

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 84, No. 5 (September, 2016), 1881–1902

THE EFFECT OF CHANGES IN RISK ATTITUDE ON STRATEGIC BEHAVIOR

JONATHAN WEINSTEIN

Washington University in St. Louis, St Louis, MO 63130, U.S.A.

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

THE EFFECT OF CHANGES IN RISK ATTITUDE ON STRATEGIC BEHAVIOR

BY JONATHAN WEINSTEIN¹

We study families of normal-form games with fixed preferences over pure action profiles but varied preferences over lotteries. That is, we subject players' utilities to monotone but nonlinear transformations and examine changes in the rationalizable set and set of equilibria. Among our results: The rationalizable set always grows under concave transformations (risk aversion) and shrinks under convex transformations (risk love). The rationalizable set reaches an upper bound under extreme risk aversion, and lower bound under risk love, and both of these bounds are characterized by elimination processes. For generic two-player games, under extreme risk love or aversion, all Nash equilibria are close to pure and the limiting set of equilibria can be described using preferences over pure action profiles.

KEYWORDS: Risk attitude, game theory, rationalizability, equilibrium.

1. INTRODUCTION

THE STANDARD SPECIFICATION of a normal-form game provides utility functions for each player, which represent Von Neumann–Morgenstern (VN–M) preferences on lotteries over action profiles. When specifying a game, we may not in fact have full knowledge of players' preferences over lotteries, but may only know their preferences over pure action profiles. This is, in fact, the typical situation in experiments on games with monetary payoffs: it is safe enough to assume that players prefer more money to less, but we have little knowledge of their preferences over lotteries. This paper will study the impact of transformations that preserve preferences over action profiles (sometimes called *ordinal preferences*) but alter preferences over lotteries. Equivalently, we look at families of games generated by monotone but nonlinear transformations of players' utilities. Within each such family, we examine the varying predictions of each of rationalizability, Nash equilibrium, and correlated equilibrium.

Rationalizability, introduced by [Bernheim \(1984\)](#) and [Pearce \(1984\)](#), is a solution concept which characterizes the implications of common knowledge of rationality. A simple example suffices to show that the set of rationalizable actions can be affected by cardinal preferences over action profiles, even with

¹I am grateful for conversations and correspondence with Pierpaolo Battigalli, Tilman Borgers, Amanda Friedenberg, and Muhamet Yildiz, anonymous referees, and seminar participants at University of Arizona, Arizona State, University of Texas, Iowa State University, the Stony Brook International Game Theory Festival, and the World Congress of the Econometric Society. An earlier version had the working title “The Behavior of Solution Concepts Under Changes in Risk Attitude.”

ordinal preferences fixed:

	L	R
U	3, 0	0, 0
M	1, 0	1, 0
D	0, 0	3, 0

In this game, M is not rationalizable for the row player, whom we call Alice; it is dominated by an equal mixture of U and D . However, suppose we subject Alice's utilities to the monotone transformation $\bar{u} = -e^{-u}$. One interpretation of this transformation is that the payoffs above are monetary, and Alice has gone from being risk-neutral to risk-averse, with a constant absolute risk aversion (CARA) specification. The new game is

	L	R
U	-0.05, 0	-1, 0
M	-0.37, 0	-0.37, 0
D	-1, 0	-0.05, 0

and in this game, M is a best reply to a uniform belief over opposing actions and all actions are rationalizable. This example helps illustrate many aspects of the general theory of the behavior of rationalizability under monotone transformations, found in Section 3:

1. All actions (for both players) except M are best replies to *pure* actions by the opposing player. Indeed, all except M would survive an iterated procedure where only best replies to pure actions survive each round; as [Bernheim \(1984\)](#), we say that such actions are *point rationalizable*. This process is not sensitive to monotone transformations of either player's utility, so all point-rationalizable actions are rationalizable under all such transformations.

2. Imagine we had introduced the second game first. The action M is then rationalizable, but not point-rationalizable. Under a transformation to the first game, it is not rationalizable. This generalizes: for any action which is *not* point-rationalizable, there is a transformation which makes it non-rationalizable. This result, [Proposition 2](#), provides a complete characterization of the actions which are rationalizable under all monotone transformations.

3. In the first game, M is not dominated (even weakly) by any pure strategy. It turns out that, as discovered by [Borgers \(1993\)](#), this is sufficient (and generically necessary) for there to exist a transformation, for example, the second game, under which M is a best reply.

4. To get from the first to the second game, we applied a monotone concave transformation to Alice's utility. Equivalently, Alice became more risk-averse.² The rationalizable set grew. This generalizes: a monotone concave transformation always makes the rationalizable set (weakly) larger, and a convex transformation makes it smaller (Proposition 1). Loosely, compromise actions (such as M) can be more easily rationalized by more risk-averse players, while risky actions (such as U and D) can be rationalized for all levels of risk aversion because extreme beliefs are allowed.

5. The fact that we were able to use a CARA (exponential) transformation to make M rationalizable (or not) is also completely general (Propositions 2 and 3). Our results specify tight upper and lower bounds for the rationalizable set as the game undergoes monotone transformations, and shows that both bounds are achieved under CARA transformations, or indeed any family of transformations which is "unboundedly concave" (or convex) as defined later.

After covering the effect of monotone transformations on rationalizability (Section 3), we move on to Nash equilibria (Section 4). Here we show, for generic two-player games, that under extreme risk love or aversion, the NE set converges to a set of pure profiles, and we characterize these limiting sets. We also discuss the behavior of correlated equilibria (Section 5). In Sections 3–5, each transformation of the game is analyzed as a complete-information game, implicitly assuming that the new preferences are common knowledge among the players. In Section 6, we take the alternate approach of analyzing Bayesian games, under the assumption that preferences over pure profiles are common knowledge but preferences over lotteries are not.

2. PRELIMINARIES

A *finite complete-information normal-form game* (henceforth "game") is a triple $G = (N, A, u)$ where $N = \{1, \dots, n\}$ is a finite set of players, $A = \prod_{i \in N} A_i$ is a finite set of (pure) action profiles, and $u = (u_1, \dots, u_n)$ is a profile of VN–M utilities for each player over A , $u_i : A \rightarrow \mathbb{R}$. Utilities are extended linearly to lotteries over A as usual. Call a game G *generic* if we always have $u_i(a) \neq u_i(a')$ when $a \neq a'$. As usual, $A_{-i} = \prod_{j \neq i} A_j$; also, $\Delta(A_{-i})$ is the set of possible beliefs for player i over opposing action profiles, where correlated beliefs are allowed. Let $\Sigma_i = \Delta(A_i)$ be the mixed strategies of player i . For a probability distribution π , let $\text{supp}(\pi)$ denote the support of π . We call a product set $B = B_1 \times \dots \times B_n$, where $B_i \subseteq A_i$, a *block*, and the *block game* G_B is simply the restriction of G to action sets B_i .

²Note that while the concept of risk aversion is usually defined in the context of monetary payoffs, the concept of a preference relation being more risk-averse than another can be applied in the case of purely abstract outcomes. Suppose Alice and Anne have the same preferences over certain outcomes. We can say that Anne is more risk-averse than Alice if whenever Alice prefers a sure thing to a lottery, so does Anne. When their preferences have VN–M representations, this corresponds to Anne's utility being a monotone concave transformation of Alice's.

For any $C \subseteq A_{-i}$, $\Delta(C) \subseteq \Delta(A_{-i})$ are those beliefs assigning positive probability only to elements of C . We define (correlated) rationalizability by iterated elimination of never-best-replies:³

$$S_i^0 \equiv A_i,$$

$$S_i^n = \bigcup_{\sigma_{-i} \in \Delta(S_{-i}^{n-1})} \text{BR}_i(\sigma_{-i}), \quad n \geq 1,$$

$$S_i^\infty \equiv \bigcap_{n=0}^\infty S_i^n,$$

where $\text{BR}_i(\sigma_{-i})$ is the set of pure best replies to a belief $\sigma_{-i} \in \Delta(A_{-i})$. Here Pearce’s lemma tells us that this procedure is equivalent (step-by-step) to eliminating actions which are strictly dominated by some pure or mixed strategy.

