

The refined best-response correspondence in normal form games

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

February 14, 2014

Abstract

This paper provides an in-depth study of the (most) refined best-response correspondence introduced by Balkenborg, Hofbauer, and Kuzmics (2013). An example demonstrates that this correspondence can be very different from the standard best-response correspondence. In two-player games, however, the refined best-response correspondence of a given game is the same as the best-response correspondence of a slightly modified game. The modified game is derived from the original game by reducing the payoff by a small amount for all pure strategies that are weakly inferior. Weakly inferior strategies, for two-player games, are pure strategies that are either weakly dominated or are equivalent to a proper mixture of pure strategies. Fixed points of the refined best-response correspondence are not equivalent to any known Nash equilibrium refinement. A class of simple communication games demonstrates the usefulness and intuitive appeal of the refined best-response correspondence.

Keywords: best-response correspondence, persistent equilibria, Nash equilibrium refinements, strict and weak dominance, strategic stability, fictitious play
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

[‡]Center for Mathematical Economics, Bielefeld University, ckuzmics@uni-bielefeld.de

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

1 Introduction

The refined best-response correspondence was introduced by Balkenborg, Hofbauer, and Kuzmics (2013) in an effort to find the smallest face of the polyhedron of mixed strategy profiles that can be termed evolutionary stable under some reasonable dynamic learning process. Balkenborg, Hofbauer, and Kuzmics (2013) demonstrate that these faces are such that they are minimally asymptotically stable under a particular “smallest” deterministic dynamical system, which is a differential inclusion based on the, so-termed, (most) refined best-response correspondence. If the best-response dynamics, introduced by Gilboa and Matsui (1991), Matsui (1992), and Hofbauer (1995), can be described as a gradual process in which agents who revise their strategy always switch to a best response, the refined best-response dynamics can be described as a gradual process in which revising agents always switch to a best response that is also a unique best response against a strategy profile arbitrarily close to the current one.

The lattice theorem of Balkenborg, Hofbauer, and Kuzmics (2013, Theorem 1) implies that of all best-response-like dynamics the refined best-response dynamics has the fewest stationary points, as the refined best-response correspondence has the fewest fixed points of all best-response-like correspondences.¹ It furthermore implies that the refined best-response dynamics has the most (asymptotically) stable points (all of course fixed points of the refined best-response correspondence, and thus Nash equilibria) of all best-response-like dynamics as it has the fewest solution trajectories of all best-response-like dynamics. Of all best-response-like dynamics the refined best-response dynamics is the one that makes the most Nash equilibria (yet, only Nash equilibria) stable. The refined best-response dynamics, of all best-response-like dynamics, therefore provides the closest justification, based on learning dynamics, of the general practice of using Nash equilibrium as the solution concept for games, while at the same time it allows us to identify Nash equilibria that can never be made stable under any best-response-like dynamics.

The refined best-response dynamics is a very reasonable dynamic learning process in the spirit of fictitious play. Balkenborg, Hofbauer, and Kuzmics (2013) provide a sketch of a micro-foundation for this dynamics as follows. 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. 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. If these beliefs are sufficiently diverse, only a vanishing fraction of individuals adopt a strategy that is best only on a set of states with Lebesgue measure zero. This gives rise to the refined best-response dynamics.

¹What we here call a best-response-like dynamics or correspondence is what in Balkenborg, Hofbauer, and Kuzmics (2013) is formally defined and termed a generalized best-response dynamics or correspondence.

The objective of this paper is to study properties of the refined best-response correspondence and its fixed points (as they are stationary points of the refined best-response dynamics) in detail for all normal form games that satisfy a mild restriction.

We show by example that the refined best-response correspondence, while, by definition, it shares many properties with the best-response correspondence, such as being upper-hemi continuous, closed- and convex-valued, and having a product structure, is not generally like a best-response correspondence. There are games with a refined best-response correspondence, for which there is no game that has this refined best-response correspondence as its best-response correspondence. Thus, even if you have studied the best-response correspondence for all games, you have not automatically covered all refined best-response correspondences.

The example that demonstrates this fundamental difference between refined best-response and best-response correspondence is a three-player game. For two-player games we show that refined best-response correspondences are like best-response correspondences. For every two-player game with its refined best-response correspondence, there is another such game with a best-response correspondence that coincides with the given refined best-response correspondence. This result is shown by characterizing strategies that are never refined best responses in terms of a local form of weak dominance, which we call weak inferiority. We then characterize weakly inferior strategies for two-player games, as those and only those pure strategies that are either weakly dominated or equivalent to a proper mixture of pure strategies.²

While there are refined best-response correspondences that are unlike any best-response correspondence in three or more player games, we can show that for all generic normal form games, the refined best-response and best-response correspondences coincide, nevertheless. This tells us that the refined best-response correspondence is of interest only in non-generic games. Of course, many games of interest, such as extensive form games, cheap-talk games, allocation games (e.g. auctions), and generally any games with many pure strategies yet only few outcomes (e.g. win, draw, or lose as in many parlour-games), have non-generic reduced normal form representations, in which the refined best-response correspondence would typically not be identical to the best-response correspondence.

We then proceed to partially characterize fixed points of the refined best-response correspondence and show by means of examples that there is no systematic relationship between these fixed points and known refinements of Nash equilibrium. A fixed point of the refined best-response correspondence does not have to be perfect (Selten 1975), persistent (Kalai and Samet 1984), proper (Myerson 1978), or strategically stable (Kohlberg and Mertens 1986). Conversely, a strategy profile that is perfect, persistent, proper, or an element of a strategically stable set need not be a fixed point of the refined best-response correspondence. Finally, we apply the refined best-response correspondence in a class of simple communication games, that perhaps best demonstrates its intuitive appeal and usefulness.

The paper proceeds as follows. Section 2 defines the very general class of games

²A proper mixture of pure strategies places positive weight on at least two pure strategies.

we study and defines the refined best-response correspondence. Section 3 analyzes the differences between best-response and refined best-response correspondences. Section 4 analyzes fixed-points of the best-response correspondence. Section 5 provides a simple direct proof of the statement that every persistent retract (Kalai and Samet 1984) contains a strategically stable set in the sense of Kohlberg and Mertens (1986). Section 6 further illustrates the differences between the refined best-response and best-response correspondences, and in particular the usefulness and intuitive appeal of the refined best-response correspondence, for a class of games of independent economic interest, namely simple games of cheap-talk communication. Section 7 concludes.

2 Preliminaries

Let $\Gamma = (I, S, u)$ be a finite n -player normal form game, where $I = \{1, \dots, n\}$ is the set of players, $S = \times_{i \in I} S_i$ is the set of pure strategy profiles, and $u : S \rightarrow \mathbb{R}^n$ the payoff function³. Let $\Theta_i = \Delta(S_i)$ denote the set of player i 's mixed strategies, and let $\Theta = \times_{i \in I} \Theta_i$ denote the set of all mixed strategy profiles. Let $\text{int}(\Theta) = \{x \in \Theta : x_{is} > 0 \ \forall s \in S_i \ \forall i \in I\}$ denote the set of all completely mixed strategy profiles.

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

As in Balkenborg, Hofbauer, and Kuzmics (2013) we shall restrict attention to games with a normal form in which the complement of the set of mixed-strategy profiles $\Psi = \{x \in \Theta \mid \mathcal{B}(x) \text{ is a singleton}\}$ has Lebesgue measure 0. We denote this class by \mathcal{G}^* . A game in \mathcal{G}^* is, therefore, such that to almost all strategy profiles all players have a unique best response. As shown in Balkenborg, Hofbauer, and Kuzmics (2013) if a game is not in this class \mathcal{G}^* it must be such that at least one player has two (own-payoff) equivalent pure strategies. 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)).

For games in \mathcal{G}^* let $\sigma : \Theta \Rightarrow \Theta$ be the *refined best-response correspondence* as defined in Balkenborg, Hofbauer, and Kuzmics (2013) as follows. For $x \in \Theta$ let the set of pure *refined best responses* be given by

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

Then, again abusing notation slightly, $\sigma_i(x) = \Delta(\mathcal{S}_i(x))$ and $\sigma(x) = \times_{i \in I} \sigma_i(x) \ \forall x \in \Theta$.

For $x \in \Theta$ a strategy $s_i \in \mathcal{S}_i(x)$ is thus a best response against x that is also a best response for an open subset of any neighborhood of x .⁴

³The function u will also denote the expected utility function in the mixed extension of the game Γ .

⁴Balkenborg (1992) calls strategies $s_i \in \mathcal{S}_i(x)$ semi-robust best responses. This is inspired by Okada (1983) who calls a strategy a robust best response to strategy profile x if it is a best response for an open neighborhood of x . One could call a strategy robust if it is a robust best response against some strategy profile. Any pure strategy that is either a robust best response or a semi-robust best response against some strategy profile x is, thus, a robust strategy. Note that, while every strategy profile $x \in \Theta$ has a semi-robust best response for all players, it may not have a robust best response.

3 The difference between the best-response and the refined best-response correspondence

This section defines and discusses notions of strict and weak local dominance (applied globally), that will be useful in understanding the difference between the best-response and the refined best-response correspondences. We term these notions *strict* and *weak inferiority*.⁵ They are such that, naturally, every strictly dominated strategy is strictly inferior, every weakly dominated strategy is weakly inferior, and every strictly inferior strategy is weakly inferior.

