex.[28] Formally, we define the $(k-1)$-dimensional unit simplex as[29]

$$\Delta^{k-1}=\left\{\boldsymbol{x}\in\mathbb{R}_+^k\colon \sum_{j=1}^{k} x_j=1.\right\} \tag{13}$$

For example the vectors $(^1/_3, 0, ^2/_3)$ and $(0, 0, 1)$ are members of the two-dimensional unit simplex. The vectors $(1, 1, 1)$ and $(^2/_3, ^2/_3, -^1/_3)$ are not. Figure 3 displays the zero-, one-, and two-dimensional simplices.[30]

Mixed strategies are chosen from the unit simplex

An alternative to conceiving of a mixed strategy for player i as a function $\sigma_i\colon S_i\to[0,1]$ is to think of it as a $\#S_i$-vector of probabilities. For example, let $m=\#S_i$ be the number of pure strategies available to player i, and then index player i's m pure strategies with a superscript: $s_i^1, \ldots, s_i^k, \ldots, s_i^m$. Then we can write the mixed strategy $\sigma_i\in\Sigma_i$ as the m-tuple

$$\sigma_i=(\sigma_i(s_i^1), \ldots, \sigma_i(s_i^m)). \tag{14}$$

We know 1 for every pure strategy $s_i\in S_i$, $\sigma_i(s_i)\in[0,1]$ and 2 that (11) holds. Referring to definition (13), then, we see that player i's mixed strategy σ_i belongs to the $(\#S_i-1)$-dimensional unit simplex. (In (13), let $k=\#S_i$ and, $\forall j\in\{1, \ldots, m\}$, let $x_j=\sigma_i(s_i^j)$.)

---

been specified, the remaining component is already determined by the requirement that the components sum up to unity.

28 It is useful elsewhere as well. For example in general equilibrium theory prices are often normalized to lie within a unit simplex. Any unit simplex is convex, and this allows the invocation of Brouwer's or Kakutani's fixed-point theorem in order to prove the existence of an equilibrium price vector.

29 Recall that $\mathbb{R}_+^k$ is the nonnegative orthant of $\mathbb{R}^k$, viz. $\{\boldsymbol{x}\in\mathbb{R}^k\colon \forall i\in\{1,\ldots,k\}, x_i\geq 0\}$.

30 To be perfectly clear… in Figure 3b the one-dimensional simplex is only the line segment connecting $(0,1)\to(1,0)$. In Figure 3c the two-dimensional simplex is the triangular planar region.

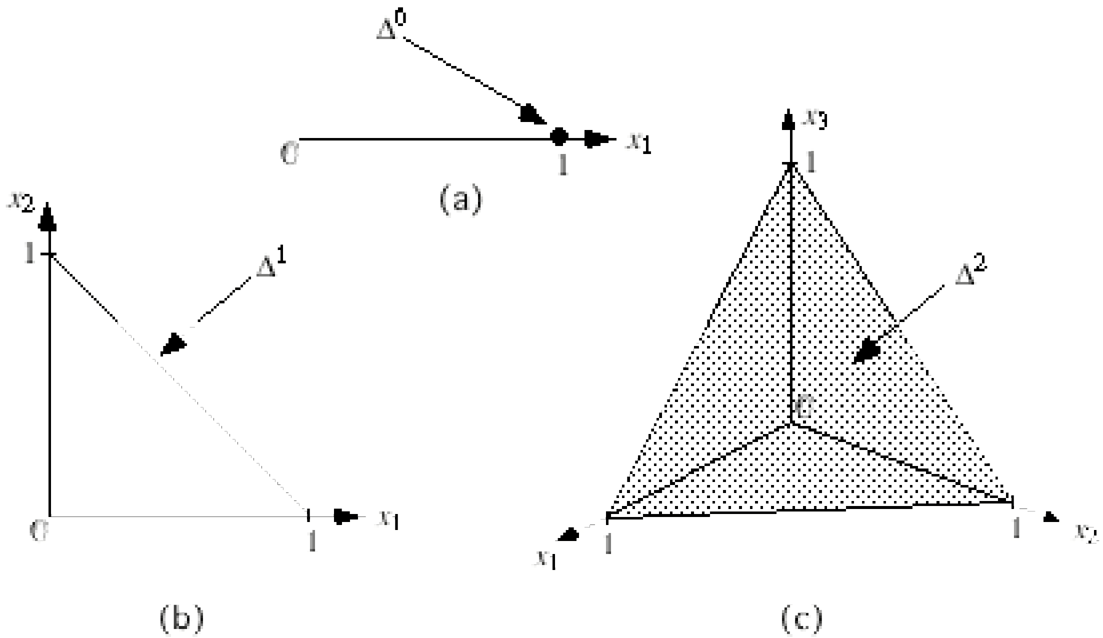

Figure 3: The (a) zero-, (b) one-, and (c) two-dimensional simplices.

Recall that a player's pure-strategy space S_i is the set of possible pure strategies she can choose and that these strategies can be very complex descriptions of contingent actions. A player's *mixed-strategy space* Σ_i is simply the set of $\#S_i$-vectors belonging to the $(\#S_i-1)$-dimensional simplex $\Delta^{\#S_i-1}$.[31]

Pure strategies are degenerate mixed strategies

Choosing a pure strategy s_i is equivalent to choosing the mixed strategy (i.e. probability distribution over pure strategies) which results with probability one in s_i. Therefore we see that every pure strategy "is" a mixed strategy.[32] We also say that a pure strategy is a *degenerate* mixed strategy.

To avoid the abuse of notation which results from writing s_i for the degenerate mixed strategy which plays the pure strategy s_i with certainty, we can define $\delta_i(s_i)\in\Delta(S_i)$ to be the player-i mixed strategy which puts unit weight on $s_i\in S_i$ and zero weight on every other player-i pure strategy $s_i'\in S_i\backslash\{s_i\}$. We formally define, for every player-i pure strategy $s_i\in S_i$, the degenerate probability distribution $\delta_i(s_i)$ by specifying, for each player-i pure strategy $s_i'\in S_i$, the probability $\delta_i(s_i)(s_i')\in[0,1]$ which it attaches to s_i':

$$\delta_i(s_i)(s_i')=\begin{cases}1, & s_i'=s_i,\\ 0, & s_i'\neq s_i.\end{cases} \tag{15}$$

(So we see that, for every $s_i\in S_i$, $\delta_i(s_i)\colon S_i\to\{0,1\}$. Alternatively, we can write $\delta_i\colon S_i\to\Delta(S_i)$ or $\delta_i\colon S_i^2\to[0,1]$.)

---

31 It might appear at first that a player's mixed-strategy space Σ_i is simpler than her pure-strategy space S_i. Keep in mind however that Σ_i is defined in terms of S_i and so inherits all of S_i's complexity.

32 I put "is" in quotes because mathematically they are different objects. The pure strategy and the corresponding mixed strategy which puts all probability on that pure strategy are equivalent in the sense that they result in exactly the same action by the player.

