

Refined best reply correspondence and dynamics

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

January 27, 2012

Abstract

We call a correspondence, defined on the set of mixed strategy profiles, a generalized best reply correspondence if it (1) has a product structure, (2) is upper hemi-continuous, (3) always includes a best reply to any mixed strategy profile, and (4) is convex- and closed-valued. For each generalized best reply correspondence, we define a generalized best reply dynamics as a differential inclusion based on it. We call a face of the set of mixed strategy profiles a minimally asymptotically stable face (MASF) if it is asymptotically stable under some such dynamics and no subface of it is asymptotically stable under any such dynamics. The set of such correspondences (and dynamics) is endowed with the partial order of point-wise set inclusion and, under a mild condition on the normal form of the game at hand, forms a complete lattice with meets based on point-wise intersections. The refined best reply correspondence is then defined as the smallest element of the set of all generalized best reply correspondences. We find that every persistent retract (Kalai and Samet 1984) contains an MASF. Furthermore, persistent retracts are minimal CURB sets (Basu and Weibull 1991) based on the refined best reply correspondence. Conversely, every MASF must be a prep set (Voorneveld 2004), based again, however, on the refined best reply correspondence.

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

[‡]Institute of Mathematical Economics, Bielefeld University, christoph.kuzmics@uni-bielefeld.de.

We would like to thank Carlos Alos-Ferrer, Pierpaolo Battigalli, Eddie Dekel, Amanda Friedenberg, Drew Fudenberg, Klaus Ritzberger, Karl Schlag, Mark Voorneveld, Jörgen Weibull, and editor Martin Osborne as well as four anonymous referees for helpful comments and suggestions.

1 Introduction

Evolutionary support for Nash equilibrium behavior in general, finite n -player games is mixed. On the one hand, strict Nash equilibria (which necessarily must be in pure strategies) are evolutionarily stable (multipopulation ESS; see, e.g. Weibull (1995, Definition 5.1)) and asymptotically stable under the multipopulation replicator dynamics. In fact, strict Nash equilibria are the only asymptotically stable states under the multipopulation replicator dynamics and other related imitation-based dynamics, as shown by Ritzberger and Weibull (1995). Of course, many games of interest do not have a strict Nash equilibrium.

On the other hand, non-strict pure Nash equilibria and properly mixed Nash equilibria do not have a lot of evolutionary support in general games. To demonstrate the latter point, Hofbauer and Swinkels (1995) and Hart and Mas-Colell (2003) consider a class of finite normal form games, in which each game has a unique but mixed Nash equilibrium; they show that any “reasonable” deterministic continuous-time dynamic process must fail to make this unique Nash equilibrium asymptotically stable in at least one of these games.

We thus have to abandon the hope of obtaining Nash equilibria as the only outcomes of evolutionary processes. Yet this is not the end of studying the outcome of evolution. One just has to accept that evolution will lead, at least in some games, to a set of states that also includes some non-Nash equilibrium states. It may still be the case that these evolutionarily stable sets are in practice quite manageable and useful for the analysis of games. Note that switching from strategy profiles to sets of strategy profiles has also been necessary in the study of Nash equilibrium refinements (see Kohlberg and Mertens (1986)) and in the study of the consequences of common knowledge of rationality in general games (see, e.g. Bernheim (1984) and Pearce (1984)).

Given the evolutionary appeal of some pure strategy profiles, it is natural to study setwise generalizations of pure strategy profiles. A useful generalization of a pure strategy profile is given by a face (of the polyhedron of mixed strategy profiles) that is obtained simply by choosing a subset of pure strategies for every player and then considering all independent mixtures over these subsets.

We are not the first to propose to study the evolutionary stability properties of such faces. Indeed, Ritzberger and Weibull (1995) identify faces that are asymptotically stable under a large class of imitation-based dynamics. These faces are spanned by what Ritzberger and Weibull (1995) call “cuwbr sets”, which are product sets of pure strategy profiles that are closed under weakly better replies. Unless a minimal cuwbr set is a singleton, it does not contain a strict Nash equilibrium, but must contain a (possibly mixed) Nash equilibrium.

There are two (related) drawbacks of Ritzberger and Weibull’s (1995) result. First, in many games even the smallest cuwbr sets are very large sets. Thus, their predictive power is limited. Second, and this is a possible reason for their limited predictive power, the dynamics on which these sets are based, while plausible in some settings, are not necessarily the most plausible in games with highly rational and highly informed players. More rational and more informed players might not adapt their strategies toward better

replies so very gradually as is implicit in the class of dynamics of Ritzberger and Weibull (1995). One alternative with sharper predictions (i.e. smaller asymptotically stable sets) and more plausible adjustment behavior for highly rational and highly informed human beings is the best reply dynamics of Gilboa and Matsui (1991), Matsui (1992), and Hofbauer (1995) (in the spirit of fictitious play). To see how Ritzberger and Weibull's (1995) imitation-based dynamics and the best reply dynamics differ in the sharpness of their predictions, consider the following 2-player game.

	<i>H</i>	<i>T</i>	<i>C</i>
<i>H</i>	4,0	0,4	2 ϵ , ϵ
<i>T</i>	0,4	4,0	2 ϵ , ϵ
<i>C</i>	ϵ ,2 ϵ	ϵ ,2 ϵ	ϵ , ϵ

Game 1: A game to demonstrate the difference between imitation-based and best reply dynamics.

For $\epsilon \in (0, 2)$, Game 1, which is matching pennies with an additional strategy, has a unique smallest cuwbr set: the set of all strategy profiles. It is easy to see that both *H* and *T* need to be in any minimal cuwbr set for both players. However, the unique minimal cuwbr set also includes pure strategy *C* even if ϵ is very close to zero. To see this, note that when play is, for instance, sufficiently close to (H, H) , strategy *C* is better than strategy *H* for player 2 (and thus better than the average strategy employed by all individuals in player population 2). Under any dynamics considered in Ritzberger and Weibull (1995), the share of *C* strategists in population 2 must then grow for some finite amount of time.

Yet for all $\epsilon \in (0, 2)$, strategy *C* is strictly dominated for both players by the mixed strategy that puts equal weight on *H* and *T*. Thus, if a human being were to play this game, and were told the current state of play and allowed to change her behavior, it seems unlikely that she would choose strategy *C*. Indeed, under the best reply dynamics, strategy *C* will never be adopted by any revising agent. Thus, the best reply dynamics will eliminate *C* from any initial state. The unique minimal asymptotically stable face under the best reply dynamics is the face spanned by the unique minimal CURB set (as defined by Basu and Weibull (1991)) $\{H, T\} \times \{H, T\}$.

Hurkens (1995) analyzes a stochastic variant of the best reply dynamics. To be more precise, he studies a stochastic version of fictitious play, in which players play best replies to samples from their memory as in the model of Young (1993). Hurkens (1995) shows that the limiting invariant distribution of the resulting Markov chain attaches probability 1 to the set of all minimal CURB sets. Analogously, one can prove that CURB sets are asymptotically stable under the best reply dynamics.¹

Even when an imitation-based dynamics and the best reply dynamics yield the same collection of asymptotically stable sets of states, their vector fields are very different. To see this, consider the replicator dynamics and the best reply dynamics for the following simple 2-player, 2-strategy game.

¹This follows from Lemma 7 in this paper.

	<i>L</i>	<i>R</i>
<i>T</i>	1,1	1,1
<i>B</i>	0,0	1,1

Game 2: A game to demonstrate the behavioral differences in imitation-based and best reply dynamics.




Figure 1: The vector fields of the replicator and best reply dynamics for Game 2.

Note that pure strategies *B* and *L* are weakly dominated for players 1 and 2, respectively. Figure 1 sketches the vector fields of the two dynamics for this game, where p denotes the proportion of *T* in player population 1 and q the proportion of *R* in player population 2. The replicator dynamics (Taylor 1979) takes play from an interior state to a possibly (weakly) dominated Nash equilibrium state on the boundary.² Thus, different initial states, under the replicator dynamics, lead to different, and often weakly dominated, Nash equilibrium outcomes. A consequence of this fact is that all Nash equilibrium boundary states are Lyapunov stable, yet none are asymptotically stable (not even the undominated equilibrium (T,R)).

The best reply dynamics, on the other hand, converges from any interior state in a straight line to the undominated Nash equilibrium (T,R) .³ However, there is something that the above picture does not show. Because the best reply dynamics is a differential inclusion, there can be several trajectories emanating from the same point. This happens here precisely for all the Nash equilibria. Here, this implies that there are solutions to the best reply dynamics that move gradually along the boundary. Thus, (T,R) is *not* Lyapunov stable because there are trajectories starting arbitrarily nearby and leaving any neighborhood. Only the whole Nash equilibrium component is asymptotically stable. Notice how central the point (T,R) is for the dynamics. Any trajectory starting in a completely mixed strategy combination and converging to the Nash equilibrium component

²The replicator dynamics is calculated as $\dot{p} = p(1-p)(1-q)$ and $\dot{q} = q(1-q)(1-p)$. Division yields, for $p \neq 0$, the separable ordinary differential equation, $\frac{dq}{dp} = \frac{q}{p}$, which, for the initial point $(p(0), q(0))$, has the unique solution $q(t) = \frac{q(0)}{p(0)}p(t)$.

