games

MDPI

Article

Epistemically Robust Strategy Subsets

Geir B. Asheim [1,*], Mark Voorneveld [2] and Jörgen W. Weibull [2,3,4]

1 Department of Economics, University of Oslo, P.O. Box 1095 Blindern, NO-0317 Oslo, Norway
2 Department of Economics, Stockholm School of Economics, Box 6501, SE-113 83 Stockholm, Sweden; nemv@hhs.se (M.V.); nejw@hhs.se (J.W.W.)
3 Institute for Advanced Study in Toulouse, 31000 Toulouse, France
4 Department of Mathematics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden
* Correspondence: g.b.asheim@econ.uio.no; Tel.: +47-455-051-36

Academic Editors: Paul Weirich and Ulrich Berger
Received: 31 August 2016; Accepted: 17 November 2016; Published: 25 November 2016

Abstract: We define a concept of *epistemic robustness* in the context of an epistemic model of a finite normal-form game where a player type corresponds to a belief over the profiles of opponent strategies and types. A Cartesian product X of pure-strategy subsets is epistemically robust if there is a Cartesian product Y of player type subsets with X as the associated set of best reply profiles such that the set Y_i contains all player types that believe with sufficient probability that the others are of types in Y_{-i} and play best replies. This robustness concept provides epistemic foundations for set-valued generalizations of strict Nash equilibrium, applicable also to games without strict Nash equilibria. We relate our concept to closedness under rational behavior and thus to strategic stability and to the best reply property and thus to rationalizability.

Keywords: epistemic game theory; epistemic robustness; rationalizability; closedness under rational behavior; mutual p-belief

JEL Classification Numbers: C72; D83

1. Introduction

In most applications of noncooperative game theory, Nash equilibrium is used as a tool to predict behavior. Under what conditions, if any, is this approach justified? In his Ph.D. thesis, Nash [1] suggested two interpretations of Nash equilibrium, one rationalistic, in which all players are fully rational, know the game, and play it exactly once. In the other, "mass action" interpretation, there is a large population of actors for each player role of the game, and now and then exactly one actor from each player population is drawn at random to play the game in his or her player role, and this is repeated (i.i.d.) indefinitely over time. Whereas the latter interpretation is studied in the literature on evolutionary game theory and social learning, the former—which is the interpretation we will be concerned with here—is studied in a sizeable literature on epistemic foundations of Nash equilibrium. It is by now well-known from this literature that players' rationality and beliefs or knowledge about the game and each others' rationality in general do not imply that they necessarily play a Nash equilibrium or even that their conjectures about each others' actions form a Nash equilibrium; see Bernheim [2], Pearce [3], Aumann and Brandenburger [4].

The problem is not only a matter of coordination of beliefs (conjectures or expectations), as in a game with multiple equilibria. It also concerns the fact that, in Nash equilibrium interpreted as an equilibrium in belief (see [4], Theorems A and B), beliefs are supposed to correspond to *specific* randomizations over the others' strategies. In particular, a player might have opponents with multiple pure strategies that maximize their expected payoffs, given their equilibrium beliefs. Hence, for these opponents, any randomization over their pure best replies maximizes their expected payoffs. Yet in Nash equilibrium, the player is assumed to have a belief that singles out a randomization over the best replies of her opponents that serves to keep this player indifferent across the support of her equilibrium strategies, and ensures that none of the player's other strategies are better replies. In addition, a player's belief concerning the behavior of others assigns positive probability *only* to best replies; players are not allowed to entertain any doubt about the rationality of their opponents.

Our aim is to formalize a notion of epistemic robustness that relaxes these requirements. In order to achieve this, we have to move away from point-valued to set-valued solution concepts. In line with the terminology of epistemic game theory, let a player's epistemic *type* correspond to a belief over the profiles of opponent strategies and types. Assume that the epistemic model is complete in the sense that all possible types are represented in the model. Let non-empty Cartesian products of (pure-strategy or type) subsets be referred to as (strategy or type) *blocks* [5]. Say that a strategy block $X = X_1 \times \cdots \times X_n$ is epistemically robust if there exists a corresponding type block $Y = Y_1 \times \cdots \times Y_n$ such that: for each player i,

(I) the strategy subset X_i coincides with the set of best replies of the types in Y_i;
(II) the set Y_i contains all player types that believe with sufficient probability that the others are of types in Y_{-i} and play best replies.

Here, for each player, (II) requires the player's type subset to be robust in the sense of including all possible probability distributions over opponent pure-strategy profiles that consist of best replies to the beliefs of opponent types that are included in the opponents' type subsets, even including player types with a smidgen of doubt that only these strategies are played. In particular, our epistemic model does not allow a player to pinpoint a specific opponent type or a specific best reply for an opponent type that has multiple best replies. The purpose of (I) is, for each player, to map this robust type subset into a robust subset of pure strategies by means of the best reply correspondence.

Consider, in contrast, the case where point (II) above is replaced by:

(II′) the set Y_i contains *only* player types that believe with probability 1 that the others are of types in Y_{-i} and play best replies.

Tan and Werlang [6] show that the strategy block X is a *best reply set* [3] if there exists a corresponding type block Y such that (I) and (II′) hold for all players. This epistemic characterization of a best reply set X explains why, for each player i, all strategies in X_i are included. In contrast, the concept of epistemic robustness explains why all strategies outside X_i are excluded, as a rational player will never choose such a strategy, not even if the player with small probability believes that opponents will not stick to their types Y_{-i} or will not choose best replies.

Any strict Nash equilibrium, viewed as a singleton strategy block, is epistemically robust. In this case, each player has opponents with unique pure strategies that maximize their expected payoffs, given their equilibrium beliefs. The player's equilibrium strategy remains her unique best reply, as long as she is *sufficiently sure* that the others stick to their unique best replies. By contrast, non-strict pure-strategy Nash equilibria by definition have 'unused' best replies and are consequently not epistemically robust: a player, even if she is sure that her opponents strive to maximize their expected payoffs given their equilibrium beliefs, might well believe that her opponents play such alternative best replies.

In informal terms, our Proposition 1 establishes that epistemic robustness is sufficient and necessary for the non-existence of such 'unused' best replies. Consequently, epistemic robustness captures, through restrictions on the players' beliefs, a property satisfied by strict Nash equilibria, but not by non-strict pure-strategy Nash equilibria. The restrictions on players' beliefs implied by epistemic robustness can be imposed also on games without strict Nash equilibria. Indeed, our Propositions 2–5 show how epistemic robustness is achieved by variants of CURB sets. A CURB set (mnemonic for 'closed under rational behavior') is a strategy block that contains, for each player, all best replies to all probability distributions over the opponent strategies in the block[1]. Hence, if a player believes that her opponents stick to strategies from their components of a CURB set, then she'd better stick to her strategies as well.

A strategy block is fixed under rational behavior (FURB; or 'tight' CURB in the terminology of Basu and Weibull [7]) if each player's component not only contains, but is identical with the set of best replies to all probability distributions over the opponent strategies in the block. Basu and Weibull [7] show that minimal CURB (MINCURB) sets and the unique largest FURB set are important special cases of FURB sets. The latter equals the strategy block of rationalizable strategies [2,3]. At the other extreme, MINCURB is a natural set-valued generalization of strict Nash equilibrium. The main purpose of this paper is to provide epistemic foundations for set-valued generalizations of strict Nash equilibrium. Our results are not intended to advocate any particular point- or set-valued solution concept, only to propose a definition of epistemic robustness and apply this to some set-valued solution concepts currently in use[2].

In order to illustrate our line of reasoning, consider first the two-player game

	l	c
u	3,1	1,2
m	0,3	2,1

In its unique Nash equilibrium, player 1's equilibrium strategy assigns probability $2/3$ to her first pure strategy and player 2's equilibrium strategy assigns probability $1/4$ to his first pure strategy. However, even if player 1's belief about the behavior of player 2 coincides with his equilibrium strategy, $(1/4, 3/4)$, player 1 would be indifferent between her two pure strategies. Hence, any pure or mixed strategy would be optimal for her, under the equilibrium belief about player 2. For all other beliefs about her opponent's behavior, only one of her pure strategies would be optimal, and likewise for player 2. The unique CURB set and unique epistemically robust subset in this game is the full set $S = S_1 \times S_2$ of pure-strategy profiles.

