

Balkenborg, Dieter; Hofbauer, Josef; Kuzmics, Christoph

Working Paper

The refined best-response correspondence in normal form games

Working Papers, No. 466

Provided in Cooperation with:

Center for Mathematical Economics (IMW), Bielefeld University

Suggested Citation: Balkenborg, Dieter; Hofbauer, Josef; Kuzmics, Christoph (2012) : The refined best-response correspondence in normal form games, Working Papers, No. 466, Bielefeld University, Institute of Mathematical Economics (IMW), Bielefeld, <https://nbn-resolving.de/urn:nbn:de:0070-pub-26717370>

This Version is available at:

<https://hdl.handle.net/10419/81096>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Working Papers

Institute of
Mathematical
Economics

466


April 2012

The refined best-response correspondence in normal form games

Dieter Balkenborg, Josef Hofbauer and Christoph Kuzmics



IMW · Bielefeld University
Postfach 100131
33501 Bielefeld · Germany



email: imw@wiwi.uni-bielefeld.de
<http://www.imw.uni-bielefeld.de/research/wp466.php>
ISSN: 0931-6558

The refined best-response correspondence in normal form games

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

April 23, 2012

Abstract

This paper provides an in-depth study of the (most) refined best reply correspondence introduced by Balkenborg, Hofbauer, and Kuzmics (2012). An example demonstrates that this correspondence can be very different from the standard best reply correspondence. In two-player games, however, the refined best reply correspondence of a given game is the same as the best reply 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 other pure strategies. Fixed points of the refined best reply 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 reply correspondence.

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

JEL codes: C62, C72, C73

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

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

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

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

1 Introduction

The minimal refined best reply correspondence was introduced by Balkenborg, Hofbauer, and Kuzmics (2012) 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 (2012) 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 reply correspondence. If the best reply 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 reply, the refined best reply dynamics can be described as a gradual process in which revising agents always switch to a best reply that is also a unique best reply against a strategy profile arbitrarily close to the current one.

The lattice theorem of Balkenborg, Hofbauer, and Kuzmics (2012, Theorem 1) implies that of all best-reply-like dynamics the refined best reply dynamics has the fewest stationary points, as the refined best reply correspondence has the fewest fixed points of all best-reply-like correspondences.¹ It furthermore implies that the refined best reply dynamics has the most (asymptotically) stable points (all of course fixed points of the refined best reply correspondence, and thus Nash equilibria) of all best-reply-like dynamics as it has the fewest solution trajectories of all best-reply-like dynamics. Of all best-reply-like dynamics the refined best reply dynamics is the one that makes the most Nash equilibria (yet, only Nash equilibria) stable. The refined best reply dynamics, of all best-reply-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-reply-like dynamics.

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

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

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

¹What we here call a best-reply-like dynamics or correspondence is what in Balkenborg, Hofbauer, and Kuzmics (2012) is formally defined and termed a generalized best reply dynamics or correspondence.

equivalent to a proper mixture of other pure strategies.

While there are refined best reply correspondences that are unlike any best reply correspondence in three or more player games, we can show that for all generic normal form games, refined best reply and best reply correspondences are identical, nevertheless. This tells us that the refined best reply correspondence is of interest only in non-generic games. Of course, many games of interest, such as extensive form games or cheap-talk games, have non-generic reduced normal form representations, in which the refined best reply correspondence would typically not be identical to the best reply correspondence.

We then proceed to partially characterize fixed points of the refined best reply 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 reply 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 strategically stable need not be a fixed point of the refined best reply correspondence. Finally, we apply the refined best reply 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 we study and defines the minimal refined best reply correspondence. Section 3 analyzes the differences between best reply and refined best reply correspondences. Section 4 analyzes fixed-points of the best reply 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 reply and best reply correspondences, and in particular the usefulness and intuitive appeal of the refined best reply 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 function². 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 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 $\beta(x) = \times_{i \in I} \beta_i(x)$.

