Stochastic dominance equilibria in two-person noncooperative games

Andres Perea   Hans Peters*   Tim Schulteis
Dries Vermeulen

February 2005

Abstract

Two-person noncooperative games with finitely many pure strategies and ordinal preferences over pure outcomes are considered, in which probability distributions resulting from mixed strategies are evaluated according to t-degree stochastic dominance. A t-best reply is a strategy that induces a t-degree stochastically undominated distribution, and a t-equilibrium is a pair of t-best replies. The paper provides a characterization and existence proofs of t-equilibria in terms of representing utility functions, and shows that for t becoming large—which can be interpreted as the players becoming more risk averse—behavior converges to a specific form of max-min play. More precisely, this means that in the limit each player puts all weight on a strategy that maximizes the worst outcome for the opponent, within the supports of the strategies in the limiting sequence of t-equilibria.

1   Introduction

In order for a mixed strategy Nash equilibrium in a two-person noncooperative game to make sense, a minimal condition is that the players have exact knowledge about cardinal, von Neumann-Morgenstern utility representations of their preferences over the outcomes of the game. This, however, is a strong assumption: it presumes that such representations exist and are known to the players, and it results, generically, in sharp point-beliefs about the mixed strategy of the opponent.

   In this paper we start by assuming that the players have complete knowledge only about the ordinal preferences over the pure outcomes in the game.

A natural way, then, to evaluate probability distributions over the outcomes induced by mixed strategies is to order them by first-degree stochastic dominance: probability distributions that shift more probability to better outcomes are considered more attractive. Since the stochastic dominance ordering is not complete, we call a mixed strategy of a player a best reply against the strategy of the opponent if it induces an undominated probability distribution over the outcomes. An equilibrium is a pair of best replies. Fishburn (1978) established that the set of equilibria is equal to the union of all sets of Nash equilibria, taken over all possible utility representations of the preferences. This is an intuitive result in view of the familiar characterization of first-degree stochastic dominance which says that a distribution is undominated if and only if it maximizes expected utility for at least one utility representation of the ordinal preferences.

We take this analysis a good deal further by assuming that the players have more specific knowledge about their risk attitudes or, equivalently, about the derived classes of utility functions that may represent their preferences. Specifically, we study so-called t-equilibria, where the natural number t is the degree of stochastic dominance used to evaluate probability distributions. As is well-known, a distribution is second-degree stochastically undominated if and only if it maximizes expected utility for at least one *concave* utility representation of the ordinal preferences. Loosely speaking, higher degrees of stochastic dominance correspond to higher degrees of risk aversion.

After preliminaries about stochastic dominance, games and equilibria in Sections 2 and 3, we consider an example in Section 4 which nicely illustrates these concepts and the main results of the paper. These results are, first, a characterization of t-degree stochastic dominance in terms of representing utility functions and existence of t-equilibria in Section 5, and, second, limit behavior as the degree of stochastic dominance t goes to infinity, in Section 6.

Existence of t-equilibria can be established directly by using a fixed point argument (this is done in Appendix B) or indirectly by using representation by utility functions and existence of Nash equilibrium (Section 5).

It follows from the results of Section 5 that the sets of characterizing utility functions become smaller as t grows, as already indicated by the transition from $t = 1$ to $t = 2$. Consequently, the best reply correspondences and sets of t-equilibria decrease as well. In Section 6 we provide a complete characterization of the sets of pure strategies that can serve as supports for t-equilibria as t becomes large. In the limit, such equilibria converge to max-min play, in the sense that each player plays a pure strategy that, among the strategies in the supports, maximizes the worst outcome for the opponent. Observe that this is very different from what is usually meant by max-min play, namely that players maximize their own worst outcomes. Max-min play in the present setting is closer to equilibrium play: for large t, a player puts probability close to 1 on the pure strategy maximizing the worst outcome for the opponent among the strategies in the support of the opponent's mixed strategy, in order to keep all these strategies undominated. The intuition for this is that, as t becomes large,

the opponent attaches increasing weights to worse outcomes, and to compensate for this a player should put low weights on those own strategies that possibly result in these worse outcomes for the opponent.

Section 7 concludes the paper with a brief discussion of related literature and directions for further research. Appendix A collects the proofs of Section 2, and Appendix B gives an independent existence result of t-equilibria by a fixed point argument.

2 Stochastic dominance

Let $\ell \geq 1$ be an integer and let $O = \{1, \ldots, \ell\}$ be a set of ℓ *alternatives*. For $1 \leq k < l \leq \ell$ we assume that a decision maker strictly prefers alternative l to alternative k.

For a probability distribution $r = (r_1, \ldots, r_\ell)$ on O (so each l occurs with probability r_l) we define, recursively, for each $l \in \{1, \ldots, \ell\}$, $F_r^0(l) = r_l$ and

$$F_r^t(l) = \sum_{i=1}^{l} F_r^{t-1}(i) \ \ (t \geq 1).$$

So F_r^1 is the cumulative distribution function of $F_r^0 = r$ and, similarly, F_r^t 'accumulates' the weights assigned by F_r^{t-1}. For probability distributions r and s on O, r *t-th degree stochastically dominates* s if

$$F_r^t(l) \leq F_s^t(l) \text{ for every } l \in \{1, \ldots, \ell\}$$

with at least one inequality strict. Observe that the latter follows if $r \neq s$. For $t = 1$, this relation means that r puts more probability on better alternatives than s. It is well known that this is equivalent to the expected utility under r being at least as large as the expected utility under s for every utility representation of the preference relation. For second degree stochastic dominance, an analogous equivalence holds if we restrict to concave utility functions, or, more generally, utility functions with non-increasing differences between adjacent alternatives. Therefore, first degree stochastic dominance is a purely ordinal concept but this is not the case for second and higher degrees. Note that t-th degree stochastic dominance implies $t + 1$-th degree stochastic dominance. In a relative sense, a similar relation holds between $t + 1$-th and t-th degree stochastic dominance as between second and first degree stochastic dominance. Thus, higher degree stochastic dominance can be associated with increased risk aversion of decision makers.