Definition 1 Let $\Gamma = (I, S, u) \in \mathcal{G}^*$. A strategy s_i is *strictly inferior* if $s_i \notin \mathcal{B}_i(x)$ for any $x \in \Theta$. A strategy $s_i \in S_i$ is *weakly inferior* if there is no open subset of strategy profiles $U \subset \Theta$ such that $s_i \in \mathcal{B}_i(x)$ for all $x \in U$.

In other words, a strictly inferior strategy is never a best response, whereas a weakly inferior strategy, may be a best response against some strategy profiles, but is never a refined best response. Another equivalent statement is that a weakly inferior strategy w_i is such that if $w_i \in \mathcal{B}_i(x)$ then $\mathcal{B}_i(x)$ is not a singleton. That is, a weakly inferior strategy is never the only best response. Note that every game in \mathcal{G}^* has at least one strategy for each player that is not weakly inferior.

Suppose we now consider a fixed strategy profile $x \in \Theta$ and player i 's best responses to x , given by $\mathcal{B}_i(x)$. We would like to know which of these best responses are also refined best responses at this given strategy profile x (i.e. are in $\mathcal{S}_i(x)$). We must, of course, have that any weakly inferior strategy w_i satisfies $w_i \notin \mathcal{S}_i(x)$. Can there be another pure best response in $\mathcal{B}_i(x)$ that is not in $\mathcal{S}_i(x)$? For two-player games, the answer is “No”. However, for three or more player games, the answer is “Yes.” The crucial difference between two- and more-player games is that for two-player games the set of strategy profiles for which a player is indifferent between two different pure strategies is a hyperplane, while for three- (or more-) player games it is some non-linear hypersurface. This in turn implies that the set of strategy profiles against which a given player's given pure strategy is a best response is a convex set in the two-player case, but may well be a non-convex set in the three-player case. How this difference matters for the refined best-response correspondence is made clear in examples below.

3.1 Two-player games

The following theorem states that in two-player games, not only is a weakly inferior strategy never a refined best response, but also any best response to a given strategy profile $x \in \Theta$ that is not a refined best response must be weakly inferior. In other words, for two-player games, the refined best response correspondence is completely understood even locally once we know all weakly inferior strategies.


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

Theorem 1 Let $\Gamma = (I, S, u) \in \mathcal{G}^*$ be a two-player normal form game. A strategy is a pure refined best response, $s_i \in \mathcal{S}_i(x)$, if and only if it is a best response, $s_i \in \mathcal{B}_i(x)$, and is not weakly inferior.

We now provide a complete characterization of weakly inferior strategies.

Theorem 2 Let $\Gamma = (I, S, u) \in \mathcal{G}^*$ be a two-player normal form game. A pure strategy is weakly inferior if and only if it is weakly dominated or equivalent to a proper mixture of pure strategies.

While the proofs of Theorems 1 and 2 are given in Appendix B, we here provide an intuitive sketch of the argument. The results are, of course, similar to Pearce's (1984) result, also to be found in Myerson (1991, Theorems 1.6 and 1.7), that in two-player games strictly dominated strategies are exactly those strategies that are never best responses, and that weakly dominated strategies are exactly those that are never a best response to a completely mixed strategy. The proofs of Theorems 1 and 2 in Appendix B, however, does not follow the proof given by Pearce (1984), which is based on the minmax theorem for zero-sum games, but on the sketch of the proof based on the separating hyperplane theorem as given, for instance, by Fudenberg and Tirole (1991, pp. 50-52).



Game 1, Figure 1, and Figure 2 (from left to right): A two-player game to illustrate the proof of Theorems 1 and 2. Payoffs are given only for player 1, who chooses the row. The first figure plots the payoff of player 1's pure strategies against all mixed strategies of player 2. The second figure plots player 1's strategies in the space of payoffs against the two opponent strategies. The x-axis is the payoff against pure strategy F , while the y-axis is the payoff against pure strategy G . Dots represent the five pure strategies, while the solid lines represent not strictly dominated payoff-tuples that can be achieved by appropriate mixtures of player 1's pure strategies.

Consider the two-player game given as Game 1 and the two figures given as Figures 1 and 2, which are simple variations of the game and pictures in Fudenberg and Tirole (1991, p. 50-51, Figures 2.2 and 2.3). We shall first explain the reasoning behind Theorem 2. Clearly pure strategies A and B are unique best responses against some opponent

strategies. Thus, both strategies are refined best responses. Refined best responses must be best responses against an open set of opponent mixed strategies. A mixed strategy of player 2 can be identified in the Figure 2 by the orthogonal vector to a downward sloping straight line, such as the two dashed lines. In fact, as there is an open set of straight lines going through point B , there is an open set of opponent strategy profiles against which strategy B is a (unique) best response. This can also be seen in Figure 1.

Now turn to the weakly dominated strategy D . The only downward sloping straight line through point D in Figure 2 that does not properly run through the convex hull of payoff tuples is the line with infinite slope. Infinite slope reflects the fact that in order to make strategy D a best response the opponent must not play strategy G with a positive probability. Thus, the fact that there is no open set of downward sloping lines that go through point D and are tangential to the convex hull of payoff tuples, implies that there is no open set of opponent strategy profiles that makes strategy D a best response. This can also be seen in Figure 1.

Now turn to strategy C , which is equivalent to an equal mix of pure strategies A and B . Note that, just as in the case of weakly dominated strategies, in Figure 2 there is only a single line through point C in the picture that is also tangential to the convex hull of payoff tuples. The difference to weakly dominated strategies is that this single line does not have infinite slope. Yet, rotate the line in any way, while keeping it fixed at point C , and it will properly run through the convex hull of payoff tuples. So also in this case there is no open set of opponent strategy profiles that would make strategy C a best response. This can also be seen in Figure 1.

To understand the reasoning behind Theorem 1 note that every strategy in a two-player game can be best only in a convex set. Figure 2 demonstrates this nicely. Consider player 1's strategy B . The dashed lines can be identified with different mixed opponent strategies against which strategy B is a best response. There is a minimal slope and a maximal slope, such that for all slopes inbetween strategy B is a best response. Thus, the set of opponent (mixed) strategies against which B is a best response is convex. Note that this is so, for all strategies of player 1. Consider now the mixed strategy of player 2 in which she places equal weight on her two pure strategies, denoted by x_2^* . Player 1's strategy C is a best response to x_2^* , but not to any mixed strategy nearby. Is it then possible that strategy C is a best response against some other (mixed) strategy of the opponent? No. If there is a unique mixed strategy of player 2 in the neighborhood of x_2^* against which C is best, then, as best response sets must be convex, strategy C cannot be a best response to any other (mixed) strategy of the opponent. The convexity of the best response sets in two-player games derives from the fact that, in these games, the space of strategy profiles for which a player is indifferent between two pure strategies is a hyperplane. For three-player games, this is typically not the case as we demonstrate in the next subsection. For two-player games, however, we can say even more.

Theorem 3 *Let $\Gamma = (I, S, u) \in \mathcal{G}^*$ be a two-player game with refined best-response correspondence $\sigma(\Gamma)$. Then there is a game $\Gamma' = (I, S, u') \in \mathcal{G}^*$ with payoff function $u' : S \rightarrow \mathbb{R}^2$ such that its best-response correspondence $\beta(\Gamma') \equiv \sigma(\Gamma)$.*

Proof: Let Γ' be such that, for every $i \in I$, every $s_{-i} \in S_{-i}$, and every weakly inferior $w_i \in S_i$, $u'_i(w_i, s_{-i}) = u_i(w_i, s_{-i}) - \delta$ for some $\delta > 0$. Then, for this game Γ' no weakly inferior strategy is ever a best response. Thus, by Theorem 1, $\sigma(\Gamma) \equiv \sigma(\Gamma') \equiv \beta(\Gamma')$. QED

Theorem 3 is useful as it tells us that in two-player games, the structure of fixed points of σ is the same as the structure of Nash equilibria. In particular, it implies that, in two-player games, there are only finitely many components of fixed points of σ . More precisely, applying the results in Jansen, Jurg, and Vermeulen (2002) we have the following.

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

3.2 Games with more than two players

In this subsection we turn to games with three or more players. We demonstrate by example that neither Theorem 1 nor Theorem 3 extend to games with more than two players.⁶ The refined best response correspondence can, in such games, be different from any best response correspondence. Nevertheless, it can be shown that for generic games (in the normal form) the refined best-response correspondence is identical to the best-response correspondence. Note, however, that many games of interest, such as extensive form games, cheap talk games, allocation games (e.g. auctions), and generally any games with many pure strategies yet only few outcomes (e.g. win, draw, or lose as in many parlour-games) are typically not generic in the space of all normal form games.

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

Game 2: A game where the refined best-response correspondence is not the best-response correspondence of a modified game. Payoffs are given only for player 3 who chooses matrix.

⁶Theorem 2 does also not extend to games with more than two players. One direction is, of course, true. That is that any pure strategy that is weakly dominated or equivalent to a proper mixture of pure strategies is weakly inferior. But there may well be additional weakly inferior strategies. The reason is well-known. In three player games an undominated strategy may still be never a best response (as players here always know, or believe, if you will, that opponents cannot correlate their strategy choices). For a textbook example see Ritzberger (2002, Example 5.7).





