



HAL
open science

Epistemically stable strategy sets

Geir B. Asheim, Mark Voorneveld, Jörgen Weibull

► **To cite this version:**

Geir B. Asheim, Mark Voorneveld, Jörgen Weibull. Epistemically stable strategy sets. 2009. hal-00440098

HAL Id: hal-00440098

<https://hal.science/hal-00440098>

Preprint submitted on 9 Dec 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



ÉCOLE POLYTECHNIQUE
CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

EPISTEMICALLY STABLE STRATEGY SETS

Geir B. ASHEIM
Mark VOORNEVELD
Jörgen W. WEIBULL

December 2009

Cahier n° 2009-53

DEPARTEMENT D'ECONOMIE

Route de Saclay
91128 PALAISEAU CEDEX
(33) 1 69333033

<http://www.enseignement.polytechnique.fr/economie/>
<mailto:chantal.poujouly@polytechnique.edu>

EPISTEMICALLY STABLE STRATEGY SETS¹

Geir B. ASHEIM²
Mark VOORNEVELD³
Jörgen W. WEIBULL⁴

October 2009

Cahier n° 2009-53

Abstract : This paper provides a definition of epistemic stability of sets of strategy profiles, and uses it to characterize variants of curb sets in finite games, including the set of rationalizable strategies and minimal curb sets.

Keywords: Epistemic game theory; epistemic stability; rationalizability; closedness under rational behavior; mutual p-belief.

JEL Classification Numbers: C72; D83

¹This manuscript is a slightly revised version of “Epistemically robust sets”. We have replace the phrase “epistemically robust” by “epistemically stable” and added a comment on its relationship to strategic stability. We thank Itai Arieli, Stefano Demichelis, Daisuke Oyama and Olivier Tercieux for helpful comments and suggestions. Financial support from the Knut and Alice Wallenberg Foundation and the Wallander/Hedelius Foundation is gratefully acknowledged.

² Department of Economics, University of Oslo, P.O. Box 1095 Blindern, NO-0317 Oslo, Norway (e-mail: g.b.asheim@econ.uio.no).

³ Department of Economics, Stockholm School of Economics, Box 6501, SE-113 83 Stockholm, Sweden (e-mail: nemv@hhs.se).

⁴ Department of Economics, Stockholm School of Economics, Box 6501, SE-113 83 Stockholm, Sweden and École Polytechnique, Department of Economics, 91128 Palaiseau Cedex, France (e-mail: nejw@hhs.se).

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 (1950) 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 — more standard one in economics — 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 (1984), Pearce (1984), Aumann and Brandenburger (1995).

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, each player’s belief is supposed to correspond to *specific* randomizations over the others’ strategies. In particular, given her beliefs, a player may have multiple pure strategies that maximize her expected payoff. Hence, any randomization over these is a best reply. Yet in Nash equilibrium, each player’s belief singles out those randomizations over the others’ pure best replies that serve to keep their opponents indifferent across their mixed-strategy supports. In addition, each 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 fellow players.

Our aim is to formalize a notion of epistemic stability that relaxes these requirements. In order to achieve this, we have to move away from point-valued to set-valued solutions. Roughly speaking, we define a set X of pure strategy profiles as epistemically stable if there exists a corresponding set Y of profiles of “player types” such that:

- (i) The strategies in X coincide with the best replies of the player types in Y .

- (ii) The set Y contains any player type that believes with sufficient probability that the others are of types in Y and choose best replies.

While (ii) specifies a stable set of beliefs, (i) specifies a stable set of strategies in response to it.

Any strict Nash equilibrium, viewed as a singleton product set, is epistemically stable in this sense. Each player is then believed by the others to choose her unique best reply to the others' actions. To deviate to any other action would be strictly worse, and remains so, as long as the player is *sufficiently sure* that the others stick to their actions. By contrast, non-strict Nash equilibria by definition have alternative best replies and are consequently not epistemically stable: players who strive to maximize their expected payoffs might well choose such alternative best replies even if they are sure that others are playing their equilibrium strategies. As will be shown below, every epistemically stable set contains at least one strategically stable set.

The notion of persistent retracts (Kalai and Samet (1984)) goes part of the way towards epistemic stability. 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 to small belief perturbations, but admit alternative best replies outside the set.