Add a third pure strategy for each player to obtain the two-player game

	l	c	r
u	3,1	1,2	0,0
m	0,3	2,1	0,0
d	5,0	0,0	6,4

(1)

Strategy profile $x^* = (x_1^*, x_2^*) = \left(\left(\frac{2}{3}, \frac{1}{3}, 0\right), \left(\frac{1}{4}, \frac{3}{4}, 0\right)\right)$ is a Nash equilibrium (indeed a perfect and proper equilibrium). However, if player 2's belief concerning the behavior of 1 coincides with x_1^*, then 2 is indifferent between his pure strategies l and c, and if 1 assigns equal probability to these two pure strategies of player 2, then 1 will play the unique best reply d, a pure strategy outside the support of the equilibrium[3]. Moreover, if player 2 expects 1 to reason this way, then 2 will play r: the smallest epistemically robust subset containing the support of the mixed equilibrium x^* is the entire pure

1 CURB sets and variants were introduced by Basu and Weibull [7] and have since been used in many applications. Several classes of adaptation processes eventually settle down in a minimal CURB set; see Hurkens [8], Sanchirico [9], Young [10], and Fudenberg and Levine [11]. Moreover, minimal CURB sets give appealing results in communication games [12,13] and network formation games [14]. For closure properties under generalizations of the best reply correspondence, see Ritzberger and Weibull [15].

2 Clearly, if a strategy block is *not* epistemically robust, then our concept does not imply that players should or will avoid strategies in the block.

3 We emphasize that we are concerned with rationalistic analysis of a game that is played once, and where players have beliefs about the rationality and beliefs of their opponents. If the marginal of a player's belief on an opponent's strategy set is non-degenerate—so that the player is uncertain about the behavior of the opponent—then this can be interpreted as the player believing that the opponent is playing a mixed strategy.

strategy space. By contrast, the pure-strategy profile (d, r) is a strict equilibrium. In this equilibrium, no player has any alternative best reply and each equilibrium strategy remains optimal also under some uncertainty as to the other player's action: the set $\{d\} \times \{r\}$ is epistemically robust. In this game, all pure strategies are rationalizable, $S = S_1 \times S_2$ is a FURB set, and the game's unique MINCURB set (thus, the unique minimal FURB set) is $T = \{d\} \times \{r\}$. These are also the epistemically robust subsets; in particular, $\{u, m\} \times \{l, c\}$ is not epistemically robust.

Our results can be described as follows. First, the intuitive link between strict Nash equilibria and our concept of epistemic robustness in terms of ruling out the existence of 'unused' best replies is formalized in Proposition 1: a strategy block X is *not* epistemically robust if and only if for each type block Y raised in its defense—so that X is the set of best reply profiles associated with Y—there is a player i and a type t_i with a best reply outside X_i, even if t_i believes with high probability that his opponents are of types in Y_{-i} and play best replies. Second, in part (a) of Proposition 2, we establish that epistemically robust strategy blocks are CURB sets. As a consequence (see [15]), every epistemically robust strategy block contains at least one strategically stable set in the sense of Kohlberg and Mertens [16]. In part (b) of Proposition 2, although not every CURB set is epistemically robust (since a CURB set may contain non-best replies), we establish that every CURB set contains an epistemically robust strategy block and we also characterize the largest such subset. As a by-product, we obtain the existence of epistemically robust strategy blocks in all finite games. Third, in Proposition 3, we show that a strategy block is FURB if and only if it satisfies the definition of epistemic robustness with equality, rather than inclusion, in (II). FURB sets thus have a clean epistemic robustness characterization in the present framework. Fourth, in Proposition 4, instead of starting with strategy blocks, we start from a type block and show how an epistemically robust strategy block can be algorithmically obtained; we also show that this is the smallest CURB set that contains all best replies for the initial type block. Fifth, Proposition 5 shows how MINCURB sets (which are necessarily FURB and hence epistemically robust) can be characterized by initiating the above algorithm with a single type profile, while no proper subset has this property. We argue that this latter result shows how MINCURB sets capture characteristics of strict Nash equilibrium.

As our notion of epistemic robustness checks for player types with 'unused' best replies on the basis of their beliefs about the opponents' types and rationality, we follow, for instance, Asheim [17] and Brandenburger, Friedenberg, and Keisler [18], and model players as having beliefs about the opponents without modeling the players' actual behavior. Moreover, we consider epistemic models that are complete in the sense of including all possible beliefs. In these respects, our modeling differs from that of Aumann and Brandenburger [4]'s characterization of Nash equilibrium. In other respects, our modeling resembles that of Aumann and Brandenburger [4]. They assume that players' beliefs about opponent play is commonly known. Here, we require the existence of a type block Y and consider, for each player i, types of player i who believe that opponent types are in Y_{-i}. In addition, as do Aumann and Brandenburger [4], we consider types of players that believe that their opponents are rational.

The notion of persistent retracts [19] goes part of the way towards epistemic robustness. These are product sets requiring the presence of *at least one* best reply to arbitrary beliefs *close to* the set. In other words, they are robust against small belief perturbations, but admit alternative best replies outside the set, in contrast to our concept of epistemic robustness. Moreover, as pointed out by (van Damme [20] Section 4.5) and Myerson and Weibull [5], persistence is sensitive to certain game details that might be deemed strategically inessential.

The present approach is related to Tercieux [21]'s analysis in its motivation in terms of epistemic robustness of solution concepts and in its use of p-belief. His epistemic approach, however, is completely different from ours. Starting from a two-player game, he introduces a Bayesian game where payoff functions are perturbations of the original ones and he investigates which equilibria are robust against this kind of perturbation. Zambrano [22] studies the stability of non-equilibrium concepts in terms of mutual belief and is hence more closely related to our analysis. In fact, our

Proposition 3 overlaps with but is distinct from his main results. Also Hu [23] restricts attention to rationalizability, but allows for p-beliefs, where $p < 1$. In the games considered in Hu [23], pure strategy sets are permitted to be infinite. By contrast, our analysis is restricted to finite games, but under the weaker condition of mutual, rather than Hu [23]'s common, p-belief of opponent rationality and of opponents' types belonging to given type sets.

The remainder of the paper is organized as follows. Section 2 contains the game theoretic and epistemic definitions used. Section 3 characterizes variants of CURB sets in terms of epistemic robustness. An appendix contains proofs of the propositions.

2. The Model

2.1. Game Theoretic Definitions

Consider a finite normal-form game $\langle N, (S_i)_{i\in N}, (u_i)_{i\in N}\rangle$, where $N = \{1, \ldots, n\}$ is the non-empty and finite set of players. Each player $i \in N$ has a non-empty, finite set of pure strategies S_i and a payoff function $u_i : S \to \mathbb{R}$ defined on the set $S := S_1 \times \cdots \times S_n$ of pure-strategy profiles. For any player i, let $S_{-i} := \times_{j\neq i} S_j$. It is over this set of *other* players' pure-strategy combinations that player i will form his or her probabilistic beliefs. These beliefs may, but need not be, product measures over the other player's pure-strategy sets. We extend the domain of the payoff functions to probability distributions over pure strategies as usual.

For each player $i \in N$, pure strategy $s_i \in S_i$, and probabilistic belief $\sigma_{-i} \in \mathcal{M}(S_{-i})$, where $\mathcal{M}(S_{-i})$ is the set of all probability distributions on the finite set S_{-i}, write

$$u_i(s_i, \sigma_{-i}) := \sum_{s_{-i}\in S_{-i}} \sigma_{-i}(s_{-i}) u_i(s_i, s_{-i}).$$

Define i's *best reply correspondence* $\beta_i : \mathcal{M}(S_{-i}) \to 2^{S_i}$ as follows: for all $\sigma_{-i} \in \mathcal{M}(S_{-i})$,

$$\beta_i(\sigma_{-i}) := \{s_i \in S_i \mid u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i}) \text{ for all } s_i' \in S_i\}.$$

Let $\mathcal{S} := \{X \in 2^S \mid \varnothing \neq X = X_1 \times \cdots \times X_n\}$ denote the collection of strategy blocks. For $X \in \mathcal{S}$, we abuse notation slightly by writing, for each $i \in N$, $\beta_i(\mathcal{M}(X_{-i}))$ as $\beta_i(X_{-i})$. Let $\beta(X) := \beta_1(X_{-1}) \times \cdots \times \beta_n(X_{-n})$. Each constituent set $\beta_i(X_{-i}) \subseteq S_i$ in this strategy block is the set of best replies of player i to all probabilistic beliefs over the others' strategy choices $X_{-i} \subseteq S_{-i}$.

Following Basu and Weibull [7], a set $X \in \mathcal{S}$ is:

closed under rational behavior (CURB) if $\beta(X) \subseteq X$;

fixed under rational behavior (FURB) if $\beta(X) = X$;

minimal CURB (MINCURB) if it is CURB and does not properly contain another one: $\beta(X) \subseteq X$ and there is no $X' \in \mathcal{S}$ with $X' \subsetneq X$ and $\beta(X') \subseteq X'$.

Basu and Weibull [7] call a FURB set a 'tight' CURB set. The reversed inclusion, $X \subseteq \beta(X)$, is the *best reply property* ([3] p. 1033). It is shown in (Basu and Weibull [7] Propositions 1 and 2) that a MINCURB set exists, that all MINCURB sets are FURB, and that the block of rationalizable strategies is the game's largest FURB set. While Basu and Weibull [7] require that players believe that others' strategy choices are statistically independent, $\sigma_{-i} \in \times_{j\neq i}\mathcal{M}(S_j)$, we here allow players to believe that others' strategy choices are correlated, $\sigma_{-i} \in \mathcal{M}(S_{-i})$[4]. Our results carry over—with minor modifications in the proofs—to the case of independent strategy choices. Thus, in games with more than two players, the present definition of CURB is somewhat more demanding than that in Basu and Weibull [7], in the

---

4 In doing so, we follow (Osborne and Rubinstein [24] Chapter 5).

sense that we require closedness under a wider space of beliefs. Hence, the present definition may, in games with more than two players, lead to different MINCURB sets[5].