Figure 3 and Figure 4 (left to right): In Figure 3 the regions where strategies E and F of player 3 are best responses in this game are indicated in the square of strategy profiles of players 1 and 2. The probability with which player 1 chooses B is indicated vertically downwards in the graph while the probability of player 2 choosing D is indicated horizontally. In the shaded area between the two branches of the hyperbola E is the best response for player 3, outside it is F . The lower branch of the hyperbola intersects the square only in the point (B, D) , indicating that F is a best response against (B, D) , but not a refined best response. Figure 4 is the analogue picture to Figure 1 and plots the payoff to player 3 for her two pure strategies, against all mixed strategy profiles of the opponents.

The three-player game, given above as Game 2, which is in \mathcal{G}^* , demonstrates that neither Theorem 1 nor Theorem 3 holds generally for games with more than two players. Here and in the following three-player games, player 1 chooses the row, player 2 the column and player 3 the matrix. In this example we specify the payoffs of player 3 only. As indicated in Figure 3, note that against opponent strategy profiles $(1/2A + 1/2B, C)$, $(A, 1/2C + 1/2D)$, and $(2/3A + 1/3B, 2/3C + 1/3D)$ (among others) both E and F are refined best responses. However, against (A, C) F is the only best response and against (B, D) E is the only refined best response. Nearby the latter strategy profile there is no open set in the square of the opponents' mixed strategy profiles where F is a best response. Thus, strategy F while it is a best response and not weakly inferior is nevertheless not a refined best response. This demonstrates that Theorem 1 does not extend to three or more player games.

Now assume there exists another game with the same strategies for which the best response mapping for player 3 is identical to the *refined* best response correspondence of the given game. This implies that player 3 must remain indifferent between E and F against the strategy profiles $(1/2A + 1/2B, C)$, $(A, 1/2C + 1/2D)$, and $(2/3A + 1/3B, 2/3C + 1/3D)$. Moreover, F must be a best response against (A, C) , but

not against (B, D) . This implies

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

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

The fact that Theorem 3 does not extend to three or more player games is quite remarkable. It implies that, although the refined best-response correspondence σ satisfies many properties that the best-response correspondence satisfies, such as being upper hemicontinuous, closed- and convex-valued, and having a product structure, it is nevertheless, at least in some cases, not like any best-response correspondence. Thus knowing that the best-response correspondence satisfies a certain property does not immediately imply that the refined best-response correspondence does satisfy this property as well.

Nevertheless, and given the above example perhaps a little surprisingly, we can show that in almost all games (whether two players or more) the refined best-response correspondence is equal to the best-response correspondence. Remember, however, that many games of interest in \mathcal{G}^* (derived for instance from an extensive form) are not among these generic normal form games.

Theorem 4 *For generic normal form games a pure strategy is a refined best response if and only if it is a best response (i.e., $s_i \in \mathcal{S}_i(x) \Leftrightarrow s_i \in \mathcal{B}_i(x)$). That is we have $\sigma \equiv \beta$ for generic normal form games.*

Theorem 4 was originally established in Balkenborg (1992). Given that persistent retracts are minimal CURB sets based on the refined best-response correspondence σ , see Balkenborg, Hofbauer, and Kuzmics (2013, Lemma 3), also originally shown in Balkenborg (1992), Theorem 4 implies that generically persistent retracts coincide with minimal CURB sets. This fact has been used by Voorneveld (2005) to show that generically persistent retracts coincide with his minimal prep sets.

While the proof of Theorem 4 can be found in Appendix A we conclude this section with a discussion of the intuition as well as the difficulties behind this result.

Consider first two-player games. Theorem 1 states that the refined best responses to a strategy profile x are all those best responses to x that are not weakly inferior. Let us now highlight another implication of the proof of Theorem 3. If we modify the payoff of any weakly inferior strategy by uniformly subtracting some positive real number, we obtain another game in which the refined best-response correspondence is identical to the best-response correspondence. It is also immediate that if we uniformly add a small real number to all weakly inferior strategies we obtain yet another game in which the

refined best-response correspondence is identical to the best-response correspondence.⁷ Thus, any small uniform subtraction or addition in payoffs to weakly inferior strategies leads to a game in which the refined best-response correspondence is identical to the best-response correspondence. Also we could dispense with the word “uniform”. If payoff reductions/additions are positive but possibly different for different pure strategy profiles of the opponents, again in the new game the refined best-response correspondence is identical to the best-response correspondence. We, thus, have that for any game in which the refined best-response correspondence is different from the best-response correspondence, and for any open set of games around this game, there is a game in which both the refined and original best-response correspondence coincide. Thus, we have that the set of two-player games in which the two correspondences coincide is dense in the space of two-player games. On the other hand, it is also easy to see that for any game in which the refined best-response correspondence coincides with the best-response correspondence, there is an open set of games around this game in which this is still true. This is almost a proof, except that we still have not established that the set of games in which the refined best-response correspondence coincides with the best-response correspondence has not only positive but “full” measure.

The remaining problem is that in the above argument we are always only changing the payoffs of the weakly inferior strategies. A payoff change to a weakly inferior strategy could be “compensated” by a similar payoff change for the other strategies. Yet, it seems it would be “unlikely” that “random” changes to payoffs in a game would lead to a game in which the refined best-response correspondence is different from the best-response correspondence. In other words it remains to be shown that the equation $\beta(\Gamma) = \sigma(\Gamma)$ is satisfied everywhere except on a lower dimensional subset of the space of games. In order to do this we appeal to an implication of Sard’s theorem known as the transversality theorem (see Guillemin and Pollack (1974)).

For three-player games (or games with more players) there is even an additional difficulty. Compare Figures 1 and 4. For two-player games, as can be nicely seen in Figure 1, if a strategy is a best response to some strategy profile x , but not a refined best response, this strategy can not be a refined best response anywhere (this is also the essence of Theorem 1). For three-player games this local property of not being a refined best response does not extend globally as Figure 4 demonstrates. Strategy F is not a refined best response against (B, D) but is a refined best response in, of course, an open set of strategy profiles around (A, C) . Thus, small payoff changes will not only affect whether or not a given strategy is a local best response against some given strategy profile x but may also have a possibly different effect on whether this strategy is a local best response against some other strategy profile, far away from x . It is still easy to see that a small reduction in payoffs to a strategy which is locally not a refined best response to some strategy profile x will lead to a new game, in which this strategy is not even a best response. Yet, in three player games this may come with some additional consequences, which are not necessarily clear. With the knowledge we built so far it still seems possible

⁷Note that the best-response correspondence in the former case is typically not identical to the best-response correspondence in the latter case.

that one could construct an example of a game-form in which small payoff changes, no matter in which direction, somehow always compensate each other in a way that every such perturbed game has some strategy profile in which there is a best response that is not a refined best response. Our proof shows that such examples cannot be constructed.

4 Nash equilibrium versus best-response refinements

This section provides a few results relating fixed points of the refined best-response correspondence to (refinements of) Nash equilibria. Given that every game has a fixed point of the refined best-response correspondence we concentrate our comparison to well-known refinements of Nash equilibrium that also have an existence property.

Proposition 1 *Let Γ be a finite two-player game in \mathcal{G}^* . Let $x \in \Theta$ be a fixed point of the refined best-response correspondence σ . Then $x_{iw_i} = 0$ for every weakly inferior $w_i \in S_i$.*

Proof: Let $x \in \sigma(x)$. By Theorem 1 $w_i \notin S_i(x)$ for any weakly inferior $w_i \in S_i$. But then no $y \in \Theta$ with $y_{iw_i} > 0$ can be in $\sigma(x)$. QED

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

Definition 2 *A (possibly mixed) strategy profile $x \in \Theta$ is a (trembling-hand normal form) perfect (Nash) equilibrium if there is a sequence $\{x_t\}_{t=1}^\infty$ of completely mixed strategy profiles (i.e., each $x_t \in \text{int}(\Theta)$) such that x_t converges to x and $x \in \beta(x_t)$ for all t .*

We obtain the following Proposition.

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

Proof: Pure fixed point of σ are undominated by Proposition 1. An undominated Nash equilibrium of a two-player game is perfect (see e.g. van Damme (1991, Theorem 3.2.2)). QED

In what follows we demonstrate by examples that this is the strongest statement one can make. Proposition 2 is not true if we replace “pure” with “mixed”, “2-player” with “3 or more player”, or “perfect” with “proper”, “KM-stable”, or “persistent”, where proper is defined in Myerson (1978), persistent is defined as a Nash equilibrium of a persistent retract by Kalai and Samet (1984), and KM-stable is (strategically) stable in the sense of Kohlberg and Mertens (1986) (a minimal set satisfying their Property S). See Section 5 for a definition of KM-stable as well as persistent retracts.

To first see that Proposition 2 cannot be generalized to mixed fixed points of the refined best-response correspondence, nor to games with more than two players, consider

the following immediate characterization of fixed points of σ . For $x_i \in \Theta_i$ let $C(x_i) = \{s_i \in S_i | x_{is_i} > 0\}$ denote the carrier (or support) of x_i .

