Games with Procedurally Rational Players

By Martin J. Osborne and Ariel Rubinstein*

We study interactive situations in which players are boundedly rational. Each player, rather than optimizing given a belief about the other players' behavior, as in the theory of Nash equilibrium, uses the following choice procedure. She first associates one consequence with each of her actions by sampling (literally or virtually) each of her actions once. Then she chooses the action that has the best consequence. We define a notion of equilibrium for such situations and study its properties. (*JEL* C72)

Economists' interest in game theory was prompted by dissatisfaction with the assumption underlying the notion of competitive equilibrium that each economic agent ignores other agents' actions when making choices. Game theory analyzes the interaction of agents who "think strategically," making their decisions rationally after forming beliefs about their opponents' moves, beliefs that are based on an analysis of the opponents' interests.

While game theory departs in this way from the theory of competitive equilibrium, the two theories share some basic assumptions. First, each player in a game, like each agent in a competitive market, is a fully rational decision maker: she has well-defined preferences, and chooses her best strategy according to these preferences, given her environment (defined by the prices in a competitive equilibrium and the other players' strategies in a game). Second, a player in a game and an agent in a competitive market comprehends the situation she faces: she understands how the outcome depends on her action and her environment. Third, the interpretation of many game-theoretic solution concepts (Nash equilibrium included) requires that, as in the theory of competitive equilibrium, when choosing an action each player knows the relevant environmental parameters. A player does not deduce the other players' actions from a theory of how these players behave, but is assumed, in equilibrium, simply to know these actions.

Unease with these assumptions shared by game theory and the theory of competitive equilibrium has arisen both because the assumptions are inconsistent with evidence about human decision-making (see, for example, Colin Camerer's 1995 survey) and because models that embody them appear to be incapable of fully explaining phenomena like advertising, incomplete contracts, the consulting industry, and the variety of performances of agents with the same preferences and information, a variety that appears to be due to differences in cognitive skills. This unease has led some researchers, notably Herbert A. Simon (see for example, 1955, 1982), to construct models in which the players are "boundedly rational."

In this paper we join game theory with a model of boundedly rational decision-making. We construct a static model of an interacting set of players, each of whom uses a decision-making procedure that departs systematically from the tenets of "rationality." We assume, as in standard game-

* Osborne: Department of Economics, McMaster University, Hamilton, L8S 4M4, Canada (e-mail: Osborne@McMaster.ca; http://www.socsci.mcmaster.ca/~econ/faculty/osborne/); Rubinstein: School of Economics, Tel Aviv University, Tel Aviv, 69978 Israel (e-mail: rariel@ccsg.tau.ac.il; http://www.princeton.edu/~ariel). We thank Carolyn Pitchik, Phil Reny, Al Roth, and two anonymous referees for helpful comments. Osborne gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada. Part of this research was conducted while Rubinstein was a visiting scholar at the Russell Sage Foundation, New York.

theoretic models, that a player makes choices on the basis of an association between her actions and outcomes. In the standard models a rational player forms such an association from a belief about the other players' actions and her understanding of the consequence of each action profile. We assume, by contrast, that a player constructs an action-consequence correspondence directly, in a manner we describe below. Each player either does not know the relationship between the other players' actions, her own action, and the outcome, or does not appreciate its significance. We consider the characteristics of an equilibrium in which each player's association of consequences with actions is consistent with all the players' behavior. Two examples illustrate the idea.

You are new to town and are planning your route to work. How do you decide which road to take? You know that other people use the roads, but have no idea which road is most congested. One procedure is to try each route once (or several times) and then permanently adopt the one that was (on average) best. The outcome of this procedure is stochastic: you may sample the route that is, in fact, the best on a day when a baseball game congests it. Once you select your route, you become part of the environment that determines other drivers' choices.

You have to bargain and must choose whether to be tough or soft. Your decision depends on the consequence you perceive for each strategy. You may base your perception on your experience in similar situations. This experience may have taught you, for example, to associate victory with toughness and defeat with softness, or the reverse, or it may have taught you that the outcome is independent of your bargaining stance. Our approach assumes that your perceptions reflect the equilibrium behavior of the various potential bargainers, behavior that may lead other bargainers, who have had different experiences, to different conclusions.

The simplest version of our model applies to a situation in which there are two players, each player's set of actions is A, and each player's payoff is $u(x, y)$ when she chooses $x \in A$ and her opponent chooses $y \in A$. The solution we study is a probability distribution α^* on A with the property that $\alpha^*(x)$ is the probability that a player finds the action x to be the best (breaking ties equiprobably) when she samples each action once and her opponent's action is determined by α^*. We refer to such a probability distribution α^* as an $S(1)$-equilibrium. (Later in the paper we consider asymmetric two-player games and extensive games, as well as procedures in which each player samples each action many times, rather than only once.)

Example 1: Suppose that each player has two actions, a and b, and payoffs as given in the following figure. A row corresponds to a player's action, and a column to one of her opponent's actions. For example, if a player chooses the action a and her opponent chooses the action b then the player's payoff $u(a, b)$ is 4.

	a	b
a	2	4
b	3	1

Denote by $\alpha(x)$ the probability that a player's action is x, for $x = a, b$. When she samples a, a player obtains the payoff 2 with probability $\alpha(a)$ and the payoff 4 with probability $\alpha(b)$; when she samples b, she obtains the payoff 3 with probability $\alpha(a)$ and the payoff 1 with probability $\alpha(b)$. Thus her payoff when she samples a exceeds her payoff when she samples b if and only if either her opponent's action is b when she samples a (in which case a is better than b regardless of her opponent's action when she samples b), or her opponent's action is a when she samples a and b when she samples b. Thus after she samples each of her actions once, the probability that the outcome she associates with a is superior to the outcome she associates with b is $\alpha(b) + \alpha(a)\alpha(b)$. In an $S(1)$-equilibrium this probability is equal to the probability $\alpha(a)$ that a player chooses a. Thus in an $S(1)$-equilibrium $\alpha(a) = \alpha(b) +$

$\alpha(a)\alpha(b) = 1 - \alpha(a) + \alpha(a)(1 - \alpha(a))$, so that $\alpha(a) = \frac{1}{2}(\sqrt{5} - 1) \approx 0.62$.

