
The Refined Best Reply Correspondence and Backward Induction

Dieter Balkenborg

University of Exeter

Christoph Kuzmics

University of Graz

Josef Hofbauer

University of Vienna

Abstract. *Fixed points of the (most) refined best reply correspondence, introduced in Balkenborg et al. (2013), in the agent normal form of extensive form games with perfect recall have a remarkable property. They induce fixed points of the same correspondence in the agent normal form of every subgame. Furthermore, in a well-defined sense, fixed points of this correspondence refine even trembling hand perfect equilibria, while, on the other hand, reasonable equilibria that are not weak perfect Bayesian equilibria are fixed points of this correspondence.*

JEL classification: C62, C72, C73.

Keywords: Subgame perfection; Nash equilibrium refinements; backward induction; sequential rationality.

1. INTRODUCTION

What constitutes a ‘reasonable’ backward induction solution to an extensive form game is debatable. The first concept developed to eliminate some commonly agreed-on unreasonable Nash equilibria is Selten’s (1965) subgame perfection. Yet, there are many subgame perfect equilibria that researchers agree on to be not reasonable. Selten (1975) introduced the concept of an extensive-form trembling hand perfect equilibrium, a trembling hand perfect equilibrium of the agent normal form of the game, to further eliminate unreasonable equilibria. This concept, while in many ways a fine one, was ‘recast and slightly weakened’ by Kreps and Wilson (1982) in their sequential equilibrium. Kreps and Wilson (1982) also coined the terms ‘assessment’ and ‘sequential rationality’, now to be found in every textbook on game theory. Kreps and Wilson (1982) thus offer a reinterpretation of (a slightly weaker notion than) extensive-form trembling hand perfect equilibria, in terms of what players believe about what happened so far in the game (when it is their turn to move) and what they should choose given that and given the likely continuation of other players after them. One commonly agreed-on problem of the concept of sequential equilibrium is that, while it is based on these assessments, their justification is still derived from trembles. Fudenberg and Tirole (1991b) introduced the notion of *perfect Bayesian equilibrium* which does not require trembles. A variant of this notion, now often called weak perfect Bayesian equilibrium, see e.g., Mas-Collel *et al.* (1995,

Definition 9.C.3) or Ritzberger (2002, Definition 6.2), has been defined for all extensive form games. Being weak perfect Bayesian is now commonly considered a minimal requirement for a reasonable solution to extensive form games. Battigalli (1996) demonstrates that Fudenberg and Tirole's (1991b) 'generally reasonable extended assessment[s] ... may violate independence, full consistency, and invariance with respect to interchanging of essentially simultaneous moves', where full consistency is the criterion an assessment has to justify in addition to the strategy profile being sequentially rational to obtain a sequential equilibrium.

Efforts have since been made to find conditions on assessments (not based on trembles) such that ultimately we obtain sequential equilibria, if not extensive-form trembling hand equilibria (Battigalli, 1996; Kohlberg and Reny, 1997).

On the other hand, there are strategy profiles that are not trembling hand equilibria (or not even weak perfect Bayesian) and yet are not unreasonable. In all these solution concepts, assessments generally need to be justified by the same strategy-profile for all players. It is not a priori clear why this is a reasonable requirement. Indeed Battigalli (1996), Bonanno (1995), and Fudenberg and Tirole (1991a, pp. 332–333) argue that this is not a necessary requirement of a reasonable solution and allow for *heterogeneous* assessments. We even argue (see the game in Figure 2) that a player might randomize (at least in the mind of her opponents) over two or more pure strategies, where each pure strategy is independently justifiable by a fully consistent (or some such requirement) assessment, but the mixture itself is not.

In this paper, we propose a new equilibrium concept for extensive form games and discuss its relationship with the afore-mentioned standard equilibrium concepts. The equilibrium we propose is derived from fixed points of the refined best reply correspondence, introduced by Balkenborg *et al.* (2013), for the agent normal form of the given extensive form game.

We first prove that fixed points of the refined best reply correspondence, introduced by Balkenborg *et al.* (2013), in the agent normal form of an extensive form game¹ have the remarkable conceptual consistency property that they automatically induce fixed points of the same correspondence in the agent normal form of every subgame.

We then prove that fixed points of this refined best reply correspondence, in the agent normal form, satisfy exactly the properties discussed above. Each pure strategy that is used by some player with positive probability must be justified by, in fact, a stronger requirement than full consistency (even stronger than trembling hand perfection). However, there does not need to be one assessment justifying all these pure strategies.

