Maastricht University

Invariance properties of persistent equilibria and related solution concepts

Citation for published version (APA):

Vermeulen, A. J., Jansen, M. J. M., & Balkenborg, D. (2001). Invariance properties of persistent equilibria and related solution concepts. *Mathematical Social Sciences*, *41*, 111-130. https://doi.org/10.1016/S0165-4896(00)00056-1

Document status and date:
Published: 01/01/2001

DOI:
10.1016/S0165-4896(00)00056-1

Document Version:
Publisher's PDF, also known as Version of record

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the “Taverne” license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy
If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Download date: 25 Apr. 2024

ELSEVIER

Mathematical Social Sciences 41 (2001) 111–130

mathematical social sciences

www.elsevier.nl/locate/econbase

Invariance properties of persistent equilibria and related solution concepts

Dieter Balkenborg[b], Mathijs Jansen[a,*], Dries Vermeulen[a]

[a]*Maastricht University, Department of Quantitative Economics, P.O. Box 616, 6200 MD Maastricht, The Netherlands*

[b]*Department of Economics, School of Business and Economics, University of Exeter, Streatham Court, Rennes Drive, Exeter, Devon EX4 4PU, UK*

Received 2 April 1998; received in revised form 30 September 1999; accepted 25 January 2000

Abstract

Kohlberg and Mertens (1986) argued that a solution concept to a game should be invariant under the addition or deletion of an equivalent strategy and not require the use of weakly dominated strategies. In this paper we study which of these requirements are satisfied by Kalai and Samet's (1984) concepts of persistent equilibria and persistent retracts. While none of these concepts has all the invariance properties, we show that a slight rephrasing of the notion of a persistent retract leads to a notion satisfying them all. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Noncooperative games; Solutions; Invariance; Ordinality

JEL classification: C72

1. Introduction

Kalai and Samet (1984) introduced the notion of a persistent retract to select among the Nash equilibria of a game. They call a Nash equilibrium contained in a persistent retract a persistent equilibrium. However, persistent retracts are also of interest in themselves because they can be viewed as sets of rules that are self-enforcing in a strong sense. By a rule we mean hereby a restriction of the possible behavior of a player to a subset of his strategies, not necessarily to a single strategy. One can think for instance of

*Corresponding author. Tel.: +31-43-388-3834; fax: +31-43-388-4874.
E-mail address: j.jansen@ke.unimaas.nl (M. Jansen).

0165-4896/01/$ – see front matter © 2001 Elsevier Science B.V. All rights reserved.
PII: S0165-4896(00)00056-1

traffic rules like ‘‘drive on the right side of the road’’. Traffic rules restrict the behavior of car drivers, but they are not intended to determine their behavior in all circumstances, as a single strategy would. Similarly participation in a market typically requires adherence to certain rules which still leave room for individual decisions. We cannot expect rational players to follow a set of rules unless it is self enforcing in the sense that no player has an incentive to violate the rules as long as he expects everyone else to follow the rules. The idea to consider self-enforcing sets of rules rather than just self enforcing strategy profiles, i.e. Nash equilibria, is appealing because it will often be impossible or undesirable to agree on social conventions that perfectly determine behaviour. A set of rules is strongly self-enforcing if no player wants to violate them even if he has a slight uncertainty about everyone else following the rules. Whereas Nash equilibria are self-enforcing considering strongly self-enforcing sets of rules leads to a different, set valued concept, the persistent retracts. The latter are more akin to some set-valued concepts discussed in the literature like the primitive formations used in the theory of equilibrium selection by Harsanyi and Selten (1988), the curb sets of Basu and Weibull (1991) and the sets closed under better replies by Ritzberger and Weibull (1995). Balkenborg (1992) discusses a unified approach to compare these concepts. Whereas simple learning processes do often not converge to Nash equilibrium, see e.g. Fudenberg and Levine (1998), there are a number of positive convergence results for these set valued concepts, see e.g. Hurkens (1995), Sanchirico (1996) and Ritzberger and Weibull (1995). Hurkens (1995) in particular discusses a simple stochastic learning process which always converges to a persistent retract.

What distinguishes persistent retracts from the other set valued concepts are their invariance properties. Kohlberg and Mertens (1986) argued that a solution concept should not depend on modelling details. In particular they made the following two requirements on solution concepts for normal form games: (a) The solutions should be invariant to the addition or deletion of payoff-equivalent strategies. (b) The solutions should not prescribe the use of weakly dominated strategies. Mertens (1987) went further and argued that we lack a theory of ordinal equivalence of games and that the solutions we consider should respect ordinal equivalence. In the theory he developed ordinal invariance can be derived from invariance requirements closely related to those above. However, (b) is replaced by the slightly different requirement that the solutions should only depend on the ‘‘admissible best reply structure’’ of a game which we will use here.

Regardless of whether one considers these requirements as indispensable or not, it is certainly useful to know whether a solution concept satisfies them or how it violates them. We discuss here the invariance properties of persistent equilibria and related concepts. For other solution concepts we refer to Mertens (1987) and Vermeulen and Jansen (1997). Some results on the invariance properties of persistent retracts and equilibria are already given in Mertens (1987). They are included in this paper for completeness.

To obtain existence of persistent equilibria, Kalai and Samet require a persistent retract to be convex. If one is primarily interested in self-enforcing sets of rules, this restriction is not necessary. Without the convexity requirement one obtains a slightly different concept which we call ‘‘persistent set’’. However, convexity still plays an

indirect role, via the formation of beliefs. We will show that persistent retracts can be viewed as the set of beliefs associated with persistent sets.

To select among equilibria as sharply as possible, the definition of a persistent retract entails a minimality condition. As in Basu and Weibull (1991) one can replace minimality by a weaker condition called ''tightness'' without altering our main conclusions.

	Persistent or tight absorbing sets	Persistent or tight absorbing retracts	Persistent equilibria
Weak invariance	Yes	Yes	Yes
Invariance	Yes	No	No
abr-invariance	Yes	Yes	Yes for $n = 2$ No for $n > 2$

This table summarizes the results. We distinguish weak invariance, (full) invariance and *abr*-invariance. Weak invariance refers to the behavior of a solution when a payoff-equivalent strategy is deleted, invariance to the behavior when such a strategy is added or deleted. *abr*-invariance refers to invariance under the admissible best reply structure.

All concepts satisfy weak invariance. For persistent retracts and equilibria this was first noted by Mertens (1987). However, he showed also that they do not satisfy invariance. In contrast we show here that the notions of tight absorbing or persistent sets we introduce satisfy all the invariance properties and are hence ordinally invariant in the sense of Mertens (1987). Mertens observed that persistent retracts satisfy *abr*-invariance. We show in addition that persistent equilibria satisfy *abr*-invariance for two player games, but not for games with more than two players.

In the remainder of this paper we will discuss cell by cell each entry in the table. Some basic terminology for normal form games is fixed in Section 2. In Section 3 we introduce the solution concepts. Loosely speaking our main result in this section states that persistent retracts correspond to the beliefs associated with a persistent set. The proof is given in the Appendix. Section 4 provides the definitions of the various invariance properties along with a discussion of the underlying motivations. In Sections 5 and 6 these properties are studied for the various solution concepts. Section 7 concludes.