A *nondegenerate* mixed strategy is one that is not pure. [Therefore $\exists s_i \in S_i$ such that $\sigma_i(s_i) \in (0,1)$.] A *completely mixed* or *strictly positive* strategy is one which puts positive weight on every strategy in the player's strategy space; i.e. σ_i is completely mixed if $\forall s_i \in S_i,\ \sigma_i(s_i) > 0$.[33] (Therefore $\operatorname{supp} \sigma_i = S_i$.)

Example (continued)

Let's return to our game between Robin and Cleever to see how these mixed-strategy concepts are represented in a less notational and less abstract setting. In Figure 4, I have indicated each player's mixed strategies with bracketed probabilities attached to the pure strategies.

		C[leever/olumn]		
		l: [p]	m: [q]	r: [$1-p-q$]
R[obin/ow]	U: [t]	2,8	7,7	0,3
	D: [$1-t$]	1,5	7,4	9,6

Figure 4: Robin vs. Cleever with mixed strategies denoted.

Robin has two pure strategies, so the cardinality of her strategy space is 2; i.e. $\#S_{\mathrm{R}} = 2$. Therefore her mixed strategies lie in a one-dimensional unit simplex; they can be described by a single parameter t. We can write any of Robin's mixed strategies as an ordered pair which specifies the probability with which she would choose Up and Down, respectively, i.e. in the form

$$\sigma_{\mathrm{R}} = (t, 1-t).$$

Alternatively we could write

$$\mathbb{P}(\text{Robin chooses Up}) = \sigma_{\mathrm{R}}(U) = t,$$

$$\mathbb{P}(\text{Robin chooses Down}) = \sigma_{\mathrm{R}}(D) = 1-t.$$

Robin's mixed-strategy space, i.e. the set of all possible mixed strategies for her, is

$$\Sigma_{\mathrm{R}} = \{(t, 1-t) \colon t \in [0,1]\}.$$

Note that we have already seen the graph of Σ_{R} in Figure 3b.

Cleever has three pure strategies, therefore a mixed strategy for him belongs to the two-dimensional unit simplex and takes the form

$$\sigma_{\mathrm{C}} = (p, q, 1-p-q),$$

where

33 The notion of completely mixed strategies is used when discussing some equilibrium refinements, e.g. sequential equilibrium and trembling-hand perfection. When all players choose completely mixed strategies, there is a positive probability of reaching any given node in the game tree. Therefore no node is off the path. (These comments will make more sense after we encounter games in extensive form.)

$$\mathbb{P}(\text{Cleever chooses left}) = \sigma_C(l) = p,$$

$$\mathbb{P}(\text{Cleever chooses middle}) = \sigma_C(m) = q,$$

$$\mathbb{P}(\text{Cleever chooses right}) = \sigma_C(r) = 1 - p - q.$$

His mixed-strategy space is:

$$\Sigma_C = \{(p, q, 1 - p - q) \colon p, q \geq 0,\ p + q \leq 1\}.$$

We have already seen the graph of Σ_C in Figure 3c.

Payoffs to mixed-strategy profiles

We have noted that $u_i(\boldsymbol{s})$ is player i's payoff when the players choose their parts of the pure-strategy profile $\boldsymbol{s} \in S$. Because the players are not randomizing their actions when they play $\boldsymbol{s}$, the resultant payoff vector is a certain, deterministic number. Now we ask the question: When the players execute the mixed-strategy profile $\boldsymbol{\sigma} \in \Sigma$, what is the payoff to player i? Right away we see a problem even with the way this question is phrased. It doesn't make sense to ask *ex ante* what *the* payoff to player i is, because her payoff depends on the precise pure strategies realized as the result of the individuals' randomizations.

We could ask then: What is the *distribution* of payoffs player i would receive if the players executed the mixed-strategy profile $\boldsymbol{\sigma}$? Fortunately, we have no need for such a complicated answer. Because our utility functions $u_i \colon S \to \mathbb{R}$ are assumed to be of the von Neumann–Morgenstern variety, we know that each player's preferences over distributions of von Neumann–Morgenstern utilities can be represented by her *expected utility*. Now we need only ask: What is player i's expected payoff given that the players choose $\boldsymbol{\sigma} \in \Sigma$? We will simplify notation by using the same function name u_i to represent the expected utility to player i from a mixed-strategy profile $\boldsymbol{\sigma} \in \Sigma$ as we used above for pure-strategy profiles. I.e. we write $u_i(\boldsymbol{\sigma})$, where $u_i \colon \Sigma \to \mathbb{R}$.[34]

The probability of a pure-strategy profile $\boldsymbol{s}$

How do we calculate this expected utility for player i? We need to weight player i's payoff to each arbitrary pure-strategy profile $\boldsymbol{s} = \langle s_i, \boldsymbol{s}_{-i} \rangle_i \in S$ by the probability that the profile $\boldsymbol{s}$ will be realized when the players randomize according to the mixed-strategy profile $\boldsymbol{\sigma} \in \Sigma$. Because the players' randomizations are independent of one another's, the probability that $\boldsymbol{s} = (s_1, \ldots, s_n)$ will occur is the

---

34 It is a formal convenience here to use the same function name for two different functions with distinct domains. Although this may appear abusive *prima facie*, the more complete justification is the following: Let A, B, and C be sets such that $A \cap B = \varnothing$. Let $g \colon A \to C$ and $h \colon B \to C$ be functions. Then we can define $f \colon (A \cup B) \to C$ by

$$f(x) = \begin{cases} g(x), & x \in A, \\ h(x), & x \in B. \end{cases}$$

I.e. we can define $u_i \colon (S \cup \Sigma) \to \mathbb{R}$. When the argument supplied to u_i is an element of S (respectively, Σ), the function is evaluated using the restriction of u_i to S (respectively, Σ).

product of the probabilities that each player j will play s_j. The probability, according to $\boldsymbol{\sigma}$, that j will play $s_j \in S_j$ is $\sigma_j(s_j)$. Therefore the probability that $\boldsymbol{s}$ will occur when the players randomize according to $\boldsymbol{\sigma}$ is the product of probabilities

$$\begin{aligned}\mathbb{P}(\boldsymbol{s}\text{ is played}) &= \mathbb{P}((s_1,\ldots,s_n)\text{ is played}) = \mathbb{P}(1\text{ plays }s_1)\cdot\cdots\cdot\mathbb{P}(n\text{ plays }s_n) \\ &= \sigma_1(s_1)\cdot\cdots\cdot\sigma_n(s_n) = \prod_{j=1}^{n}\sigma_j(s_j) = \prod_{j\in I}\sigma_j(s_j).\end{aligned} \tag{16}$$

Expected payoff to a mixed-strategy profile $\boldsymbol{\sigma}$

To complete our calculation of i's expected utility when the mixed-strategy profile $\boldsymbol{\sigma}$ is played, we must look at every possible pure-strategy profile $\boldsymbol{s}\in S$, find i's deterministic payoff for this pure-strategy profile, and weight this payoff according to the profile's probability of occurrence as given by (16). The weighted sum over all these possible pure-strategy profiles is our desired expected payoff. I.e. the expected payoff to player i when the players participate in the mixed-strategy profile $\boldsymbol{\sigma}$ is[35]

$$u_i(\boldsymbol{\sigma}) = \sum_{\boldsymbol{s}\in S}\mathbb{P}(\boldsymbol{s})u_i(\boldsymbol{s}) = \sum_{\boldsymbol{s}\in S}\left(\prod_{j\in I}\sigma_j(s_j)\right)u_i(\boldsymbol{s}). \tag{17}$$