Fishburn (1976, 1980) characterizes stochastic dominance in terms of utility functions and in terms of moments of distributions. Below, we provide a characterization of stochastically undominated distributions in terms of utility functions for the context of this paper.

3

Denote $F_r^t = (F_r^t(1), \ldots, F_r^t(\ell))$ and let $A = [a_{ij}]$ be the $\ell \times \ell$-matrix with

$$a_{ij} = \left\{ \begin{array}{ll} 1 & \text{if } i \leq j \\ 0 & \text{if } i > j \end{array} \right. \quad \text{for all } i, j \in \{1, \ldots, \ell\}.$$

Write $A^t = A \cdot A \cdot \ldots \cdot A$ (t times, $t \in \mathbb{N}$). The following lemma gives a convenient representation of F_r^t.

Lemma 2.1 $F_r^t = rA^t$ for every $t \in \mathbb{N}$.

Denoting the element in row i and column j of A^t by a_{ij}^t, we derive the following expression for this number.

Lemma 2.2 Let $t \in \mathbb{N}$. Then

$$a_{ij}^t = \left\{ \begin{array}{ll} \frac{(j-i+t-1)!}{(j-i)!(t-1)!} & \text{if } i \leq j \\ 0 & \text{if } i > j \end{array} \right. .$$

The following lemma applies to t approaching infinity.

Lemma 2.3 Let $i, i', j \in \{1, \ldots, \ell\}$ with $i < i' \leq j$. Then $a_{ij}^t \geq a_{i'j}^t$ for every $t \in \mathbb{N}$ and $\lim_{t \to \infty} a_{ij}^t / a_{i'j}^t = \infty$.

For every $t \in \mathbb{N}$ let[1]

$$U^t := \{u \in \mathbb{R}^\ell \mid u = -A^t c \text{ for some } c \in \mathbb{R}^\ell, c > 0\}.$$

An element u of U^t can be interpreted as a utility function representing the implicit preference ordering by the assignment $i \mapsto u_i$, since $u_1 < u_2 < \cdots < u_\ell$. Call a probability distribution r over O *t-undominated* if there is no probability distribution s such that s t-th degree stochastically dominates r. Then we have the following result, which adapts Fishburn (1976) to our context.

Proposition 2.4 *The probability distribution r over O is t-undominated if and only if there is a $u \in U^t$ such that $\sum_{l=1}^{\ell} r_l u_l \geq \sum_{l=1}^{\ell} s_l u_l$ for all probability distributions s over O.*

Note that U^1 contains essentially any utility representation of σ. This is consistent with remarks made earlier. The set U^t is decreasing in t.

Proofs of Lemmas 2.1–2.3 and Proposition 2.4 can be found in Appendix A.

---

[1] For vectors x and y, $x > y$ [$x \geq y$] means $x_i > y_i$ [$x_i \geq y_i$] for every coordinate i. Similarly for $x < y$, $x \leq y$.

3 Two-person games and t-equilibria

Consider two players. Player 1 has pure strategy set $M = \{1, \ldots, m\}$. A (mixed) strategy for player 1 is a probability distribution over M. Denote the set of strategies for player 1 by Δ^M. A pure strategy i is identified with the mixed strategy $e^i \in \Delta^M$, where $e^i_k = 1$ if $k = i$ and $e^i_k = 0$ otherwise. Similarly, player 2 has pure strategy set $N = \{1, \ldots, n\}$ and (mixed) strategy set Δ^N. A pure strategy j is identified with the mixed strategy $e^j \in \Delta^N$. If player 1 plays pure strategy i and player 2 pure strategy j, then the alternative o_{ij} results. If player 1 plays $p \in \Delta^M$ and player 2 plays $q \in \Delta^N$, then o_{ij} results with probability $p_i q_j$. Let $O := \{o_{ij} \mid i \in M, \ j \in N\}$ and assume that players 1 and 2 have preference relations represented, respectively, by bijections $\sigma, \tau : M \times N \rightarrow \{1, \ldots, mn\}$. For instance, if $1 \leq k < l \leq mn$ then player 1 strictly prefers o_{ij} to $o_{i'j'}$, where $\sigma(i, j) = l$ and $\sigma(i', j') = k$.

For $p \in \Delta^M$ and $q \in \Delta^N$ we denote by $pq\sigma$ the vector of probabilities with l-th coordinate $pq\sigma_l = p_i q_j$ such that $\sigma(i, j) = l$, for all $l \in \{1, \ldots, mn\}$. We assume that the players evaluate strategies according to a stochastic dominance criterion. More precisely, let $t \in \mathbb{N}$ and fix a strategy $q \in \Delta^N$ for player 2. Then a strategy $p \in \Delta^M$ of player 1 results in the weight vector $F^t_{pq\sigma}$, which depends on σ and assigns weight $F^t_{pq\sigma}(\sigma(i, j))$ to alternative o_{ij}.

We call p a t-*best reply against* q if there is no $p' \in \Delta^M$ such that $p'q$ t-th degree stochastically dominates pq. The definition of a t-best reply q against p is analogous. A pair $(p, q) \in \Delta^M \times \Delta^N$ is a t-*equilibrium* if p is a t-best reply against q and vice versa. By E^t we denote the set of t-equilibria.

4 An example

The example presented here is illustrative of the main results of this paper, namely (i) existence and characterization of t-equilibria; and (ii) asymptotic behavior for t approaching infinity.

Let $m = n = 2$ and consider the game

$$\begin{bmatrix} o_{11} & o_{12} \\ o_{21} & o_{22} \end{bmatrix}$$

where the rows are the pure strategies of player 1, the columns those of player 2, and the preferences are given by $\sigma(1, 2) = 1$, $\sigma(2, 2) = 2$, $\sigma(2, 1) = 3$, and $\sigma(1, 1) = 4$ for player 1 and $\tau(1, 2) = 1$, $\tau(1, 1) = 2$, $\tau(2, 1) = 3$, and $\tau(2, 2) = 4$ for player 2.