2.2. Epistemic Definitions

The epistemic analysis builds on the concept of player types, where a type of a player is characterized by a probability distribution over the others' strategies and types.

For each $i \in N$, denote by T_i player i's non-empty type space. The *state space* is defined by $\Omega := S \times T$, where $T := T_1 \times \cdots \times T_n$. For each player $i \in N$, write $\Omega_i := S_i \times T_i$ and $\Omega_{-i} := \times_{j \neq i} \Omega_j$. To each type $t_i \in T_i$ of every player i is associated a probabilistic belief $\mu_i(t_i) \in \mathcal{M}(\Omega_{-i})$, where $\mathcal{M}(\Omega_{-i})$ denotes the set of Borel probability measures on Ω_{-i} endowed with the topology of weak convergence. For each player i, we thus have the player's pure-strategy set S_i, type space T_i and a mapping $\mu_i : T_i \to \mathcal{M}(\Omega_{-i})$ that to each of i's types t_i assigns a probabilistic belief, $\mu_i(t_i)$, over the others' strategy choices and types. Assume that, for each $i \in N$, μ_i is continuous and T_i is compact. The structure $(S_1, \ldots, S_n, T_1, \ldots, T_n, \mu_1, \ldots, \mu_n)$ is called an *S-based (interactive) probability structure*. Assume in addition that, for each $i \in N$, μ_i is onto: all Borel probability measures on Ω_{-i} are represented in T_i. A probability structure with this additional property is called *complete*.[6] The completeness of the probability structure is essential for our analysis and results. In particular, the assumption of completeness is invoked in all proofs.

For each $i \in N$, denote by $\mathbf{s}_i(\omega)$ and $\mathbf{t}_i(\omega)$ i's strategy and type in state $\omega \in \Omega$. In other words, $\mathbf{s}_i : \Omega \to S_i$ is the projection of the state space to i's strategy set, assigning to each state $\omega \in \Omega$ the strategy $s_i = \mathbf{s}_i(\omega)$ that i uses in that state. Likewise, $\mathbf{t}_i : \Omega \to T_i$ is the projection of the state space to i's type space. For each player $i \in N$ and positive probability $p \in (0, 1]$, the *p-belief operator* B_i^p maps each event (Borel-measurable subset of the state space) $E \subseteq \Omega$ to the set of states where player i's type attaches at least probability p to E:

$$B_i^p(E) := \{\omega \in \Omega \mid \mu_i(\mathbf{t}_i(\omega))(E^{\omega_i}) \geq p\},$$

where $E^{\omega_i} := \{\omega_{-i} \in \Omega_{-i} \mid (\omega_i, \omega_{-i}) \in E\}$. This is the same belief operator as in Hu [23][7]. One may interpret $B_i^p(E)$ as the event 'player i believes E with probability at least p'. For all $p \in (0, 1]$, B_i^p satisfies $B_i^p(\varnothing) = \varnothing$, $B_i^p(\Omega) = \Omega$, $B_i^p(E') \subseteq B_i^p(E'')$ if $E' \subseteq E''$ (monotonicity), and $B_i^p(E) = E$ if $E = \text{proj}_{\Omega_i} E \times \Omega_{-i}$. The last property means that each player i always p-believes his own strategy-type pair, for any positive probability p. Since also $B_i^p(E) = \text{proj}_{\Omega_i} B_i^p(E) \times \Omega_{-i}$ for all events $E \subseteq \Omega$, each operator B_i^p satisfies both positive ($B_i^p(E) \subseteq B_i^p(B_i^p(E))$) and negative ($\neg B_i^p(E) \subseteq B_i^p(\neg B_i^p(E))$) introspection. For all $p \in (0, 1]$, B_i^p violates the truth axiom, meaning that $B_i^p(E) \subseteq E$ need not hold for all $E \subseteq \Omega$. In the special case $p = 1$, we have $B_i^p(E') \cap B_i^p(E'') \subseteq B_i^p(E' \cap E'')$ for all $E', E'' \subseteq \Omega$. Finally, note that $B_i^p(E)$ is monotone with respect to p in the sense that, for all $E \subseteq \Omega$, $B_i^{p'}(E) \supseteq B_i^{p''}(E)$ if $p' < p''$.

We connect types with the payoff functions by defining i's *choice correspondence* $C_i : T_i \to 2^{S_i}$ as follows: For each of i's types $t_i \in T_i$,

$$C_i(t_i) := \beta_i(\text{marg}_{S_{-i}} \mu_i(t_i))$$

---

5 We also note that a pure strategy is a best reply to some belief $\sigma_{-i} \in \mathcal{M}(S_{-i})$ if and only if it is not strictly dominated (by any pure or mixed strategy). This follows from Lemma 3 in Pearce [3], which, in turn, is closely related to (Ferguson [25] p. 86, Theorem 1) and (van Damme [26] Lemma 3.2.1).

6 An adaptation of the proof of (Brandenburger, Friedenberg, and Keisler [18] Proposition 7.2) establishes the existence of such a complete probability structure under the assumption that, for all $i \in N$, player i's type space T_i is Polish (separable and completely metrizable). The exact result we use is Proposition 6.1 in an earlier working paper version [27]. Existence can also be established by constructing a universal state space [28,29].

7 See also Monderer and Samet [30].

consists of i's best replies when player i is of type t_i. Let $\mathcal{T} := \{Y \in 2^T \mid \varnothing \neq Y = Y_1 \times \cdots \times Y_n\}$ denote the collection of type blocks. For any such set $Y \in \mathcal{T}$ and player $i \in N$, write $C_i(Y_i) := \bigcup_{t_i \in Y_i} C_i(t_i)$ and $C(Y) := C_1(Y_1) \times \cdots \times C_n(Y_n)$. In other words, these are the choices and choice profiles associated with Y. If $Y \in \mathcal{T}$ and $i \in N$, write

$$[Y_i] := \{\omega \in \Omega \mid \mathbf{t}_i(\omega) \in Y_i\}.$$

This is the event that player i is of a type in the subset Y_i. Likewise, write $[Y] := \bigcap_{i \in N}[Y_i]$ for the event that the type profile is in Y. Finally, for each player $i \in N$, write R_i for the event that player i uses a best reply:

$$R_i := \{\omega \in \Omega \mid \mathbf{s}_i(\omega) \in C_i(\mathbf{t}_i(\omega))\}.$$

One may interpret R_i as the event that i is rational: if $\omega \in R_i$, then $\mathbf{s}_i(\omega)$ is a best reply to $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega))$.

3. Epistemic Robustness

We define a strategy block $X \in \mathcal{S}$ to be *epistemically robust* if there exists a $\bar{p} < 1$ such that, for each probability $p \in [\bar{p}, 1]$, there is a type block $Y \in \mathcal{T}$ (possibly dependent on p) such that

$$C(Y) = X \tag{2}$$

and

$$B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) \subseteq [Y_i] \quad \text{for all } i \in N\,. \tag{3}$$

Hence, epistemic robustness requires the existence of a type block Y satisfying, for each player i, that X_i is the set of best replies of the types in Y_i, and that every type of player i who p-believes that opponents are rational and of types in Y_{-i} is included in Y_i. Condition (2) is thus not an equilibrium condition as it is not interactive: it relates each player's type subset to the same player's strategy subset. The interactivity enters through condition (3), which relates each player's type subset to the type subsets of the other players. For each $p < 1$, condition (3) allows each player i to attach a positive probability to the event that others do not play best replies and/or are of types outside Y. It follows from the monotonicity of $B_i^p(\cdot)$ with respect to p that, for a fixed type block Y, if inclusion (3) is satisfied for $p = \bar{p}$, then inclusion (3) is satisfied also for all $p \in (\bar{p}, 1]$.

Note that if condition (2) is combined with a variant of condition (3), with the weak inclusion reversed and p set to 1, then we obtain a characterization of Pearce [3]'s best reply set; see [6].

In line with what we mentioned in the introduction, we can now formally show that if $s \in S$ is a strict Nash equilibrium, then $\{s\}$ is epistemically robust. To see this, define for all $i \in N$, $Y_i := \{t_i \in T_i \mid C_i(t_i) = \{s_i\}\}$. Since the game is finite, there is, for each player $i \in N$, a $\underline{p}_i \in (0, 1)$ such that $\beta_i(\sigma_{-i}) = \{s_i\}$ for all $\sigma_{-i} \in \mathcal{M}(S_{-i})$ with $\sigma_{-i}(\{s_{-i}\}) \geq \underline{p}_i$. Let $\underline{p} = \max\{\underline{p}_1, \ldots, \underline{p}_n\}$. Then it holds for each $p \in [\underline{p}, 1]$:

$$B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) \subseteq B_i^p(\{\omega \in \Omega \mid \forall j \neq i, \mathbf{s}_j(\omega) \in X_j\}) \subseteq [Y_i] \quad \text{for all } i \in N\,.$$

Thus, by condition (2) and condition (3), $\{s\}$ is epistemically robust.