Payoff to i from $\boldsymbol{\sigma}$ is linear in any one player's mixing probabilities

We can single out for special attention any player $k\in I$ in our calculation of player i's payoff to a mixed-strategy profile $\boldsymbol{\sigma}\in\Sigma$ and rewrite (17) as

$$\begin{aligned}u_i(\boldsymbol{\sigma}) &= \sum_{s_k\in S_k}\sum_{\boldsymbol{s}_{-k}\in S_{-k}}\left(\sigma_k(s_k)\prod_{j\in I\setminus\{k\}}\sigma_j(s_j)\right)u_i(\langle s_k,\boldsymbol{s}_{-k}\rangle_k) \\ &= \sum_{s_k\in S_k}\sigma_k(s_k)\left[\sum_{\boldsymbol{s}_{-k}\in S_{-k}}\left(\prod_{j\in I\setminus\{k\}}\sigma_j(s_j)\right)u_i(\langle s_k,\boldsymbol{s}_{-k}\rangle_k)\right] \\ &= \sum_{s_k\in S_k}c_k(s_k,\boldsymbol{\sigma}_{-k})\sigma_k(s_k),\end{aligned} \tag{18}$$

where $\forall k\in I$, c_k: $S_k\times\Sigma_{-k}\to\mathbb{R}$ is defined by

35 You may be more familiar with writing a summation (and similar remarks hold for products) in the form $\Sigma_{k=1}^{m}x_k$, i.e. with an integer index k to indicate a particular element of a finite set of objects $X=\{x_1,\ldots,x_m\}$ to be added. We will often find it more convenient—e.g. when there is no natural indexing scheme—to write this summation as $\Sigma_{x\in X}x$. This simply means to form a sum whose terms consist of every element of X represented once. This is equivalent to the indexed formalism, for both summations and products, because both (finite) addition and multiplication are commutative operations.

$$c_k(s_k, \boldsymbol{\sigma}_{-k}) \equiv \sum_{\boldsymbol{s}_{-k} \in S_{-k}} \left(\prod_{j \in I \setminus \{k\}} \sigma_j(s_j) \right) u_i(\langle s_k, \boldsymbol{s}_{-k} \rangle_k). \tag{19}$$

Equations (18) and (19) say that player i's expected payoff in the mixed-strategy profile $\boldsymbol{\sigma}$ is a linear function of player k's mixing probabilities $\{\sigma_k(s_k)\}_{s_k \in S_k}$. (To see this note that, for each $s_k \in S_k$, the corresponding coefficient $c_k(s_k, \boldsymbol{\sigma}_{-k})$ is independent of $\sigma_k(s_k')$ for all $s_k' \in S_k$.) This observation will be relevant when we determine a player's best-response correspondence. (Note that this analysis includes the case $k = i$.)

Player i's payoff to a pure strategy s_i against a deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i}$

We will see that it will be very useful to determine player i's payoff against a deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ when player i herself chooses some pure strategy $s_i \in S_i$.

To represent the mixed-strategy profile $\boldsymbol{\sigma} \in \Sigma$ induced by this combination, we again extend the domain of the $\langle \cdot, \cdot \rangle_i$ function to include $S_i \times \Sigma_{-i}$ so that we can make sense of the expression "$\boldsymbol{\sigma} = \langle s_i, \boldsymbol{\sigma}_{-i} \rangle_i$." We stipulate the restriction of $\langle \cdot, \cdot \rangle_i$ to $S_i \times \Sigma_{-i}$ to be a function $\langle \cdot, \cdot \rangle_i : S_i \times \Sigma_{-i} \to \Sigma$ defined by $\langle s_i, \boldsymbol{\sigma}_{-i} \rangle_i = \langle \delta_i(s_i), \boldsymbol{\sigma}_{-i} \rangle_i$.[36,37] In other words, we replace the pure strategy s_i with the degenerate mixed strategy $\delta_i(s_i)$ which puts all of its weight on s_i.

We calculate the expected payoff to player i when she plays the pure strategy $s_i' \in S_i$ against the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ using (18), (19), and (15), where we let $k \leftarrow i$ and let $\boldsymbol{\sigma} = \langle s_i', \boldsymbol{\sigma}_{-i} \rangle_i$; i.e. $\sigma_i = \delta_i(s_i')$:

$$\begin{aligned} u_i(\langle s_i', \boldsymbol{\sigma}_{-i} \rangle_i) = u_i(\langle \delta_i(s_i'), \boldsymbol{\sigma}_{-i} \rangle_i) &= \sum_{s_i \in S_i} c_i(s_i, \boldsymbol{\sigma}_{-i}) \sigma_i(s_i) \\ &= \sum_{s_i \in S_i} \left(\delta_i(s_i')(s_i) \right) c_i(s_i, \boldsymbol{\sigma}_{-i}) = c_i(s_i', \boldsymbol{\sigma}_{-i}) \\ &= \sum_{\boldsymbol{s}_{-i} \in S_{-i}} \left(\prod_{j \in I \setminus \{i\}} \sigma_j(s_j) \right) u_i(\langle s_i', \boldsymbol{s}_{-i} \rangle_i). \end{aligned} \tag{20}$$

Note then that we have shown that $\forall s_i \in S_i$, $c_i(s_i, \boldsymbol{\sigma}_{-i}) = u_i(\langle s_i, \boldsymbol{\sigma}_{-i} \rangle_i)$. Therefore from (18), letting $k \leftarrow i$, we can rewrite player i's expected payoff to a mixed strategy $\sigma_i \in \Sigma_i$ against the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ as:

---

36 Let $f: X \to Z$ be a function and let $Y \subset X$ be a subset of X. Then we can define a function $\tilde{f}$, the *restriction* of f to Y, as a function whose domain is Y and which agrees with f for all points in Y. I.e. $\tilde{f}: Y \to Z$ and $\forall x \in Y, \tilde{f}(x) = f(x)$.