³To see this, note that for interior points $(p(0), q(0))$ the differential equation for the best reply dynamics (1), defined below, is given by $\dot{p} = 1-p$ and $\dot{q} = 1-q$ with unique solution $p(t) = 1 - (1-p(0))e^{-t}$ and $q(t) = 1 - (1-q(0))e^{-t}$.

must either converge to (T,R) or go through that point before it converges to any other Nash equilibrium.

In this paper we take the view that the fact that only the whole Nash equilibrium component is asymptotically stable is due to inessential and superfluous trajectories. In fact, if one took the view that rational and informed human beings would never adopt dominated strategies, by, for instance, only playing cautious best responses (as defined by Pearce (1984)), then the resulting cautious best reply dynamics would make (T,L) the unique minimal asymptotically stable face.

Hurkens (1995) considers a second stochastic variant of the best reply dynamics, in which individuals restrict themselves to playing semi-robust best replies, as defined by Balkenborg (1992). A definition of semi-robust best replies is also given in this paper. At this point it suffices to say that the set of semi-robust best replies to a particular strategy profile is a (sometimes proper) subset of the set of all best replies to this strategy profile. Hurkens (1995) shows that the limiting invariant distribution of the resulting Markov chain in this second model attaches probability 1 to the set of persistent retracts (as defined by Kalai and Samet (1984)). For the simple 2-player, 2-strategy game above, this implies that only the undominated pure Nash equilibrium (T,L) receives limiting probability 1. Persistent retracts are faces that are typically smaller, and never larger, than minimal CURB sets. To be more precise, every minimal CURB set contains a (possibly much smaller) persistent retract. Similarly every minimal cuwbr set contains a (possibly much smaller) minimal CURB set.

Motivated by the results of Ritzberger and Weibull (1995) and Hurkens (1995), we are interested in this paper in identifying and characterizing the smallest faces that are evolutionarily stable under some reasonable dynamics (appropriate for highly rational and highly informed human beings). We are interested in this for the following reason: Suppose a pure Nash equilibrium is not such a smallest evolutionarily stable face. Then there is not a single dynamics, from the whole class of dynamics we consider, under which this Nash equilibrium is evolutionarily stable. Similarly, if a face is such that no subface of it is such a smallest evolutionarily stable face, then, again, this face is not evolutionarily stable under any dynamics in the large class we consider.

We restrict attention to best-reply-like dynamics. This is motivated by the intuitive appeal of best responding for highly rational and informed human beings as well as by the fact that Ritzberger and Weibull's (1995) asymptotically stable minimal cuwbr sets are typically much larger than the minimal CURB sets or persistent retracts which Hurkens (1995) identifies as the outcome of his two variants of Young's (1993) model of best reply learning. We define and study a large class of *generalized best reply dynamics*, which is supposed to contain all reasonable best reply like dynamics.⁴

We can thus define a minimally asymptotically stable face (MASF) as a face that is asymptotically stable under some generalized best reply dynamics, with the additional property that it does not contain a proper subface that is also asymptotically stable under some (possibly different) generalized best reply dynamics. Note that it may well seem

⁴This class also contains some unreasonable dynamics. It will become clear in the analysis, however, that this does not pose a problem for the interpretation of our results.

possible that the dynamics that makes one MASF asymptotically stable in one game is quite different from the dynamics that makes another MASF asymptotically stable in another game. Our first main result, however, shows that this is not possible. In fact, under a mild restriction on the class of games we can study, there is a single dynamics, the same for all games, which determines which faces are MASFs and which are not. That is, a face is a MASF if and only if it is a minimally asymptotically stable face under this particular dynamics. We are thus justified in terming this dynamics the (most) *refined best reply dynamics*.

The refined best reply dynamics is a very reasonable and intuitive dynamics (for highly rational and highly informed individuals). The following micro-story is adapted from Hofbauer's (1995) story, which gives rise to the best reply dynamics. For every player position, there is a large population of individuals. Time is continuous and runs from zero to infinity. Individuals play pure strategies. At time zero, individuals' behavior is given by some arbitrary frequency distribution of pure strategies, with one distribution for each population. In every short time interval, a small fraction of individuals is given the opportunity to revise their strategy. When doing so, individuals know the aggregate distribution of play (the state of play). If there is a unique best reply, a revising individual adopts it. If there are multiple best replies, a revising individual considers them all, but adopts only one of those that are also unique best replies in an open set of nearby states of play. One could call this a cautious, myopically rational individual.⁵ One similar alternative story could be that revising individuals do not know the exact state of play. Different individuals have different beliefs (that are close to the truth) about the aggregate play. Any individual's belief over play in any two opponent populations i and j is assumed to be statistically independent.⁶ If these beliefs are sufficiently diverse, only a vanishing fraction of individuals will adopt a strategy that is best only on a set of states with Lebesgue measure zero. This again gives rise to the refined best reply dynamics.⁷

Interestingly, the refined best reply dynamics is based on Balkenborg's (1992) semi-robust best replies, and thus is in some sense analogous to the second stochastic model of learning studied by Hurkens (1995). This then gives rise to the question whether MASFs are exactly those faces which Hurkens (1995) identifies as the long-run outcome of his learning process, namely, persistent retracts. Somewhat surprisingly, the answer to this question is no. There are MASFs in some games that are proper subfaces of persistent retracts, as we demonstrate by example. In an effort to at least partially characterize MASFs in terms of concepts known from the literature, we ultimately show that every persistent retract contains a MASF and every MASF must be an appropriate version of a prep set (as defined by Voorneveld (2004)). The appropriate version is not the original prep set, which is based on the best reply correspondence, but is instead based on the

⁵The behavior of a revising individual is reminiscent of, yet not completely the same as, playing Pearce's (1984) cautious response.

⁶This statistical independence implies that beliefs can be represented as a state of play, an independent vector of distributions over strategies, one for each player position.

⁷The approach in Roth and Sandholm (2011) can be used to obtain approximations of the refined best reply dynamics by discrete time stochastic processes.

refined best reply correspondence, which underlies the refined best reply dynamics.

Methodologically, this paper overlaps in part with Balkenborg (1992). In order to analyze the properties of persistent retracts, that paper 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 the class of generalized best reply correspondences in section 3, where we also prove their lattice structure and the existence of a unique smallest element. In section 4 we study the notion of a CURB set (Basu and Weibull (1991)) and a prep set⁸ (Voorneveld (2004)) for all generalized best reply correspondences, and we study their relationship. We also prove that CURB sets based on the refined best reply correspondence coincide with Kalai and Samet’s (1984) persistent retracts. Finally, Section 5 provides the main results: persistent retracts are asymptotically stable under the refined best reply dynamics and thus contain an MASF, and every MASF must be a tight prep set based on the refined best reply correspondence. Section 6 concludes. The paper has one appendix, where we show in which sense our restriction to games with generically unique best replies is not essential.

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 $u : S \rightarrow \mathbb{R}^n$ is the payoff function.⁹ 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.¹⁰ 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

⁸A learning model that leads to the original prep sets of Voorneveld (2004) is given by Kets and Voorneveld (2008).

⁹The function u will also denote the expected utility function in the mixed extension of the game Γ .

¹⁰Generally, let $\Delta(K)$ for any finite set K denote the set of all probability distributions over K .

concepts primarily for pure strategies.

Let $\Psi = \{x \in \Theta : \mathcal{B}(x) \text{ is a singleton}\}$. Notice that the unique best reply against a strategy combination in Ψ is necessarily a *pure* strategy. Throughout this paper we 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 3. 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 is because player 1 has two own-payoff equivalent pure strategies.

	<i>C</i>	<i>D</i>
<i>A</i>	1,1	1,0
<i>B</i>	1,0	1,1

Game 3: 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 nonempty 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 any 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 reply 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 for 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 that 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 $\Gamma \in \mathcal{G}^*$. 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 .*

Games in the class \mathcal{G}^* are essentially those that do not have own-payoff equivalent pure strategies for any player. The semi-reduced normal form of a game is usually obtained by removing all payoff equivalent strategies. In the appendix we argue that the games in which there are own-payoff equivalent pure strategies that are not payoff-equivalent are exceptional. Hence, the restriction to games in the class \mathcal{G}^* made throughout the paper is

essentially the restriction to the semi-reduced normal form in the sense of Kohlberg and Mertens (1986).¹¹ Since we are primarily interested in the best reply 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 that projects onto it.