Proposition 3 *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 $\{s_i\} = \mathcal{B}_i(y)$ for all $y \in U^{s_i}$.*

Suppose $x \in \sigma(x)$. Consider player i . Then for all $s_i \in C(x_i)$ let U^{s_i} denote this open set in which s_i is best and let V^{s_i} denote its closure. Now a necessary condition for x to be perfect is that $\bigcap_{i \in I} \bigcap_{s_i \in C(x_i)} V^{s_i} \cap \int(\Theta) \neq \emptyset$.⁸ However, this is not necessarily the case. Consider the two-player Game 3 taken from Hendon, Jacobson, and Sloth (1996). For this game σ and β are identical. The mixed strategy profile $x^* = ((0, 1/2, 1/2); (1/2, 0, 1/2))$ is a Nash equilibrium, hence a fixed point of β , hence of σ , that, as Hendon, Jacobson, and Sloth (1996) point out is not perfect.⁹ In this fixed point of σ player 2 uses his pure strategies D and F only. D is best in the open set $U^D = \{x \in \Theta | x_{1C} > \frac{1}{2}\}$, while F is best in the open set $U^F = \{x \in \Theta | x_{1B} > \frac{1}{2}\}$. There are no bigger open sets with the same property. Yet the intersection of the closure of the two sets contains no interior point (no completely mixed strategy). Hence, x^* is not perfect.

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

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

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

Game 4: A three-player game in which a pure fixed point of σ is not perfect.

The same logic also underlies the fact that Proposition 2 does not extend to games with more than two players. Consider the three-player game given as Game 4. Pure strategy profile (A, C, E) is a Nash equilibrium. Player 1's strategy A is a best response against $x \in \Theta$ if and only if $3x_{2D} \geq x_{3F}$. Player 2's strategy C is a best response against $x \in \Theta$ if and only if $x_{3F} \geq 3x_{1B}$. Player 3's strategy E is a best response against $x \in \Theta$ if and only if $x_{1B} \geq 3x_{2D}$. Thus, each strategy A, C, E is best in an open set of strategy profiles with the closure containing (A, C, E) and (A, C, E) is a fixed point of σ . However, there is not a single strategy profile, except (A, C, E) , against which all three are best at the same time. To see this use the first inequality in the third to obtain $x_{1B} \geq x_{3F}$. Now use this in the second inequality to obtain $x_{3F} \geq 3x_{3F}$, which is only satisfied at $x_{3F} = 0$. But then, by the same inequalities, we must also have $x_{2D} = 0$ and $x_{1B} = 0$. There is, thus, no strategy profile except (A, C, E) itself against which (A, C, E) is a best response for all three players. This in turn implies that (A, C, E) is not perfect.

⁸Note that this condition is, for instance, satisfied, for the mixed equilibrium in the two-player game of matching pennies.

⁹In fact, this can be seen directly from the observation that player 2's strategy $(1/2, 0, 1/2)$ is weakly dominated by strategy E .

	A	B	C
A	1,1	0,0	-9,-9
B	0,0	0,0	-7,-7
C	-9,-9	-7,-7	-7,-7

Game 5: A two-player game in which a pure fixed point of σ is not proper nor KM-stable.

	A	B	C
A	0,0	0,0	1,-1
B	0,0	1,1	0,-1
C	-1,1	-1,0	-1,-1

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

To see that a pure fixed point of the refined best-response correspondence does not have to be proper or KM-stable even in two-player games consider the symmetric two-player Game 5. This game is from Myerson (1978), who uses it to illustrate the difference between perfect and proper equilibrium. Note that strategy profile (B, B) is a Nash equilibrium. Note that strategy B for each player is a best response if and only if the opponent strategy satisfies $x_A \leq 2x_C$. Thus, B is best against an open set of strategy profiles with closure containing (B, B) , and (B, B) is, therefore, a fixed point of the refined best-response correspondence. As Myerson (1978) shows, however, (B, B) is not a proper equilibrium.¹⁰ Note, furthermore, that strategy profile (A, A) (which is also a fixed point of σ) is a strict Nash equilibrium and, thus, a singleton KM-stable set (or strictly perfect). For (B, B) to be in a KM-stable set we would have to have that it is also strictly perfect. Otherwise the minimality requirement of the KM-stability definition would only pick up (A, A) as a KM-stable set. It is, however, easy to see that (B, B) is not strictly perfect.¹¹ Consider trembles (a tremble for a given pure strategy is the minimal probability with which a player must play this pure strategy, see Section 5 for a definition) such that player 2's trembling probability for strategy A is more than twice as large as the trembling probability for strategy C . As C is never a best response for player 2 she will use it only with minimal trembling probability. Thus, she will use strategy A with a probability that exceeds twice that of strategy C . But then player 1's unique best response is A . Any perturbed game with such trembles only has one equilibrium, and that is close to (A, A) . Thus, (B, B) is not in a KM-stable set.

To see that a pure fixed point of the refined best-response correspondence does not have to be persistent even in two-player games consider the symmetric two-player Game 6. Strategy C is strictly dominated. Strategies A and B are both refined best responses against A . Thus, (A, A) is a fixed point of the refined best-response correspondence. However, the unique persistent retract (minimal absorbing retract, see Kalai and Samet (1984)) is the set $\{(B, B)\}$. Thus (B, B) is the only persistent equilibrium.

¹⁰Strategy C yielding the lowest possible payoff must be played with much smaller probability than strategy A in any ϵ -proper equilibrium.

¹¹That (B, B) is not strictly perfect also follows from the fact, shown in Vermeulen (1996), that in 3×3 -games strictly perfect equilibria are proper.

	A	B	C
A	2,2	1,2	1,2
B	2,1	2,2	0,0
C	2,1	0,0	2,2

Game 7: A two-player game in which a pure perfect (and proper and KM-stable) equilibrium is not a fixed point of σ .

	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

Game 8: A three-player game in which a pure persistent equilibrium is not a fixed point of σ .

We now turn to the question whether some version of the converse of Proposition 2 could be true. We first demonstrate that the direct converse is not true and that even a strengthening of the converse in which we replace “perfect” with “proper” or “singleton KM-stable set” is not true.

To see this consider the symmetric two-player Game 7. In this game strategy A is equivalent to the equal mixture of pure strategies B and C . However, A is a best response only on a thin set of mixed-strategy profiles. In fact, A is best against any $x \in \Theta$ in which the opponent uses $x_B = x_C$, the set of which is a thin set. Thus, this game is in \mathcal{G}^* . In this game (A, A) constitutes a perfect equilibrium. In fact every mixed strategy profile $((\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}); (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}))$ is a strictly perfect equilibrium, and hence, constitutes a singleton KM-stable set. But none of these equilibria, except the one with $\alpha = 0$, are fixed points of σ , due to the fact that A is only best on a thin set (it is a weakly inferior strategy).

The following version of a converse can be established.

Proposition 4 *Let Γ be a 2-player game in \mathcal{G}^* . Then every (pure or mixed) persistent equilibrium $x \in \Theta$ is a fixed point of the refined best-response correspondence σ .*

Proof: A persistent equilibrium $x \in \Theta$, by definition, is a Nash equilibrium contained in a persistent retract (or, equivalently contained in a minimal σ -CURB set, as defined in Section 5). Therefore, x must place positive probability only on those pure strategies that are refined best responses against some strategy profile. These pure strategies are, thus, not weakly inferior. By Theorem 1 all these pure strategies must then be a refined best response against x . QED

Proposition 4 cannot be extended to games with more than two players. Consider Game 8, taken from Kalai and Samet (1984). The strategy profile (A, C, E) is persistent (see Kalai and Samet (1984)) but is not a fixed point of σ . To see this note that player 1’s strategy A is never best for nearby strategy profiles. The one pure strategy combination (of players 2 and 3) against which A is better than B is (D, F) which for nearby (to (A, C, E)) strategy profiles will always have lower probability than the outcomes (C, F) and (D, E) against which B is better than A .

Having, thus, established that fixed points of the refined best-response correspondence have little relationship with well-known refinements of Nash equilibrium, we now demonstrate that the two are, however, not completely incompatible either.

Proposition 5 *Let Γ be a game in \mathcal{G}^* . Then there is a fixed point of the refined best-response correspondence σ that is also a proper equilibrium.*

Proof: The proof requires only a small modification of Myerson's (1978) proof of the existence of proper equilibrium. Recall that a proper equilibrium is the limit point of a sequence of ϵ -proper equilibria. For $\epsilon > 0$ an ϵ -proper equilibrium is a completely mixed strategy profile x^ϵ with the property that, for every player $i \in I$ and for all pure strategies $s_i, s'_i \in S_i$, whenever $u_i(s'_i, x_{-i}^\epsilon) < u_i(s_i, x_{-i}^\epsilon)$ then $x_i^\epsilon(s'_i) \leq \epsilon x_i^\epsilon(s_i)$. Here we need to require, in addition, that an ϵ -proper equilibrium have the property that any pure strategy $s_i \in S_i$ that is not a refined best response to x^ϵ receive weight less than ϵ , i.e. $x_i^\epsilon(s_i) \leq \epsilon$. It is straightforward to verify that Myerson's (1978) proof goes through unchanged, with the result that we obtain existence of a proper equilibrium that is also a fixed point of σ . QED

5 σ -CURB sets and strategic stability

Balkenborg, Hofbauer, and Kuzmics (2013) 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. This equivalence allows us to provide a relatively simple proof of the fact that every persistent retract contains a strategically stable set in the sense of Kohlberg and Mertens (1986), also known as KM-stable set.