37 To complete the definition of $\langle \cdot, \cdot \rangle_i$ we also extend its domain to include the set $\Sigma_i \times S_{-i}$, i.e. where player i chooses a mixed strategy and the other players choose pure strategies. (This will be handy in our later analysis of strategic dominance.) We provide the obvious definition for $\langle \cdot, \cdot \rangle_i$: $\Sigma_i \times S_{-i} \to \Sigma$: $\langle a_i, (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n) \rangle_i = (\delta_1(b_1), \ldots, \delta_{i-1}(b_{i-1}), a_i, \delta_{i+1}(b_{i+1}), \ldots, \delta_n(b_n))$. Therefore now $\langle \cdot, \cdot \rangle_i$: $[(S_i \cup \Sigma_i) \times (S_{-i} \times \Sigma_{-i})] \to (S \cup \Sigma)$.

$$u_i(\langle\sigma_i,\boldsymbol{\sigma}_{-i}\rangle_i)=\sum_{s_i\in S_i} c_i(s_i,\boldsymbol{\sigma}_{-i})\sigma_i(s_i), \tag{21}$$

or, for more convenient future reference:

$$u_i(\langle\sigma_i,\boldsymbol{\sigma}_{-i}\rangle_i)=\sum_{s_i\in S_i}\sigma_i(s_i)u_i(\langle s_i,\boldsymbol{\sigma}_{-i}\rangle_i)=\sum_{s_i\in \operatorname{supp}\sigma_i}\sigma_i(s_i)u_i(\langle s_i,\boldsymbol{\sigma}_{-i}\rangle_i). \tag{22}$$

In other words a player's payoff to a mixed strategy (against some fixed deleted mixed strategy profile) is a convex combination of the payoffs to the pure strategies in the mixed strategy's support (against that deleted mixed strategy profile). (The set $\{\sigma_i(s_i): s_i\in\operatorname{supp}\sigma_i\}$ of coefficients is a set of convex coefficients because they are nonnegative and sum to unity.)

Example (continued):

Let's employ the very notational expression (17) to compute the payoff to Robin for an arbitrary mixed-strategy profile $\boldsymbol{\sigma}=(\sigma_{\mathrm{R}},\sigma_{\mathrm{C}})$. Note that the summation $\Sigma_{s\in S}$ of (17) generates six terms, viz. one for each member of

$$S=S_{\mathrm{R}}\times S_{\mathrm{C}}=\{(U,l),(U,m),(U,r),(D,l),(D,m),(D,r)\}.$$

For each of these terms the $\Pi_j\,\sigma_j(s_j)$ product multiplies two factors: $\sigma_{\mathrm{R}}(s_{\mathrm{R}})$ and $\sigma_{\mathrm{C}}(s_{\mathrm{C}})$. For example, when $s=(D,l)$,

$$\Pi_j\,\sigma_j(s_j)=\sigma_{\mathrm{R}}(D)\sigma_{\mathrm{C}}(l)=(1-t)p.$$

This product, then, is the weight attached to $u_{\mathrm{R}}((D,l))=1$ when we calculate Robin's expected payoff when the mixed-strategy profile $\boldsymbol{\sigma}$ is played.

We can use the game matrix from Figure 4 to easily compute the probability coefficients associated with each pure-strategy profile. See Figure 5. This matrix of probability coefficients was formed by multiplying the mixing probability of Robin's associated with a cell's row by the mixing probability of Cleever's associated with that cell's column. Inspection of Figure 5 shows quickly what we had already determined: that the probability coefficient corresponding to (D,l) is $(1-t)p$.

		C[leever/olumn]		
		l: $[p]$	m: $[q]$	r: $[1-p-q]$
R[obin/ow]	U: $[t]$	tp	tq	$t(1-p-q)$
	D: $[1-t]$	$(1-t)p$	$(1-t)q$	$(1-t)(1-p-q)$

Figure 5: The probability coefficients which weight the pure-strategy profile payoffs in the calculation of the expected payoff to an arbitrary mixed-strategy profile.

To compute Robin's expected payoff to the mixed-strategy profile $\boldsymbol{\sigma}$, then, we multiply her payoff in each cell by the probability coefficient given in that cell of Figure 5, and then sum over all the cells. For example, consider the mixed-strategy profile $\boldsymbol{\sigma}=(\sigma_{\mathrm{R}},\sigma_{\mathrm{C}})$ where Robin mixes between Up and Down according to $\sigma_R = (^1/_6, ^5/_6)$ and Cleever mixes according to $\sigma_C = (^9/_{10}, 0, ^1/_{10})$. You can easily verify that Robin's expected payoff for this mixed-strategy profile is

$$u_R(\sigma) = \frac{1}{6} \cdot \frac{9}{10} \cdot 2 + \frac{5}{6} \cdot \frac{9}{10} \cdot 1 + \frac{5}{6} \cdot \frac{1}{10} \cdot 9 = \frac{9}{5}.$$

The best-response correspondence

We have earlier considered player i's problem of deciding on a *best-response pure strategy* $s_i^* \in S_i$ to some deleted pure-strategy profile [i.e. $(n-1)$-tuple] $\boldsymbol{s}_{-i} \in S_{-i}$ of pure-strategy choices by her opponents. In her calculations for a fixed $\boldsymbol{s}_{-i}$ she was certain that player $j \in I \backslash \{i\}$ would play a particular $s_j \in S_j$.

Now we ask: given that all other players but i are playing the deleted *mixed-strategy* profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$, what pure strategy is best for i? The answer to this question is i's best-response correspondence $\mathrm{BR}_i \colon \Sigma_{-i} \twoheadrightarrow S_i$, which maps the space of deleted mixed-strategy profiles Σ_{-i} into subsets of the space of i's pure strategies S_i. (This definition of the best-response correspondence BR_i is a generalization and replacement of the earlier definition which considered only pure-strategies by the other players.)

Formally we write player i's problem as finding, for every deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$, the set $\mathrm{BR}_i(\boldsymbol{\sigma}_{-i}) \subset S_i$ of pure strategies for player i :

$$\mathrm{BR}_i(\boldsymbol{\sigma}_{-i}) = \arg\max_{s_i \in S_i} u_i(\langle s_i, \boldsymbol{\sigma}_{-i} \rangle_i). \tag{23}$$

Nonemptiness of $\mathrm{BR}_i(\boldsymbol{\sigma}_{-i})$ (i.e. the existence of a best-response pure strategy) is guaranteed for each $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ because S_i is a nonempty and finite set.

Example (continued)

Let's compute Cleever's and Robin's best-response correspondences to the other's arbitrary mixed strategy. Any mixed strategy Robin chooses can be described as a choice of $t \in [0, 1]$. We first seek Cleever's best-response correspondence $\mathrm{BR}_C(t)$, which specifies all of Cleever's pure strategies which are best responses to Robin's mixed strategy $\sigma_R = (t, 1-t)$.

To determine this correspondence we compute Cleever's payoffs to each of his three pure strategies against Robin's arbitrary mixed strategy t. Each pure-strategy choice by Cleever corresponds to a column in Figure 4. We then look at the second element of each ordered pair in that column, because that component corresponds to Cleever's payoff, and weight each one by the probability that its row will be chosen by Robin, viz. by t and $(1-t)$ for Up and Down, respectively. This process yields:

$$u_C(l; t) = 8t + 5(1-t) = 5 + 3t,$$

$$u_C(m; t) = 7t + 4(1-t) = 4 + 3t,$$

$$u_C(r; t) = 3t + 6(1-t) = 6 - 3t.$$

We plot Cleever's three pure-strategy payoffs as functions of Robin's mixed strategy $t \in [0, 1]$ in Figure 6.