2. Preliminaries

For a natural number n, a (*finite n-person normal form*) *game* is a pair $\Gamma = \langle A, u\rangle$ such that $A := \prod_i A_i$ is a product of n non-empty, finite sets and $u = (u_i)_{i \in N}$ is an n-tuple of functions $u_i: A \to \mathrm{IR}$. A_i is called the set of *pure strategies* of player i and u_i is his *payoff function*.

As usual, a game $\Gamma = \langle A, u\rangle$ will be identified with its mixed extension. For this game, the *mixed strategies* of player i are the elements of the set $\Delta(A_i)$ of probability

distributions on A_i. By abuse of notation we will identify a pure strategy $a_i \in A_i$ with the mixed strategy in $\Delta(A_i)$ that puts all weight on a_i. So, A_i will simply be viewed as a subset of $\Delta(A_i)$. Pure strategy profiles will be denoted by $a \in A$.

For a (mixed) strategy profile $x = (x_i)_{i \in N} \in \Delta_A := \prod_{i \in N} \Delta(A_i)$ we define, as usual, the (expected) payoff function of player i by $u_i(x) := \Sigma_{a \in A} \prod_j x_j(a_j) u_i(a)$ where $x_j(a_j)$ is the probability with which a_j is played in x_j. Furthermore, $(x|y_i) \in \Delta_A$ denotes the strategy profile where player i uses $y_i \in \Delta(A_i)$ and his opponents use their strategies in $x \in \Delta_A$. For a strategy profile $x \in \Delta_A$

$$\beta_i(x) := \{ y_i \in \Delta(A_i) | u_i(x|y_i) \geq u_i(x|z_i) \quad \text{for all} \quad z_i \in \Delta(A_i)\}$$

is the set of *best replies* (of player i) to x. The strategy profiles in the set $\beta(x) := \prod_i \beta_i(x)$ are called best replies to x. A strategy profile $x \in \Delta_A$ is called a *Nash equilibrium* of Γ if $x \in \beta(x)$.

3. Persistent equilibria and related solution concepts

A non-empty subset $R \subset \Delta_A$ of strategy profiles is called a *product set* if $R = \prod_i R_i$ with $R_i \subset \Delta(A_i)$ for each player i. If the product set R is convex we call it a *retract*. Notice that a product set is convex if and only if each of its components is convex.

Let U be an open subset in Δ_A. We say that a set of strategies R_i of player i *absorbs* U if for every strategy profile in U there is a best reply against this profile in R_i. A product set R is said to *absorb* U, if each of its components R_i absorbs U.

Definition. A product set R is an absorbing set if it absorbs some convex open neighborhood of itself.

Remark 1. *The preceding definition requires a convex neighborhood for the following reason: a belief of a player $i \neq j$ over player j's strategy choice is a probability distribution over the mixed strategy space $\Delta(A_j)$. For the purpose of utility maximization this belief can be identified with the expected mixed strategy it induces. Player i is certain that player j will make a strategy choice in $R_j \subseteqq \Delta(A_j)$ if the support of his belief is in R_j. The expected mixed strategy will hence be in the closed convex hull of R_j. If we consider a small open neighborhood of this set, which we can itself assume to be convex, we allow for the possibility that player i is slightly uncertain about whether player j chooses a strategy in R_i or not. Thus a product set R is absorbing if no player i has an incentive to violate the rule "choose in R_i", as long as he is almost certain that his opponents follow the rule "choose in R_j".*

We can now proceed with the definition of the solution concepts considered in this paper:

Definitions. A *persistent set* is an absorbing product set which does not properly contain another absorbing product set.

A *persistent retract* (Kalai and Samet, 1984) is an absorbing retract which does not properly contain another absorbing retract.

A *persistent equilibrium* (Kalai and Samet, 1984) is a Nash equilibrium contained in a persistent retract.

Kalai and Samet (1984) showed that every game has a persistent retract and that every persistent retract contains a Nash equilibrium. In particular, a persistent equilibrium always exists.

If the primary interest is not in selecting among Nash equilibria, but in strongly self-enforcing sets of rules one can replace the minimality condition in the definition of a persistent set/retract as follows (compare Basu and Weibull, 1991):

Definition. An absorbing set (retract) R is called *tight* if no open convex neighborhood of R is absorbed by a product set (retract) properly contained in R.

Our main result in this section states that the persistent retracts correspond to the sets of beliefs associated with persistent sets. We need the following additional notation. The convex hull of a set X in some Euclidean space is denoted by $\mathrm{conv}(X)$. Conversely, the set of extremal points of a convex set Y (i.e. the set of points in the closure of Y which cannot be written as proper convex combinations of other points in the closure of Y) is denoted by $\mathrm{ext}(Y)$.

Proposition 1. *Let Γ be a game and let $R \subset \Delta_A$ be a product set. Then the following three statements are equivalent*

(*i*) *R is a persistent retract.*

(*ii*) *R is the convex hull of a persistent set.*

(*iii*) *R is closed and convex and* $\mathrm{ext}(R)$ *is a persistent set.*

The same holds if "persistent" is replaced by "tight absorbing" in the previous statement.

The proof, which uses and extends several central findings in Kalai and Samet (1984), is given in Appendix A. The central tool will be a careful analysis of the *minimal* sets of strategies of a player that absorb an open set.

It will further be shown in the appendix that every absorbing set contains a persistent set. From the above proposition it follows that every absorbing retract contains a persistent retract. Hence the persistent sets (retracts) are exactly the *minimal* tight absorbing sets (retracts).

Call two strategies of a player $x_i, x_i' \in \Delta(A_i)$ *player-equivalent* if they yield the same payoff for the player regardless of the strategies played by the opponents, i.e., if we have for all $y \in \Delta_A$

$$u_i(y|x_i) = u_i(y|x_i').$$

We may still have $u_j(y|x_i) \neq u_j(y|x_i')$ for some other player $j \neq i$.

There will be no player-equivalent pure strategies in our examples. The following further consequence from the results in the appendix is therefore useful when studying them:

Proposition 2. *A tight absorbing set is finite. If no player has two distinct player-equivalent pure strategies, then a tight absorbing set contains only pure strategy combinations.*

Example 1. We consider the following three games

	L	R
T	3,1	0,0
B	0,0	1,1

	L	R
T'	3,0	−1,0
T	3,1	0,0
B	0,0	1,1

	L	R
T	1,−1	−1,1
B	−1,1	1,−1

The Battle-of-the-Sexes game depicted on the left has two persistent sets, namely $\{(T, L)\}$ and $\{(B, R)\}$. There is another tight absorbing set, namely the set of all pure strategy combinations $\{T, B\} \times \{L, R\}$, but it is not minimal and hence not persistent.

These solutions do not change if we add a weakly dominated strategy T' for player 1 to obtain the game in the middle. Although (T, L) is no longer a strict equilibrium, i.e., it does not have a unique best reply, it is still a persistent equilibrium: As long as player 1 is almost certain that player 2 chooses L, he has no reason to switch to T'. However, the Nash equilibrium (T', L) is not even contained in any tight absorbing retract.