The paper proceeds as follows. Section 2 provides the setup and the definition of the refined best reply correspondence. The conceptual subgame consistency property, as well as a full characterization of fixed points of the refined best reply correspondence of the agent normal form, are both stated and proven in Section 3. Section 4 concludes with a discussion of the degree of 'reasonableness' of fixed points of the refined best reply correspondence. The paper has two

1. Such fixed points can be understood as the potential convergence points of a most refined learning dynamics (Balkenborg *et al.*, 2013).

appendices. Appendix A demonstrates how, not surprisingly, none of our results extend to fixed points of the refined best reply correspondence for the reduced normal form of an extensive form game. Appendix B demonstrates that rationalizability based on the refined best reply correspondence in the agent normal form has no connection to forward induction reasoning and, thus, very little in common, with the notion of extensive form rationalizability of Battigalli (1997) and Pearce (1984).

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 function.² For any finite set K with $|K|$ elements let

$$\Delta(K) = \{x \in \mathbb{R}^{|K|} \mid x_k \geq 0 \text{ for all } k \in K \text{ and } \sum_{k \in K} x_k = 1\}$$

denote the set of all probability distributions over K . 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 $x_i \in \Theta_i$ (for any $i \in I$) let $C(x_i) \subset S_i$ denote the support (or carrier) of mixed strategy x_i of player i .

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

As in Balkenborg *et al.* (2013), we shall restrict attention to games with a normal form in which the set

$$\Psi = \{x \in \Theta \mid \mathcal{B}(x) \text{ is a singleton}\}$$

of mixed-strategy profiles is dense in Θ . We denote this class by \mathcal{G}^* .

For every extensive form game with an agent normal form that is in this class \mathcal{G}^* , it is true that every subgame has an agent normal form that is in this class \mathcal{G}^* . This follows from the proof of Proposition 1.

For normal form games, the class \mathcal{G}^* is essentially the class of games in which no player has two or more pure strategies that are payoff equivalent for her, see Balkenborg *et al.* (2013, Proposition 1). We are here interested in extensive form games and their agent normal form. One can show that the agent normal form of an extensive form game being in class \mathcal{G}^* is a generic property of extensive form games, meaning that it is true for a set of extensive form games that is open and dense in the space of all extensive form games (with the same game tree, i.e., just varying payoffs, not strategy sets). This follows from Balkenborg *et al.* (2013, Proposition 3).

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

Refined Best-Response and Backward Induction

For games in \mathcal{G}^* and for mixed strategy profile $x \in \Theta$ let

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

denote the set of pure (most) refined best replies. Furthermore, for $x \in \Theta$ let $\sigma_i(x) = \Delta(\mathcal{S}_i(x))$ and $\sigma(x) = \times_{i \in I} \sigma_i(x)$ for all $x \in \Theta$. We then call, following Balkenborg *et al.* (2013), σ the *refined best reply correspondence*.³

An equivalent definition, in words, of a pure refined best reply to a strategy profile x is that it is not only a best reply against x , but is also the unique best reply against all strategy profiles in an open subset of any neighborhood of x . This open subset necessarily includes (mostly) completely mixed strategy profiles. A mixed refined best reply to strategy profile x is then simply any probability distribution over the set of pure refined best replies to x .

3. RESULTS

3.1. A consistency property

Recall that the agent normal form is derived from an extensive form by replacing each player with a set of agents, one for each information set of this player. These agents then choose independently of each other. Fixed points of σ in the agent normal form have a surprising property. They induce fixed points of σ in every subgame. This is much more than saying that fixed points of σ are subgame perfect. One might call it a *conceptual consistency* property.

Proposition 1. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of a given extensive form game. Then if a strategy profile x is a fixed point of σ it is also a fixed point of σ in the agent normal form of every subgame of this extensive form game.

Proof. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of the given extensive form game. Let Θ denote the space of mixed strategies. Let $x \in \Theta$ be a fixed point of σ . Consider agent $i \in I$, where I is the set of all agents and consider a pure action $s_i \in C(x_i)$. By definition there is a sequence of $x^t \in \Theta$ such that $x^t \rightarrow x$ and $\mathcal{B}_i(x^t) = \{s_i\}$.

Now consider any subgame in which agent i also moves. Let $\hat{\Gamma} = (\hat{I}, \hat{\mathcal{S}}, \hat{u})$ denote its agent normal form. Obviously $\hat{I} \subset I$, with $i \in \hat{I}$, and for all $j \in \hat{I}$ we have $\hat{\mathcal{S}}_j = \mathcal{S}_j$, and \hat{u} is defined accordingly. Now consider the projection of every $x^t \in \Theta$ onto the reduced game $\hat{\Gamma}$. Let it be denoted by $\hat{x}^t \in \hat{\Theta}$. Hence we simply have that $\hat{x}_j^t = x_j^t$ for all $j \in \hat{I}$.