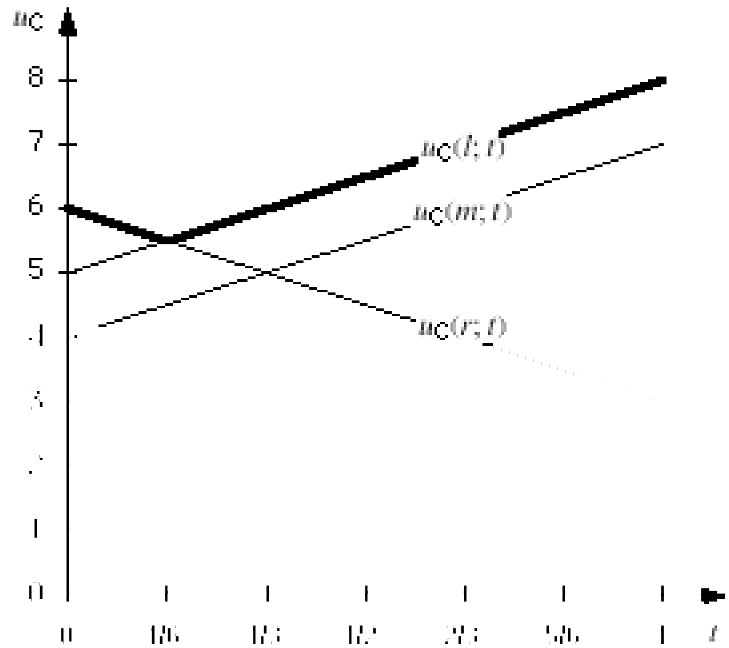

Figure 6: Cleever's pure-strategy payoffs as functions of Robin's mixed strategy.

We first observe that Cleever's payoff to left is strictly above his payoff to middle. (I.e. $\forall t \in [0, 1]$, $u_C(l; t) > u_C(m; t)$.) We will see later that this means that middle is *strictly dominated* by left. We also see that, when $t < {}^1/_6$, right supplies Cleever with a higher payoff than either of the other strategies. When $t > {}^1/_6$, left provides the highest payoff. When $t = {}^1/_6$, both left and right provide Cleever with a payoff of $5\frac{1}{2}$. In the first two cases Cleever has a unique best response to Robin's mixed strategy. In the last case two of Cleever's strategies are best responses.

To summarize we can write Cleever's best-response correspondence as:

$$\mathrm{BR}_C(t) = \begin{cases} \{r\}, & t \in \left[0, 1/6\right), \\ \{r, l\}, & t = 1/6, \\ \{l\}, & t \in \left(1/6, 1\right]. \end{cases}$$

We can represent this best-response correspondence graphically by mapping the relevant intervals describing Robin's mixed strategy into pure-strategy choices by Cleever. (See Figure 7.)

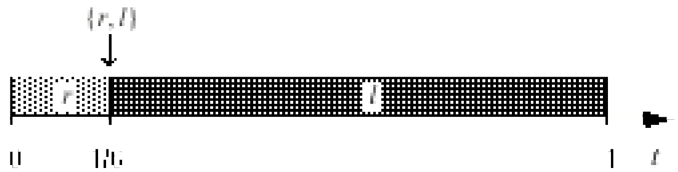

Figure 7: Cleever's best-response correspondence for three subsets of Robin's mixed-strategy space.

The *upper envelope* of these three payoff functions is indicated by heavier line segments in Figure 6.[38] This represents the expected payoff Cleever would receive, as a function of Robin's mixed strategy t, if Cleever played a best response to this mixed strategy.

Now we'll determine Robin's best-response correspondence as a function of Cleever's mixed strategy σ_{C}. Because Cleever has three pure strategies to choose from, we need two parameters to describe his arbitrary mixed strategy, viz. p and q. Analogously as we did above, we compute Robin's expected payoff to each of her two pure strategies as a function of Cleever's mixed-strategy parameters:

$$u_{\mathrm{R}}(U;p,q)=2p+7q+0\cdot(1-p-q)=2p+7q,$$

$$u_{\mathrm{R}}(D;p,q)=1\cdot p+7q+9(1-p-q)=9-8p-2q.$$

Robin should weakly prefer to play Up whenever

$$u_{\mathrm{R}}(U;p,q)\geq u_{\mathrm{R}}(D;p,q),$$

which occurs when

$$q\geq 1-\frac{10}{9}p.$$

The isosceles right triangle in Figure 8 represents Cleever's mixed-strategy space in the following sense: Every mixed strategy of Cleever's can be represented by a (p,q) pair satisfying $p,q\geq 0$ and $p+q\leq 1$. Therefore there is a one-to-one correspondence between Σ_{C} and the points in that triangle.[39] Also marked is the line segment of mixed strategies of Cleever's at which Robin is indifferent between playing Up and Down. On that line segment Robin's best-response correspondence contains both pure strategies. Points above that line segment represent Cleever mixed strategies against which Robin strictly prefers to play Up; below that she strictly prefers to play Down. Robin's best-response correspondence can be written as

---

38 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: $\bar{f}(x)\equiv\max\{f_1(x),\ldots,f_n(x)\}$.

39 The isosceles triangle is just the projection of Σ_{C} $(=\Delta^2)$ onto the pq-axis. [In Figure 3(c) just drop each point on the shaded triangle perpendicularly down onto the x_1x_2-axis to see where the isosceles triangle in Figure 8 comes from.]

$$\mathsf{BR}_{\mathsf{R}}(p,q)=\begin{cases}\{U\}, & q>1-\dfrac{10}{9}p,\\[2ex] \{U,D\}, & q=1-\dfrac{10}{9}p,\\[2ex] \{D\}, & q<1-\dfrac{10}{9}p,\end{cases}$$

with the understanding that $(p,q)\in\mathbb{R}_+^2$ and $p+q\leq 1$.

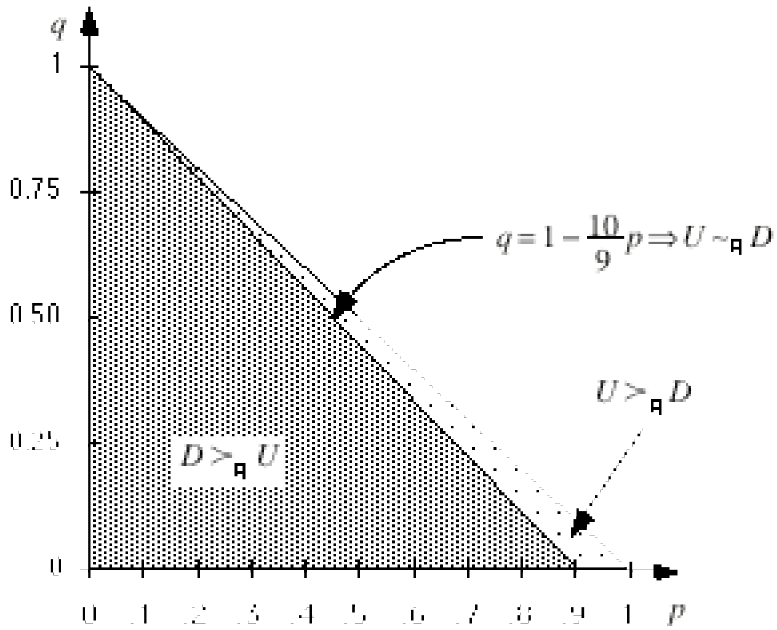

Figure 8: Robin’s best-response correspondence for three subsets of Cleever’s mixed-strategy space.