We define P^k to be actions surviving k rounds of pure best reply; that is, in each round, only actions that are a best reply to a profile of remaining pure strategies survive. Symbolically,

$$P_i^0 \equiv A_i,$$

$$P_i^n = \bigcup_{a_{-i} \in P_{-i}^{n-1}} \text{BR}_i(a_{-i}), \quad n \geq 1,$$

$$P_i^\infty \equiv \bigcap_{n=0}^\infty P_i^n.$$

It is immediate that $P^n \subseteq S^n$ for all $n \in \mathbb{Z} \cup \infty$. As Bernheim (1984), we call P^∞ the *point-rationalizable* actions.

The following notion was introduced by Borgers (1993): An action $a_i \in A_i$ is *B-dominated*⁴ if for all non-empty subsets $T \subseteq A_{-i}$, there is an $a_i^T \in A_i$ such that for all $a_{-i} \in T$, $u(a_i^T, a_{-i}) \geq u(a_i, a_{-i})$, with strict inequality for some $a_{-i} \in T$. Being *B-dominated* is implied by strict dominance by a pure action, and implies weak dominance by a pure action (the case $T = A_{-i}$), but is not equivalent to either of these.⁵ Though a weakly dominated action can be a best reply, a *B-dominated* action cannot. To see this, let a_i be any *B-dominated* action. It cannot be a best reply to any belief $\sigma_{-i} \in \Delta(A_{-i})$, because if we let T be

³Bernheim and Pearce originally defined rationalizability as constrained to independent beliefs over opposing actions, a significant constraint when $n > 2$. In recent years, allowing correlation has become more popular, in part due to the arguments of Aumann (1987), and we do so throughout. Here, and elsewhere, we suppress the dependence of the solution concept on the game G when the intended game is clear in context, but include it, writing $S_i^\infty(G)$, when necessary.

⁴Borgers calls this simply “dominated.”

⁵For generic games as defined above, *B-dominance* is equivalent to dominance (strict or weak) by a pure action.

the support of σ_{-i} , a_i^T will outperform a_i against σ_{-i} . Let $B_i^k \subseteq A_i$ be actions surviving k rounds of iterated B -dominance, and $B_i^\infty \equiv \bigcap_{n=0}^\infty B_i^n$. Iterating the argument above shows that $S^n \subseteq B^n$ for all $n \in \mathbb{Z} \cup \infty$.

We write $G \sim G'$ if G and G' have identical player and action sets, and players have identical *ordinal preferences*, that is, identical preferences over pure action profiles. Now, we have argued earlier that $P^n(G) \subseteq S^n(G) \subseteq B^n(G)$, and it is immediate from the definitions that P^k and B^k only depend on ordinal preferences, that is, they are constant on each equivalence class under \sim . So, we have the sandwich relationship:⁶

$$(1) \quad \forall G' \sim G, \forall n \in \mathbb{Z} \cup \infty: \quad P^n(G) \subseteq S^n(G') \subseteq B^n(G).$$

Borgers showed that the upper bound is tight, in the sense that every action in $B^n(G)$ can be rationalized by some beliefs consistent with common knowledge of ordinal preferences. We will show, in Proposition 3, that it is tight in the slightly stronger sense that there is a specific complete-information game $G' \sim G$ such that $S^\infty(G') = B^\infty(G)$. The constructed G' can be described simply as any transformation of G in which players are sufficiently risk-averse. In a complementary fashion, Proposition 2 states that any G' in which players are sufficiently risk-loving achieves the lower bound $S^\infty(G') = P^\infty(G)$.

To formalize the ideas of increasing risk aversion or risk love, we introduce the following definitions:

Fixing $G = (N, A, u)$, let $G^r = (N, A, u^r)$, $r \in (0, \infty)$ be a family of games with $G^r \sim G$. We say that this family is *unboundedly convex* if:

1. For any $r > s$ and all $i \in N$, $u_i^r = f_{i,r,s} \circ u_i^s$ for some increasing, convex $f_{i,r,s}$.
2. For any $\pi, \pi' \in \Delta(A)$ with $\max_{a \in \text{supp}(\pi)} u_i(a) > \max_{a \in \text{supp}(\pi')} u_i(a)$, $u_i^r(\pi) > u_i^r(\pi')$ for all sufficiently large r .

Condition 2 is equivalent to the condition that for every $i \in N$ and $a, a', a'' \in A$ with $u_i(a) > u_i(a') > u_i(a'')$:

$$\lim_{r \rightarrow \infty} \frac{u_i^r(a') - u_i^r(a'')}{u_i^r(a) - u_i^r(a')} = 0.$$

In particular, letting $u_i^r(a) = e^{ru_i(a)}$ produces such a family. In words, Condition 2 says that for large enough r , preferences are dominated by the likelihood of the largest prize.

Similarly, we say that the family $G^{-r} = (N, A, u^{-r})$, $r \in (0, \infty)$ with $G^{-r} \sim G$ is *unboundedly concave* if:

⁶In fact, this relationship holds for a weaker notion of similarity, in which ordinal preferences are fixed conditional on each opposing profile. That is, let $(N, A, u) \approx (N, A, u')$ if $\forall i, a_i, a'_i, a_{-i} : u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i}) \Leftrightarrow u'_i(a_i, a_{-i}) > u'_i(a'_i, a_{-i})$. Then P^k and B^k are well-defined on equivalence classes under \approx , so the relationship (1), as well as Propositions 2 and 3, are still true with \approx replacing \sim .

1. For any $r > s$ and all $i \in N$, $u_i^{-r} = f_{i,r,s} \circ u_i^{-s}$ for some increasing, concave $f_{i,r,s}$.

2. For any $\pi, \pi' \in \Delta(A)$ with $\min_{a \in \text{supp}(\pi)} u_i(a) > \min_{a \in \text{supp}(\pi')} u_i(a)$, $u_i^{-r}(\pi) > u_i^{-r}(\pi')$ for all sufficiently large r .

In words, Condition 2 says that for large enough r , preferences are dominated by the smallest possible prize in each lottery, approaching a “max–min” criterion. Condition 2 is equivalent to the condition that for every $i \in N$ and $a, a', a'' \in A$ with $u_i(a) > u_i(a') > u_i(a'')$:

$$\lim_{r \rightarrow \infty} \frac{u_i^r(a) - u_i^r(a')}{u_i^r(a') - u_i^r(a'')} = 0.$$

In particular, letting $u_i^r(a) = -e^{-ru_i(a)}$ produces such a family. If the original game G has monetary payoffs, these utilities are of course well-known as representing Constant Absolute Risk Aversion (CARA) with parameter r . The convex case, $u_i^r(a) = e^{ru_i(a)}$, could correspondingly be termed constant absolute risk love. I suggest that the reader keep these examples in mind when interpreting the coming results.

Finally, the phrase “for all sufficiently large r ,” occasionally abbreviated to “for large r ,” means that there is an R such that the statement in question is true for all $r > R$.

