

An aerial photograph of the University of Exeter campus, showing green lawns, trees, and buildings, partially framed by a blue curved graphic element.

**Economics Department
Discussion Papers Series**

ISSN 1473 – 3307

**Refined best-response correspondence
and dynamics**

Dieter Balkenborg, Josef Hofbauer and Christoph Kuzmics

Paper number 08/06

Refined best-response correspondence and dynamics

Dieter Balkenborg*, Josef Hofbauer† and Christoph Kuzmics‡

September 5, 2008

Abstract

We characterize the smallest faces of the polyhedron of strategy profiles that could possibly be made asymptotically stable under some reasonable deterministic dynamics. These faces are Kalai and Samet's (1984) persistent retracts and are spanned by Basu and Weibull's (1991) CURB sets based on a natural (and, in a well-defined sense, minimal) refinement of the best-reply correspondence. We show that such a correspondence satisfying basic properties such as existence, upper hemi-continuity, and convex-valuedness exists and is unique in most games. We introduce a notion of rationalizability based on this correspondence and its relation to other such concepts. We study its fixed-points and their relations to equilibrium refinements. We find, for instance, that a fixed point of the refined best reply correspondence in the agent normal form of any extensive form game constitutes a perfect Bayesian equilibrium, which is weak perfect Bayesian in every subgame. Finally, we study the index of its fixed point components.

Keywords: Evolutionary game theory, best response dynamics, CURB sets, persistent retracts, asymptotic stability, Nash equilibrium refinements, learning

JEL codes: C62, C72, C73

*Department of Economics, School of Business and Economics, University of Exeter, D.G.Balkenborg@exeter.ac.uk

†Department of Mathematics, University of Vienna, Josef.Hofbauer@univie.ac.at

‡Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, c-kuzmics@kellogg.northwestern.edu.

We would like to thank Eddie Dekel, Carlos Alos-Ferrer, Amanda Friedenberg, Drew Fudenberg, Klaus Ritzberger, Karl Schlag, Mark Voorneveld, and Jörgen Weibull for helpful comments and suggestions.

1 Introduction

Experiments, in which subjects play relatively simple finite normal form or extensive form games, often focus on testing what one might call economic theory. Economic theory in such cases can be said to be the combination of game theory and, importantly, the assumed (monotonic) one-to-one link between utility and material pay-offs. Often the reaction to finding violations of this economic theory is the introduction of preferences which not only depend on a player's material payoff but also perhaps other players' material payoffs or even utilities. Often maintaining that players (who are assumed, in contrast to us researchers, to know their co-players' preferences) would play some highly sophisticated notion of equilibrium (a prediction of game theory), such as subgame perfection or sequential equilibrium for extensive form games or an undominated equilibrium in normal form games.

Few experimental papers endeavor to test the predictions of game theory on their own. However, what are the predictions of game theory really? One prediction is that play will be in Nash equilibrium, but sometimes we even refine that to Selten (1965)'s subgame perfect equilibrium or even Kreps and Wilson (1982)'s sequential equilibrium in extensive form games and undominated or Selten (1975)'s trembling-hand perfect equilibrium in both normal and extensive form games. But does game theory really predict even Nash equilibrium behavior? Justifying Nash equilibrium behavior or any of its refinement is very hard. In a truly one-shot game even if we assume that players are rational and have common knowledge of rationality we can only really expect players to play some strategy within the set of rationalizable strategies, see Bernheim (1984) and Pearce (1984). To then justify equilibrium behavior we would have to argue that players' beliefs are somehow aligned. In a truly one-shot game, however, there is no reason to believe that this would be the case. That is why, for instance, the coordination game is such an interesting game, precisely because we often see that players are not able to coordinate on a Nash equilibrium if the game is only played once.

We only have hope of further pinning down what players might be doing in a game, beyond that they might play a rationalizable strategy, if the game is played repeatedly by various people, so that learning (or evolution) can take place. Models of learning were developed virtually at the same time as Nash proposed his solution concept. Even Nash had an evolutionary interpretation of his solution concept in mind (for the appropriate quote from Nash's PhD thesis see page 1 in Ritzberger and Weibull (1995)). These models, however, principally failed to provide a justification for equilibrium behavior.

There are then two main avenues of research. One is to find processes which do lead to equilibrium behavior in some sense, and then question the reasonability or plausibility of these processes after the fact. Alternatively one could just accept that equilibrium behavior can not be so easily justified and ask the question what can be justified instead. The two most striking results in the second type of literature are due to Hurkens (1995) and Ritzberger and Weibull (1995). In a stochastic best-reply model a la Young (1993) Hurkens (1995) shows that the only candidates for stochastically stable states are those within Basu and Weibull's (1991) CURB sets. Using results of Balkenborg (1992), Hurkens (1995) furthermore shows that a specially refined stochastic best-reply dynamics leads to play eventually being within Kalai and Samet's (1984) persistent retracts. Ritzberger and Weibull (1995) show that under any general payoff-positive dynamics minimal asymptotically stable faces are spanned by pure strategy sets which are closed under weakly better replies. Now these sets can be very large.

In this paper we are after the following. What is the smallest possible set of states one could

still call asymptotically stable under some plausible dynamic? In order to find an answer to this question we restrict attention to best-reply dynamics as opposed to better-reply dynamics. This is, of course, a strong assumption about the rationality of individuals. We then go further in asking whether we could reasonably restrict players to play only a subset of best-replies.

We, in fact, study refinements of the best-reply correspondence which satisfy 5 conditions. A refinement must be a subset of the best-reply correspondence, be never empty valued, be convex-valued, have a product structure, and be upper hemi-continuous. Under certain mild conditions on the normal form game at hand there is a unique minimal such refined correspondence, which we characterize. This is in some sense the opposite exercise undertaken by Ritzberger and Weibull (1995). They find sets which are asymptotically stable under a wide variety of deterministic dynamics, while we here investigate sets which are asymptotically stable under only the, in a well-specified sense, most selective of deterministic dynamics. In this sense, we characterize the smallest faces which one could justifiably call evolutionarily stable. The main result of this paper is that these smallest evolutionarily stable faces coincide with Kalai and Samet's (1984) persistent retracts, which again coincide with Basu and Weibull's (1991) CURB sets adapted for the minimal refined best-reply correspondence (Theorems 2, 5, and 6). This result is analogous to Hurkens's (1995) result that persistent retracts are the only candidates for stochastically evolutionarily stable states in a particular stochastic model of best-reply learning à la Young (1993). On the "way" to this result, in addition, we find a series of interesting results about the underlying minimal refined best-reply correspondence, its fixed points, and notions of rationalizability based on it. One result, for instance, is that a pure fixed point of this minimal refined best-reply correspondence induces a perfect Bayesian equilibrium (with weak perfect Bayesian equilibria in every subgame) in the extensive form with the given normal form as its agent normal form (Proposition 17).

Methodologically there is some overlap with Balkenborg (1992) who, in order to analyze the properties of persistent retracts, studies the "semi-robust best reply correspondence", which differs from the refined best reply correspondences considered here by not being convex valued. Balkenborg, Jansen, and Vermeulen (2001) analyze the invariance of persistent retracts and equilibria using "sparse strategy selections". These are particularly useful when no unique minimal refined best reply correspondence exists.

The paper proceeds as follows. We first define the class of games we study in section 2. We then define what we call a refinement of the best-reply correspondence in section 3, where we also characterize the, in the given class of games, unique minimal such refinement. In section 4 we, to some extent, characterize fixed-points of the minimal refined best-reply correspondence. In section 5 we discuss the concept of rationalizability based on the refined best-reply correspondence and its relationship to other notions of rationalizability and Dekel and Fudenberg's (1990) $S^\infty W^1$ elimination procedure. In section 6 we study the notion of a CURB set (Basu and Weibull (1991)) for the refined best-reply correspondence and prove that it coincides with Kalai and Samet's (1984) notion of an absorbing retract. In section 7 we discuss extensive form games. Section 8 provides a micro-story, similar in spirit to Björnerstedt and Weibull (1996), leading to a differential inclusion based on the refined best-reply correspondence, before we prove the main result in section 9. Section 10 discusses index-theory based on the minimal refined best-reply correspondence. Section 11 concludes.

The paper has three appendices. Appendix A shows in which sense our restriction to games with generically unique best replies is not essential. The result in Appendix B implies that many concepts we consider here coincide in generic normal form games. Appendix C studies

the special structure of the minimal refined best reply correspondence in two-player games.

2 Preliminaries

Let $\Gamma = (I, S, u)$ be a finite n -player normal form game, where $I = \{1, \dots, n\}$ is the set of players, $S = \times_{i \in I} S_i$ is the set of pure strategy profiles, and $u : S \rightarrow \mathbb{R}^n$ the payoff function¹. Let $\Theta_i = \Delta(S_i)$ denote the set of player i 's mixed strategies, and let $\Theta = \times_{i \in I} \Theta_i$ denote the set of all mixed strategy profiles. Let $\text{int}(\Theta) = \{x \in \Theta : x_{is} > 0 \ \forall s \in S_i \ \forall i \in I\}$ denote the set of all completely mixed strategy profiles.

A strategy profile $x \in \Theta$ may also represent a population state in an evolutionary interpretation of the game in the following sense. Each player $i \in I$ is replaced by a population of agents playing in player position i and x_{is_i} denotes the proportion of players in population i who play pure strategy $s_i \in S_i$.

For $x \in \Theta$ let $\mathcal{B}_i(x) \subset S_i$ denote the set of pure-strategy best-replies to x for player i . Let $\mathcal{B}(x) = \times_{i \in I} \mathcal{B}_i(x)$. Let $\beta_i(x) = \Delta(\mathcal{B}_i(x)) \subset \Theta_i$ denote the set of mixed-strategy best-replies to x for player i . Let $\beta(x) = \times_{i \in I} \beta_i(x)$.

Two strategies $x_i, y_i \in \Theta_i$ are **own-payoff equivalent** (for player i) if $u_i(x_i, x_{-i}) = u_i(y_i, x_{-i})$ for all $x_{-i} \in \Theta_{-i} = \times_{j \neq i} \Theta_j$ (see Kalai and Samet (1984)). In contrast, Kohlberg and Mertens (1986) call two strategies $x_i, y_i \in \Theta_i$ **payoff equivalent** if $u_j(x_i, x_{-i}) = u_j(y_i, x_{-i})$ for all $x_{-i} \in \Theta_{-i}$ and for all players $j \in I$. We will use these concepts primarily for pure strategies.

Let $\Psi = \{x \in \Theta \mid \mathcal{B}(x) \text{ is a singleton}\}$. Notice that the unique best reply against a strategy combination in Ψ is necessarily a *pure* strategy. Throughout this paper we will restrict attention to games Γ for which this set Ψ is dense in Θ . Let this set of games be denoted by $\mathcal{G}^*$. A game $\Gamma \notin \mathcal{G}^*$ is given by Game 1. Player 1's best reply set is $\{A, B\}$ for any (mixed) strategy of player 2. Hence, $\beta(x)$ is never a singleton and $\Psi = \emptyset$ is not dense in Θ . This has to do with the fact that player 1 has two own-payoff equivalent pure strategies.

	C	D
A	1,1	1,0
B	1,0	1,1

Game 1: A Game in which Ψ is not dense in Θ .

Proposition 1 demonstrates that without equivalent strategies Ψ is dense in Θ . The following lemma, due to Kalai and Samet (1984), is used in the proof of Proposition 1.

Lemma 1 *Let U be a non-empty open subset of Θ . Then two strategies $x_i, y_i \in \Theta_i$ are own-payoff equivalent (for player i) if and only if $u_i(x_i, z_{-i}) = u_i(y_i, z_{-i})$ for all $z \in U$.*

Proposition 1 *Let Γ be without own-payoff equivalent pure strategies. Then Ψ is dense in Θ ; i.e., $\Gamma \in \mathcal{G}^*$.*

Proof: Suppose Ψ is not dense in Θ . Then there is an open set U in Θ such that for all $y \in U$ the pure best-response set $\mathcal{B}(y)$ is not a singleton, i.e., has at least two elements. Without loss of generality, due to the finiteness of S , we can assume that there are two pure strategy-profiles

¹The function u will also denote the expected utility function in the mixed extension of the game Γ .

$s_i, t_i \in S_i$ such that $s_i, t_i \in \mathcal{B}_i(y)$ for all $y \in U$ and some player $i \in I$. But then by Lemma 1, s_i and t_i are own-payoff equivalent for player i . QED

Note that the converse of Proposition 1 is not true. Consider two own-payoff equivalent strategies which are strictly dominated by another strategy. If these are the only equivalent strategies in Γ then Ψ is still dense in Θ . However, the following proposition is immediate. Call $x_i \in \Theta_i$ a **robust best reply** against $x \in \Theta$ if x_i is a best reply against all strategy combinations in a neighborhood of x . Call $x_i \in \Theta_i$ a **robust strategy** if x_i is a robust best reply against *some* strategy combination $x \in \Theta$. This terminology is inspired by Okada (1983).

Proposition 2 *Let Γ be such that Ψ is dense in Θ . Let $s_i \in S_i$ be a robust strategy. Then player i has no distinct own-payoff equivalent strategy to s_i in S_i .*

Adapting a notion of Brandenburger and Friedenberg (2007) for perfect information games, let a normal form game satisfy the Single Payoff Condition (SPC) if all own-payoff equivalent pure strategies are also payoff equivalent. Not every game satisfies the SPC: a player other than i might not be indifferent between player i 's own-payoff equivalent strategies (as is the case in Game 1). Thus, our restriction that the game should have no own-payoff equivalent strategies for any player i is a stronger requirement than saying the game has to be in semi-reduced normal form (see e.g., page 147 in Ritzberger (2002)). However, games not satisfying the SPC are exceptional. Trivially, for generic normal form games there are neither own-payoff nor payoff equivalent strategies and hence the SPC holds.² This is of little interest because most important classes of normal form games such as normal forms of extensive form games or of finitely repeated games are non-generic. Requiring genericity conflicts with imposing any additional structure on the class of games considered.³ In Appendix A we identify a condition on a class of normal form games which implies that the SPC holds generically *within this class*. This condition is shown to be satisfied by the classes of normal forms of extensive form games, of finitely repeated games and of cheap talk games.

Hence the restriction to games in the class $\mathcal{G}^*$ made throughout the paper is essentially the restriction to the semi-reduced normal form⁴ in the sense of Kohlberg and Mertens (1986). Since we are primarily interested in the best response correspondence this restriction is largely without loss of generality. In fact, every trajectory of the best reply dynamics of the reduced form of a normal form game corresponds in a canonical fashion to a family of trajectories in the original game which projects onto it.

We now turn to the notion of inferior strategies.⁵

Definition 1 *A strategy $s_i \in S_i$ is strictly inferior if for every $x \in \Theta$ there is a $t_i \in S_i$ such that $u_i(s_i, x_{-i}) < u_i(t_i, x_{-i})$. A strategy $w_i \in S_i$ is weakly inferior if for every $x \in \Theta$ there is a $t_i \in S_i$, $t_i \neq w_i$ such that $u_i(w_i, x_{-i}) \leq u_i(t_i, x_{-i})$ and $u_i(w_i, y_{-i}) < u_i(t_i, y_{-i})$ for some $y \in \Theta$.*

²Throughout this paper we say that a property holds generically in some open subset of a Euclidean space if it fails only on a lower dimensional semi-algebraic or (in Appendix A) semi-analytic set. Such lower-dimensional sets always have measure zero, are nowhere dense and hence meagre (first category).

³For an illuminating discussion on this point see Mertens (2004).

⁴In particular, we are, for instance, not ruling out games with weakly dominated strategies.

⁵Our notions of strict and weak inferiority are motivated by, but not identical to, the notion of inferior choices in Harsanyi and Selten (1988).

For games in $\mathcal{G}^*$ a strictly inferior strategy s_i is such that $s_i \notin \mathcal{B}_i(x)$ for any $x \in \Theta$, while a weakly inferior strategy w_i is such that if $w_i \in \mathcal{B}_i(x)$ then $\mathcal{B}_i(x)$ is not a singleton. Note that every game in $\mathcal{G}^*$ has at least one strategy for each player which is not weakly inferior. If a strategy is strictly (weakly) dominated then it is strictly (weakly) inferior. The converse is not true (see Example 5.7 in Ritzberger (2002) for a strategy which is strictly inferior but not strictly dominated).