Obviously we have that $\hat{x}^t \rightarrow \hat{x}$, where \hat{x} is the projection of x onto the subgame. Furthermore we must have that $\hat{\mathcal{B}}_i(\hat{x}^t) \subset \mathcal{B}_i(x^t)$ for every x^t . This is so because either agent i 's information is reached in the whole game under strategy profile x^t , in which case player i 's best responses cannot change in the subgame, or player i 's information set is not reached under x^t , and, hence, every strategy of player i is a best reply against x^t in the full game. But now given that

3. See Balkenborg *et al.* (2015) for a detailed discussion of the refined best reply correspondence in general normal form games.

$\mathcal{B}_i(x^t) = \{s_i\}$ for all x^t we must also have $\hat{\mathcal{B}}_i(\hat{x}^t) = \{s_i\}$ for all \hat{x}^t . Thus, $s_i \in \hat{\mathcal{S}}_i(\hat{x})$ and, as this is true for all $s_i \in C(x_i)$ and all players $i \in \hat{I}$ we have that \hat{x} is a fixed point of σ in $\hat{\Gamma}$. ■

The standard best reply correspondence obviously does not have this conceptual consistency property. The well-known entry deterrence game, see e.g., Mas-Collel *et al.* (1995, example 9.B.1), serves to illustrate that. This can also be seen in the game given in Figure 1 below. This is the reason why Selten (1965) introduced the notion of subgame perfection. One can ask whether there are other generalized best reply correspondences, as defined in Balkenborg *et al.* (2013, Definition 1), that also have this conceptual consistency property (for all games!). However, a generalized best reply correspondence, as defined in Balkenborg *et al.* (2013, Definition 1), is only defined for a given game. One can define the (most) refined best reply correspondence generally, because, by Balkenborg *et al.* (2013, Theorem 1) there is a unique smallest (or most refined) generalized best reply correspondence in all games (in \mathcal{G}^*). If one wants to study a particular generalized best reply correspondence in all games, one first needs to define systems of generalized best reply correspondences, where a system is a selection from the set of generalized best reply correspondences for every game. One such system is the system of most refined best reply correspondences (for all games), a system we have here somewhat loosely described simply by the term refined best reply correspondence. Another such system is the system of standard best reply correspondences (for all games). If we impose no structure on what systems are allowed, then there are of course systems of generalized best reply correspondences that differ from the system of most refined best reply correspondences in at least one game (in which however the fixed points of this correspondence are the same as those of the most refined correspondence) such that this system also has the conceptual consistency property. One could investigate whether a reasonable class of systems of generalized best reply correspondence is such that only the system of the most refined best reply correspondence has the conceptual consistency property among all systems in this class. We shall not pursue this here. Instead, we offer an example of a particular system of generalized best reply correspondences that is more refined than the system of best reply correspondences but still does not satisfy the conceptual consistency property. The game to demonstrate this claim also helps to illustrate Proposition 1. It is a three player extension of the well-known entry-deterrence game.

Consider the game of Figure 1. First, note that the strategy profile (O, F, I) is a Nash equilibrium of the game that does not induce Nash equilibria in all subgames. For instance, the third player's choice I is strictly dominated in the subgame that begins at her information set. This again demonstrates that the best reply correspondence does not have the conceptual consistency property.

Now consider the following *admissible best reply correspondence*. For $x \in \Theta$ and $i \in I$ let $\tilde{\mathcal{B}}_i(x) \subset S_i$ denote the set of pure and admissible (i.e., not weakly dominated) best replies against x . Then let $\tilde{\beta}_i(x) = \Delta(\tilde{\mathcal{B}}_i(x))$ and $\tilde{\beta}(x) = \times_{i \in I} \tilde{\beta}_i(x)$ for all $x \in \Theta$.

We now show that the pure strategy profile (O, F, N) is a fixed point of $\tilde{\beta}$. First, note that players 2 and 3 are playing best responses as their strategy choice has no impact on their payoffs given player 1's choice of strategy O . In other




Figure 1 A game that illustrates the conceptual consistency property of the refined best reply correspondence stated in Proposition 1. It also demonstrates that other generalized best reply correspondences, such as the standard best reply correspondence and the admissible best reply correspondence do not satisfy the conceptual consistency property.

words, the subgame starting with player 2’s information set is unreached in this strategy profile. Player 1 is then playing the unique best reply to player 2’s (and 3’s) strategy choice. Player 1’s strategy *O* is thus also admissible. Player 2’s strategy *F* is a unique best reply to players 1 and 3 playing *I* each and thus also admissible. Finally player 3’s strategy *N* is admissible as it weakly dominates the only other pure strategy of player 3’s, strategy *I*.