Finally, the unique persistent retract of the Matching-Pennies game on the right is the whole strategy space $\Delta\{T, B\} \times \Delta\{L, R\}$. Consequently the unique mixed strategy equilibrium of this game is persistent.

4. Invariance

For the purposes of this paper we will call any map which assigns to each game Γ a collection of subsets of Δ_A a *solution concept*. Kohlberg and Mertens (1986) called a solution concept invariant if, roughly speaking, the addition or deletion of pure strategies that are convex combinations of other pure strategies does not affect the solutions of a game. To make this precise, we have to compare (solution) sets contained in strategy spaces of different dimensions. The method suggested by Mertens (1987) uses reduction maps to identify strategy profiles of the original game Γ and the game Γ', where "duplicate" strategies are deleted. This method can be formalized as follows.

Let $\Gamma = \langle A, u\rangle$ and $\Gamma' = \langle B, v\rangle$ be two games. A map $f = (f_i)_{i \in N}$ from Δ_A to Δ_B is called a *reduction map* from Γ to Γ' (denoted by $\Gamma \xrightarrow{pay}_f \Gamma'$) if for every player $i \in N$

1. f_i: $\Delta(A_i) \to \Delta(B_i)$ is affine and onto
2. $u_i = v_i \circ f$.

The function $v_i \circ f$ denotes the composition of v_i and f. The function f links strategy profiles in the two games that are yielding each player the same payoff.

Definition. A solution concept σ is *weakly invariant* if for any triplet (Γ, Γ', f) with $\Gamma \overset{pay}{\to}_f \Gamma'$,

$$\sigma(\Gamma') = \{f(S) | S \in \sigma(\Gamma)\}.$$

This property states that every image under f of a solution of the game Γ must be a solution of the game Γ' and, conversely, that every solution of Γ' is the image under f of a solution of Γ.

Definition. A solution concept σ is called *invariant* if it is weakly invariant and, moreover, for all $T \in \sigma(\Gamma')$,

$$f^{-1}(T) = \cup\{S \in \sigma(\Gamma) | f(S) = T\}.$$

This additional property states that every strategy profile of the game Γ whose image under f is an element of a solution T of Γ' must be an element of a solution S of Γ whose image under f equals T.

In the following example we will clarify these concepts using a number of pictures. In the example payoff-equivalent strategies play an important role. We call two strategies x_i and y_i of player 1 payoff-equivalent if for all $z \in \Delta_A$ and all players j,

$$u_j(z | x_i) = u_j(z | y_i).$$

Note that in this situation no player cares whether player i plays x_i or y_i.

Example 2. In the picture below the triangle represents the strategy space of a game Γ. In fact one can think of it as being the strategy space of a game in which player 1 has three pure strategies while player 2 has only one pure strategy.

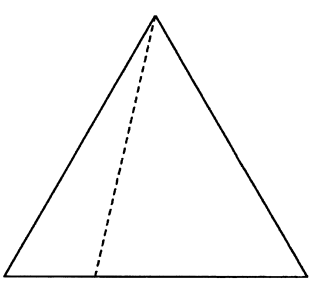

Now suppose that the topmost pure strategy is payoff-equivalent with the strategy at the bottom connected to it by the dotted line. So, the game Γ can be thought of as being

constructed from a 'smaller' game Γ' (in which player one has two pure strategies which correspond with the two endpoints of the bottom line) by the addition of that particular mixed strategy as a pure strategy. Then all strategies on the dotted line are payoff-equivalent.

Moreover, every pair of strategies that are on a line parallel to this dotted line are also payoff-equivalent. Thus, each strategy in the triangle is payoff-equivalent with (exactly) one strategy on the bottom line. Consequently, we can identify each mixed strategy with a unique mixed strategy in the smaller game. The mapping f that assigns to a mixed strategy in the triangle the unique mixed strategy on the bottom line is (part of) a reduction map.

In the next picture the fat line depicts a solution set S of Γ.

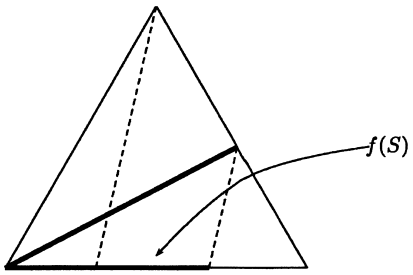

Now weak invariance roughly states that the collection $f(S)$ of strategy profiles that results from this identification must be a solution set of the smaller original game. (It also states that for each solution set T of the smaller game we conversely have a solution set S of the larger game such that $f(S) = T$.)

Invariance does do more than that though. Since in our example we have that the topmost strategy is a copy of a strategy at the bottom we know that in particular the strategies x and y indicated below are payoff-equivalent.

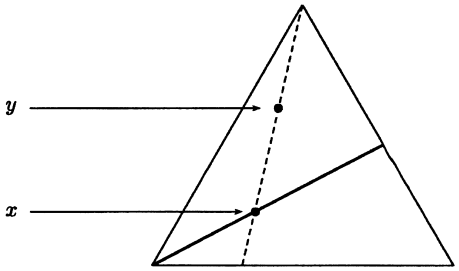

A first, naive, intuition of invariance now is the idea that the set S^* in the second picture below, obtained by the interchange of x and y (i.e. it is the union of the punctured fat line and the single strategy y_i), should also be a solution set. After all, from a pure game-theoretic point of view, we can argue as follows. Given that we, players, consider payoff-equivalent strategies as 'identical', and all of us do so, then it should not matter whether x or y is played. So, if we think that S is a good solution, than we must consider the set S^* equally good.

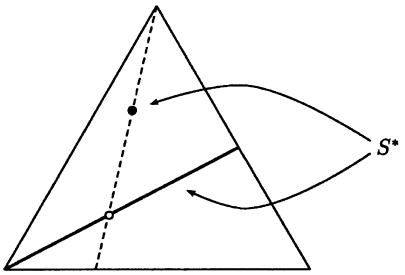

This requirement is very strong though, since it implies that at least some solution sets are not closed as we can clearly see in the above figure. Even worse, there is no solution we know of that satisfies this requirement.

There are several ways to get around this technical problem (there seems to be no clear game theoretic argument against this stronger requirement). One of them is the following.

Let S be a solution set of Γ. Further, let x be a strategy profile in S and let y be a strategy profile that is payoff-equivalent with x (i.e., x_i is payoff-equivalent with y_i for every i). Then there exists a solution set S^* such that y is an element of S^* *and moreover:* a strategy profile z is payoff-equivalent to an element of S if and only if it is payoff-equivalent to an element of S^*. We express the latter condition by saying that S is payoff-equivalent to S^*.