We concentrate on player 1. The matrix A^t $(t \geq 1)$ can be computed using

Lemma 2.2. This results in

$$A^t = \begin{bmatrix} 1 & t & \frac{1}{2}t(t+1) & \frac{1}{6}t(t+1)(t+2) \\ 0 & 1 & t & \frac{1}{2}t(t+1) \\ 0 & 0 & 1 & t \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Consider strategies $p = (p_1, p_2)$ and $q = (q_1, q_2)$ for players 1 and 2, respectively. In order to examine t-best replies of player 1 we compute (cf. Lemma 2.1)

$$F^t_{pq\sigma} = (p_1 q_2, p_2 q_2, p_2 q_1, p_1 q_1) A^t.$$

Dropping the part that does not depend on p and which therefore is not needed to compute t-best replies of player 1, this results in the vector

$$p_1 \left(q_2, (t-1)q_2, \frac{1}{2}(t^2 - t)q_2 - q_1, \frac{1}{6}(t^3 - t)q_2 - q_1(t-1) \right). \qquad (1)$$

For $t = 1$, (1) reduces to $p_1(q_2, 0, -q_1, 0)$. Since player 1 wants to 'minimize' this vector, the t-best responses are $p_1 = 0$ if $q_1 = 0$, $p_1 = 1$ if $q_1 = 1$, and any $0 \le p_1 \le 1$ if $0 < q_1 < 1$. With a similar argument for player 2 (not reproduced here) we find the set E^1: it contains the two pure Nash equilibria of the game, resulting in o_{11} and in o_{22}, and *all* strategy combinations where no player plays a pure strategy. This is no surprise: in general, E^1 consists of all strategy combinations that are a Nash equilibrium for at least one choice of utility functions representing σ and τ. This is a consequence of the familiar characterization of first degree stochastic dominance using utility representations, mentioned in Section 2. See also Fishburn (1978), where the result is derived formally.

Next, consider $t = 2$, so best replies are second degree stochastically undominated. By substituting $t = 2$ in (1) it follows that for $q_1 = 1$ the 2-best reply is $p_1 = 1$, for $1 > q_1 > \frac{1}{2}$ any $0 \le p_1 \le 1$ is a 2-best reply, and for $q_1 \le \frac{1}{2}$ the 2-best reply is $p_1 = 0$. Again after a similar argument for player 2 it follows that E^2 consists of the two pure Nash equilibria of the game plus the set

$$\{(p, q) \mid \frac{1}{2} < q_1 < 1, \ 0 < p_1 < \frac{1}{2}\}.$$

In these mixed strategy equilibria player 1 puts a larger weight on row 2. Row 2 is player 1's max-min pure strategy: he prefers the worst alternative in row 2, o_{22}, to the worst alternative in row 1, o_{12}. Thus, one might be tempted to conclude that a higher t leads to max-min play. This, however, is deceptive. As will turn out later, what is important is that row 2 is the max-min row from the point of view of player 2: player 2 prefers the worst alternative (for him) of row 2, o_{21}, to the worst alternative of row 1, o_{12}. (A similar consideration holds for the strategy of player 2.)

Observe also that the 2-best reply correspondences of the players are not upper semi-continuous (their graphs are not closed).

For $t > 2$, let $\hat{q}_1 = (t^3 - t)/(t^3 + 5t - 6)$. For $0 \leq q_1 \leq \hat{q}_1$ the t-best reply is $p_1 = 0$, for $\hat{q}_1 < q_1 < 1$ any p_1 is a t-best reply, and for $q_1 = 1$ the t-best reply is $p_1 = 1$. The t-equilibria are again the two pure Nash equilibria in the game together with the collection

$$\{(p, q) \mid \frac{t^3 - t}{t^3 + 5t - 6} < q_1 < 1,\ 0 < p_1 < 1 - \frac{t^3 - t}{t^3 + 5t - 6}\},$$

For $t \to \infty$ these mixed strategy t-equilibria converge to the pure strategy combination of row 2 and column 1.

5   Existence and characterization of t-equilibria

The existence of t-equilibria can be proved directly by applying a fixed point argument to the best-reply correspondences. This proof is not completely straightforward since the best reply correspondences do not have to be upper semi-continuous, see the example in the previous section, so that the argument has to be applied to a suitable sub-correspondence. For the details, see Appendix B.

Alternatively, t-equilibria can be characterized as Nash equilibria for suitably chosen utility functions, by an argument analogous to the proof of Proposition 2.4. Existence then follows from the standard existence result for Nash equilibrium. This is the approach taken here. For $t = 1$, this has already been done in Fishburn (1978).

In the next lemma we consider the game as defined in Section 3. The bijection σ represents the preference relation of player 1. The set U^t was defined in Section 2.

Lemma 5.1 *Let $p \in \Delta^M$, $q \in \Delta^N$, and $t \geq 1$. Then p is a t-best reply against q if and only if there is a $u^t \in U^t$ such that*

$$\sum_{i=1}^{m} \sum_{j=1}^{n} p_i q_j u_{\sigma(i,j)}^t \geq \sum_{i=1}^{m} \sum_{j=1}^{n} p_i' q_j u_{\sigma(i,j)}^t$$

for all $p' \in \Delta^M$.

Proof. p is a t-best reply against q if and only if there is no $p' \in \Delta^M$ such that $F_{p'q\sigma}^t \leq F_{pq\sigma}^t$ with at least one coordinate strictly smaller. This is the case if and only if

$$\{x \in I\!\!R^{mn} \mid x \leq F_{pq\sigma}^t\} \cap \{x \in I\!\!R^{mn} \mid x = F_{p'q\sigma}^t \text{ for some } p' \in \Delta^M\} = \{F_{pq\sigma}^t\}.$$

By a separation argument analogous to the one in the proof of Proposition 2.4 this, in turn, holds if and only if there is a $c \in \mathbb{R}^{mn}$ with $c > 0$ such that $F^t_{pq\sigma} \cdot c \leq F^t_{p'q\sigma} \cdot c$ for all $p' \in \Delta^M$. By Lemma 2.1 this is equivalent to $(pq\sigma)A^t c \leq (p'q\sigma)A^t c$ for all $p' \in \Delta^M$. The proof is complete by taking $u^t := -A^t c$. □

Lemma 5.1 can be formulated for player 2 in an analogous way. Then t-equilibria can be characterized as follows.