Jansen, Jurg, and Borm (1994) have shown that persistent retracts contain a KM-stable set for all two-player games. Mertens (1991) showed, for general n -player games, that every persistent retract contains an M-stable set with the corollary that every persistent retract also contains a KM-stable set. The proof is somewhat involved. One of the authors of this paper showed in his PhD thesis, Balkenborg (1992), that every persistent retract contains a strategically stable set in the sense of Hillas (1990). From this it also follows that every persistent retract contains a KM-stable set. Both results are cited, without proof, in van Damme (2002, Theorem 12 (iv)), who also writes that "... it can be shown that each persistent retract contains a stable set of equilibria. (This is easily seen for stability as defined by Kohlberg and Mertens ...)". This "easy proof", however, to the best of our knowledge, has not been written down anywhere. We provide it here.

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

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

Definition 3 *Let $\Gamma = (I, S, u)$ be a normal form game. For $i \in I$ let $\eta_i : S_i \rightarrow \mathbb{R}$ be such*

that

$$\eta_i(s_i) > 0 \quad \forall s_i \in S_i$$

and

$$\sum_{s_i \in S_i} \eta_i(s_i) < 1.$$

Then $\eta = (\eta_1, \dots, \eta_n)$ is a **tremble**. Let $\Theta_i(\eta) = \{x \in \Theta_i | x_i(s_i) \geq \eta_i(s_i) \quad \forall s_i \in S_i\}$. Then $\Gamma(\eta) = (I, \Theta(\eta), u)$ is the η -**perturbed game**.

The following defines property S of Kohlberg and Mertens (1986) for a set of strategy profiles without the requirement of it being a subset of the set of Nash equilibria, before defining Kohlberg and Mertens's (1986) concept of strategic stability.

Definition 4 Let Γ be a finite normal form game. Let $Q \subset \Theta$ be a closed subset of the set of mixed strategy profiles. Q is **prestable** if for all $\epsilon > 0$ there is a $\delta > 0$ such that for all trembles η with $\eta_i(s_i) < \delta$ for all $s_i \in S_i$ and for all $i \in I$ there is a Nash equilibrium, x^η , of the perturbed game $\Gamma(\eta)$ such that $\|x^\eta - x\| < \epsilon$ for some $x \in Q$. Such a set Q is **KM-stable** if it is prestable and does not properly contain another prestable set.

Note that minimality requires that a KM-stable set consists exclusively of perfect equilibria.

Proposition 6 Let $\Gamma = (I, S, u)$ be a normal form game. Every σ -CURB set of Γ contains a KM-stable set.

Proof: It is sufficient to show that a σ -CURB set is prestable. Let R be a σ -CURB set. Fix a tremble η and the associated perturbed game $\Gamma(\eta)$ with the set of restricted strategy profiles $\Theta(\eta)$. Define $\Theta^*(R) = \{x \in \Theta(\eta) | x_{is} = \eta_{is} \text{ if } s \notin R_i\}$, a compact and convex subset of $\Theta(\eta)$. For $x \in \Theta^*(R)$ let $\sigma^*(x) = \{y \in \Theta(\eta) | y_{is} = \eta_{is} \text{ if } s \notin S_i(x)\}$. Thus, σ^* is an upper hemi-continuous correspondence defined on a convex compact set. By Kakutani's fixed point theorem σ^* has a fixed point. By the assumption that R is a σ -CURB set and the fact that σ is upper hemi-continuous, we have that there is a neighborhood U of $\Theta(R)$ such that $\sigma(U) \subset \Theta(R)$. Thus, as long as η is sufficiently close to zero, such that $\Theta^*(R) \subset U$, this fixed point of σ^* is a Nash equilibrium of the perturbed game. Thus, every sufficiently close perturbed game has a Nash equilibrium close to the σ -CURB set. QED

Given the interpretation of Balkenborg, Hofbauer, and Kuzmics (2013) that σ -CURB sets are asymptotically stable sets under the refined best-response dynamics, this result is reminiscent to the result by Swinkels (1993) that every convex asymptotically stable set of states under some reasonable deterministic dynamics, in which every Nash equilibrium is stationary, contains a hyper-stable set. Of course, not every Nash equilibrium is stationary under the refined best-response dynamics.

6 A cheap-talk example

This section presents a very simple, perhaps the simplest, class of cheap-talk games, or sender-receiver games.¹² All games (in particular also generic games) within this class are non-generic in the space of normal form games. We here show that the use of the refined best reply correspondence greatly simplifies and clarifies the analysis for this class of games.

Suppose there are two states of the world a and b . State a realizes with probability $\rho \in (0, 1)$. Player 1 (the sender) is informed about the state of the world, player 2 (the receiver) is not. Player 1 can, in each state, send one of two messages m or n . Player 2 upon observing the message sent must choose one of two actions A or B . Thus, both players have four pure strategies. Player 1's strategy space is $S_1 = \{(m_a, m_b), (m_a, n_b), (n_a, m_b), (n_a, n_b)\}$, where strategy (m_a, n_b) , for instance, stands for "send message m in state a and message n in state b ". Player 2's strategy space is $S_2 = \{(A_m, A_n), (A_m, B_n), (B_m, A_n), (B_m, B_n)\}$, where strategy (A_m, B_n) , for instance, stands for "choose action A when message m is received and action B when message n is received".

There are only four possible outcomes: action A is chosen when the state is a , action A is chosen when the state is b , action B is chosen when the state is a , and action B is chosen when the state is b . Denote the set of these four outcomes by $X = \{aA, aB, bA, bB\}$. The two players have preferences over these four outcomes given by utility levels u_x^i for all $x \in X$ and $i \in \{1, 2\}$. Let these games be called *simple communication games*. The general payoffs for such games are given in Game 9.

	A_m, A_n	A_m, B_n	B_m, A_n	B_m, B_n
m_a, m_b	$\rho u_{aA}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aA}^i + (1 - \rho) u_{bB}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bB}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bB}^i$
m_a, n_b	$\rho u_{aA}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aA}^i + (1 - \rho) u_{bB}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bB}^i$
n_a, m_b	$\rho u_{aA}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aA}^i + (1 - \rho) u_{bB}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bB}^i$
n_a, n_b	$\rho u_{aA}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bB}^i$	$\rho u_{aA}^i + (1 - \rho) u_{bA}^i$	$\rho u_{aB}^i + (1 - \rho) u_{bB}^i$

Game 9: Payoffs for simple communication games.

Note that for generic choices of payoffs over outcomes in X the simple communication game is in our class \mathcal{G}^* . In what follows we shall assume that for both players $u_{aA}^i \neq u_{aB}^i$ and $u_{bA}^i \neq u_{bB}^i$. Consider first, player 1, the sender. There are only two substantially different cases to be considered. Case 1 is such that player 1 prefers one of the two actions in every state. Without loss of generality this can be action A . Case 2 is such that player 1 prefers different actions in different states. Without loss of generality she could prefer action A in state a and action B in state b .

¹²The games here are much simpler than those of Crawford and Sobel (1982). For a discussion of communication in simple games see Farrell and Rabin (1996). See Gordon (2006) for a discussion of persistent retracts in the cheap-talk games of Crawford and Sobel (1982).

Claim 1 Suppose $u_{aA}^1 > u_{aB}^1$ and $u_{bA}^1 > u_{bB}^1$ (i.e. the sender prefers action A in both states). Then her strategies m_a, n_b and n_a, m_b (i.e. those strategies, in which she conditions her message on the state) are weakly inferior, but not weakly dominated (i.e. are equivalent to a proper mixture of the other two pure strategies).

Claim 2 Suppose $u_{aA}^1 > u_{aB}^1$ and $u_{bA}^1 < u_{bB}^1$ (i.e. the sender prefers different actions in different states). Then her strategies m_a, m_b and n_a, n_b (i.e. those strategies, in which she does not condition her message on the state) are weakly inferior, but not weakly dominated (i.e. are equivalent to a proper mixture of the other two pure strategies).

Proof of Claims 1 and 2: Considering the payoffs in Game 9 it is apparent that all of the sender's strategies yield the same payoff against receiver strategy A_m, A_n (in which the receiver ignores the message anyway and always plays A). Similarly, all of the sender's strategies yield the same payoff against receiver strategy B_m, B_n (in which the receiver ignores the message anyway and always plays B). It is also apparent that all of the sender's strategies yield the same payoff against the receiver's mixed strategy that places equal weight on A_m, B_n and B_m, A_n . Thus, against any mixed strategy of the form $(\alpha, \frac{1}{2} - \alpha, \frac{1}{2} - \alpha, \alpha)$ all of the sender's four pure strategies are equally good. The sender, therefore, does not have any weakly dominated strategies. Now suppose $u_{aA}^1 > u_{aB}^1$ and $u_{bA}^1 > u_{bB}^1$. That is the sender prefers action A in both states. Then strategies m_a, m_b and n_a, n_b provide more "extreme" payoffs than strategies m_a, n_b and n_a, m_b : The highest (lowest) possible payoff under strategies m_a, m_b and n_a, n_b is higher (lower) than the highest (lowest) possible payoff under strategies m_a, n_b and n_a, m_b . From this fact and the symmetry inherent in the game, it follows that strategies m_a, n_b and n_a, m_b are equivalent to a proper mixture of strategies m_a, m_b and n_a, n_b , and, thus, weakly inferior strategies by Theorem 2. In the case that $u_{aA}^1 > u_{aB}^1$ and $u_{bA}^1 < u_{bB}^1$, we have the reverse result that now strategies m_a, n_b and n_a, m_b provide more "extreme" payoffs than strategies m_a, m_b and n_a, n_b . Thus, in this case, strategies m_a, m_b and n_a, n_b are equivalent to a proper mixture of strategies m_a, n_b and n_a, m_b , and, thus, weakly inferior strategies. QED

For player 2, the receiver, the following is true. The proof of these claims is straightforward and omitted.