3 Generalized best reply correspondences

Definition 1 *A correspondence $\tau : \Theta \rightrightarrows \Theta$ is a generalized best reply correspondence if all of the following hold:*

1. $\tau(x) = \times_{i \in I} \tau_i(x) \forall x \in \Theta$, where $\tau_i : \Theta \rightrightarrows \Theta_i$ for all $i \in I$.
2. τ is upper hemi-continuous¹² at all $x \in \Theta$.
3. $\tau_i(x) \cap \beta_i(x) \neq \emptyset \forall x \in \Theta, \forall i \in I$.
4. $\tau(x)$ is convex and closed for all $x \in \Theta$.

Note that property 3 immediately implies that $\tau(x) \neq \emptyset$. Thus, a generalized best reply correspondence has the same basic technical properties as the best reply correspondence β ; and moreover, it is minimally connected to the best reply correspondence β by the requirement that at least one best reply to some given strategy profile x (i.e. an element of $\beta(x)$) is also available in $\tau(x)$.

In words, a generalized best reply correspondence has four properties. It must have a product structure, since we want players to choose independently. At least one best reply must always be available to players. It must be point-wise closed and convex. Closedness is more of a technical requirement, but the requirement of convexity derives from the desire to have players randomize arbitrarily between their generalized best replies. Finally, we require a generalized best reply correspondence to be upper hemi-continuous. This is an important technical requirement as it guarantees that such a generalized best reply correspondence has a fixed point and that the differential inclusion based on it always has a solution. In terms of player behavior, this translates to the requirement that if one were to slightly perturb the current strategy profile of a player's opponents, that player will not choose a new strategy that was not formerly in the set of generalized best replies.

A subclass of the set of generalized best reply correspondences which is of independent interest is the one that is based on pure strategies only, in the following sense. A correspondence $\tau : \Theta \rightrightarrows \Theta$ is a *generalized best reply correspondence based on pure strategies* if it is a generalized best reply correspondence, and if property 4 is replaced by the more

¹¹In particular, we are not ruling out games with, for instance, weakly dominated strategies.

¹²Following Aliprantis and Border (1999, ch.17.2), or Ritzberger (2002, Definition 5.8), the correspondence τ is upper hemi-continuous at x if for every open set $V \subset \Theta$ with $\tau(x) \subset V$, there is an open subset $U \subset \Theta$ with $x \in U$ such that for all $y \in U$: $\tau(y) \subset V$.

stringent property 4* such that $\tau_i(x) = \Delta(T_i(x))$ for some $T_i(x) \subset S_i$ for all $x \in \Theta$ and for all $i \in I$. Then property 3 is equivalent to $T_i(x) \cap \mathcal{B}_i(x) \neq \emptyset$ for all $x \in \Theta$ and for all $i \in I$.

Let $\mathcal{T} = \mathcal{T}(\Gamma)$ denote the set of all generalized best reply correspondences (of a game Γ), and let \mathcal{T}^{PS} denote the subset of all generalized best reply correspondences based on pure strategies.

One natural example of a correspondence which is in \mathcal{T} but not in \mathcal{T}^{PS} is the correspondence of all mixed, weakly better replies. This correspondence is given by $\tau = \times_{i \in I} \tau_i$ with $\tau_i(x) = \{y_i \in \Theta_i : u_i(y_i, x_{-i}) \geq u_i(x_i, x_{-i})\}$.

One example of a correspondence in \mathcal{T}^{PS} is, of course, the best reply correspondence itself. For another example, let $T_i(x) = \{s_i \in S_i : u_i(s_i, x_{-i}) \geq u_i(x_i, x_{-i})\}$. That is $T_i(x)$ is the set of all weakly better replies to x_{-i} given x_i . The resulting correspondence is that of all mixtures of pure, weakly better replies (see Ritzberger and Weibull (1995)). Another example, closely connected to the $S^\infty W$ -procedure of Dekel and Fudenberg (1990), can be found by letting $T_i(x)$ be the set of all pure best replies, except weakly dominated ones.

The following example of a correspondence in \mathcal{T}^{PS} is key to the subsequent analysis in this paper.

Definition 2 For games in \mathcal{G}^* , and for $x \in \Theta$, let

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

Then $\sigma_i(x) = \Delta(\mathcal{S}_i(x))$ and $\sigma(x) = \times_{i \in I} \sigma_i(x) \forall x \in \Theta$. We call this correspondence $\sigma : \Theta \Rightarrow \Theta$ the (most) refined best reply correspondence.

The set $\mathcal{S}_i(x)$ in the above definition is the set of pure semi-robust best replies, as defined in Balkenborg (1992).

Given two correspondences $\tau, \tau' \in \mathcal{T}$, let $\tau \subset \tau'$ if $\tau(x) \subset \tau'(x)$ for all $x \in \Theta$. The set \mathcal{T} endowed with this notion of “smaller than” is a partial order (see, e.g. Davey and Priestley (2002, Definition 1.2)). Given two correspondences $\tau, \tau' \in \mathcal{T}$, let $\tau'' = \tau \wedge \tau'$ if $\tau''_i(x) = \tau_i(x) \cap \tau'_i(x)$ and $\tau''(x) = \times_{i \in I} \tau''_i(x)$ for all $x \in \Theta$.

The first theorem of this paper demonstrates that for games in \mathcal{G}^* the set \mathcal{T} of generalized best reply correspondences has a lot of structure, and especially, a smallest element. In fact, the set \mathcal{T} is a complete lattice, which means that every subset of \mathcal{T} has an infimum (meet) and a supremum (join) (see, e.g. Davey and Priestley (2002, Definition 2.4)).

Theorem 1 Let $\Gamma \in \mathcal{G}^*$. Then:

1. \mathcal{T} is a complete lattice.
2. \mathcal{T} has a unique smallest element, which is given by σ , the refined best reply correspondence.
3. $\sigma \in \mathcal{T}^{PS}$.

Proof: Let \mathcal{T}' be a nonempty subset of \mathcal{T} . Define τ^* by $\tau^*(x) = \bigcap_{\tau \in \mathcal{T}'} \tau(x)$ for all $x \in \Theta$. We claim that $\tau^* \in \mathcal{T}$. Since the intersection of products, convex sets, and closed sets is itself a convex and closed product, τ^* has properties (1) and (4) of a generalized best reply correspondence. Any $\tau \in \mathcal{T}'$ is in fact compact valued. The set Θ is compact and Hausdorff and therefore a regular topological space (Aliprantis and Border 1999, Theorem 2.48). Hence, τ^* is upper hemi-continuous (Aliprantis and Border 1999, Theorem 17.25.3), i.e. it satisfies property (2) of a generalized best reply correspondence.

It remains for us to show that τ^* has property (3) of a generalized best reply correspondence. Since $\Gamma \in \mathcal{G}^*$, $\beta(x)$ is by definition a singleton for all $x \in \Psi$, and is given by the pure strategy that is the unique element of $\mathcal{B}(x)$. Thus, for all $x \in \Psi$, we have $\beta(x) \subset \tau(x)$ for any $\tau \in \mathcal{T}$. For $x \in \Psi$, let $\mathcal{S}_i(x) = \mathcal{B}_i(x)$ for all $i \in I$. Since $\Gamma \in \mathcal{G}^*$, it follows that Ψ is dense in Θ . For $x \notin \Psi$, let $U \subset \Theta$ be a neighborhood of x , let $\mathcal{S}_i^U(x) = \bigcup_{x' \in U \cap \Psi} \mathcal{B}_i(x')$, and let $\mathcal{S}_i(x) = \bigcap_{(U \text{ neighborhood of } x)} \mathcal{S}_i^U(x)$. Let $\sigma_i(x) = \Delta(\mathcal{S}_i(x))$ and let $\sigma(x) = \times_{i \in I} \sigma_i(x) \forall x \in \Theta$. Thus, σ is the refined best reply correspondence.

Then, by properties (2) and (4) of any $\tau \in \mathcal{T}$, we must have $\sigma(x) \subset \tau(x)$ for all $x \notin \Psi$, which implies that this holds for all $x \in \Theta$, and for all $\tau \in \mathcal{T}$. Thus, $\sigma \subset \tau^*$, and, as a result, τ^* satisfies property (3). It follows immediately that τ^* is the infimum of \mathcal{T}' in \mathcal{T} . The supremum of \mathcal{T}' is easily seen to be the intersection of all upper bounds of \mathcal{T}' in \mathcal{T} . The set \mathcal{T} is therefore a complete lattice, which proves part 1 of Theorem 1.

A complete lattice has a lowest element. Given that $\sigma \subset \tau$ for all $\tau \in \mathcal{T}$ and that $\sigma \in \mathcal{T}$, this lowest element must be σ . This proves part 2. Part 3 follows immediately from the construction of σ . QED

Note that Theorem 1 also implies that the set \mathcal{T}^{PS} is a complete lattice with the same smallest element, σ .