3. BEHAVIOR OF RATIONALIZABILITY

We discussed the intuition for the following result in the [Introduction](#):

PROPOSITION 1: *Let $G = (N, A, u)$ be a finite complete-information normal-form game, and let $G' = (N, A, v)$ where $v_i = f_i \circ u_i$ for some strictly increasing, convex functions $f_i : \mathbb{R} \rightarrow \mathbb{R}$. Then $S^k(G') \subseteq S^k(G)$ for each $k \in \mathbb{N} \cup \{\infty\}$. The reverse inclusion holds for increasing, concave f_i .*

PROOF: The case $k = 0$ is trivial. Suppose the claim is true for k , but false for $k + 1$, so that we have $a_i \in S_i^{k+1}(G') - S_i^{k+1}(G)$. Since $S^{k+1}(G') \subseteq S^k(G') \subseteq S^k(G)$, we have $a_i \in S_i^k(G) - S_i^{k+1}(G)$, so by the equivalence of rationalizability with iterated dominance, there must be a mixture $\sigma_i \in \Delta(S_i^k(G))$ with $u_i(\sigma_i, a_{-i}) > u_i(a_i, a_{-i})$ for all $a_{-i} \in S_{-i}^k(G)$. Then convexity and monotonicity of f_i give

$$v_i(\sigma_i, a_{-i}) \geq f_i(u_i(\sigma_i, a_{-i})) > f_i(u_i(a_i, a_{-i})) = v_i(a_i, a_{-i})$$

for all $a_{-i} \in S_{-i}^k(G)$, a fortiori for all $a_{-i} \in S_{-i}^k(G')$. This contradicts $a_i \in S_i^{k+1}(G')$ and completes the induction step. The case $k = \infty$ follows immediately from the finite case.

The concave case follows from the fact that the inverse⁷ of an increasing, concave function is increasing and convex. Q.E.D.

In contemporaneous work, Battigalli, Cerreia-Vioglio, Maccheroni and Marinacci (2016) studied the effects of both risk attitude and ambiguity attitude on the set of justifiable actions in decision problems. Their main result is a more general form of Proposition 1, concerning a concave transformation of the DM's attitude towards ambiguity (in the smooth model) or risk, and including the case of infinite action spaces with appropriate continuity and compactness assumptions. They showed that concave transformations always expand the set of possible actions, with a similar proof relying on an extended version of Pearce's lemma which they established.

The next proposition shows that for extreme risk love, the elimination process for rationalizability becomes precisely the process for point-rationalizability. That is, in our notation, S^k becomes P^k . The proof will use the characterization of S^∞ via iterated dominance; our strategy is to show that for a very "risk-loving" player, a uniform mixture over all actions dominates any action which is not a best reply to any pure strategy. Viewed through the lens of best replies, the intuition is that extremely risk-loving players essentially use a "maximax" criterion: they will always "bet" on a single action profile by the other players and play a pure best reply to that profile.

PROPOSITION 2: *Let $G = (N, A, u)$ be a game and $G^r \sim G$ be any unboundedly convex family of transformations of G . Then:*

1. *As $r \rightarrow \infty$, $S^\infty(G^r)$ shrinks monotonically to $P^\infty(G)$.*
- 2.

$$\bigcap_{G^r \sim G} S^\infty(G^r) = P^\infty(G).$$

PROOF: Monotonicity is immediate from Proposition 1. We will show that for sufficiently large r , $S^k(G^r) = P^k(G)$ for all $k \in \mathbb{N} \cup \{\infty\}$. Fix an action $a_i \notin P_i^\infty(G)$ and suppose it is eliminated at step $k + 1$, that is, $a_i \in P_i^k(G) - P_i^{k+1}(G)$. This means that for every $a_{-i} \in P_{-i}^k(G)$, there is an action $a_i(a_{-i}) \in P_i^k(G)$ with $u_i(a_i(a_{-i}), a_{-i}) > u_i(a_i, a_{-i})$. Let σ_i be a mixture putting equal weight on each action in $P_i^k(G)$. Then by the definition of unbounded convexity, for each $a_{-i} \in P_{-i}^k(G)$ we have, for sufficiently large r , $u_i^r(\sigma_i, a_{-i}) > u_i^r(a_i, a_{-i})$. Since the game is finite, we can find r large enough for every such $a_{-i} \in P_{-i}^k(G)$. Again by finiteness, for r large enough, all $a_i \notin P_i^\infty(G)$ will be dominated at the appropriate stage. We have just showed that for sufficiently large r , only pure best replies survive at each stage of the rationalizability process in G^r , which is the claim $S^k(G^r) = P^k(G)$ for finite k . The case $k = \infty$ follows immediately.

⁷The inverse may not be defined on all of \mathbb{R} , but will be defined on the full interval of payoffs in our finite game, and this suffices.

In part 2, the inclusion $P^\infty(G) \subseteq \bigcap_{G' \sim G} S^\infty(G')$ is immediate, since best replies to pure action profiles are preserved under monotone transformations, and the opposite inclusion is implied by part 1. *Q.E.D.*

Next, we establish the upper bound on the actions that may be rationalizable in some transformed game. As mentioned earlier, this upper bound is achieved in the limit as the players become extremely risk-averse.⁸

PROPOSITION 3: *Let $G = (N, A, u)$ be a game and $G^{-r} \sim G$ be any unboundedly concave family of transformations of G . Then:*

1. *As $r \rightarrow \infty$, $S^\infty(G^{-r})$ grows monotonically to $B^\infty(G)$.*
- 2.

$$\bigcup_{G' \sim G} S^\infty(G') = B^\infty(G).$$

PROOF: Monotonicity is immediate from Proposition 1. We established the inclusion $S^\infty(G') \subseteq B^\infty(G)$ for any $G' \sim G$ in the preliminaries. Since the game is finite, we must have some k such that $B^k(G) = B^\infty(G)$, so that $B^\infty(G)$ is a fixed point of the B -elimination process. We will show that for some r , $B^\infty(G)$ is also a fixed point of the process defining $S^\infty(G^{-r})$, which implies $B^\infty(G) \subseteq S^\infty(G^{-r})$ and hence the desired results.

Suppose to the contrary, that $B^\infty(G)$ is not a fixed point for any r . Let \tilde{G}^{-r} be the game with payoffs as in G^{-r} and action sets restricted to those in $B^\infty(G)$, so that no action is B -dominated in \tilde{G}^{-r} . By monotonicity, there must be an action $\hat{a}_i \in B_i^\infty(G)$ which is never a best reply in \tilde{G}^{-r} for any r , no matter how large. Applying Pearce’s lemma, for each $r > 0$ let π^r be a mixture on $B_i^\infty(G)$ which strictly dominates \hat{a}_i in \tilde{G}^{-r} . The fact that \hat{a}_i is not B -dominated in \tilde{G}^{-r} means that there is a non-empty $T \subseteq B_{-i}^\infty(G)$ such that every $a'_i \in B_i^\infty(G)$ either satisfies (i) $u(a'_i, a_{-i}) = u(\hat{a}_i, a_{-i})$ for all $a_{-i} \in T$ or (ii) $u(a'_i, a_{-i}) < u(\hat{a}_i, a_{-i})$ for some $a_{-i} \in T$. By strict dominance, the mixture π_r cannot have only actions of type (i); by throwing out the actions of type (i) and renormalizing, we can obtain for each r a mixture ϕ^r which assigns positive mass only to actions a'_i satisfying (ii) and satisfies $u_i^{-r}(\phi^r, a_{-i}) > u_i^{-r}(\hat{a}_i, a_{-i})$ for all $a_{-i} \in T$.

Then, fixing any $\varepsilon \in (0, 1/|A_i|)$, there must be some $a'_i \in B_i^\infty(G)$ and $a_{-i} \in T$ for which $\phi^r(a'_i) > \varepsilon$ for an unbounded set of r and $u_i(a'_i, a_{-i}) < u_i(\hat{a}_i, a_{-i})$. Let \bar{a} be player i ’s most preferred outcome, and let μ be the lottery with $\mu(\bar{a}) = 1 - \varepsilon$, $\mu(a'_i, a_{-i}) = \varepsilon$. Clearly, $u_i^{-r}(\phi^r, a_{-i}) < u_i^{-r}(\mu)$ for the unbounded set of r where $\phi^r(a'_i) > \varepsilon$. Also, by definition of unbounded concavity, for all sufficiently large r we have $u_i^{-r}(\mu) < u_i^{-r}(\hat{a}_i, a_{-i})$. Then for some r , $u_i^{-r}(\phi^r, a_{-i}) < u_i^{-r}(\hat{a}_i, a_{-i})$, contradicting the definition of ϕ^r . *Q.E.D.*

⁸The $n = 1$ version of this result, without mention of monotonicity, was discovered in contemporaneous work by Gafarov and Salcedo (2014).