Corollary 5.2 *Let σ and τ represent the preferences of players 1 and 2, respectively. Let $t \geq 1$, $p^* \in \Delta^M$, and $q^* \in \Delta^N$. Then (p^*, q^*) is a t-equilibrium if and only if there are $u, v \in U^t$ such that (p^*, q^*) is a Nash equilibrium for the payoff functions $O \to \mathbb{R}$ defined by $o_{ij} \mapsto u_{\sigma(i,j)}$ and $o_{ij} \mapsto v_{\tau(i,j)}$ for players 1 and 2, respectively.*

Since Nash equilibria always exist, Corollary 5.2 implies existence of t-equilibria.

Corollary 5.3 $E^t \neq \emptyset$ *for every $t \geq 1$.*

6 Limiting behavior of equilibria

The example in Section 4 suggests some kind of max-min behavior of the players in a t-equilibrium for t going to infinity. In this section we consider this in detail. The setting is the general game model as defined in Section 3. Unless stated otherwise, the number t is arbitrary but fixed. The preferences of the players are represented by the bijections σ for player 1 and τ for player 2.

For $p \in \Delta^M$, the *support* of p is the set

$$\mathrm{supp}(p) = \{i \in M \mid p_i > 0\}.$$

For $q \in \Delta^N$, $\mathrm{supp}(q)$ is defined in the same way. We start with an auxiliary result.

Lemma 6.1 *Let $p^t \in \Delta^M$ and let $q^t \in \Delta^N$. Let $p \in \Delta^M$ and $q \in \Delta^N$ with $\mathrm{supp}(p) \subseteq \mathrm{supp}(p^t)$ and $\mathrm{supp}(q) \subseteq \mathrm{supp}(q^t)$. Then*

(i) if p^t is a t-best reply against q^t, then p is a t-best reply against q^t;

(ii) if q^t is a t-best reply against p^t, then q is a t-best reply against p^t.

Proof. Apply Lemma 5.1. □

Lemma 6.2 *Let $I \subseteq M$, $J \subseteq N$, and let $(p^t, q^t)_{t \in \mathbb{N}}$ be a sequence of pairs of mixed strategies such that $I = \mathrm{supp}(p^t)$ and $J = \mathrm{supp}(q^t)$ for all $t \in \mathbb{N}$.*

(i) *Let p^t be a t-best reply against q^t for every $t \in \mathbb{N}$. Then, for every $i \in M$, there is a $j \in J$ such that $\sigma(i,j) < \sigma(i',j)$ for all $i' \in I \backslash \{i\}$.*

(ii) *Let q^t be a t-best reply against p^t for every $t \in \mathbb{N}$. Then, for every $j \in N$, there is an $i \in I$ such that $\tau(i,j) < \tau(i,j')$ for all $j' \in J \backslash \{j\}$.*

Proof. We only prove (i), the proof of (ii) is analogous. Suppose (i) were not true. Then there is an $\hat{i} \in M$ such that for every $j \in J$

$$I_j := \{i \in I \mid \sigma(i,j) < \sigma(\hat{i},j)\} \neq \emptyset.$$

For every $j \in J$, choose an $i_j \in I_j$. Let $\hat{I} := \{i_j \mid j \in J\}$, $s := |\hat{I}|$, [2] and define $\hat{p} \in \Delta^M$ by $\hat{p}_i = 1/s$ if $i \in \hat{I}$ and $\hat{p}_i = 0$ otherwise. We will show that, for t sufficiently large,

$$F^t_{e^{\hat{i}} q^t \sigma} \leq F^t_{\hat{p} q^t \sigma}. \tag{2}$$

Since, clearly, the two probability distributions $e^{\hat{i}} q^t \sigma$ and $\hat{p} q^t \sigma$ in (2) are different, this means that at least one of the inequalities must be strict if (2) holds. Since $\text{supp}(\hat{p}) \subseteq I = \text{supp}(p^t)$ for all t, (2) contradicts Lemma 6.1. This proves (i).

We are left to prove (2), hence we are left to prove

$$F^t_{e^{\hat{i}} q^t \sigma}(k) \leq F^t_{\hat{p} q^t \sigma}(k) \text{ for all } k = 1, \dots, mn. \tag{3}$$

Fix $k \in \{1, \dots, mn\}$ and define $J_k := \{j \in J \mid \sigma(\hat{i},j) \leq k\}$. Then

$$
\begin{aligned}
F^t_{e^{\hat{i}} q^t \sigma}(k) &= \sum_{l=1}^{mn} (e^{\hat{i}} q^t \sigma)_l a^t_{lk} \\
&= \sum_{l=1}^{k} (e^{\hat{i}} q^t \sigma)_l a^t_{lk} \\
&= \sum_{j \in J_k} q^t_j a^t_{\sigma(\hat{i},j),k}.
\end{aligned}
\tag{4}
$$

Here, the first equality follows from Lemma 2.1, the second equality from Lemma 2.2, and the last equality by Lemma 2.2 and the definition of J_k.

Now

$$
\begin{aligned}
F^t_{\hat{p} q^t \sigma}(k) &= \sum_{l=1}^{k} (\hat{p} q^t \sigma)_l a^t_{lk} \\
&\geq \sum_{j \in J_k} q^t_j \hat{p}_{i_j} a^t_{\sigma(i_j,j),k}
\end{aligned}
$$

---

[2] $|\cdot|$ denotes the cardinality of a set.

$$= \sum_{j \in J_k} \frac{1}{s} q_j^t a_{\sigma(i_j,j),k}^t$$

$$\geq \sum_{j \in J_k} \frac{1}{s} q_j^t a_{\sigma(\hat{\imath},j)-1,k}^t. \tag{5}$$

The first equality follows again by Lemmas 2.1 and 2.2. The first inequality follows since some terms are left out. The second equality follows by definition of $\hat{p}$ since $i_j \in \hat{I}$ for every $j \in J$. The last inequality follows by the first statement in Lemma 2.3 since $\sigma(i_j, j) \leq \sigma(\hat{\imath}, j) - 1$ for every $j \in J$.