Lemma 2 *A strategy is robust if and only if it is not weakly inferior.*

Proof: If $s_i \in S_i$ is not weakly inferior, then there exists an $x_{-i} \in \Theta_{-i}$ such that, for all $t_i \in S_i$, $u_i(s_i, x_{-i}) > u_i(t_i, x_{-i})$ or $u_i(s_i, y_{-i}) \geq u_i(t_i, y_{-i})$ holds for all $y_{-i} \in \Theta_{-i}$. In the latter case t_i is weakly dominated by s_i . By continuity s_i is a best reply in a neighborhood of x_{-i} and hence a robust strategy. Conversely, if $s_i \in S_i$ is robust it is a best reply on a non-empty open set in Θ . Any strategy which is not own-payoff equivalent to s_i can, by Lemma 1, be a best reply jointly with s_i only on a closed, nowhere dense set. There are only finitely many pure strategies to consider and any mixed strategy is a best reply only against strategy profiles against which all pure strategies in its support are also best replies. Thus, there exists a non-empty open set in Θ_i such that a strategy of player i is a best response against a strategy profile in this set if and only if it is own-payoff equivalent to s_i . Hence s_i is not weakly inferior. QED

3 Refined best-reply correspondences

A correspondence $\tau : \Theta \Rightarrow \Theta$ is a *refined best-reply correspondence*⁶ if

1. $\tau(x) = \times_{i \in I} \tau_i(x) \forall x \in \Theta$, where $\tau_i : \Theta \Rightarrow \Theta_i$ for all $i \in I$,
2. $\tau_i(x) \subset \beta_i(x) \forall x \in \Theta, \forall i \in I$,
3. $\tau_i(x) \neq \emptyset \forall x \in \Theta, \forall i \in I$,
4. $\tau(x)$ is convex and closed for all $x \in \Theta$,
5. $\tau(x)$ is upper-hemi continuous at all $x \in \Theta$.

Note first that if we replaced property 5 by requiring τ to have a closed graph, then we could omit the requirement of closedness in property 4. Note, furthermore, that if $\beta(x)$ is a singleton then so must be any $\tau(x)$ (with $\tau(x) = \beta(x)$) by properties 2 and 3. In games in $\mathcal{G}^*$ we thus must have $\tau(x) = \beta(x)$ for all $x \in \Psi$ for any refined best-reply correspondence τ . If the best-response at x , $\beta(x)$ is a singleton, then it must be a pure strategy profile. For $x \notin \Psi$ the set $\tau(x)$ must include all pure strategies which are best responses to some nearby $x' \in \Psi$ by property 5⁷. For such x any $\tau(x)$ must then also include all convex combinations of all pure strategies in $\tau(x)$ by property 4. For games in $\mathcal{G}^*$, therefore, the unique minimal such refined best-reply correspondence, denoted $\sigma : \Theta \Rightarrow \Theta$, can be found in the following way. For $x \in \Theta$ let

$$\mathcal{S}_i(x) = \{s_i \in S_i \mid \exists \{x_t\}_{t=1}^\infty \in \Psi : x_t \rightarrow x \wedge \mathcal{B}_i(x_t) = \{s_i\} \forall t\}.$$

⁶The difference between such refined best replies and the notion of admissible best replies developed in Mertens (2004) is subtle because the latter do not have to satisfy our assumptions 1 and 4.

⁷Strategies that are unique best replies to some x are called *inducible* in von Stengel and Zamir (2004).

For games in $\mathcal{G}^*$ the set $\mathcal{S}_i(x)$ is the set of pure semi-robust best replies defined in Balkenborg (1992).⁸ Let $\mathcal{S}(x) = \times_{i \in I} \mathcal{S}_i(x)$. These observations prove the following theorem.

Theorem 1 *Let $\Gamma \in \mathcal{G}^*$. The unique minimal refined best-reply correspondence is given by σ , defined such that for any $x \in \Theta$, $\sigma(x) = \times_{i \in I} \Delta(\mathcal{S}_i(x))$.*

Note that for games not in $\mathcal{G}^*$ there may well be multiple minimal refined best-reply correspondences. For the most part in this paper we will study games in $\mathcal{G}^*$ only.

The following Proposition says that for generic normal form games the minimal refined best-reply correspondence coincides everywhere with the best-reply correspondence, i.e., $\sigma \equiv \beta$. Remember, however, that many games of interest in $\mathcal{G}^*$ (derived for instance from an extensive form) are not among these generic normal form games.

Proposition 3 *For generic normal form games a strategy is a semi-robust best reply if and only if it is a best reply (i.e., $s_i \in \mathcal{S}_i(x) \Leftrightarrow s_i \in \mathcal{B}_i(x)$). Consequently, the best reply correspondence is minimal (i.e., $x_i \in \sigma_i(x) \Leftrightarrow x_i \in \beta_i(x)$) in these games.*

For the proof see appendix B. As an immediate consequence of Lemma 2 we have:

Lemma 3 *Let $w_i \in S_i$ be weakly inferior for player i . Then $w_i \notin \mathcal{S}_i(x)$ (i.e., w_i is not a semi-robust best reply) for any $x \in \Theta$.*

In two-player games the following reverse statement can be established. Its proof can be found in appendix C.

Proposition 4 *In a two-player game a strategy is a semi-robust best reply, $s_i \in \mathcal{S}_i(x)$, if and only if it is a best reply, $s_i \in \beta_i(x)$, and is not weakly inferior.*

For two-player games, in fact, we can completely characterize weakly inferior strategies. The proof is again in appendix C.

Proposition 5 *In a two-player game a strategy is weakly inferior if and only if it is weakly dominated or equivalent to a proper mixture of strategies which are not equivalent.*

Proposition 4 implies the following Proposition.

Proposition 6 *Let $\Gamma = (I, S, u) \in \mathcal{G}^*$ be a two-player game with minimal refined best-reply correspondence $\sigma(\Gamma)$. Then there is a game $\Gamma' = (I', S', u') \in \mathcal{G}^*$ with $I' = I$, $S' = S$, and a payoff function $u' : S \rightarrow \mathbb{R}^2$ such that its best-reply correspondence $\beta(\Gamma') \equiv \sigma(\Gamma)$.*

Proof: Let Γ' be such that, for every $i \in I$, every $s_{-i} \in S_{-i}$, and every weakly inferior $w_i \in S_i$, $u'_i(w_i, s_{-i}) = u_i(w_i, s_{-i}) - \delta$ for some $\delta > 0$. Then, for this game Γ' no weakly inferior strategy is ever a best reply. Thus, by Proposition 4, $\sigma(\Gamma) \equiv \sigma(\Gamma') \equiv \beta(\Gamma')$. QED


Proposition 6 is useful as it tells us that in two-player games, the structure of fixed points of σ is the same as the structure of Nash equilibria. In particular, it implies that, in two-player games, there are only finitely many components of fixed points of σ . More precisely, applying the results in Jansen, Jurg, and Vermeulen (2002) we have the following.

⁸Balkenborg (1992) defines a best reply against x_{-i} as semi-robust if it is a robust best reply against a sequence of strategy combinations converging to x_i

Corollary 1 Let $\Gamma = (I, S, u) \in \mathcal{G}^*$ be a two-player game with minimal refined best-reply correspondence σ . Then the set of fixed points of σ is the union of finitely many polytopes and hence the union of finitely many connected components.

	C	D
A	0	0
B	0	0
E		

	C	D
A	1	-1
B	-1	0
F		



Game 2 and Figure 1: A game where the refined best reply correspondence is not the best reply correspondence of a modified game. In Figure 1 the regions where strategies E and F of player 3 are best replies in this game are indicated in the square of strategy profiles of players 1 and 2. The probability with which player 1 chooses B is indicated vertically downwards in the graph while the probability of player 2 choosing D is indicated horizontally. In the shaded area between the two branches of the hyperbola E is the best reply for player 3, outside it is F . The lower branch of the hyperbola intersects the square only in the point (B, D) , indicating that F is a best reply against (B, D) , but not semi-robust.

Proposition 6 does not extend to games with three or more players, as the $2 \times 2 \times 2$ game 2 shows. Here and in the following games, player 1 chooses the row, player 2 the column and player 3 the matrix. In this example we specify the payoffs of player 3 only. As indicated in Figure 1, note that against opponent strategy profiles $(1/2A + 1/2B, C)$, $(A, 1/2C + 1/2D)$, and $(2/3A + 1/3B, 2/3C + 1/3D)$ (among others) both E and F are semi-robust best replies. However, against (A, C) F is the only best reply and against (B, D) E is the only semi-robust best reply. Nearby the latter strategy profile there is no open set in the square of the opponents' mixed strategy profiles where F is a best response.

Now assume there exists another game with the same strategies for which the best response mapping for player 3 is identical to the *minimal refined* best response correspondence of the given game. This implies that player 3 must remain indifferent between E and F against the strategy profiles $(1/2A + 1/2B, C)$, $(A, 1/2C + 1/2D)$, and $(2/3A + 1/3B, 2/3C + 1/3D)$. Moreover, F must be a best response against (A, C) , but not against (B, D) . This implies

$$\begin{aligned}
& \frac{1}{2} (u_3(A, C, E) - u_3(A, C, F)) - \frac{1}{2} (u_3(B, C, E) - u_3(B, C, F)) = 0 \\
& \frac{1}{2} (u_3(A, C, E) - u_3(A, C, F)) - \frac{1}{2} (u_3(A, D, E) - u_3(A, D, F)) = 0 \\
& \frac{4}{9} (u_3(A, C, E) - u_3(A, C, F)) - \frac{2}{9} (u_3(B, C, E) - u_3(B, C, F)) \\
& - \frac{2}{9} (u_3(A, D, E) - u_3(A, D, F)) + \frac{1}{9} (u_3(B, D, E) - u_3(B, D, F)) = 0
\end{aligned}$$

We conclude that $u_3(B, D, E) - u_3(B, D, F) = 0$, and, thus F is a best response against (B, D) , a contradiction.

4 Nash equilibrium versus best-reply refinements

This section provides a few results relating fixed points of the (minimal) refined best-reply correspondence to (refinements of) Nash equilibria.

Proposition 7 *Let Γ be a finite two-player game in $\mathcal{G}^*$. Let $x \in \Theta$ be a fixed point of the refined best-reply correspondence σ . Then $x_{iw_i} = 0$ for every weakly inferior $w_i \in S_i$.*

Proof: Immediate from Lemma 3: Let $x \in \sigma(x)$. By Lemma 3 $w_i \notin S_i(x)$ for any weakly inferior $w_i \in S_i$. But then no $y \in \Theta$ with $y_{iw_i} > 0$ can be in $\sigma(x)$. QED

Selten (1975) introduced the concept of a (trembling-hand) perfect (Nash) equilibrium. A useful characterization of a perfect equilibrium in normal form games is given in the following lemma, which is also due to Selten (1975) (see also Proposition 6.1 in Ritzberger (2002) for a textbook treatment).

Lemma 4 *A (possibly mixed) strategy profile $x \in \Theta$ is a (trembling-hand) perfect (Nash) equilibrium if there is a sequence $\{x_t\}_{t=1}^\infty$ of completely mixed strategy profiles (i.e., each $x_t \in \text{int}(\Theta)$) such that x_t converges to x and $x \in \beta(x_t)$ for all t .*

Not every fixed point of σ is necessarily a trembling-hand perfect equilibrium, even in 2-player games. To see this consider Game 3, taken from Hendon, Jacobson, and Sloth (1996). For this game σ and β are identical. The mixed strategy profile $x^* = ((0, 1/2, 1/2); (1/2, 0, 1/2))$ is a Nash equilibrium, hence a fixed point of β , hence of σ , which, as Hendon, Jacobson, and Sloth (1996) point out is not perfect. Indeed, while the two pure strategies in the support of x_2^* , i.e., strategies D and F are not weakly dominated, the mixture x_2^* is weakly dominated by the pure strategy E . As weakly dominated strategies in a 2-player game cannot be perfect (see e.g., Theorem 3.2.2 in van Damme (1991)), x^* is not perfect.

	D	E	F
A	0,0	0,1	0,0
B	2,0	2,1	0,2
C	0,2	0,1	2,0

Game 3: A game in which a fixed point of σ is not perfect.

	D	E	F
A	2,2	1,2	1,2
B	2,1	2,2	0,0
C	2,1	0,0	2,2

Game 4: A game in which a perfect equilibrium (and, in fact, KM-stable equilibrium) is not a fixed point of σ .

Proposition 8 *Let Γ be a 2-player game in $\mathcal{G}^*$. Then every pure fixed-point, $s \in S$, of the refined best-reply correspondence, σ , is a perfect equilibrium.*

Proof: Every pure fixed point of σ is undominated by Proposition 7. In two-player games every undominated Nash equilibrium is (trembling-hand) perfect. QED

The converse of Proposition 8 is not true. Consider Game 4. In this game strategy A (and similarly D) is equivalent to the mixture of pure strategies B and C (E and F respectively). However, A is a best-reply only on a thin set of mixed-strategy profiles. In fact, A is best against any $x \in \Theta$ in which $x_{2E} = x_{2F}$, the set of which is a thin set. By Proposition 2 this game is in

$\mathcal{G}^*$. In this game (A, D) constitutes a perfect equilibrium. In fact every mixed strategy profile $((\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}); (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}))$ is a strictly perfect equilibrium, and hence, constitutes a singleton KM-stable set. But none of these equilibria, except the one with $\alpha = 0$, are fixed points of σ , due to the fact that A (and D) is only best on a thin set (it is a weakly inferior strategy).

Proposition 8 cannot be generalized to general n -player games. To see this consider the following immediate characterization of fixed points of σ . For $x_i \in \Theta_i$ let $C(x_i) = \{s_i \in S_i | x_{is_i} > 0\}$ denote the carrier (or support) of x_i .

Lemma 5 *Strategy profile $x \in \Theta$ satisfies $x \in \sigma(x)$ if and only if for all $i \in I$ and for all $s_i \in C(x_i)$ there is an open set $U^{s_i} \subset \Theta$, with x in the closure of U^{s_i} , such that $s_i \in \mathcal{B}_i(y)$ for all $y \in U^{s_i}$.*

Suppose $x \in \sigma(x)$. Consider player i . Then for all $s_i \in C(x_i)$ let U^{s_i} denote this open set in which s_i is best. Now if $\bigcap_{i \in I} \bigcap_{s_i \in C(x_i)} U^{s_i} \neq \emptyset$, then x is also trembling-hand perfect. However, this is not necessarily the case. We already saw this for the 2-player Game 3. In the fixed point of σ , $x^* = ((0, 1/2, 1/2); (1/2, 0, 1/2))$, player 2 uses his pure strategies D and F only. D is best in the open set $U^D = \{x \in \Theta | x_{1C} > \frac{1}{2}\}$, while F is best in the open set $U^F = \{x \in \Theta | x_{1B} > \frac{1}{2}\}$. There are no bigger open set with the same property. Yet the two sets have an empty intersection. Hence, x^* is not perfect. The extensive form game in Figure 5 demonstrates that for games with more than 2 players this phenomenon may even occur for pure fixed points of σ .

In section 6 we prove that CURB sets (Basu and Weibull (1991)) based on σ give rise to absorbing retracts (Kalai and Samet (1984)) and minimal such sets give rise to persistent retracts. One might think that fixed points of σ will have some relation to persistent equilibria (Nash equilibria in a persistent retract, Kalai and Samet (1984)). This is not true, though. Note first that the mixed equilibrium in the coordination game is not persistent and is a fixed point of σ . Consider Game 5 taken from Kalai and Samet (1984). The equilibrium (B, D, E) is perfect and proper but not persistent as Kalai and Samet (1984) point out. It is also a fixed point of σ . To see this, note that E is weakly dominant for player 3 and that B and D are best (for players 1 and 2, respectively) against all nearby strategy profiles in which player 2 chooses strategy C with smaller probability than player 3 chooses F and player 1 chooses A with smaller probability than player 3 chooses F .

	C	D		C	D
A	1,1,1	0,0,0	A	0,0,0	0,0,0
B	0,0,0	0,0,0	B	0,0,0	1,1,0
E			F		

Game 5: A game in which a pure fixed point of σ is not persistent.

Game 6, taken from Kalai and Samet (1984), demonstrates that there are persistent equilibria which are not fixed points of σ . The strategy profile (A, C, E) is persistent (see Kalai and Samet (1984)) but is not a fixed point of σ . To see this note that player 1's strategy A is never best for nearby strategy profiles. The one pure strategy combination (of players 2 and 3) against which A is better than B is (D, F) which for nearby (to (A, C, E)) strategy profiles will always have lower probability than the outcomes (C, F) and (D, E) against which B is better than A .