As in Balkenborg, Hofbauer, and Kuzmics (2012) we shall restrict attention to games with a normal form in which the set of mixed-strategy profiles $\Psi = \{x \in \Theta \mid \mathcal{B}(x) \text{ is a singleton}\}$ is dense in Θ . In other words Ψ has Lebesgue measure 1. 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 (2012) if a game is not in this class \mathcal{G}^* it must be such that at least one player has two own-payoff equivalent strategies. Two strategies $x_i, y_i \in \Theta_i$

²The function u will also denote the expected utility function in the mixed extension of the game Γ .

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 reply correspondence* as defined in Balkenborg, Hofbauer, and Kuzmics (2012) and as follows. For $x \in \Theta$ let the set of pure *refined best replies* be given by

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

Then $\sigma_i(x) = \Delta(\mathcal{S}_i(x))$ 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 reply against x that is also a best reply for an open subset of any neighborhood of x .³

3 The difference between the best reply and the refined best reply correspondence

This section defines and discusses notions of strict and weak local dominance, that will be useful in understanding the difference between the best reply and the refined best reply correspondences. We term these notions *strict* and *weak inferiority*.⁴ 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 *A strategy $s_i \in S_i$ is strictly inferior if for every $x \in \Theta$ there is a $t_i \in S_i$ such that $u_i(s_i, x_{-i}) < u_i(t_i, x_{-i})$. A strategy $w_i \in S_i$ is weakly inferior if for every $x \in \Theta$ there is a $t_i \in S_i$, $t_i \neq w_i$ such that $u_i(w_i, x_{-i}) \leq u_i(t_i, x_{-i})$ and $u_i(w_i, y_{-i}) < u_i(t_i, y_{-i})$ for some $y \in \Theta$.*

A strictly inferior strategy is, thus, defined as a strategy that for every strategy profile $x \in \Theta$ is locally strictly dominated by some other pure strategy, but not necessarily uniformly dominated by a single strategy. A weakly inferior strategy is such that for every strategy profile $x \in \Theta$ it is either locally strictly dominated or if it is a best response then there is another best response to x that is better than the weakly inferior strategy against some other strategy profile $y \in \Theta$.

One can express inferior strategies in terms of the best response correspondence as follows. A strictly inferior strategy s_i is such that $s_i \notin \mathcal{B}_i(x)$ for any $x \in \Theta$. That is, a strictly inferior strategy is never a best response. 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.

For games in the class \mathcal{G}^* another alternative characterization of weakly inferior strategies can be given. For such games a weakly inferior strategy is never best on an open set of strategy profiles. This result is given in Proposition 1 below. In order to prove it, the following lemma, due to Kalai and Samet (1984), is useful.

³Balkenborg (1992) calls strategies $s_i \in \mathcal{S}_i(x)$ semi-robust best replies. This is inspired by Okada (1983) who calls a strategy a robust best reply to strategy profile x if it is a best reply for an open neighborhood of x . One could call a strategy robust if it is a robust best reply against some strategy profiles. Any pure strategy that is either a robust best reply or a semi-robust best reply against some strategy profile x is, thus, a robust strategy. Note that, while every strategy profile $x \in \Theta$ has a semi-robust best reply for all players, it may not have a robust best reply.

⁴Our notions of strict and weak inferiority are motivated by, but not identical to, the notion of inferior choices in Harsanyi and Selten (1988).

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

Proposition 1 *Let $\Gamma = (I, S, u) \in \mathcal{G}^*$. A strategy $s_i \in S_i$ is not weakly inferior if and only if it is a best reply for an open subset of strategy profiles.*

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

Note that every game in \mathcal{G}^* has at least one strategy for each player that is not weakly inferior.

We now turn to the question as to how these notions of inferiority, especially of the weak variety, help us to understand the distinction between the best reply and refined best reply correspondences.

It is immediate to see that a strictly inferior strategy is a strategy that is never a best response, i.e. not in $\mathcal{B}_i(x)$ for any $x \in \Theta$ and vice versa. We now ask the question whether a weakly inferior strategy is a strategy that is never a refined best-response, i.e. not in $\mathcal{S}_i(x)$ for any $x \in \Theta$, and vice versa. It is immediate from Proposition 1 that one direction is true: a weakly inferior strategy is a never a refined best-response. This result we state as Proposition 2 below. However, the other direction, somewhat surprisingly, is not generally true as we demonstrate by example below. In two-player games, however, one can show that both directions hold, which we state as Theorem 1.

