Refined best-response correspondence and dynamics

Dieter Balkenborg[*], Josef Hofbauer[†], and Christoph Kuzmics[‡]

February 18, 2007

Abstract

We study a natural (and, in a well-defined sense, minimal) refinement of the best-reply correspondence. Its fixed points, notions of rationalizability and CURB sets based on this correspondence are defined and characterized. Finally minimally asymptotically stable faces under the induced refined best-reply dynamics are characterized.

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 College London, j.hofbauer@ucl.ac.uk

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

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 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 normal-form games. But does game theory really even predict 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 (see e.g. Footnote 1 in Weibull (1995)). These models, however, principally failed to provide 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 (1996). 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 (1991)'s CURB sets. Using results of Balkenborg (1992) Hurkens (1995) furthermore shows that a specially refined stochastic best-reply dynamic leads to play eventually being within Kalai and Samet (1984)'s persistent retracts. Ritzberger and Weibull (1996) show that under any general payoff-positive dynamics minimally 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? We restrict attention to best-reply dynamics as opposed to better-reply dynamics. This is, of course, a reasonably strong assumption about the rationality of individuals. We then go further, however, 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 (1996). 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 evolutionary stable. The main result of this paper is that these smallest evolutionary stable faces coincide with Kalai and Samet (1984)'s persistent retracts, which again coincide with Basu and Weibull (1991)'s CURB sets adapted for the minimal refined best-reply correspondence. This result is analogous to Hurkens (1995)'s result that persistent retracts are the only candidates for stochastically evolutionary stable states in a particular stochastic model of best-reply learning a 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 striking result, for instance, is that a pure fixed point of this minimal refined best-reply correspondence induces a perfect Bayesian equilibrium in any extensive form with the given normal form as its induced normal form. This is somewhat reminiscent of the statement that a proper equilibrium (Myerson (1978)) induces a sequential equilibrium in any extensive form with the given normal form.

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. In this section we also characterize the, in the given class of games, unique minimal such refinement and to an extent characterize its fixed-points. In section 4 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 (1990)'s $S^{\infty}W^{1}$ elimination procedure. In section 5 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 (1984)' notion of an absorbing retract. In section 6 we discuss implications of the results of the previous sections for extensive form games. Section 7 provides a micro-story, similar in spirit to Björnerstedt and Weibull (1996), leading a deterministic differential inclusion based on the refined best-reply dynamics, before we finally prove the main result in section 8.

2 Preliminaries

Let $\Gamma = (I, S, u)$ be a finite n-player normal form game, where $I = \{1, ..., n\}$ is the set of players, $S = \times_{i \in I} S_i$ is the set of pure strategy profiles, and $u : S \to \mathbb{R}^n$ the payoff function[1]. 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 $\text{int}(\Theta)$ denote the relative interior of Θ, i.e. $\text{int}(\Theta) = \{x \in \Theta : \; x_{is} > 0 \; \forall s \in S_i \; \forall i \in I\}$, i.e. 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. Player i's (mixed) strategy $x_i \in \Theta$ then represents the vector of proportions of people playing the various pure strategies available to players in player population i, i.e. x_{is} denotes the proportion of players in population i who play pure strategy $s \in S_i$.

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 $\beta(x) = \times_{i \in I} \beta_i(x)$.

The following definition can be found in Kalai and Samet (1984). Two strategies $x_i, y_i \in \Theta_i$ are **equivalent** (for player i) if $u_i(x_i, x_{-i}) = u_i(y_i, x_{-i})$

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

for all $x_{-i} \in \Theta_{-i} = \times_{j \neq i} \Theta_j$.

Let $\Psi = \{x \in \Theta | \ \mathcal{B}(x) \text{ is a singleton}\}$. 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 in Table 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 $\Psi = \emptyset$, which is obviously not dense in Θ. This is to do with the fact that player 1 has two equivalent strategies.