As discussed in the previous section, this result includes the main result of [Borgers \(1993\)](#). Our proof is made more straightforward by focusing on a family of concave transformations and exploiting Pearce's lemma, while Borgers explicitly constructed a transformation and belief to make an arbitrary action in B^∞ a best reply.

4. BEHAVIOR OF NASH EQUILIBRIUM

The next natural line of inquiry concerns the behavior of the set of Nash equilibria as we apply monotone transformations to a game. Here also, there are some striking regularities when we look at a family of convex or concave transformations (such as the exponential family). Restricting attention to generic two-player games, we will be able to state a precise description of the limiting set of equilibria under extreme risk love (Proposition 4) and extreme risk aversion (Proposition 5). Please note that while our results show existence of these limits, if we allow preferences to go to the extreme of infinite risk aversion (max–min preferences), then Nash equilibria may not exist. Propositions 4 and 5 do not hold for more than two players; counterexamples are at the end of this section.

In order to pin down the limiting behavior of the set of Nash equilibria under extreme risk love or aversion, we will ask which “blocks” $S \times T$ are possible supports of Nash equilibria, where $S \subseteq A_1$, $T \subseteq A_2$. For the case of risk love, the key will be the following:

DEFINITION 1: A block $S \times T$ has the *point best-reply property* if every action in S is a best reply (in the full game) to some action in T and vice versa.

In generic games, the point best-reply property implies that $|S| = |T|$; indeed, the players' best-reply correspondences must give (possibly distinct) bijections between S and T . Notice that $P^\infty(G)$ will always have the point best-reply property, and a singleton block has the property if and only if it is a pure NE.

PROPOSITION 4: Fix a generic two-player game G , an unboundedly convex family of transformations G^r , and subsets $S \subseteq A_1$, $T \subseteq A_2$:

1. If $S \times T$ has the point best-reply property, then:
 - (a) For all sufficiently large r , there is a NE of G^r with support $S \times T$.
 - (b) As $r \rightarrow \infty$, the equilibria (σ_1^r, σ_2^r) of G^r with support $S \times T$ converge to the pure profile $(\tilde{a}_1, \tilde{a}_2)$ where

$$\tilde{a}_1 = \operatorname{argmin}_{a_1 \in S} \left[\max_{a_2 \in T} u_2(a_1, a_2) \right],$$

$$\tilde{a}_2 = \operatorname{argmin}_{a_2 \in T} \left[\max_{a_1 \in S} u_1(a_1, a_2) \right].$$

2. If $S \times T$ lacks the point best-reply property, then for all sufficiently large r , there is no NE of G^r with support $S \times T$.

The second item follows from the same argument as the proof of Proposition 2; indeed, if an a_1 is not a best reply to any $a_2 \in T$, then for large r it is not a best reply in G^r to any belief over T . The first item, existence, is shown with the aid of a general lemma on existence of Nash equilibrium with given supports. An intuition for our identification of the limiting equilibria is that in the limit, each player is mainly interested in chasing the largest payoff possible, so to make (say) Player 2 indifferent, Player 1 must offer a very high probability that the least enticing of his best responses is optimal. All of these results are restricted to the two-player case. At the end of this section, see counterexamples for $n = 3$ and a discussion.

Keeping our focus on generic two-player games, a similar result obtains for extreme risk aversion, with a slightly more complicated statement. It turns out that for extremely risk-averse players, what kills an equilibrium on $S \times T$ is the presence of an action a_1 which is not the *worst* reply among $S \cup \{a_1\}$ to any $a_2 \in T$ (or the analogue with roles reversed). Indeed, one direction of our proof verifies that for a sufficiently risk-averse player, such an a_1 dominates a uniform mixture over S when Player 2 is restricted to T , preventing any full-support mixture on S from being a best reply. Accordingly, we make the following definitions:

DEFINITION 2: A *never-worst deviation* from a block $S \times T$ is an $a_1 \in A_1$ such that

$$\forall a_2 \in T : \exists a'_1 \in S : u_1(a_1, a_2) > u_1(a'_1, a_2)$$

or an a_2 with roles reversed. We say $S \times T$ *allows no safety* if there is no never-worst deviation from $S \times T$.

Notice that an $a_1 \in S$ is a never-worst deviation from $S \times T$ exactly when it is not a point worst reply in the block game $G_{S \times T}$. In particular, under genericity, if $|S| > |T|$, there will be a never-worst deviation $a_1 \in S$. Therefore, when the block $S \times T$ allows no safety, we must have $|S| = |T|$, and the “worst-reply correspondences” of the block game must each form a bijection between S and T . Observe also that a singleton block allows no safety exactly when it is a pure NE.

PROPOSITION 5: Fix a generic two-player game G , an unboundedly concave family of transformations G^{-r} , and subsets $S \subseteq A_1, T \subseteq A_2$:

1. If $S \times T$ allows no safety, then:
 - (a) For all sufficiently large r , a NE of G^{-r} with support $S \times T$ exists.

(b) As $r \rightarrow \infty$, the equilibria $(\sigma_1^{-r}, \sigma_2^{-r})$ of G^{-r} with supports S, T converge to $(\tilde{a}_1, \tilde{a}_2)$ where

$$\tilde{a}_1 := \operatorname{argmax}_{a_1 \in S} \left[\min_{a_2 \in T} u_2(a_1, a_2) \right],$$

$$\tilde{a}_2 := \operatorname{argmax}_{a_2 \in T} \left[\min_{a_1 \in S} u_1(a_1, a_2) \right].$$

2. Otherwise, for all sufficiently large r , no NE of G^{-r} with support $S \times T$ exists.

Let $\text{NE}(G)$ be the set of NE profiles of G , and apply the Hausdorff metric⁹ to compact sets of (mixed) profiles. From Propositions 4 and 5, we can immediately deduce the following:

PROPOSITION 6: For every generic two-player game G , the sets $\text{NE}(G^r)$, $\text{NE}(G^{-r})$ each converge to a set of pure profiles as $r \rightarrow \infty$. The limiting sets are the same for all $G' \sim G$ and do not depend on the particular families of transformations G^r, G^{-r} .

That is, for extreme risk love or extreme risk aversion, the set of equilibria approaches a set of extreme points in the space of mixed strategies. These sets can easily be found by applying Propositions 4 and 5: first check each square block $S \times T$ to see if it has an equilibrium in the limit; then each proposition tells us which extreme point is the limiting equilibrium in this block. The criteria determining these limits depend only on ordinal data, that is, on preferences over pure action profiles.

Since pure NE are preserved by monotone transformations, we know that in a game with no pure NE, the limit points will not be pure NE, as in this example, a generic form of Matching Pennies:

	L	R
U	2, -2	-1, 1
D	-2, 2	1, -1

This generic 2×2 game has no pure equilibria, and one mixed, for any monotone transformation. The limiting NE are $\text{NE}(G^r) \rightarrow \{(U, R)\}$, $\text{NE}(G^{-r}) \rightarrow \{(D, R)\}$ as $r \rightarrow \infty$. Note that the “limiting preference” of extreme risk aversion would be max–min decision-makers, players whose utility is simply that of the worst outcome with positive probability. These preferences are discontinuous, and are easily seen in this example to admit no equilibrium, because

⁹The Hausdorff distance between two sets is the supremum of the distances from each point in each set to the nearest point in the other set.

no mixture makes either player indifferent.¹⁰ The limit of $\text{NE}(G^{-r})$ is a good candidate for a solution concept for this case.