The converse of Theorem 1 is, in fact, also true, in the following sense. For any game $\Gamma \notin \mathcal{G}^*$, \mathcal{T} is not a lattice and does not have a unique smallest element. To see this, consider any game $\Gamma \notin \mathcal{G}^*$. This game must have at least two own-payoff equivalent pure strategies for some player that are simultaneous best replies in an open set of strategy profiles. Thus, one can construct $\tau \in \mathcal{T}$ based on only one of these two pure strategies, and $\tau' \in \mathcal{T}$ based only on the other, such that $\tau(x) \cap \tau'(x) = \emptyset$ for some $x \in \Theta$ and, as a result, $\tau \wedge \tau' \notin \mathcal{T}$.

Theorem 1 justifies our calling σ the *(most) refined best reply correspondence*, since σ is the unique smallest generalized best reply correspondence and obviously satisfies $\sigma \subset \beta$.

This refined best reply correspondence σ coincides with the best reply correspondence β almost everywhere (i.e. for all $x \in \Psi$, which is dense in Θ , given $\Gamma \in \mathcal{G}^*$). Furthermore, for strategy profiles $x \notin \Psi$, it is constructed in a minimal way to ensure upper hemi-continuity by requiring that $\sigma(x)$ includes those and only those pure strategies that are best replies to some nearby $x' \in \Psi$.¹³ For such x , any $\sigma(x)$ must then also include all convex combinations of all pure strategies in $\tau(x)$ by property 4.

¹³Strategies that are unique best replies to some x are called *inducible* in von Stengel and Zamir (2010).

We now provide a brief partial characterization of the refined best reply correspondence in terms of well-known objects from the theory of games. A detailed characterization of the refined best reply correspondence, its fixed points, and other objects based on it can be found in our companion paper, Balkenborg, Hofbauer, and Kuzmics (2009). For 2-player games the refined best reply correspondence includes those and only those best replies that are not weakly dominated and are not equivalent to a mixture of other pure strategies. The following example (from Balkenborg, Hofbauer, and Kuzmics (2009)) demonstrates part of this claim.

	<i>D</i>	<i>E</i>	<i>F</i>
<i>A</i>	2,2	1,2	1,2
<i>B</i>	2,1	2,2	0,0
<i>C</i>	2,1	0,0	2,2

Game 4: A Game with equivalent mixed strategies.

In this game, strategy *A* is equivalent to the mixture of pure strategies *B* and *C* (and strategy *D* to the mixture of pure strategies *E* and *F*). 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. Thus, we have $y_{1A} = 0$ for all $y \in \sigma_1(x)$ and for all $x \in \Theta$, i.e. *A* is never in the set of refined best replies. This implies, for instance, that there are strategically stable equilibria in the sense of Kohlberg and Mertens (1986) that are not fixed points of the refined best reply correspondence. For games with more than two players, the set of refined best replies at a given strategy profile x may well be a proper subset of the set of best replies that are not weakly dominated and not equivalent to a mixed strategy. For a thorough discussion of this, we refer the reader to our companion paper, Balkenborg, Hofbauer, and Kuzmics (2009).

4 τ -CURB and τ -prep sets

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

Let $A \subset \Theta$. For any $\tau \in \mathcal{T}$, let $\tau(A) = \times_{i \in I} \tau_i(A)$ with $\tau_i(A) = \bigcup_{x \in A} \tau_i(x)$. The following definition is a generalized version of Basu and Weibull's (1991) CURB sets.

For $\tau \in \mathcal{T}$, a strategy selection R is a τ -CURB set if $\tau(\Theta(R)) \subset \Theta(R)$. It is a *tight τ -CURB set* if, in addition, $\tau(\Theta(R)) \supset \Theta(R)$, and hence, $\tau(\Theta(R)) = \Theta(R)$.¹⁴ It is a *minimal τ -CURB set* if it does not properly contain another τ -CURB set.

¹⁴Note that, for any $\tau \in \mathcal{T}$, the set of tight τ -CURB sets, together with the empty set, also forms a finite and thus complete lattice. This follows from the fact that the set of all pure strategy selections is a lattice if we include the empty set. Then τ is an order-preserving function from the set of subsets of the set of pure strategy sets to itself, i.e. if $R \subset R'$, then $\tau(R) \subset \tau(R')$. Thus, by Tarski's fixed point

These definitions, while well-defined for all $\tau \in \mathcal{T}$, are more natural for $\tau \in \mathcal{T}^{PS}$ since they are based on pure strategies. In fact, if $\tau \in \mathcal{T} \setminus \mathcal{T}^{PS}$, then there are typically not many tight τ -CURB sets.

Note that if $\tau = \beta$, we obtain the original definition of CURB sets of Basu and Weibull (1991). If τ is the correspondence of all mixtures of pure weakly better replies, as defined in Section 3, τ -CURB sets are Ritzberger and Weibull's (1995) *cwubr-sets* (closed under weakly better replies). If, as mentioned in Section 3, $T_i(x)$ is the set of all pure best replies, except weakly dominated ones, and if τ is the correspondence in \mathcal{T}^{PS} based on these $T_i(x)$, then τ -CURB sets are Basu and Weibull's (1991) *CURB* sets*.

The lattice structure of \mathcal{T} allows us to compare CURB sets based on different generalized best reply correspondences.

Lemma 2 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau, \tau' \in \mathcal{T}$ with $\tau \subset \tau'$. Then any τ' -CURB set is also a τ -CURB set. Furthermore, any τ -CURB set for some $\tau \in \mathcal{T}$ is also a σ -CURB set.*

Proof: For the first part, let R be a τ' -CURB set. Then, by definition, $\tau'(\Theta(R)) \subset \Theta(R)$. But since $\tau \subset \tau'$, we have $\tau(\Theta(R)) \subset \tau'(\Theta(R)) \subset \Theta(R)$. The second part follows from the first part of this lemma and the second part of Theorem 1. QED

Lemma 2 implies that the smallest of all τ -CURB sets for any $\tau \in \mathcal{T}$ are the minimal σ -CURB sets. And it turns out that we can characterize these minimal σ -CURB sets.

The following definitions are due to Kalai and Samet (1984). A set $\varphi \subset \Theta$ is a *retract* if $\varphi = \times_{i \in I} \varphi_i$, where $\varphi_i \subset \Theta_i$ is nonempty, compact, and convex. A set $\varphi \subset \Theta$ *absorbs* another set $\varphi' \subset \Theta$ if for all $x \in \varphi'$, we have $\beta(x) \cap \varphi \neq \emptyset$. A retract φ 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.

Lemma 3 *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 $\Theta(R)$. It follows 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 $U^r = \{y \in U : r \in \mathcal{B}(y)\}$. It is obvious that $\bigcup_{r \in R} U^r = U$.

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 $s_i \in \beta(y)$ for all $y \in O$ for some open set O , whose closure 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 that 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.

theorem, the set of all fixed points of τ also forms a lattice. These are tight τ -CURB sets (and the empty set).

" \Rightarrow ": Suppose $R \subset S$ is a σ -CURB set. Suppose also 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 also 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 $r_i \in R_i$ such that $r_i \in \beta(y_U)$ at least for a subsequence of all such y_U (converging to x), which again gives rise to a contradiction. QED

Lemma 3 immediately implies the following theorem.

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

Theorem 2, together with Lemma 2, implies that the smallest τ -CURB sets for any $\tau \in \mathcal{T}$ are Kalai and Samet's (1984) persistent retracts.

The largest tight β -CURB set is the set of rationalizable strategies (Bernheim (1984) and Pearce (1984)). We can similarly define, for any $\tau \in \mathcal{T}$, the set of τ -rationalizable strategies as the largest tight τ -CURB set.

Alternatively, we can define τ -rationalizable strategies, more in the original spirit of Bernheim (1984) and Pearce (1984), in the following way. We shall do this only for $\tau \in \mathcal{T}^{PS}$. If $\tau \in \mathcal{T}^{PS}$ then there are correspondences T_i for each player such that $\tau_i(x) = \Delta(T_i(x))$ for all $x \in \Theta$ and for all $i \in I$. For $A \subset \Theta$, let $T_i(A) = \bigcup_{x \in A} T_i(x)$. Let $\tau_i(A) = \Delta(T_i(A))$. Let $\tau(A) = \times_{i \in I} \tau_i(A)$. For $k \geq 2$, let $\tau^k(A) = \tau(\tau^{k-1}(A))$. For $A = \Theta$, $\tau^k(A)$ is a decreasing sequence. Let $\tau^\infty(\Theta) = \bigcap_{k=1}^{\infty} \tau^k(\Theta)$. A pure strategy profile $s \in S$ is τ -rationalizable if it is an element of the strategy selection $R \subset S$ that satisfies $\Theta(R) = \tau^\infty(\Theta)$.