Full epistemic stability is achieved by variants of CURB sets. A CURB set — mnemonic for ‘closed under rational behavior’ — is a Cartesian product of pure-strategy sets, one for each player, that includes all best replies to all probability distributions over the strategies in the set.¹ Hence, if a player believes that her opponents stick to strategies from their components of a CURB set, then her component contains all her best replies, so she'd better stick to her strategies as well.

A Cartesian product of pure-strategy sets is fixed under rational behavior (FURB) if each player's component not only contains, but is identical with the set of best replies to all probability distributions over the set. Hence, FURB sets are the natural set-valued generalization of strict Nash equilibria. Basu and Weibull (1991) — who

¹ CURB sets and variants were introduced by Basu and Weibull (1991) and became of importance in the literature on strategy adaptation in finite games. Several classes of adaptation processes eventually settle down in a minimal CURB set; cf. Hurkens (1995), Sanchirico (1996), Young (1998), and Fudenberg and Levine (1998). Such sets also give appealing results in communication games (Hurkens, 1996; Blume, 1998) and network formation games (Galeotti, Goyal, and Kamphorst, 2006). For closure properties under generalizations of the best-response correspondence, see Ritzberger and Weibull (1995).

refer to FURB sets as ‘tight’ CURB sets — show that minimal CURB sets and the product set of rationalizable strategies (Bernheim, 1984; Pearce, 1984) are important special cases of FURB sets.

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 stable set 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 $x^* = (x_1^*, x_2^*) = ((\frac{2}{3}, \frac{1}{3}, 0), (\frac{1}{4}, \frac{3}{4}, 0))$ 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. Moreover, if player 2 expects 1 to reason this way, then 2 will play r : the smallest epistemically stable set containing the support of the mixed equilibrium x^* is the entire pure 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 stable. In this game, all pure strategies are rationalizable, $S = S_1 \times S_2$ is a FURB set, and the game’s unique minimal CURB set and unique minimal FURB set is $T = \{d\} \times \{r\}$.

Our results on epistemic stability can be heuristically described as follows. By Proposition 1(a), epistemically stable sets must be CURB sets. Conversely, although CURB sets² may involve strategies that are not best replies — e.g., strategies that are strictly dominated — every CURB set contains an epistemically stable subset. Proposition 1(b) characterizes the largest one, whereas the smallest one(s), minimal CURB sets, receive special attention in Proposition 3. Proposition 2 establishes that CURB sets can be characterized in terms of epistemic stability, by removing player types that do not believe with sufficient probability that the others choose best replies. Proposition 3 starts with an algorithm (Prop. 3(a)) to generate epistemically stable sets from any product set of types: epistemic stability requires including all beliefs over the opponents’ best replies, and any beliefs over opponents’ types that has such beliefs over their opponents, and so on. After all these beliefs have been included, the corresponding product set of best responses to it is epistemically stable and indeed the smallest CURB set containing the best responses to the type set one started with. With this algorithm in hand, minimal CURB sets, the prime focus of attention in applications of CURB sets (recall footnote 1), can be characterized by means of a path-independence property: a product set X of pure strategies is a minimal CURB set if and only if it is the outcome of the algorithm, whenever you initiate it with a type profile assigning probability one to something from X being played.

As our notion of epistemic stability implies stability to alternative best replies, it is natural to follow, for instance, Asheim (2006) and Brandenburger, Friedenberg, and Keisler (2008), and model players as having beliefs about the opponents without modeling the players’ actual behavior. Moreover, we consider complete epistemic models. In these respects, our modeling differs from that of Aumann and Brandenburger’s (1995) characterization of Nash equilibrium. In its motivation in terms of epistemic stability of solution concepts and in its use of p -belief, the present approach is related to Tercieux’s (2006) analysis. 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 to this kind of perturbation. Zambrano (2008) studies the stability of non-equilibrium concepts in terms of mutual belief and is hence more closely related to our analysis. In contrast with our approach, however, Zambrano (2008) restricts attention to rationalizability and probability-1

²Recall that the entire pure strategy space of a game is a CURB set.

beliefs. His main result follows from our Proposition 2. Also Hu (2007) restricts attention to rationalizability, but allows for p -beliefs, where $p < 1$. In the games considered in Hu (2007), 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'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 gives the characterizations of variants of CURB sets. Proofs of the propositions are provided in the appendix.