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

Table 1: A game in which Ψ is not dense in Θ.

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

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

Theorem 1 *Let Γ be without equivalent 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 $s_i, t_i \in S_i$ 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 equivalent for player i. QED

Note that the opposite of Theorem 1 is not true. Consider two 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 theorem is immediate.

Theorem 2 *Let Γ be such that Ψ is dense in Θ. Let $s_i \in S_i$ be a best-reply on an open subset of Θ. Then player i has no equivalent strategy to s_i in S_i.*

Note that the restriction that a game should have no equivalent strategies is not a severe one. In particular we are not ruling out games with weakly dominated strategies.

Definition 1 *A strategy* $s_i \in S_i$ *is a strict never best reply if for every* $x \in \Theta$ *there is a* $t_i \in S_i$ *such that* $u_i(s_i, x_{-i}) < u_i(t_i, x_{-i})$*.*

In other words a strict never best reply s_i is such that $s_i \notin \mathcal{B}_i(x)$ for any $x \in \Theta$.

Definition 2 *A strategy* $w_i \in S_i$ *is a weak never best reply if for every* $x \in \Theta$ *there is a* $t_i \in S_i$*,* $t_i \neq w_i$ *such that* $u_i(w_i, x_{-i}) \leq u_i(t_i, x_{-i})$*.*

In other words a weak never best reply 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 a weak never best reply. Games not in $\mathcal{G}^*$, however, may be such that all strategies are weak never best replies[2]. A weak never best reply is, in fact, a strategy which is weakly dominated by a set of strategies, as defined in Balkenborg (1992).

Of course, if a strategy is strictly dominated then it is a strict never best-reply. If a strategy is weakly dominated then it is a weak never best-reply. The reverse is not true (see Example 5.7 in Ritzberger (2002) for a strategy which is a strict never best reply but not strictly dominated).

3 Refined best-reply correspondences

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

1. $\tau(x) = \times_{i \in I} \tau_i(x) \ \forall \ x \in \Theta$,
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-valued for all $x \in \Theta$,
5. $\tau(x)$ is upper-hemi continuous at all $x \in \Theta$.

Note 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 5[3]. For such x any $\tau(x)$ must then also

[2] Consider, for instance, the game in which every player has at least two strategies, yet receives a payoff of 1, regardless of what other players choose.

[3] Strategies that are unique best replies to some $x \in \Psi$ were called *inducible* in von Stengel and Zamir (2004). We might call the pure strategy profiles in $\sigma(x)$ the *inducible* or *indispensable* best replies to x.

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 $\sigma : \Theta \rightrightarrows \Theta$, can be found in the following way. For $x \in \Theta$ let

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

The set $\mathcal{S}_i(x)$ is in fact the set of pure semi-robust best replies as defined by Balkenborg (1992). Let $\mathcal{S}(x) = \times_{i \in I} \mathcal{S}_i(x)$. From the observations above we then obtain the following theorem.

Theorem 3 *Let $\Gamma \in \mathcal{G}^*$. The unique minimal refined best-reply correspondence is given by σ, defined such that for any $x \in \Theta$, $\sigma(x) = \Theta[\mathcal{S}(x)] = \times_{i \in I} \Delta(\mathcal{S}_i(x))$.*

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

The next lemma is immediate.

Lemma 2 *Let $w_i \in S_i$ be a weak never best reply for player i. Then $w_i \notin \mathcal{S}_i(x)$ for any $x \in \Theta$.*

Proof: By the definition of a weak never best reply $w_i \notin \beta_i(x)$ for any $x \in \Psi$, but only strategies in $\beta_i(y)$ for some $y \in \Psi$ can be in $\sigma_i(x)$. QED

This, in turn, leads to an immediate theorem.