Example 2: If the payoffs are given in the following figure

	a	b
a	4	1
b	3	2

then the probability that the outcome a player associates with a is better than that she associates with b is $\alpha(a)$, the probability that when she samples a her opponent chooses a. Thus every distribution α is an $S(1)$-equilibrium.

In what circumstances might a player choose her action using the procedure we assume? Suppose that her information about the structure of the interaction in which she is involved is poor: she knows her set of actions (the set of routes she could take, in the route-planning example), but little else. The standard approach calls for her to form beliefs about the various uncertainties she faces (her payoffs, the number of other players, the other players' actions, etc.). Under conditions of poor information, it seems more natural for her to act on the basis of a direct link between actions and consequences, a link she may construct by trying each action once, for example. Indeed, if she does not even know how her own payoffs depend on all the players' actions, then sampling her actions seems her only sensible alternative. If her information is better, she may still be attracted to such a procedure—she may, for example, fail to recognize the connection between her action, her opponents' actions, and an outcome of the game— in which case she exhibits "bounded rationality."

We suggest two interpretations for an $S(1)$-equilibrium.

Experimentation.—There is a large population of individuals, pairs of whom are occasionally matched and interact. When entering the population, a player chooses her action after sampling each alternative once, picking the action that yields the highest payoff. An equilibrium corresponds to a steady state in which the probability that a new player chooses any given action is equal to the fraction of the population that currently chooses that action. Although the model is static, under this interpretation the equilibrium we study is defined by the property that entry into the population has no effect on the distribution of actions.

Virtual experimentation.—A player tries to connect actions with consequences. She does not know her opponent's behavior or the consequence of any pair of actions, but uses some mental process to construct an association of consequences with actions. Possible inputs into this process include, for example, the player's experience in similar games, her observations of other players, and information learned in school. The process may yield random results, but is not arbitrary. For example, a player does not conclude that the action a yields the consequence c if no player in such a situation ever takes an action b for which $f(a, b) = c$, where $f(a, b)$ is the consequence of the action pair (a, b). We assume, in fact, that the process yields an action-consequence relationship that is correct, given the other player's behavior. (This assumption is analogous to the assumption in the theory of Nash equilibrium that a player's belief about the other player's mixed strategy is correct.) Precisely, a player associates the consequence c with the action a with probability $\pi_a(c\,|\,\alpha^*) = \Sigma_{\{b: f(a,b) = c\}}\, \alpha^*(b)$ when the other player's action is determined by the probability distribution α^*. Under this interpretation, an equilibrium is a probability distribution α^* on A such that for each action a, $\alpha^*(a)$ is the probability that a yields the highest expected payoff when the probability that a player associates the consequence c with a is $\pi_a(c\,|\,\alpha^*)$.

A variant of the first interpretation is that a player obtains information about the consequences of her actions by observing other players' experience, rather than by trying actions herself. A difficulty with this variant is that some actions may not be taken by any player. Nevertheless, the variant is appropriate for cases in which every action is taken by some player.

Our approach contrasts with the conventional game-theoretic notion of Nash equilibrium. According to the notion of Nash equilibrium, each player constructs a probabilistic belief about the consequences of her actions by first constructing a probabilistic belief about the other player's behavior and then choosing the action that yields the best distribution of consequences according to her preferences. The Nash equilibrium requirement is that the belief be consistent with the other player's behavior.

We retain the assumption that each player constructs a probabilistic belief about the consequences of her actions and takes an action that is best given these beliefs. We also retain the equilibrium principle: the beliefs are consistent with the other player's behavior. But our analysis differs in a key respect from that of Nash equilibrium: each player associates a consequence with each action *independently.* She does not derive her action from a single conjecture about the other player's behavior, as in the theory of Nash equilibrium. To appreciate this point, note that our solution differs from one in which a player, in order to determine her action, selects a single representative action for her opponent, on the basis of which she compares the outcome of her own actions. Such a procedure, unlike the one we study, requires that a player recognize the connection between her action, her opponent's action, and an outcome of the game.

The boundedly rational decision-making procedure embodied in the notion of $S(1)$-equilibrium is one of many possible procedures. Why does our model have appeal?

Interpretability.—Our solution concept has natural interpretations. In the two interpretations we have presented, it requires each player to know much less about the interaction than does Nash equilibrium: each player must know only her own set of actions.

Tractability.—The solution concept may be applied to a variety of games. Although the calculation of the $S(1)$-equilibria of a two-player symmetric game involves solving a system of polynomial inequalities of degree k (where k is the number of actions), for all the examples we present, the equilibria may be easily analyzed.

Gradations of reasoning.—A natural generalization of the decision-making procedure used by players in an $S(1)$-equilibrium is that in which each player samples each of her actions K times, rather than only once. The larger the value of K, the more detailed the players' reasoning processes. Thus for any given game we can examine how the outcome depends on the players' sophistication.

Discrimination.—In interesting examples, the set of $S(1)$-equilibria is small; in some cases it is smaller than the set of Nash equilibria.

Two related lines of research also span game theory and bounded rationality. The literature on games and automata (especially that concerning repeated games) models the idea that more complex strategies are more costly to use. Players behave rationally, but explicitly consider the complexity costs when choosing a strategy. (This literature is surveyed in Rubinstein, 1998 Chs. 8 and 9.) In the literature on evolutionary game theory, players do not explicitly deliberate; rather, each player's action is determined by an automatic behavioral rule that may change over time as a result of evolutionary forces. The theory explores the implications of such automatic behavior and relates the outcomes to those that emerge under conventional game-theoretic solution concepts.