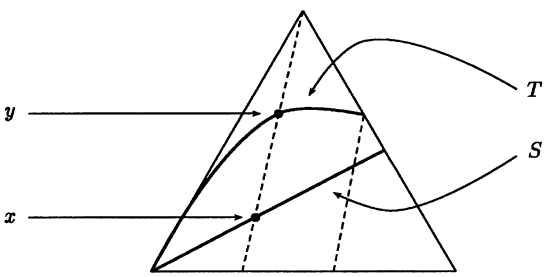

In the above figure T does contain the strategy (profile) y. However, it does not satisfy the additional condition, since it does not intersect as many equivalence classes as S does. The set S^* in the figure below does satisfy both conditions.

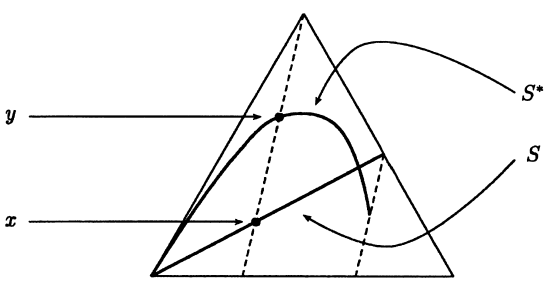

Finally we remark that invariance is (if we disregard relabeling of pure strategies) equivalent to weak invariance and the type of ‘‘shift’’ invariance mentioned above.

For a more detailed elaboration we refer to Mertens (1987) and Vermeulen and Jansen (2000).

5. Invariance of solution concepts

The two invariance properties will now be checked for the solution concepts introduced in the previous section. We will repeatedly use the following (straightforward) Lemma:

Lemma 1. *If f is a reduction map from a game Γ to a game Γ', then y is a best reply to x in the game Γ if and only if $f(y)$ is a best reply to $f(x)$ in the game Γ'.*

Let (Γ, Γ', f) be a triplet with $\Gamma \overset{pay}{\rightarrow}_f \Gamma'$. Let V' be a convex open neighborhood of a product set R' in the strategy space $\Delta_{A'}$ of the game Γ' and let V be a convex neighborhood of a product set R in the strategy space Δ_A of the game Γ such that $f(V) = V'$ and $f(R) = R'$.

Lemma 2. *R absorbs V in the game Γ if and only if R' absorbs V' in Γ'. Consequently, R is absorbing if and only if R' is absorbing.*

Proof. (a) Suppose that R' absorbs V' in Γ'. We will show that R absorbs V. To this end, take a strategy profile $x \in V$. Then obviously $f(x) \in V'$. So, since R' absorbs V', there exists a best reply $y' \in R'$ to $f(x)$. Furthermore, there exists a strategy profile $y \in R$ with $f(y) = y'$ since $f(R) = R'$. Then, by Lemma 1, y is a best reply to x.

(b) Conversely, suppose that R absorbs V in Γ. We will show that R' absorbs V'. To this end, let x' be a strategy profile in V'. Then there exists a strategy profile x in V with $f(x) = x'$, since $f(V) = V'$. Furthermore, there exists a best reply $y \in R$ to x, since R absorbs V. Then $f(y) \in R'$ is a best reply to x' by Lemma 1.

(c) To prove the second statement of the Lemma, suppose that R' absorbs a convex open neighborhood V' of itself. Let R be a product set with $f(R) = R'$. Then $V := f^{-1}(V')$ is a convex open neighborhood of R, since Δ_A is convex and f is affine and onto. Furthermore, $f(V) = V'$. Hence, R absorbs V by (a). Conversely, suppose that the product set R absorbs a convex open neighborhood V of itself. Then $V' := f(V)$ is a convex open neighborhood of the product set $R' = f(V)$. Hence, R' absorbs V' by (b). □

Proposition 3. *Persistent sets and tight absorbing sets are invariant.*

Proof. (a) Let (Γ, Γ', f) be a triplet with $\Gamma \overset{pay}{\rightarrow}_f \Gamma'$. Assume first that R is a tight absorbing or persistent set of the game Γ. Then the product set $R' := f(R)$ is absorbing by Lemma 2. In order to show that R' is also tight absorbing or persistent, suppose that

$S' \subset R'$ is a product set absorbing a convex neighborhood V' of R'. First note that $S := f^{-1}(S') \cap R$ is a product set since both $f^{-1}(S')$ and R are product sets. Furthermore, since $S' \subset R'$ and $R' = f(R)$, we obtain $f(S) = S'$. Therefore, by Lemma 2, S is a product set absorbing the convex open neighborhood $V := f^{-1}(V')$ of R. If R is a persistent set, it cannot contain a proper product set absorbing V and hence $S = R$ since $S \subset R$ by definition. If R is a tight absorbing set, we obtain $S = R$ as well, since S absorbs any sufficiently small neighborhood of R. We conclude in both cases that $S' = f(S) = f(R) = R'$. Hence, R' is also either tight absorbing or persistent.

(b) Suppose now that R' is a tight absorbing or persistent set of the game Γ'. Fix $x_0 \in f^{-1}(R')$ and define $y_0 := f(x_0) \in R'$. Choose any product set R in Δ_A which contains x_0 and for which $f_{|R}: R \to R'$ is a bijection. By Lemma 2, R is an absorbing set of Γ. In order to show that R is tight absorbing or persistent, let $S \subset R$ absorb a convex neighborhood of R. Then $f(S) \subset R'$ absorbs the open neighborhood $f(V)$ of R' by Lemma 2. As in (a) we conclude $f(S) = R'$ since R' is tight absorbing or persistent. Since $S \subseteqq R$ and since $f_{|R}: R \to R'$ is a bijection we conclude $R = S$. Therefore R is tight absorbing or persistent. □

Mertens (1987) already noticed that persistent retracts and persistent equilibria are weakly invariant.

Proposition 4. *Tight absorbing retracts, persistent retracts and persistent equilibria are weakly invariant.*

Proof. Notice that Nash equilibria are invariant by Lemma 1. We will only give a proof for tight absorbing retracts, since the proof for persistent retracts is completely analogous. Let (Γ, Γ', f) be a triplet with $\Gamma \xrightarrow{pay}_f \Gamma'$.

(a) Suppose that R is a persistent retract of Γ. Then $\mathrm{ext}(R)$ is a persistent set of Γ by Proposition 1. So, by Proposition 3 we know that $f(\mathrm{ext}(R))$ is a persistent set of Γ'. Hence,

$$f(R) = f(\mathrm{conv}(\mathrm{ext}(R))) = \mathrm{conv}(f(\mathrm{ext}(R)))$$

is a persistent retract of Γ' by Proposition 1.

(b) Suppose that R' is a persistent retract of Γ'. Then $\mathrm{ext}(R')$ is a persistent set of Γ' with convex hull R' by Proposition 1. So, there exists a persistent set S of Γ with $f(S) = \mathrm{ext}(R')$ by Proposition 3. Hence, $\mathrm{conv}(S)$ is a persistent retract of Γ by Proposition 1, while $f(\mathrm{conv}(S)) = \mathrm{conv}(f(S)) = \mathrm{conv}(\mathrm{ext}(R')) = R'$. □

The next example shows that persistent or tight absorbing retracts are **not** invariant.