We thus have notions of rationalizability for any generalized best reply correspondence (that is based on pure strategies).

Lemma 2 immediately implies the following result.

Corollary 1 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau, \tau' \in \mathcal{T}$ such that $\tau \subset \tau'$. Then every τ -rationalizable strategy is also τ' -rationalizable. In particular, every σ -rationalizable strategy is also τ -rationalizable for any $\tau \in \mathcal{T}$.*

This corollary thus states that the smallest set of τ -rationalizable strategies is obtained when $\tau = \sigma$.

It turns out that not only such τ -CURB sets play a role in our analysis of generalized best reply dynamics in the next section, but also τ -versions (especially the σ -version) of Voorneveld's (2004) prep sets. We shall define τ -prep sets only for $\tau \in \mathcal{T}^{PS}$. Let

$\tau \in \mathcal{T}^{PS}$. A strategy selection R is a τ -prep set if for all $x \in \Theta(R)$ and for all $i \in I$, it is true that $\tau_i(x) \cap \Delta(R_i) \neq \emptyset$. A τ -prep set is minimal if it does not properly contain any other τ -prep set. Thus, any pure fixed point of τ is a minimal τ -prep set, just as every pure Nash equilibrium is a minimal (β -)prep set. We shall call a τ -prep set *tight* if for every $s_i \in R_i$, it is the case that $s_i \in \tau_i(x)$ for some $x \in \Theta(R)$. Minimal τ -prep sets are necessarily tight.

Analogously to Lemma 2, we can compare τ -prep sets for different τ 's in \mathcal{T}^{PS} . However, the comparison is reversed, as the next lemma states.

Lemma 4 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau, \tau' \in \mathcal{T}^{PS}$ with $\tau \subset \tau'$. Then any τ -prep set is also a τ' -prep set. Furthermore, any σ -prep set is also a τ -prep set for every $\tau \in \mathcal{T}^{PS}$.*

Proof: For the first part, let R be a τ -prep set. Thus, by definition, for all $x \in \Theta(R)$ and for all $i \in I$, it must be the case that $\tau_i(x) \cap \Theta(R)_i \neq \emptyset$. But then, as $\tau \subset \tau'$, it follows that for all $x \in \Theta(R)$ and for all $i \in I$, we have $\tau'_i(x) \cap \Theta(R)_i \neq \emptyset$. The second part follows from the first part of this lemma and the second part of Theorem 1. QED

Thus, Lemmas 2 and 4 imply that the smaller $\tau \in \mathcal{T}$ is (in the lattice) the more τ -CURB sets and the fewer τ -prep sets there are.

5 Generalized best reply dynamics

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. This best reply dynamics is given by this differential inclusion:

$$\dot{x} \in \beta(x) - x. \tag{1}$$

Similarly, we can define generalized τ -best reply dynamics, for some $\tau \in \mathcal{T}$, given by this differential inclusion:

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

The best reply dynamics (1) is obviously a special case of (2) for $\tau = \beta$. A solution to (2) is an absolutely continuous function $\xi(t, x_0)$ through initial state $x_0 \in \Theta$, defined for at least all $t \geq 0$, that satisfies (2) for almost all t .¹⁵

Since the right hand side of (2) is upper-hemi continuous with compact and convex values, the existence of at least one Lipschitz-continuous solution $\xi(t, x_0)$ through each

¹⁵Gilboa and Matsui (1991) and Matsui (1992) additionally require 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 (e.g. 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 for when these constitute all solutions. See also Hofbauer and Sigmund (1998), Cressman (2003), and Sandholm (2010).

initial state x_0 is guaranteed for any $\tau \in \mathcal{T}$ (see Aubin and Cellina (1984, Chapter 2, Section 1, Theorem 3, p. 98)). In general, several solutions can exist through a given initial state.

The objects we aim to identify and characterize in this paper can now be defined as follows.

Definition 3 *Let $\Gamma \in \mathcal{G}^*$. A face $\Theta(R)$ (spanned by a strategy selection R) is a minimal asymptotically stable face (MASF) if there is a $\tau \in \mathcal{T}$ such that $\Theta(R)$ is asymptotically stable¹⁶ under $\dot{x} \in \tau(x) - x$ and for all proper subfaces $\Theta(R') \subset \Theta(R)$, where R' is a strategy selection; and for all $\tau' \in \mathcal{T}$, $\Theta(R')$ is not asymptotically stable under $\dot{x} \in \tau'(x) - x$.*

Since $\sigma(x) \subset \tau(x)$ for all x and for all $\tau \in \mathcal{T}$, every solution of the σ -best reply dynamics

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

is also a solution of the τ -best reply dynamics (2). This means that any substantial difference in the dynamics is due to the multiplicity of trajectories. In fact, this is worth stating as a lemma.

Lemma 5 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau, \tau' \in \mathcal{T}$ such that $\tau \subset \tau'$. Let $x^0 \in \Theta$ be an arbitrary initial state. Then every solution to $\dot{x} \in \tau(x) - x$ through x^0 is also a solution to $\dot{x} \in \tau'(x) - x$ through x^0 .*

Proof: This follows immediately from the partial order on \mathcal{T} . QED

This lemma immediately implies another lemma.

Lemma 6 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau, \tau' \in \mathcal{T}$ such that $\tau \subset \tau'$. Let $A \subset \Theta$. If A is asymptotically stable under $\dot{x} \in \tau'(x) - x$, then A is also asymptotically stable under $\dot{x} \in \tau(x) - x$. Furthermore, if A is asymptotically stable under $\dot{x} \in \tau(x) - x$ for some $\tau \in \mathcal{T}$, then A is also asymptotically stable under $\dot{x} \in \sigma(x) - x$.*

Proof: The first part follows directly from Lemma 5. The second part follows from the first part and the second part of Theorem 1. QED

¹⁶ We call A asymptotically stable if it is Lyapunov stable and attractive. A set A is Lyapunov stable if for every neighborhood U of A , there exists a neighborhood V of A such that all solutions $\xi(t, x_0)$ with $x_0 \in V$ satisfy $\xi(t, x_0) \in U$ for all $t \geq 0$. A is attractive if there is a neighborhood U of A such that for every solution $\xi(t, x_0)$ with $x_0 \in U$ its ω -limit set is contained in A : $\bigcap_{T \geq 0} \overline{\{\xi(t, x_0) : t \geq T\}} \subset A$.

Note that in contrast to Definition IX on p. 339 in Benaïm, Hofbauer, and Sorin (2005), we drop here the requirement (i) of invariance. A set A is invariant if for every $x_0 \in A$, there is a complete solution $\xi(t, x_0) \in A$ (i.e., defined for all positive and negative times $t \in \mathbb{R}$). Consider as an example the matching pennies game with the best reply dynamics. Then the only invariant set is the unique Nash equilibrium which is the minimal asymptotically stable set of this game. In particular, there are no invariant faces under the best reply dynamics. The only CURB set is the whole strategy space. But this is not invariant, only forward-invariant. Therefore, it is not reasonable to require invariance in the definition of MASF.

Lemma 6 thus implies that minimal asymptotically stable faces (MASF) are those and only those faces that are the smallest faces that are asymptotically stable under the (most) refined best reply dynamics $\dot{x} \in \sigma(x) - x$. In the remainder of this section we provide partial characterizations of such faces. We first show that for any $\tau \in \mathcal{T}^{PS}$, the τ -best reply dynamics converges to the set of τ -rationalizable strategies.¹⁷ Furthermore, every τ -CURB set is asymptotically stable under this dynamics. In particular, Basu and Weibull's (1991) CURB sets are asymptotically stable under the best reply dynamics, and Kalai and Samet's (1984) persistent retracts are asymptotically stable under the refined best reply dynamics. These results are similar to the results of Hurkens (1995), who showed that for a stochastic learning model à la Young (1993), recurrent sets coincide with either CURB sets or persistent retracts, depending on the details of the model. These results are also similar to that of Ritzberger and Weibull (1995), who show that any strategy selection that is closed under weakly better replies is asymptotically stable under any deterministic payoff-positive dynamics.

We then, however, we give an example of a game in which a proper sub-face of a persistent retract is asymptotically stable under the refined best reply dynamics. That is, an MASF can be smaller than a persistent retract. We show that a necessary condition for a face to be asymptotically stable under the refined best reply dynamics (i.e. to be a MASF) is that it constitute a tight σ -prep set (and must thus be a τ -prep set for all $\tau \in \mathcal{T}$).