If $J_k = \emptyset$ then (3) follows immediately from (4) and (5). Otherwise, by the second statement in Lemma 2.3 there is a t sufficiently large such that for every $k = 1, 2, \ldots, mn$ and $j \in J_k$ we have

$$\frac{1}{s} q_j^t a_{\sigma(\hat{\imath},j)-1,k}^t \geq q_j^t a_{\sigma(\hat{\imath},j),k}^t. \tag{6}$$

Then (3) follows from (4), (5), and (6). $\qquad \square$

Lemma 6.2(i) implies that for every row (pure strategy) i in I there must be a column (pure strategy) j in J such that the resulting outcome o_{ij} is the worst outcome for player 1 in that column restricted to the rows in I. In turn, this implies $|I| \leq |J|$. Similarly, Lemma 6.2(ii) implies $|J| \leq |I|$. So we have the following result.

Corollary 6.3 *Let $I \subseteq M$, $J \subseteq N$, and let $(p^t, q^t)_{t \in \mathbb{N}}$ be a sequence of t-equilibria such that $I = \mathrm{supp}(p^t)$ and $J = \mathrm{supp}(q^t)$ for all $t \in \mathbb{N}$. Then $|I| = |J|$.*

If $|I| = |J| = 1$ in Corollary 6.3, then the sequence of t-equilibria reduces to the constant Nash equilibrium in which player 1 picks the best element from the column played by player 2 and player 2 picks the best element from the row played by player 1.

The next result implies that in Corollary 6.3 the t-equilibria must converge to pure strategy combinations.

Lemma 6.4 *Let $I \subseteq M$, $J \subseteq N$, and let $(p^t, q^t)_{t \in \mathbb{N}}$ be a sequence of t-equilibria such that $I = \mathrm{supp}(p^t)$ and $J = \mathrm{supp}(q^t)$ for all $t \in \mathbb{N}$. Let $\hat{\imath} \in I$ and $\hat{\jmath} \in J$ such that $\min\{\tau(i,j) \mid j \in J\} < \min\{\tau(\hat{\imath}, j) \mid j \in J\}$ for all $i \in I \backslash \{\hat{\imath}\}$ and $\min\{\sigma(i,j) \mid i \in I\} < \min\{\sigma(i, \hat{\jmath}) \mid i \in I\}$ for all $j \in J \backslash \{\hat{\jmath}\}$. Then (i) $\lim_{t \to \infty} p_{\hat{\imath}}^t = 1$ and (ii) $\lim_{t \to \infty} q_{\hat{\jmath}}^t = 1$.*

Proof. We only prove (ii), the proof of (i) is analogous. By Lemma 6.2 and Corollary 6.3 we may renumber the strategies of the players such that:

(a) $I = J = \{1, \ldots, s\}$ for some $s \geq 1$;

(b) $\sigma(j,j) < \sigma(i,j)$ for every $j \in J$ and $i \in I\backslash\{j\}$;

(c) $\sigma(1,1) < \sigma(2,2) < \ldots < \sigma(s,s)$.

Note that $\hat{\jmath} = s$. Let $s' \in \{1, \ldots, s-1\}$ arbitrary. To prove (ii), it is sufficient to prove that $\lim_{t \to \infty} q_{s'}^t = 0$. Suppose that this is not the case. Then we may assume that there is an $\alpha > 0$ such that $q_{s'}^t \geq \alpha$ for all t (otherwise there is a subsequence with this property and we can apply the following argument to this subsequence). Define the strategy $\hat{p} \in \Delta^M$ by $\hat{p}_i = 1/s'$ for $i = 1, \ldots, s'$ and $\hat{p}_i = 0$ otherwise. We will show that $\hat{p}q^t\sigma$ is t-dominated by $e^{s'+1}q^t\sigma$ for sufficiently large t, which contradicts Lemma 6.1 and therefore completes the proof. So we are left to show that for t sufficiently large

$$F_{e^{s'+1}q^t\sigma}^t(k) \leq F_{\hat{p}q^t\sigma}^t(k) \text{ for every } k = 1, \ldots, mn. \tag{7}$$

(Since the probability distributions $e^{s'+1}q^t\sigma$ and $\hat{p}q^t\sigma$ are clearly different, at least one of the inequalities in (7) must be strict.)

Let $k \in \{1, \ldots, mn\}$. By (b) and Lemma 2.3 we can choose t_1 such that for all $t \geq t_1$ we have

$$a_{\sigma(j,j),k}^t \geq s\, a_{\sigma(s'+1,j),k}^t \text{ for all } j = 1, \ldots, s'-1. \tag{8}$$

Also by (b) and Lemma 2.3 we can choose t_2 such that for all $t \geq t_2$ we have

$$\alpha a_{\sigma(s',s'),k}^t \geq 2s\, q_{s'}^t a_{\sigma(s'+1,s'),k}^t. \tag{9}$$

By (b), (c), and Lemma 2.3 we can choose t_3 such that for all $t \geq t_3$ we have

$$\alpha a_{\sigma(s',s'),k}^t \geq 2s \sum_{j \in J:\ j \geq s'+1} q_j^t a_{\sigma(s'+1,j),k}^t. \tag{10}$$

Then, for $t \geq \max\{t_1, t_2, t_3\}$, we have

$$
\begin{aligned}
F_{e^{s'+1}q^t\sigma}^t(k) &= \sum_{j \in J:\ \sigma(s'+1,j) \leq k} q_j^t a_{\sigma(s'+1,j),k}^t \\
&= \sum_{j \in J:\ j < s',\ \sigma(s'+1,j) \leq k} q_j^t a_{\sigma(s'+1,j),k}^t \\
&\quad + \sum_{j \in J:\ j \geq s',\ \sigma(s'+1,j) \leq k} q_j^t a_{\sigma(s'+1,j),k}^t \\
&\leq \sum_{j \in J:\ j < s'} \frac{1}{s} q_j^t a_{\sigma(j,j),k}^t + \frac{1}{2s}\alpha a_{\sigma(s',s'),k}^t + \frac{1}{2s}\alpha a_{\sigma(s',s'),k}^t \\
&\leq F_{\hat{p}q^t\sigma}^t(k),
\end{aligned}
$$

where the first inequality follows from (8)–(10). This implies (7) and completes the proof of the lemma. $\square$

Lemma 6.4 has the following converse.