Example 3. (Mertens, 1987) Persistent retracts and equilibria are not invariant. Consider the games in the following figure:

	L	R
T	$1,-1$	$-1,1$
B	$-1,1$	$1,-1$

	L	R
T	$1,-1$	$-1,1$
B	$-1,1$	$1,-1$
B'	$0,0$	$0,0$

The first game Γ' is again Matching pennies, where the unique persistent (tight absorbing) retract of the first game is $\Delta(A_1') \times \Delta(A_2')$. In the second game Γ the mixed strategy $(\frac{1}{2}, \frac{1}{2})$ of player 1 is duplicated. Let $f = (f_1, f_2)$: $\Delta_A \to \Delta_{A'}$ be the affine map where f_2 equals the identity on $\Delta(A_2) = \Delta(A_2')$ and where for $(q_1, q_2, q_3) \in \Delta(A_1)$

$$f_1(q_1, q_2, q_3) := (q_1 + \tfrac{1}{2}q_3, q_2 + \tfrac{1}{2}q_3).$$

(q_3 denotes the probability for the added strategy.) We have $\Gamma \xrightarrow{pay}_f \Gamma'$. Propositions 1 and 3 imply that $\{q \in \Delta(A_1) | q_3 = 0\} \times \Delta(A_2)$ is the only persistent or tight absorbing retract of the second game. $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is the only persistent equilibrium of the game Γ'. $((0, 0, 1), (\frac{1}{2}, \frac{1}{2}))$ is a Nash equilibrium which is mapped by f onto the persistent equilibrium in Γ, although it does itself not belong to a persistent retract or even a tight absorbing retract. Hence neither persistent or tight absorbing retracts nor persistent equilibria are invariant. Although in this example the original persistent (tight absorbing) retract remains persistent (or tight absorbing), we believe that the violation of invariance here is 'severe'. Since $(0, 0, 1)$ and $(\frac{1}{2}, \frac{1}{2}, 0)$ are payoff-equivalent strategies of player 1, neither player cares whether $(0, 0, 1)$ or $(\frac{1}{2}, \frac{1}{2}, 0)$ is played by 1 *no matter what player 2 does*. So, from a game-theoretical point of view, these strategies are indistinguishable. Hence, since $(\frac{1}{2}, \frac{1}{2}, 0)$ is part of a solution – a persistent retract/equilibrium in this case – one would want $(0, 0, 1)$ also to be in a solution. This is not the case. The main reason why persistent retracts lack invariance is the identification of mixed strategies[2] with beliefs mentioned in Remark 1. In this case $(\frac{1}{2}, \frac{1}{2}, 0)$ is an element of the (unique) persistent retract because it reflects a belief over T and B. There are convincing arguments why T and B should be in the solution (persistent set) of the game. Hence, the belief $\frac{1}{2}T + \frac{1}{2}B$ is credible. However, the strategy B', although payoff-equivalent to $(\frac{1}{2}, \frac{1}{2}, 0)$, is not a convex combination of likewise justified (pure) strategies. Therefore $(0, 0, 1)$ is not in the persistent set and hence persistent retracts lack invariance. See also the discussion of Example 4. □

6. *abr*-Invariance

In order to define the third notion of invariance considered in this paper, we first introduce an equivalence relation for games with the **same** strategy spaces. Two games $\Gamma = \langle A, u \rangle$ and $\Gamma^* = \langle A, u^* \rangle$ are called *admissible-best-reply-equivalent* (*abr*-equivalent)

[2]Our main argument here is that we think this identification is not as inoccuous as is commonly assumed. Notice e.g. the two different interpretations of the equilibrium $((\frac{1}{2}, 0, \frac{1}{2}, 0), c)$ in Example 4.

if, for every completely mixed strategy profile $x \in \Delta_A$, the set of best replies to x in the game Γ coincides with the set of best replies to x in the game Γ^*.

Definition. A solution concept σ is called *abr-invariant* if

$$\sigma(\Gamma) = \sigma(\Gamma^*)$$

for any pair of *abr*-equivalent games Γ and Γ^*.

Remark 2. *Abr-invariance requires first of all that a solution depends only on the best-reply structure of a game and not on the actual utilities of the players. For instance, the solution concept which assigns to a game the set of all Pareto-efficient Nash equilibria is not abr-invariant.*

Secondly, a solution may actually depend only on the admissible best reply correspondence of a game. Hereby a strategy x_i is called an admissible best reply against a strategy combination y if there exist a sequence of strategy combinations $(y_k)_{k\geq 1}$ converging to y such that x_i is a best reply against each strategy combination in the sequence. (*Notice the similarity to the definition of a normal-form perfect equilibrium in Selten*, 1975.) *Because a weakly dominated strategy is never a best reply against a completely mixed strategy combination, abr-invariance implies that a solution does not use weakly dominated pure strategies.* (*One could always decrease the payoffs to such a strategy such that it becomes strictly dominated. The resulting game would be abr-equivalent.*) *An abr-invariant solution concept may however contain weakly dominated mixed strategies, as we will show in Example* 4 *below. Thus abr-invariance differs slightly from the admissibility requirement discussed in Kohlberg and Mertens* (1986). *See Mertens* (1987) *and Vermeulen and Jansen* (1997) *for further discussions on the notion of abr-invariance.*

Lemma 3. *Suppose a product set P absorbs a non-empty open set V of strategy combinations in a game Γ. Then P also absorbs V in any abr-equivalent game Γ^*.*

Proof. Let V^0 be the set of all completely mixed strategy combinations in V. Then V^0 is open and dense in V. Since P absorbs V^0 in Γ and Γ^* and since the best reply correspondence is upper hemi-continuous, it follows that P absorbs V in Γ^*. □

Using the lemma it follows straightforwardly from the definitions:

Proposition 5. *Tight absorbing sets or retracts and persistent sets or retracts are abr-invariant.*

Proposition 6. *Persistent equilibria are abr-invariant for two-player games.*

Proof. Let $x = (x_1, x_2)$ be a persistent equilibrium of the game Γ. We have to show that x is also a persistent equilibrium of any *abr*-equivalent game Γ^*. By Proposition 1 each strategy x_i is in the convex hull of a set R_i, where $R_1 \times R_2$ is a persistent set of Γ. By Proposition 5, $R_1 \times R_2$ is also a persistent set in Γ^*. Thus it remains to show that x is a

Nash equilibrium of Γ^*. For this purpose we write each x_i as a convex combination $\Sigma_{z_i \in R_i'} \alpha_i(z_i) \cdot z_i$ where R_i' is a subset of R_i and each coefficient $\alpha_i(z_i)$ is strictly positive. The claim follows, if we can show that every $z_i \in R_i'$ is a best reply to x in Γ^*.