Claim 3 Suppose $u_{aA}^2 > u_{aB}^2$ and $u_{bA}^2 > u_{bB}^2$ (i.e. the receiver prefers action A in both states). Then her strategy A_m, A_n weakly dominates all other strategies.

Claim 4 Suppose $u_{aA}^2 > u_{aB}^2$ and $u_{bA}^2 < u_{bB}^2$ (i.e. the receiver prefers different actions in different states). Suppose further that $\rho u_{aA}^2 + (1 - \rho)u_{bA}^2 > \rho u_{aB}^2 + (1 - \rho)u_{bB}^2$ (i.e. the receiver with her a-priori information prefers action A over B). Then strategy B_m, B_n is weakly dominated by A_m, A_n . The remaining three strategies are not weakly inferior.

The analogue result holds for the case $\rho u_{aA}^2 + (1 - \rho)u_{bA}^2 < \rho u_{aB}^2 + (1 - \rho)u_{bB}^2$.

Claim 5 Suppose $u_{aA}^2 > u_{aB}^2$ and $u_{bA}^2 < u_{bB}^2$ (i.e. the receiver prefers different actions in different states). Suppose further that $\rho u_{aA}^2 + (1 - \rho)u_{bA}^2 = \rho u_{aB}^2 + (1 - \rho)u_{bB}^2$ (i.e.

the receiver with her a-priori information is indifferent between actions A and B). Then strategies B_m, B_n and A_m, A_n are weakly inferior, but not weakly dominated (i.e. are equivalent to a proper mixture of the other two pure strategies).

The first two claims can be summarized as follows. If the sender always wants the same action implemented, regardless of the state, then the sender finds those pure strategies of hers in which she conditions her message on the state weakly inferior (or never a refined best response). This means, even without thinking about what the receiver does if the sender avoids weakly inferior strategies she will not even contemplate sending different messages in different states. Of course, and now depending on the receiver's preferences, she might want to randomize between which message she sends.

If the sender, on the other hand, would like to see different actions implemented in different states, and if she avoids weakly inferior strategies, she will only consider strategies in which she conditions her message on the state and will disregard those strategies of hers that do not reveal any information in the first place. Again, this does not depend on the receiver's preferences. Of course, if the receiver also wants to choose different actions in different states, but the opposite action than the sender prefers, then the sender may randomize between her non-weakly inferior actions, in order to confuse the receiver.

Simple communication games are, thus, very intuitively and simply solvable using the refined best-response correspondence. Using refinements of Nash equilibrium will typically not do very much in these games. To perhaps see this best consider the special case of what is essentially a coordination (or common interest) game, given as Game 10, for which $\rho = \frac{1}{2}$ and $u_{AA}^i = u_{BB}^i = 1$ and $u_{AB}^i = u_{BA}^i = -1$ for both $i \in \{1, 2\}$.

	A_m, A_n	A_m, B_n	B_m, A_n	B_m, B_n
m_A, m_B	0,0	0,0	0,0	0,0
m_A, n_B	0,0	1,1	-1,-1	0,0
n_A, m_B	0,0	-1,-1	1,1	0,0
n_A, n_B	0,0	0,0	0,0	0,0

Game 10: A simple communication game with common interest.

Note that Game 10 is a symmetric game and such that Claims 2 and 5 apply.¹³ There are no weakly dominated strategies. Yet, the sender's pure strategies (m_A, n_B) and (n_A, m_B) are (in fact unique) best responses against appropriate pure strategies of the opponent, (A_m, B_n) and (B_m, A_n) , respectively. The sender's pure strategies (m_A, m_B) and (n_A, n_B) are best responses against any proper mixture of all opponent strategies, in which pure strategies (m_A, n_B) and (n_A, m_B) receive equal weight. By symmetry, the same arguments apply to the receiver's strategies. Thus, any strategy profile $(x_1, x_2, x_3, x_4), (y_1, y_2, y_3, y_4)$ with $x_2 = x_3$ and $y_2 = y_3$ is a Nash equilibrium of this game. Every Nash equilibrium of this sort is a singleton strategically stable set in the sense of

¹³Note that even though players in this game have own-payoff equivalent pure strategies, this game is in the class \mathcal{G}^* .

Kohlberg and Mertens (1986). To see this note that any completely mixed Nash equilibrium is always a singleton KM-stable set as such an equilibrium is also an equilibrium of a sufficiently slightly perturbed game. To see that even $(1, 0, 0, 0)$ for both players is a KM-stable set, note that arbitrarily close to it there is a completely mixed Nash equilibrium: for instance, when both players choose $(1 - 3\epsilon, \epsilon, \epsilon, \epsilon)$. In addition to this continuum of Nash equilibria, there are two additional ones: $(0, 1, 0, 0)$ for both players and $(0, 0, 1, 0)$ for both players. These are also singleton KM-stable sets.

Yet, the sender's pure strategies (m_A, m_B) and (n_A, n_B) are both equivalent to (each other and to) an even mixture of pure strategies (m_A, n_B) and (n_A, m_B) . Thus, the sender's pure strategies (m_A, m_B) and (n_A, n_B) are weakly inferior and never refined best responses. This game has only three fixed points of the refined best-response correspondence: the two pure informative equilibria $(m_A, n_B), (A_m, B_n)$ and $(n_A, m_B), (B_m, A_n)$ and the mixed equilibrium $(0, \frac{1}{2}, \frac{1}{2}, 0)$ for both players. Thus, the refined best-response correspondence by removing weakly inferior strategies turns this game into the simple coordination it essentially is. This, in turn, greatly simplifies the analysis. Instead of a continuum of singleton KM-stable sets or a continuum of proper equilibria we just have three fixed points of the refined best-response correspondence, one of which is unstable under the refined best-response dynamics.

7 Conclusion

We studied the refined best-response correspondence in normal form games as introduced by Balkenborg, Hofbauer, and Kuzmics (2013) as the basis for a dynamic learning model. We show by example that the refined best-response correspondence can be unlike any best-response correspondence. In two-player games, however, the refined best-response correspondence coincides with the best-response correspondence of a slightly modified game. The modification is such that all pure weakly inferior strategies, as we define them, receive a uniform payoff reduction. In two-player games we show that pure weakly inferior strategies are those and only those strategies that are either weakly dominated or equivalent to a proper mixture of pure strategies. While in general n -player games, we cannot provide such a simple characterization, we show that for generic normal form games refined best-response and best-response correspondences coincide. Of course, many interesting games, such as cheap talk games or extensive form games, are non-generic in the space of all normal form games.

The fixed points of the refined best-response correspondence are the stationary points of the refined best-response dynamics of Balkenborg, Hofbauer, and Kuzmics (2013). They are therefore the only candidates for convergence points of this dynamic process as well as the only candidates for (Lyapunov or asymptotically) stable points under this dynamic process. We show by examples that the set of fixed points of the refined best-response correspondence is neither a subset nor a superset of the set of perfect equilibria (Selten 1975), proper equilibria (Myerson 1978), persistent equilibria (Kalai and Samet 1984), or strategically stable equilibria (Kohlberg and Mertens 1986).

We demonstrated the usefulness and intuitive appeal of the refined best-response correspondence over the best-response correspondence in a simple class of communication games.

There are still many open questions. We have, for instance, refrained in this paper from discussing refined rationalizable strategies. That is, strategies which do survive the iterated elimination of never refined best responses. These would be of interest, as the refined best-response dynamics converges to the set of refined rationalizable strategies in every game (Balkenborg, Hofbauer, and Kuzmics 2013). It is fairly easy to see that the set of refined rationalizable strategies must be a sometimes proper subset of the set of strategies which survive the $S^\infty W$ -procedure of one round of elimination of all weakly dominated strategies and then the iterated elimination of strictly dominated strategies.¹⁴ This is true, for instance, when a game has strategies that are weakly inferior but not weakly dominated. On the other hand iterated admissibility, for which an epistemic derivation has been given by Brandenburger, Friedenberg, and Keisler (2008), is sometimes a subset and sometimes a superset of the set of refined rationalizable strategies. We would find it of interest, to understand better the differences between the various concepts of rationalizability and especially the reasons behind these differences.

Taking our class of simple communication games as a starting point we would also find it of interest to study other classes of games, in which the set of outcomes is much smaller than the set of strategy profiles. We believe that the study of the refined best-response correspondence could be fruitful in many such cases. One example of such a class is the class of extensive form games. Another is the class of communication games with more states and strategies. These are topics we endeavor to address in future work.