	C	D		C	D
A	0,0,0	0,0,1	A	0,1,0	1,0,0
B	0,1,0	1,0,1	B	1,0,1	0,1,0
E			F		

Game 6: A game in which a pure persistent equilibrium is not a fixed point of σ .

5 σ -Rationalizability

A set $R \subset S$ is a **strategy selection** if $R = \times_{i \in I} R_i$ and $R_i \subset S_i$, $R_i \neq \emptyset$ for all i . For a strategy selection R let $\Theta(R) = \times_{i \in I} \Delta(R_i)$ denote set of independent strategy mixtures of the pure strategies in R . A set $\Psi \subset \Theta$ is a **face** if there is a strategy selection R such that $\Psi = \Theta(R)$. Note that $\Theta = \Theta(S)$. Note also that $\beta(x) = \Theta(\mathcal{B}(x))$ and $\sigma(x) = \Theta(\mathcal{S}(x))$.

For $A \subset \Theta$ let $\mathcal{B}_i(A) = \{s_i \in S_i \mid s_i \in \mathcal{B}_i(x) \text{ for some } x \in A\}$ denote the set of all pure best-replies for player i to all strategy profiles in set A . Let $\beta_i(A) = \Delta(\mathcal{B}_i(A))$ denote the convex hull of this set $\mathcal{B}_i(A)$. Let $\beta(A) = \times_{i \in I} \beta_i(A)$. For $k \geq 2$ let $\beta^k(A) = \beta(\beta^{k-1}(A))$. For $A = \Theta$, $\beta^k(A)$ is a decreasing sequence, and we denote $\beta^\infty(\Theta) = \bigcap_{k=1}^\infty \beta^k(\Theta)$. A pure strategy profile $s \in S$ is **rationalizable** if it is an element of the strategy selection $R \subset S$ which satisfies $\Theta(R) = \beta^\infty(\Theta)$ (Bernheim (1984) and Pearce (1984); see also Ritzberger (2002), Definition 5.3 for a textbook treatment).

The same can be done with the refined best-reply correspondence σ . For $A \subset \Theta$ let $\mathcal{S}_i(A) = \{s_i \in S_i \mid s_i \in \mathcal{S}_i(x) \text{ for some } x \in A\}$ denote the set of all pure refined best-replies for player i to all strategy profiles in set A . Let $\sigma_i(A) = \Delta(\mathcal{S}_i(A))$. Let $\sigma(A) = \times_{i \in I} \sigma_i(A)$. For $k \geq 2$ let $\sigma^k(A) = \sigma(\sigma^{k-1}(A))$. For $A = \Theta$, $\sigma^k(A)$ is again a decreasing sequence, and we denote $\sigma^\infty(\Theta) = \bigcap_{k=1}^\infty \sigma^k(\Theta)$. A pure strategy profile $s \in S$ is **σ -rationalizable** if it is an element of the strategy selection $R \subset S$ which satisfies $\Theta(R) = \sigma^\infty(\Theta)$.

Using Propositions 3 and 4 we can provide two propositions for generic normal form games and for two-player games, respectively.

Proposition 9 *In generic normal form games a strategy is σ -rationalizable if and only if it is rationalizable.*

Proof: This follows immediately from Proposition 3.

QED

Proposition 10 *In a two-player game the σ -rationalizable strategies are the strategies that survive one round of elimination of weakly inferior pure strategies followed by the iterated elimination of strictly dominated strategies.*

Proof: This follows from Proposition 4 and the fact that in two-player games strictly inferior strategies are the same as strictly dominated strategies.

QED

For general games in $\mathcal{G}^*$, by the fact that $\sigma(x) \subset \beta(x)$ for all $x \in \Theta$, we obviously have that the set of σ -rationalizable strategies is a subset of the set of rationalizable strategies. However, we can say more. Let $\tilde{\Gamma} = (I, S, \tilde{u})$ denote the game derived from Γ by defining $\tilde{u}_i(s_i, s_{-i}) = u_i(s_i, s_{-i}) - \delta$, for a fixed positive δ , if $s_i \in S_i$ is weakly, and not strictly, inferior in Γ and $\tilde{u}_i(s_i, s_{-i}) = u_i(s_i, s_{-i})$ otherwise. Every pure strategy which is weakly inferior in Γ is, therefore, strictly inferior in $\tilde{\Gamma}$. Let $\tilde{\beta}$ denote the best-reply correspondence of $\tilde{\Gamma}$. Then we have the following lemma.

Lemma 6 For $\tilde{\Gamma}$ and $\tilde{\beta}$ defined as above we have $\sigma(x) \subset \tilde{\beta}(x)$ for all $x \in \Theta$.

Proof: Follows immediately from Lemma 3. QED

The refined best-reply set $\sigma(x)$ may, for some games Γ and some $x \in \Theta$, be a proper subset of $\tilde{\beta}(x)$. To see this consider Game 7, taken from van Damme (1991), Figure 2.2.1; see also exercise 6.10 in Ritzberger (2002). In this game strategies D and F are strictly inferior for players 2 and 3, respectively. There are no strategies which are weakly but not strictly inferior. Hence, $\tilde{\beta}(x) = \beta(x)$ for any $x \in \Theta$. Player 1's strategy B is (the unique) best strategy when players 2 and 3 play D and F , respectively. Both A and B are best when players 2 and 3 play C and E , respectively. However, for any (mixed) strategy profile, $y \in \Theta$ in which players 2 and 3 play close to C and E , A is the unique best reply. Hence, $B \notin \mathcal{S}_1(x)$ for any $x \in \Theta$ for which $x_{2C} = 1$ and $x_{3E} = 1$. Therefore, $\mathcal{S}_1(x) = \{A\}$ is indeed a proper subset of $\mathcal{B}_1(x) = \{A, B\}$ for any such x , and, hence, $\sigma(x)$ is a proper subset of $\tilde{\beta}(x)$ for any such x . In fact, this game is usually used to illustrate that in 3-player games an undominated Nash equilibrium, (B, C, E) , need not be perfect, as is indeed the case here.

	C	D		C	D
A	1,1,1	1,0,1	A	1,1,0	0,0,0
B	1,1,1	0,0,1	B	0,1,0	1,0,0
	E			F	

Game 7: A game in which for some $x \in \Theta$, $\sigma(x)$ is a proper subset of $\tilde{\beta}(x)$.

Let $\tilde{\beta}^\infty$ be defined analogously to β^∞ . We call a pure strategy $s \in S$ **Dekel-Fudenberg rationalizable** (or DF-rationalizable⁹) if it is an element of the strategy selection $R \subset S$ which satisfies $\Theta(R) = \tilde{\beta}^\infty(\Theta)$.

Proposition 11 Let $\Gamma \in \mathcal{G}^*$. Every σ -rationalizable strategy for Γ is DF-rationalizable.

Note that for two-player games Proposition 10 actually implies that a strategy is σ -rationalizable if and only if it is DF-rationalizable. In fact for two-player games without pure strategies which are equivalent to a proper mixture of two or more pure strategies we have that, by Proposition 5, a strategy is σ -rationalizable if and only if it survives the DF-procedure (see footnote 9).

Game 7 illustrates that the set of σ -rationalizable strategies, here $\{A\} \times \{C\} \times \{E\}$, may well be a proper subset of the set of DF-rationalizable strategies, here $\{A, B\} \times \{C\} \times \{E\}$.

There are a variety of refinements of the concept of (uncorrelated) rationalizability of Bernheim (1984) and Pearce (1984). The ones we are aware of are **cautious rationalizability** (Pearce (1984)), **perfect rationalizability** (Bernheim (1984)), **proper rationalizability** (Schuhmacher (1999)), **trembling-hand perfect rationalizability**, and **weak perfect rationalizability** (both Herings and Vannetelbosch (1999)).

⁹Dekel and Fudenberg (1990) allow players to hold beliefs which are arbitrary distributions over the set of possible opposition play. This gives rise to what one might call DF correlated rationalizability (see Ritzberger (2002), p.209, for a discussion of rationalizability versus correlated rationalizability; see also Börgers (1994) and Brandenburger (1992) for epistemic conditions under which DF correlated rationalizability is obtained). A strategy is DF correlated rationalizable if and only if it survives the DF-procedure (or $S^\infty W$ -procedure), i.e., one round of elimination of all pure weakly dominated strategies and then the iterated deletion of all pure strictly dominated strategies. The set of DF-rationalizable strategies is contained in the set of correlated DF rationalizable strategies.

Herings and Vannetelbosch (1999) study the relationship between all these concepts. They find that perfect and proper rationalizability both imply weakly perfect rationalizability and provide counter-examples to every other possible set-inclusion. We do not want to go into the various definitions here now, but will just point out how these concepts are related to σ -rationalizability as defined in this paper.

In Game 4 all of the above refinements of rationalizability yield the whole strategy set, while σ -rationalizability leads to the smaller set $\{B, C\} \times \{E, F\}$. In Game 8, trembling hand perfect rationalizability yields, with $\{A\} \times \{D\}$, a subset of the set of σ -rationalizable strategies, $\{A, B\} \times \{D, E\}$. In the game, derived from Game 8 by replacing C and F with strictly dominated strategies, and not changing the payoffs other strategies obtain against C and F , the set of cautiously rationalizable strategies, $\{A\} \times \{D\}$, is a proper subset of the set of σ -rationalizable strategies, again given by $\{A, B\} \times \{D, E\}$. In the semi-reduced normal form, Game 10, of the extensive form game given in Figure 3, the set of properly rationalizable strategies, $\{A\} \times \{F\}$, is smaller than the set of σ -rationalizable strategies, $\{A, B, C\} \times \{D, F\}$. While we thus have no systematic relationship between the concepts of cautious, trembling hand perfect, proper, and σ -rationalizability, it may well be the case that perfect and weakly perfect rationalizability, both as defined in Herings and Vannetelbosch (1999), are, sometimes strictly, weaker criteria than σ -rationalizability. This issue is open.

To illustrate that σ -rationalizability does not always allow the iterated deletion of weakly dominated strategies, unlike trembling-hand perfect rationalizability, consider Game 8 from Samuelson (1992). In this game strategies C and F are weakly dominated, and, hence, not σ -rationalizable. In the semi-reduced game without strategies C and F , strategies B and E are now weakly dominated, and, hence, not trembling-hand perfect rationalizable. However, B (the analogue holds for E) is a best reply against completely mixed strategy profiles close to D , in which the weight on F is greater than the weight on E . Hence, $\mathcal{S}_1(x) = \{A, B\}$ for any such $x \in \Theta$. Hence, B is σ -rationalizable.

	D	E	F
A	1,1	1,1	2,1
B	1,1	0,0	3,1
C	1,2	1,3	1,1

Game 8: A game in which the set of σ -rationalizable strategies includes an iteratively weakly dominated strategy.

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

Game 9: A game in which there are SAS not in any minimal σ -CURB set.

In some special contexts σ -rationalizability does allow the iterated deletion of weakly dominated strategies. See section 7.

6 σ -CURB sets

The following definitions are due to Basu and Weibull (1991). A strategy selection R is a **CURB set** if $\mathcal{B}(\Theta(R)) \subset R$. It is a **tight CURB set** if, in addition $\mathcal{B}(\Theta(R)) \supset R$, and, hence, $\mathcal{B}(\Theta(R)) = R$. It is a **minimal CURB set** if it does not properly contain another CURB set.

We define σ -CURB sets in a similar fashion. A strategy selection R is a **σ -CURB set** if $\mathcal{S}(\Theta(R)) \subset R$. It is a **tight σ -CURB set** if, in addition $\mathcal{S}(\Theta(R)) \supset R$, and, hence, $\mathcal{S}(\Theta(R)) = R$. It is a **minimal σ -CURB set** if it does not properly contain another σ -CURB set.

Note that every CURB set is a σ -CURB set. In fact even every Basu and Weibull's (1991) CURB*-set, a CURB set without weakly dominated strategies, is a σ -CURB set. Game 7 illustrates that a σ -CURB set may well be a proper subset of even a minimal CURB*-set. In this game the unique minimal CURB*-set (and minimal CURB set) is the set $\{A, B\} \times \{C\} \times \{E\}$, while the unique minimal σ -CURB set is the set $\{A\} \times \{C\} \times \{E\}$.

The following definitions are due to Kalai and Samet (1984). A set $\Psi \subset \Theta$ is a **retract** if $\Psi = \times_{i \in I} \Psi_i$, where $\Psi_i \subset \Theta_i$ is nonempty, compact, and convex. A set $\Psi \subset \Theta$ **absorbs** another set $\Psi' \subset \Theta$ if for all $x \in \Psi'$ we have that $\beta(x) \cap \Psi \neq \emptyset$. A retract Ψ is an **absorbing retract** if it absorbs a neighborhood of itself. It is a **persistent retract** if it does not properly contain another absorbing retract. Kalai and Samet (1984) show that, for games without equivalent strategies, and, hence, for games in $\mathcal{G}^*$, persistent retracts have to be faces.

Theorem 2 *Let $\Gamma \in \mathcal{G}^*$. A strategy selection $R \subset S$ is a σ -CURB set if and only if $\Theta(R)$ is an absorbing retract.*

Proof: " $\Leftarrow$ ": Let the strategy selection $R \subset S$ be such that $\Theta(R)$ is an absorbing retract, i.e., it absorbs a neighborhood of itself. Let U be such a neighborhood of $\Theta(R)$. We then have that for every $y \in U$ there is an $r \in R$ such that $r \in \mathcal{B}(y)$. For all $r \in R$ let $U^r = \{y \in U \mid r \in \mathcal{B}(y)\}$. We obviously have $\bigcup_{r \in R} U^r = U$. Suppose R is not a σ -CURB set. Then there is a player $i \in I$ and a pure strategy $s_i \in S_i \setminus R_i$ such that $s_i \in \mathcal{S}_i(x)$ for some $x \in \Theta(R)$. By the definition of $\mathcal{S}_i$ we must then have that $s_i \in \beta(y)$ for all $y \in O$ for some open set O which contains x . But then, by the finiteness of R , there is a strategy profile $r \in R$ such that U^r and O have an intersection which contains an open set. On this set s_i and r_i are now both best replies. But then, by Lemma 1, s_i and r_i are equivalent for player i , which, by Proposition 2, contradicts our assumption. " $\Rightarrow$ ": Suppose $R \subset S$ is a σ -CURB set. Suppose that $\Theta(R)$ is not an absorbing retract. Then for every neighborhood U of $\Theta(R)$ there is a $y_U \in U$ such that $\beta(y_U) \cap \Theta(R) = \emptyset$. In particular for every such y_U there is a player $i \in I$ and a pure strategy $s_i \in S_i \setminus R_i$ such that $s_i \in \mathcal{B}_i(y_U)$. By the finiteness of the number of players and pure strategies and by the compactness of Θ , this means that there is a convergent subsequence of $y_U \in \text{int}(\Theta)$ such that $y_U \rightarrow x$ for some $x \in \Theta(R)$ and there is an $i \in I$ and an $s_i \in S_i \setminus R_i$ such that $s_i \in \mathcal{B}_i(y_U)$ for all such y_U . Now one of two things must be true. Either s_i is a best-reply in an open set with closure intersecting $\Theta(R)$, in which case $s_i \in R_i$ given the definition of σ and a σ -CURB set, which gives rise to a contradiction. Or there is no open set with closure intersecting $\Theta(R)$ such that s_i is best on the whole open set, in which case there must be a strategy $r_i \in R_i$ which is such that $r_i \in \beta(y_U)$ at least for a subsequence of all such y_U , which again gives rise to a contradiction. QED

Theorem 2 allows us to derive two corollaries.

Corollary 2 *In generic normal form games a strategy selection R is a minimal curb set if and only if $\Theta(R)$ is a persistent retract.*

Proof: This follows from Theorem 2 and Proposition 3. QED

Corollary 3 *Consider a strategy selection R in a two-payer game in $\mathcal{G}^*$. Then R is a minimal curb set of the game where all weakly inferior strategies have been eliminated if and only if $\Theta(R)$ is a persistent retract.*

Proof: This follows from Theorem 2 and Proposition 4.

QED

Theorem 2 together with the following lemma can be used to provide an alternative and simple proof of Kalai and Samet (1984)'s proposition that two distinct persistent retracts have an empty intersection. Let R be a strategy selection. Let $C_i(R) = R_i \cup \mathcal{S}_i(\Theta(R))$. Let $C(R) = \times_{i \in I} C_i(R)$. This makes $C(R)$ a strategy selection. Let $C^0(R) = R$ and, for $k \geq 1$, $C^k(R) = C(C^{k-1}(R))$. Finally let $C^\infty(R) = \bigcup_{k=0}^\infty C^k(R)$.