Theorem 4 *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 weak never best reply $w_i \in S_i$.*

Proof: Immediate from Lemma 2: Let $x \in \sigma(x)$. By Lemma 2 $w_i \notin \mathcal{S}_i(x)$ for any weak never best reply $w_i \in S_i$. 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 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 3 *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 $x_t \in int(\Theta)$) 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 the Game given in Table 2, 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. By Theorem 3.2.2 in van Damme (1991) x^*, being weakly dominated, cannot be 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

Table 2: A game in which a fixed point of σ is not perfect.

Theorem 5 *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 Theorem 4. But then in two-player games every undominated Nash equilibrium is (trembling-hand) perfect (see Theorem 3.2.2 in van Damme (1991)).

The reverse of Theorem 5 is not true. Consider the game given in Table 3. 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 Theorem 2 this game is in $\mathcal{G}^*$. Now, 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 perfect equilibrium. In fact they are also all KM-stable. 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 in fact a weak never best-reply.

Theorem 5 cannot be generalized to general n-player games. To see this consider the following 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 4 *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 $s_i \in B_i(y)$ for all $y \in U^{s_i}$.*

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

Table 3: A game in which a perfect equilibrium (and, in fact, KM-stable equilibrium) is not a fixed point of σ.

Proof: Immediate.

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 does not necessarily have to be the case. We already saw this for the 2-player game given in Table 2. 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 $U^F = \{x \in \Theta | x_{1B} > \frac{1}{2}\}$. These two sets are such that there is no bigger open set with the same property and they have an empty intersection. Hence, x^* is not perfect. The extensive form game in Figure 4 demonstrates that for games with more than 2 players this phenomenon may even occur for pure fixed points of σ.

In section 5 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. In section 8 we show that these persistent retracts are asymptotically stable under our refined best-reply dynamic. So 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 the game given in Table 4 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.

The game given in Table 5, 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 combina-

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

E

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

F

Table 4: A game in which a pure fixed point of σ is not persistent.

tion (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.

4 $\boldsymbol{\sigma}$-Rationalizability

A set $R \subset S$ is a **selection** if $R = \times_{i \in I} R_i$ and $R_i \subset S_i$, $R_i \neq \emptyset$ for all i. For a 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 selection R such that $\Psi = \Theta(R)$. Note that $\Theta = \Theta(S)$. Note also that $\beta(x) = \Theta(\mathcal{B}(x))$ and $\sigma(x) = \Theta(S(x))$.

For $A \subset \Theta$ let $\mathcal{B}_i(A) = \{s_i \in S_i | s_i \in \mathcal{B}_i(x) \text{ for some } x \in A\}$ 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 $\beta^k(A) = \beta\left(\beta^{k-1}(A)\right)$. 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 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 $\mathcal{S}_i(A) = \{s_i \in S_i : s_i \in \mathcal{S}_i(x) \text{ for some } x \in A\}$ 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 $\sigma^k(A) = \sigma\left(\sigma^{k-1}(A)\right)$. 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 **$\boldsymbol{\sigma}$-rationalizable** if it is an element of the selection $R \subset S$ which satisfies $\Theta(R) = \sigma^\infty(\Theta)$.

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 a weak, and not strict, never best reply in Γ and $\tilde{u}_i(s_i, s_{-i}) =$

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

E

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

F

Table 5: A game in which a pure persistent equilibrium in not a fixed point of σ.

$u_i(s_i, s_{-i})$ otherwise. Every pure strategy which is a weak never best-reply in Γ is, therefore, a strict never best reply in $\tilde{\Gamma}$. Let $\tilde{\beta}$ denote the best-reply correspondence of $\tilde{\Gamma}$. Then we have the following lemma.

Lemma 5 *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 2. 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 the game given in Table 6, 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 strict never best replies for players 2 and 3, respectively. There are no strategies which are weak but not strict never best replies. Hence, $\tilde{\beta}(x) = \beta(x)$ for any $x \in \Theta$. Player 1's strategy B is (the unique) best strategy when player's 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.