As an immediate consequence of Proposition 1 we have:

Proposition 2 *Let $w_i \in S_i$ be weakly inferior for player i . Then $w_i \notin \mathcal{S}_i(x)$ (i.e., w_i is not a refined best reply) for any $x \in \Theta$.*

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

Theorem 1 *In a two-player game a strategy is a pure refined best reply, $s_i \in \mathcal{S}_i(x)$, if and only if it is a best reply, $s_i \in \mathcal{B}_i(x)$, and is not weakly inferior.*

Theorem 1 implies the following Theorem.

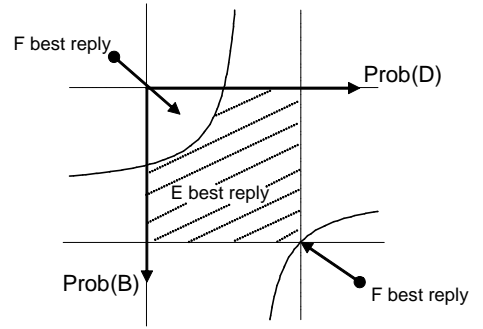
Theorem 2 *Let $\Gamma = (I, S, u) \in \mathcal{G}^*$ be a two-player game with minimal refined best reply correspondence $\sigma(\Gamma)$. Then there is a game $\Gamma' = (I', S', u') \in \mathcal{G}^*$ with $I' = I$, $S' = S$, and a payoff function $u' : S \rightarrow \mathbb{R}^2$ such that its best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

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

Theorem 2 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 reply correspondence σ . Then the set of fixed points of σ is the union of finitely many polytopes and hence the union of finitely many connected components.*

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



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

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

must be a best response against (A, C) , but not against (B, D) . This implies

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

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

The fact that Theorem 2 does not extend to 3 or more player games is quite remarkable. It implies that, although the refined best reply correspondence σ satisfies many properties that the best reply correspondence satisfies, such as being upper hemi continuity, closed- and convex-valued, and having a product structure, it is nevertheless, at least in some cases, not like any best reply correspondence. Thus knowing that the best reply correspondence satisfies a certain property does not immediately imply that the refined best reply 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 reply correspondence is equal to the best reply 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 3 *For generic normal form games a pure strategy is a refined best reply if and only if it is a best reply (i.e., $s_i \in \mathcal{S}_i(x) \Leftrightarrow s_i \in \mathcal{B}_i(x)$). That is we have $\sigma \equiv \beta$ for generic normal form games.*

For the proof see appendix A.

Theorem 3 was originally established in Balkenborg (1992). Given that persistent retracts are minimal CURB sets based on the refined best reply correspondence σ , see Balkenborg, Hofbauer, and Kuzmics (2012, Lemma 3), also originally shown in Balkenborg (1992), Theorem 3 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.


Having established that, for two-player games, weakly inferior strategies are exactly those strategies that are never refined best replies, we now provide a complete characterization of weakly inferior strategies.

Theorem 4 *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 proof of Theorem 4 is given in Appendix B, we here provide an intuitive sketch of the argument. The result is, of course, similar to Pearce's (1984) results, 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 replies, and that weakly dominated strategies are exactly those that are never a best reply to a properly mixed strategy. The proof of Theorem

4 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).

| | F | G |
|---|---|---|
| A | 3 | 1 |
| B | 1 | 3 |
| C | 2 | 2 |
| D | 3 | 0 |
| E | 0 | 3 |



Game 2 and Figure 2: A two-player game to illustrate the proof of Theorem 4. Payoffs are given only for player 1, who chooses the row. The 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 2, which are simple variations of the game and picture in Fudenberg and Tirole (1991, p. 50, Figure 2.2). Clearly pure strategies A and B are unique best responses against some opponent strategies. Thus, both strategies are refined best replies. Refined best replies 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 reply.