In Section I we describe the model in detail and provide a handful of applications to symmetric two-player games. We extend the model to asymmetric games in Section II, and discuss the appropriateness of applying the solution concept to extensive games in Section III. We study the $S(K)$-equilibrium concept and compare its limit (as K goes to infinity) with Nash equilibrium in Section IV. In Section V we define a general solution concept we call *procedural equilibrium* that encompasses the notion of $S(1)$-equilibrium and other solution concepts that have been proposed in work on game theory and bounded rationality (in particular those of Robert W. Rosenthal [1989] and Hsiao-Chi Chen et al. [1997]).

I. Symmetric Two-Player Games

Each of two players has a finite set A of actions and obtains the payoff $u(a, b)$ when she chooses the action a and the other player chooses b. We interpret the payoff function u as a representation of each player's ordinal preferences over the set $A \times A$ of outcomes.

Let α be a probability distribution on A. For every $a \in A$, let $v(a, \alpha)$ be the random variable that yields the number $u(a, b)$ with probability $\alpha(b)$ for each $b \in A$. Denote by $w(a, \alpha)$ the probability that the action a yields the highest payoff (a is the "winner") when each random variable $v(x, \alpha)$ is drawn once, independently. That is, $w(a, \alpha)$ is the probability that $v(a, \alpha) > v(x, \alpha)$ for all x plus the weighted sum of the probabilities that a is tied for the highest realized payoff, with weights equal to the reciprocal of the numbers of tied alternatives.

Definition 1: An $S(1)$-*equilibrium* is a probability distribution α on the set A of actions with the property that

$$w(a, \alpha) = \alpha(a) \text{ for every action } a.$$

That is, an $S(1)$-equilibrium has the following (fixed-point) property. Suppose that for any actions a and b, the probability with which a player associates the payoff $u(a, b)$ with her action a is the probability that the other player chooses b, namely $\alpha(b)$. Then α is an $S(1)$-equilibrium if for each action a the probability that, given this association, a yields the highest payoff, is precisely $\alpha(a)$.

Note that an $S(1)$-equilibrium requires that the model specify each player's ranking of all outcomes, whereas the notion of pure Nash equilibrium requires only each player's ranking of the outcomes for each fixed action of her opponent. At the same time, an $S(1)$-equilibrium does not require any information about the players' preferences between lotteries, whereas the notion of a mixed strategy Nash equilibrium does. In particular, if, for some action y, we add the same number to $u(x, y)$ for every action x, then the sets of pure and mixed Nash equilibria remain the same, while the set of $S(1)$-equilibria may change. On the other hand, any change in the payoffs that maintains their order has no effect on the sets of pure Nash equilibria and $S(1)$-equilibria, but may change the set of mixed strategy Nash equilibria.

It is immediate that α is an $S(1)$-equilibrium in which $\alpha(a) = 1$ for some $a \in A$ if and only if (α, α) is a strict pure strategy Nash equilibrium.

PROPOSITION 1: *Every symmetric game has an* $S(1)$-*equilibrium.*

PROOF:

Define the function H from the simplex in $\mathbb{R}^A$ to itself to assign to α the vector defined by $H(\alpha)(a) = w(a, \alpha)$ for each $a \in A$. This function is continuous and hence by Brouwer's fixed-point theorem has a fixed point, which is an $S(1)$-equilibrium.

Can a dominated action be used with positive probability in an $S(1)$-equilibrium? If a dominates b then b may still do better than a for some realizations of the other player's action, so one cannot conclude immediately that a dominated action is not used.

Example 3: Consider games of the form

	a	b
a	1	x
b	0	x

in which the action a weakly dominates the action b and a player's actions tie if the other player chooses the *dominated* action. (There are five such games, differing in the ranking of x relative to 0 and 1.) The action b ties with a if the outcome (a, b) is associated with a and the outcome (b, b) is associated with b. The only other cases in which b is a winner are those in which either $x \geq 1$, $v(a, \alpha) = 1$, and $v(b, \alpha) = x$, or $x \leq 0$, $v(a, \alpha) = x$, and $v(b, \alpha) = 0$. Thus in an $S(1)$-equilibrium we have $1 - \alpha(a) \leq \frac{1}{2}(1 - \alpha(a))^2 + \alpha(a)(1 - \alpha(a))$, and hence $\alpha(a) = 1$: the only $S(1)$-equilibrium assigns probability one to a.

Example 4: Consider games of the form

	a	b
a	x	1
b	x	0

in which the action a weakly dominates the action b and a player's actions tie if the other player chooses the *dominating* action. (Again, there are five such games.) The unique symmetric Nash equilibrium action a is not an $S(1)$-equilibrium. Denoting $p = \alpha(a)$, the $S(1)$-equilibria are:

- If $x > 1$ or $x < 0$ then $p = \frac{1}{2}p^2 + (1-p)^2 + (1-p)p$, so that $p = 2 - \sqrt{2} \approx 0.59$.
- If $x = 0$ or $x = 1$ then $p = \frac{1}{2}p + 1 - p$, so that $p = \frac{2}{3}$.
- If $0 < x < 1$ then $p = 1 - p + (1-p)p + \frac{1}{2}p^2$, so that $p = \sqrt{3} - 1 \approx 0.73$.

While a weakly dominated action may thus be used with positive probability in an $S(1)$-equilibrium, a *strictly* dominated action is never used in an $S(1)$-equilibrium of a *two*-action game by the following argument. Let the action a strictly dominate the action b. The number $u(b, y)$ is greater than or equal to $u(a, x)$ in at most one of the two cases in which $x \neq y$. Thus for equilibrium we need $1 - \alpha(a) \leq \alpha(a)(1 - \alpha(a))$, and hence $\alpha(a) = 1$.