Also, as discussed in the introduction, non-strict pure-strategy Nash equilibria have 'unused' best replies. Our first result demonstrates that epistemic robustness is sufficient and necessary for the non-existence of such 'unused' best replies.

Proposition 1. *The following two statements are equivalent:*

(a) $X \in \mathcal{S}$ *is* not *epistemically robust.*

(b) *For all $\bar{p} < 1$, there exists $p \in [\bar{p}, 1]$ such that if $Y \in \mathcal{T}$ satisfies $C(Y) = X$, then there exist $i \in N$ and $t_i \in T_i$ such that $C(t_i) \not\subseteq X_i$ and $[\{t_i\}] \subseteq B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right)$.*

Hence, while an epistemically robust subset is defined by a *set* of profiles of player types, it suffices with one player and one possible type of this player to determine that a strategy block is not epistemically robust.

We now relate epistemically robust subsets to CURB sets. To handle the fact that all strategy profiles in any epistemically robust subset are profiles of best replies, while CURB sets may involve strategies that are not best replies, introduce the following notation: For each $i \in N$ and $X_i \subseteq S_i$, let

$$\beta_i^{-1}(X_i) := \{\sigma_{-i} \in \mathcal{M}(S_{-i}) \mid \beta_i(\sigma_{-i}) \subseteq X_i\}$$

denote the pre-image (upper inverse) of X_i under player i's best reply correspondence[8]. For a given subset X_i of i's pure strategies, $\beta_i^{-1}(X_i)$ consists of the beliefs over others' strategy profiles having the property that all best replies to these beliefs are contained in X_i.

Proposition 2. *Let $X \in \mathcal{S}$.*

(a) *If X is epistemically robust, then X is a CURB set.*
(b) *If X is a CURB set, then $\times_{i \in N} \beta_i(\beta_i^{-1}(X_i)) \subseteq X$ is epistemically robust. Furthermore, it is the largest epistemically robust subset of X.*

Claim (a) implies that every epistemically robust subset contains at least one strategically stable set, both as defined in Kohlberg and Mertens [16] and as defined in Mertens [32], see Ritzberger and Weibull [15] and Demichelis and Ritzberger [33], respectively[9]. Claim (a) also implies that subsets of epistemically robust sets need not be epistemically robust. Concerning claim (b), note that $\times_{i \in N} \beta_i(\beta_i^{-1}(S_i))$ equals the set of profiles of pure strategies that are best replies to some belief. Hence, since for each $i \in N$, both $\beta_i(\cdot)$ and $\beta_i^{-1}(\cdot)$ are monotonic with respect to set inclusion, it follows from Proposition 2(b) that any epistemically robust subset involves only strategies surviving one round of strict elimination. Thus, $\times_{i \in N} \beta_i(\beta_i^{-1}(S_i))$ is the largest epistemically robust subset, while the characterization of the smallest one(s) will be dealt with by Proposition 5.

Our proof shows that Proposition 2(a) can be slightly strengthened, as one only needs the robustness conditions with $p = 1$; as long as there is a $Y \in \mathcal{T}$ such that $C(Y) = X$ and condition (3) holds with $p = 1$, X is CURB.[10] Moreover, although epistemic robustness allows that $Y \in \mathcal{T}$ depends on p, the proof of (b) defines Y independently of p.

The following result shows that FURB sets are characterized by epistemic robustness when player types that do *not* believe with sufficient probability that the others play best replies are removed:

Proposition 3. *The following two statements are equivalent:*

(a) *$X \in \mathcal{S}$ is a FURB set.*
(b) *There exists a $\bar{p} < 1$ such that, for each probability $p \in [\bar{p}, 1]$, there is a type block $Y \in \mathcal{T}$ satisfying condition (2) such that condition (3) holds with equality.*

The block of rationalizable strategies [2,3] is the game's largest FURB set [7]. Thus, it follows from Proposition 3 that epistemic robustness yields a characterization of the block of rationalizable strategies,

---

8 Harsanyi and Selten [31] refer to such pre-images of strategy sets as *stability sets*.
9 In fact, these inclusions hold under the slightly weaker definition of CURB sets in Basu and Weibull [7], in which a player's belief about other players is restricted to be a product measure over the others' pure-strategy sets.
10 In the appendix we also prove that if $p \in (0, 1]$ and $Y \in \mathcal{T}$ are such that $C(Y) = X$ and (3) holds for all $i \in N$, then X is a p-best reply set in the sense of Tercieux [21].

without involving any explicit assumption of common belief of rationality. Instead, only mutual p-belief of rationality and type sets are assumed. Proposition 3 also applies to MINCURB sets, as these sets are FURB. In particular, it follows from Propositions 2(a) and 3 that a strategy block is MINCURB if and only if it is a minimal epistemically robust subset[11].

As much of the literature on CURB sets (recall footnote 1) focuses on minimal ones, we now turn to how smallest CURB sets can be characterized in terms of epistemic robustness. This characterization is presented through Propositions 4 and 5.

Proposition 4 starts from an arbitrary block Y of types and generates an epistemically robust subset by including all beliefs over the opponents' best replies, and all beliefs over opponents' types that have such beliefs over their opponents, and so on. Formally, define for any $Y \in \mathcal{T}$ the sequence $\langle Y(k) \rangle_k$ by $Y(0) = Y$ and, for each $k \in \mathbb{N}$ and $i \in N$,

$$[Y_i(k)] := [Y_i(k-1)] \cup B_i^1 \left(\bigcap\nolimits_{j \neq i} (R_j \cap [Y_j(k-1)]) \right) . \tag{4}$$

Define the correspondence $E : \mathcal{T} \to 2^S$, for any $Y \in \mathcal{T}$, by

$$E(Y) := C \left(\bigcup\nolimits_{k \in \mathbb{N}} Y(k) \right) .$$

We show that the strategy block $E(Y)$ of best replies is epistemically robust and is the smallest CURB set that includes $C(Y)$.[12]

Proposition 4. *Let $Y \in \mathcal{T}$. Then $X = E(Y)$ is the smallest* CURB *set satisfying $C(Y) \subseteq X$. Furthermore, $E(Y)$ is epistemically robust.*

Remark 1. *If the strategy block $C(Y)$ contains strategies that are not rationalizable, then $E(Y)$ will not be* FURB*. Therefore, the epistemic robustness of $E(Y)$ does not follow from Proposition 3; its robustness is established by invoking Proposition 2(b).*

Note that if a strategy block X is epistemically robust, then there exists a type block Y satisfying condition (2) such that condition (3) is satisfied for $p = 1$. Thus, $X = C(Y) = E(Y)$, showing that all epistemically robust strategy blocks can be obtained using the algorithm of Proposition 4.

The final Proposition 5 shows how MINCURB sets can be characterized by epistemically robust subsets obtained by initiating the algorithm of Proposition 4 with a single type profile: a strategy block X is a MINCURB set if and only if (a) the algorithm leads to X from a single type profile, and (b) no single type profile leads to a strict subset of X.

Proposition 5. *$X \in \mathcal{S}$ is a* MINCURB *set if and only if there exists a $t \in T$ such that $E(\{t\}) = X$ and there exists no $t' \in T$ such that $E(\{t'\}) \subsetneq X$.*

Strict Nash equilibria (interpreted as equilibria in beliefs) satisfy 'coordination', in the sense that there is mutual belief about the players' sets of best replies, 'concentration', in the sense that each player has only one best reply, and epistemic robustness (as defined here), implying that each player's set of beliefs about opponent choices contains all probability distributions over opponent strategies that are best replies given their beliefs. In Proposition 5, starting with a single type profile t that corresponds to 'coordination', using the algorithm of Proposition 4 and ending up with $E(\{t\}) = X$ ensures epistemic

11 We thank Peter Wikman for this observation.

12 For each strategy block $X \in \mathcal{S}$, there exists a unique smallest CURB set $X' \in \mathcal{S}$ with $X \subseteq X'$ (that is, X' is a subset of all CURB sets X'' that include X). To see that this holds for all finite games, note that the collection of CURB sets including a given block $X \in \mathcal{S}$ is non-empty and finite, and that the intersection of two CURB sets that include X is again a CURB set including X.

robustness, while the non-existence of $t' \in T$ such that $E(\{t'\})$ is a proper subset of X corresponds to 'concentration'. Hence, these three characteristics of strict Nash equilibria characterize MINCURB sets in Proposition 5.