A On the generic equivalence of best responses and refined best responses

This appendix provides a proof of Theorem 4, which is organized in a number of steps: We will first fix some notations for the mappings and various submanifolds to be considered. Step 1 argues that the embedding of the uncorrelated strategy combinations into the set of beliefs has nice differentiability properties. Step 2 invokes the transversality theorem (see Guillemin and Pollack (1974)) to show that for generic payoff functions we obtain the needed transversality conditions.¹⁵ Step 3 argues that we can restrict attention to completely mixed strategy combinations of the opponents. If the player is indifferent between several of his strategies against a given completely mixed strategy combination, step 4 shows how we can construct an arbitrarily nearby strategy combination, against which the player strictly prefers a given one among these strategies. Step 5 completes the

¹⁴For this procedure see e.g. Dekel and Fudenberg (1990), Börgers (1994), and Ben Porath (1997).

¹⁵This transversality theorem is a straightforward consequence of Sard's theorem. If one assumes an algebraic map and uses in its proof in Guillemin and Pollack (1974) the algebraic version of Sard's theorem in Bochnak, Coste, and Roy (1998) one obtains a stronger version of the transversality theorem where the conclusion of the theorem holds outside a lower dimensional semi-algebraic set.

argument.

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

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

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

The multilinear function q_i describes the first step in the computation mentioned above. While $x_{-i} \in \prod_{j \neq i} \mathbb{R}^{S_j}$ denotes the usual strategy combinations of the opponents, we use $\chi_{-i} \in \mathbb{R}^{S_{-i}}$ to describe a “correlated strategy of the opponents”, i.e., a belief over the set S_{-i} of pure strategy combinations of the opponents. q_i maps mixed strategy combinations to such beliefs.

Let L_i be the vector space of all linear mappings

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

If $\chi_{-i} \in \mathbb{R}^{S_{-i}}$ is a belief of player i and $s_i \in S_i$ a pure strategy player i chooses, then $(v_i(\chi_{-i}))(s_i)$ is the payoff player i expects with his strategy choice. In this context a vector $z \in \mathbb{R}^{S_i}$ represents the various gains a player could make, not probabilities. The linear function v_i describes for every s_i the second step in the computation of the expected payoff. Any $v_i \in L_i$ corresponds via $u_i(s_{-i}, s_i) = (v_i(s_{-i}))(s_i)$ uniquely to a payoff function

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

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

For $T_i \subseteq S_i$ set $Z_i(T_i) = \{z \in \mathbb{R}^{S_i} \mid \forall s_i, t_i \in T_i : z(s_i) = z(t_i)\}$. $v_i(\chi_{-i}) \in Z_i(T_i)$ means that player i is indifferent between all his strategies in T_i when his belief is χ_{-i} . Let $X_j := \{x_j \in \mathbb{R}^{S_j} \mid \sum_{s_j \in S_j} x_j(s_j) = 1\}$ be the affine space generated by player j 's strategy simplex for $j \neq i$ and let $X_{-i} := \prod_{j \neq i} X_j$.

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

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

and

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

The sets $\Theta_{-i} \cap X_{-i}(T_{-i})$ describe the various faces of the polyhedron Θ_{-i} . The strategies of player j with support in T_j have $X_j(T_j)$ as their affine hull.

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

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

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

is well defined, is one-to-one, maps $X_{-i}(T_{-i})$ onto a closed set, and has a derivative $dq_i|_{x_{-i}}$ with maximal rank everywhere.¹⁶

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

Proof: The family of linear maps L_i defines a mapping

$$V_i : L_i \times \mathbb{R}^{S_{-i}} \rightarrow \mathbb{R}^{S_i} \quad (1)$$

$$(v_i, \chi_{-i}) \mapsto v_i(\chi_{-i}) . \quad (2)$$

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

$$dV_i|_{(v_i, \chi_{-i})} : T_{v_i}L_i \times T_{\chi_{-i}}\mathbb{R}^{S_{-i}} \cong L_i \times \mathbb{R}^{S_{-i}} \rightarrow \mathbb{R}^{S_i} \quad (3)$$

$$(\nu_i, \xi_{-i}) \mapsto \nu_i(\chi_{-i}) + v_i(\xi_{-i}) . \quad (4)$$

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

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

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

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

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

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

Proof: Suppose x_{-i} is in the boundary of Θ_{-i} . For each $j \neq i$ define $T_j := \{s_j \in S_j \mid x_j(s_j) \neq 0\}$. Thus $x_{-i} \in X(T_{-i}) \cap \Theta_{-i}$. If $T_j = S_j$, x_j is in the relative interior of Θ_j . By assumption $T_j \neq S_j$ for at least one opponent j . Fix $j^* \neq i$ with $T_{j^*} \neq S_{j^*}$ and $t_{j^*} \notin T_{j^*}$. Set $\tilde{T}_j := T_j$ for $i \neq j \neq j^*$ and $\tilde{T}_{j^*} := T_{j^*} \cup \{t_{j^*}\}$. We show that $O_{-i} \cap Y(T_i)$ contains some $y_{-i} \in \Theta_{-i} \cap X(\tilde{T}_{-i})$ such that $\tilde{T}_j = \{s_j \in S_j \mid y_j(s_j) \neq 0\}$ for all $j \neq i$. In other words: y_{-i} is in the relative interior of the face $\Theta_{-i} \cap X(\tilde{T}_{-i})$. The claim then follows by induction.

¹⁶The result is well known, e.g., in algebraic geometry: q_i defines the so-called Segre-embedding. The result is needed in algebraic geometry to show that the product of projective spaces can itself be embedded into a projective space, i.e., is projective-algebraic.

The transversality conditions imply that the submanifolds $X_{-i}(T_{-i})$ and $Y(T_i) \cap X_{-i}(\tilde{T}_{-i})$ meet transversally in $X_{-i}(\tilde{T}_{-i})$ (see (Guillemin and Pollack 1974, Exercise 2.3.7)). Since $X_{-i}(T_{-i})$ has codimension 1 in $X_{-i}(\tilde{T}_{-i})$ it follows with arguments as in the next step that $X_{-i}(\tilde{T}_{-i}) \cap Y(T_i) \cap \{y_{-i} \mid y_{j*}(t_{j*}) > 0\} \cap O_{-i}$ intersects the relative interior of $X_{-i}(\tilde{T}_{-i}) \cap \Theta_{-i}$.

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

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

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

$$\lambda : Y_i(T_i \setminus \{s_i\}) \cap X_{-i} \rightarrow \mathbb{R} \tag{5}$$

$$y_{-i} \rightarrow (v_i \circ q_i)(y_{-i})(s_i) - (v_i \circ q_i)(y_{-i})(t_i) \tag{6}$$

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

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

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

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

B Refined best responses in two-player games

This appendix provides proofs of Theorems 1 and 2. In the case of two player games the payoff function is linear in the mixed strategy choice of the opponent. This allows the use of convex analysis (see Rockafellar (1970)) to study the best-response correspondence of a player. The most direct consequence is the convexity of the region where a strategy is

a best response. From this Theorem 1 follows immediately, the arguments are given after Lemma 1 below. More work is needed to obtain Theorem 2. We use conjugate functions and provide the proof after Lemma 2.

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

$$x_2 = (x_2^1, x_2^2, \dots, x_2^{K-1}) \in \mathbb{R}^{K-1} \quad (7)$$

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

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

$$f(x_2) = \begin{cases} \max_{s_1 \in S_1} u_1(s_1, x_2) & \text{for } x_2 \in \Theta_2 \\ +\infty & \text{else} \end{cases} \quad (8)$$

Because u_1 is linear in x_2 , f is, in the terminology of Rockafellar (1970) a proper convex polyhedral function. A key idea explored in the following is that the strategies of player 1 that are refined best responses, correspond to the maximal compact faces of the epigraph

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

of f , which is a convex (but not compact) polyhedron. Duality theory allows us to identify these faces with the extreme points of the epigraph F^* of the conjugate function f^* of f . This will be used in the proof of theorem Theorem 2.

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

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

$$G(x_1) = \{(x_2, \alpha) \in \Theta_2 \times \mathbb{R} \mid x_1 \in \beta(x_2) \text{ and } \alpha = u_1(x_1, x_2)\} \quad (9)$$

and the set $H(x_1) = \{x_2 \in \Theta_2 \mid x_1 \in \beta(x_2)\}$, the projection of $G(x_1)$ onto Θ_2 . $H(x_1)$ is the region where x_1 is a best response.

Lemma 1 *The region $H(x_1)$ is a convex polyhedron.*

Proof: $G(x_1)$ is a face of the epigraph $\{(x_2, \alpha) \in \Theta_2 \times \mathbb{R} \mid f(x_2) \leq \alpha\}$ of the function f , which is a convex polyhedron. $G(x_1)$ is hence a convex polyhedron. $H(x_1)$ is the image of the convex polyhedron $G(x_1)$ under the linear projection mapping and hence also a convex polyhedron. QED

Clearly x_1 is *not weakly inferior* if and only if the convex polyhedron $H(x_1)$ has non-empty interior $H^\circ(x_1)$. Moreover, $H(x_1)$ is the closure of $H^\circ(x_1)$ if $H^\circ(x_1)$ is not empty. Therefore, if x_1 is not weakly inferior and a best response against x_2 , then x_2 is in the closure of the open set $H^\circ(x_1)$ and so x_1 is a refined best response against x_2 . Given Definition 1 this implies immediately Theorem 1.