Theorem 3 *Let $\Gamma \in \mathcal{G}^*$. Let $\tau \in \mathcal{T}^{PS}$. Let R be the strategy selection of S that spans the set of τ -rationalizable strategies, i.e., $\Theta(R) = \tau^\infty(\Theta)$. Let $s_i \in S_i \setminus R_i$. Then $x_{is_i}(t) \rightarrow 0$ for any solution $x(\cdot)$ to $\dot{x} \in \tau(x) - x$ 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 $\tau^\infty(\Theta) = \bigcap_{k=1}^{\infty} \tau^k(\Theta)$. Let R^k denote the strategy selection of S that spans $\tau^k(\Theta)$, i.e., $\Theta(R^k) = \tau^k(\Theta)$. For $k = 1$ consider an arbitrary strategy $s_i \in S_i \setminus R_i^1$. Then, by definition, $s_i \notin \tau_i(x)$ for any $x \in \Theta$. Hence, this strategy's growth rate according to $\dot{x} \in \tau(x) - x$ is $\dot{x}_{is_i} = 0 - x_{is_i}$, and therefore

$$x_{is_i}(t) = e^{-t}x_{is_i} \quad (4)$$

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 $s_i \in S_i \setminus R_i^{k-1}$; that is, for any such s_i , we have $x_{is_i}(t) \rightarrow 0$ for any solution $x(\cdot)$ to $\dot{x} \in \tau(x) - x$ 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 τ , for $s_i \in S_i \setminus R_i^k$ it is the case that $s_i \notin \tau_i(x(t))$, provided ϵ is small enough (or t large enough). But then for all $t \geq T$, it follows once again that $\dot{x}_{is_i} = 0 - x_{is_i}$ and hence, that $x_{is_i}(t)$ shrinks exponentially to zero. QED

¹⁷For $\beta = \tau$, this result seems to be well understood. For a related, but weaker, statement about the iterated elimination of strictly dominated strategies, see Sandholm (2010, Theorem 7.4.2).

Lemma 7 *Let $\Gamma \in \mathcal{G}^*$, let R be a strategy selection, and let $\tau \in \mathcal{T}^{PS}$. If R is a τ -CURB set, then $\Theta(R)$ is asymptotically stable under $\dot{x} \in \tau(x) - x$.*

Proof: Let U be a neighborhood of $\Theta(R)$. For U sufficiently small, it follows from the definition of a τ -CURB set and the upper hemi-continuity of τ that, for any $x \in U$ and any $i \in I$, $s_i \in \tau_i(x)$ implies $s_i \in R_i$. Hence, for any $x \in U$, we must have $\dot{x}_{is_i} = -x_{is_i}$ for all $i \in I$ and $s_i \notin R_i$. But then it must follow that $\|x(t) - \Theta(R)\|_\infty$ shrinks exponentially to zero for all $x(0) \in U$. QED

A corollary to Lemma 7 is the following:

Corollary 2 *Let $\Gamma \in \mathcal{G}^*$. A robust equilibrium point (Okada (1983)) is asymptotically stable under the refined best reply dynamics (3).*

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.

Note also that a game could well have asymptotically stable sets under $\dot{x} \in \sigma(x) - x$ that are proper subsets of persistent retracts, but are not faces. The unique Nash equilibrium of Matching Pennies is an example. It turns out, however, that there may even be faces that are proper subsets of persistent retracts and yet are asymptotically stable under $\dot{x} \in \sigma(x) - x$.

Game 5: Consider the following 4-player game $\Gamma = (I, S, u)$ with $I = \{1, 2, 3, 4\}$, $S_1 = \{H_1, T_1\}$, $S_2 = \{H_2, T_2\}$, $S_3 = \{D, U\}$, and $S_4 = \{A, B\}$. The utility functions are given as follows.

Players 1 and 2 are playing Matching Pennies and do not care about other players' strategies. That is, their payoffs are given by

	H_2	T_2
H_1	2,0	0,2
T_1	0,2	2,0

for any strategy pair of players 3 and 4. Player 3's strategy U is strictly dominated by D . In other words, $u_3(s_1, s_2, U, s_4) = 0$, while $u_3(s_1, s_2, D, s_4) = 1$ for all $s_1 \in S_1$, $s_2 \in S_2$, and $s_4 \in S_4$. Player 4's payoffs are more interesting. Strategy A provides player 4 a payoff of zero regardless of all other players' strategies; that is, $u_4(s_1, s_2, s_3, A) = 0$ for all $s_1 \in S_1$, $s_2 \in S_2$, and $s_3 \in S_3$.

Player 4's payoffs from strategy B are given as follows: $u_4(T_1, T_2, U, B) = 1$, $u_4(T_1, T_2, D, B) = 0$, $u_4(H_1, T_2, D, B) = -1$, $u_4(T_1, H_2, D, B) = -1$, $u_4(H_1, T_2, U, B) = 0$, $u_4(T_1, H_2, U, B) = 0$, and finally, $u_4(H_1, H_2, D, B) = u_4(H_1, H_2, U, B) = -2$.

Claim 1 *The unique minimal σ -CURB set in this game is the face spanned by the strategy selection $\{H_1, T_1\} \times \{H_2, T_2\} \times \{D\} \times \{A, B\}$.*

Proof: The fact that player 1's part of a minimal σ -CURB set is $\{H_1, T_1\}$, and player 2's part is $\{H_2, T_2\}$ follows immediately from the matching pennies structure of their payoffs.

The fact that player 3's part of a minimal σ -CURB set is $\{D\}$ follows from the fact that D strictly dominates U . Thus, all that remains to be shown is that player 4's part must be $\{A, B\}$.

Obviously A is the unique best reply against (H_1, H_2, D) . Thus, A is included in player 4's part of the minimal σ -CURB set. Furthermore, a short calculation shows that $u_4(x_1, x_2, x_3, B) > u_4(x_1, x_2, x_3, A) = 0$ if and only if $x_3(U) > [x_1(H_1) + x_2(H_2)] / [1 - x_1(H_1)x_2(H_2)]$. The strategy profiles satisfying this condition form an open set, which contains the strategy profile (T_1, T_2, D, A) in its boundary. Therefore, B is also included in player 4's part of the minimal σ -CURB set. QED

Claim 2 *In this game, the face spanned by the strategy selection $\{H_1, T_1\} \times \{H_2, T_2\} \times \{D\} \times \{A\}$ is asymptotically stable under $\dot{x} \in \sigma(x) - x$.*

Proof: Consider a solution, ζ , to $\dot{x} \in \sigma(x) - x$ through some initial point x^0 close to this face. From the matching pennies structure of the payoffs of players 1 and 2, it follows that $\zeta_{1H_1}(t, x^0) \rightarrow \frac{1}{2}$ as $t \rightarrow \infty$. Also, $\zeta_{2H_2}(t, x^0) \rightarrow \frac{1}{2}$ as $t \rightarrow \infty$. Furthermore, we must have $\zeta_{3D}(t, x^0) \rightarrow 1$ as $t \rightarrow \infty$ by the fact that D strictly dominates U . In fact, we must have $\zeta_{3D}(t, x^0) > \zeta_{3D}(t', x^0)$ if $t > t'$.

Note that this game has only one Nash equilibrium $(\frac{1}{2}H_1 + \frac{1}{2}T_1, \frac{1}{2}H_2 + \frac{1}{2}T_2, D, A)$. Given the above observations, it must be true that $\zeta_{4A}(t, x^0) \rightarrow 1$ as $t \rightarrow \infty$ for all $x^0 \in \Theta$.

The only thing left to show is that the face spanned by $\{H_1, T_1\} \times \{H_2, T_2\} \times \{D\} \times \{A\}$ is also Lyapunov stable. In order to show this, we must prove that for any neighborhood V of the face, there is another neighborhood $U \subset V$ such that any solution to $\dot{x} \in \sigma(x) - x$ with initial state in U must stay in V for all $t \geq 0$.

Let V be a neighborhood such that for any $x \in V$, we have $x_3(U) < \epsilon$ for some $\epsilon > 0$. Let $E(\epsilon) \subset \Theta$ denote the ϵ -box around state (T_1, T_2, D, A) , i.e. $x \in E$ if and only if $x_1(H_1) \leq \epsilon$, $x_2(H_2) \leq \epsilon$, $x_3(U) \leq \epsilon$, and $x_4(B) \leq \epsilon$.

Let $x^0 \in V \setminus E(2\epsilon)$, i.e. x^0 is a state outside the 2ϵ -box around (T_1, T_2, D, A) . Given the matching pennies structure of the game between players 1 and 2, it is easy to see that $\zeta(t, x^0) \notin E(\epsilon)$ for any $t \geq 0$. Thus, by the argument in the proof of claim 1, B is never best against $\zeta(t, x^0)$ for any $t \geq 0$, and thus $\zeta(t, x^0)_4(B)$ will shrink to zero as $t \rightarrow \infty$.