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 $N = \{1, \dots, n\}$ 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 $u_i : S \rightarrow \mathbb{R}$ defined on the set $S := S_1 \times \dots \times S_n$ 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 an arbitrary Polish (separable and completely metrizable) space F , let $\mathcal{M}(F)$ denote the set of Borel probability measures on F , endowed with the topology of weak convergence. For each player $i \in N$, pure strategy $s_i \in S_i$, and probabilistic belief $\sigma_{-i} \in \mathcal{M}(S_{-i})$, 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* $\beta_i : \mathcal{M}(S_{-i}) \rightarrow 2^{S_i}$ as follows: For all $\sigma_{-i} \in \mathcal{M}(S_{-i})$,

$$\beta_i(\sigma_{-i}) := \{s_i \in S_i \mid u_i(s_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i}) \forall s'_i \in S_i\}.$$

Let $\mathcal{S} := \{X \in 2^S \mid \emptyset \neq X = X_1 \times \dots \times X_n\}$ denote the collection of non-empty Cartesian products of subsets of the players' strategy sets. 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 $\beta(X) := \beta_1(X_{-1}) \times \dots \times \beta_n(X_{-n})$. Each constituent set $\beta_i(X_{-i}) \subseteq S_i$ in this

Cartesian product 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 (1991), 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 $X' \subset X$ and $\beta(X') \subseteq X'$.

Basu and Weibull (1991) call a FURB set a 'tight' CURB set. The reversed inclusion, $X \subseteq \beta(X)$, is sometimes referred to as the 'best response property' (Pearce, 1984, p. 1033). It is shown in Basu and Weibull (1991, Prop. 1 and 2) that a MINCURB set exists, that all MINCURB sets are FURB, and that the product set of rationalizable strategies is the game's largest FURB set. While Basu and Weibull (1991) 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, $\sigma_{-i} \in \mathcal{M}(S_{-i})$.³ Thus, in games with more than two players, the present definition of CURB is somewhat more demanding than that in Basu and Weibull (1991), in the 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.⁴

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 Polish space of types. The *state space* is defined by $\Omega := S \times T$, where $T := T_1 \times \dots \times T_n$. 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 Borel probability measure $\mu_i(t_i) \in \mathcal{M}(\Omega_{-i})$. For each player i , we thus have the

³Our results carry over — with minor modifications in the proofs — to the case of independent strategies.

⁴ 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 (1984), which, in turn, is closely related to Lemma 3.2.1 in van Damme (1983).

player's pure-strategy set S_i , type space T_i and a mapping $\mu_i : T_i \rightarrow \mathcal{M}(\Omega_{-i})$ that to each of i 's types t_i assigns a probabilistic belief, $\mu_i(t_i)$, over the others' strategy choices and types. The structure $(S_1, \dots, S_n, T_1, \dots, T_n, \mu_1, \dots, \mu_n)$ is called an S -based (interactive) probability structure. Assume 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 property is called *complete*.
- μ_i is continuous.
- T_i is compact.

An adaptation of the proof of Brandenburger, Friedenberg, and Keisler (2008, Proposition 7.2) establishes the existence of such a complete probability structure.⁵

In the setting to be developed here, we consider players who choose best replies to their beliefs — but need not believe that all other players do so, only that this is sufficiently likely.

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, $\mathbf{s}_i : \Omega \rightarrow S_i$ 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, $\mathbf{t}_i : \Omega \rightarrow T_i$ 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 (2007).⁶ 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 $B_i^p(\emptyset) = \emptyset$, $B_i^p(\Omega) = \Omega$, $B_i^p(E') \subseteq B_i^p(E'')$ 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

⁵The exact result we use is Proposition 6.1 in an earlier working paper version (Brandenburger, Friedenberg, and Keisler, 2004). Existence can also be established by constructing a universal state space (cf. Mertens and Zamir, 1985; Brandenburger and Dekel, 1993).

⁶See also Monderer and Samet (1989).

($\neg B_i^p(E) \subseteq B_i^p(\neg B_i^p(E))$) 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 $E', E'' \subseteq \Omega$.

Define i 's choice correspondence $C_i : T_i \rightarrow 2^{S_i}$ as follows: For each of i 's types $t_i \in T_i$,

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

consists of i 's best replies when player i is of type t_i . Let \mathcal{T} denote the collection of non-empty Cartesian products of subsets of the players' type spaces:

$$\mathcal{T} := \{Y \in 2^T \mid \emptyset \neq Y = Y_1 \times \cdots \times Y_n\}.$$

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 stability

We define a product set $X \in \mathcal{S}$ of strategies to be *epistemically stable* if there exists a $\bar{p} < 1$ such that, for all probabilities $p \in [\bar{p}, 1]$, there is a set of type profiles $Y \in \mathcal{T}$ such that

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

and

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

Condition (2) states that the strategies in X are precisely those that rational players whose types are in Y may use. 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 .