In order to illustrate Propositions 4 and 5, consider the Nash equilibrium x^* in game (1) in the introduction. This equilibrium corresponds to a type profile (t_1, t_2) where t_1 assigns probability $1/4$ to (l, t_2) and probability $3/4$ to (c, t_2), and where t_2 assigns probability $2/3$ to (u, t_1) and probability $1/3$ to (m, t_1). We have that $C(\{t_1, t_2\}) = \{u, m\} \times \{l, c\}$, while the full strategy space S is the smallest CURB set that includes $C(\{t_1, t_2\})$. Proposition 4 shows that $C(\{t_1, t_2\})$ is not epistemically robust, since it does not coincide with the smallest CURB set that includes it. Recalling the discussion from the introduction: if player 2's belief concerning the behavior of 1 coincides with x_1^*, then 2 is indifferent between his pure strategies l and c, and if 1 assigns equal probability to these two pure strategies of player 2, then 1 will play the unique best reply d, a pure strategy outside the support of the equilibrium. Moreover, if player 2 expects 1 to reason this way, then 2 will play r. Hence, to assure epistemic robustness, starting from type set $\{t_1, t_2\}$, the repeated inclusion of all beliefs over opponents' best replies eventually leads to the smallest CURB set, here S, that includes the Nash equilibrium that was our initial point of departure. By contrast, for the type profile (t'_1, t'_2) where t'_1 assigns probability 1 to (r, t'_2) and t'_2 assigns probability 1 to (d, t'_1) we have that $C(\{t'_1, t'_2\}) = \{(d, r)\}$ coincides with the smallest CURB set that includes it. Thus, the strict equilibrium (d, r) to which (t'_1, t'_2) corresponds is epistemically robust, when viewed as a singleton set. Furthermore, by Proposition 5, $\{(d, r)\}$ is the unique MINCURB set.

Acknowledgments: We thank four anonymous referees, Itai Arieli, Stefano Demichelis, Daisuke Oyama, Olivier Tercieux, Peter Wikman, and seminar participants in Montreal, Paris, Singapore and Tsukuba for helpful comments and suggestions. Voorneveld's research was supported by the Wallander-Hedelius Foundation under grant P2010-0094:1. Weibull's research was supported by the Knut and Alice Wallenberg Research Foundation, and by the Agence Nationale de la Recherche, chaire IDEX ANR-11-IDEX-0002-02.

Author Contributions: The authors contributed equally to this work.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix

Proof of Proposition 1. Let $\mathcal{T}(X) := \{Y \in \mathcal{T} \mid C(Y) = X\}$ denote the collection of type blocks having the property that X is the strategy block of best replies. By the completeness of the probability structure, we have that $\mathcal{T}(X)$ is non-empty if and only if $X \subseteq \times_{i \in N} \beta_i(\beta_i^{-1}(S_i))$. Furthermore, by completeness, if $\mathcal{T}(X)$ is non-empty, then $\mathcal{T}(X)$ has a largest element, $\bar{Y}(X)$, which is constructed by letting $\bar{Y}_i(X) = \{t_i \in T_i \mid C_i(t_i) \subseteq X_i\}$ for all $i \in N$.

By the definition of epistemic robustness, a strategy block $X \in \mathcal{S}$ is *not* epistemically robust if and only if, for all $\bar{p} < 1$, there exists $p \in [\bar{p}, 1]$ such that for all $Y \in \mathcal{T}(X)$, there exists $i \in N$ such that

$$B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) \nsubseteq [Y_i] \,.$$

Hence, $X \in \mathcal{S}$ is *not* epistemic robust if and only if

(∗) $X \nsubseteq \times_{i \in N} \beta_i(\beta_i^{-1}(S_i))$ so that $\mathcal{T}(X) = \emptyset$, or

(∗∗) $X \subseteq \times_{i \in N} \beta_i(\beta_i^{-1}(S_i))$ so that $\mathcal{T}(X) \neq \emptyset$, and, for all $\bar{p} < 1$, there exists $p \in [\bar{p}, 1]$ such that if $Y \in \mathcal{T}(X)$, then there exist $i \in N$ and $t_i \notin Y_i$ such that

$$[\{t_i\}] \subseteq B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) .$$

(b) implies (a). Assume that, for all $\bar{p} < 1$, there exists $p \in [\bar{p}, 1]$ such that if $Y \in \mathcal{T}(X)$, then there exist $i \in N$ and $t_i \in T_i$ such that $C(t_i) \nsubseteq X_i$ and $[\{t_i\}] \subseteq B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right)$. Note that if $Y \in \mathcal{T}(X)$ and $C(t_i) \nsubseteq X_i$, then $t_i \notin Y_i$.

Either $\mathcal{T}(X) = \emptyset$, so that $(*)$ is satisfied, or $\mathcal{T}(X) \neq \emptyset$ and, for all $\bar{p} < 1$, there exists $p \in [\bar{p}, 1]$ such that if $Y \in \mathcal{T}(X)$, then there exist $i \in N$ and $t_i \notin Y_i$ such that $[\{t_i\}] \subseteq B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right)$, so that $(**)$ is satisfied.

(a) implies (b). Assume that $(*)$ or $(**)$ is satisfied.

Assume that $(*)$ is satisfied, and fix $p < 1$. Then, it holds trivially that if $Y \in \mathcal{T}(X)$, then there exist $i \in N$ and $t_i \in T_i$ such that $C(t_i) \nsubseteq X_i$ and $[\{t_i\}] \subseteq B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right)$.

Assume that $(**)$ is satisfied. Then, since $\bar{Y}(X) \in \mathcal{T}(X)$, it must also hold that for all $\bar{p} < 1$, there exist $p(\bar{p}) \in [\bar{p}, 1]$, $i(\bar{p}) \in N$ and $t_{i(\bar{p})}(\bar{p}) \notin \bar{Y}_{i(\bar{p})}(X)$ such that $[\{t_{i(\bar{p})}(\bar{p})\}] \subseteq B_{i(\bar{p})}^{p(\bar{p})} \left(\bigcap_{j \neq i(\bar{p})} (R_j \cap [\bar{Y}_j(X)]) \right)$. By the definition of $\bar{Y}(X)$, $C(t_{i(\bar{p})}(\bar{p})) \nsubseteq X_{i(\bar{p})}$. It is sufficient to construct, for all $\bar{p} < 1$ and $Y \in \mathcal{T}(X)$, a type $t_{i(\bar{p})} \in T_i$ such that $C(t_{i(\bar{p})}) = C(t_{i(\bar{p})}(\bar{p}))$ and $[\{t_{i(\bar{p})}\}] \subseteq B_{i(\bar{p})}^{p(\bar{p})} \left(\bigcap_{j \neq i(\bar{p})} (R_j \cap [Y_j]) \right)$.

For all $s_{-i(\bar{p})} \in X_{-i(\bar{p})}$ with $\text{marg}_{S_{-i(\bar{p})}} \mu_{-i(\bar{p})}(t_{i(\bar{p})}(\bar{p}))(s_{-i(\bar{p})}) > 0$, select $t_{-i(\bar{p})} \in Y_{-i(\bar{p})}$ such that $s_j \in C_j(t_j)$ for all $j \neq i(\bar{p})$ (which exists since $C(Y) = X$) and let

$$\mu_{i(\bar{p})}(t_{i(\bar{p})})(s_{-i(\bar{p})}, t_{-i(\bar{p})}) = \text{marg}_{S_{-i(\bar{p})}} \mu_{-i(\bar{p})}(t_{i(\bar{p})}(\bar{p}))(s_{-i(\bar{p})}) .$$

For all $s_{-i(\bar{p})} \notin X_{-i(\bar{p})}$ with $\text{marg}_{S_{-i(\bar{p})}} \mu_{-i(\bar{p})}(t_{i(\bar{p})}(\bar{p}))(s_{-i(\bar{p})}) > 0$, select $t_{-i(\bar{p})} \in Y_{-i(\bar{p})}$ arbitrary and let again

$$\mu_{i(\bar{p})}(t_{i(\bar{p})})(s_{-i(\bar{p})}, t_{-i(\bar{p})}) = \text{marg}_{S_{-i(\bar{p})}} \mu_{-i(\bar{p})}(t_{i(\bar{p})}(\bar{p}))(s_{-i(\bar{p})}) .$$