In games with more than two actions, however, strictly dominated actions may appear with positive probability in $S(1)$-equilibria.

Example 5 (Voluntary exchange): Each of two people has two books. Each person attaches the value 1 to each of her books and the value 3 to each of the other person's books. Exchange is voluntary: each person decides how many books (0, 1, or 2) to give to her partner, without any commitment to receive anything in return. Thus the payoff function is given by $u(a, b) = 3b + 2 - a$, yielding the following game.

	0	1	2
0	2	5	8
1	1	4	7
2	0	3	6

An $S(1)$-equilibrium α satisfies

$$\alpha(0) = \alpha(0)^3 + \alpha(1)(1 - \alpha(2))^2 + \alpha(2)$$

$$\alpha(1) = \alpha(0)\alpha(1)(1 - \alpha(2)) + \alpha(2)(1 - \alpha(2))$$

$$\alpha(2) = \alpha(0)^2\alpha(1) + \alpha(2)(1 - \alpha(2))^2.$$

One solution is the no-exchange equilibrium $(\alpha(0), \alpha(1), \alpha(2)) = (1, 0, 0)$; another solution is $(\alpha(0), \alpha(1), \alpha(2)) \approx (0.52, 0.28, 0.2)$, in which both dominated actions are used with positive probability. Thus in an $S(1)$-equilibrium players may give away goods without any commitment to receive goods in return.

Although a weakly dominated action may occur in an $S(1)$-equilibrium with positive probability, the following result shows that the probability assigned to any such action is less than the probability assigned to any action that dominates it.

PROPOSITION 2: *If the action a weakly dominates the action b then $\alpha(a) \geq \alpha(b)$ in any $S(1)$-equilibrium α.*

PROOF:

We need to show that for every probability distribution α on A we have $w(a, \alpha) \geq w(b, \alpha)$. Now, for any action y, the probability $w(y, \alpha)$ that y is the winner is the weighted sum of the probabilities that y is the winner over all possible associations x between the player's action y and the action x_y of the other player. Precisely, let $W(x)$ be the set of winning actions given the association x:

$$W(x) = \{z \in A : u(z, x_z) \geq u(y, x_y) \quad \text{for all } y\}.$$

Then

$$w(y, \alpha) = \sum_x \Pr(x|\alpha) \frac{\delta(y, W(x))}{|W(x)|},$$

where the sum ranges over all possible associations, $\Pr(x|\alpha)$ is the probability of the association x given the other player's choices are given by α, and $\delta(y, W(x))$ is 1 if $y \in W(x)$ and 0 otherwise. Now, let x' differ from x only in that $x'_a = x_b$ and $x'_b = x_a$. The mapping $x \to x'$ is one-to-one and onto, and $\Pr(x|\alpha) = \Pr(x'|\alpha)$. Since a dominates b, if $b \in W(x)$ then $a \in W(x')$ and $|W(x')| \leq |W(x)|$. Thus, $w(a, \alpha) \geq w(b, \alpha)$.

Example 6 (Duplicated actions): Duplicating actions affects the $S(1)$-equilibrium outcome. Consider the game

1	⋯	1	4
⋮	⋱	⋮	⋮
1	⋯	1	4
3	⋯	3	2

where the first action is repeated K times. Denoting by p the equilibrium probability assigned to the bottom action, we have $p = (1 - p)^K$. This equation has a unique solution since the left-hand side is increasing and the right-hand side is decreasing. Thus there is a unique $S(1)$-equilibrium, with a value of p converging to 0 as $K \to \infty$.

The following result gives a useful condition for an $S(1)$-equilibrium to assign probability zero to an action.

PROPOSITION 3: *If there are distinct actions a and b such that $u(a, x) > u(b, y)$ for all x and all $y \neq b$, and an action c with $c \neq b$ such that $u(c, b) \geq u(b, b)$, then in any $S(1)$-equilibrium α we have $\alpha(b) = 0$.*

PROOF:

The action b wins over a only if the outcome (b, b) is associated with b. Even then, if $\alpha(b) > 0$ the action b does not always win, since with positive probability the outcome (c, b) is associated with the action c. Thus, for equilibrium we need $\alpha(b) < \alpha(b)$ if $\alpha(b) > 0$. Hence $\alpha(b) = 0$.

Example 7 (Bertrand competition): Two sellers of a good compete for buyers, each of whose reservation values is 1. The cost of production is 0. The strategic variable is price. Each consumer buys one unit of the good from the cheaper seller, splitting her demand if the prices are equal. Denote by $p_1, \ldots, p_K$, with $0 < p_1 < p_2 < \cdots < p_K < 1$, the K possible prices, and by $u(p, p')$ the profit of a seller who charges the price p when her opponent charges the price p'. We have $u(p, p') = 0$ if $p > p'$ and $u(p, p') > 0$ if $p \leq p'$. Assume that the price gradations are fine enough that $u(p_k, p_k) < u(p_{k-1}, p_k)$ for each k.

Letting $a = p_1$, $b = p_K$, and $c = p_{K-1}$, we conclude from Proposition 3 that $\alpha(p_K) = 0$ in any $S(1)$-equilibrium. Continuing by induction we conclude that the only $S(1)$-equilibrium assigns probability one to the action p_1, the unique Nash equilibrium price.

Example 8 (Hotelling's location game): Two rivals compete for the hearts of a population of $2K + 1$ people; one person is located at each of the positions $1, \ldots, 2K + 1$ on a straight line segment. Each competitor chooses a location, one of the $2K + 1$ positions. Each member of the population patronizes the competitor closest to her; if the competitors are equally distant, she divides her patronage equally between the two. Thus, for example, $K = 2$ generates the following game.