For $x^0 \in E(2\epsilon)$, things are different. For some such initial states x^0 , the path $\zeta(t, x^0)$ can go through $E(\epsilon)$ for some time. For part of this time, indeed, player 4's strategy B could indeed be best and could grow. However, there is an upper bound on the time period that $\zeta(t, x^0)$ spends within $E(\epsilon)$. This upper bound depends on ϵ . For ϵ small enough, the direction that $\zeta(t, x^0)$ takes for players 1 and 2 must be towards (T_1, H_2) (irrespective of what players 3 and 4 do). Given that T_2 is thus not best anywhere for player 2 in this ϵ -box around (T_1, T_2, D, A) , this means that T_2 will shrink (or H_2 will grow). In fact, we must have $\zeta(t, x^0)_2(T_2) = x_2^0(T_2)e^{-t}$. In the worst case, we thus have $\zeta(T, x^0) \notin E(\epsilon)$ for all $t \geq T$ if $T = -\ln(1 - \epsilon)$. Therefore, $-\ln(1 - \epsilon)$ is the longest time possible that $\zeta(t, x^0)$ stays within $E(\epsilon)$ for any $x^0 \in E(2\epsilon)$. For at least part and at most all of this time, B could be (uniquely) best for player 4. Thus, $\zeta(t, x^0)_4(B)$ could grow for up to this amount of time. Given $x^0 \in V$ and thus $x_4^0(B) \leq \epsilon$, we

have $\zeta(t, x^0)_4(B) \leq 1 - (1 - \epsilon)e^{\ln(1-\epsilon)} = 1 - (1 - \epsilon)^2 < 2\epsilon$. As $\epsilon \rightarrow 0$, this expression tends to zero as well. This implies that for any neighborhood V of the face spanned by $\{H_1, T_1\} \times \{H_2, T_2\} \times \{D\} \times \{A\}$, there is another neighborhood U such that $x_3(U) \leq \epsilon$ for some $\epsilon > 0$ small enough and for all $x \in V$, such that $\zeta(t, x^0) \in V$ for all $t \geq 0$ and all $x^0 \in U$. QED

Note that $\sigma = \beta$ in this game. Therefore, the converse of Lemma 7 is not true for the best reply dynamics either. The difference in the seemingly discrepant conclusions of Claims 1 and 2 is driven by the following observation: B is the unique best reply for player 4 in a cone-shaped set of opponent strategy profiles with apex (T_1, T_2, D) . Hence, B is included in each σ -CURB set. Still, the face spanned by the smaller strategy selection $\{H_1, T_1\} \times \{H_2, T_2\} \times \{D\} \times \{A\}$ is asymptotically stable: the trajectories starting in this cone leave it quickly; moreover, the closer the starting point is to the (smaller) face, the more quickly the trajectory leaves. Thus, there is a vanishing amount of time in which B can grow.

Lemma 8 *Let $\Gamma \in \mathcal{G}^*$ and let R be a strategy selection. If $\Theta(R)$ is asymptotically stable under the τ -best reply dynamics $\dot{x} \in \tau(x) - x$, for $\tau \in \mathcal{T}^{PS}$, then R is a τ -prep set. Furthermore, if for all strategy selections R' which are proper subsets of R , the face $\Theta(R')$ is not asymptotically stable under $\dot{x} \in \tau(x) - x$ (i.e. $\Theta(R)$ is a minimal asymptotically stable set), then R is a tight τ -prep set.*

Proof: Since $\tau \in \mathcal{T}^{PS}$, then for every $x \in \Theta$, there is a $T_i(x)$ such that $\tau_i(x) = \Delta(T_i(x))$ for all $i \in I$. Now suppose $\Theta(R)$ is asymptotically stable under $\dot{x} \in \tau(x) - x$ and R is not a τ -prep set. Then there is an $x \in \Theta(R)$ and a player $i \in I$ such that $T_i(x) \cap R_i = \emptyset$. But then, by the upper hemi continuity of τ , there is a neighborhood of x such that $T_i(x') \subset T_i(x)$ for all x' in this neighborhood. Thus, a solution to $\dot{x} \in \tau(x) - x$ starting from a state in the interior of this neighborhood must spend a finite amount of time within this neighborhood. In this finite amount of time, all strategies in R_i must shrink. The face $\Theta(R)$ is therefore not asymptotically stable, which is a contradiction, proving the first part of this lemma. To prove the second part, suppose $\Theta(R)$ is asymptotically stable and R is not a tight τ -prep set. Then there is a player $i \in I$ and $s_i \in R_i$ such $s_i \notin T_i(x)$ for any $x \in \Theta(R)$. Let R' be derived from R by letting $R'_i = R_i \setminus \{s_i\}$ and $R'_j = R_j$ for all $j \neq i$. Then $\Theta(R')$ is also asymptotically stable under $\dot{x} \in \tau(x) - x$. QED

Not every τ -prep set is asymptotically stable under $\dot{x} \in \tau(x) - x$. To see this for $\tau = \sigma$, consider the following game.

	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 6: This game is taken from Samuelson (1992).

This game is symmetric. Note that strategies C and F are weakly dominated. The unique persistent retract of this game is spanned by $\{A, B\} \times \{D, E\}$. There are three singleton

σ -prep sets $\{(A, D)\}$, $\{(A, E)\}$, and $\{(B, D)\}$. None of these are asymptotically stable under the refined best reply dynamics. Thus, none of these are MASF. To see this, note that $\dot{x} \in \sigma(x) - x$ has a solution starting at (A, D) that gradually takes play towards (A, E) .

We are finally in a position to prove our main theorem.

Theorem 4 *Let $\Gamma \in \mathcal{G}^*$. If R is a strategy selection such that $\Theta(R)$ is a persistent retract (minimal σ -CURB set), then $\Theta(R)$ contains an MASF. Conversely, if R is a strategy selection such that $\Theta(R)$ is an MASF, then R is a tight σ -prep set.*

Proof: For the first part, suppose that R is a strategy selection such that $\Theta(R)$ is a persistent retract (minimal σ -CURB set). Then the fact that $\Theta(R)$ is asymptotically stable under $\dot{x} \in \sigma(x) - x$ follows from Lemma 6 and Lemma 7. Thus, either $\Theta(R)$ is an MASF or it contains one. The second part follows from Lemma 8. QED

If one is interested in (partially) characterizing not only MASFs but also their asymptotically stable subsets, one approach could be to define an appropriate version of Balkenborg and Schlag's (2007) strict equilibrium sets, based on the refined best reply correspondence. We conjecture that these sets, appropriately defined, are asymptotically stable under $\dot{x} \in \sigma(x) - x$.

6 Conclusion

In this paper we endeavor to find the smallest faces of the set of mixed strategy profiles that can justifiably be called evolutionarily stable. In order to do so, we introduce generalizations of the best reply correspondence that satisfy four reasonable criteria. We define a generalized best reply dynamics as an appropriate differential inclusion based on the respective generalized best reply correspondence. We define a minimal asymptotically stable face (MASF) as a face of the set of mixed strategy profiles that is asymptotically stable under some such generalized, deterministic evolutionary process, and with the additional property that it does not properly contain another face that is also asymptotically stable under some (possibly different) generalized best reply dynamics.

We show (Theorem 1) that the set of all generalized best reply correspondences (and hence, dynamics) is, for most games, a complete lattice with a unique smallest element. We call this smallest element the refined best reply correspondence (and dynamics).

The main results in this paper offer a partial characterization of minimal asymptotically stable faces (MASFs). We show (Theorem 2) that every persistent retract (Kalai and Samet (1984)) is the same as a minimal CURB set (Basu and Weibull (1991)) based, however, on the refined best reply correspondence, and contains an MASF (Theorem 4). An MASF can be a proper subface of a persistent retract, as we show by example, and must be (Theorem 4) a tight prep set (Voorneveld (2004)) based again on the refined best reply correspondence. Thus, MASFs are sets somewhere "in-between" CURB and prep sets based on the refined best reply correspondence.

Our findings are thus reminiscent of Hurkens’s (1995) findings that a stochastic best reply learning process based on semi-robust best replies results in play that eventually leads to a persistent retract. This suggests that while it is difficult to justify Nash equilibrium behavior, either epistemically or through evolution or learning, let alone any of its point-wise refinements or even set-wise refinements such as Kohlberg and Mertens’s (1986) strategically stable sets, there are relatively small sets of strategy profiles that can be justified through learning. Furthermore, every persistent retract, or CURB set based on the refined best reply correspondence, contains such a set. In addition, for a face to be an MASF it is a necessary condition that this face be spanned by a tight prep-set based, again, on the refined best reply correspondence.

We thus suggest that in applied game theory work, MASFs, or at least appropriate variations of CURB sets, persistent retracts, and prep sets, which have not yet been used to a great extent,¹⁸ may be very apt choices for a solution concept.