Then $\text{marg}_{S_{-i(\bar{p})}} \mu_{-i(\bar{p})}(t_{i(\bar{p})})(s_{-i(\bar{p})}) = \text{marg}_{S_{-i(\bar{p})}} \mu_{-i(\bar{p})}(t_{i(\bar{p})}(\bar{p}))(s_{-i(\bar{p})})$, implying that $C(t_{i(\bar{p})}) = C(t_{i(\bar{p})}(\bar{p}))$. Furthermore, by the construction of $t_{i(\bar{p})}$:

$$\begin{aligned}
&\mu_{i(\bar{p})}(t_{i(\bar{p})}) \left(\{(s_{-i(\bar{p})}, t_{-i(\bar{p})}) \in S_{-i(\bar{p})} \times Y_{-i(\bar{p})} \mid s_j \in C_j(t_j) \text{ for all } j \neq i(\bar{p})\} \right) \\
= &\ \mu_{i(\bar{p})}(t_{i(\bar{p})}) \left(X_{-i(\bar{p})} \times T_{-i(\bar{p})} \right) \\
= &\ \mu_{i(\bar{p})}(t_{i(\bar{p})}(\bar{p})) \left(X_{-i(\bar{p})} \times T_{-i(\bar{p})} \right) \\
\geq &\ \mu_{i(\bar{p})}(t_{i(\bar{p})}(\bar{p})) \left(\{(s_{-i(\bar{p})}, t_{-i(\bar{p})}) \in S_{-i(\bar{p})} \times \bar{Y}_{-i(\bar{p})}(X) \mid s_j \in C_j(t_j) \text{ for all } j \neq i(\bar{p})\} \right) \\
\geq &\ p(\bar{p}),
\end{aligned}$$

since $C(Y) = X = C(\bar{Y}(X))$.[13] Thus, $[\{t_{i(\bar{p})}\}] \subseteq B_{i(\bar{p})}^{p(\bar{p})} \left(\bigcap_{j \neq i(\bar{p})} (R_j \cap [Y_j]) \right)$. □

Proof of Proposition 2. *Part (a).* By assumption, there is a $Y \in \mathcal{T}$ with $C(Y) = X$ such that for each $i \in N$, $B_i^1 \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) \subseteq [Y_i]$.

Fix $i \in N$, and consider any $\sigma_{-i} \in \mathcal{M}(X_{-i})$. Since $C(Y) = X$, it follows that, for each $s_{-i} \in S_{-i}$ with $\sigma_{-i}(s_{-i}) > 0$, there exists $t_{-i} \in Y_{-i}$ such that, for all $j \neq i$, $s_j \in C_j(t_j)$. Hence, since the probability structure is complete, there exists a

$$\omega \in B_i^1 \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) \subseteq [Y_i]$$

13 To see that the first equality in the expression above holds, note first that, since $C(Y) = X$,

$$\{(s_{-i(\bar{p})}, t_{-i(\bar{p})}) \in S_{-i(\bar{p})} \times Y_{-i(\bar{p})} \mid s_j \in C_j(t_j) \text{ for all } j \neq i(\bar{p})\} \subseteq X_{-i(\bar{p})} \times T_{-i(\bar{p})} .$$

However, by construction, for any $(s_{-i(\bar{p})}, t_{-i(\bar{p})}) \in X_{-i(\bar{p})} \times T_{-i(\bar{p})}$ assigned positive probability by $\mu_{i(\bar{p})}(t_{i(\bar{p})})$, it is the case that $t_{-i(\bar{p})} \in Y_{-i(\bar{p})}$ and $s_j \in C_j(t_j)$ for all $j \neq i(\bar{p})$. Hence, the two sets are given the same probability by $\mu_{i(\bar{p})}(t_{i(\bar{p})})$.

with $\text{marg}_{S_{-i}}\, \mu_i(\mathbf{t}_i(\omega)) = \sigma_{-i}$. So

$$\beta_i(X_{-i}) := \beta_i(\mathcal{M}(X_{-i})) \subseteq \bigcup\nolimits_{t_i \in Y_i} \beta_i(\text{marg}_{S_{-i}}\, \mu_i(t_i)) := C_i(Y_i) = X_i\,.$$

Since this holds for all $i \in N$, X is a CURB set.

Part (b). Assume that $X \in \mathcal{S}$ is a CURB set, i.e., X satisfies $\beta(X) \subseteq X$. It suffices to prove that $\times_{i \in N} \beta_i(\beta_i^{-1}(X_i)) \subseteq X$ is epistemically robust. That it is the largest epistemically robust subset of X then follows immediately from the fact that, for each $i \in N$, both $\beta_i(\cdot)$ and $\beta_i^{-1}(\cdot)$ are monotonic with respect to set inclusion.

Define $Y \in \mathcal{T}$ by taking, for each $i \in N$, $Y_i := \{t_i \in T_i \mid C_i(t_i) \subseteq X_i\}$. Since the probability structure is complete, it follows that $C_i(Y_i) = \beta_i(\beta_i^{-1}(X_i))$. For notational convenience, write $X'_i = \beta_i(\beta_i^{-1}(X_i))$ and $X' = \times_{i \in N} X'_i$. Since the game is finite, there is, for each player $i \in N$, a $\underline{p}_i \in (0,1)$ such that $\beta_i(\sigma_{-i}) \subseteq \beta_i(X'_{-i})$ for all $\sigma_{-i} \in \mathcal{M}(S_{-i})$ with $\sigma_{-i}(X'_{-i}) \geq \underline{p}_i$. Let $\underline{p} = \max\{\underline{p}_1, \ldots, \underline{p}_n\}$.

We first show that $\beta(X') \subseteq X'$. By definition, $X' \subseteq X$, so for each $i \in N$: $\mathcal{M}(X'_{-i}) \subseteq \mathcal{M}(X_{-i})$. Moreover, as $\beta(X) \subseteq X$ and, for each $i \in N$, $\beta_i(X_i) := \beta_i(\mathcal{M}(X_{-i}))$, it follows that $\mathcal{M}(X_{-i}) \subseteq \beta_i^{-1}(X_i)$. Hence, for each $i \in N$,

$$\beta_i(X'_i) := \beta_i(\mathcal{M}(X'_{-i})) \subseteq \beta_i(\mathcal{M}(X_{-i})) \subseteq \beta_i(\beta_i^{-1}(X_i)) = X'_i\,.$$

For all $p \in [\underline{p}, 1]$ and $i \in N$, we have that

$$\begin{aligned}
& B_i^p\left(\bigcap\nolimits_{j \neq i}(R_j \cap [Y_j])\right) \\
= \;& B_i^p\left(\bigcap\nolimits_{j \neq i}\{\omega \in \Omega \mid \mathbf{s}_j(\omega) \in C_j(\mathbf{t}_j(\omega)) \subseteq X'_j\}\right) \\
\subseteq \;& \left\{\omega \in \Omega \mid \mu_i(\mathbf{t}_i(\omega))\{\omega_{-i} \in \Omega_{-i} \mid \text{for all } j \neq i,\ \mathbf{s}_j(\omega) \in X'_j\} \geq p\right\} \\
\subseteq \;& \{\omega \in \Omega \mid \text{marg}_{S_{-i}}\, \mu_i(\mathbf{t}_i(\omega))(X'_{-i}) \geq p\} \\
\subseteq \;& \{\omega \in \Omega \mid C_i(\mathbf{t}_i(\omega)) \subseteq \beta_i(X'_{-i})\} \\
\subseteq \;& \{\omega \in \Omega \mid C_i(\mathbf{t}_i(\omega)) \subseteq X'_{-i}\} \;=\; [Y_i],
\end{aligned}$$

using $\beta(X') \subseteq X'$. □

For $X \in \mathcal{S}$ and $p \in (0,1]$, write, for each $i \in N$,

$$\begin{aligned}
\beta_i^p(X_{-i}) := \{s_i \in S_i \mid \exists \sigma_{-i} \in \mathcal{M}(S_{-i}) \text{ with } \sigma_{-i}(X_{-i}) \geq p \\
\text{such that } u_i(s_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i})\ \forall s'_i \in S_i\}\,.
\end{aligned}$$

Let $\beta^p(X) := \beta_1^p(X_{-1}) \times \cdots \times \beta_n^p(X_{-n})$. Following Tercieux [21], a set $X \in \mathcal{S}$ is a *p-best reply set* if $\beta^p(X) \subseteq X$.

Claim: *Let $X \in \mathcal{S}$ and $p \in (0,1]$. If $Y \in \mathcal{T}$ is such that $C(Y) = X$ and condition (3) holds for each $i \in N$, then X is a p-best reply set.*

Proof. By assumption, there is a $Y \in \mathcal{T}$ with $C(Y) = X$ such that for each $i \in N$, $B_i^p\left(\bigcap_{j \neq i}(R_j \cap [Y_j])\right) \subseteq [Y_i]$.

Fix $i \in N$ and consider any $\sigma_{-i} \in \mathcal{M}(S_{-i})$ with $\sigma_{-i}(X_{-i}) \geq p$. Since $C(Y) = X$, it follows that, for each $s_{-i} \in X_{-i}$, there exists $t_{-i} \in Y_{-i}$ such that $s_j \in C_j(t_j)$ for all $j \neq i$. Hence, since the probability structure is complete, there exists a

$$\omega \in B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) \subseteq [Y_i]$$

with $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega)) = \sigma_{-i}$. Thus, by definition of $\beta_i^p(X_{-i})$:

$$\beta_i^p(X_{-i}) \subseteq \bigcup_{t_i \in Y_i} \beta_i(\text{marg}_{S_{-i}} \mu_i(t_i)) := C_i(Y_i) = X_i \,.$$

Since this holds for all $i \in N$, X is a p-best reply set. □

Proof of Proposition 3. *(b) implies (a).* By assumption, there is a $Y \in \mathcal{T}$ with $C(Y) = X$ such that for all $i \in N$, $B_i^1 \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) = [Y_i]$.