11

Lemma 6.5 *Let $I \subseteq M$ and $J \subseteq N$ satisfy conditions (i) and (ii) in Lemma 6.2, and let $t \in \mathbb{N}$. Then there are $p^t \in \Delta^M$ and $q^t \in \Delta^N$ with $\mathrm{supp}(p^t) = I$, $\mathrm{supp}(q^t) = J$, and $(p^t, q^t) \in E^t$.*

Proof. Note that, as before, (i) and (ii) in Lemma 6.2 imply $|I| = |J|$. If $|I| = |J| = 1$, then there is a pure Nash equilibrium $(p^t, q^t) \in E^t$ with supports I and J. Assume now that $|I| = |J| \geq 2$. As in the proof of Lemma 6.4 we may renumber the pure strategies of the players such that

  (a) $I = J = \{1, \ldots, s\}$ for some $s \geq 2$;

  (b) $\sigma(j, j) < \sigma(i, j)$ for every $j \in J$ and $i \in I \backslash \{j\}$;

  (c) $\sigma(1, 1) < \sigma(2, 2) < \ldots < \sigma(s, s)$.

Define $q^t \in \Delta^N$ with $\mathrm{supp}(q^t) = \{1, \ldots, s\}$ such that

$$q_j^t / q_{j-1}^t = m \, n \, a_{1,mn}^t \text{ for every } j = 2, \ldots, s. \tag{11}$$

We will show that every $p \in \Delta^M$ with $\mathrm{supp}(p) \subseteq I$ is a t-best reply against q^t. Since we can define p^t analogously and show that every $q \in \Delta^N$ with $\mathrm{supp}(q) \subseteq J$ is a t-best reply against p^t, the proof is complete.

So let $p \in \Delta^M$ with $\mathrm{supp}(p) \subseteq I$. Assume, contrary to what we wish to prove, that there is a $p' \in \Delta^M$ such that pq^t is t-dominated by $p'q^t$. We first argue that without loss of generality $\mathrm{supp}(p) \cap \mathrm{supp}(p') = \emptyset$. For, suppose that i is an element in this intersection, and let $\alpha := \min\{p_i, p_i'\}$. Define $\bar{p} := 1/(1 - \alpha)\,(p - \alpha e^i)$ and $\bar{p}' := 1/(1 - \alpha)\,(p' - \alpha e^i)$. Then $\bar{p}, \bar{p}' \in \Delta^M$, and $\bar{p}q^t$ is still t-dominated by $\bar{p}'q^t$, whereas $i \notin \mathrm{supp}(\bar{p}) \cap \mathrm{supp}(\bar{p}')$ and $\mathrm{supp}(\bar{p}) \subseteq I$.

So assume that $\mathrm{supp}(p) \cap \mathrm{supp}(p') = \emptyset$, and take $\hat{\imath} \in \mathrm{supp}(p')$ such that $p_{\hat{\imath}}' \geq 1/m$. If $\hat{\imath} \in I$ then let $\hat{\jmath} := \hat{\imath}$. Then, by condition (b), $\sigma(\hat{\imath}, \hat{\jmath}) < \sigma(i, \hat{\jmath})$ for all $i \in I \backslash \{\hat{\imath}\}$, hence for all $i \in \mathrm{supp}(p)$ in particular. If $\hat{\imath} \in M \backslash I$ then by condition (i) in Lemma 6.2 we can take $\hat{\jmath} \in J$ such that $\sigma(\hat{\imath}, \hat{\jmath}) < \sigma(i, \hat{\jmath})$ for all $i \in I$, hence for all $i \in \mathrm{supp}(p)$. Together with conditions (b) and (c) this implies

$$\sigma(\hat{\imath}, \hat{\jmath}) < \sigma(i, j) \text{ for all } i \in \mathrm{supp}(p) \text{ and } j \in \{\hat{\jmath}, \ldots, s\}. \tag{12}$$

Let $k := \sigma(\hat{\imath}, \hat{\jmath})$. Then

$$
\begin{aligned}
F_{pq^t\sigma}^t(k) &= \sum_{l=1}^{k} (pq^t\sigma)_l a_{lk}^t \\
&= \sum_{j=1}^{\hat{\jmath}-1} q_j^t \left(\sum_{i \in I: \ p_i > 0, \ \sigma(i,j) < k} p_i \, a_{\sigma(i,j),k}^t \right) \\
&\leq \sum_{j=1}^{\hat{\jmath}-1} q_j^t a_{1,mn}^t
\end{aligned}
$$

$$
\begin{aligned}
&< \quad n q_{\hat{j}-1}^t \, a_{1,mn}^t \\
&= \quad q_{\hat{j}}^t / m \\
&\leq \quad (1/m) \, q_{\hat{j}}^t a_{\sigma(\hat{i},\hat{j}),k}^t \\
&\leq \quad F_{p'q^t\sigma}^t(k)
\end{aligned} \qquad (13)
$$

where the first equality follows from Lemma 2.1; the second equality by (12); the first inequality by Lemma 2.2; the second (strict) inequality and the third equality by (11); the third inequality since $a_{\sigma(\hat{i},\hat{j}),k}^t \geq 1$; and the final inequality by Lemma 2.1 and the choice of $p_{\hat{i}}' \geq 1/m$.

Since (13) contradicts the assumption that pq^t is t-dominated by $p'q^t$, the proof of the lemma is complete. $\qquad \square$

Lemmas 6.2, 6.4, and 6.5, and Corollary 6.3 can be summarized as follows.