Best-response mixed strategies

Above we defined a player’s best *pure-strategy* response(s) to a given deleted profile of other players’ mixed strategies. (To be even more precise… we determined which of a player’s pure strategies were best responses *within her set of pure strategies*.) But how do we know that a player’s best response is a pure strategy? Could she do better by playing a mixed strategy? We will see that, *given her opponents’ strategies*, a player would never strictly prefer to play a mixed strategy over one of her pure-strategy best responses.[40] In fact the only time when—again, against a particular $\boldsymbol{\sigma}_{-i}\in\Sigma_{-i}$—a player would even be willing to mix is when her best-response correspondence for that deleted strategy profile contains more than one pure strategy; i.e. when $\#\mathsf{BR}_i(\boldsymbol{\sigma}_{-i})>1$. When that is true, she is willing to put positive

40 That doesn’t mean that mixed strategies aren’t useful. Even if a player is indifferent between playing a mixed strategy and a pure strategy against some particular set of opponents’ strategies, playing a mixture has the effect of making her opponents uncertain about what she will do. This can cause them to choose their strategies more beneficially for her.

weight on any pure-strategy best response.

Formally… a player-i mixed strategy $\sigma_i^* \in \Sigma_i$ is a best response for player i against the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ if

$$\sigma_i^* \in \operatorname*{arg\,max}_{\sigma_i \in \Sigma_i} u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i). \tag{24}$$

It might help slice through the notational fog if we index with k the possible pure strategies for i and let $m = \#S_i$ be the number of pure strategies for i. For every $k \in \{1, \ldots, m\}$, let 1 s_i^k be i's k-th pure strategy, 2 p_i^k be the probability of s_i^k according to the mixed strategy $\sigma_i \in \Sigma_i$ (i.e. $p_i^k \equiv \sigma_i(s_i^k)$), and 3 u_i^k be the payoff to i against $\boldsymbol{\sigma}_{-i}$ when she plays her k-th pure strategy s_i^k; i.e. $u_i^k \equiv u_i(\langle s_i^k, \sigma_{-i} \rangle_i)$. Now we can write the maximization problem for a player seeking an optimal mixed-strategy response as that of choosing an m-vector $\boldsymbol{p}_i = (p_i^1, \ldots, p_i^m) \in \Delta^{m-1}$ of probabilities which solves

$$\max_{\boldsymbol{p}_i \in \Delta^{m-1}} p_i^1 u_i^1 + p_i^2 u_i^2 + \cdots + p_i^m u_i^m. \tag{25}$$

Note, as we showed before in (18), that player i's payoff is linear in her mixing probabilities $p_i^1, \ldots, p_i^m$. And as we showed in (22) player i's payoff to her mixed strategy $\boldsymbol{p}_i$ is a convex combination of her pure strategy payoffs $\{u_i^k\}_{k=1}^m$.

In the above optimization problem player i must assign probabilities to her different pure-strategy choices in such a way as to maximize the probability-weighted sum of her payoffs from those pure strategies. In this type of problem probability is a scarce resource: the more of it one pure strategy gets, the less another receives. Therefore you want to put your probability where it counts the most.

Consider the case in which one pure strategy s_i^k is strictly better than any of the other pure strategies; i.e. u_i^k is strictly larger than all of the other pure-strategy payoffs. Then the pure strategy s_i^k should receive all of the probability; i.e. p_i^k should be unity and all of the other probabilities should be zero. (Otherwise the objective function could be increased by shifting probability away from a pure strategy whose payoff is lower.) This would correspond to playing the pure strategy s_i^k.

Now consider the case in which several pure strategies are best. E.g. s_i^k and s_i^r, $k \neq r$, both result in the payoff $\bar{u}_i \equiv u_i^k = u_i^r$ which is strictly larger than all of the other pure-strategy payoffs. We should definitely not waste probability on any of these other, low-performance pure strategies, because we could increase our expected payoff by shifting that probability to either of these best pure strategies. However, we are indifferent to how much of our probability we assign to the various best pure strategies. E.g., because s_i^k and s_i^r both result in the payoff $\bar{u}_i$, it does not matter whether we put all of the probability on s_i^k, put all of the probability on s_i^r, or split the probability by putting $\alpha \in [0, 1]$ on s_i^k and $(1-\alpha)$ on s_i^r. [In this third case, our expected payoff is still $\alpha u_i^k + (1-\alpha) u_i^r = \alpha \bar{u}_i + (1-\alpha)\bar{u}_i = \bar{u}_i$.] In this last case we see that the best-response mixed-strategy correspondence for this deleted mixed-strategy profile contains a continuum of mixed strategies corresponding to all the possible mixtures over the best pure strategies.

So we see that any mixed strategy which allocates probability only to best-response pure strategies is a best-response mixed strategy and vice versa.[41] We can express this conclusion in the following theorem:

Theorem The player-i mixed strategy $\sigma_i^* \in \Sigma_i$ is a best-response for player i to the deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ if and only if

$$\operatorname{supp} \sigma_i^* \subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i}). \tag{26}$$

Sketch of Proof To prove $\mathcal{P} \Leftrightarrow \mathcal{Q}$ (i.e. "$\mathcal{P}$ if and only if $\mathcal{Q}$"), where $\mathcal{P}$ and $\mathcal{Q}$ are propositions, we must prove both A $\mathcal{P} \Rightarrow \mathcal{Q}$ (i.e. "$\mathcal{P}$ only if $\mathcal{Q}$") and B $\mathcal{Q} \Rightarrow \mathcal{P}$ (i.e. "$\mathcal{P}$ if $\mathcal{Q}$").[42]

A: (σ_i^* is a best response to $\boldsymbol{\sigma}_{-i}$) $\Rightarrow$ $\operatorname{supp} \sigma_i^* \subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$.