Now consider the induced strategy profile (*F*, *N*) in the subgame that starts with player 2’s information set. While *F* and *N* are still admissible strategies, player 2’s strategy *F* is, in this subgame, no longer a best reply to player 3’s strategy *N*. Thus, the strategy profile (*F*, *N*) is not a fixed point of $\tilde{\beta}$ in this subgame. Thus, the admissible best reply correspondence does not have the conceptual consistency property.

Finally, consider the refined best reply correspondence, σ . By Proposition 1 it has the conceptual consistency property. To see this in this game we note that the only fixed point of σ is the backward induction solution (*I*, *A*, *N*). Let us just argue why, for instance, (*O*, *F*, *N*) is not a fixed point of σ . Strategy *N* weakly dominates strategy *I* for player 3, so player 3’s pure strategy *N* is the unique best reply to any completely mixed strategy profile. Pure strategy *N* is thus player 3’s unique refined best reply to strategy profile (*O*, *F*, *N*). What about player 2’s choice of *F*? Note that in all completely mixed strategy profiles in which player 3 puts probability close to 1 on *N*, pure strategy *A* is player 2’s unique best reply. Therefore, only pure strategy *A* is a refined best reply to strategy profile (*O*, *F*, *N*). Thus, (*O*, *F*, *N*) is not a fixed point of the refined best reply correspondence.

3.2. Relationship to standard solution concepts

In fact, a fixed point of σ of the agent normal form of an extensive form game is both stronger in some respects and weaker in others than a sequential equilibrium. To clarify these issues, we separate the equilibrium definitions of the various equilibrium concepts into two parts.

For a given strategy profile $x \in \Theta$ in the agent normal form of an extensive form game Γ let μ be a system of probability distributions, one for each

information set, where each probability distribution is a probability distribution over the set of nodes of one information set. This system μ , which must be derived from x using Bayes' rule whenever possible, is often called a **system of beliefs** and the pair (x, μ) an **assessment**. Recall that $C(x_i) \subset \Theta_i$ denotes the support of mixed strategy x_i of player i . Given assessment (x, μ) we call $s_i \in C(x_i)$ **sequentially rational** if it maximizes the conditional expected payoff given (x, μ) at agent i 's (only) information set.

Definition 1 Consider a strategy profile $x \in \Theta$ in the agent normal form of an extensive form game.

It is **very weakly idio-justifiable** if for every $s_i \in C(x_i)$ and every player $i \in I$ there is a system of beliefs $\mu = \mu(s_i)$ such that s_i is sequentially rational given (x, μ) .

It is **weakly idio-justifiable** if for every $s_i \in C(x_i)$ and every player $i \in I$ there is a sequence of assessments (x^k, μ^k) such that the sequence $x^k \in \text{int}(\Theta)$ converges to x , μ^k is the appropriate and unique system of beliefs derived from x^k using Bayes' rule, and the sequence μ^k converges to $\mu = \mu(s_i)$ such that s_i is sequentially rational for player i given (x, μ) .

It is **strongly idio-justifiable** if it is weakly idio-justifiable and each $s_i \in C(x_i)$ is also sequentially rational for every assessment along the sequence.

It is **very strongly idio-justifiable** if for every $s_i \in C(x_i)$ and every player $i \in I$ there is an open set $U^x \subset \text{int}(\Theta)$, with closure containing x , such that s_i is sequentially rational also for all assessments (x', μ') , where $x' \in U^x$ and μ' derived from x' using Bayes' rule.

A strategy profile is **very weakly, weakly, strongly, or very strongly pan-justifiable** if it is very weakly, weakly, strongly, or very strongly idio-justifiable, respectively, and, in addition, the assessment, or the sequence or open set of assessments, that respectively justifies any one of the various pure strategies $s_i \in C(x_i)$ for the various players $i \in I$ is the same for all $s_i \in C_i(x_i)$ and all players $i \in I$.

Note that the definition of a very weakly pan-justifiable strategy profile is exactly that of a **weak perfect Bayesian equilibrium** as given in e.g., Definition 9.C.3 in Mas-Collel *et al.* (1995), or Definition 6.2 in Ritzberger (2002). Similarly, the definition of a weakly pan-justifiable strategy profile is exactly that of a **sequential equilibrium** as given in Kreps and Wilson (1982, p. 872). Finally, the definition of a strongly pan-justifiable strategy profile is exactly that of the equivalent definition of Selten's (1975) extensive-form **trembling hand perfect equilibrium** given in e.g., Ritzberger (2002, Proposition 6.1.c).