Theorem 6.6 *If $I \subseteq M$, $J \subseteq N$, and if $(p^t, q^t)_{t \in \mathbb{N}}$ is a sequence of t-equilibria such that $I = \mathrm{supp}(p^t)$ and $J = \mathrm{supp}(q^t)$ for all $t \in \mathbb{N}$, then (i) and (ii) in Lemma 6.2 hold, $|I| = |J|$, and the sequence of t-equilibria converges to the pure strategy combination $(\hat{i}, \hat{j})$, where $\hat{i}$ and $\hat{j}$ are as in Lemma 6.4. Conversely, if $\emptyset \neq I \subseteq M$ and $\emptyset \neq J \subseteq N$ satisfy (i) and (ii) in Lemma 6.2, then a sequence of t-equilibria with supports I and J for players 1 and 2, respectively, exists.*

As announced earlier, the results in this section imply that, as t becomes large, the equilibrium behavior of the players converges to max-min play in a specific sense. Take any sequence of t-equilibria with (without loss of generality) constant supports I and J of the players' strategies. Then, in the limit, player 1 puts all weight on that strategy (row) in I in which the worst outcome for player 2 with respect to the strategies (columns) in J is maximal among all rows in I; and player 2 puts all weight on that column in J in which the worst outcome for player 1 with respect to the rows in I is maximal among all columns in J. The next examples illustrate this further.

Example 6.7 Consider the following 3×3 game, in which the numbers express the ordinal preferences of the players:

$$
\begin{bmatrix}
1,1 & 6,5 & 9,4 \\
5,6 & 2,2 & 8,7 \\
4,9 & 7,8 & 3,3
\end{bmatrix}.
$$

In this game, $I := M = \{1,2,3\}$ and $J := N = \{1,2,3\}$ clearly satisfy (i) and (ii) in Lemma 6.2. Obviously, $\hat{i} = \hat{j} = 3$. For every $t \in \mathbb{N}$ the t-equilibrium used in the proof of Lemma 6.5 is defined by

$$
p^t = q^t = \left(\frac{1}{1+A+A^2}, \frac{A}{1+A+A^2}, \frac{A^2}{1+A+A^2} \right)
$$

where $A = 9(7+t)!/8!(t-1)!$. In this equilibrium the weights on the first two rows (and columns) converge to 0 and p_1^t/p_2^t (and q_1^t/q_2^t) converges to zero as well. That is, the weight on the first row (column) goes to 0 much faster than the weight on the second row (column). The latter phenomenon is not a necessary one: in the present example, for instance, it is also possible to have a sequence of t-equilibria with equal weights on the first two rows (columns). It can be verified that taking $p_1^t = p_2^t = q_1^t = q_2^t =: \alpha^t$ such that $(1-2\alpha^t)/\alpha^t > a_{1,3}^t + a_{1,8}^t$ for every $t \in I\!\!N$ is again a t-equilibrium. Finally, $I = M$ and $J = N$ are the only subsets of pure strategies satisfying (i) and (ii) in Lemma 6.2, hence the only supports of t-equilibria. Hence, in the limit each player plays his third strategy, resulting in the 'payoffs' $(3,3)$.

Example 6.8 Consider the following 3×3 game:

$$\begin{bmatrix} 4,4 & 8,5 & 3,6 \\ 5,8 & 7,7 & 2,9 \\ 6,3 & 9,2 & 1,1 \end{bmatrix}.$$

The following combinations satisfy (i) and (ii) in Lemma 6.2:

(a) $I = \{3\}$, $J = \{1\}$, resulting in $(3,1)$ in the limit;

(b) $I = \{1\}$, $J = \{3\}$, resulting in $(1,3)$ in the limit;

(c) $I = \{1,3\}$, $J = \{2,3\}$, resulting in $(1,2)$ in the limit;

(d) $I = \{2,3\}$, $J = \{1,3\}$, resulting in $(2,1)$ in the limit;

(e) $I = M$, $J = N$, resulting in $(2,2)$ in the limit.

This means that the 'payoff pairs' that can arise as limits of t-equilibria are $(6,3)$, $(3,6)$, $(8,5)$, $(5,8)$, and $(7,7)$.

7   Concluding remarks

There is quite some literature on noncooperative games with only ordinal preferences: many economic games (for instance, Cournot or Bertrand oligopoly games) belong to this category, but also games used for implementing social choice correspondences, to name just a few examples. However, apart from Fishburn (1978) the only references to ordinal games with mixed strategies that we know of are Börgers (1993) and Rothe (1995). Börgers (1993) proposes a definition of rationalizability in which only ordinal preferences over outcomes are assumed to be common knowledge. Rothe (1995) considers equilibrium selection in 2×2-games under the first-degree stochastic dominance criterion.

There are several possible extensions of the present model that we list here without any implication about their interest: extension to more than 2 players; to non-strict preferences; to the players being characterized by different t's;

and to letting t vary continuously instead of in discrete steps. More generally, another extension is obtained by letting U_σ and U_τ denote classes of utility functions representing the ordinal preferences of players 1 and 2, and by considering (U_σ, U_τ)-equilibria, i.e., strategy combinations that are Nash equilibrium for at least one pair of utility functions from U_σ and U_τ.

A Remaining proofs

Proof of Lemma 2.2. The proof is by induction on t. For $t = 1$ the formula holds by definition of A. Let the formula be true for all $k < t$, where $t \geq 2$. Then for all $i, j \in \{1, \ldots, \ell\}$

$$a_{ij}^t = \sum_{l=1}^{\ell} a_{il}^{t-1} a_{lj} = \sum_{l=i}^{\ell} a_{il}^{t-1} a_{lj}, \tag{14}$$

where the second equality holds by induction. If $i > j$ then every $a_{lj} = 0$ in the RHS of (14) since $l > j$. If $i \leq j$ then (14) implies by induction

$$a_{ij}^t = \sum_{l=i}^{j} a_{il}^{t-1} = \sum_{l=i}^{j} \frac{(t-2+l-i)!}{(l-i)!(t-2)!}. \tag{15}$$

We are done if we can prove

$$\sum_{l=i}^{j} \frac{(t-2+l-i)!}{(l-i)!(t-2)!} = \frac{(t-1+j-i)!}{(j-i)!(t-1)!}.$$