	1	2	3	4	5
1	5	2	3	4	5
2	8	5	4	5	6
3	7	6	5	6	7
4	6	5	4	5	8
5	5	4	3	2	5

We claim that in any $S(1)$-equilibrium, $\alpha(1) = \alpha(2K + 1) = 0$. Assume without loss of generality that $\alpha(1) \geq \alpha(2K + 1)$. Action 1 is a winner only if in the outcome associated with it the other player chooses an extreme position, and in the outcome associated with the middle action both players choose the middle position. Hence $\alpha(1) \leq [\alpha(1) + \alpha(2K + 1)]\alpha(K + 1)/2 \leq \alpha(1)\alpha(K + 1)$, so that $\alpha(1) = 0$. An inductive argument implies that the only $S(1)$-equilibrium is the pure strategy that attaches probability one to the middle action, which is also the unique Nash equilibrium action.

II. Asymmetric Games

The definition of an $S(1)$-equilibrium may straightforwardly be extended to situations in which the players' roles are asymmetric. Let $\langle N, (A_i), (u_i)\rangle$, be a strategic game, where N is a finite set of players, A_i is the finite set of actions of player i, and u_i is a payoff function that represents player i's ordinal preferences over the set of outcomes. Let α be a profile of mixed strategies. For each $a_i \in A_i$ let $v_i(a_i, \alpha_{-i})$ be the random variable that, for each a_{-i}, yields the number $u_i(a)$ with probability $\Pi_{j \in N\setminus\{i\}}\alpha_j(a_j)$.

Denote by $w(a_i, \alpha_{-i})$ the probability that the action a_i yields the highest payoff when each of the random variables $v_i(x, \alpha_{-i})$ for $x \in A_i$ is drawn independently.

Definition 2: An $S(1)$-*equilibrium* is a profile α of mixed strategies with the property that

$$w(a_i, \alpha_{-i}) = \alpha_i(a_i)$$

for every player i and action $a_i \in A_i$.

That is, an $S(1)$-equilibrium is a profile α, where α_i is a probability distribution over player i's set of actions, with the following (fixed-point) property. Suppose that for any action profile a, the probability with which player i associates the payoff $u_i(a)$ with her action a_i is the probability that the actions chosen by the other players are given by a_{-i}, namely $\Pi_{j \in N\setminus\{i\}}\alpha_j(a_j)$. Then α is an $S(1)$-equilibrium if for every player i and each action a_i the probability is precisely $\alpha_i(a_i)$ that, given this association, a_i yields player i the highest payoff.

It is easy to verify that Propositions 1 and 2 hold for asymmetric games. The next example illustrates the notion of $S(1)$-equilibrium in an asymmetric game.

Example 9 (Choosing reservation values): A seller owns an indivisible good that she can trade to a buyer; the good is worth 4 to the buyer and 0 to the seller. The seller announces a price and the buyer announces the maximum price she is willing to pay, where each price is 1, 2, or 3; each player is committed to her price during a bargaining session. In the bargaining session the buyer has superior bargaining power; if the buyer's reservation price is at least the seller's price then the good is traded at the seller's price.

The game is as follows.

		Buyer		
		1	2	3
	1	1,3	1,3	1,3
Seller	2	0,0	2,2	2,2
	3	0,0	0,0	3,1

Every pair (a, a) is a Nash equilibrium. However, the unique $S(1)$-equilibrium is (3, 3). To see this, let (α_1, α_2) be an $S(1)$-equilibrium. We first argue that $\alpha_1(1) = 0$ (i.e., the seller announces the price 1 with probability zero). Suppose to the contrary that $\alpha_1(1) > 0$. A necessary condition for the buyer's action 1 to be a winner is that in the outcome the buyer associates with it, the seller chooses 1, an event with probability $\alpha_1(1)$. This condition is not sufficient, since if the buyer associates the outcome in which the seller chooses 1 with either or both of her actions 2 and 3 then the buyer's action 1 ties with these actions. Thus $\alpha_2(1) < \alpha_1(1)$. We need also $\alpha_1(1) \leq \alpha_2(1)$, since a necessary condition for the seller's action 1 to be a winner is that in the outcome the seller associates with her action 2, the buyer chooses 1, an event with probability $\alpha_2(1)$. These two inequalities are inconsistent, so we conclude that $\alpha_1(1) = 0$. It follows that $\alpha_2(1) = 0$, since the buyer's action 1 is never a winner when $\alpha_1(1) = 0$. By induction we conclude that in the only $S(1)$-equilibrium each player chooses the price 3. (This argument extends straightforwardly to the case in which there is any finite number of possible prices, rather than only three.)

III. Extensive Games

The concept of an $S(1)$-equilibrium can be applied to the reduced strategic form of an extensive game. We present two illustrative examples and discuss the appropriateness of the solution concept in this case.

Example 10:

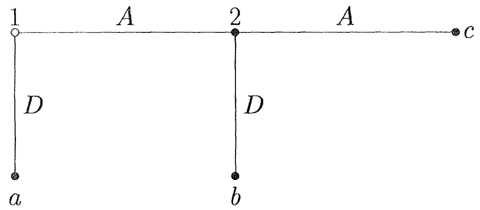

Assume that $c \succ_1 a$ and $c \succ_2 b$, so that the unique subgame-perfect equilibrium yields the outcome c. Denote by p the probability that player 1 chooses D and by q the probability that player 2 chooses D. The action D wins for player 1 only if player 1 associates the outcome b with her action A, so that $p \leq q$. The action D is a winner for player 2 only if either she associates the outcome a with both her actions, in which case the actions tie, or she associates a with D and c with A (if $a \succsim_2 c$) or b with D and a with A (if $b \succsim_2 a$). Since $c \succ_2 b$, we do not have both $a \succsim_2 c$ and $b \succsim_2 a$, so the probability q that D is a winner is at most $\frac{1}{2}p^2 + p(1-p)$. Since $p \leq q$ we conclude that $p = q = 0$. Thus every such game has a unique $S(1)$-equilibrium, which coincides with the unique subgame-perfect equilibrium. (Note that if $a \succsim_1 b$ then the game has a Nash equilibrium yielding the outcome a, different from the unique subgame-perfect equilibrium outcome.)