Lemma 7 *Let R be a strategy selection. Then the above defined set $C^\infty(R)$ is a σ -CURB set.*

Proof: By construction.

QED

Note that the set $C^\infty(R)$ is not necessarily a tight nor a minimal σ -CURB set.

Proposition 12 *Let R and R' be two minimal σ -CURB sets. Suppose $R \cap R' \neq \emptyset$. Then $R = R'$.*

Proof: Suppose $s \in R \cap R'$. Then $\{s\}$ is a strategy selection and $C^\infty(\{s\}) \subset R$ by the fact that R is a σ -CURB set. But, analogously, $C^\infty(\{s\}) \subset R'$. By Lemma 7 $C^\infty(\{s\})$ is a σ -CURB set. By the minimality of both R and R' we have the desired result.

QED

Note that the same construction (replacing $\mathcal{S}_i$ with $\mathcal{B}_i$ in the definition of $C_i(R)$) can be used to show that two minimal CURB sets are either disjoint or identical.

We conclude this section with comparing σ -CURB sets to two other recent solution concepts based on strategy selections. Voorneveld (2004) introduces the concept of a prep set (for preparation). A prep set is a strategy selection $R \subset S$ such that $\mathcal{B}_i(\Theta(R)) \cap R_i \neq \emptyset$. Hence every player has at least one best response in her part of the prep set against all beliefs over opponents' play within their parts of the prep set. Generically, see Voorneveld (2005), prep sets and CURB sets are the same, which, by Corollary 2 implies that generically also prep sets and σ -CURB sets (or persistent retracts) are the same. The proof of this last statement was first given in Balkenborg (1992). We reproduce it here in Appendix B. However, other than in generic games, prep sets and σ -CURB sets do not have much in common. Every pure strategy Nash equilibrium, weakly inferior or not, is a minimal prep set. Eliminating weakly inferior (or at least weakly dominated) strategies was one of our objectives with this paper, and no minimal σ -CURB set includes one.

Another recent solution concept based on strategy selections is that of a self admissible set or SAS introduced by Brandenburger, Friedenberg, and Keisler (2008) motivated by epistemic considerations of common knowledge of rationality with the avoidance of weakly dominated strategies built-in. SAS are defined for two-player games only. A strategy selection $R \subset S$ is an SAS if three conditions are satisfied. First, every strategy $s_i \in R_i$ has to be admissible (not weakly dominated). Second, every strategy $s_i \in R_i$ has to be admissible in the restricted game given by strategy space $S_i \times R_{-i}$, where R_{-i} , here, is the single opponent's strategy selection. Third, if $s_i \in R_i$ and s_i is own-payoff equivalent to a proper mixture of two or more pure strategies in S_i then all these pure strategies also have to be in R_i . Recall that, by Proposition 5, weakly inferior strategies in two-player games are those strategies which are either weakly dominated or equivalent to a proper mixture of pure strategies. Recall, further, that by Proposition 4, a best reply which is not weakly inferior is a semi-robust best reply, i.e. in $\mathcal{S}_i(x)$. This implies the following interesting observation.

Proposition 13 *In two-player games, a singleton strategy selection (i.e. a pure strategy profile) is a singleton SAS if and only if it is a fixed point of σ .*

Proof: Let $\{s\}$, for $s \in S$, be a singleton SAS. Let s_i be an arbitrary player i 's part of it. Then the first condition implies that s_i is not weakly dominated, and the third condition implies that it cannot be equivalent to a proper mixture of pure strategies. By Proposition 5 it is, therefore, not weakly inferior, the game being a two-player game. By the second condition of an SAS s must be a Nash equilibrium, i.e. s_i a best reply to s_{-i} . But then by Proposition 4 s_i is a semi-robust best reply, and, hence, s is a fixed point of σ . To see the converse, note that if s is a fixed point of σ it must be such that s_i is not weakly dominated and not equivalent to a proper mixture of pure strategies by Proposition 5, thus satisfying conditions 1 and 3 of an SAS. A fixed point of σ is a best reply to itself and, hence, condition 2 of an SAS is also satisfied. QED

It is, however, not true that σ -CURB sets (not even minimal ones) are also SAS. To see this consider Game 8. The unique minimal σ -CURB set is $\{A, B\} \times \{D, E\}$, yet it is not SAS as strategy B is weakly dominated in the restricted game. It does, however, contain an SAS. In fact three of them: $\{A, D\}$, $\{A, E\}$, and $\{B, D\}$ are all SAS. There are also SAS outside minimal σ -CURB sets. To see this consider Game 9 taken from Figure 2.7 in Brandenburger, Friedenberg, and Keisler (2008). Here, the unique minimal σ -CURB set is $\{D\} \times \{R\}$, which is also an SAS, but there are two more SAS: $\{U\} \times \{L\}$ and $\{U\} \times \{L, R\}$.

In any case exploring the connection with SAS in more detail seems of interest. Also extending SAS concepts and their epistemic foundations to 3 or more player games, perhaps with the requirement not of avoiding weakly dominant strategies, but perhaps weakly inferior strategies could be fruitful.

7 Extensive form games

In this section we investigate what the various concepts based on the refined best-reply correspondence give rise to in extensive form games. We will look at both the agent normal form as well as the semi-reduced normal form.

We first consider extensive form games of perfect information (EFGOPI). Note that the agent normal form of such games is in $\mathcal{G}^*$ as long as no player has 2 or more equivalent actions at any of her information sets (which here are singletons, i.e., nodes). Not every normal form derived from even a generic extensive form game of perfect information (GEFGOPI) is in $\mathcal{G}^*$. Consider the 1-player extensive form game, given in Figure 2, in which at node 1 the player has two choices, L and R , where L terminates the game, while R leads to a second node, where the player again faces two choices l and r . The two pure strategies Ll and Lr are obviously equivalent. The semi-reduced normal form has been introduced to eliminate exactly this type of equivalences. The semi-reduced normal form of any GEFGOPI is again in $\mathcal{G}^*$.




Figure 2: A 1-player extensive form game.

Proposition 14 *Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of a GEFGOPI. Then only the (unique) subgame-perfect strategy profile is σ -rationalizable.*

Proof: Consider a final node. A strategy, available to the player, say, i at this node, which is not subgame perfect is weakly dominated. Hence, it can not be in $\mathcal{S}_i(x)$ for any $x \in \Theta$. So it is not in $\sigma(\Theta)$. Now consider an immediate predecessor node to the above final node. A non-subgame perfect strategy at this node can only be a best-reply if the behavior at the following nodes is non-subgame perfect. For any $x \in \Theta$ in a neighborhood of $\sigma(\Theta)$ this is still true. Hence, any such non-subgame perfect strategy at this node can not be in $\sigma^2(\Theta)$. This argument can be reiterated any finite number of times. QED

Proposition 15 *Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of a GEFGOPI. The only fixed point of σ for this game is the (unique) subgame perfect equilibrium.*

Proof: Every fixed point of σ is in the set of σ -rationalizable strategies. This set, by Proposition 14, only consists of the subgame perfect equilibrium. QED




Figure 3: A centipede game.


	D	E	F
A	3,0	3,0	3,0
B	4,3	1,2	1,2
C	4,3	0,1	2,4

Game 10: The normal form game of the centipede game in Figure 3.

None of the above propositions is true for the semi-reduced normal form. Consider the centipede game (Figure 8.2.2 in Cressman (2003)) given here in Figure 3 with semi-reduced normal form given as Game 10, where player 1's strategies are $A = Ll|Lr$, $B = Rl$, and $C = Rr$, while player 2's strategies are $D = Ll|Lr$, $E = Rl$, and $F = Rr$. This game is a GEFGOPI and, hence, has a unique subgame perfect equilibrium, which is (Lr, Rr) . The set of σ -rationalizable strategies is $\{A, B, C\} \times \{D, F\}$, a lot more than just the subgame perfect strategy-profile. Also the non-subgame perfect, and, hence, non weak-perfect Bayesian and non-sequential, Nash equilibrium (B, D) is a fixed point of σ . So indeed, fixed points of σ in a given normal form game need not

induce sequential or even weak perfect Bayesian equilibria in every extensive form game with this semi-reduced normal form.

Also not every sequential equilibrium is necessarily a fixed point of σ . The game given in Figure 4, Figure 13 in Kreps and Wilson (1982), has a sequential equilibrium (L, r) which is not a fixed point of σ (it is not perfect). Here the agent normal form and the semi-reduced normal form are the same and given as Game 11.



	l	r
L	1,1	1,1
R	2,0	-1,-1

Figure 4: A game with a sequential equilibrium (L, r) which is not a fixed point of σ .

Game 11: The normal form of the game in Figure 4.

There are even extensive form games in agent normal form in which a fixed point of σ is not a sequential equilibrium. Consider the game in Figure 5. The Nash equilibrium (A, R, r) is a fixed point of σ , but is not sequential and, hence, not extensive form trembling hand perfect.




Figure 5: A game in which there is a fixed point of σ in the agent normal form which is not sequential (and, hence, not extensive form trembling hand perfect).

To see that (A, R, r) is a fixed point of σ we need to check that each strategy choice is a best reply in an open set around (A, R, r) . For player 1's choice A this is definitely true as A weakly dominates both B and C . Player 2's choice R is best as long as player 1 is sufficiently more likely to tremble to C than to B . In fact the probability of C has to be at least twice that of B . Player 2's payoffs are unaffected by player 3's choice. Player 3's choice r is best as long as player 1 trembles sufficiently more to B than to C . In fact the probability of B has to be at least twice that of C . This is true for whatever player 2 does. Hence, for each player's strategy choice there is an open set of strategy profiles around (A, R, r) against which the player's choice is a best reply. Hence, (A, R, r) is indeed a fixed point of σ . However, these open sets (for players 2 and 3) are mutually exclusive. This in turn means that there is no system of consistent beliefs for players 2 and 3 which make both choices R and r best replies simultaneously. Player 2's belief that sustains the (A, R, r) equilibrium is such that his first node has conditional probability of at most $1/3$. Player 3's belief that sustains the (A, R, r) equilibrium is such that her first node

has conditional probability of at least $2/3$. But in a sequential equilibrium these two beliefs would have to coincide. Thus this (A, R, r) is not sequential (and not trembling-hand perfect).

The following proposition states that any pure fixed point of σ in the agent normal form of an extensive form game induces a (weak) perfect Bayesian equilibrium (see e.g., Definition 9.C.3 in Mas-Collel, Whinston, and Green (1995) or Definition 6.2 in Ritzberger (2002)) in this extensive form game.

Proposition 16 *Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of an extensive form game. Then if a pure strategy profile s is a fixed point of σ it induces a perfect Bayesian equilibrium in this extensive form game.*

Proof: Let s_i be player i 's part of the pure strategy profile s . Given s is a fixed point of σ we have by Lemma 5 that there is an open set $U^{s_i} \subset \Theta$ with its closure containing s , such that $s_i \in \mathcal{B}_i(y)$ for any $y \in U^{s_i}$. But then there is a sequence of completely mixed $y_t \in \text{int}\Theta$ such that y_t converges to s and $s_i \in \mathcal{B}_i(y_t)$ for all these y_t . Being completely interior every such y_t induces a unique probability distribution over all nodes in all information sets of every player. In particular also for player i . But then there is a unique consistent belief for player i , μ_t , given y_t . But then the sequence μ_t must have a convergent subsequence, which converges to some feasible belief μ , consistent with s wherever possible, and such that s is optimal given belief μ . Hence, s is a perfect Bayesian equilibrium. QED

Whether this result is true for mixed fixed-points is not clear. The difficulty here is that if the support of x_i , $C(x_i)$, contains two pure strategies, one may not be able to find one belief μ justifying both pure strategies. It may be the case that both pure strategies are justifiable, but only with different beliefs. But then x would not be a weak perfect Bayesian equilibrium.

However, fixed points of σ in the agent normal form have another surprising property. They induce fixed points of σ in every subgame.

Proposition 17 *Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of a given extensive form game. Then if a strategy profile x is a fixed point of σ it is also a fixed point of σ in the agent normal form of every subgame of this extensive form game.*

Proof: Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of the given extensive form game. Let Θ denote the space of mixed strategies. Let $x \in \Theta$ be a fixed point of σ . Consider player (agent) $i \in I$, where I is the set of all agents. Player i only moves once, i.e. has only one information set. By the definition of $\sigma_i(x) = \Delta(\mathcal{S}_i(x))$, player i , in $x_i \in \sigma_i(x)$ is exclusively randomizing over pure strategies, each of which are unique best replies to some $x' \in \Theta$ (possibly different for different $s_i \in \mathcal{S}_i$) in a neighborhood of x . In fact, for each $s_i \in \mathcal{S}_i(x)$, and, hence, for each $s_i \in C(x_i)$ we have that there is a sequence of $x^t \in \Theta$ such that $x^t \rightarrow x$ and $\mathcal{B}_i(x^t) = \{s_i\}$. Now consider any subgame in which player i also moves. Let $\Gamma' = (I', S', u')$ denote its agent normal form. Obviously $I' \subset I$ and for all $i \in I'$ we have $S'_i = S_i$ and u' is defined accordingly. Now for every $s_i \in C(x_i)$ consider the projection of $x^t \in \Theta$ onto the reduced game Γ' . Let it be denoted by $\hat{x}^t \in \Theta'$. Hence we simply have that $\hat{x}_i^t = x_i^t$ for all $i \in I'$. Now consider for the exact sequence of $x^t \in \Theta$ such that $x^t \rightarrow x$ and $\mathcal{B}_i(x^t) = \{s_i\}$ its projection. Obviously we have that $\hat{x}^t \rightarrow \hat{x}$ and also we must have that $\mathcal{B}'_i(\hat{x}^t) \subset \mathcal{B}_i(x^t)$. This is so because either in x player i 's information is reached, in which case player i 's best responses cannot change in the subgame, or player i 's information set is not reached in x , and, hence, every strategy of player i is a best response

against x in the full game. But now given $\mathcal{B}_i(x^t) = \{s_i\}$ we must also have $\mathcal{B}'_i(\hat{x}^t) = \{s_i\}$ for the whole sequence of $\hat{x}^t$. But this is nothing but saying that $s_i \in \mathcal{S}'_i(\hat{x})$ and, as this is true for all $s_i \in C(x_i)$ and all players $i \in I'$ we have that $\hat{x}$ is a fixed point of σ in Γ' . QED

Proposition 17 therefore implies that every fixed point of σ in the agent normal form of an extensive form game is automatically subgame perfect. Propositions 17 and 16 together imply that every **pure** fixed point of σ in the agent normal form of an extensive form game is weak perfect Bayesian in every subgame.

8 Micro models of learning

In this section we shall sketch two micro models of learning in the spirit of the models in Björnerstedt and Weibull (1996) which give rise to the (unrefined) best-reply dynamics (1) and the minimal refined best reply dynamics (2), respectively. In section 9 we shall discuss properties of the latter dynamics in some detail.

8.1 A micro model leading to the (unrefined) best-reply dynamics

This subsection provides a model of individual learning which gives rise to the original, unrefined best-reply dynamics (1). Suppose there is a continuum of agents for each player $i \in I$. Players only play pure strategies. Then a (mixed) strategy-profile $x \in \Theta$ represents a state in the following sense. For player population $i \in I$, x_{is} denotes the proportion of agents in this population who play pure strategy $s \in S_i$. Over time agents review their strategies at a given rate, $r = 1$, which we will assume fixed and the same for all agents in all populations. Any agent, in any population, who is reviewing her strategy is assumed to switch to any pure best reply against the current state x . If the agent is currently already playing a best reply the agent may nevertheless switch to an alternative best reply if there is one. Suppose $s \in S_i$ is such that $s \notin \mathcal{B}_i(x)$. Then every reviewing s -strategist will switch away from strategy s to a best-reply, while no other agent will switch to s either. Hence, $\dot{x}_{is} = -x_{is}$. Now suppose $\{s\} = \mathcal{B}_i(x)$, i.e., s is the unique best reply to current state $x \in \Theta$. Then every reviewing s -strategist will remain to be one, while every other reviewing agent will switch to s . Hence, $\dot{x}_{is} = \sum_{t \neq s} x_{it} = 1 - x_{is}$.