Proposition 1(a) establishes that epistemically stable sets are necessarily CURB sets. Proposition 1(b) establishes that any CURB set X contains an epistemically stable subset, and also characterizes the largest such subset.

Denote, for each $i \in N$ and $X_i \subseteq S_i$ the pre-image (upper inverse) of X_i under player i 's best response correspondence by

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

For a given subset X_i of i 's pure strategies, $\beta_i^{-1}(X_i)$ consists of beliefs over others' strategy profiles having the property that all best replies to these beliefs are contained in X_i .

Proposition 1 *Let $X \in \mathcal{S}$.*

- (a) *If X is epistemically stable, 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 stable. Furthermore, it is the largest epistemically stable subset of X .*

Claim (a) implies that every epistemically stable set contains at least one strategically stable set, both as defined in Kohlberg and Mertens (1986) and as defined in Mertens (1989), see Ritzberger and Weibull (1995) and Demichelis and Ritzberger (2003), respectively.⁷ Concerning claim (b), we note that $\times_{i \in N} \beta_i(\beta_i^{-1}(S_i))$ equals the set of profiles of 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 w.r.t. set inclusion, it follows from Proposition 1(b) that any epistemically stable set involves only strategies surviving one round of strict elimination.

Our proof shows that Proposition 1 can be slightly strengthened. For (a), one only needs the stability conditions with $p = 1$; as long as there is a $Y \in \mathcal{T}$ such that $C(Y) = X$ and (3) holds, X is CURB.⁸ Moreover, although epistemic stability allows that $Y \in \mathcal{T}$ depends on p , the proof of (b) defines Y independently of p .

Also FURB sets can be characterized in terms of epistemic stability:

⁷In fact, these inclusions hold under the slightly weaker definition of CURB sets in Basu and Weibull (1991), in which a player's belief about other players is restricted to be a product measure over the others' pure-strategy sets.

⁸In 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 response set in the sense of Tercieux (2006).

Proposition 2 *$X \in \mathcal{S}$ is a FURB set if and only if X is epistemically stable and (3) holds with equality.*

As a corollary, Proposition 2 characterizes the set of rationalizable strategy profiles (Bernheim, 1984; Pearce, 1984), since this is the game's largest FURB set (Basu and Weibull, 1991), without involving any explicit assumption of common belief of rationality; only mutual p -belief of rationality and type sets are assumed. Proposition 2 generalizes the main result of Zambrano (2008) to p -belief for p sufficiently close to 1. Proposition 2 also applies to MINCURB sets, as these sets are FURB.

By Proposition 1, the smallest epistemically stable sets are exactly the game's MINCURB sets. As much of the literature on CURB sets (recall footnote 1) focuses on minimal ones, we now turn to an epistemic characterization of MINCURB sets. The characterization has two parts. The first part starts from an arbitrary product set Y of types and generates an epistemically stable set by including all beliefs over the opponents' best replies, and any beliefs over opponents' types that has such beliefs over their opponents, and so on. The so obtained product set of best replies is epistemically stable and is the smallest CURB set containing $C(Y)$. The second part characterizes MINCURB sets in terms of a path independence condition: a product set of pure strategies X is a MINCURB set if and only if it is the output of the algorithm in the first part, *whenever* the algorithm starts from a singleton set consisting of a profile of types that assign probability one to strategies in X .

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 $i \in N$,

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

Define the correspondence $E : T_i \rightarrow 2^{S_i}$, for any $Y \in \mathcal{T}$, by

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

Note that for each set $X \in \mathcal{S}$ in any finite game, there exists a unique smallest CURB set $X' \in \mathcal{S}$ with $X \subseteq X'$ (that is, X' is a subset of all other CURB sets X'' , if any, with $X \subseteq X''$).⁹

⁹To see that this holds for all finite games, note that the collection of CURB sets containing a given set $X \in \mathcal{S}$ is non-empty and finite, and that the intersection of two CURB sets containing X is again a CURB set containing X .

Proposition 3 (a) Let $Y \in \mathcal{T}$. Then $X = E(Y)$ is the smallest CURB set satisfying $C(Y) \subseteq X$. Furthermore, $E(Y)$ is epistemically stable.