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

Adapting a notion of Brandenburger and Friedenberg (2007) for perfect information games, we let a normal form game satisfy the Single Payoff Condition (SPC) if all own-payoff equivalent pure strategies are also payoff equivalent.¹⁹ Not every game satisfies the SPC: a player other than i might not be indifferent between player i ’s own-payoff equivalent strategies. This is the case in Game 3. For another example consider the following game between a completely deaf person and her neighbor. If the completely deaf person does not care whether there is a loud radio playing in her apartment or not, while her neighbor does care, the game of how loud the deaf person plays her radio does not satisfy the SPC.

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., p. 147 in Ritzberger (2002)). However, games not satisfying the SPC are exceptional. For instance, the deaf person in the previous example would normally be expected to care at least a little bit about the effect of the loud music on her neighbors. This may be because she wants to get along with them, or, to the contrary, because she likes to annoy them. Either way, any slight such change of her preferences would eliminate the violation of the SPC. It seems intuitive, therefore, that a typical game should satisfy the SPC. For this consideration it does not matter from which class of games, normal form games, extensive form games, repeated games or any other family of games the game is taken, as long as the family does not itself systematically violate the SPC. However, if the notion of ”typical” is formalized mathematically as genericity, being generic in any of these classes mean something very different.

¹⁸Some notable exceptions are Kalai and Samet (1985), Balkenborg (1993), Blume (1994, 1996), Hurkens (1996), van Damme and Hurkens (1996), and Gordon (2006).

¹⁹Battigalli’s (1997) notion of games without relevant ties is closely connected.

Trivially, for generic normal form games, there are neither own-payoff equivalent nor payoff equivalent strategies and hence the SPC holds. 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.²⁰ In this appendix 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. Thus, the restriction to games satisfying the SPC and hence, after the identification of payoff equivalent strategies, to games in \mathcal{G}^* , is not a severe one.

Definition 4 *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 equation*

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

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

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

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

Definition 5 *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*

$$u_i(s, \mu) \tag{5}$$

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 a given $s \in \mathcal{S}$.²¹ 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 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 $s_{-i} \in S_{-i}$.

²⁰For an illuminating discussion on this point, see Mertens (2004).

²¹An analytic function is a function that is locally described by power series. The notion covers most functions arising in applications, and in particular, linear and rational functions or functions like e^x or $\ln(x)$. In our examples the functions are always linear.

Proposition 3 *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

$$u_j(s_i, s_{-i}, \mu) = u_j(s'_i, s_{-i}, \mu) \quad (6)$$

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 (6) for *some* $i \in I$, $s_i, s'_i \in S_i$, $s_{-i} \in S_{-i}$, and $j \in I$ implies the identity (6) for *all* $i \in I$, $s_i, s'_i \in S_i$, $s_{-i} \in S_{-i}$, and $j \in I$ and also for all $\mu \in O$. 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

$$m = (m_1, \dots, m_n) \in M = \times_{i \in I} M_i, \quad (7)$$

they simultaneously and independently choose 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

$$a = (a_1, \dots, a_n) \in A = \times_{i \in I} A_i, \quad (8)$$

of which, 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

$$u_i(s, \mu) = \mu_{a,i}, \quad (9)$$

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

$$u_i(s_i, s_{-i}, \mu) = u_i(s'_i, s_{-i}, \mu) \quad (10)$$

can hold for *all* μ only 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,

$$u_j(s_i, s_{-i}, \mu) = \mu_{a,j} = u_j(s'_i, s_{-i}, \mu) \quad (11)$$

for all j and μ . Thus Proposition 3 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

$$u_i(s, \mu) = \sum_{t \in T} p_t \mu_{t,i}, \quad (12)$$

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

$$u_i(s_i, s_{-i}, \mu) = \sum_{t \in T} p_t \mu_{t,i} = \sum_{t \in T} p'_t \mu_{t,i} = u_i(s'_i, s_{-i}, \mu) \quad (13)$$

can hold for *all* $\mu \in R^{T \times I}$ only 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

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

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 the two strategy combinations $s = (s_i, s_{-i})$ and $s' = (s'_i, s_{-i})$,

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

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

$$u_i(s, \mu) = \sum_{a \in A} k_{s,a} \mu_{a,i} = \sum_{a \in A} k_{s',a} \mu_{a,i} = u_i(s', \mu) \quad (16)$$

Again, the SPC holds generically in repeated games.

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

$$u_i(s_i, s_{-i}, \mu) = u_i(s'_i, s_{-i}, \mu). \quad (17)$$

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.

References

- ALIPRANTIS, C. D., AND K. C. BORDER (1999): *Infinite-dimensional analysis*. Springer, Berlin, 2 edn.
- AUBIN, J.-P., AND A. CELLINA (1984): *Differential Inclusions*. Springer, Berlin.
- BALKENBORG, D. (1992): *The Properties of Persistent Retracts and Related Concepts*. Ph.D. thesis, University of Bonn.
- (1993): “Strictness, evolutionary stability and repeated games with common interests,” CARESS WP 93-20.
- BALKENBORG, D., J. HOFBAUER, AND C. KUZMICS (2009): “The refined best-response correspondence in normal form games,” Unpublished manuscript.
- BALKENBORG, D., M. JANSEN, AND D. VERMEULEN (2001): “Invariance properties of persistent equilibria and related solution concepts,” *Mathematical Social Sciences*, 41(1), 111–130.
- BALKENBORG, D., AND K. SCHLAG (2007): “On the evolutionary selection of sets of Nash equilibria,” *Journal of Economic Theory*, 133(1), 295–315.
- BASU, K., AND J. W. WEIBULL (1991): “Strategy subsets closed under rational behavior,” *Economics Letters*, 36, 141–46.
- BATTIGALLI, P. (1997): “On rationalizability in extensive games,” *Journal of Economic Theory*, 74, 40–61.
- 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.
- BLUME, A. (1994): “Equilibrium refinements in sender-receiver games,” *Journal of Economic Theory*, 64, 66–77.
- (1996): “Neighborhood stability in sender-receiver games,” *Games and Economic Behavior*, 13, 2–25.
- BRANDENBURGER, A., AND A. FRIEDENBERG (2007): “Axioms for Backward Induction,” Unpublished manuscript.
- 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.
- DAVEY, B., AND H. PRIESTLEY (2002): *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 2 edn.
- 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.
- GORDON, S. (2006): “Iteratively stable cheap talk equilibria,” Unpublished manuscript.
- GUNNING, R., AND H. ROSSI (1965): *Analytic Functions of Several Complex Variables*. Prentice Hall, Englewood Cliffs, N.J.
- HART, S., AND A. MAS-COLELL (2003): “Uncoupled dynamics do not lead to Nash equilibrium,” *American Economic Review*, 93, 1830–1836.
- 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.
- HOFBAUER, J., AND J. SWINKELS (1995): “A universal Shapley example,” Unpublished manuscript.
- HURKENS, S. (1995): “Learning by forgetful players,” *Games and Economic Behaviour*, 11, 304–29.
- (1996): “Multi-sided pre-play communication by burning money,” *Journal of Economic Theory*, 69, 186–97.
- KALAI, E., AND D. SAMET (1984): “Persistent equilibria in strategic games,” *International Journal of Game Theory*, 13, 129–44.
- (1985): “Unanimity games and Pareto optimality,” *International Journal of Game Theory*, 14, 41–50.
- KETS, W., AND M. VOORNEFELD (2008): “Learning to be prepared,” *International Journal of Game Theory*, 37, 333–52.
- KOHLBERG, E., AND J.-F. MERTENS (1986): “On the strategic stability of equilibria,” *Econometrica*, 54, 1003–37.

- MATSUI, A. (1992): “Best response dynamics and socially stable strategies,” *Journal of Economic Theory*, 57, 343–62.
- 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, Oxford.
- RITZBERGER, K., AND J. W. WEIBULL (1995): “Evolutionary selection in normal form games,” *Econometrica*, 63, 1371–1399.
- ROTH, G., AND W. SANDHOLM (2011): “Stochastic Approximation with Constant Step Size and Differential Inclusions,” Unpublished manuscript.
- SAMUELSON, L. (1992): “Dominated strategies and common knowledge,” *Games and Economic Behavior*, 4, 284–313.
- SANDHOLM, W. H. (2010): *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, M.A.
- TAYLOR, P. D. (1979): “Evolutionary stable strategies with two types of players,” *Journal of Applied Probability*, 16, 76–83.
- VAN DAMME, E. E. C., AND S. HURKENS (1996): “Commitment robust equilibria and endogenous timing,” *Games and Economic Behaviour*, 15, 290–311.
- VON STENGEL, B., AND S. ZAMIR (2010): “Leadership games with convex strategy sets,” *Games and Economic Behavior*, 69, 446–457.
- VOORNEVELD, M. (2004): “Preparation,” *Games and Economic Behavior*, 48, 403–14.
- WEIBULL, J. W. (1995): *Evolutionary Game Theory*. MIT Press, Cambridge, Mass.
- YOUNG, H. P. (1993): “The evolution of conventions,” *Econometrica*, 61, 57–84.