Fix $i \in N$. Since $C(Y) = X$, and the probability structure is complete, there exists, for any $\sigma_{-i} \in \mathcal{M}(S_{-i})$, an

$$\omega \in B_i^1 \left(\bigcap_{j \neq i} (R_j \cap [Y_j]) \right) = [Y_i]$$

with $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega)) = \sigma_{-i}$ if and only if $\sigma_{-i} \in \mathcal{M}(X_{-i})$. Thus,

$$\beta_i(X_{-i}) := \beta_i(\mathcal{M}(X_{-i})) = \bigcup_{t_i \in Y_i} \beta_i(\text{marg}_{S_{-i}} \mu_i(t_i)) := C_i(Y_i) = X_i \,.$$

Since this holds for all $i \in N$, X is a FURB set.

(a) implies (b). Assume that $X \in \mathcal{S}$ satisfies $X = \beta(X)$. Since the game is finite, there exists, for each player $i \in N$, a $\underline{p}_i \in (0,1)$ such that $\beta_i(\sigma_{-i}) \subseteq \beta_i(X_{-i})$ if $\sigma_{-i}(X_{-i}) \geq \underline{p}_i$. Let $\underline{p} = \max\{\underline{p}_1, \ldots, \underline{p}_n\}$.

For each $p \in [\underline{p}, 1]$, construct the sequence of type blocks $\langle Y^p(k) \rangle_k$ as follows: for each $i \in N$, let $Y_i^p(0) = \{t_i \in T_i \mid C_i(t_i) \subseteq X_i\}$. Using continuity of μ_i, the correspondence $C_i : T_i \rightrightarrows S_i$ is upper hemi-continuous. Thus $Y_i^p(0) \subseteq T_i$ is closed, and, since T_i is compact, so is $Y_i^p(0)$. There exists a closed set $Y_i^p(1) \subseteq T_i$ such that

$$[Y_i^p(1)] = B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j^p(0)]) \right) .$$

It follows that $Y_i^p(1) \subseteq Y_i^p(0)$. Since $Y_i^p(0)$ is compact, so is $Y_i^p(1)$. By induction,

$$[Y_i^p(k)] = B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j^p(k-1)]) \right) \tag{A1}$$

defines, for each player i, a decreasing chain $\left\langle Y_i^p(k) \right\rangle_k$ of compact and non-empty subsets: $Y_i^p(k+1) \subseteq Y_i^p(k)$ for all k. By the finite-intersection property, $Y_i^p := \bigcap_{k \in \mathbb{N}} Y_i^p(k)$ is a non-empty and compact subset of T_i. For each k, let $Y^p(k) = \times_{i \in N} Y_i^p(k)$ and let $Y^p := \bigcap_{k \in \mathbb{N}} Y^p(k)$. Again, these are non-empty and compact sets.

Next, $C(Y^p(0)) = \beta(X)$, since the probability structure is complete. Since X is FURB, we thus have $C(Y^p(0)) = X$. For each $i \in N$,

$$[Y_i^p(1)] \subseteq \{\omega \in \Omega \mid \text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega))(X_{-i}) \geq p\} \,,$$

implying that $C_i(Y_i^p(1)) \subseteq \beta_i(X_{-i}) = X_i$ by the construction of $\underline{p}$. Moreover, since the probability structure is complete, for each $i \in N$ and $\sigma_{-i} \in \mathcal{M}(X_{-i})$, there exists $\omega \in [Y_i^p(1)] = B_i^p(\bigcap_{j \neq i}(R_j \cap [Y_j^p(0)]))$ with $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega)) = \sigma_{-i}$, implying that $C_i(Y_i^p(1)) \supseteq \beta_i(X_{-i}) = X_i$. Hence, $C_i(Y_i^p(1)) = \beta_i(X_{-i}) = X_i$. By induction, it holds for all $k \in N$ that $C(Y^p(k)) = \beta(X) = X$. Since $\left\langle Y_i^p(k) \right\rangle_k$ is a decreasing chain, we also have that $C(Y^p) \subseteq X$. The converse inclusion follows by upper hemi-continuity of the correspondence C. To see this, suppose that $x^o \in X$ but $x^o \notin C(Y^p)$. Since $x^o \in X$,

$x^o \in C(Y^p(k))$ for all k. By the Axiom of Choice: for each k, there exists a $y_k \in Y^p(k)$ such that $(y_k, x^o) \in graph(C)$. By the Bolzano–Weierstrass Theorem, we can extract a convergent subsequence for which $y_k \to y^o$, where $y^o \in Y^p$, since Y^p is closed. Moreover, since the correspondence C is closed-valued and upper hemi-continuous, with S compact (it is in fact finite), $graph(C) \subseteq T \times S$ is closed, and thus $(y^o, x^o) \in graph(C)$, contradicting the hypothesis that $x^o \notin C(Y^p)$. This establishes the claim that $C(Y^p) \subseteq X$.

It remains to prove that, for each $i \in N$, condition (3) holds with equation for Y^p. Fix $i \in N$, and let

$$E_k = \bigcap\nolimits_{j \neq i} (R_j \cap [Y_j^p(k)]) \qquad \text{and} \qquad E = \bigcap\nolimits_{j \neq i} (R_j \cap [Y_j^p]) \,.$$

Since, for each $j \in N$, $\langle Y_j^p(k) \rangle_k$ is a decreasing chain with limit Y_j^p, it follows that $\langle E_k \rangle_k$ is a decreasing chain with limit E.

To show $B_i^p(E) \subseteq [Y_i^p]$, note that by (A1) and monotonicity of B_i^p, we have, for each $k \in \mathbb{N}$, that

$$B_i^p(E) \subseteq B_i^p(E_{k-1}) = [Y_i^p(k)] \,.$$

As the inclusion holds for all $k \in \mathbb{N}$:

$$B_i^p(E) \subseteq \bigcap\nolimits_{k \in \mathbb{N}} [Y_i^p(k)] = [Y_i^p] \,.$$

To show $B_i^p(E) \supseteq [Y_i^p]$, assume that $\omega \in [Y_i^p]$.[14] This implies that $\omega \in [Y_i^p(k)]$ for all k, and, using (A1): $\omega \in B_i^p(E_k)$ for all k. Since $E_k = \Omega_i \times \text{proj}_{\Omega_{-i}} E_k$, we have that $E_k^{\omega_i} = \text{proj}_{\Omega_{-i}} E_k$. It follows that

$$\mu_i(\mathbf{t}_i(\omega))(\text{proj}_{\Omega_{-i}} E_k) \geq p \quad \text{for all } k \,.$$

Thus, since $\langle E_k \rangle_k$ is a decreasing chain with limit E,

$$\mu_i(\mathbf{t}_i(\omega))(\text{proj}_{\Omega_{-i}} E) \geq p \,.$$

Since $E = \Omega_i \times \text{proj}_{\Omega_{-i}} E$, we have that $E^{\omega_i} = \text{proj}_{\Omega_{-i}} E$. Hence, the inequality implies that $\omega \in B_i^p(E)$. □

Proof of Proposition 4. Let $X \in \mathcal{S}$ be the smallest CURB set containing $C(Y)$: (i) $C(Y) \subseteq X$ and $\beta(X) \subseteq X$ and (ii) there exists no $X' \in \mathcal{S}$ with $C(Y) \subseteq X'$ and $\beta(X') \subseteq X' \subsetneq X$. We must show that $X = E(Y)$.

Consider the sequence $\langle Y(k) \rangle_k$ defined by $Y(0) = Y$ and condition (4) for each $k \in \mathbb{N}$ and $i \in N$. We show, by induction, that $C(Y(k)) \subseteq X$ for all $k \in \mathbb{N}$. By assumption, $Y(0) = Y \in \mathcal{T}$ satisfies this condition. Assume that $C(Y(k-1)) \subseteq X$ for some $k \in \mathbb{N}$, and fix $i \in N$. Then, $\forall j \neq i$, $\beta_j(\text{marg}_{S_{-j}} \mu_j(\mathbf{t}_j(\omega))) \subseteq X_j$ if $\omega \in [Y_j(k-1)]$ and $\mathbf{s}_j(\omega) \in X_j$ if, in addition, $\omega \in R_j$. Hence, if $\omega \in B_i^1(\bigcap_{j \neq i}(R_j \cap [Y_j(k-1)]))$, then $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega)) \in \mathcal{M}(X_{-i})$ and $C_i(\mathbf{t}_i(\omega)) \subseteq \beta_i(X_{-i}) \subseteq X_i$. Since this holds for all $i \in N$, we have $C(Y(k)) \subseteq X$. This completes the induction.