Note that this is a general phenomenon: Pure strategies, which are equivalent to mixed strategies, and, hence, are weak never best replies can never appear in the refined best-reply correspondence.

Let $\tilde{\beta}^{\infty}$ be defined analogously to β^{∞}. We call a pure strategy $s \in S$ **Dekel-Fudenberg rationalizable** (or DF-rationalizable[4]) if it is an element of the selection $R \subset S$ which satisfies $\Theta(R) = \tilde{\beta}^{\infty}(\Theta)$.

[4]Dekel and Fudenberg (1990) in fact allow players to hold beliefs which are arbitrary distributions over the set of possible opposition play. This gives rise to what one might call Dekel-Fudenberg 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 Dekel-Fudenberg correlated rationalizability is obtained). A strategy is Dekel-Fudenberg correlated rationalizable if

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

E

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

F

Table 6: A game in which for some $x \in \Theta$, $\sigma(x)$ is a proper subset of $\tilde{\beta}(x)$.

Theorem 6 *Let $\Gamma \in \mathcal{G}^*$. Every σ-rationalizable strategy for Γ is DF-rationalizable.*

The game in Table 6 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)).

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 the game given in Table 3 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 the game given in Table 7, 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 the game in Table 7 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 reduced normal form game, given in Table 8, of the extensive form game given in Figure 2, the set of properly rationalizable strategies, $\{A\} \times \{F\}$,

and only if it survives the Dekel-Fudenberg 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 Dekel-Fudenberg rationalizable strategies is obviously contained in the set of correlated Dekel-Fudenberg rationalizable strategies.

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 the game given in Table 7 from Samuelson (1992). In this game strategies C and F are weakly dominated, and, hence, not σ-rationalizable. In the 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

Table 7: A game in which the set of σ-rationalizable strategies includes an iteratively weakly dominated strategy.

In some special contexts σ-rationalizability does allow the iterated deletion of weakly dominated strategies. See section 6.

5 σ-CURB sets

The following definitions are due to Basu and Weibull (1991). A 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.

Again we can define strong CURB sets in a similar fashion. A 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 (1991)'s CURB*-set, a CURB set without weakly dominated strategies, is a σ-CURB set. The game given in Table 6 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. The following theorem is proven in Balkenborg (1992). We here give an alternative proof.

Theorem 7 *Let $\Gamma \in \mathcal{G}^*$. A selection $R \subset S$ is a σ-CURB set if and only if $\Theta(R)$ is an absorbing retract.*

Proof: "$\Leftarrow$": Let the 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 $U^r = \{y \in U | r \in \mathcal{B}(y)\}$. 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 whose closure includes 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 Theorem 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 $y_U \to x$ 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

6 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 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 1, 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 reduced normal form has been introduced to eliminate exactly this type of equivalences. The reduced normal form of any GEFGOPI is again in $\mathcal{G}^*$.

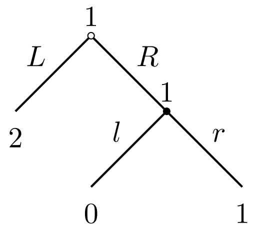

Figure 1: A 1-player extensive form game.

Theorem 8 *Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of a GEFGOPI. Then only the subgame-perfect strategy profile is 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

Theorem 9 *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 Theorem 8, only consists of the subgame perfect equilibrium. QED

None of the above theorems is true for the reduced normal form. Consider the centipede game (Figure 8.2.2 in Cressman (2003)) given here in Figure 2. This game is a GEFGOPI and, hence, has a unique subgame perfect equilibrium, which is (Lr, Rr). Note that this is, of course, also the unique sequential (Kreps and Wilson (1982)) and unique weak perfect Bayesian equilibrium.