Note that when a possibly mixed strategy profile x is idio-justifiable (of some degree) it means that **each of its parts**, i.e., each pure strategy in the support of some player's strategy part of x , is **idiosyncratically** justifiable through its **very own** assessment, or sequence of assessments, or open set of assessments (depending on which degree of idio-justifiability we speak of). Thus, even for a single player, we might have two different assessments justifying two different pure strategies in the support of this player's mixed strategy. When a possibly mixed strategy profile x is pan-justifiable (of some degree) it means that it is idio-justifiable (of the same degree) and, in addition, there is a single assessment, or

sequence of assessments, or open set of assessment, that justifies (of the same degree) **all** pure strategies in the support of **any** player's part of x .

Obviously very strong idio-justifiability implies strong idio-justifiability, which implies weak idio-justifiability, which, in turn, implies very weak idio-justifiability. The same is true for the respective four notions of pan-justifiability. Also any level of pan-justifiability implies the same level of idio-justifiability by definition.

The following proposition states that fixed points of the refined best reply correspondence are exactly those strategy profiles that are very strongly idio-justifiable.

Proposition 2. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of an extensive form game. A strategy profile x is a fixed point of σ if and only if it is very strongly idio-justifiable.

Proof. Let s_i be in the support of player i 's part of x , i.e., $s_i \in C(x_i)$. Strategy profile x is a fixed point of σ if and only if, by the definition of σ , there is an open set $U^x \subset \text{int}(\Theta)$ with closure containing x , such that $s_i \in \mathcal{B}_i(\gamma)$ for any $\gamma \in U^x$. Note that this set U^x in general depends on the player i and on pure strategy s_i and can well be different for different pure strategies and different players. For all these γ there is a unique system of beliefs μ^γ derived from γ using Bayes' rule. Again, the thus derived beliefs μ^γ can be different for different pure strategies $s_i' \in C(x_i)$ and for different players. Then $s_i \in \mathcal{B}_i(\gamma)$ if and only if it is also sequentially rational given (γ, μ^γ) .

The following proposition has two parts. First, it states that even the strongest well-known degree of pan-justifiability, i.e., extensive-form trembling hand perfection, does not imply the strongest notion of idio-justifiability, i.e., being a fixed point of σ . In other words, an extensive-form trembling hand perfect equilibrium is not necessarily a fixed point of the refined best reply correspondence. Second, it states that the strongest form of idio-justifiability, i.e., being a fixed point of σ , does not imply even the weakest notion of pan-justifiability, i.e., being weak perfect Bayesian. In other words, a fixed point of the refined best reply correspondence is not necessarily a weak perfect Bayesian equilibrium. ■

Proposition 3. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of an extensive form game.

1. There are games with an extensive-form trembling hand perfect equilibrium that is not a fixed point of the refined best reply correspondence. I.e., strong pan-justifiability does not imply very strong idio-justifiability.
2. There are games with a fixed point of the refined best reply correspondence that is not a weak perfect Bayesian equilibrium. Thus, very strong idio-justifiability does not imply very weak pan-justifiability.

Proof. The proof is done by constructing two counterexamples. As every normal form game is also a (trivial) extensive form game, the first statement follows from the normal form game given in figure 7 (also figure 8) of Balkenborg *et al.* (2015), which shows a trembling hand perfect equilibrium that is not a fixed point of σ . For the second statement, consider the game given by Figure 2.

Consider the strategy profile x^* , in which player 1 plays A and player 2 puts probability $\frac{1}{2}$ each on L and R . Note first that this is indeed a Nash equilibrium as A is best when player 2 behaves thus, and player 2 is indifferent between all strategies when player 1 chooses A . To see that this is also a fixed point of σ , note that A continues to be the unique best reply of player 1 in a small enough neighborhood of player 2's strategy. Strategy L for player 2 is a unique best reply for some open set of mixed strategy profiles close to x^* , namely all these that put sufficiently higher probability on player 1's strategy B than on C . Similarly, player 2's strategy R is the unique best reply for some open set of mixed strategy profiles close to x^* , namely all these that put sufficiently higher probability on player 1's strategy C than on B . Thus $x^* \in \sigma(x^*)$. However, x^* is not weak perfect Bayesian. There is in fact no assessment of player 2 that would make player 1's best reply randomizing equally between L and R . This is so, because to induce player 2 to randomize she has to be indifferent between both strategies, which is only true if her assessment is that both nodes in her information set are equally likely. In that case, however, player 2's strategy M dominates. ■

3.3. Pure fixed points

Corollary 1. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of an extensive form game. If a pure strategy profile s is a fixed point of σ then it induces a weak perfect Bayesian equilibrium in every subgame.