(b) $X \in \mathcal{S}$ is a MINCURB set if and only if for each $t \in T$ with

$$\forall i \in N : \quad \text{marg}_{S_{-i}} \mu_i(t_i)(X_{-i}) = 1, \quad (5)$$

it holds that $E(\{t\}) = X$.

Remark 1 If the set $C(Y)$ in claim (a) includes strategies that are not rationalizable, then $E(Y)$ will not be furb. Therefore, the epistemic stability of $E(Y)$ does not follow from Proposition 2: its stability is established by invoking Proposition 1(b).

In order to illustrate Proposition 3, 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 3(a) shows that $C(\{t_1, t_2\})$ is not epistemically stable, 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 stability, 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 stable, when viewed as a singleton set.

Appendix

Proof of Proposition 1. 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

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

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

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

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 stable. That it is the largest epistemically stable 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 w.r.t. 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 $\underline{p} = \max\{\underline{p}_1, \dots, \underline{p}_n\}$.

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

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

using $\beta(X') \subseteq X'$. ■

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

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

Let $\beta^p(X) := \beta_1^p(X_{-1}) \times \cdots \times \beta_n^p(X_{-n})$. Following Tercieux (2006), a set $X \in \mathcal{S}$ is a p -best response 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 (3) holds for each $i \in N$, then X is a p -best response 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 $\text{marg}_{S_{-i}} \mu_i(\mathbf{t}_i(\omega)) = \sigma_{-i}$. So, 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 response set. ■

Proof of Proposition 2. (If) 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 $\sigma_{-i} \in \mathcal{M}(X_{-i})$. So

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

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

(Only if) 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 $\underline{p} = \max\{\underline{p}_1, \dots, \underline{p}_n\}$.

For each $p \in [\underline{p}, 1]$, construct the sequence of Cartesian products of type subsets $\langle Y^p(k) \rangle_k$ 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,

$$[Y_i^p(k)] = B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j^p(k-1)]) \right). \quad (6)$$

defines, for each player i , a decreasing chain $\langle Y_i^p(k) \rangle_k$ 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 $\omega \in [Y_i^p(1)] = B_i^p \left(\bigcap_{j \neq i} (R_j \cap [Y_j^p(0)]) \right)$ 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 $k \in \mathbb{N}$ that $C(Y^p(k)) = \beta(X) = X$. Since $\langle Y_i^p(k) \rangle_k$ 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 $(y_k, x^o) \in \text{graph}(C)$. By the Bolzano-Weierstrass Theorem, we can extract a convergent subsequence for which $y_k \rightarrow y^o$, where $y^o \in Y^p$, since Y^p is closed. Moreover, since the correspondence C is closed-valued and u.h.c., with S compact (it is in fact finite), $\text{graph}(C) \subseteq T \times S$ is closed, and thus $(y^o, x^o) \in \text{graph}(C)$, 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 $i \in N$, (3) holds with equation for Y^p . Fix $i \in N$, and let

$$E_k = \bigcap_{j \neq i} (R_j \cap [Y_j^p(k)]) \quad \text{and} \quad E = \bigcap_{j \neq i} (R_j \cap [Y_j^p]).$$

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 (6) 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 $k \in \mathbb{N}$:

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

To show $B_i^p(E) \supseteq [Y_i^p]$, assume that $\omega \in [Y_i^p]$.¹⁰ This implies that $\omega \in [Y_i^p(k)]$ for all k , and, using (6): $\omega \in B_i^p(E_k)$ 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 3. *Part (a).* 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 $\beta(X') \subseteq X' \subset X$. We must show that $X = E(Y)$.

Consider the sequence $\langle Y(k) \rangle_k$ defined by $Y(0) = Y$ and (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 w.r.t. 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

¹⁰We 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 (1989).

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 $X' \subset X$ would hold, then this would contradict that there exists no $X' \in \mathcal{S}$ with $C(Y) \subseteq X'$ such that $\beta(X') \subseteq X' \subset X$. Hence, $X = C(\bigcup_{k \in \mathbb{N}} Y(k)) = E(Y)$.

Write $X = E(Y)$. To establish that X is epistemically stable, by Proposition 1(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

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

since $C_i(\cdot)$ is monotonic w.r.t. set inclusion.

Part (b). (Only if) Let $X \in \mathcal{S}$ be a MINCURB set. Let $t \in T$ satisfy (5). By construction, $C(\{t\}) \subseteq X$. By part (a), $E(\{t\})$ 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.