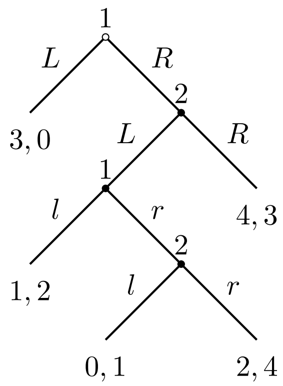

Figure 2: A centipede game.

The reduced normal form of this game is given in Table 8 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 $F = Rr$. 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 do not induce sequential or even weak perfect Bayesian equilibria in every extensive form game with this reduced normal form.

Also not every sequential equilibrium is necessarily a fixed point of σ. The game given in Figure 3, 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 reduced normal form are the same and given in Table 9.

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 4. The Nash equilibrium (A, R, r) is a fixed point of σ, but 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

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

Table 8: The normal form game of the centipede game in Figure 2.

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 theorem is reminiscent of the result that a proper equilibrium (Myerson (1978)) of a normal form game induces a sequential equilibrium in any extensive form with this normal form. The result here is, however, only for pure strategy profiles, and states that any pure fixed point of σ induces a (weak) perfect Bayesian equilibrium (see e.g. Definition 6.2 in Ritzberger (2002)) in every extensive form game with this reduced normal form.

Theorem 10 *Let $\Gamma \in \mathcal{G}^*$ be a normal form game. Then if a pure strategy profile s is a fixed point of σ it induces a perfect Bayesian equilibrium in any extensive form game with this normal form game as its reduced normal form.*

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 4 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

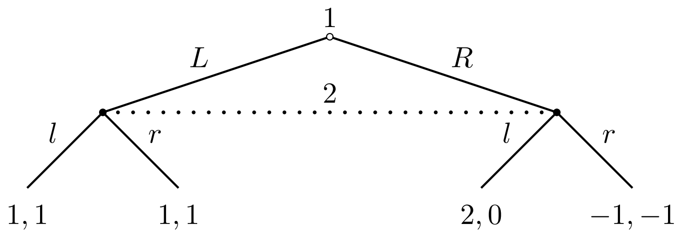

Figure 3: A game with a sequential equilibrium (L, r) which is not a fixed point of σ.

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 say $C(x_i)$ 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 perfect Bayesian equilibrium.

7 A micro model leading to the minimal refined best-reply dynamics

In section 8 we will finally consider the *refined best-reply dynamics*

$$\dot{x} \in \sigma(x) - x, \tag{1}$$

where σ is as in section 3. We discuss the properties of this dynamic in detail in section 8. In this section we provide a micro-motivation for this refined best-reply dynamic (1). To do this we first consider a micro model leading to the best-reply dynamic process (2), similar in spirit to some of the models in Björnerstedt and Weibull (1996); see also section 4.4 in Weibull (1995). 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

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

Table 9: The normal form game of the game in Figure 3.

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. For a moment let $\alpha \in [0, 1]$ denote the fraction of reviewing agents, whatever their strategy, who switch to s. Hence, $\dot{x}_{is} = (1 - x_{is})\alpha - x_{is}(1 - \alpha)$, which leads to $\dot{x}_{is} = \alpha - x_{is}$. Given $\alpha \in [0, 1]$ can take on any value the combinations of the three cases above leads to the specification in equation (2).

A similar micro story is also sketched in Gilboa and Matsui (1991). 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 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 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

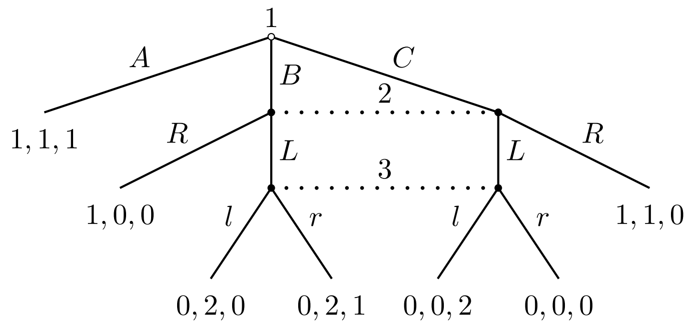

Figure 4: A game in which there is a fixed point of σ in the agent normal form which is not sequential (and, hence, not extensive form trembling hand perfect).

density is 0 only on a set with Lebesgue-measure 0. This means there is heterogeneity in agents' prior belief. 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. $U_\epsilon^x = \{y \in \Theta | ||x-y||_\infty \leq \epsilon\}$, 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.