Proof. This follows immediately from Propositions 1 and 2 and the realization that there is only one strategy in the support of each players's x_i . ■

3.4. Refined rationalizability and perfect information games

Define a notion of rationalizability (Bernheim, 1984; Pearce, 1984) based on σ as follows.⁴ For $A \subset \Theta$ let $S_i(A) = \bigcup_{x \in A} S_i(x)$. Let $\sigma_i(A) = \Delta(S_i(A))$. Let




Figure 2 A game in which there is a fixed point of σ in the agent normal form (i.e., a very strongly idio-justifiable equilibrium) which is not weak perfect Bayesian (i.e., not even very weakly pan-justifiable).

4. We use what Brandenburger and Dekel (1987) call independent rationalizability as opposed to correlated rationalizability. See also the notion of Δ -rationalizability of Battigalli and Siniscalchi (2003), with Δ here being the restriction to beliefs having a product structure.

Refined Best-Response and Backward Induction

$\sigma(A) = \times_{i \in I} \sigma_i(A)$. For $k \geq 2$ let $\sigma^k(A) = \sigma(\sigma^{k-1}(A))$. For $A = \emptyset$, $\sigma^k(A)$ is a decreasing sequence, and we denote $\sigma^\infty(\Theta) = \bigcap_{k=1}^{\infty} \tau^k(\Theta)$ the set of σ -rationalizable strategies.

Proposition 4. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of an extensive form game of perfect information in which for each player there are no two final nodes at which this player receives the same payoff. Then only the (unique) subgame-perfect strategy profile is σ -rationalizable.

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

Corollary 2. Let $\Gamma \in \mathcal{G}^*$ be the agent normal form of an extensive form game of perfect information in which for each player there are no two final nodes at which this player receives the same payoff. Then the only fixed point of σ for this game is the (unique) subgame perfect equilibrium.

Proof. Every fixed point of σ must be in the set of σ -rationalizable strategies. This set, by Proposition 4, only consists of the subgame perfect equilibrium.

This corollary also follows from Proposition 1 and the fact that σ is a refinement of β . ■

4. CONCLUSION

In this paper, we propose a new equilibrium concept for extensive form games. It is given by the fixed points of the refined best reply correspondence, as defined by Balkenborg *et al.* (2013), of the agent normal form of the extensive form game.

We show that fixed points of the refined best reply correspondence, unlike those of the best reply correspondence, induce fixed points of the refined best reply correspondence in every subgame. Thus, these fixed points are subgame perfect equilibria. Moreover, every pure strategy that any player uses with positive probability in any fixed point of this correspondence is what we term very strongly idio-justifiable. This means that there is an open set of strategy profiles near to the fixed point such that the given pure strategy of the given player is a unique best reply to all strategy profiles in this open set. In this sense, every pure strategy that any player uses with positive probability in any fixed point of this

correspondence is more strongly justified than those in weak perfect Bayesian, sequential, or even extensive-form trembling hand perfect equilibria. This is true for every individual strategy used by every individual player. However, the justification for all these strategies in the support of the fixed point and for all players does not have to be the same. This is the reason for our choice of the ‘idio’ prefix that is supposed to reflect that indeed every individual strategy may have its idiosyncratic justification. Most well-known equilibrium notions, such as weak perfect Bayesian, sequential, or extensive-form trembling hand perfect equilibrium are such that while each strategy individually is less justified than those in a fixed point of the refined best reply correspondence, all strategies in one equilibrium, however, have a common justification, i.e., are justified by the same belief system, or sequence of belief systems, etc., which we denote by adding a ‘pan’ prefix, where pan signifies (one for) all.

The sense in which fixed points of the refined best reply correspondence are weaker than even weak perfect Bayesian equilibria is that we cannot necessarily guarantee that the system of beliefs that justifies any single pure strategy in its support is the same that justifies other pure strategies in its support. This is reminiscent of Fudenberg and Levine’s (1993) self-confirming equilibria, in which players can disagree about other players’ strategies. Note, however, that for fixed points of the refined best reply correspondence of the agent normal form, players can only disagree about how to interpret apparent deviations from the prescribed strategy profile. Fixed points of the refined best reply correspondence are Nash equilibria, while even the most stringent of self-confirming equilibria, which are the rationalizable self-confirming equilibria of Dekel *et al.* (1999), are not necessarily Nash equilibria. In some sense, thus, fixed points of refined best reply correspondence of the agent normal form are highly justifiable self-confirming equilibria within the bounds of being Nash equilibria as well.