Suppose, finally, that $s \in \mathcal{B}_i(x)$ and $\mathcal{B}_i(x)$ is not a singleton. Then reviewing s -strategists may or may not switch to something else, while all other reviewing agents may or may not switch to s . We could here specify for every such situation (or state $x \in \Theta$) a fixed particular probability distribution over best replies with probability mass function $p_i(x) \in \beta_i(x)$, which individuals use to randomly choose one of their best replies given state x . For instance, for a given state x and a given player i with $\mathcal{B}_i(x)$ containing at least two elements we could specify, for a given $s \in \mathcal{B}_i(x)$, that $p_i(x)(s) = \alpha \in [0, 1]$. In this case we would have $\dot{x}_{is} = (1 - x_{is})\alpha - x_{is}(1 - \alpha)$, which leads to $\dot{x}_{is} = \alpha - x_{is}$.

Different choices of $p_i(x)$ would then generate different systems of differential **equations**. Making a particular choice for all $p_i(x)$ functions is problematic in two ways, one conceptually, and one technically. The conceptual problem with choosing a particular set of $p_i(x)$ functions is that it is simply arbitrary. Why would one choice of $p_i(x)$ functions be preferred over another? The technical problem is that a particular choice of $p_i(x)$ typically leads to a discontinuous vector field, which generally does not admit solutions for all initial conditions. Both problems are solved if we simply allow at every state $x \in \Theta$ incremental changes towards all best responses.

But this then is nothing but the system of differential inclusions

$$\dot{x} \in \beta(x) - x, \quad (1)$$

which is the best reply dynamics of Gilboa and Matsui (1991).

8.2 A micro model leading to the refined best-reply dynamics

In section 9 we will finally consider the *refined best-reply dynamics*

$$\dot{x} \in \sigma(x) - x, \quad (2)$$

where σ is as in section 3. In this section we provide a micro-motivation for this refined best-reply dynamics (2) based on the above given micro model which leads to the unrefined best-reply dynamics (1).

In order to motivate their best-reply dynamics Gilboa and Matsui (1991) also sketch a micro model of learning. In Gilboa and Matsui (1991)'s story, however, it is assumed that agents do not exactly know the current state, or, as Gilboa and Matsui (1991) call it, the current behavior pattern. In fact they assume that "... , there is a limitation [on the agents part] of recognizing the current behavior pattern ..." and that agents choose a "... best response to a possibly different behavior pattern which is in the ϵ -neighborhood of the current one." (Gilboa and Matsui (1991), p. 863).

Let us here now also assume that agents do not exactly know the current state $x \in \Theta$, but we will force them to hold a belief about the current state, drawn from some (absolutely continuous) distribution over the intersection of Θ and an ϵ -ball around x . Agents then choose a best reply to their belief.

To make this precise, we assume that, at some time t , every reviewing agent always holds a prior belief $\mu_0 \in \Theta$ about x where each agent's μ_0 is independently drawn from a distribution F on Θ , where F is an arbitrary distribution with a density that is positive almost everywhere, i.e., this density is 0 only on a set with Lebesgue-measure 0. This means there is heterogeneity in agents' prior beliefs. Every agent then learns what a proportion of $1 - \epsilon$ of all agents in every population are doing and updates her belief accordingly. This updated belief μ_1 then has a distribution which has support only within an ϵ -ball, U_ϵ^x , around the true state x . This ϵ -ball is with respect to the sup-norm, i.e., $U_\epsilon^x = \{y \in \Theta \mid \|x - y\|_\infty \leq \epsilon\}$, where $\|\cdot\|_\infty$ denotes the sup (or max) norm. The density of this posterior distribution is then positive almost everywhere within U_ϵ^x , i.e., within U_ϵ^x it is 0 only on a set with Lebesgue-measure 0 again.

An agent after updating her belief will then play a best response to her belief. Given the heterogeneity in beliefs different agents might choose different best responses. In fact a particular choice of F , the distribution agents independently choose their priors from, for any given i and any given state x , leads to a particular distribution $p_i(x)$ which reviewing individuals choose their best replies from, where $p_i(x)$ is very much as in the micro story sketched in the previous subsection. However, now not all such $p_i(x) \in \beta_i(x)$ are possible. In fact, given the essentially full support assumption for F , pure strategies which are best replies to x only on a thin set (Lebesgue-measure 0), such as $\Psi \cap U_\epsilon^x$, within the ϵ -ball around x will only be chosen by a vanishing fraction of reviewing agents for all such prior distributions F and, hence, have to receive probability 0 in $p_i(x)$. In fact we must have $p_i(x) \in \sigma_i(x)$.

Different choices of F will lead to different systems of refined best-response differential **equations**. Again both for conceptual and technical reasons it seems more reasonable to allow all


such differentials towards refined best responses. This, if ϵ is small enough, then leads to the differential inclusion (2), which we call the minimal refined best-reply dynamics.

9 The refined best-reply dynamics: Results

Gilboa and Matsui (1991), Matsui (1992) and Hofbauer (1995) introduced and studied the continuous time best reply dynamics (1), which is, modulo a time change, equivalent to Brown (1951)'s continuous time version of fictitious play. A solution to (1) is an absolutely continuous function $x(t)$, defined for at least all $t \geq 0$, that satisfies (1) for almost all t .¹⁰ General theory guarantees the existence of at least one solution $\xi(t, x_0)$ through each initial state x_0 . In general, several solutions can exist through a given initial state. In some games, there appear to be too many of them.

For Game 12, within the component of NE any function $x(t)$ with $-x_i \leq \dot{x}_i \leq 1 - x_i$ (i.e., which does not move too quickly) is a solution while all nearby interior solutions move straight to AC , see Figure 6. We can, in fact, eliminate these unnatural solutions by looking at the refined best reply dynamics (2).

	C	D
A	1,1	1,1
B	1,1	0,0



Game 12: A game in which the BR dynamics seems to have ‘too many’ solutions.

Figure 6: Best response dynamics of Game 12.

Since the right hand side of (2) is UHC with compact and convex values, existence of at least one Lipschitz-continuous solution $x(t)$ through each initial state $x(0)$ is again guaranteed.


The mathematical motivation to consider this refined inclusion is the classical approach (due to Filippov, see Aubin and Cellina (1984)) to regularize a differential equation with a piecewise smooth right hand side. In our case this means, we view the best reply dynamics (1) as a piecewise linear differential equation, defined for x in the open dense set Ψ only. In this approach one considers at each point of discontinuity (i.e., $x \notin \Psi$) the convex hull of all limit points of nearby values. This leads to the smallest UHC correspondence with compact convex values that contains the graph of the given discontinuous single-valued function. Applying this

¹⁰Gilboa and Matsui (1991) and Matsui (1992) require additionally the right differentiability of solutions. Hofbauer (1995) argued that all solutions in the sense of differential inclusions should be admitted. This is natural for applications to discrete approximations (fictitious play, see Hofbauer and Sorin (2006)) or stochastic approximations, see Benaim, Hofbauer, and Sorin (2005). Note that any absolutely continuous solution is automatically Lipschitz, since the right hand side of (1) is bounded. Hofbauer (1995) also provides an explicit construction of all piecewise linear solutions (for 2 person games) and provides conditions when these constitute all solutions. See also Hofbauer and Sigmund (1998) and Cressman (2003).

idea to games (in the class $\mathcal{G}^*$) leads to the correspondence σ and the dynamics (2) instead of the classical best reply correspondence β and the dynamics (1).

As another simple example take Game 13 which is the restriction of Game 9 without the strictly dominated strategy M . This game is also known as the ultimatum minigame, see Cressman (2003). This game has a strict Nash equilibrium (D, R) and a component of Nash equilibria, where the column player uses the weakly dominated strategy L . The strict equilibrium is the only fixed point of σ and the global attractor for (2), see Figure 7.

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



Game 13: Ultimatum minigame.

Figure 7: Best response dynamics of Game 13.

We show now that the refined best-reply dynamics converges to the set of σ -rationalizable strategies, and that every σ -CURB set is asymptotically stable under this dynamics. The proofs are the same as the proofs of the statements that every solution of the best-reply dynamic (1) converges to the set of rationalizable strategies and that every CURB set is asymptotically stable under the best-reply dynamics. These results are analogous to the results of Hurkens (1995), who for a stochastic learning model a la Young (1993) showed that recurrent sets coincide with CURB sets or persistent retracts depending on the details of the model.

Theorem 3 *Let $\Gamma \in \mathcal{G}^*$. Let R be the strategy selection of S which spans the set of σ -rationalizable strategies, i.e., $\Theta(R) = \sigma^\infty(\Theta)$. Let $s_i \in S_i \setminus R_i$. Then $x_{is_i}(t) \rightarrow 0$ for any solution $x(\cdot)$ to (2) for any initial state $x(0) \in \Theta$.*

Proof: The proof is by induction on k , the iteration in the deletion process, i.e., the k in $\sigma^\infty(\Theta) = \bigcap_{k=1}^\infty \sigma^k(\Theta)$. Let R^k denote the strategy selection of S which spans $\sigma^k(\Theta)$, i.e., $\Theta(R^k) = \sigma^k(\Theta)$. For $k = 1$ consider an arbitrary strategy $s_i \in S_i \setminus R_i^1$. By definition then $s_i \notin S_i(x)$ for any $x \in \Theta$. Hence its growth rate according to (2) is $\dot{x}_{is_i} = 0 - x_{is_i}$, and therefore

$$x_{is_i}(t) = e^{-t} x_{is_i} \quad (3)$$

for all $t \geq 0$, i.e., $x_{is_i}(t)$ shrinks exponentially to zero. This proves the statement of the theorem for $s_i \in S_i \setminus R_i^1$. Now assume the statement of the theorem is true for $s_i \in S_i \setminus R_i^{k-1}$, i.e., for any such s_i we have that $x_{is_i}(t) \rightarrow 0$ for any solution $x(\cdot)$ to (2) for any initial state $x(0) \in \Theta$. Then for any such s_i and for any $x(0) \in \Theta$ there is a finite T such that $x_{is_i}(t) < \epsilon$ for all $t \geq T$. Now by the definition of σ , $s_i \in S_i \setminus R_i^k$ implies that $s_i \notin S_i(x(t))$ provided ϵ is small enough (or t large enough). But then for all $t \geq T$ we again have that $\dot{x}_{is_i} = 0 - x_{is_i}$ and, hence, that $x_{is_i}(t)$ shrinks exponentially to zero. QED

Given a differential inclusion on Θ and the corresponding multivalued semi-flow $\Phi^t : \Theta \rightarrow \Theta$ ($t \geq 0$), see Benaim, Hofbauer, and Sorin (2005), we define its **global attractor** as the set $\mathcal{A} = \bigcap_{t \geq 0} \Phi^t(\Theta)$. It consists of all $x \in \Theta$ for which there is a full solution $x(t) \in \Theta$ (i.e., defined for all $t \in \mathbb{R}$) with $x(0) = x$. Since a forward solution exists always, the restrictive requirement is that this can be extended backwards for all negative times, without leaving Θ . The global attractor is compact, invariant and asymptotically stable, indeed the maximal subset of Θ with these properties. For applications of this useful concept to the best-reply dynamics see Hofbauer (1995), Cressman (2003), and Hofbauer and Sorin (2006).

For Game 13 (see also Figure 7) the global attractor of the best reply dynamics (1) consists of all Nash equilibria and the line segment from $(U, \frac{1}{2}(L+R))$ to (D, R) which is the connecting orbit between the component and the strict Nash equilibrium. This is also the global attractor of both (1) and (2) for the larger Game 9. However, for Game 13 and the refined best reply dynamics (2) the global attractor consists of the strict equilibrium (D, R) only.

Theorem 4 *Let $\Gamma \in \mathcal{G}^*$. Then the global attractor of (2) is contained in the set of σ -rationalizable strategies $\sigma^\infty(\Theta)$.*

Proof: Let $\mathcal{A}$ be the global attractor of (2). As in the previous proof we show by induction on k : If $s_i \in S_i \setminus R_i^k$, and $x \in \mathcal{A}$ then $x_{is_i} = 0$. Indeed, this follows since (3) must hold for all $t \in \mathbb{R}$, and $x_{is_i}(t)$ would be unbounded for $t \rightarrow -\infty$ otherwise. Hence $\mathcal{A} \subset \sigma^\infty(\Theta)$. QED

Theorem 5 *Let $\Gamma \in \mathcal{G}^*$ and R a σ -CURB set. Then $\Theta(R)$ is asymptotically stable under (2).*

Proof: By the definition of σ and a σ -CURB set we have that for any $x \in U$ where U is a sufficiently small neighborhood of $\Theta(R)$ it is true that for any $i \in I$ $s_i \in \mathcal{S}_i(x)$ implies $s_i \in R_i$. Hence, for any $x \in U$ we must have that $\dot{x}_{is_i} = -x_{is_i}$ for all $i \in I$ and $s_i \notin R_i$. But then we must have that $\|x(t) - \Theta(R)\|_\infty$ shrinks exponentially to zero for all $x(0) \in U$. QED

Theorem 5 has important consequences. It provides us with the main theorem of this paper. Consider a differential inclusion $\dot{x} \in \tau(x) - x$, where $\tau : \Theta \Rightarrow \Theta$ is a correspondence which satisfies properties 1 (product structure), 3 (non-empty valued), 4 (convex valued), and 5 (upper hemi-continuous) and property 2 is replaced by $\tau_i(x) \cap \beta_i(x) \neq \emptyset \forall x \in \Theta, \forall i \in I$ (i.e. $\tau(x)$ includes at least one best reply to x). Then analogous to Theorem 5 we have that every τ -CURB set (defined analogously to a σ -CURB set) is asymptotically stable under its induced differential inclusion $\dot{x} \in \tau(x) - x$. In addition, for every correspondence τ' satisfying the same 5 properties and, in addition, $\tau'(x) \subset \tau(x)$ for all $x \in \Theta$ (i.e. τ' is a refinement of τ) we have that every τ -CURB set is also a τ' -CURB set. Thus every τ -CURB set is also asymptotically stable under the differential inclusion $\dot{x} \in \tau'(x) - x$. However, for every correspondence τ satisfying the above 5 properties we have that $\sigma(x) \subset \tau(x)$ for all $x \in \Theta$ by the minimality of σ among refined best-reply correspondences and by the property of τ that $\tau_i(x) \cap \beta_i(x) \neq \emptyset \forall x \in \Theta, \forall i \in I$, which σ also satisfies. Thus every face that is asymptotically stable under some differential inclusion $\dot{x} \in \tau(x) - x$, where τ satisfies the 5 properties above, is also asymptotically stable under the minimal refined best-reply correspondence. We thus have the following main theorem.

Theorem 6 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau : \Theta \Rightarrow \Theta$ be a correspondence which satisfies properties 1 (product structure), 3 (non-empty valued), 4 (convex valued), and 5 (upper hemi-continuous)*

and property 2 is replaced by $\tau_i(x) \cap \beta_i(x) \neq \emptyset \forall x \in \Theta, \forall i \in I$ (i.e. $\tau(x)$ includes at least one best reply to x). Then every asymptotically stable face under the induced differential inclusion $\dot{x} \in \tau(x) - x$ is also asymptotically stable under the minimal refined best-response dynamics (2).

A corollary of Theorem 5, combined with Theorem 2, is that Kalai and Samet's (1984) absorbing retracts are asymptotically stable under the refined best-reply dynamics (2).

Corollary 4 *Let $\Gamma \in \mathcal{G}^*$. Let R be a strategy selection such that $\Theta(R)$ is an absorbing retract. Then $\Theta(R)$ is asymptotically stable under (2).*

It is not true that the refined best-reply dynamics necessarily converges to a persistent retract. A simple example is a coordination game where there are some solutions that end up in the completely mixed equilibrium. Worse, in Game 5, the non-persistent equilibrium (B, D, E) attracts an open set of initial values.

In some games there are sets which are proper subsets of persistent retracts which are asymptotically stable. A simple example is the matching pennies game where the unique equilibrium is asymptotically stable, in fact the global attractor. (This holds for any two-person zero-sum game with a unique interior equilibrium, see Brown (1951) and Hofbauer and Sorin (2006).) Another example is Game 8. Here, the unique persistent retract is the set of σ -rationalizable strategies $\sigma^\infty(\Theta) = \Delta(\{A, B\}) \times \Delta(\{D, E\})$. The subset $\mathcal{A} = \{x \in \sigma^\infty(\Theta) | x_{1B}x_{2E} = 0\}$, which is not a retract, is also asymptotically stable, being the global attractor. Note also that the refined best-response correspondence of this game constrained to the set $\Delta(\{A, B\}) \times \Delta(\{D, E\})$ is exactly the same as the best-response correspondence in Game 12.