Now turn to the weakly dominated strategy D. The only downward sloping straight line through point D in the picture 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 opponent strategy profiles that makes strategy D a best reply.

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, 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 reply.

4 Nash equilibrium versus best reply refinements

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

Proposition 3 *Let Γ be a finite two-player game in \mathcal{G}^* . Let $x \in \Theta$ be a fixed point of the refined best reply correspondence σ . Then $x_{iw_i} = 0$ for every weakly inferior $w_i \in S_i$.*

Proof: Immediate from Proposition 2: Let $x \in \sigma(x)$. By Proposition 2 $w_i \notin \mathcal{S}_i(x)$ 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 .

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

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

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

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

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

Proposition 4 Let Γ be a 2-player game in \mathcal{G}^* . Then every pure fixed-point, $s \in S$, of the refined best reply correspondence, σ , is a perfect equilibrium.

Proof: Every pure fixed point of σ is undominated by Proposition 3. In two-player games every undominated Nash equilibrium is perfect. QED

The converse of Proposition 4 is not true. Consider Game 4. In this game strategy A (and similarly D) is equivalent to the mixture of pure strategies B and C (E and F respectively). However, A is a best reply only on a thin set of mixed-strategy profiles. In fact, A is best against any $x \in \Theta$ in which $x_{2E} = x_{2F}$, the set of which is a thin set. Thus, this game is in \mathcal{G}^* . In this game (A, D) constitutes a perfect equilibrium. In fact every mixed strategy profile $((\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}); (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2}))$ is a strictly perfect equilibrium, and hence, constitutes a singleton KM-stable set, where KM-stable is (strategically) stable in the sense of Kohlberg and Mertens (1986) (a minimal set satisfying their Property S). See next section for a definition. But none of these equilibria, except the one with $\alpha = 0$, are fixed points of σ , due to the fact that A (and D) is only best on a thin set (it is a weakly inferior strategy).

Proposition 4 cannot be generalized to general n -player games. To see this consider the following immediate characterization of fixed points of σ . For $x_i \in \Theta_i$ let $C(x_i) = \{s_i \in S_i | x_{is_i} > 0\}$ denote the carrier (or support) of x_i .

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

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

Balkenborg, Hofbauer, and Kuzmics (2012) prove that CURB sets (Basu and Weibull (1991)) based on σ give rise to absorbing retracts (Kalai and Samet (1984)) and minimal such sets give rise to persistent retracts. One might think that fixed points of σ will have some relation to persistent equilibria (Nash equilibria in a persistent retract, Kalai and Samet (1984)). This is not true, though. Note first that the mixed equilibrium in the coordination game is not persistent and is a fixed point of σ . Even for pure strategy fixed points of σ this is not true. Consider Game 5 taken from Kalai and Samet (1984). The equilibrium (B, D, E) is perfect and proper but not persistent as Kalai and Samet (1984) point out. It is also a fixed point of σ . To see this, note that E is weakly dominant for player 3 and that B and D are best (for players 1 and 2, respectively) against all nearby strategy profiles in which player 2 chooses strategy C with smaller probability than player 3 chooses F and player 1 chooses A with smaller probability than player 3 chooses F .

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

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

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

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

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

5 σ -CURB sets and strategic stability

As already mentioned Balkenborg, Hofbauer, and Kuzmics (2012) 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 KM-stable set.

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

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

The following defines property S of Kohlberg and Mertens (1986) without the requirement of 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 $\Theta^*(R) = \{x \in \Theta(\eta) \mid x_{is} = \eta_{is} \text{ if } s \notin R_i\}$, a compact and convex subset of $\Theta(\eta)$. For $x \in \Theta^*(R)$ let $\sigma^*(x) = \{y \in \Theta(\eta) \mid y_{is} = \eta_{is} \text{ if } s \notin \mathcal{S}_i(x)\}$. 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 (2012) that σ -Curb sets are asymptotically stable sets under the refined best reply dynamics, this result is reminiscent to the result of 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 reply dynamics.