First note that each $z_i \in R_i'$ is a best reply to x in Γ. Moreover, every strategy $z_i \in R_i'$ is robust by Proposition 7 (in the appendix), i.e., it is a best reply against a non-empty open subset of strategies in Γ. Then it must also be a best reply against some completely mixed strategy combination y in Γ. So, since for a two-player game the function $u_i(y'|z_i)$ is linear in y', the strategy z_i must be a best reply in Γ against each of the completely mixed strategy profiles $y(\lambda) := (1-\lambda) \cdot x + \lambda \cdot y$ with $0 < \lambda \leq 1$. Consequently, it is a best reply against each $y(\lambda)$ $(0 < \lambda \leq 1)$ in the *abr*-equivalent game Γ^*. By continuity, z_i is also a best reply against $x = y(0)$, which concludes the proof. □

Example 4. Although persistent retracts and (for two-player games) persistent equilibria are *abr*-invariant, they may use weakly dominated strategies: for the game

	a	*b*	*c*
A	6,0	0,0	0,1
B	4,0	4,0	0,1
C	0,0	6,0	0,1
D	3,0	3,0	0,1

the unique persistent set is $\{A, B, C\} \times \{c\}$ and the unique persistent retract is the convex hull of this set. Thus the equilibrium $((\frac{1}{2}, 0, \frac{1}{2}, 0), c)$ is persistent, although it is not perfect because the strategy $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ is weakly dominated by B.[3]

The fourth pure strategy D, payoff equivalent to $(\frac{1}{2}, 0, \frac{1}{2}, 0)$, is added to demonstrate to what extent persistent retracts may violate invariance. Note that the strategy D is not in the persistent retract. To understand why $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ but not D can belong to a persistent retract it is useful to distinguish between beliefs over (possibly mixed) strategies[4] and the induced "expected" mixed strategy. Suppose player 2 is certain that player 1 is choosing in the persistent set. However, he is uncertain about player 1's strategy choice and assigns equal probabilities to strategies A and C and zero probability to strategy B. Then, although player 2 is certain that player 1 is not choosing a weakly dominated strategy, the "expected" mixed strategy $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ induced by his belief is

[3]Kalai and Samet (1984) showed in a three-person game that persistent equilibria do not have to be perfect. They also showed that there is at least one perfect persistent equilibrium.

[4]In the interpretation indicated above, where the mixed strategy $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ merely reflects the uncertainty of player 2, player 2's belief is very "sound" and there is no obvious reason why the equilibrium should be discarded. This interpretation of an equilibrium is in fact the one favoured by Aumann and Brandenburger (1995) and it also underlies Harsanyi's (1973) justification for equilibria in mixed strategies. The other interpretation is the one where player 1 actually chooses the weakly dominated strategy $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ and player 2 is certain that he does. For this interpretation the "trembling hand" argument underlying Selten's (1975) notion of perfection yields a compelling argument to discard the equilibrium.

weakly dominated. In contrast, the only belief inducing the expected strategy D is the belief that player 1 is choosing D with certainty, which is inconsistent with the assumption that player 1 is not choosing a weakly dominated strategy. In general, a weakly dominated strategy may belong to a persistent retract provided it represents a belief which does not violate the assumption that players do not choose weakly dominated strategies on purpose. Since only weakly dominated mixed strategies, but not weakly dominated pure strategies, can satisfy this condition, the first may belong to a persistent retract, but not the latter. Since one can always add a pure strategy to duplicate any mixed one, invariance is necessarily violated.

In consequence, if one insists on the convexity requirement one can not avoid that the solution may contain weakly dominated strategies and may violate invariance. On the other hand, if one takes the view that persistent sets are the ''proper'' solution concept and that persistent retracts merely describe the beliefs associated with these solutions, invariance and very strong admissibility properties are naturally satisfied. □

Finally, the solution concept which assigns to every game the set of all persistent equilibria is **not** *abr*-invariant for games with more than two players:

Counterexample. Consider the following three-person game Γ:

e_1^1:

	e_3^1	e_3^2
e_2^1	1,0,0	−2,0,0
e_2^2	−2,0,0	1,1,1

e_1^2:

	e_3^1	e_3^2
e_2^1	0,1,1	0,0,0
e_2^2	0,0,0	0,0,0

e_1^3:

	e_3^1	e_3^2
e_2^1	−1,1,1	0,0,0
e_2^2	0,0,0	1,0,0

Player 1 picks the matrix, player 2 the row and player 3 the column. The second game Γ^* differs from Γ in the sense that player 1's payoffs are slightly changed if he takes his third strategy:

e_1^3:

	e_3^1	e_3^2
e_2^1	−2,1,1	0,0,0
e_2^2	0,0,0	2,0,0

Obviously, (e_1^1, e_2^2, e_3^2) is a Nash equilibrium of the first, but not of the second game. We will show next that Γ and Γ^* are *abr*-equivalent and then that the whole set of strategy profiles is a persistent retract in both games. Hence, the above strategy profile is a persistent equilibrium of Γ, but not of the *abr*-equivalent game Γ^*.

Γ and Γ^* are *abr* -equivalent: Since the payoffs for players 2 and 3 are not changed in the two games it suffices to show that the set of pure best replies of player 1 against any completely mixed strategy combination is the same in the two games. We can denote

a strategy combination for the opponents by $(p,\ q) \in [0,1]^2$, where $p = \text{prob}(e_2^2)$ and $q = \text{prob}(e_3^2)$. The following Figure illustrates for both games how player 1's payoff with one of his pure strategies e_1^l ($l \in \{1, 2, 3\}$) varies with the strategy combinations $(p,\ q)$ in the unit square $[0,\ 1]^2$.

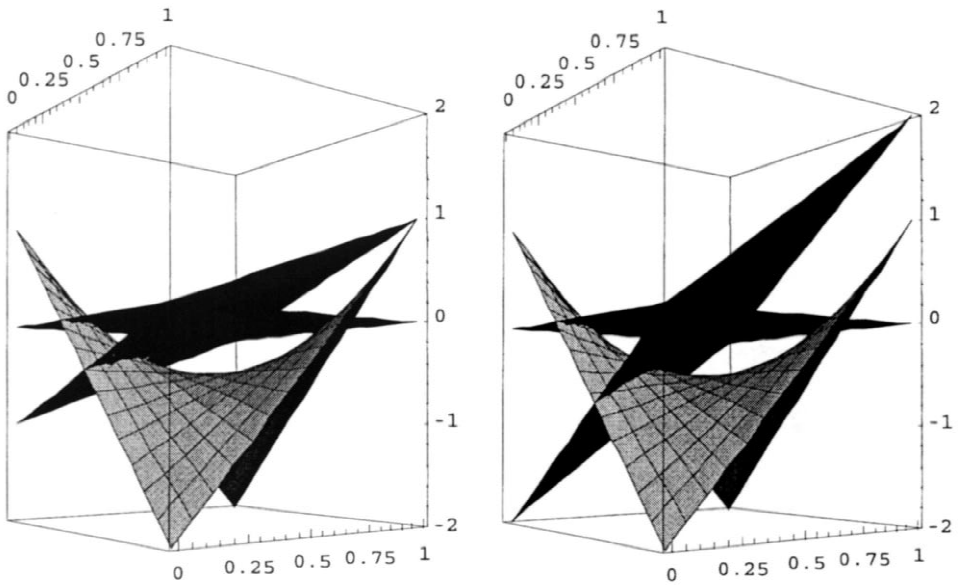

The graphic on the left illustrates this for the game Γ. The payoff with e_1^1 against $(p,\ q) \in [0,\ 1]^2$ is calculated as $6pq - 3p - 3q + 1$. The graph of this function is the quadratic surface in the picture. With e_1^2 player 1 obtains the constant payoff 0, which yields the horizontal plane. With e_1^3 he obtains $p + q - 1$, which yields the non-horizontal plane. When changing to Γ^*, the payoffs for e_1^1 and e_1^2 are not altered, while those for e_1^3 are changed to $2(p + q - 1)$. The corresponding plane is flipped upward in the graphic on the right and no longer intersects the quadratic surface above the point (1, 1).