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

Game 14: A game in which the global attractor is much smaller than the set of σ -rationalizable strategies.




Figure 8: Best response dynamics of Game 14, ignoring the strictly dominated strategy M .

In Game 14, the global attractor is even contained in a proper subspace of the set of σ -rationalizable strategies. In this game, M is strictly dominated. The correspondences σ and β coincide. The set of (σ) -rationalizable strategies equals $\Delta(\{U, D\}) \times \Delta(\{L, R\})$. The game restricted to this face is a 'chain store' game as in Cressman (2003). The global attractor equals the set of fixed points of σ (which equals the set of Nash equilibria), $\{D\} \times \Delta(\{\frac{1}{2}(L + R), R\})$. Most orbits converge to (D, R) , see Figure 8. Note that only part of the face $\{D\} \times \Delta(\{L, R\})$ is invariant under (2). Starting in (D, L) leads first into the face $\Delta(\{U, D\}) \times \Delta(\{L, R\})$ and only in the limit to (D, R) .

A subset of Θ which is spanned by a strategy selection R , however, if it is asymptotically stable under (2), then R must be a σ -CURB set and, hence, $\Theta(R)$ an absorbing retract.

Theorem 7 *Let $\Gamma \in \mathcal{G}^*$. Let R be a strategy selection such that $\Theta(R)$ is asymptotically stable under (2). Then R is a σ -CURB set, and, hence, $\Theta(R)$ is an absorbing retract.*

Proof: Let $\Gamma \in \mathcal{G}^*$. Let R be a strategy selection such that $\Theta(R)$ is asymptotically stable under (2). Suppose R is not a σ -CURB set. Then there is an $x \in \Theta(R)$ and a player $i \in I$ such that there is a $s_i \in \mathcal{S}_i(x) \setminus R_i$. By the definition of $\mathcal{S}_i$ this implies that $s_i \in \mathcal{B}_i(y)$ for all y in an open set $U \subset \Theta$ containing x . By the definition of $\mathcal{G}^*$ there then must be an open subset U' of U such that $\mathcal{B}_i(y)$ is a singleton for all $y \in U'$, and, hence, $\mathcal{B}_i(y) = \{s_i\}$. But then for every solution $y(\cdot)$ to (2) starting at any $y \in U'$ we must have that $\dot{y}_{i,s_i}(t) = 1 - y_{i,s_i}(t)$, which must be close to 1. Hence, for all initial states $y \in U'$ $y_{i,s_i}(t)$ initially grows exponentially, and, hence, $\Theta(R)$ is not asymptotically stable under (2), providing a contradiction. QED

Theorem 7 implies that a minimal asymptotically stable face of Θ must be a persistent retract or minimal σ -CURB set. One cannot replace face with closed and convex set in this statement. A simple example is again the matching pennies game.

A corollary to Theorems 5 and 7 is the following.

Corollary 5 *Let $\Gamma \in \mathcal{G}^*$. A state $x \in \Theta$ is asymptotically stable under (2) if and only if it is a robust equilibrium point (Okada (1983)).*

This follows from the fact that a robust equilibrium point is a singleton persistent retract. Note that in games in $\mathcal{G}^*$ a robust equilibrium point must be a pure strategy profile.

10 Index theory

The simplest definition of the index of an isolated Nash equilibrium or a component of Nash equilibria C is via the best reply correspondence: $i(C)$ is defined as the Lefschetz fixed point index of C under the best reply correspondence. Earlier definitions use the replicator dynamics and Brouwer degree, as in Hofbauer and Sigmund (1998) (chapter 13.2-4) or in Ritzberger (2002) (section 6.5), which leads to an explicit formula for regular equilibria in terms of the sign of a certain determinant.

For a correspondence $F : \Theta \rightarrow \Theta$ satisfying properties 3-5 from section 3 and an open subset $U \subset \Theta$ such that ∂U contains no fixed points of F , the Lefschetz fixed point index $\Lambda(F, U)$ is defined and satisfies a number of properties, in particular the five axioms in section 4 of McLennan (1989). If C is a component of Nash equilibria then define the **index of C** as $i(C) = \Lambda(\beta, U)$ where U is an open neighborhood of C such that $C = \{x \in \bar{U} : x \in \beta(x)\}$, i.e., all Nash equilibria in the closure of U are in C . The additivity axiom shows that this definition does not depend on the choice of U . We then define the **σ -index of C** as $i_\sigma(C) = \Lambda(\sigma, U)$ where C is a closed set (e.g., a component of fixed points of σ , a component of Nash equilibria, or a persistent retract) and U an open neighborhood of C such that $C = \{x \in \bar{U} : x \in \sigma(x)\}$.

The continuity axiom (Axiom 2 in section 4 of McLennan (1989)) implies that $i(C) = i_\sigma(C)$ holds for every component C of Nash equilibria.

Proposition 18 *Every Nash equilibrium component with nonzero index contains a fixed point of σ .*

Proof: Let C be a Nash equilibrium component. Then $i(C) = i_\sigma(C) = i_\sigma(\{x \in C : \sigma(x) = x\})$. Hence, if $i(C) \neq 0$, C contains a fixed point of σ . QED

Theorem 8 *Every Nash equilibrium component C contains finitely many components C_i of σ fixed points. Furthermore, $i(C) = \sum i_\sigma(C_i)$.*

Proof: For two person games the first statement follows from Corollary 1. For more players we need the following lemma which generalizes a corresponding statement about the set of Nash equilibria in Kohlberg and Mertens (1986). The second statement follows then from the additivity axiom. QED

Lemma 8 *For every game in $\mathcal{G}^*$ the set of fixed points of σ is a semi-algebraic set and therefore has only finitely many components.*

Proof: We use the Tarski-Seidenberg principle and some of its consequences, as outlined in Bochnak, Coste, and Roy (1998), to show that the set of fixed points of σ is semi-algebraic and hence the result follows from Theorem 2.4.5 in Bochnak, Coste, and Roy (1998). We notice first that the set $\Psi(s_i)$ consisting of all $x \in \Theta$ for which a pure strategy $s_i \in S_i$, $i \in I$, is a unique best reply is an open semi-algebraic set because it consists of all $x \in \Theta$ for which the inequality $u_i(s_i, x_{-i}) > u_i(t_i, x_{-i})$ holds for all $t_i \neq s_i$. The closure of this set $\bar{\Psi}_i(s_i)$ is the set where s_i is a semi-robust best reply. It is a semi-algebraic set by Proposition 2.2.2 in Bochnak, Coste, and Roy (1998). For every $s_i \in S_i$, $i \in I$, we conclude that the set of all $x \in \Theta$ for which $x_{is_i} = 0$ or $x \in \bar{\Psi}_i(s_i)$ is semi-algebraic because it is a union of semi-algebraic sets. The intersection of all these sets over varying $i \in I$ and $s_i \in S_i$, which is the set of fixed points of σ , is hence also semi-algebraic. QED

	D	E	F
A	1,2	1,0	1,0
B	1,0	1,2	0,0
C	0,2	1,2	2,0

Game 15: A game with a unique component of Nash equilibria, which "breaks up" into two (much smaller) subcomponents of fixed points of σ .

	D	E	F
A	1,1	0,-1	-1,1
B	-1,0	0,0	-1,0
C	1,-1	0,-1	-2,-2

Game 16: A game from Kohlberg and Mertens (1986).

Game 15 provides an example of a Nash equilibrium component C that contains more than one component of σ fixed points. In this game there is a single closed connected component of Nash equilibria. It is the union of two line segments, $\Delta(A, \frac{A+B}{2}) \times D$ and $\frac{A+B}{2} \times \Delta(D, E)$, and the triangle $\Delta(\frac{A+B}{2}, B, C) \times E$. Note that strategy F is strictly dominated and B is weakly dominated. There are only two fixed points of σ . These are (A,D) (which is asymptotically stable under (2) and therefore has σ -index 1) and (C,E) (with σ -index 0). The global attractor for (2) consists of these two equilibria and a connecting orbit.

Proposition 19 *Every σ -CURB set has σ -index 1.*

Proposition 20 *If a component C of fixed points of σ is asymptotically stable under the refined best reply dynamics (2) then $i_\sigma(C) = \chi(C)$.*

This is a multi-valued version of Theorem 1 of Demichelis and Ritzberger (2003). It follows from a more general result of DeMichelis and Sorin (personal communication).

Consider now Game 16, an example from Kohlberg and Mertens (1986) (p. 1034), see also Demichelis and Ritzberger (2003) (Example 3, p. 71). In this symmetric game, the second and the third strategy are weakly dominated. Hence AD is the unique perfect equilibrium and the unique fixed point of σ . The game has a unique NE component. It is a hexagon spanned by the sequence of pure Nash equilibria AD–AF–BF–BE–CE–CD–AD. Hence its Euler characteristic is 0, but its index is 1. Hence it is not asymptotically stable by Theorem 1 of Demichelis and Ritzberger (2003). From an evolutionary perspective this game presents a puzzle: There is a unique component of equilibria, yet it cannot be asymptotically stable. In this game every asymptotically stable set under any reasonable dynamics must contain non-equilibrium strategies.¹¹ However, AD is the only σ -rationalizable strategy, and therefore, by Theorem 4, AD is the global attractor under the refined best-reply dynamics (2). Thus, the refined best-reply dynamics resolves the above puzzle.

Finally, we conjecture that every σ -CURB set, and more generally, every set with nonzero σ -index contains a strategically stable set in the sense of Mertens.

11 Conclusion

In this paper we endeavored to find the smallest sets (corresponding to strategy selections) of strategy profiles that can be justifiably called evolutionarily stable. To be more precise we characterize the, in a well-specified sense, smallest set of strategy profiles which is asymptotically stable under a deterministic evolutionary process. In order to do so we introduce refinements of the best-reply correspondence which satisfy 5, we believe reasonable, criteria. We believe a player might reasonably restrict attention to choosing best-replies from this refined set. The criteria are as follows. A refined best-reply correspondence must have a product structure, as we want players to choose independently. It must be a point-wise subset of the best-reply correspondence, as we want players to only choose best responses. It must be non-empty-valued, thereby not allowing a player to not choose at all. It must be point-wise closed and convex. Closedness is more of a technical requirement, but convexity derives from the desire to have players randomize arbitrarily between their refined best-responses. Finally we require a refined best-reply correspondence to be upper-hemi continuous. This is an important technical requirement as it guarantees that such a refined best-reply correspondence has a fixed point and the differential inclusion based on it, i.e., the refined best-reply dynamics, always has a solution. In terms of player behavior it translates to the requirement that if one were to perturb the current strategy profile of the opponents a little bit this player will not choose a new strategy which was not formerly in the set of refined best-replies.

The existence of a refined best-reply correspondence immediately follows from the observation that the usual best-reply correspondence does satisfy all 5 criteria and hence is also a refined best-reply correspondence. A game could then in principle have many such refined best-reply

¹¹In many dynamics, e.g. (1), the global attractor in this game contains besides the Nash equilibria a continuum of orbits connecting the nonperfect equilibria to the perfect one, AD. This forms a two dimensional surface spanned by the hexagon. This filled-in hexagon has Euler characteristic 1.

correspondences, in fact it could even have (infinitely) many minimal such refined best-reply correspondences. Under mild conditions, which are satisfied, for instance, when no player has any two equivalent strategies, then this game has a unique minimal refined best-reply correspondence. For these games we then studied this minimal refined best-reply correspondence and the refined best-reply dynamics based on it.

While there are many results in the paper, characterizing fixed points of this correspondence as well as a notion of rationalizability based on it among other things, the main result, to our minds, is that the smallest asymptotically stable faces under any differential inclusion with reasonable properties are those which are asymptotically stable based on the minimal refined best-reply correspondence (Theorem 6) and these are Kalai and Samet's (1984) persistent retracts, which in turn are CURB sets (Basu and Weibull (1991)) based on the refined best-reply correspondence (Theorems 2 and 5). This reinforces Hurkens's (1995) finding that a stochastic best-reply learning process based on semi-robust best-replies also leads to play eventually leading to a persistent retract. Altogether this suggest that while it is difficult to justify Nash equilibrium behavior, either epistemically or through evolution or learning, yet alone any of its point-wise refinements or even set-wise refinements such as Kohlberg and Mertens's (1986) strategically stable sets, there are sets of strategy profiles which are justifiable through learning and the smallest such sets (if we restrict attention to faces) are Kalai and Samet's (1984) persistent retracts, which in turn are CURB sets (Basu and Weibull (1991)) based on the refined best-reply correspondence. Hence, we suggest that in applied game theory work CURB sets and persistent retracts, which are as of now not used to a great extent, may be very apt choices for a solution concept.

References

- AUBIN, J.-P., AND A. CELLINA (1984): *Differential Inclusions*. Springer.
- BALKENBORG, D. (1992): *The Properties of Persistent Retracts and Related Concepts*. Ph.D. thesis, University of Bonn.
- BALKENBORG, D., M. JANSEN, AND D. VERMEULEN (2001): "Invariance properties of persistent equilibria and related solution concepts," *Mathematical Social Sciences*, 41(1), 111–130.
- BASU, K., AND J. W. WEIBULL (1991): "Strategy subsets closed under rational behavior," *Economics Letters*, 36, 141–46.
- BENAIM, M., J. HOFBAUER, AND S. SORIN (2005): "Stochastic approximation and differential inclusions," *SIAM Journal on Control and Optimization*, 44, 328–348.
- BERNHEIM, B. D. (1984): "Rationalizable strategic behavior," *Econometrica*, 52, 1007–29.
- BJÖRNERSTEDT, J., AND J. W. WEIBULL (1996): "Nash equilibrium and evolution by imitation," in *The Rational Foundations of Economic Behavior. Proceedings of the IEA Conference held in Turin, Italy.*, ed. by K. J. A. et al., pp. 155–71. MacMillan Press Ltd, London.
- BOCHNAK, J., M. COSTE, AND M. ROY (1998): *Real Algebraic Geometry*. Springer.
- BÖRGER, T. (1994): "Weak dominance and approximate common knowledge," *Journal of Economic Theory*, 64, 265–276.

- BRANDENBURGER, A. (1992): “Economic Analysis of Markets and Games,” in *Economic Analysis of Markets and Games*, ed. by P. Dasgupta, D. Gale, O. Hart, and E. Maskin. MIT Press, Cambridge.
- BRANDENBURGER, A., AND A. FRIEDENBERG (2007): “Axioms for Backward Induction,” Unpublished manuscript.
- BRANDENBURGER, A., A. FRIEDENBERG, AND H. J. KEISLER (2008): “Admissibility in Games,” *Econometrica*, 76, 307–353.
- BROWN, G. W. (1951): “Iterative solutions of games by fictitious play,” in *Activity Analysis of Production and Allocation*, ed. by T. C. Koopmans, pp. 374–376. Wiley, New York.
- CRESSMAN, R. (2003): *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge, Mass.
- DEKEL, E., AND D. FUDENBERG (1990): “Rational behavior with payoff uncertainty,” *Journal of Economic Theory*, 52, 243–67.
- DEMICHIELIS, S., AND K. RITZBERGER (2003): “From evolutionary to strategic stability,” *Journal of Economic Theory*, 113, 5175.
- GILBOA, I., AND A. MATSUI (1991): “Social stability and equilibrium,” *Econometrica*, 59, 859–67.
- GUILLEMIN, V., AND A. POLLACK (1974): *Differential Topology*. Prentice Hall, Englewood Cliffs, N.J.
- GUNNING, R., AND H. ROSSI (1965): *Analytic Functions of Several Complex Variables*. Prentice Hall.
- HARSANYI, J., AND R. SELTEN (1988): *A general theory of equilibrium selection in games*. MIT Press, Cambridge, Massachusetts.
- HENDON, E., H. J. JACOBSON, AND B. SLOTH (1996): “Fictitious play in extensive form games,” *Games and Economic Behavior*, 15, 177–202.
- HERINGS, P. J.-J., AND V. J. VANNETELBOSCH (1999): “Refinements of rationalizability for normal-form games,” *International Journal of Game Theory*, 28, 53–68.
- HOFBAUER, J. (1995): “Stability for the best response dynamics,” Unpublished manuscript.
- HOFBAUER, J., AND K. SIGMUND (1998): *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.
- HOFBAUER, J., AND S. SORIN (2006): “Best response dynamics for continuous zero-sum games,” *Discrete and Continuous Dynamical Systems*, B 6, 215–224.
- HURKENS, S. (1995): “Learning by forgetful players,” *Games and Economic Behaviour*, 11, 304–29.