The conditional proposition $\mathcal{P} \Rightarrow \mathcal{Q}$ is equivalent to $\sim\mathcal{Q} \Rightarrow \sim\mathcal{P}$.[43] Therefore to prove A we assume that $\operatorname{supp} \sigma_i^* \not\subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$ and try to deduce that σ_i^* is not a best response to $\boldsymbol{\sigma}_{-i}$. The fact that $\operatorname{supp} \sigma_i^* \not\subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$ implies that $\exists s_i' \in S_i \backslash \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$ such that $\sigma_i^*(s_i') > 0$. We show that σ_i^* is not a best response to $\boldsymbol{\sigma}_{-i}$ by exhibiting a mixed strategy $\sigma_i^\dagger \in \Sigma_i$ such that $u_i(\langle \sigma_i^\dagger, \boldsymbol{\sigma}_{-i} \rangle_i) > u_i(\langle \sigma_i^*, \boldsymbol{\sigma}_{-i} \rangle_i)$. To construct the better mixed strategy $\sigma_i^\dagger$, we arbitrarily pick some best-response pure strategy $\tilde{s}_i \in \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$ on which to shift all the probability which the original mixed strategy σ_i^* bestowed upon the non–best-response strategy s_i'. Formally… define $\sigma_i^\dagger \colon S_i \to [0, 1]$ for all $s_i \in S_i$ by

$$\boldsymbol{\sigma}_i^\dagger(s_i) = \begin{cases} 0, & s_i = s_i', \\ \boldsymbol{\sigma}_i^*(\tilde{s}_i) + \boldsymbol{\sigma}_i^*(s_i'), & s_i = \tilde{s}_i, \\ \boldsymbol{\sigma}_i^*(s_i), & s_i \notin \{s_i', \tilde{s}_i\}. \end{cases}$$

I leave it as an exercise for you to show that indeed $u_i(\langle \sigma_i^\dagger, \boldsymbol{\sigma}_{-i} \rangle_i) > u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i)$.

B: $\operatorname{supp} \sigma_i^* \subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i}) \Rightarrow$ (σ_i^* is a best response to $\boldsymbol{\sigma}_{-i}$).

We first observe that every player-i pure-strategy best response yields player i the same expected utility; i.e. $\forall s_i', s_i'' \in \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$, $u_i(\langle s_i', \boldsymbol{\sigma}_{-i} \rangle_i) = u_i(\langle s_i'', \boldsymbol{\sigma}_{-i} \rangle_i)$.[44] Denote by $\bar{u}_i \in \mathbb{R}$, this common expected utility; i.e. $\forall s_i \in \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$, $u_i(\langle s_i, \boldsymbol{\sigma}_{-i} \rangle_i) = \bar{u}_i$.

Now we show that $\bar{u}_i$ is an upper bound on the utility which any mixed strategy can achieve. To see this we refer to (22) which shows that the payoff to player i from any mixed strategy $\sigma_i \in \Sigma_i$ against $\boldsymbol{\sigma}_{-i}$ is a convex combination of the payoffs to the pure strategies in the support of σ_i. A convex combination

41 We have already argued that a best-response pure-strategy exists. The degenerate mixed strategy which puts unit weight on any such pure-strategy best response exists and is a mixed-strategy best response. Therefore a mixed-strategy best response exists.

42 A *proposition* is a statement which is either true or false.

43 For any proposition $\mathcal{P}$, we denote by $\sim\mathcal{P}$ the *negation* of $\mathcal{P}$. $\sim\mathcal{P}$ is also a proposition. Its truth value is the opposite of the truth value of $\mathcal{P}$.

44 If one yielded a strictly higher expected utility, the other would not be a best response.

of a set of real numbers must be weakly less than the maximum of that set.[45] Therefore any mixed strategy σ_i must yield an expected utility such that $u_i(\langle \sigma_i, \boldsymbol{\sigma}_{-i} \rangle_i) \le \bar{u}_i$.

Any mixed strategy which yields an expected utility of $\bar{u}_i$ must be a best response. (If it were not, there would exist another mixed strategy which yielded a higher utility, but this would contradict that $\bar{u}_i$ is an upper bound.)

Now use $\mathsf{supp}\, \sigma_i{}^* \subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$ and (22), to show that $u_i(\langle \sigma_i{}^*, \boldsymbol{\sigma}_{-i} \rangle_i) = \bar{u}_i$. Therefore $\sigma_i{}^*$ is a best-response mixed strategy. ☺[46]

The expected payoff to i from playing such a best-response mixed strategy is exactly the expected payoff she would receive from playing any one of her best pure strategies. Therefore a player never strictly prefers to mix rather than to play one of her best pure strategies against a particular profile of opponents' strategies.

For a given deleted profile of opponents' mixed strategies $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$, we now know which mixed strategies are best responses given the pure-strategy best-responses $\mathsf{BR}_i(\boldsymbol{\sigma}_{-i})$. This gives us an alternative and often useful way to graphically represent the players' best-responses, viz. in terms of the mixing probabilities which are optimal given the opponents' mixed strategies.

We can now define player i's mixed-strategy best-response correspondence $\mathsf{MBR}_i\colon \Sigma_{-i} \twoheadrightarrow \Sigma_i$, which specifies, for any deleted mixed-strategy profile $\boldsymbol{\sigma}_{-i} \in \Sigma_{-i}$ by i's opponents, a set $\mathsf{MBR}_i(\boldsymbol{\sigma}_{-i}) \subset \Sigma_i$ of player-i mixed strategies which are best responses to $\boldsymbol{\sigma}_{-i}$. This definition follows directly from (26):

$$\mathsf{MBR}_i(\boldsymbol{\sigma}_{-i}) = \{\sigma_i \in \Sigma_i\colon \mathsf{supp}\, \sigma_i \subset \mathsf{BR}_i(\boldsymbol{\sigma}_{-i})\}. \tag{27}$$

Example:

We previously determined Robin's and Cleever's (pure-strategy) best-response correspondences BR_R and BR_C. We can use each player's pure-strategy best-response correspondence to express the corresponding mixed-strategy best response correspondence, viz. MBR_R and MBR_C, respectively.

Every mixed-strategy for Robin can be written as an element of the one-dimensional simplex $\{(t, 1-t)\colon t \in [0, 1]\}$, where we adopt the convention that the mixed strategy $(t, 1-t)$ corresponds to playing Up with probability t. Robin's mixed-strategy best-response correspondence is:

45 More formally, you can show the following: For some integer n, let $\{x_1, \ldots, x_n\}$ and $\{\alpha_1, \ldots, \alpha_n\}$ be sets of real numbers such that $(\alpha_1, \ldots, \alpha_n) \in \Delta^{n-1}$; i.e. the $\{\alpha_j\}$ are convex coefficients. Then $\alpha_1 x_1 + \cdots + \alpha_n x_n \le \max\{x_1, \ldots, x_n\}$.

46 This smiley-face symbol indicates the end of the proof, as in Aumann and Sorin [1989]Aumann and Sorin [1989: 14] .

$$\text{MBR}_R(p,q)=\begin{cases}\{(1,0)\}, & q>1-\frac{10}{9}p,\\ \{(t,1-t): t\in[0,1]\}, & q=1-\frac{10}{9}p,\\ \{(0,1)\}, & q<1-\frac{10}{9}p,\end{cases}$$

with the understanding that $(p,q)\in\mathbb{R}_+^2$ and $p+q\leq 1$.

Every mixed strategy for Cleever can be written as an ordered triple belonging to the two-dimensional simplex $\{(p,q,1-p-q): p,q\geq 0, p+q\leq 1\}$, where we adopt the convention that the mixed strategy $(p,q,1-p-q)$ corresponds to playing left and middle with probabilities p and q, respectively. Cleever's mixed-strategy best-response correspondence is:

$$\text{MBR}_C(t)=\begin{cases}\{(0,0,1)\}, & t\in\left[0,1/6\right),\\ \{(p,0,1-p): p\in[0,1]\}, & t=1/6,\\ \{(1,0,0)\}, & t\in\left(1/6,1\right].\end{cases}$$

Example:

Consider the two-player game of Figure 9. Each player has two pure strategies, so each player's mixed strategy can be described by a single number on the unit interval. I'll assign p and q to Row for Up and to Column for left, respectively.

	l: $[q]$	r: $[1-q]$
U: $[p]$	**2**, −1	−1, **2**
D: $[1-p]$	−3, **2**	**0**, −1

Figure 9: A simple two-player game.