(If) For each $t \in T$ satisfying (5), $C(\{t\}) \subseteq E(\{t\}) = X$, so X is a CURB set. To show that X is a *minimal* CURB set, suppose — to the contrary — that there is a CURB set $X' \subset X$. Let $t' \in T$ be such that $\text{marg}_{S_{-i}} \mu_i(t'_i)(X'_{-i}) = 1$ for each $i \in N$. By construction, $C(\{t'\}) \subseteq X'$, so X' is a CURB set containing $C(\{t'\})$. By part (a), $E(\{t'\})$ is the smallest CURB set containing $C(\{t'\})$. Moreover, as $X' \subset X$, t' satisfies (5), so $X' \supseteq E(\{t'\}) = X$, contradicting that $X' \subset X$. ■

References

- Asheim, G.B. (2006), The Consistent Preferences Approach to Deductive Reasoning in Games. Springer, Dordrecht.
- Aumann, R.J. and Brandenburger, A. (1995), Epistemic conditions for Nash equilibrium, *Econometrica* **63**, 1161–1180.
- Basu, K. and Weibull, J.W. (1991), Strategy subsets closed under rational behavior, *Economics Letters* **36**, 141–146.

- Bernheim, D. (1984), Rationalizable strategic behavior, *Econometrica* **52**, 1007–1028.
- Blume, A. (1998), Communication, risk, and efficiency in games, *Games and Economic Behavior* **22**, 171–202.
- Brandenburger, A. and Dekel, E. (1993), Hierarchies of beliefs and common knowledge, *Journal of Economic Theory* **59**, 189–198.
- Brandenburger, A., Friedenberg, A., and Keisler, H.J. (2004), Admissibility in games, mimeo, NYU, Washington University and University of Wisconsin-Madison, December 12, 2004.
- Brandenburger, A., Friedenberg, A., and Keisler, H.J. (2008), Admissibility in games, *Econometrica* **76**, 307–352.
- van Damme, E. (1983), Refinements of the Nash Equilibrium Concept. Springer Verlag, Berlin.
- Demichelis, S. and Ritzberger, K. (2003), From Evolutionary to Strategic Stability, *Journal of Economic Theory* **113**, 51–75.
- Fudenberg, D. and Levine, D.K. (1998), The Theory of Learning in Games. MIT Press, Cambridge, MA.
- Galeotti, A., Goyal, S., and Kamphorst, J.J.A. (2006), Network formation with heterogeneous players, *Games and Economic Behavior* **54**, 353–372.
- Hu, T.-W. (2007), On p -rationalizability and approximate common certainty of rationality, *Journal of Economic Theory* **136**, 379–391.
- Hurkens, S. (1995), Learning by forgetful players, *Games and Economic Behavior* **11**, 304–329.
- Hurkens, S. (1996), Multi-sided pre-play communication by burning money, *Journal of Economic Theory* **69**, 186–197.
- Kalai, E. and Samet, D. (1984), Persistent equilibria in strategic games, *International Journal of Game Theory* **14**, 41–50.
- Kohlberg, E. and Mertens J.-F. (1986), On the strategic stability of equilibria, *Econometrica* **54**, 1003–1037.

- Mertens, J.-F. (1989), Stable equilibria - a reformulation, part I: Definition and basic properties, *Mathematics of Operations Research* **14**, 575-625.
- Mertens, J.-F. and Zamir, S. (1985), Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory* **14**, 1-29.
- Monderer, D. and Samet, D. (1989), Approximating common knowledge with common beliefs, *Games and Economic Behavior* **1**, 170-190.
- Nash, J.F. (1950), Non-cooperative games, Ph.D. thesis, Princeton University.
- Pearce, D.G. (1984), Rationalizable strategic behavior and the problem of perfection, *Econometrica* **52**, 1029-1050.
- Ritzberger, K. and Weibull, J. (1995), Evolutionary selection in normal-form games, *Econometrica* **63**, 1371-1399.
- Sanchirico, C. (1996), A probabilistic model of learning in games, *Econometrica* **64**, 1375-1394.
- Tercieux, O. (2006), p -Best response set, *Journal of Economic Theory* **131**, 45-70.
- Young, P.H. (1998), Individual Strategy and Social Structure. Princeton University Press, Princeton, NJ.
- Zambrano, E. (2008), Epistemic conditions for rationalizability. *Games and Economic Behavior* **63**, 395-405.