We show this again by induction. For $j = i$ it is immediate. Let the equality hold for $i, \ldots, j-1$, then

$$\sum_{l=i}^{j} \frac{(t-2+l-i)!}{(l-i)!(t-2)!} = \frac{(t-1+j-1-i)!}{(j-1-i)!(t-1)!} + \frac{(t-2+j-i)!}{(j-i)!(t-2)!}$$

$$= \frac{(t-1+j-i-1)!(j-i) + (t-2+j-i)!(t-1)}{(j-i)!(t-1)!}$$

$$= \frac{(t-1+j-i)!}{(j-i)!(t-1)!}.$$

This completes the proof. $\qquad\square$

Proof of Lemma 2.1. The proof is by induction on t. For $t = 1$ the identity in the lemma holds by definition. Assume it holds for every $k < t$ $(t \geq 2)$. Let $l \in \{1, \ldots, \ell\}$. Then

$$F_r^t = F_r^{t-1} A = r A^{t-1} A = r A^t,$$

where the second equality follows by induction. □

Proof of Lemma 2.3. By Lemma 2.2,

$$
\begin{aligned}
a_{ij}^t / a_{i'j}^t &= \frac{(t-1+j-i)!}{(j-i)!(t-1)!} \cdot \frac{(j-i')!(t-1)!}{(t-1+j-i')!} \\
&= \frac{(j-i+1) \cdot (j-i+2) \cdots (j-i+t-1)}{(j-i'+1) \cdot (j-i'+2) \cdots (j-i'+t-1)} \\
&= \frac{(j-i'+t) \cdot (j-i'+t+1) \cdots (j-i+t-1)}{(j-i'+1) \cdot (j-i'+2) \cdots (j-i))},
\end{aligned}
$$

hence $a_{ij}^t \geq a_{i'j}^t$ and $\lim_{t \to \infty} a_{ij}^t / a_{i'j}^t = \infty$. □

Proof of Proposition 2.4 By Lemma 2.1, the probability distribution r is t-undominated if and only if

$$
\{x \in I\!R^\ell \mid x \leq rA^t\} \cap \{x \in I\!R^\ell \mid x = sA^t \text{ for some } s\} = \{rA^t\}.
$$

By a standard separation argument it follows that the two sets on the left-hand side of this identity can be separated by a hyperplane through rA^t with a nonnegative normal c such that $sA^t c \geq rA^t c$ for all probability distributions s on O. Since the second set is a polytope, this normal can be chosen positive (see Shapley, 1959, for a detailed argument). This implies $\sum_{l=1}^\ell r_l u_l \geq \sum_{l=1}^\ell s_l u_l$ for all probability distributions s on O, where $u := -A^t c \in U^t$. Also the converse of this argument holds, so that the proof of the lemma is complete. □

B Existence of t-equilibria

We provide an existence proof of t-equilibria by applying the Kakutani fixed point theorem to subsets of the t-best reply correspondences—recall from Section 4 that these correspondences themselves are not necessarily upper semi-continuous.

Consider the two-person game defined in Section 3. Fix a strategy $q \in \Delta^n$ for player 2. By Lemma 2.1 it follows that

$$
F_{pq\sigma}^t = \sum_{i=1}^m p_i F_{e^i q\sigma}^t \tag{16}
$$

for every $p \in \Delta^m$. Hence

$$
\{F_{pq\sigma}^t \mid p \in \Delta^m\} = \text{conv}\{F_{e^i q\sigma}^t \mid i \in M\}, \tag{17}
$$

where 'conv' denotes 'the convex hull of'. The t-best replies against q are those strategies p that generate points on the 'southwest corner' of the polytope in

(17), i.e., each point for which there is no other point in the polytope that has every coordinate smaller or equal and, thus, t-degree stochastically dominates it. These undominated points do not necessarily lie in the same face of the polytope and therefore a convex combination of t-best replies against q is not necessarily a t-best reply itself against q. So t-best reply correspondences are not necessarily convex-valued, in contrast with Nash best reply correspondences. For applying a fixed point argument this lack of convexity is not a real problem since from (17) it is easily seen that the set of t-best replies is homeomorphic to a convex set. In the proof of the existence result we apply Kakutani to the correspondences obtained by minimizing sums of coordinates on the polytopes in (17).

Theorem B.1 $E^t \neq \emptyset$ for every $t \geq 1$.

Proof. Let $t \geq 1$. For every $q \in \Delta^n$ define

$$\mu_1(q) = \arg \min \left\{ \sum_{l=1}^{mn} F_{pq\sigma}^t(l) \mid p \in \Delta^m \right\}.$$

Then it is easy to verify that the correspondence $q \mapsto \mu_1(q)$ is nonempty- and convex-valued, and upper semi-continuous. For player 2, we define an analogous correspondence $\Delta^m \ni p \mapsto \mu_2(p)$ with the same properties. Now the correspondence

$$\mu : \Delta^m \times \Delta^n \ni (p, q) \mapsto \mu_1(q) \times \mu_2(p)$$

satisfies all conditions needed to apply the Kakutani fixed point theorem. Hence, there exists a pair (p^*, q^*) with $p^* \in \mu_1(q^*)$ and $q^* \in \mu_2(p^*)$. Since, for any p and q, all elements of $\mu_1(q)$ are t-best replies against q and all elements of $\mu_2(p)$ are t-best replies against p, this holds in particular for p^* and q^*. Hence, $(p^*, q^*) \in E^t$. $\square$

References

Börgers, T. (1993): "Pure Strategy Dominance," *Econometrica*, 61, 423–430.

Fishburn, P.C. (1976): "Continua of Stochastic Dominance Relations for Bounded Probability Distributions," *Journal of Mathematical Economics*, 3, 295–311.

Fishburn, P.C. (1978): "Non-Cooperative Stochastic Dominance Games," *International Journal of Game Theory*, 7, 51–61.

Fisburn, P.C. (1980): "Stochastic Dominance and Moments of Distributions," *Mathematics of Operations Research*, 5, 94–100.

Rothe, J. (1995): "Equilibrium Selection in Ordinal Games," mimeo, London School of Economics.

Shapley, L.S. (1959): "Equilibrium points in games with vector payoffs," *Naval Research Logistics Quarterly*, 6, 57–61