Since e_1^3 yields strictly negative payoffs in the triangle $\{(p,\ q)|\ p + q < 1\}$ and payoff 0 along the diagonal $\{(p,\ q)|\ p + q = 1\}$ in both games, the best replies are the same in the region $\{(p,\ q)|p + q \leq 1\}$. It therefore remains to study the best replies on the triangle $\{(p,\ q)|1 < p + q\}$. Here e_1^3 yields strictly positive payoffs and is hence better than e_2^2. To compare with the payoffs for e_1^1 in Γ we calculate

$$\begin{aligned} 6pq - 3p - 3q + 1 < p + q - 1 \quad &\Leftrightarrow \quad pq - \tfrac{2}{3}p - \tfrac{2}{3}q + \tfrac{1}{3} < 0 \\ &\Leftrightarrow \quad pq - \tfrac{2}{3}p - \tfrac{2}{3}q + \tfrac{4}{9} < \tfrac{1}{9} \\ &\Leftrightarrow \quad (p - \tfrac{2}{3})(q - \tfrac{2}{3}) < \tfrac{1}{9}. \end{aligned}$$

The inequality is satisfied if $\frac{2}{3} \leq p \leq 1$, $\frac{2}{3} \leq q \leq 1$ and $(p, q) \neq (1,1)$. If $p < \frac{2}{3}$, the first factor in the last inequality is negative and the second factor positive since $q > 1 - p$. Similarly the inequality holds for $q < \frac{2}{3}$. Thus e_1^3 is the unique best reply on the region $\{(p, q) | 1 < p + q\}$ except for (1, 1), where e_1^1 is also a best reply. When changing to Γ^* the payoffs for e_1^3 are further increased on the region and e_1^3 becomes the unique best reply even against (1,1).

In summary, the best replies on $\{(p, q) | 0 < p < 1\} \times \{(p, q) | 0 < q < 1\}$ are the same in both games, i.e., the two games are *abr*-equivalent.

Determination of the persistent retract: It remains to be shown that the whole strategy space is a persistent retract in Γ and Γ^*. Since both games are *abr*-equivalent, it suffices to show that any persistent set of Γ^* must contain all pure strategy combinations. Since the games have no player-equivalent strategies, a persistent set can only contain pure strategy combinations.

Let $R = R_1 \times R_2 \times R_3$ be a persistent set of Γ^*. Let $V = V_1 \times V_2 \times V_3$ be a convex open neighborhood of R which is absorbed by R.

(i) Suppose R_1 contains e_1^1. Consider a strategy profile $(e_1^1, x_3) \in R_1 \times V_3$ with x_3 completely mixed. The unique best reply of player 2 against such a profile is e_2^2. Hence $e_2^2 \in R_2$. Similarly, $e_3^2 \in R_3$. Furthermore, e_1^3 is the unique best reply of player 1 to (e_2^2, e_3^2). Hence, $e_1^3 \in R_1$. But then it follows as before that $e_2^1 \in R_2$ and $e_3^1 \in R_3$. Therefore $R_2 = A_2$ and $R_3 = A_3$. One calculates that e_1^2 is the unique best reply against the strategy pair $((\frac{3}{5}, \frac{2}{5}), (\frac{3}{5}, \frac{2}{5})) \in V_2 \times V_3$. It is therefore contained in R_1. Thus $R = A_1 \times A_2 \times A_3$.

(ii) Suppose now that $e_1^2 \in R_1$ or that $e_1^3 \in R_1$. Then $(e_2^1, e_3^1) \in R_2 \times R_3$ and therefore $e_1^1 \in R_1$. Step (i) implies $R = A_1 \times A_2 \times A_3$, which concludes the discussion of the example. □

7. Conclusions

We have argued that persistent sets (or tight absorbing sets) can be viewed as strongly self-enforcing sets of rules. We have seen that these solution concepts satisfy all the invariance properties discussed here.

Mertens (1987) has noted that persistent retracts do not satisfy the invariance requirement related to the addition of a payoff-equivalent strategy. However, we have shown that persistent retracts correspond to the sets of beliefs associated with a persistent set. We therefore think that the lack of invariance of persistent retracts is caused by the identification of beliefs and mixed strategies.

The lack of invariance properties of persistent equilibria is more severe. It occurs because the concept does not refine among the equilibria within a persistent retract, although it restricts the set of equilibria to be considered. To satisfy all invariance properties it is hence necessary to amend further refinement requirements to the concept. For instance, Mertens (1992) has shown that every persistent retract contains a strategically stable set as reformulated in Mertens (1989). Such ‘‘persistent stable sets’’ satisfy all the invariance properties.

Appendix A

In the following we fix an arbitrary game Γ. We need some additional terminology before stating our first Lemma. A strategy $x_i \in \Delta(A_i)$ is called a *robust best reply* of player i against the strategy combination $y \in \Delta_A$, if x_i is a best reply against all strategy combinations in a neighborhood of x in Δ_A. Any strategy which is a robust best reply against *some* strategy combination is called *robust*. We call x_i a *semi-robust best reply* against y, if there exists a sequence $(y_k)_{k\geq 0}$ converging to y such that x_i is a robust best reply against each y_k. Finally, a *face* of the strategy simplex $\Delta(A_i)$ is the convex hull of a non-empty subset of pure strategies.

Lemma 4. (*i*) *The sets of robust best replies against a strategy combination $y \in \Delta_A$ is either empty or a full player-equivalence class that forms a face of the strategy simplex $\Delta(A_i)$.*

(*ii*) *There exists a semi-robust best reply against every strategy combination $y \in \Delta_A$.*

(*iii*) *If x_i is a semi-robust best reply against y then there is a strategy combination z arbitrarily close to y such that all best replies against z are player-equivalent to x_i.*

Proof. (i) Suppose x_i is a robust best reply against y. If x_i' is player-equivalent to x_i, then x_i' is a best reply whenever x_i is and therefore x_i' is a robust best reply against y. If x_i' is a robust best reply against y then x_i' is player-equivalent to x_i by Lemma 4 in Kalai and Samet (1984). Hence the set of all robust best replies against y is a player-equivalence class. Moreover, if x_i' is a robust best reply against y and a_i a pure strategy in the support of x_i', then a_i is a best reply against a strategy-combination whenever x_i' is. Consequently a_i is also a robust best reply against y and therefore player-equivalent to x_i' and hence to x_i. It follows that the player equivalence class of x_i is a face of $\Delta(A_i)$.