Secondly, since the sequence $\langle Y(k) \rangle_k$ is non-decreasing and $C(\cdot)$ is monotonic with respect to set inclusion, and the game is finite, there exist a $k' \in \mathbb{N}$ and some $X' \subseteq X$ such that $C(Y(k)) = X'$ for all $k \geq k'$. Let $k > k'$ and consider any player $i \in N$. Since the probability structure is complete, there exists, for each $\sigma_{-i} \in \mathcal{M}(X'_{-i})$ a state $\omega \in [Y_i(k)]$ with $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega)) = \sigma_{-i}$, implying that $\beta_i(X'_{-i}) \subseteq C_i(Y_i(k)) = X'_i$. Since this holds for all $i \in N$, $\beta(X') \subseteq X'$. Therefore, if $X' \subsetneq X$ would hold, then this would contradict that there exists no $X' \in \mathcal{S}$ with $C(Y) \subseteq X'$ such that $\beta(X') \subseteq X' \subsetneq X$. Hence, $X = C(\bigcup_{k \in \mathbb{N}} Y(k)) = E(Y)$.

---

[14] We thank Itai Arieli for suggesting this proof of the reversed inclusion, shorter than our original proof. A proof of both inclusions can also be based on property (8) of Monderer and Samet [30].

Write $X = E(Y)$. To establish that X is epistemically robust, by Proposition 2(b), it is sufficient to show that

$$X \subseteq \times_{i \in N} \beta_i(\beta_i^{-1}(X_i))\,,$$

keeping in mind that, for all $X' \in \mathcal{S}$, $X' \supseteq \times_{i \in N} \beta_i(\beta_i^{-1}(X'_i))$.

Fix $i \in \mathbb{N}$. Define $Y'_i \in \mathcal{T}$ by taking $Y'_i := \{t_i \in T_i \mid C_i(t_i) \subseteq X_i\}$. Since the probability structure is complete, it follows that $C_i(Y'_i) = \beta_i(\beta_i^{-1}(X_i))$. Furthermore, for all $k \in \mathbb{N}$, $Y(k) \subseteq Y'$ and, hence, $\bigcup_{k \in \mathbb{N}} Y(k) \subseteq Y'$. This implies that

$$X_i = C\left(\bigcup\nolimits_{k \in \mathbb{N}} Y(k)\right) \subseteq C_i(Y'_i) = \beta_i(\beta_i^{-1}(X_i)),$$

since $C_i(\cdot)$ is monotonic with respect to set inclusion. □

Proof of Proposition 5. *(Only if)* Let $X \in \mathcal{S}$ be a MINCURB set. Let $t \in T$ satisfy $\mathrm{marg}_{S_{-i}}\, \mu_i(t_i)(X_{-i}) = 1$ for all $i \in N$. By construction, $C(\{t\}) \subseteq X$, as X is a CURB set. By Proposition 4, $E(\{t\})$ is the *smallest* CURB set with $C(\{t\}) \subseteq E(\{t\})$. But then $E(\{t\}) \subseteq X$. The inclusion cannot be strict, as X is a MINCURB set. Hence, there exists a $t \in T$ such that $E(\{t\}) = X$. Moreover, as $E(\{t'\})$ is a CURB set for all $t' \in T$ and X is a MINCURB set, there exists no $t' \in T$ such that $E(\{t'\}) \subsetneq X$.

(If) Assume that there exists a $t \in T$ such that $E(\{t\}) = X$ and there exists no $t' \in T$ such that $E(\{t'\}) \subsetneq X$. Since $E(\{t\}) = X$, it follows from Proposition 4 that X is a CURB set. To show that X is a *minimal* CURB set, suppose—to the contrary—that there is a CURB set $X' \subsetneq X$. Let $t' \in T$ be such that $\mathrm{marg}_{S_{-i}}\, \mu_i(t'_i)(X'_{-i}) = 1$ for each $i \in N$. By construction, $C(\{t'\}) \subseteq X'$, so X' is a CURB set containing $C(\{t'\})$. By Proposition 4, $E(\{t'\})$ is the smallest CURB set containing $C(\{t'\})$. However, by assumption, there exists no $t' \in T$ such that $E(\{t'\}) \subsetneq X$, so it must be that $E(\{t'\}) \supseteq X$. This contradicts $X' \subsetneq X$. □

References

1. Nash, J.F. Non-Cooperative Games. Ph.D. Thesis, Princeton University, Princeton, NJ, USA, 1950.
2. Bernheim, D. Rationalizable strategic behavior. *Econometrica* **1984**, *52*, 1007–1028.
3. Pearce, D.G. Rationalizable strategic behavior and the problem of perfection. *Econometrica* **1984**, *52*, 1029–1050.
4. Aumann, R.J.; Brandenburger, A. Epistemic conditions for Nash equilibrium. *Econometrica* **1995**, *63*, 1161–1180.
5. Myerson, R.; Weibull, J.W. Tenable strategy blocks and settled equilibria. *Econometrica* **2015**, *83*, 943–976.
6. Tan, T.; Werlang, S.R.C. The Bayesian foundations of solution concepts of games. *J. Econ. Theory* **1988**, *45*, 370–391.
7. Basu, K.; Weibull, J.W. Strategy subsets closed under rational behavior. *Econ. Lett.* **1991**, *36*, 141–146.
8. Hurkens, S. Learning by forgetful players. *Games Econ. Behav.* **1995**, *11*, 304–329.
9. Sanchirico, C. A probabilistic model of learning in games. *Econometrica* **1996**, *64*, 1375–1394.
10. Young, P.H. *Individual Strategy and Social Structure*; Princeton University Press: Princeton, NJ, USA, 1998.
11. Fudenberg, D.; Levine, D.K. *The Theory of Learning in Games*; MIT Press: Cambridge, MA, USA, 1998.
12. Blume, A. Communication, risk, and efficiency in games. *Games Econ. Behav.* **1998**, *22*, 171–202.
13. Hurkens, S. Multi-sided pre-play communication by burning money. *J. Econ. Theory* **1996**, *69*, 186–197.
14. Galeotti, A.; Goyal, S.; Kamphorst, J.J.A. Network formation with heterogeneous players. *Games Econ. Behav.* **2006**, *54*, 353–372.
15. Ritzberger, K.; Weibull, J.W. Evolutionary selection in normal-form games. *Econometrica* **1995**, *63*, 1371–1399.
16. Kohlberg, E.; Mertens J.-F. On the strategic stability of equilibria. *Econometrica* **1986**, *54*, 1003–1037.
17. Asheim, G.B. *The Consistent Preferences Approach to Deductive Reasoning in Games*; Springer: Dordrecht, The Netherlands, 2006.
18. Brandenburger, A.; Friedenberg, A.; Keisler, H.J. Admissibility in games. *Econometrica* **2008**, *76*, 307–352.
19. Kalai, E.; Samet, D. Persistent equilibria in strategic games. *Int. J. Game Theory* **1984**, *14*, 41–50.

20. Van Damme, E. Strategic Equilibrium. In *Handbook of Game Theory*; Aumann, R.J., Hart, S., Eds.; Chapter 41; Elsevier: Amsterdam, The Netherlands, 2002; Volume 3.
21. Tercieux, O. *p*-Best response set. *J. Econ. Theory* **2006**, *131*, 45–70.
22. Zambrano, E. Epistemic conditions for rationalizability. *Games Econ. Behav.* **2008**, *63*, 395–405.
23. Hu, T.-W. On *p*-rationalizability and approximate common certainty of rationality. *J. Econ. Theory* **2007**, *136*, 379–391.
24. Osborne, M.J.; Rubinstein, A. *A Course in Game Theory*; The MIT Press: Cambridge, MA, USA, 1994.
25. Ferguson, T.S. *Mathematical Statistics, A Decision Theoretic Approach*; Academic Press: New York, NY, USA, 1967.
26. Van Damme, E. *Refinements of the Nash Equilibrium Concept*; Springer: Berlin, Germany, 1983.
27. Brandenburger, A.; Friedenberg, A.; Keisler, H.J. *Admissibility in Games*; New York University: New York, NY, USA; Washington University: St. Louis, MO, USA; University of Wisconsin-Madison: Madison, WI, USA, 12 December 2004.
28. Brandenburger, A.; Dekel, E. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* **1993**, *59*, 189–198.
29. Mertens, J.-F.; Zamir, S. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory* **1985**, *14*, 1–29.
30. Monderer, D.; Samet, D. Approximating common knowledge with common beliefs. *Games Econ. Behav.* **1989**, *1*, 170–190.
31. Harsanyi, J.C.; Selten, R. *A General Theory of Equilibrium Selection in Games*; The MIT Press: Cambridge, MA, USA, 1988.
32. Mertens, J.-F. Stable equilibria—A reformulation, part I: Definition and basic properties. *Math. Oper. Res.* **1989**, *14*, 575–625.
33. Demichelis, S.; Ritzberger, K. From Evolutionary to Strategic Stability. *J. Econ. Theory* **2003**, *113*, 51–75.