Example 11 (The centipede game): In Rosenthal's (1981) centipede game two players move alternately, starting with player 1 in period 1. In any period, the player whose turn it is to move may stop the game or pass the move to the other player. If neither player has stopped the game after each player has had T opportunities to do so, then the game ends. For each player, the outcome when she stops the game in period t is better than that in which the other player stops the game in period $t + 1$ (or the game ends), but worse than any outcome that is reached if in period $t + 1$ the other player passes the move to her. In particular, all the outcomes in period 3 and beyond are better for both players than the outcome in which player 1 stops the game immediately. Nevertheless, for any value of T the game has a unique Nash equilibrium, in which player 1 stops the game immediately.

As the number of periods goes to infinity in this game, the probability assigned by any $S(1)$-equilibrium to player 1's stopping the game immediately goes to 0. To see this, denote by p and q the $S(1)$-equilibrium probabilities that players 1 and 2 stop the game at their first opportunities. Player 1's reduced strategy of stopping immediately wins if and only if she associates with all her other reduced strategies the outcome that player 2 stops the game at her first opportunity, an event with probability q^T. For player 2's reduced strategy of stopping at her first opportunity to be a winner it is necessary that either she associates with this reduced strategy the outcome that she indeed stops at her first opportunity (namely, player 1 does not stop the game immediately) or she associates with all reduced strategies the outcome that player 1 stops the game immediately, in which case all her reduced strategies tie. Thus $q \leq 1 - p + p^{T+1}/(T+1)$. Any pair (p, q) satisfying the two relations $p = q^T$ and $q \leq 1 - p + p^{T+1}/(T+1)$ has the property that $p \to 0$ as $T \to \infty$.

This result has appeal as a potential resolution of the paradoxical aspects of the centipede game. To examine this resolution, consider the following case $T = 1$.

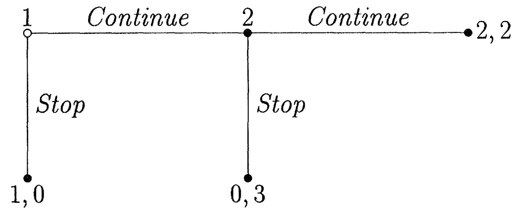

Player 2 finds the strategy *Continue* better than the strategy *Stop* when the scenario that she constructs for her strategy *Continue* is that player 1 chooses *Continue* as well, and the scenario that she constructs for her strategy *Stop* is that player 1 chooses *Stop* as well. Although she has to move only if player 1 does not stop the game, she is allowed to include in the scenario she attaches to each strategy the possibility that player 1 did stop the game. That is, the logic of an $S(1)$-equilibrium entails her comparing the strategies *Continue* and *Stop* without noticing that either of these strategies is executed only if player 1 chooses *Continue.* Thus an $S(1)$-equilibrium makes sense in the centipede game only if both players fail to understand the structure of the game, as captured by its extensive form.

IV. Limits of $S(K)$-Equilibria

Return, for simplicity, to the world of symmetric strategic games. Consider a modification of the notion of an $S(1)$-equilibrium in which each player associates with each action a distribution over the set of consequences, rather than a deterministic consequence. The distribution of consequences each player associates with each of her actions results from sampling the action K times. A player chooses an action by comparing these distributions. Thus if $K \geq 2$, her payoffs must have more than ordinal significance. Consequently we equip the players with vNM payoffs rather than ordinal payoffs as for the case $K = 1$.

An $S(K)$-equilibrium is a profile of mixed strategies, each of which is consistent with each player's choosing the action that yields the highest expected payoff given the distributions she associates with her actions.

More precisely, let (A, u) be a symmetric two-player game in which u is a vNM payoff function. Let α be a probability distribution on A. For every $a \in A$, let $v(a, \alpha, K)$ be the random variable equal to the sum of K independent random variables, each of which yields the number $u(a, b)$ with probability $\alpha(b)$ for each $b \in A$. The random variable $v(a, \alpha, K)$ is a player's total payoff from K plays of the action a when the other player's action is determined by the probability distribution α. Denote by $w(a, \alpha, K)$ the probability that the action a obtains the highest score, assuming that ties are broken by an equiprobability rule. That is, $w(a, \alpha, K)$ is the probability that $v(a, \alpha, K) > v(x, \alpha, K)$ for all x plus the weighted sum of the probabilities that x is tied for the highest score, with weights equal to the reciprocals of the numbers of tied alternatives.

Definition 3: For any positive integer K, an $S(K)$-*equilibrium* is a probability distribution α on the set A of actions with the property that

$$w(a, \alpha, K) = \alpha(a) \text{ for every action } a.$$

That is, an $S(K)$-equilibrium differs from an $S(1)$-equilibrium only in that each player's choice is based on K samples of each action rather than only one. Propositions 1 and 2 hold for $S(K)$-equilibrium.

The distributions over consequences that each player associates with her actions reflect more accurately the other player's behavior the larger is K. This suggests that there is a relationship between the limit of $S(K)$-equilibria as K increases and the mixed strategy Nash equilibria.