This means that 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.

Assuming that the prior distribution is arbitrary (and potentially different at every point in time), but satisfies the restrictions posed, and provided that ϵ is small enough, exactly how small depends on the game, the induced learning dynamics can again be written as (1).

8 The refined best-reply dynamics: Results

Gilboa and Matsui (1991), Matsui (1992) and Hofbauer (1995) introduced the continuous time best reply dynamics, which modulo a time change is equivalent to the continuous time version of fictitious play. This best-reply dynamic is given by the differential inclusion

$$\dot{x} \in \beta(x) - x. \tag{2}$$

A solution to (2) is an absolutely continuous function $x(t)$, defined for at least $t \geq 0$, that satisfies (2) for almost all t. [5] 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.

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

Table 10: A game in which the BR dynamics seems to have 'too many' solutions.

For the game given in Table 10, 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. There is a deeper reason for this seeming anomaly. To explain it we need to consider payoff perturbations[6].

The perturbed game, given in Table 11, has unique solutions (from most initial conditions). The limits of these solutions, as $\epsilon \to 0$ must be solutions of (2) of the unperturbed game, by elementary upper hemi-continuity (UHC) properties. If we chose a sequence $\epsilon_n \to 0$ with $\epsilon_{2n} > 0$ and $\epsilon_{2n+1} < 0$, we can obtain any zig-zag solution in the limit. Hence these many irregular solutions within the NE component are a consequence of continuous dependence and payoff perturbations.

Only if the payoffs of the game are kept fixed, many of these solutions become dispensable. The main reason for this is that for some states x some of the best replies are dispensable, in fact leading to our refined best-reply dynamic 1, reproduced here:

$$\dot{x} \in \sigma(x) - x, \tag{3}$$

[5] Gilboa and Matsui (1991) and Matsui (1992) require additionally 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 (2) 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).

[6] On first glance it might be natural to dismiss all non-constant solutions through a NE. But in the above game the solutions then violate the continuous dependence on initial conditions - an extremely useful property.

	C	D
A	1,1	1,1 + ϵ
B	1 + ϵ,1	0,0

Table 11: A perturbed game where solutions move away from AC.

Since the right hand side is UHC with compact and convex values, existence of at least one Lipschitz-continuous solution $\zeta(t, x_0)$ through each initial state x_0 is guaranteed.

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

In this section it is now shown that this refined best-reply dynamic converges to the set of σ-rationalizable strategies, and that every σ-CURB set is asymptotically stable under this dynamic. The proofs are the same as the proofs of the statements that every solution of the best-reply dynamic (2) converges to the set of rationalizable strategies and that every CURB set is asymptotically stable under the best-reply dynamic. 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 11 *Let $\Gamma \in \mathcal{G}^*$. Let R be the 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 $\zeta_{is_i}(t, x_0) \to 0$ for any solution ζ to (1) 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 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 $s_i \notin \mathcal{S}_i(x)$ for any $x \in \Theta$. Hence its growth rate according to (1) is $\dot{x}_{is_i} = 0 - x_{is_i}$, and, hence, x_{is_i} 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 $\zeta_{is_i}(t, x_0) \to 0$ for any solution ζ to (1) 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 $\zeta_{is_i}(t, x_0) < \epsilon$ for all $t \geq T$. Now by the definition of σ, $s_i \in S_i \setminus R_i^k$ implies that $s_i \notin \mathcal{S}_i(\zeta(t, x_0))$ 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} shrinks exponentially to zero. QED

Theorem 12 *Let $\Gamma \in \mathcal{G}^*$. Let R be a σ-CURB set. Then $\Theta(R)$ is asymptotically stable under (1).*

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 $||\zeta(t, x_0) - \Theta(R)||_\infty$ shrinks exponentially to zero for all $x_0 \in U$. QED

A corollary of Theorem 12, combined with Theorem 7, is that Kalai and Samet (1984)'s persistent retracts are asymptotically stable under the refined best-reply dynamic (1).