(ii) Consider any open neighborhood of y. By Lemma 5 in Kalai and Samet (1984) there exists a strategy combination y' in this neighborhood such that all best replies against y' are player-equivalent. Their proof reveals moreover that all best replies against y' are robust best replies against y'. We can consequently find a sequence $(y^k)_{k\geq 1}$ converging to y such that the set of best replies against each y^k is equal to the set of robust best replies against y^k. Since the set of best replies against y^k is a face of the strategy simplex and since $\Delta(A_i)$ has only finitely many faces, we can choose the sequence $(y^k)_{k\geq 1}$ such that the set of (robust) best replies against y^k consists of the same face F of strategies for all $k \geq 1$. Therefore each $x_i \in F$ is a semi-robust best reply against y.

(iii) Let x_i be a semi-robust best reply against y and choose an open neighborhood U of y. By definition, there exists $z \in U$ such that x_i is a robust best reply against z and therefore an open neighborhood $V \subseteqq U$ of z such that x_i is a best reply against each $z' \in V$. By Lemma 5 in Kalai and Samet (1984) we can choose $z' \in V$ such that all best replies against z' are player-equivalent to x_i. □

By the previous lemma, the set of robust strategies of a player is a non-empty disjoint union of faces of his strategy simplex. Each such face is a player-equivalence class. We denote the collection of these faces by $\mathscr{F}_i$. Notice that each $F \in \mathscr{F}_i$ will consist of a

single pure strategy, if the player has no pure player-equivalent strategies. We call a non-empty subset of strategies of a player *sparse*, if it contains at most one strategy from each of these faces. This is equivalent to: The non-empty set consists of robust strategies that are not player-equivalent. Since the strategy simplex has only finitely many faces, a sparse set of strategies is finite. Since these faces are all disjoint, no strategy in a sparse set can be a convex combination of other strategies in the set.

Let U be a non-empty open subset of Δ_A. We say that a subset of strategies for a player *tightly absorbs* U if it absorbs U while no proper subset absorbs U. A product set *tightly absorbs* U, if each of its components does.

Lemma 5. *Let U be a non-empty open subset of Δ_A. Then*:

(*a*) *Every strategy set of a player that absorbs U contains a subset that tightly absorbs U.*

(*b*) *A strategy set that tightly absorbs U is sparse.*

Proof. Suppose R_i absorbs U. Let R_i' denote the set of strategies in R_i that are semi-robust best replies against some strategy combination in U. Lemma 4 (ii) and (iii) implies that R_i' absorbs U. Let R_i'' be a set obtained by selecting for each face $F \in \mathscr{F}_i$ a single strategy in $F \cap R_i'$, if this intersection is not empty. Then R_i'' is sparse and absorbs U by construction. Moreover, the following argument shows that it tightly absorbs U: Let $x_i \in R_i''$. Then x_i is a semi-robust best reply against some $y \in U$. By Lemma 4 (iii) we can find $y' \in U$ such that all best replies against U are player-equivalent. Therefore x_i is the only best reply against y' contained in R_i''.

Suppose that moreover R_i tightly absorbs U. Then it follows $R_i = R_i''$. Hence R_i is sparse. □

The definitions imply:

Proposition 7. *A non-empty product set $R = \prod_{i=1}^{n} R_i$ in Δ_A is a tight absorbing set if and only if R tightly absorbs each sufficiently small convex open neighborhood of itself. In particular, each component R_i is sparse.*

Proposition 8. *Every absorbing set contains a persistent set.*

Proof. Let R be an absorbing set. Using Lemma 5 (a) we can find a product set $R' \subseteqq R$ that tightly absorbs some convex open neighborhood U of R. R' is finite since each of its components is sparse. If R' is not an absorbing set, then it will properly contain another absorbing set. If the latter is not absorbing, it will again properly contain some absorbing set etc. Since R' is finite, we will find after finitely many steps an absorbing set that does not properly contain an absorbing set, i.e., a persistent set. □

We have now all the tools to give the Proof of Proposition 1. We only give the proof for persistency. The proof for tight absorbing sets/retracts is similar.

(i) implies (ii): Let R be a persistent retract. By Proposition 8, R contains a persistent set $R' \subseteqq R$. Then $\mathrm{conv}(R') \subseteqq R$ is an absorbing retract and therefore $R = \mathrm{conv}(R')$.

(ii) implies (iii): Let $R = \mathrm{conv}(R')$, where R' is a persistent set. Since each R_i' is a sparse, it consists of finitely many strategies lying on disjoint faces of $\Delta(A_i)$. We obtain $\mathrm{ext}(R_i) = \mathrm{ext}(\mathrm{conv}(R_i')) = R_i'$ for each player i and hence $\mathrm{ext}(R) = R'$. Moreover, R is closed and convex.

(iii) implies (i): Since R is convex, $\mathrm{ext}(R) = R'$ and R' is finite, we obtain $R = \mathrm{conv}(R')$. An absorbing retract R'' contained in R must contain a persistent set by Proposition 7 But this persistent set must be R'. Hence $\mathrm{conv}(R') \subset R'' \subset R = \mathrm{conv}(R')$ since $\mathrm{conv}(R')$ and R are closed. Therefore R is a persistent retract. □

References

Aumann, R.J., Brandenburger, A., 1995. Epistemic conditions for Nash equilibrium. Econometrica 63, 1161–1181.

Balkenborg, D., 1992. The properties of persistent retracts and related concepts. PhD thesis, University of Bonn.

Basu, K., Weibull, J.W., 1991. Strategy subsets closed under rational behavior. Economic Letters 36, 141–146.

Fudenberg, D., Levine, D.K., 1998. The Theory of Learning in Games. MIT Press, Cambridge, MA.

Harsanyi, J.C., 1973. Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points. International Journal of Game Theory 2, 1–23.

Harsanyi, J.C., Selten, R., 1988. A General Theory of Equilibrium Selection in Games. MIT Press, Cambridge MA.

Hurkens, S., 1995. Learning by forgetful players. Games and Economic Behavior 11, 304–329.

Kalai, E., Samet, D., 1984. Persistent equilibria in strategic games. International Journal of Game Theory 13, 129–144.

Kohlberg, E., Mertens, J.-F., 1986. On the strategic stability of equilibria. Econometrica 54, 1003–1037.

Mertens, J.-F., 1987. Ordinality in non cooperative games. Discussion paper, CORE.

Mertens, J.-F., 1989. Stable equilibria: a reformulation. Part I. Mathematics of Operations Research 14, 575–624.

Mertens, J.-F., 1992. The small worlds axiom for stable equilibria. Games and Economic Behavior 4, 553–564.

Ritzberger, K., Weibull, J.W., 1995. Evolutionary selection in normal form games. Econometrica 63, 1371–1399.

Sanchirico, C.W., 1996. A probabilistic model of learning in games. Econometrica 64, 1375–1393.

Selten, R., 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4, 25–55.

Vermeulen, A., Jansen, M.J.M., 1997. On the invariance of solutions of finite games. Mathematical Social Sciences 33, 251–267.

Vermeulen, A., Jansen, M.J.M., 2000. Ordinality of solutions of noncooperative games. Journal of Mathematical Economics 33, 13–34.