- JANSEN, M., P. JURG, AND D. VERMEULEN (2002): “On the set of equilibria of a bimatrix game: A survey,” in *Chapters in Game Theory*, ed. by P. Borm, and H. Peters, pp. 121–142. Kluwer, Amsterdam.
- KALAI, E., AND D. SAMET (1984): “Persistent equilibria in strategic games,” *International Journal of Game Theory*, 13, 129–44.
- KOHLBERG, E., AND J.-F. MERTENS (1986): “On the strategic stability of equilibria,” *Econometrica*, 54, 1003–37.
- KREPS, D. M., AND R. WILSON (1982): “Sequential equilibria,” *Econometrica*, 50, 863–94.
- MAS-COLLEL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*. Oxford University Press, Oxford, UK.
- MATSUI, A. (1992): “Best response dynamics and socially stable strategies,” *Journal of Economic Theory*, 57, 343–62.
- MCLENNAN, A. (1989): “Fixed points of contractible valued correspondences,” *International Journal of Game Theory*, 18, 175–184.
- MERTENS, J. (2004): “Ordinality in non-cooperative games,” *International Journal of Game Theory*, 32, 387–430.
- OKADA, A. (1983): “Robustness of equilibrium points in strategic games,” Working Paper B-137, Tokyo Center for Game Theory.
- PEARCE, D. G. (1984): “Rationalizable strategic behavior and the problem of perfection,” *Econometrica*, 52, 1029–51.
- RITZBERGER, K. (2002): *Foundations of Non-Cooperative Game Theory*. Oxford University Press.
- RITZBERGER, K., AND J. W. WEIBULL (1995): “Evolutionary selection in normal form games,” *Econometrica*, 63, 1371–1399.
- ROCKAFELLAR, R. T. (1970): *Convex Analysis*. Princeton University Press, Princeton, New Jersey.
- SAMUELSON, L. (1992): “Dominated strategies and common knowledge,” *Games and Economic Behavior*, 4, 284–313.
- SCHUHMACHER, F. (1999): “Proper rationalizability and backward induction,” *International Journal of Game Theory*, 28, 599–615.
- SELTEN, R. (1965): “Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit,” *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–324.
- (1975): “Re-examination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory*, 4, 25–55.

- VAN DAMME, E. E. C. (1991): *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin, Heidelberg.
- VON STENGEL, B., AND S. ZAMIR (2004): “Leadership with commitment to mixed strategies,” CDAM Research Report LSE-CDAM-2004-01; to appear in *Games and Economic Behavior*.
- VOORNEVELD, M. (2004): “Preparation,” *Games and Economic Behavior*, 48, 403–14.
- (2005): “Persistent retracts and preparation,” *Games and Economic Behavior*, 51, 228–32.
- YOUNG, H. P. (1993): “The evolution of conventions,” *Econometrica*, 61, 57–84.

A On the generic equivalence of own-payoff equivalence and payoff equivalence

The condition that own-payoff equivalence implies payoff equivalence for pure strategies is closely related to the *Single Payoff Condition (SPC)* of Brandenburger and Friedenberg (2007) for perfect information games and it seems apt to use the same name here. This appendix spells out a simple condition under which SPC holds generically in a given class of games. We then provide a series of examples to illustrate, using this result, that the restriction to games satisfying SPC and hence, after the identification of payoff equivalent strategies, to games in $\mathcal{G}^*$, is not a severe one.

Definition 2 *A normal form game satisfies the Single Payoff Condition (SPC) if the following holds for all players $i \in I$. Two strategies $s_i, s'_i \in S_i$ satisfy the equations*

$$u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$ only if also the equations

$$u_j(s_i, s_{-i}) = u_j(s'_i, s_{-i})$$

holds for all $j \in I$ and all $s_{-i} \in S_{-i}$.

Definition 3 *For a given set of strategy combinations $\mathcal{S}$ consider a family of normal form games $\{\Gamma^\mu\}_{\mu \in O}$ given by utility functions*

$$u_i(s, \mu) \tag{4}$$

for $s \in \mathcal{S}$ and $i \in I$, which depend on a vector of parameters μ taken from a nonempty open set O in some Euclidian space $\mathbb{R}^k$. We call the family analytic if all $u_i(s, \mu)$ are analytic functions in μ for given $s \in \mathcal{S}$.¹² We say that the family satisfies the Functional Single Payoff Condition if the following holds for all players $i \in I$. Two strategies $s_i, s'_i \in S_i$ satisfy the functional identities

$$u_i(s_i, s_{-i}, \mu) = u_i(s'_i, s_{-i}, \mu) \quad \text{for all } \mu \in O$$

for all $s_{-i} \in S_{-i}$ only if also the functional identities

$$u_j(s_i, s_{-i}, \mu) = u_j(s'_i, s_{-i}, \mu) \quad \text{for all } \mu \in O$$

hold for all $j \in I$ and all $s_{-i} \in S_{-i}$.

Proposition 21 *Suppose the analytic family of games $\{\Gamma^\mu\}_{\mu \in O}$ satisfies the Functional Single Payoff Condition. Then for generic $\mu \in O$ the game Γ^μ satisfies the Single Payoff Condition.*

Proof: Fix $i \in I$, $s_i, s'_i \in S_i$, $s_{-i} \in S_{-i}$ and $j \in I$ such that $u_j(s_i, s_{-i}, \mu)$ and $u_j(s'_i, s_{-i}, \mu)$ are distinct as functions in μ . Then the set of parameter values μ for which

$$u_j(s_i, s_{-i}, \mu) = u_j(s'_i, s_{-i}, \mu) \tag{5}$$

¹²An analytic function is a function that is locally described by power series. The notion covers most functions arising in applications, in particular linear and rational functions or functions like e^x or $\ln(x)$. In our examples the functions are always linear.

is a closed lower dimensional analytic set because the function is analytic (see e.g., Gunning and Rossi (1965)). Because there are finitely many choices of $i \in I$, $s_i, s'_i \in S_i$, $s_{-i} \in S_{-i}$ and $j \in I$ to consider we find that for μ outside a lower dimensional analytic subset D of O the identity (5) for *some* $i \in I$, $s_i, s'_i \in S_i$, $s_{-i} \in S_{-i}$ and $j \in I$ implies the identity (5) for *all* $i \in I$, $s_i, s'_i \in S_i$, $s_{-i} \in S_{-i}$ and $j \in I$ and also $\mu \in O$. In particular, the SPC condition holds for all $\mu \notin D$. QED

Example 1 In a **cheap talk game** players first send simultaneously and independently public messages m_i from message spaces M_i . After all players have received the combination of messages

$$m = (m_1, \dots, m_n) \in M = \times_{i \in I} M_i \quad (6)$$

they choose simultaneously and independently actions $a_i \in A_i$. A pure strategy in such a game consists of a message m_i and a function $f_i : M \rightarrow A_i$. The play of any strategy combination s will result in a combination of messages $m \in M$ and a combination of actions

$$a = (a_1, \dots, a_n) \in A = \times_{i \in I} A_i, \quad (7)$$

where, in a cheap talk game, only the latter is payoff relevant. In this example the parameter space is $\mathbb{R}^{A \times I}$. For $\mu \in \mathbb{R}^{A \times I}$ we define the utility function by

$$u_i(s, \mu) = \mu_{a,i} \quad (8)$$

where a is the combination of actions induced by s . The utility function is then for each $s \in S$ the projection onto a particular component of the vector μ . The identity

$$u_i(s_i, s_{-i}, \mu) = u_i(s'_i, s_{-i}, \mu) \quad (9)$$

can only hold for *all* μ if both functions project onto the same component of μ , i.e., if the play of both (s_i, s_{-i}) and (s'_i, s_{-i}) results in the same combination of actions a , although in possibly different combinations of messages. (If (s_i, s_{-i}) and (s'_i, s_{-i}) would result in different combinations of actions a and a' the equality would not hold in the game where all players get 1 after a and 0 after a' .) If this is the case then, by construction,

$$u_j(s_i, s_{-i}, \mu) = \mu_{a,j} = u_j(s'_i, s_{-i}, \mu) \quad (10)$$

for all j and μ . Thus Proposition 21 applies and we conclude that the SPC holds generically in cheap talk games.

Example 2 In an **extensive game without chance moves** the play of any pure strategy combination results in a terminal node $t \in T$. In this case the parameter space for a given extensive form is $\mathbb{R}^{T \times I}$ and the utility function is $u_i(s, \mu) = \mu_{t,i}$ if s induces t . The same arguments as for cheap talk games imply that the SPC holds in generic extensive form games with no random moves. Notice, though, that almost no extensive game with the extensive form of a cheap talk game is itself a cheap talk game. Hence the previous result is not a special case of this one.

Example 3 In an extensive game with chance moves the parameter space remains as in the previous example, but the utility function becomes

$$u_i(s, \mu) = \sum_{t \in T} p_t \mu_{t,i} \quad (11)$$

where p_t is the probability with which terminal node t is reached when the pure strategy combination s is played. Clearly, the equation

$$u_i(s_i, s_{-i}, \mu) = \sum_{t \in T} p_t \mu_{t,i} = \sum_{t \in T} p'_t \mu_{t,i} = u_i(s'_i, s_{-i}, \mu) \quad (12)$$

can only hold for *all* $\mu \in R^{T \times I}$ if $p_t = p'_t$ for all $t \in T$. Thus the SPC holds for generic extensive form games even with chance moves.

Example 4 In a finitely repeated game with perfect monitoring, no discounting and $t \geq 0$ periods, the play of a pure strategy combination s results in a sequence $(a_1, a_2, \dots, a_t)$ of combinations of actions in the stage game. The payoff to a player can be written as

$$\sum_{a \in A} k_{s,a} \mu_{a,i} \quad (13)$$

where the parameter $\mu_{a,i}$ is player i 's payoff in the stage game from the combination of actions a and $k_{s,a}$ is the number of times a is played in the sequence $(a_1, a_2, \dots, a_t)$. If for two strategy combinations $s = (s_i, s_{-i})$ and $s' = (s'_i, s_{-i})$

$$u_i(s, \mu) = \sum_{a \in A} k_{s,a} \mu_{a,i} = \sum_{a \in A} k_{s',a} \mu_{a,i} = u_i(s', \mu) \quad (14)$$

holds for all $\mu \in R^{A \times I}$ then $k_{s,a} = k_{s',a}$ for all $a \in A$ and, hence,

$$u_i(s, \mu) = \sum_{a \in A} k_{s,a} \mu_{a,i} = \sum_{a \in A} k_{s',a} \mu_{a,i} = u_i(s', \mu) \quad (15)$$

Again, the SPC holds generically in repeated games.

Example 5 Consider finally the class of normal form games which satisfy for every $i \in I$, every $s_{-i} \in S_{-i}$ and any $s_i, s'_i \in S_i$ the equation

$$u_i(s_i, s_{-i}, \mu) = u_i(s'_i, s_{-i}, \mu) \quad (16)$$

If at least one player has two strategies, then this class does not satisfy the Functional Single Payoff Condition. Almost all games in this class violate the SPC.

B On the generic equivalence of best responses and semi-robust best responses

This section provides a proof of Proposition 3. Note first that generic normal form games are in the class $\mathcal{G}^*$ where all semi-robust responses are pure strategies. The proof of Proposition

3 is organized in a number of steps: We will first fix some notations for the mappings and various submanifolds to be considered. Step 1 argues that the embedding of the uncorrelated strategy combinations into the set of beliefs has nice differentiability properties. Step 2 invokes the transversality theorem (see Guillemin and Pollack (1974)) to show that for generic payoff functions we obtain the needed transversality conditions.¹³ Step 3 argues that we can restrict attention to completely mixed strategy combinations of the opponents. If the player is indifferent between several of his strategies against a given completely mixed strategy combination, step 4 shows how we can construct an arbitrarily nearby strategy combination, against which the player strictly prefers a given one among these strategies. Step 5 completes the argument.

For any finite set M let $\mathbb{R}^M$ be the vector space of all mappings $y : M \rightarrow \mathbb{R}$. The dimension of $\mathbb{R}^M$ is the number of elements in M .

Let $q_i : \prod_{j \neq i} \mathbb{R}^{S_j} \rightarrow \mathbb{R}^{S_{-i}}$ be the mapping defined by

$$(q_i(x_{-i}))(s_{-i}) := \prod_{j \neq i} x_j(s_j) .$$

q_i describes the first step in the computation mentioned above.

While x_{-i} denotes the usual strategy combination of the opponents, we define χ_{-i} to describe a “correlated strategy of the opponents”, i.e. a belief over S_{-i} .

Let L_i be the vector space of all linear mappings

$$v_i : \mathbb{R}^{S_{-i}} \rightarrow \mathbb{R}^{S_i} .$$

If $\chi_{-i} \in \mathbb{R}^{S_{-i}}$ is a belief and $s_i \in S_i$ a pure strategy of player i , then $(v_i(\chi_{-i}))(s_i)$ is the expected payoff for player i . v_i describes for every s_i the second step in the computation of the expected payoff. Any $v_i \in L_i$ corresponds uniquely to a payoff function

$$u_i : S \rightarrow \mathbb{R}$$

in the standard notation (and this relation is a homeomorphism).

For $T_i \subseteq S_i$ set $Z_i(T_i) = \{z \in \mathbb{R}^{S_i} \mid \forall s_i, t_i \in T_i : z(s_i) = z(t_i)\}$. Let $X_j := \{x_j \in \mathbb{R}^{S_j} \mid \sum_{s_j \in S_j} x_j(s_j) = 1\}$ for $j \neq i$ and $X_{-i} := \prod_{j \neq i} X_j$.

For $T_j \subseteq S_j$ ($j \neq i$) and $T_{-i} := \prod_{j \neq i} T_j$ set

$$X_j(T_j) := \{x_j \in X_j \mid \forall s_j \in T_j : x_j(s_j) = 0\}$$

and

$$X_{-i}(T_{-i}) := \prod_{j \neq i} X_j(T_j) .$$

The sets $\Theta_{-i} \cap X_{-i}(T_{-i})$ describe the various faces of the polyhedron Θ_{-i} .

Step 1: For all T_{-i} the mapping $q_i : X_{-i}(T_{-i}) \rightarrow \mathbb{R}^{S_{-i}} \setminus \{0\}$ is a diffeomorphism onto its image (in particular $q_i(X_{-i}(T_{-i}))$ is a closed submanifold of $\mathbb{R}^{S_{-i}} \setminus \{0\}$).

¹³This transversality theorem is a straightforward consequence of Sard’s theorem. If one assumes an algebraic map and uses in its proof in Guillemin and Pollack (1974) the algebraic version of Sard’s theorem in Bochnak, Coste, and Roy (1998) one obtains a stronger version of the transversality theorem where the conclusion of the theorem holds outside a lower dimensional semi-algebraic set.

Proof: $X_{-i}(T_{-i})$ is a closed affine submanifold in $\prod_{j \neq i} (\mathbb{R}^{S_j} \setminus \{0\})$. It is straightforward to check that

$$q_i|_{X_{-i}(T_{-i})} : X_{-i}(T_{-i}) \rightarrow \mathbb{R}^{S_{-i}} \setminus \{0\}$$

is well defined, is one-to-one, maps $X_{-i}(T_{-i})$ to a closed set, and has a derivative $dq_i|_{x_{-i}}$ with maximal rank everywhere.¹⁴

Step 2: Let $Z \subseteq \mathbb{R}^{S_i}$ and $X \subseteq \mathbb{R}^{S_{-i}} \setminus \{0\}$ be submanifolds. Then for almost every $v_i \in L_i$ the mapping $v_i|_X : X \rightarrow \mathbb{R}^{S_{-i}} \setminus \{0\}$ is transversal to Z .