Jansen, Jurg, and Born (1994) have shown that persistent retracts contain a KM-stable set for all 2-player games. Balkenborg (1992) has shown that every persistent retract contains a strategically stable set in the sense of Hillas (1990). We here provide a simple direct proof for the original concept of Kohlberg and Mertens (1986).

6 A cheap-talk example

This section presents a very simple, perhaps the simplest, class of cheap-talk games, or sender-receiver games.⁵ All games (in particular also generic games) within this class are non-generic in the space of normal form 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*. As a particular simple example consider the pure common interest game, for which $u_{AA}^i = u_{BB}^i = 1$ and $u_{AB}^i = u_{BA}^i = -1$ for both $i \in \{1, 2\}$. For the special case of $\rho = \frac{1}{2}$ the game can be represented in matrix form as follows.

| | 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 7: A simple communication game with common interest.⁶

⁵The 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).

Note that Game 7 is a symmetric game. Note that it has no weakly dominated strategies. Player 1's pure strategies (m_A, n_B) and (n_A, m_B) are (in fact unique) best replies against appropriate pure strategies of the opponent, (A_m, B_n) and (B_m, A_n) , respectively. Player 1's pure strategies (m_A, m_B) and (n_A, n_B) are best 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 player 2'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 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 lightly 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, player 1'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, by Theorem 4 player 1's pure strategies (m_A, m_B) and (n_A, n_B) are weakly inferior and, thus by Theorem 1 never refined best replies. This game has only three fixed points of the refined best reply 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.

We leave it to the reader to verify that the following statements are true for all simple communication games. None of these games has weakly dominated strategies. All 4 uninformative pure equilibria of Game 7 are also Nash equilibria in every simple communication game. In fact all are singleton strategically stable sets in the sense of Kohlberg and Mertens (1986). In fact, all these games have a continuum of singleton strategically stable sets. Yet, in all these games player 1's pure strategies (m_A, m_B) and (n_A, n_B) are each equivalent to some mixture of pure strategies (m_A, n_B) and (n_A, m_B) . Similarly, player 2's pure strategies (A_m, A_n) and (B_m, B_n) are each equivalent to some mixture of pure strategies (A_m, B_n) and (B_m, A_n) . The only fixed points of the refined best reply correspondence, in each of these games, are the (possibly mixed) Nash equilibria of the reduced game with strategies (m_A, n_B) and (n_A, m_B) only for player 1 and strategies (A_m, B_n) and (B_m, A_n) only for player 2.

7 Conclusion

We studied the refined best reply correspondence in normal form games as introduced by Balkenborg, Hofbauer, and Kuzmics (2012) as the basis for a dynamic learning model. We show by example that the refined best reply correspondence can be unlike any best reply correspondence. In two-player games, however, the refined best reply correspondence coincides with the best reply 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 other 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 reply and best reply 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 reply correspondence are the stationary points of the refined best reply dynamics of Balkenborg, Hofbauer, and Kuzmics (2012). 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 reply 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 reply correspondence over the ordinary best reply 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 replies. These would be of interest, as the refined best reply dynamics converges to the set of refined rationalizable strategies in every game (Balkenborg, Hofbauer, and Kuzmics 2012). 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.⁷ 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 reply correspondence could be fruitful in many such cases. One example of such a class is, of course, the class of extensive form games. Another would be 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 3. Note first that generic normal form games are in the class \mathcal{G}^* where all refined responses are pure strategies. The proof of Theorem 3 is organized in a number of steps: We will first fix some notations for the mappings and various submanifolds to be considered. Step 1 argues that the embedding of the uncorrelated strategy combinations into the set of beliefs has nice differentiability properties. Step 2 invokes the transversality theorem (see Guillemin and Pollack (1974)) to show that for generic payoff functions we obtain the needed transversality conditions.⁸ Step 3 argues that we can restrict attention to completely

⁷For this procedure see e.g. Dekel and Fudenberg (1990), Börgers (1994), and Ben Porath (1997).