The limiting case of extreme risk love, decision-makers whose utility equals their maximum possible payoff, or “wishful thinkers,” was studied in Yildiz (2007). He found that in examples such as this one (generic two-player games with no pure NE), even the analogue of rationalizability is empty: no profile is consistent with a common prior and common knowledge of wishful thinking. Equilibrium then does not exist either.¹¹ The limit of $\text{NE}(G^r)$ is a good candidate for a solution concept in this case.

The fact that equilibria converge to a non-equilibrium may appear to violate upper hemicontinuity¹² of the equilibrium correspondence. It does not, because payoffs are unbounded and the games G^r do not have a convergent subsequence. If we normalized payoffs to have range $[0, 1]$ after each transformation, then upper hemicontinuity would apply: we would find that for each ε , (U, R) is a ε -equilibrium of the normalized version of G^r for large enough r , and the payoff matrix would converge to

	L	R
U	1, 0	0, 0
D	0, 1	0, 0

where (U, R) is an equilibrium. Note, though, that the other pure NE of this limit game, (U, L) , is not a limit point of equilibria of G^r , and hence one cannot identify the limiting set simply by finding NE of the limit game; this is a standard failure of lower hemicontinuity of the equilibrium correspondence.

If pure NE exist, of course they will be among the limit points, but there may be others, as in this example:

	L	R
U	3, 1	0, 0
D	2, 2	1, 3

This generic 2×2 game, G , has two pure equilibria, (U, L) and (D, R) , and one mixed. Any monotone transformation, of course, preserves the pure equi-

¹⁰For more on equilibrium existence under non-standard preferences, see, for instance, Azrieli and Teper (2011).

¹¹This presumes that we define equilibrium as assigning positive probability only to utility-maximizing actions. If we instead considered all mixtures which maximize utility according to the wishful-thinking criterion, it is easy to see that all completely mixed profiles are equilibria in any game, because each player then achieves his maximal utility. This is then a useless concept.

¹²Upper hemicontinuity tells us that the graph of the correspondence from games to NE is closed, or in other words, that given a convergent sequence of NE of a convergent sequence of games, the limit of the NE is a NE of the limiting game.

libria and also has one mixed equilibrium. $NE(G^r) \rightarrow \{(U, L), (D, R), (U, R)\}$, $NE(G^{-r}) \rightarrow \{(U, L), (D, R), (D, L)\}$ as $r \rightarrow \infty$. Proposition 4 tells us that in the limit of extreme risk love, in any full-support equilibrium, each player almost always chooses the action which gives the opponent the lowest possible maximum (i.e., the minmaxing action if only pure actions are allowed). Conversely, in the limit of extreme risk aversion, the weight goes to the action which maximizes the opponent's minimum payoff.

The proofs of these results take advantage of the restriction to $n = 2$, which simplifies the construction of equilibria considerably. In the two-player case, once the support of a NE is proposed, the equilibrium conditions become linear inequalities, subject to standard techniques including duality. For $n > 2$, constructing equilibria with given support is harder because it involves non-linear conditions. Indeed, the results fail in the following examples. Here is a three-player counterexample to Proposition 4: a generic three-player game where every action is a point best reply, but there is no completely mixed equilibrium in *any* monotone convex transformation:

In	<i>L</i>	<i>R</i>	Out	<i>L</i>	<i>R</i>
<i>U</i>	4, 1, 3	2, 3, 1	<i>U</i>	8, 5, 7	6, 7, 8
<i>D</i>	1, 4, 5	3, 2, 2	<i>D</i>	5, 8, 4	7, 6, 6

In fact, any such transformation has a unique NE, in which Player 3 chooses Out for sure and the other players mix. To see this, first check that Players 1 and 2 must mix in any equilibrium (for instance, if P1 plays *U*, P2 must play *R*, whence P1 plays *D*, etc.). Next, P1 must play *U* at least half the time; if not, P2 strictly prefers *L* in both matrices, because, for instance, $pf(1) + (1-p)f(4) > pf(3) + (1-p)f(2)$ for any $p < 0.5$ and convex f . Similarly, P2 must play *R* at least half the time. It is now easy to check that P3 strictly prefers Out.

For a corresponding counterexample for Proposition 5, simply negate all payoffs. Now every action is a point worst reply, so the full block allows no safety. There can be no fully mixed equilibrium in any concave transformation, since negating would give a convex transformation of the diagrammed game, and fully mixed equilibria are preserved by negating all payoffs (since players remain indifferent to all actions).

Also, to show that Proposition 6 fails when $n > 2$, here is a generic three-player game which, under the transformations $u \mapsto e^{ru}$, has NE which do not approach a pure profile:

In	<i>L</i>	<i>R</i>	Out	<i>L</i>	<i>R</i>
<i>U</i>	4, 6, 6	1, 5, 2	<i>U</i>	8, 2, 5	5, 1, 1
<i>D</i>	3, 7, 7	2, 8, 3	<i>D</i>	7, 3, 8	6, 4, 4

Here, in G^r , any profile where $p(U) = e^{2r}/(e^{2r} + 1)$, $p(L) = 1/(e^{2r} + 1)$ is an equilibrium, for any strategy of P3. One reaction to this example is that a

stronger notion of genericity might restore convergence to pure profiles; for instance, we are capitalizing on the fact that P1 and P2's best replies are not affected by P3's action, in the original game or any exponential transformations. Notice, though, that many perturbations of the payoffs in the matrix could be undone by a different family of convex transformations, so any successful strengthening of genericity would have to include a restriction on the transformations. Note also that neither of these examples is a counterexample to the related conjecture on CE stated at the end of Section 5, which to my knowledge may be true for all n .

5. CORRELATED EQUILIBRIA AND JOINT COHERENCE

Let $J(G)$ denote the set of strategy profiles which are played with positive probability in some correlated equilibrium (CE) of G . By [Aumann \(1987\)](#), we know that these are the profiles which are consistent with common knowledge of rationality together with a common prior. The added power of the common-prior assumption gives a refinement of rationalizability; we will shortly see an example where the refinement is strict. [Nau and McCardle \(1990\)](#) called $J(G)$ the “jointly coherent” profiles and showed that players who do not always play profiles in $J(G)$ are susceptible to a form of arbitrage against the group of players: an outside observer can find a set of bets, each accepted by some player, which gives him positive payoff whenever an element of $A - J(G)$ is played and zero payoff on $J(G)$. Their result is another argument for $J(G)$ as a natural notion of joint rationality. Unlike rationalizability, joint coherence may grow under monotone convex transformations of the players' utility functions, as the following example shows. We call this game “Matching Pennies with Entry Fee”:

	H	T	S
H	10, -10	-10, 10	0, 1
T	-10, 10	10, -10	0, 1
S	1, 0	1, 0	1, 1

The intended interpretation of the game is that two players walk into a casino with \$1 each, and may either play Safe and keep their dollar, or each pay the casino \$1 for the privilege of playing Matching Pennies against each other for a \$10 stake. Although all strategies are rationalizable here, the only jointly coherent profile is (S, S) . That is, the unique CE puts probability 1 on (S, S) . This game gives an economic intuition for the impact of the common-prior assumption which is embedded in equilibrium concepts: players with a common prior

cannot *both* believe they will win the zero-sum block game on $\{H, T\} \times \{H, T\}$, an argument similar to those for no-trade theorems.¹³

However, consider the game G^1 formed by applying the monotone convex transformation $f(u) = e^u$ to both players' utilities. There is a Nash equilibrium of G^1 in which Alice mixes equally U and D and Bob mixes equally L and R , showing that $J(G) \subsetneq J(G^1)$. Indeed, risk-loving players may, even in equilibrium, pay a casino a premium to play zero-sum games such as poker against one another. In fact, $J(G^1) = A$, that is, all profiles are jointly coherent.¹⁴ Though this example is non-generic, its behavior is robust to small payoff perturbations. For generic games, Proposition 7 generalizes the example. It says that for extreme risk love, joint coherence and rationalizability coincide, which we can interpret epistemically to mean that the common-prior assumption has no bite. An intuitive justification is that extreme risk love makes one a "wishful thinker" who chases the highest possible payoff; hence a restriction to identical beliefs becomes irrelevant—only the support of beliefs matters.