We first compute Row's mixed-strategy best-response correspondence $p^*: [0,1]\twoheadrightarrow[0,1]$, where $p^*(q)$ returns the set of all optimal mixing probabilities of playing Up for a given probability q with which Column plays left. To do this we compute Row's expected payoff to each of her pure strategies as a function of Column's mixed strategy q:

$$u_R(U;q)=2q-(1-q)=3q-1,$$

$$u_R(D;q)=-3q.$$

Comparing these two pure-strategy payoffs, we see that Row strictly prefers Up when $q>{}^1\!/_6$, is indifferent between Up and Down when $q={}^1\!/_6$, and strictly prefers Down when $q<{}^1\!/_6$. Row's pure-strategy best-response correspondence is

$$\mathrm{BR}_{\mathrm{R}}(q)=\begin{cases}\{U\}, & q\in\left[0,1/6\right),\\ \{U,D\}, & q=1/6,\\ \{D\}, & q\in\left(1/6,1\right].\end{cases}$$

We can therefore write Row's mixed-strategy best-response correspondence as

$$\mathrm{MBR}_{\mathrm{R}}(q)=p^*(q)=\begin{cases}\{0\}, & q\in\left[0,1/6\right),\\ [0,1], & q=1/6,\\ \{1\}, & q\in\left(1/6,1\right].\end{cases}$$

The graph of this correspondence appears as the solid heavy line in Figure 10.

We similarly compute Column's mixed-strategy best-response correspondence $q^*(p)$, which returns the probabilities of choosing left which are optimal given the probability p with which Row chooses Up. Computing Column's pure-strategy expected payoffs against Row's mixed-strategy p we obtain

$$u_{\mathrm{C}}(l;p)=-p+2(1-p)=2-3p,$$

$$u_C(r;p)=2p-(1-p)=-1+3p.$$

Column strictly prefers left if $p<\frac{1}{2}$, is indifferent between left and right if $p=\frac{1}{2}$, and strictly prefers right if $p>\frac{1}{2}$. Column's pure-strategy best-response correspondence is

$$\mathrm{MBR}_{\mathrm{C}}(q)=q^*(p)=\begin{cases}\{U\}, & p\in\left[0,1/2\right),\\ \{U,D\}, & p=1/2,\\ \{D\}, & p\in\left(1/2,1\right].\end{cases}$$

We can write Column's mixed-strategy best-response correspondence as

$$q^*(p)=\begin{cases}\{1\}, & p\in\left[0,1/2\right),\\ [0,1], & p=1/2,\\ \{0\}, & p\in\left(1/2,1\right].\end{cases}$$

The graph of this correspondence appears as the heavy shaded line in Figure 10.

A foreshadowing: We will see later that the intersection of the graphs of these correspondences is the *Nash equilibrium* of the game.

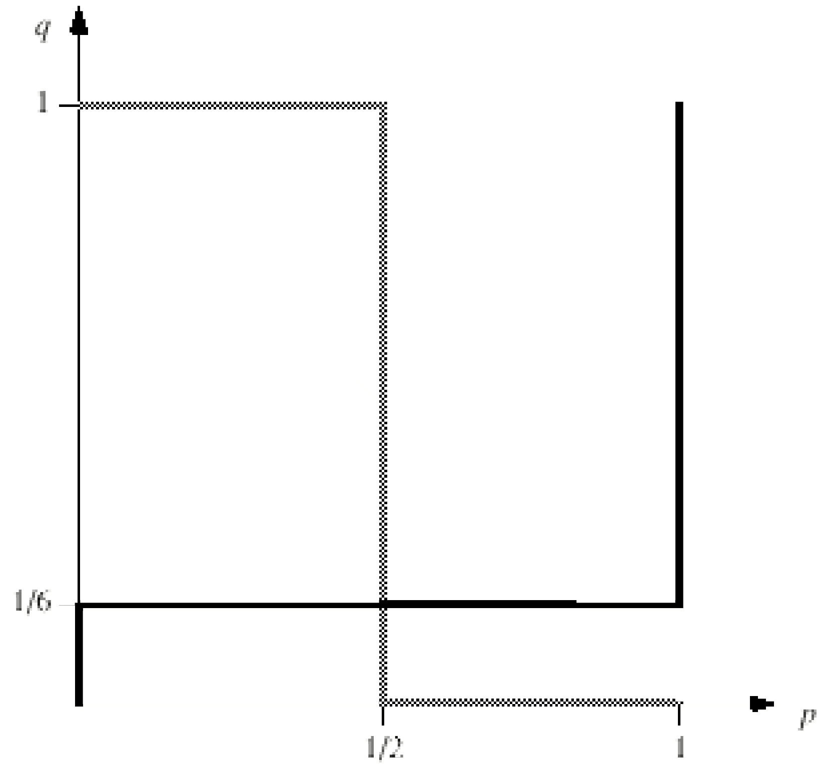

Figure 10: The mixed-strategy best-response correspondences for the game in Figure 9.

Summary

A strategic-form game is determined by a set of players, a pure-strategy space for each player, and a von Neumann-Morgenstern utility function for each player, the arguments of which are the pure strategies chosen by all the players. Such games can be conveniently represented by a matrix which includes a side for each player and whose cells contain payoff vectors.

Rational players choose actions which maximize their expected utility given their beliefs about the actions of their opponents. We set out to solve the Easy Part of Game Theory, viz. the problem of what choice a rational player would make given her beliefs about the choices of her opponents. We saw how to determine each player's pure-strategy best response(s) to the pure-strategy choices of her opponents.

We then expanded the set of choices available to the players by introducing mixed strategies, which are probability distributions over pure strategies. We saw that mixed strategies are members of a unit simplex and that pure strategies are degenerate mixed strategies. We computed players' expected payoffs to arbitrary mixed-strategy profiles by weighting the payoffs to pure-strategy profiles by the probability that each pure-strategy profile would be realized by the players' independent randomizations.

We determined each player's pure-strategy best response(s) against arbitrary mixed strategies by the opponent. We saw that, for a given deleted profile of other players' strategies, a nondegenerate mixed strategy can be a best response only when there are at least two pure-strategy best responses and that a mixed strategy is a best response if and only if it puts positive probability only upon best-response pure strategies. We then saw how the pure-strategy best-response correspondences yield the mixed-strategy best-response correspondences.

References

Aumann, Robert J. [1987] "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* **55** 1 (January): 1–18.

Aumann, Robert J. and Sylvain Sorin [1989] "Cooperation and Bounded Recall," *Games and Economic Behavior* **1** 1 (March): 5–39.