In Proposition 3 we show that fixed points of the refined best reply correspondence are not necessarily even weak perfect Bayesian. Can a non-weak perfect Bayesian equilibrium be a reasonable equilibrium? To discuss this, let us reconsider the game given in Figure 2 and the strategy profile x^* defined in the proof of Proposition 3. We can interpret this mixed strategy profile not so much as a probability distribution over the actual pure strategies chosen, but rather a belief of opponent players about this player’s choice of pure strategy. In the case at hand, the argument would be as follows. Player 1 chooses A because player 1 does not know how player 2 would interpret a deviation to either B or C , which player 2 cannot distinguish. Player 1 might think it is equally likely that player 2 will react by playing L or by playing R (because presumably player 2 has a clear assessment of player 1’s intended choice, only player 1 does not know what that is). Thus, this equilibrium, interpreted as beliefs about opponent strategies, is completely reasonable.⁵ In fact, this is a weak perfect Bayesian equilibrium in this game if we interchange the two players’ essentially simultaneous moves. Consider the thus modified game in Figure 3.

In this game x^* is weak perfect Bayesian. In fact, it is even sequential, given that every information set is reached. This demonstrates the point made by

5. Of course, this game has additional, also reasonable, equilibria, which are also fixed points of σ .

Refined Best-Response and Backward Induction




Figure 3 This game is derived from the game in Figure 2 by interchanging the two players' essentially simultaneous moves.




Figure 4 This game is derived from the game in Figure 2 by splitting player 1's move into two sequential moves.

Battigalli (1996) that weak perfect Bayesian equilibria are not invariant under the interchanging of essentially simultaneous moves. Fixed points of σ , on the other hand, naturally satisfy this invariance property as they are defined for the agent normal form of the game, which does not change when such essentially simultaneous moves are interchanged.

We should note, however, that fixed points of σ in the agent normal form are not immune to splitting information sets into parts, as the same example demonstrates (Figure 4).

Of course, one problem is that this game has a different agent normal form, as it now has three players. However, that alone is not necessarily a problem. The problem is that in any strategy profile of the agent normal form of this game, player 1's choice between B and C now has to be specified. This is equivalent to specifying player 1's possible deviation to B and C in the original game. Thus player 2 in this game is no longer free to interpret how play arrived at her information set. While it may be a deviation of the first player 1 to play X, it does not take a deviation of the second player 1 to get to this information set.

APPENDIX A.

SEMI-REDUCED NORMAL FORM

In this appendix we show by example that none of the results in this paper extend to the semi-reduced normal form even of extensive form games of perfect information in which for each player there are no two final nodes at which this player receives the same payoff.

First, note that not every normal form derived from an extensive form game, even if it is of perfect information and such that for each player there are no two final nodes at which this player receives the same payoff, is in \mathcal{G}^* . Consider the 1-player extensive form game, given in Figure A1, in which at node 1 the player has two choices, L and R , where L terminates the game, while R leads to a second node, where the player again faces two choices l and r . The two pure strategies Ll and Lr are obviously equivalent. Indeed the set Ψ as defined in the definition of the class of games \mathcal{G}^* is empty for this game, and, thus, far from dense in the space of all strategy profiles.

The semi-reduced normal form has been introduced to eliminate exactly this type of equivalences. Moreover, that the semi-reduced normal form is in \mathcal{G}^* , has been shown, see Balkenborg *et al.* (2013, Proposition 3), to be a generic property in the space of all extensive form games.

Consider the centipede game, Figure 8.2.2 in Cressman (2003), given here in Figure A2 with semi-reduced normal form given as Game 1, where player 1's strategies are $A = Ll|Lr$, $B = Rl$, and $C = Rr$, while player 2's strategies are $D = Rl|$




Figure A1 A 1-player extensive form game.




Figure A2 A centipede game. The normal form game of the centipede game in Figure A2.

Refined Best-Response and Backward Induction

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

Game 1 The normal form game of the centipede game in Figure A2.

Rr , $E = Ll$, and $F = Lr$. This game has a unique subgame perfect equilibrium, which is (Lr, Lr) . The non-subgame perfect, non weak-perfect Bayesian, and, hence, non-sequential, Nash equilibrium (B, D) is a fixed point of σ . So indeed, fixed points of σ in a given normal form game need not induce sequential or even weak perfect Bayesian equilibria in every extensive form game with this semi-reduced normal form.

APPENDIX B.

FORWARD INDUCTION

In this section, we provide a (well-known) example that demonstrates that rationalizability based on, and fixed points of, the refined best reply correspondence have no relation to forward induction solutions (Figure B1).