PROPOSITION 7: *For any generic two-player game G and unboundedly convex transformations G^r , for all sufficiently large r we have $J(G^r) = S^\infty(G^r) = P^\infty(G^r)$.*

PROOF: The latter equality is exactly Proposition 2. Now apply Proposition 4 to the block $P^\infty(G^r)$ of G^r , which has the point best-reply property. It guarantees an equilibrium σ of G^r with support $P^\infty(G^r)$ for large r . *Q.E.D.*

The proof only uses Nash equilibria, so in the proposition the extra freedom of CE over NE in the definition of $J(G)$ was not actually needed.

Now, because the set of CE is convex, we cannot hope for Proposition 6 to extend perfectly to correlated equilibria. Any time there are multiple pure NEs, for instance, the set of CEs will include the convex hull of these NEs for all r . However, we do conjecture, again for generic G , that the polytope of CEs of G^r , or G^{-r} , converges in Hausdorff topology to a convex hull of pure profiles as $r \rightarrow \infty$. This conjecture is not resolved at this time. An equivalent statement of this conjecture is that for each profile a , the maximum weight assigned to a in any CE of G^r goes to 0 or 1 as $r \rightarrow \infty$.

¹³This example has the same properties as one used in [Bernheim \(1984\)](#) to demonstrate that some rationalizable profiles are not played in any Nash equilibrium. I transformed the payoffs to give the game a better story. [Nau and McCardle](#) noted the uniqueness of CE as well as NE in [Bernheim's](#) example.

¹⁴The full-support distribution p with $p(M, L) = p(M, R) = p(U, C) = p(D, C) = \varepsilon^2$, $p(M, C) = \varepsilon$, and $p(U, L) = p(U, R) = p(D, L) = p(D, R) = 0.25 - \varepsilon/4 - \varepsilon^2$ is a CE for small enough $\varepsilon > 0$.

6. INCOMPLETE INFORMATION ON CARDINAL PAYOFFS

To this point, we have been analyzing complete-information games. The natural interpretation of the complete-information approach is that the exact preferences are commonly known to all the players, but not to a modeler who tries to predict their behavior knowing only preferences over pure action profiles. Perhaps a more natural assumption would be that the players themselves know only each other’s preferences over action profiles and not over lotteries. We show in this section that this generalization to incomplete-information games does not affect the upper or lower bounds on the rationalizable set, but does cause violations of Proposition 1 (monotonicity with respect to convex transformations).

To address this point, we will need to introduce *monotone elaborations* of a complete-information game. These are Bayesian games in which preferences over pure action profiles are fixed. Formally, the following holds:

DEFINITION: A (finite) Bayesian game is a tuple $(N, A, \hat{u}, \Theta, T, \kappa)$ where N is the set of players, $A = \prod_{i \in N} A_i$ is a finite set of (pure) action profiles, $u = (u_1, \dots, u_n)$ is a profile of VN–M utilities for each player over $A \times \Theta$, Θ is the finite set of payoff parameters, $T = \prod_{i \in N} T_i$ is the finite set of type profiles, and κ maps each type t_i to beliefs on $\Theta \times T_{-i}$.

DEFINITION: Given a complete-information game $G = (N, A, u)$, a Bayesian game $\hat{G} = (N, A, \hat{u}, \Theta, T, \kappa)$ is a *monotone elaboration* of G if, for all $i \in N$, $a, a' \in A$, and $\theta \in \Theta$,

$$u_i(a) > u_i(a') \iff \hat{u}_i(a, \theta) > \hat{u}_i(a', \theta).$$

Our solution concept for incomplete-information games will be interim correlated rationalizability (ICR), as introduced in Dekel, Fudenberg and Morris (2007): To construct the ICR sets for a Bayesian game $(N, A, \hat{u}, \Theta, T, \kappa)$, first let $ICR_i^0(t_i) = A_i$ for each i and t_i . Then define $ICR_i^n(t_i)$ iteratively by letting the set of “admissible conjectures” $C_n(t_i)$ be elements $\nu \in \Delta(\Theta \times T_{-i} \times A_{-i})$ satisfying the conditions:

1.

$$\text{marg}_{\Theta \times T_{-i}} \nu = \kappa_{t_i},$$

2.

$$\nu(\{(\theta, t_{-i}, a_{-i}) : a_{-i} \in ICR_{-i}^{n-1}(t_{-i})\}) = 1,$$

and then letting

$$ICR_i^n(t_i) = \bigcup_{\nu \in C_n(t_i)} \text{BR}_i(\text{marg}_{\Theta \times A_{-i}} \nu).$$

The following proposition states that the bounds on the rationalizable set are the same in monotone elaborations as in complete-information games.

PROPOSITION 8: *For any monotone elaboration \hat{G} of G and every type t_i of every player i , $P_i^\infty(G) \subseteq \text{ICR}_{\hat{G}}(t_i) \subseteq B_i^\infty(G)$.*

However, the following example shows that, in contrast with Proposition 1, it is possible for the ICR set of a type to grow under a convex transformation. Here, each player has just one type, who believes states θ_1 and θ_2 are equally likely:

θ_1	L	R	θ_2	L	R
U	2, 0	-1, 0	U	10.01, 0	9.9, 0
M	0, 0	0, 0	M	10, 0	10, 0
D	-1, 0	2, 0	D	9.9, 0	10.01, 0

On any belief over $\Theta \times A_2$ which has uniform marginal over Θ , M scores 5 while a 50–50 mixture of U and D scores about 5.2. Hence, M is not ICR. However, if we transform the game by sending all utilities from u to e^u , then M is a best response to the belief that Player 2 randomizes uniformly and independently of θ . This can be explained by noting two distinct effects of the transformation. Conditional on each θ , it has the usual effect of increasing risk love, and indeed, conditional on θ_1 , M is still dominated by a mixture of U and D . It also, though, has the effect of increasing the importance of the θ_2 case compared to the θ_1 case, and here this is the decisive effect; conditional on θ_2 , M is a best reply to a uniform belief, both before and after the transformation. The proof of Proposition 1 breaks down because dominance no longer only involves comparing a lottery to a certain outcome; pure-action profiles do not correspond to certain outcomes.

On the other hand, if we sent each u to e^{ru} for very large r , M would be dominated conditional on θ_2 as well as θ_1 , and indeed would not be rationalizable. The limiting behavior is as in the complete-information case, though monotonicity has failed. The regular behavior of the limiting case generalizes. We extend the concepts of unboundedly convex/concave families of transformations to Bayesian games by stipulating that conditional on each θ , the transformed payoffs are an unboundedly convex/concave family in the sense previously defined for complete-information games. (The transformation may be different for each θ .) Then we find that versions of Propositions 2 and 3 hold:

PROPOSITION 9: *Let \hat{G} be any monotone elaboration of G and let \hat{G}^r, \hat{G}^{-r} be unboundedly convex/concave families of transformations of \hat{G} . Then for all sufficiently large r , for each type t_i of each player i , $\text{ICR}_{\hat{G}^r}(t_i) = P_i^\infty(G)$ and $\text{ICR}_{\hat{G}^{-r}}(t_i) = B_i^\infty(G)$.*

The proofs of Propositions 8 and 9 are found in Appendix B.