The remainder of this section aims at proving Theorem 2. We consider again the epigraph F of the map f defined above. We notice that F is a polyhedral convex set whose compact faces are precisely the sets $G(x_1)$ with $x_1 \in \Theta_1$. The non-compact faces are of the form $F \cap (\Theta'_1 \times \mathbb{R})$, where Θ'_1 is a face of Θ_1 .

The *conjugate function* $f^* : \mathbb{R}^{K-1} \rightarrow \mathbb{R}$ of f is defined by

$$f^*(x_2^*) = \sup_{x_2 \in \mathbb{R}^{K-1}} \{x_2^* \bullet x_2 - f(x_2)\} = \max_{x_2 \in \Theta_2} \{x_2^* \bullet x_2 - f(x_2)\} < \infty, \quad (10)$$

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

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

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

$$a(x_2) = -\alpha + \sum_{k=1}^{K-1} x_2^{*k} x_2^k \quad (11)$$

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

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

Lemma 2 F^* is a polyhedral convex set generated by extreme points x_1 that are refined best responses in Θ_1 and the directions

$$-e_k = (-e_k^1, \dots, -e_k^K) \in \mathbb{R}^K \text{ with } e_k^l = \begin{cases} -1 & \text{for } k=l \\ 0 & \text{else} \end{cases} \quad (12)$$

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

$$e = (1, \dots, 1) \in \mathbb{R}^K \quad (13)$$

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

be directions in F^* . For every combination $v = -\sum_{k=1}^{K-1} \rho_k e_k + \rho_K e$ with $\rho_1, \dots, \rho_K \geq 0$, every $\lambda \geq 0$, every $(x_2^*, \alpha^*) \in F^*$ and every $x_2 \in \Theta_2$ we have conversely

$$\alpha^* + \lambda \rho_K \geq x_2' x_2 - \sum_{k=1}^{K-1} \lambda \rho_k x_2^k + \lambda \rho_K - f(x_2) \quad (14)$$

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

We have characterized the directions of F^* and must now determine the extremal points of F^* . Suppose $(\hat{x}_2^*, \hat{\alpha}^*)$ is an extremal point. Because F^* has only finitely many extremal points, these are exposed points by Straszewick's theorem (Theorem 18.6 in Rockafellar (1970)). Therefore we can find $x_2 \in \Theta_2$ such that the hyperplane $\{x_2^* \bullet x_2 = f(x_2)\}$ is a supporting hyperplane that meets F^* only in $(\hat{x}_2^*, f^*(\hat{x}_2^*))$. Because F^* has only finitely many extreme points and directions there exists an open neighborhood U of x_2 in Θ_2 for which the hyperplanes $\{x_2^* \bullet y_2 = f(y_2)\}$ are for all $y_2 \in U$ supporting hyperplanes that intersect F^* only in $(\hat{x}_2^*, f^*(\hat{x}_2^*))$. This implies that the graph of the affine function $(\hat{x}_2^*, f^*(\hat{x}_2^*))$ intersects F in a $K-1$ dimensional face. It is therefore identical to an affine function defined by a strategy x_1 in Θ_1 for which $H(x_1)$ is full dimensional. Given our identification, $(\hat{x}_2^*, f^*(\hat{x}_2^*))$ is consequently a not weakly inferior strategy in Θ_1 , which was to be shown. QED

Proof of Theorem 2:

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

However, no points on the compact faces of F^* apart from the extreme points are not weakly inferior strategies. To see this, notice that a proper mixture $x_1 = \sum_{l=1}^L \rho_l x_{1l}$ ($L > 2, \rho_l > 0, \sum_{l=1}^L \rho_l = 1$) of non-equivalent not weakly inferior strategies in Θ_1 is weakly inferior. Otherwise there would be an open set in Θ_2 on which x_1 and hence all strategies x_{1k} were best responses. They would yield identical payoffs on an open set and were hence (by Kalai and Samet (1984, Lemma 4)) all own-payoff equivalent, contradicting the assumption. Per construction such a mixture is own-payoff equivalent to a proper mixture of strategies that are pairwise not own-payoff equivalent.

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

$$u_1(x_1', x_2) = u_1(x_1, x_2) - \sum_{k=1}^{K-1} \rho_k x_2^k - \rho_K \left(1 - \sum_{k=1}^{K-1} x_2^k\right) \leq u_1(x_1, x_2), \quad (15)$$

where this inequality holds as a strict one for the k -th pure strategy of player 2 whenever $\rho_k > 0$. Thus x_1' is weakly dominated. It is a weakly inferior strategy because it is a best response only on a proper face of Θ_1 (see Pearce (1984)).

In summary, the only robust strategies in Θ_1 are the extreme points of F^* . All other strategies are proper mixtures of not own-payoff equivalent not weakly inferior strategies or are weakly dominated and therefore weakly inferior. QED

References

- BALKENBORG, D. (1992): *The Properties of Persistent Retracts and Related Concepts*. Ph.D. thesis, University of Bonn.
- BALKENBORG, D., J. HOFBAUER, AND C. KUZMICS (2013): “Refined best-response correspondence and dynamics,” *Theoretical Economics*, 8, 165–192.
- BASU, K., AND J. W. WEIBULL (1991): “Strategy subsets closed under rational behavior,” *Economics Letters*, 36, 141–46.
- BEN PORATH, E. (1997): “Rationality, Nash equilibrium and backwards induction in perfect information games,” *Review of Economic Studies*, 64, 23–46.
- BOCHNAK, J., M. COSTE, AND M. ROY (1998): *Real Algebraic Geometry*. Springer.
- BÖRGERS, T. (1994): “Weak dominance and approximate common knowledge,” *Journal of Economic Theory*, 64, 265–276.
- BRANDENBURGER, A., A. FRIEDENBERG, AND H. J. KEISLER (2008): “Admissibility in Games,” *Econometrica*, 76, 307–353.
- CRAWFORD, V., AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50, 1431–1451.
- DEKEL, E., AND D. FUDENBERG (1990): “Rational behavior with payoff uncertainty,” *Journal of Economic Theory*, 52, 243–67.
- FARRELL, J., AND M. RABIN (1996): “Cheap Talk,” *Journal of Economic Perspectives*, 10:3, 103–118.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press, Cambridge, MA.
- GILBOA, I., AND A. MATSUI (1991): “Social stability and equilibrium,” *Econometrica*, 59, 859–67.
- GORDON, S. (2006): “Iteratively stable cheap talk equilibria,” Unpublished manuscript.
- GUILLEMIN, V., AND A. POLLACK (1974): *Differential Topology*. Prentice Hall, Englewood Cliffs, N.J.
- HARSANYI, J., AND R. SELTEN (1988): *A general theory of equilibrium selection in games*. MIT Press, Cambridge, Massachusetts.
- HENDON, E., H. J. JACOBSON, AND B. SLOTH (1996): “Fictitious play in extensive form games,” *Games and Economic Behavior*, 15, 177–202.
- HILLAS, J. (1990): “On the Definition of the Strategic Stability of Equilibria,” *Econometrica*, 58(6), 1365–1390.

- HOFBAUER, J. (1995): “Stability for the best response dynamics,” Unpublished manuscript.
- JANSEN, M., A. JURG, AND P. BORM (1994): “On strictly perfect sets,” *Games and Economic Behavior*, 6, 400–415.
- JANSEN, M., P. JURG, AND D. VERMEULEN (2002): “On the set of equilibria of a bimatrix game: A survey,” in *Chapters in Game Theory*, ed. by P. Borm, and H. Peters, pp. 121–142. Kluwer, Amsterdam.
- KALAI, E., AND D. SAMET (1984): “Persistent equilibria in strategic games,” *International Journal of Game Theory*, 13, 129–44.
- KOHLBERG, E., AND J.-F. MERTENS (1986): “On the strategic stability of equilibria,” *Econometrica*, 54, 1003–37.
- MATSUI, A. (1992): “Best response dynamics and socially stable strategies,” *Journal of Economic Theory*, 57, 343–62.
- MERTENS, J.-F. (1991): “Stable Equilibria: A Reformulation. Part II. Discussion of the Definition, and Further Results,” *Mathematics of Operations Research*, 16, 694–753.
- MYERSON, R. B. (1978): “Refinements of the Nash equilibrium concept,” *International Journal of Game Theory*, 7, 73–80.
- (1991): *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.
- 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.
- ROCKAFELLAR, R. T. (1970): *Convex Analysis*. Princeton University Press, Princeton, New Jersey.
- SELTEN, R. (1975): “Re-examination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory*, 4, 25–55.
- SWINKELS, J. M. (1993): “Adjustment dynamics and rational play in games,” *Games and Economic Behavior*, 5, 455–84.
- VAN DAMME, E. E. C. (1991): *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin, Heidelberg.

- (2002): “Strategic Equilibrium,” in *Handbook in Game Theory*, ed. by R. Aumann, and S. Hart, vol. III. North Holland, Chapter 41.
- VERMEULEN, A. J. (1996): *Stability in Noncooperative Game Theory*. Department of Mathematics, University of Nijmegen, PhD Thesis.
- VOORNEVELD, M. (2005): “Persistent retracts and preparation,” *Games and Economic Behavior*, 51, 228–32.