The strategy profile (O, R) is a fixed point of σ . This is so because O is best against an open set of strategy profiles close to player 2's R and R is best against an open set of strategy profiles close to O (in which player 1 uses B sufficiently more than A). This emphasizes the fact that fixed points of σ are exclusively about backward induction, as any deviation from the equilibrium play is essentially interpreted as a tremble or mistake. This is, thus, fundamentally different from forward induction reasoning. According to forward induction reasoning O, R should not be played, because a deviation of player 1 into player 2's information set should be interpreted by player 2 as a clear attempt to go for the other equilibrium A, L . This is so, if player 2 tries to maintain as much as possible her original hypothesis that her opponent, player 1, is rational, see Battigalli and Siniscalchi (2002). For fixed points of σ any such deviation is simply understood as a mistake.

This also implies that σ -rationalizability, thus, has not much in common with extensive form rationalizability, as in Battigalli (1997) and Pearce (1984), which, as shown by Battigalli and Siniscalchi (2002), is more related to forward induction than to backward induction.




Figure B1 Battle of the Sexes with an outside option.

ACKNOWLEDGEMENTS

We thank Carlos Alos-Ferrer, Pierpaolo Battigalli, Eddie Dekel, Amanda Friedenberg, Drew Fudenberg, Olivier Gossner, Michael Greinecker, Klaus Ritzberger, Karl Schlag, Mark Voorneveld, Jörgen Weibull, and especially two anonymous referees for helpful comments and suggestions.

Address for correspondence: Christoph Kuzmics, Department of Economics, University of Graz, 8010 Graz, Austria. Tel.: +43 316 380 7111; fax: +43 316 380 69 7113; e-mail: christoph.kuzmics@uni-graz.at

REFERENCES

- Balkenborg, D., J. Hofbauer and C. Kuzmics (2013), 'Refined Best-Response Correspondence and Dynamics', *Theoretical Economics* 8, 165–192.
- Balkenborg, D., J. Hofbauer and C. Kuzmics (2015), 'The Refined Best-Response Correspondence in Normal Form Games', *International Journal of Game Theory*, 44, 165–193.
- Battigalli, P. (1996), 'Strategic Independence and Perfect Bayesian Equilibria', *Journal of Economic Theory*, 70, 201–234.
- Battigalli, P. (1997), 'On Rationalizability in Extensive Games', *Journal of Economic Theory*, 74, 40–61.
- Battigalli, P. and M. Siniscalchi (2002), 'Strong Belief and Forward-Induction Reasoning', *Journal of Economic Theory*, 106, 356–391.
- Battigalli, P. and M. Siniscalchi (2003), 'Rationalization and Incomplete Information', *Advances in Theoretical Economics*, 3, Article 3.
- Bernheim, B. D. (1984), 'Rationalizable Strategic Behavior', *Econometrica* 52, 1007–1029.
- Bonanno, G. (1995), 'A Characterization of Sequential Equilibrium', *Economic Notes* 24, 109–128.
- Brandenburger, A. and E. Dekel (1987), 'Rationalizability and Correlated Equilibria', *Econometrica*, 55, 1391–1402.
- Cressman, R. (2003), *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge, MA.
- Dekel, E., D. Fudenberg and D. K. Levine (1999), 'Payoff Information and Selfconfirming Equilibrium', *Journal of Economic Theory* 89, 165–185.
- Fudenberg, D. and D. K. Levine (1993), 'Self-Confirming Equilibrium', *Econometrica*, 61, 523–546.
- Fudenberg, D. and J. Tirole (1991a), *Game Theory*. MIT Press, Cambridge, MA.
- Fudenberg, D. and J. Tirole (1991b), 'Perfect Bayesian and Sequential Equilibrium', *Journal of Economic Theory*, 53, 236–260.
- Kohlberg, E. and P. J. Reny (1997), 'Independence on Relative Probability Spaces and Consistent Assessments in Game Trees', *Journal of Economic Theory*, 75, 280–313.
- Kreps, D. M. and R. Wilson (1982), 'Sequential Equilibria', *Econometrica*, 50, 863–894.
- Mas-Collel, A., M. D. Whinston and J. R. Green (1995), *Microeconomic Theory*. Oxford University Press, Oxford, UK.
- Pearce, D. G. (1984), 'Rationalizable Strategic Behavior and the Problem of Perfection', *Econometrica*, 52, 1029–1051.
- Ritzberger, K. (2002), *Foundations of Non-Cooperative Game Theory*. Oxford University Press, Oxford, UK.
- Selten, R. (1965), 'Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträchtigkeit', *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–324.
- Selten, R. (1975), 'Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games', *International Journal of Game Theory*, 4, 25–55.