APPENDIX A

Additional notation: Given a subset T of A_{-i} , write $\sigma_i >_T \sigma'_i$ if $u_i(\sigma_i, a_{-i}) \geq u_i(\sigma'_i, a_{-i})$ for all $a_{-i} \in T$, with strict inequality for some $a_{-i} \in T$, that is, if σ_i weakly dominates σ'_i when opposing profiles are restricted to T .

The following lemma, used in proofs of Propositions 4 and 5, is Corollary 2 in Weinstein (2015):

LEMMA 1: *A two-player game has a Nash equilibrium with supports $S \subseteq A_1, T \subseteq A_2$, if and only if there is no pair $\sigma_1 \in \Delta(A_1), \sigma'_1 \in \Delta(S)$ such that $\sigma_1 >_T \sigma'_1$ and no pair $\sigma_2 \in \Delta(A_2), \sigma'_2 \in \Delta(T)$ such that $\sigma_2 >_S \sigma'_2$. When such pairs do exist, they can be chosen with disjoint support, that is, $\sigma_1(a_1)\sigma'_1(a_1) = 0$ for all a_1 .*

No analogous lemma holds when $n > 2$; the proof relies on the fact that fixing S and T , the problem of finding an equilibrium with support $S \times T$ is a linear program with probabilities as variables. When $n > 2$, the problem is nonlinear in probabilities and the technique breaks down. Our proofs use $n = 2$ in other places as well. Counterexamples to Propositions 4 and 5 for $n > 2$ were given in Section 4.

PROOF OF PROPOSITION 4:

Claim 1(a): Suppose $S \times T$ has the point best-reply property. We show that for sufficiently large r , the condition following “if and only if” in Lemma 1 holds in G^r , whence the lemma yields an equilibrium. Fix any $a_1 \in S$; it is a best response to some $a_2 \in T$. Let a'_1 be the second-best response to a_2 , and a''_1 the worst response.¹⁵ Fix $\varepsilon \in (0, 1/|S|)$. By definition of unboundedly convex, for all sufficiently large r we have

$$(2) \quad \varepsilon u_1^r(a_1, a_2) + (1 - \varepsilon)u_1^r(a''_1, a_2) > u_1^r(a'_1, a_2).$$

Suppose $\sigma_1 >_T \sigma'_1$ with $\sigma_1 \in \Delta(A_1), \sigma'_1 \in \Delta(S)$ having disjoint support, as in the lemma. Suppose $\sigma'_1(a_1) \geq \varepsilon$. Disjoint support implies $\sigma_1(a_1) = 0$, so we have

$$u_1^r(\sigma'_1, a_2) \geq \varepsilon u_1^r(a_1, a_2) + (1 - \varepsilon)u_1^r(a''_1, a_2) > u_1^r(a'_1, a_2) \geq u_1^r(\sigma_1, a_2),$$

contradicting weak dominance. So $\sigma'_1(a_1) < \varepsilon$. Letting r be large enough that (2) holds for all $a_1 \in S, \sigma'_1(a_1) < \varepsilon < 1/|S|$ for all a_1 , so $\sum_{a_1 \in S} \sigma'_1(a_1) < 1$, a contradiction. Also take r large enough for the analogous conditions to hold with roles reversed, and we have proven the existence for such r of an equilibrium with support $S \times T$.

¹⁵If $|A_1| = 1$, the result is trivial: S, T must be singletons and form a pure-strategy NE.

Claim 1(b): Let $k = |S| = |T|$. Fix $\varepsilon > 0$, define

$$\tilde{a}_2 = \operatorname{argmin}_{a_2 \in T} \left[\max_{a_1 \in S} u_1(a_1, a_2) \right]$$

as in the proposition, and let \bar{a}_1 be the best reply to \tilde{a}_2 . Consider the set $P_\varepsilon := \{\sigma_2 \in \Delta(A_2) : \sigma_2(\tilde{a}_2) \leq 1 - \varepsilon\}$. We *claim* that for all sufficiently large r , in G^r action \bar{a}_1 is not a best reply to any belief in P_ε . To prove this secondary claim: Viewed as a subset of \mathbb{R}^k , P_ε is bounded and defined by finitely many linear inequalities, hence P_ε is a polytope and is the convex hull of its finite set of extreme points. These extreme points of S are¹⁶ the $k - 1$ points where $\sigma_2(\hat{a}_2) = 1$, for some $\hat{a}_2 \neq \tilde{a}_2$, and the $k - 1$ points where $\sigma_2(\hat{a}_2) = \varepsilon$, $\sigma_2(\tilde{a}_2) = 1 - \varepsilon$, for some $\hat{a}_2 \neq \tilde{a}_2$. Consider the mixed strategy σ_1 with $\sigma_1(\hat{a}_1) = 1/(k - 1)$ for each of the $k - 1$ actions $\hat{a}_1 \neq \bar{a}_1$. Let $\hat{\sigma}_2 \in P_\varepsilon$ be one of the extreme points of P_ε , so $\hat{\sigma}_2(\hat{a}_2) \geq \varepsilon$ for some $\hat{a}_2 \neq \tilde{a}_2$. Let \hat{a}_1 be the best response to \hat{a}_2 . Now

$$(3) \quad u_1^r(\sigma_1, \hat{\sigma}_2) > u_1^r(\bar{a}_1, \hat{\sigma}_2)$$

for all sufficiently large r , because the right-hand lottery has support contained in $\{(\bar{a}_1, \hat{a}_2), (\bar{a}_1, \tilde{a}_2)\}$, both of which have lower utility than (\hat{a}_1, \hat{a}_2) which is contained in the left-hand lottery. Thus, since the set of extreme points $\hat{\sigma}_2$ is finite, for large enough r the inequality (3) holds for *all* extreme points $\hat{\sigma}_2$ of S . We then conclude by linearity that (3) holds for all $\hat{\sigma}_2$ in the convex hull of these extreme points, which is all of P_ε . Then \bar{a}_1 is not a best reply to any belief in P_ε , as desired. This implies that, given ε , for sufficiently large r any equilibrium (σ_1, σ_2) with support $S \times T$ must have $\sigma_2 \notin P_\varepsilon$, which together with the corresponding claim for σ_1 gives the desired result.

Claim 2: Suppose $a_1 \in S$ is not a best reply to any $a_2 \in T$. Let σ_1 be an equal mixture over all of A_1 . For each $a_2 \in T$,

$$u^r(\sigma_1, a_2) > u^r(a_1, a_2)$$

for sufficiently large r , because the left-hand lottery has a higher maximum. By finiteness of T , $\sigma_1 \succ_T a_1$ in G^r for all sufficiently large r , and the desired result follows from Lemma 1. *Q.E.D.*

PROOF OF PROPOSITION 5:

Claim 1(a): Suppose $S \times T$ admits no safety. We show that for sufficiently large r , the condition following “if and only if” in Lemma 1 holds in G^{-r} , whence the lemma yields an equilibrium. Fix any $a_1 \in A_1$; it is not a never-worst deviation, so for some $a_2 \in T$, $u(a_1, a_2) < u(a'_1, a_2)$ for all $a'_1 \in S - \{a_1\}$.

¹⁶This characterization of the extreme points can be verified by straightforward linear algebra: they are the points where $k - 1$ of the defining inequalities bind. The only important thing about the extreme points for the proof is that each assigns mass at least ε to some $\hat{a}_2 \neq \tilde{a}_2$.

Fix $\varepsilon \in (0, 1/|A_1|)$. Suppose $\sigma_1 >_T \sigma'_1$ with $\sigma_1 \in \Delta(A_1)$, $\sigma'_1 \in \Delta(S)$ having disjoint support, as in the lemma. Suppose $\sigma_1(a_1) \geq \varepsilon$. Disjoint support implies $\sigma'_1(S - \{a_1\}) = 1$. Similarly to the previous proof, we can deduce from unbounded concavity that

$$u'_1(\sigma_1, a_2) < u'_1(\sigma'_1, a_2)$$

for large enough r , contradicting weak dominance. So $\sigma_1(a_1) < \varepsilon$. Letting r be large enough that $\sigma_1(a_1) < \varepsilon < 1/|A_1|$ for all $a_1 \in A_1$ gives a contradiction, so for large r no pair σ_1, σ'_1 with $\sigma_1 >_T \sigma'_1$ exists. Also take r large enough for the analogous conditions to hold with roles reversed, and we have proven the existence for such r of an equilibrium with support $S \times T$.