If a solution of the refined best-response dynamic converges it must, of course, necessarily converge to a fixed point of the refined best-reply correspondence, σ. There are fixed points of σ, however, which no solution can converge to. For mixed equilibria this is very easy to see. Consider the coordination game. The mixed Nash equilibrium is a fixed point of σ, but no solution to the refined best-reply correspondence will converge to this equilibrium, unless the initial value is exactly this equilibrium. There are also pure fixed points of the refined best-reply correspondence which (essentially) no solution can converge to. In fact the following is true.

Theorem 13 *Suppose $x^* \in \Theta$ is pure strategy profile and is such that there is an open set $O \subset \Theta$ such that for every $x \in O$ there is a solution ζ to (1) which converges to x^*. Then $x^* \in \sigma(x^*)$ and x^* is perfect.*

This theorem is not true for mixed strategy profiles x^* as the game in Table 2 illustrates. Hendon, Jacobson, and Sloth (1996) demonstrate that the best-reply dynamic, which in this game is the same as the refined best-reply dynamic, does converge to the non-perfect mixed equilibrium $x^* = ((0, 1/2, 1/2); (1/2, 0, 1/2))$ from a open set of initial values.

Also it is not true that the refined best-reply dynamic necessarily converges to a persistent retract as the game given in Table 4 demonstrates. The non-persistent equilibrium (B, D, E) is an attractor for an open set of initial values.

Note that for some games there are sets which are proper subsets of persistent retracts which are asymptotically stable. Consider the game given in

Table 7. The unique persistent retract is the set of σ-rationalizable strategies $\sigma^\infty(\Theta) = \Delta(\{A, B\}) \times \Delta(\{D, E\})$. The set $\Psi = \{x \in \sigma^\infty(\Theta) | x_{1B}x_{2E} = 0\}$, which is not a retract, is also asymptotically stable. Note also that the refined best-response correspondence of this game projected onto (or constrained to) the set $\Delta(\{A, B\}) \times \Delta(\{D, E\})$ is exactly the same as the best-response correspondence in the game given in Table 10. This suggests an alternative interpretation why one might consider the full best-response dynamic for the game in Table 10. That the payoff perturbations are due to the presence of strategies, not specified in the game, which, however, are used with very small and unknown probability.

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.

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.

Börgers, 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.

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.

GILBOA, I., AND A. MATSUI (1991): "Social stability and equilibrium," *Econometrica*, 59, 859–67.

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.

KALAI, E., AND D. SAMET (1984): "Persistent equilibria in strategic games," *International Journal of Game Theory*, 13, 129–44.

KREPS, D. M., AND R. WILSON (1982): "Sequential equilibria," *Econometrica*, 50, 863–94.

MATSUI, A. (1992): "Best response dynamics and and socially stable strategies," *Journal of Economic Theory*, 57, 343–62.

MYERSON, R. B. (1978): "Refinements of the Nash equilibrium concept," *International Journal of Game Theory*, 7, 73–80.

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 (1996): "Evolutionary selection in normal form games," *Econometrica*, 63, 1371–1399.

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 Nachfrageträgheit," *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.

Weibull, J. W. (1995): *Evolutionary Game Theory.* MIT Press, Cambridge, Mass.

Young, H. P. (1993): "The evolution of conventions," *Econometrica*, 61, 57–84.