PROPOSITION 4: *Let α be the limit as $K \to \infty$ of a subsequence of a sequence of $S(K)$-equilibria. Then (α, α) is a mixed strategy Nash equilibrium.*

PROOF:

Let $\{\alpha^{K_n}\}$ (with $K_n \to \infty$) be a sequence of $S(K_n)$-equilibria converging to α. Suppose that (α, α) is not a mixed strategy Nash equilibrium. Then there is an action a in the support of α and an action b such that the expected payoff of a against α is less than the expected payoff of b against α. Thus, for any $\varepsilon > 0$ there exists n^* large enough that for all $n > n^*$, the probability that $v(a, \alpha, K_n) \geq v(b, \alpha, K_n)$ is less than ε. Since $\alpha^{K_n} \to \alpha$, for any $\varepsilon > 0$ there exists n^{**} large enough that for all $n > n^{**}$, the probability that $v(a, \alpha^{K_n}, K_n) \geq v(b, \alpha^{K_n}, K_n)$ is less than ε, contradicting the assumption that $\alpha(a) > 0$.

Since the set of mixed strategies is compact, every sequence of $S(K)$-equilibria has a subsequence that converges. Thus if a game has a unique mixed strategy equilibrium, then by Proposition 4 every limit is that equilibrium, so we have the following corollary.

COROLLARY: *In a game with a unique mixed strategy Nash equilibrium, the equilibrium mixed strategy is the unique limit of $S(K)$-equilibria as $K \to \infty$.*

In a game with multiple mixed strategy Nash equilibria some equilibria, even pure equilibria, may not be the limits of $S(K)$-equilibria. For example, in the degenerate game

1	1
1	1

the only $S(K)$-equilibrium, for any value of K, is $(^1/_2, {}^1/_2)$. A more interesting example follows.

Example 12:

	T	M	B
T	1	1	1
M	1	0	2
B	1	2	0

Every strategy α with $\alpha(M) = \alpha(B)$ is a symmetric mixed Nash equilibrium strategy. However, for large K, all $S(K)$-equilibria are close to $(^1/_4, {}^3/_8, {}^3/_8)$. To see this, first note that for every K, if α is an $S(K)$-equilibrium then $\alpha(M) = \alpha(B)$: if, for example, $\alpha(M) > \alpha(B)$, the probability B wins is higher than the probability M wins. Now, for every probability distribution (x, y, y) on $\{T, M, B\}$ with $x < 1$, the probability that T obtains a higher score than does M converges to $^1/_2$ as $K \to \infty$ (since the probability of a tie goes to zero) and the probability that T obtains a higher score than does B similarly converges to $^1/_2$. These two events are independent, so the probability that T obtains the highest score converges to $^1/_4$. Thus as $K \to \infty$ the $S(K)$-equilibria approach $(^1/_4, {}^3/_8, {}^3/_8)$.

The next result shows that a necessary condition for a pure strategy to be the limit of $S(K)$-equilibria is that it be a strict best response to some mixed strategy.

PROPOSITION 5: *Let α^* be a limit of $S(K)$-equilibria as $K \to \infty$, with $\alpha^*(a) = 1$. Then a is a strict best response to a mixed strategy.*

PROOF:

Let $\{\alpha^{K_n}\}$ be a sequence of $S(K_n)$-equilibria converging to α^*. Suppose that a is not a strict best response to any mixed strategy. For each n, let $b^{K_n} \in A \setminus \{a\}$ be an action that yields an expected payoff against α^{K_n} at least as large as does a. Then by the law of large numbers, the limit as $K_n \to \infty$ of the probability that a yields a score higher than does b^{K_n} is at most $^1/_2$, so that for any $\varepsilon > 0$ there exists n^* large enough that $w(a, \alpha^{K_n}, K_n) \leq$

$1/2 + \varepsilon$ for all $n > n^*$, contradicting $\alpha^K \to \alpha^*$ since $\alpha^*(a) = 1$.

The condition that a be a strict best response to a mixed strategy, while necessary, is not sufficient for a distribution α^* with $\alpha^*(a) = 1$ to be a limit of $S(K)$-equilibria. For example, in the game

	T	M	B
T	1	1	1
M	1	0	2
B	1	0	2

the action T is a strict best response to the mixed strategy (0, 1, 0), but is not a limit of $S(K)$-equilibria, by the following argument. In any $S(K)$-equilibrium M and B are used with the same probability, since the probability that each of them obtains the highest score is the same. Thus, as in Example 12, the $S(K)$-equilibria approach $(1/4, 3/8, 3/8)$.

As K increases, the players' payoffs in $S(K)$-equilibria may decrease: when everyone makes decisions based on more research, the outcome may be worse for everyone. Consider the following game.

0	2
0	1

The only $S(1)$-equilibrium assigns probability $2 - \sqrt{2}$ to the top action whereas the $S(K)$-equilibrium converges, as K increases, to the distribution that assigns probability 1 to the top action.

V. A General Framework

The notion of $S(K)$-equilibrium embodies a specific class of decision-making procedures. We now suggest a way of modeling the interaction between players who use arbitrary decision-making procedures. This allows us to relate our model to others in the literature and provides a framework for future research.

Let (A, u) be a symmetric two-player game. Let α be a probability distribution on A. We can think of each player's decision as a choice from the list $(L_a(\alpha))_{a \in A}$ of lotteries, where $L_a(\alpha)$ is the probability distribution over the player's payoffs when she takes the action a and her opponent's action is determined by α. Her method of choice from any finite list $\mathcal{L}$ of lotteries may be described by a function C that assigns to $\mathcal{L}$ a probability distribution over $\mathcal{L}$ (her choice may be stochastic). We denote by $C(\mathcal{L})(L)$ the probability that C assigns to the lottery L (an element in the list $\mathcal{L}$). The function C describes a player's behavior. It may be generated by a procedure that the player uses to make her choice and may exhibit the player's rationality or reflect her bounded rationality. Note that a choice function is defined on lists, not sets. For example the choice from the list (L, L, L') may differ from the choices from either of the lists (L', L, L) and (L, L', L'). Defining a choice function in this way allows us to model procedures in which, for example, the order of the alternatives or the multiple presence of an alternative affects the choice.