Claim 1(b): For any game G , define $-G$ as the game with identical action sets and all utilities u_i replaced by $-u_i$. Notice that if G^{-r} is an unboundedly concave family of transformations of G , $-[G^{-r}]$ is an unboundedly convex family of transformations of $-G$. Any equilibrium of G^{-r} with support $S \times T$ makes players indifferent among all actions in the block, so it is also a completely mixed equilibrium of $-[G^{-r}]_{S \times T}$. Applying Proposition 4 to the completely mixed equilibria of $-[G^{-r}]_{S \times T}$ tells us that they converge to the pure strategy profile described there. It remains only to observe that

$$\begin{aligned} \operatorname{argmin}_{a_1 \in S} \left[\max_{a_2 \in T} -u_2(a_1, a_2) \right] &= \operatorname{argmax}_{a_1 \in S} \left[\min_{a_2 \in T} u_2(a_1, a_2) \right], \\ \operatorname{argmin}_{a_2 \in T} \left[\max_{a_1 \in S} -u_1(a_1, a_2) \right] &= \operatorname{argmax}_{a_2 \in T} \left[\min_{a_1 \in S} u_1(a_1, a_2) \right], \end{aligned}$$

as desired.

Claim 2: Suppose $a_1 \in A_1$ is a never-worst deviation from $S \times T$. Let σ'_1 be an equal mixture of all elements of S . For all $a_2 \in T$,

$$u^{-r}(a_1, a_2) > u^{-r}(\sigma'_1, a_2)$$

for sufficiently large r , because the right-hand lottery has a worst outcome which is less than the left-hand-side, by definition of never-worst. By finiteness of T , $a_1 >_T \sigma'_1$ in G^{-r} for all sufficiently large r , and the desired result follows by Lemma 1. *Q.E.D.*

APPENDIX B

PROOF OF PROPOSITION 8: First, we show $P_i^\infty(G) \subseteq \operatorname{ICR}_{\hat{G}}(t_i)$. We suppress the \hat{G} in future notation. It suffices to show that if, for some k , $P_i^k(G) \subseteq \operatorname{ICR}_i^k(t_i)$, the same is true for $k + 1$. Let $a_i \in P_i^{k+1}(G)$. Then $a_i \in \operatorname{BR}(a_{-i})$ for some pure profile $a_{-i} \in P_{-i}^k(G)$. By inductive hypothesis, $a_{-i} \in \operatorname{ICR}_i^k(t_{-i})$ for

all t_{-i} , implying that $\kappa_{t_i} \times \delta_{a_{-i}}$ is an admissible conjecture, where δ denotes a Dirac measure. Furthermore, because a_i is a best response to a_{-i} for every θ , the sure-thing principle tells us $a_i \in \text{BR}_i(\text{marg}_{\theta} \kappa_{t_i} \times \delta_{a_{-i}})$ as desired.

Next, we show $\text{ICR}_{\hat{G}}(t_i) \subseteq B_i^{\infty}(G)$. Following a similar idea, we show inductively that $\text{ICR}_{\hat{G}}^k(t_i) \subseteq B_i^k(G)$ for every i, t_i , and k . Suppose, towards a contradiction, that the first violation of this claim is $a_i \in B_i^k(G) - B_i^{k+1}(G)$ and $a_i \in \text{ICR}_{\hat{G}}^{k+1}(t_i)$. Let $\nu \in C_{k+1}(t_i)$ be the conjecture which rationalizes a_i . By inductive hypothesis, $\text{marg}_{A_{-i}} \nu(B_{-i}^k) = 1$. By definition of the B -process, there is then an action a'_i such that $u(a'_i, a_{-i}) \geq u(a_i, a_{-i})$ for all $a_{-i} \in \text{supp}(\text{marg}_{A_{-i}} \nu)$ with strict inequality for some such a_{-i} . Then also

$$\text{marg}_{\theta \times A_{-i}} \nu(\theta, a_{-i}) \hat{u}(a'_i, a_{-i}, \theta) \geq \text{marg}_{\theta \times A_{-i}} \nu(\theta, a_{-i}) \hat{u}(a_i, a_{-i}, \theta)$$

for every pair (θ, a_{-i}) , and again inequality must be strict somewhere, but this contradicts $a_i \in \text{BR}_i(\text{marg}_{\theta \times A_{-i}} \nu)$. Q.E.D.

PROOF OF PROPOSITION 9: For the first part, recall the proof of Proposition 2. We showed that for sufficiently large r , each action a_i removed at stage k of the P -elimination process is dominated (at stage k of the rationalizability process) by a uniform mixture of the remaining actions. Here, take r large enough that this is true for the game induced by each θ , and then the uniform mixture dominates the removed action in \hat{G} also, which suffices for the proof. Note that we needed the fact that the same mixture dominates a_i for each θ ; merely the fact that the rationalizable set is the same for each θ would not suffice.

For the second part, it suffices to show that for large enough r , for every t_i , each $a_i \in B_i^{\infty}(G)$ is a best response to a belief with support on $B_{-i}^{\infty}(G)$, satisfying the condition on marginals as in 1. Proposition 3 tells us that for large enough r , for each θ there is a belief ν_{θ} with support on $B_{-i}^{\infty}(G)$ such that, when parameter θ obtains for sure, a_i is a best response to ν_{θ} . But then it suffices to assign type t_i a belief ν given by

$$\nu(\theta, t_{-i}, a_{-i}) = \kappa_{t_i}(\theta, t_{-i}) \nu_{\theta}(a_{-i})$$

because in response to this belief, a_i is optimal conditional on each θ and hence unconditionally optimal. Q.E.D.

REFERENCES

- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18. [1884,1894]
 AZRIELI, Y., AND R. TEPER (2011): "Uncertainty Aversion and Equilibrium Existence in Games With Incomplete Information," *Games and Economic Behavior*, 73, 310–317. [1892]

- BATTIGALLI, P., S. CERREIA-VIOLIO, F. MACCHERONI, AND M. MARINACCI (2016): "A Note on Comparative Ambiguity Aversion and Justifiability," *Econometrica*, 84, 1903–1916. [1887]
- BERNHEIM, D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028. [1881, 1882, 1884, 1895]
- BORGERS, T. (1993): "Pure Strategy Dominance," *Econometrica*, 61, 423–430. [1882, 1884, 1889]
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2007): "Interim Correlated Rationalizability," *Theoretical Economics*, 2, 15–40. [1896]
- GAFAROV, B., AND B. SALCEDO (2014): "Ordinal Dominance and Risk Aversion," *Economic Theory Bulletin*, 1–12. [1888]
- NAU, R. F., AND K. F. MCCARDLE (1990): "Coherent Behavior in Non-Cooperative Games," *Journal of Economic Theory*, 50, 424–444. [1894]
- PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050. [1881]
- WEINSTEIN, J. (2015): "Best-Reply Sets," Working Paper, Available at <https://pages.wustl.edu/files/pages/imce/weinstein/best-replies.pdf>. [1898]
- YILDIZ, M. (2007): "Wishful Thinking in Strategic Environments," *Review of Economic Studies*, 74, 319–344. [1892]

Washington University in St. Louis, 1 Brookings Dr., St Louis, MO 63130, U.S.A.; j.weinstein@wustl.edu.

Co-editor Joel Sobel handled this manuscript.

Manuscript received November, 2015; final revision received May, 2016.