Proof: The family of linear maps L_i defines a mapping

$$V_i : L_i \times \mathbb{R}^{S_{-i}} \rightarrow \mathbb{R}^{S_i} \quad (17)$$

$$(v_i, \chi_{-i}) \mapsto v_i(\chi_{-i}). \quad (18)$$

The derivative of V_i at (v_i, χ_{-i}) can be computed as

$$dV_i|_{(v_i, \chi_{-i})} : T_{v_i}L_i \times T_{\chi_{-i}}\mathbb{R}^{S_{-i}} \cong L_i \times \mathbb{R}^{S_{-i}} \rightarrow \mathbb{R}^{S_i} \quad (19)$$

$$(\nu_i, \xi_{-i}) \mapsto \nu_i(\chi_{-i}) + v_i(\xi_{-i}). \quad (20)$$

If $\chi_{-i} \neq 0$ we can find for every $\zeta_i \in \mathbb{R}^{S_i}$ some $\nu_i \in L_i$ with $\nu_i(\chi_{-i}) = \zeta_i$. Then $dV_i|_{(v_i, \chi_{-i})}(\nu_i, 0) = \zeta_i$.

Because for $\chi_{-i} \in X$ the tangent space $T_{\chi_{-i}}X \subseteq \mathbb{R}^{S_{-i}}$ contains the 0-vector, $dV_i|_{(v_i, \chi_{-i})} : T_{v_i}L_i \times T_{\chi_{-i}}X \rightarrow \mathbb{R}^{S_i}$ is surjective. Thus $V_i : L_i \times X \rightarrow \mathbb{R}^{S_i}$ is transversal to Z and our claim follows from the transversality theorem.

By step 1 and step 2 almost every $v_i \in L_i$ satisfies:

⊗ For all subsets $T_i \subseteq S_i$ ($1 \leq i \leq n$) the mapping $(v_i \circ q_i)|_{X_{-i}(T_{-i})}$ is transversal to $Z_i(T_i)$.

For given v_i define $Y(T_i) = \{x_{-i} \in X_{-i} \mid (v_i \circ q_i)(x_{-i}) \in Z(T_i)\}$. $Y(T_i) \cap \Theta_{-i}$ is the set of strategy combinations of the opponents such that player i is indifferent between the strategies in T_i (i.e. they give the same payoff).

Step 3: Suppose v_i satisfies ⊗. For $T_i \subseteq S_i$ let $x_{-i} \in Y(T_i) \cap \Theta_{-i}$ and let O_{-i} be a neighborhood of x_{-i} . Then $O_{-i} \cap Y(T_i)$ contains a point in the interior of Θ_{-i} .

Proof: Suppose x_{-i} is in the boundary of Θ_{-i} . For each $j \neq i$ define $T_j := \{s_j \in S_j \mid x_j(s_j) = 0\}$. If $T_j = \emptyset$, x_j is in the relative interior of Θ_j . By assumption T_j is not empty for at least one opponent j . Fix $j^* \neq i$ with $T_{j^*} \neq \emptyset$ and $t_{j^*} \in T_{j^*}$. Set $\tilde{T}_j := T_j$ for $i \neq j \neq j^*$ and $\tilde{T}_{j^*} := T_{j^*} \setminus \{t_{j^*}\}$. We show that $O_{-i} \cap Y(T_i)$ contains some $y_{-i} \in \Theta_{-i} \cap X(\tilde{T}_{-i})$ such that $\tilde{T}_j = \{s_j \in S_j \mid y_j(s_j) = 0\}$ for all $j \neq i$. In other words: y_{-i} is in the relative interior of the face $\Theta_{-i} \cap X(\tilde{T}_{-i})$. The claim then follows by induction.

¹⁴The result is well known e.g. in algebraic geometry: q_i defines the so-called Segre-embedding. The result is needed in algebraic geometry to show that the product of projective spaces can itself be embedded into a projective space, i.e. is projective-algebraic.

The transversality conditions imply that the submanifolds $X_{-i}(T_{-i})$ and $Y(T_i) \cap X_{-i}(\tilde{T}_{-i})$ meet transversely in $X_{-i}(\tilde{T}_{-i})$. (See Guillemin and Pollak [1974, exercise 2.3.7.]) Since $X_{-i}(T_{-i})$ has codimension 1 in $X_{-i}(\tilde{T}_{-i})$ it follows that $X_{-i}(\tilde{T}_{-i}) \cap Y(T_i) \cap \{y_{-i} \mid y_{j*}(t_{j*}) > 0\} \cap O_{-i}$ intersects the relative interior of $X_{-i}(\tilde{T}_{-i}) \cap \Theta_{-i}$.

Step 4: Suppose v_i satisfies $\otimes$. For $T_i \subseteq S_i$ with $\#T_i \geq 2$ let $x_{-i} \in Y(T_i)$ be in the interior of Θ_{-i} and let O_{-i} be a neighborhood of x_{-i} . Then we can find for every $s_i \in T_i$ some $y_{-i} \in O_{-i} \cap \Theta_{-i}$ such that

$$(v_i \circ q_i)(y_{-i})(s_i) > (v_i \circ q_i)(y_{-i})(t_i) \quad \text{for all } t_i \in T_i \setminus \{s_i\}.$$

Proof: Because $v_i \circ q_i : X_{-i} \rightarrow \mathbb{R}^{S_i}$ is transversal to both $Z(T_i)$ and $Z(T_i \setminus \{s_i\})$ it follows that $v_i \circ q_i : Y(T_i \setminus \{s_i\}) \rightarrow Z(T_i \setminus \{s_i\})$ is transversal to $Z(T_i)$. From this we can deduce the existence of a tangent vector $\xi \in T_{x_{-i}}(Y(T_i \setminus \{s_i\}))$ with $d\lambda|_{x_{-i}}(\xi) = 1$, where λ is the function

$$\lambda : Y_i(T_i \setminus \{s_i\}) \cap X_{-i} \rightarrow \mathbb{R} \tag{21}$$

$$y_{-i} \rightarrow (v_i \circ q_i)(y_{-i})(s_i) - (v_i \circ q_i)(y_{-i})(t_i) \tag{22}$$

defined for arbitrary but fixed $t_i \in T_i \setminus \{s_i\}$. We can therefore select a differentiable curve

$$c : (-\epsilon, \epsilon) \rightarrow Y_i(T_i \setminus \{s_i\})$$

with $c(0) = x_{-i}$ and $(\lambda \circ c)'(0) = 1$. For sufficiently small $0 < \gamma < \epsilon$ $y_{-i} := c(\gamma)$ has the required properties.

Step 5: Suppose s_i is a pure best response against x_{-i} . For every neighborhood O_{-i} of x_{-i} the continuity of the payoff function and the two steps above can be used to find $y_{-i} \in O_{-i}$ such that $s_i \in T_i$ is the unique best response against y_{-i} . Shrinking the open sets we can find a sequence of such y_{-i} 's converging to x_{-i} . Continuity yields an open set around each element in the sequence, where s_i is the unique best response. s_i is the unique best response on the union of these sets, which is again open. Thus s_i is a semi-robust best response against x_{-i} . *QED*

C Refined best replies in two-player games

We will restrict attention to the best replies of player 1. Suppose player 2 has $K \geq 2$ strategies $s_2^1, \dots, s_2^K$. It will be convenient to identify the mixed strategies $x_2 \in \Theta_2$ with the vectors

$$x_2 = (x_2^1, x_2^2, \dots, x_2^{K-1}) \in \mathbb{R}^{K-1} \tag{23}$$

for which $x_2^k \geq 0$ for all $1 \leq k \leq K-1$ and $x_2^K := 1 - \sum_{k=1}^{K-1} x_2^k \leq 0$. Notice that the zero vector corresponds to pure strategy s_2^K .

We define the function $f : \mathbb{R}^{K-1} \rightarrow \mathbb{R}$ by

$$f(x_2) = \begin{cases} \max_{s_1 \in S_1} u_1(s_1, x_2) & \text{for } x_2 \in \Theta_2 \\ +\infty & \text{else} \end{cases} \tag{24}$$

Because u_1 is linear in x_2 , f is, in the terminology of Rockafellar (1970) a proper convex polyhedral function. Each strategy $x_1 \in \Theta_1$ defines an affine function $a : \mathbb{R}^{K-1} \rightarrow \mathbb{R}$ by $a(x_2) = u_1(x_1, x_2)$, which, for all $x_2 \in \Theta_2$, satisfies the inequality $a(x_2) \leq f(x_2)$ and $a(x_2) = f(x_2)$ holds if and only if $x_1 \in \beta_1(x_2)$.

For a strategy $x_1 \in \Theta_1$ we define the set

$$G(x_1) = \{(x_2, \alpha) \in \Theta_2 \times \mathbb{R} \mid x_1 \in \beta(x_2) \text{ and } \alpha = u_1(x_1, x_2)\} \quad (25)$$

and the set $H(x_1) = \{x_2 \in \Theta_2 \mid x_1 \in \beta(x_2)\}$, the projection of $G(x_1)$ onto Θ_2 . $H(x_1)$ is the region where x_1 is a best reply. x_1 is *robust* if $H(x_1)$ has non-empty interior $H^\circ(x_1)$. We can now describe the epigraph

$$F = \{(x_2, \alpha) \in \mathbb{R}^{K-1} \times \mathbb{R} \mid f(x_2) \leq \alpha\} \quad (26)$$

of f as follows: F is a polyhedral convex set whose compact faces are precisely the sets $G(x_1)$ with $x_1 \in \Theta_1$. The non-compact faces are of the form $F \cap (\Theta'_1 \times \mathbb{R})$, where Θ'_1 is a face of Θ_1 .

Because $H(x_1)$ is a convex polyhedron the closure of $H^\circ(x_1)$ is $H(x_1)$ if $H^\circ(x_1)$ is non-empty. Using lemma 2 this implies immediately Proposition 4.

To prove Proposition 5 we study the conjugate function $f^* : \mathbb{R}^{K-1} \rightarrow \mathbb{R}$ of f defined by

$$f^*(x_2^*) = \sup_{x_2 \in \mathbb{R}^{K-1}} \{x_2^* \bullet x_2 - f(x_2)\} = \max_{x_2 \in \Theta_2} \{x_2^* \bullet x_2 - f(x_2)\} < \infty, \quad (27)$$

where $x_2^* \bullet x_2$ denotes the usual scalar product $\sum_{k=1}^{K-1} x_2^{*k} x_2^k$. As shown for any convex polyhedral function in Rockafellar (1970), the conjugate is again a convex polyhedral function and one has $f^{**}(x_2) = f(x_2)$.

Any two strategies $x_1, x'_1 \in \Theta_1$ define the same affine function if and only if the two strategies are equivalent. Without loss of generality we can thus identify Θ_1 with a subset of the affine functions on $\mathbb{R}^{K-1}$.

Any vector (x_2^*, α) with $x_2^* \in \mathbb{R}^{K-1}$ and $\alpha \in \mathbb{R}$ defines one and only one affine function on $\mathbb{R}^{K-1}$ by

$$a(x_2) = -\alpha + \sum_{k=1}^{K-1} x_2^{*k} x_2^k \quad (28)$$

We will identify affine functions with such vectors. For instance, $e = (1, \dots, 1)$ corresponds to the function $-x_2^{K+} = -1 + \sum_{k=1}^{K-1} x_2^k$ which assigns 0 to the first $K-1$ pure strategies and -1 to the last pure strategy of player 2.

Let F^* be the epigraph of f^* .

Lemma 9 *F^* is a polyhedral convex set generated by extreme points x_1 which are robust best replies in Θ_1 and the directions*

$$-e_k = (-e_k^1, \dots, -e_k^K) \in \mathbb{R}^K \text{ with } e_k^l = \begin{cases} -1 & \text{for } k = l \\ 0 & \text{else} \end{cases} \quad (29)$$

for $k = 1, \dots, K-1$ and

$$e = (1, \dots, 1) \in \mathbb{R}^K \quad (30)$$

Proof: By definition $(x_2^*, \alpha^*) \in F^*$ if and only if $\alpha^* \geq x_2^* \bullet x_2 - f^*(x_2)$ for all $x_2 \in \Theta_2$. $v \in \mathbb{R}^K$ is a direction in F^* if and only if there exists $(x_2^*, \alpha^*) \in F^*$ such that all vectors $(x_2^*, \alpha^*) + \lambda v$ with $\lambda \geq 0$ are in F^* . We can write $v = -\sum_{k=1}^{K-1} \rho_k e_k + \rho_K e$ with $\rho_1, \dots, \rho_K \in \mathbb{R}$ since $-e_1, \dots, -e_{K-1}, e$ form a vector basis of $\mathbb{R}^K$. We must show that v is a direction in F^* if and only if all ρ_i are non-negative. Suppose that v is a direction in F^* . Let $x_2 = (0, \dots, 0) \in \Theta_2$. The condition that $(x_2^*, \alpha^*) + \lambda v \in F^*$ for all $\lambda \geq 0$ yields for this x_2 that $\alpha^* + \lambda \rho_K \geq -f(x_2)$ holds for all $\lambda \geq 0$. This can be true only if $\rho_K \geq 0$. For $e_k \in \Theta_2$ ($1 \leq k \leq K-1$) we obtain similarly $\alpha^* + \lambda \rho_K \geq x_2^{*k} - \lambda \rho_k + \lambda \rho_K - f(e_k)$ for all $\lambda \geq 0$, which can hold only if $\rho_k \geq 0$. Thus only positive combinations of $-e_1, \dots, -e_{K-1}, e$ can be directions in F^* . For every combination $v = -\sum_{k=1}^{K-1} \rho_k e_k + \rho_K e$ with $\rho_1, \dots, \rho_K \geq 0$, every $\lambda \geq 0$, every $(x_2^*, \alpha^*) \in F^*$ and every $x_2 \in \Theta_2$ we have conversely

$$\alpha^* + \lambda \rho_K \geq x_2^{*'} x_2 - \sum_{k=1}^{K-1} \lambda \rho_k x_2^k + \lambda \rho_K - f(x_2) \quad (31)$$

which proves that v is a direction in F^* .

We have characterized the directions of F^* and must now determine the extremal points of F^* . Suppose $(\hat{x}_2^*, \hat{\alpha}^*)$ is an extremal point. Because F^* has only finitely many extremal points, these are exposed points by Straszewick's theorem (Theorem 18.6 in Rockafellar (1970)). Therefore we can find $x_2 \in \Theta_2$ such that the hyperplane $\{x_2^* \bullet x_2 = f(x_2)\}$ is a supporting hyperplane which meets F^* only in $(\hat{x}_2^*, f^*(\hat{x}_2^*))$. Because F^* has only finitely many extreme points and directions there exists an open neighborhood U of x_2 in Θ_2 for which the hyperplanes $\{x_2^* \bullet y_2 = f(y_2)\}$ are for all $y_2 \in U$ supporting hyperplanes which intersect F^* only in $(\hat{x}_2^*, f^*(\hat{x}_2^*))$. This implies that the graph of the affine function $(\hat{x}_2^*, f^*(\hat{x}_2^*))$ intersects F in a $K-1$ dimensional face. It is therefore identical to a affine function defined by a strategy x_1 in Θ_1 for which $H(x_1)$ is full dimensional. Given our identification, $(\hat{x}_2^*, f^*(\hat{x}_2^*))$ is consequently a robust strategy in Θ_1 , which was to be shown. QED

The lemma implies that all extreme points and hence all the points in the compact faces of F^* are in Θ_1 .

However, no points on the compact faces of F^* apart from the extreme points are robust strategies. To see this, notice that a proper mixture $x_1 = \sum_{l=1}^L \rho_l x_{1l}$ ($L > 2, \rho_l > 0, \sum_{l=1}^L \rho_l = 1$) of non-equivalent robust strategies in Θ_1 is not robust. Otherwise there would be an open set in Θ_2 on which x_1 and hence all strategies x_{1l} were best replies. They would yield identical payoffs on an open set and were hence (by Lemma 1) all equivalent, contradicting the assumption.

Finally, consider a strategy in Θ_1 which is not on a compact face of F^* . It can be written as $x_1' = x_1 - \sum_{k=1}^K \rho_k e_k + \rho_K e$ where x_1 is on one of the compact faces of F^* and, hence, in Θ_1 , and the ρ_k are all non-negative and at least some of them are strictly positive. We obtain

$$u_1(x_1', x_2) = u_1(x_1, x_2) - \sum_{k=1}^{K-1} \rho_k x_2^k - \rho_K \left(1 - \sum_{k=1}^{K-1} x_2^k\right) \leq u_1(x_1, x_2), \quad (32)$$

where this inequality holds as a strict one for the k -th pure strategy of player 2 whenever $\rho_k > 0$. Thus x_1' is weakly dominated. It is not a robust strategy because it is a best reply only on a proper face of Θ_1 (see Pearce (1984)).

In summary, the only robust strategies in Θ_1 are the extreme point of F^* . All other strategies are proper mixtures of non-equivalent robust strategies or are weakly dominated and therefore not robust. In particular, we have proved Proposition 5.