For any choice function C, we can define an equilibrium notion.

Definition 4: Let (A, u) be a symmetric two-player game and let C be a choice function. A probability distribution α on A is a *C-procedural equilibrium* if

$$C((L_a(\alpha))_{a \in A})(L_x) = \alpha(x) \quad \text{for every } x \in A,$$

where for each action a, $L_a(\alpha)$ is the lottery that takes the value $u(a, b)$ with probability $\alpha(b)$.

Note that *the concept of Nash equilibrium does not fit into this framework:* there is no choice function C such that for every game the set of C-procedural equilibria is equal to the set of Nash equilibria. The reason is that the notion of Nash equilibrium requires a player's mixed strategy not only to be a best response to the other player's strategy, but also to serve as a belief that supports the *other* player's equilibrium choice. To illustrate this point, compare the following two symmetric games.

	a	b
a	1	1
b	1	0

	a	b
a	0	1
b	1	1

The unique mixed strategy Nash equilibria of the two games are different: (a, a) in left-hand game and (b, b) in the right-hand game. But a player's choice problem, given that her opponent chooses the equilibrium action, is the same in both games: choose from the two-element list in which each element is the degenerate lottery that yields 1 with probability 1.

An $S(1)$-equilibrium is a procedural equilibrium for the choice function C for which $C(\mathcal{L})(L)$ is the probability that the realization of the lottery L is higher than every other realization (breaking ties equiprobably) when a player tests each lottery once.

Two other papers study equilibria that can be described as procedural equilibria for other choice functions. Rosenthal (1989) investigates procedural equilibria for choice functions C_θ $(\theta > 0)$ such that

$$
\begin{aligned}
C_\theta(\mathcal{L})(L_1) &- C_\theta(\mathcal{L})(L_2) \\
&= \theta[\mathrm{E}(L_1) - \mathrm{E}(L_2)] \\
&\quad \text{if } C_\theta(\mathcal{L})(L_1) > 0 \\
&\quad \text{and } C_\theta(\mathcal{L})(L_2) > 0 \\
&\leq \theta[\mathrm{E}(L_1) - \mathrm{E}(L_2)] \\
&\quad \text{if } C_\theta(\mathcal{L})(L_1) > 0 \\
&\quad \text{and } C_\theta(\mathcal{L})(L_2) = 0
\end{aligned}
$$

for all lotteries L_1 and L_2, where $\mathrm{E}(L_i)$ is the expected value of L_i. Chen et al. (1997) study a procedural equilibrium for choice functions C_θ $(\theta > 0)$ such that

$$
\begin{aligned}
C_\theta(\mathcal{L})(L_1) / C_\theta(\mathcal{L})(L_2) \\
= [\mathrm{E}(L_1)/\mathrm{E}(L_2)]^\theta
\end{aligned}
$$

(where the expected values are required to be nonnegative). (In each case, smaller values of θ correspond to bigger departures from full rationality, $\theta = 0$ being the case in which each action is assigned the same probability. Chen et al.'s results are formally closely related to those of Richard D. McKelvey and Thomas R. Palfrey, 1995.)

The primitive of the procedure behind our choice function is a player's preference relation. Each player optimizes relative to her preferences, given the beliefs she forms; she is boundedly rational only in the way in which she constructs an action-consequence relationship. By contrast, sensible interpretations of Chen et al.'s and Rosenthal's choice functions appear to require the payoffs to be given a meaning different from that of representing preferences.

VI. Concluding Comments

In this paper we present a tractable solution concept for the analysis of situations in which boundedly rational agents interact. The notion of $S(K)$-equilibrium is a special case of a C-procedural equilibrium, a notion that can be used to investigate some classical questions posed by students of bounded rationality, in particular the effect of a player's decision-making procedure on her achievement in a context in which decision makers interact.

Our approach is static; it does not link the solution concept with any dynamics. We believe that a static approach is valuable, though certainly a dynamic justification for the solution concept we study may be interesting. An investigation of an evolutionary model in which players *occasionally* experiment is given by Drew Fudenberg and David M. Kreps (1988).

REFERENCES

Camerer, Colin. ''Individual Decision Making,'' in John H. Kagel and Alvin E. Roth, eds., *The handbook of experimental economics*. Princeton, NJ: Princeton University Press, 1995, pp. 587–703.

Chen, Hsiao-Chi; Friedman, James W. and Thisse, Jacques-François. ''Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach.'' *Games and Economic Behavior*, January 1997, *18*(1), pp. 32–54.

Fudenberg, Drew and Kreps, David M. ''A Theory of Learning, Experimentation, and Equilibrium in Games.'' Unpublished manuscript, 1988.

McKelvey, Richard D. and Palfrey, Thomas R. ''Quantal Response Equilibria for Normal Form Games.'' *Games and Economic Behavior*, July 1995, *10*(1), pp. 6–38.

Rosenthal, Robert W. ''Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox.'' *Journal of Economic Theory*, August 1981, *25*(1), pp. 92–100.

———. ''A Bounded-Rationality Approach to the Study of Noncooperative Games.'' *International Journal of Game Theory*, 1989, *18*(3), pp. 273–92.

Rubinstein, Ariel. *Modeling bounded rationality*. Cambridge, MA: MIT Press, 1998.

Simon, Herbert A. ''A Behavioral Model of Rational Choice.'' *Quarterly Journal of Economics*, February 1955, *69*(1), pp. 99–118.

———. *Models of bounded rationality*, Vol. 2. Cambridge, MA: MIT Press, 1982.