⁸This 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.

mixed strategy combinations of the opponents. If the player is indifferent between several of his strategies against a given completely mixed strategy combination, step 4 shows how we can construct an arbitrarily nearby strategy combination, against which the player strictly prefers a given one among these strategies. Step 5 completes the argument.

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

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

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

q_i describes the first step in the computation mentioned above.

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

Let L_i be the vector space of all linear mappings

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

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

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

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

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

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

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

and

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

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

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

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

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

is well defined, is one-to-one, maps $X_{-i}(T_{-i})$ to a closed set, and has a derivative $dq_i|_{x_{-i}}$ with maximal rank everywhere.⁹

⁹The 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.

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).

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

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

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

Step 4: Suppose v_i satisfies ⊗. 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

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

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

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 $0 < \gamma < \epsilon$ $y_{-i} := c(\gamma)$ 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 replies in two-player games

This appendix provides proofs of Theorems 1 and 4. In the case of two player games the payoff function is bilinear and hence linear in the mixed strategy choice of the opponent. This allows the use of convex analysis (see Rockafellar (1970)) to study the best reply correspondence of a player. The most direct consequence is the convexity of the region where a strategy is a best reply. From this Theorem 1 follows immediately, the arguments are given after Lemma 2 below. More work is needed to obtain Theorem 4. We use conjugate functions and provide the proof after Lemma 3.

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

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

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. 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 reply.

Lemma 2 *The region $H(x_1)$ is a convex polyhedron.*

Proof: $H(x_1)$ is the image of convex polyhedron $G(x_1)$ under the linear projection mapping and hence also a convex polyhedron. QED

Clearly x_1 is *robust* 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 robust and a best reply against x_2 , then x_2 is in the closure of the open set $H^\circ(x_1)$ and so x_1 is a robust best reply against x_2 . Using Proposition 1 this implies immediately Theorem 1.

The remainder of this section aims at proving Theorem 4. We consider next the epigraph

$$F = \{(x_2, \alpha) \in \mathbb{R}^{K-1} \times \mathbb{R} \mid f(x_2) \leq \alpha\} \quad (10)$$

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

$$f^*(x_2^*) = \sup_{x_2 \in \mathbb{R}^{K-1}} \{x_2^* \bullet x_2 - f(x_2)\} = \max_{x_2 \in \Theta_2} \{x_2^* \bullet x_2 - f(x_2)\} < \infty, \quad (11)$$

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

$$a(x_2) = -\alpha + \sum_{k=1}^{K-1} x_2^{*k} x_2^k \quad (12)$$

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

$$-e_k = (-e_k^1, \dots, -e_k^K) \in \mathbb{R}^K \text{ with } e_k^l = \begin{cases} -1 & \text{for } k=l \\ 0 & \text{else} \end{cases} \quad (13)$$

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

$$e = (1, \dots, 1) \in \mathbb{R}^K \quad (14)$$

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

$$\alpha^* + \lambda \rho_K \geq x_2^{*'} x_2 - \sum_{k=1}^{K-1} \lambda \rho_k x_2^k + \lambda \rho_K - f(x_2) \quad (15)$$

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 a affine function defined by a strategy x_1 in Θ_1 for which $H(x_1)$ is full dimensional. Given our identification, $(\hat{x}_2^*, f^*(\hat{x}_2^*))$ is consequently a robust strategy in Θ_1 , which was to be shown. QED

Proof of Theorem 4:

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

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

$$u_1(x_1', x_2) = u_1(x_1, x_2) - \sum_{k=1}^{K-1} \rho_k x_2^k - \rho_K \left(1 - \sum_{k=1}^{K-1} x_2^k\right) \leq u_1(x_1, x_2), \quad (16)$$

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

In summary, the only robust strategies in Θ_1 are the extreme point of F^* . All other strategies are proper mixtures of not own-payoff equivalent robust strategies or are weakly dominated and therefore not robust. 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 (2012): “Refined best-response correspondence and dynamics,” *Theoretical Economics*, forthcoming.
- 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.
- 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.
- VOORNEVELD, M. (2005): “Persistent retracts and preparation,” *Games and Economic Behavior*, 51, 228–32.