

Polyequilibrium[§]

Igal Milchtaich*

Polyequilibrium is a generalization of Nash equilibrium that is applicable to any strategic game, whether finite or otherwise, and to dynamic games, with perfect or imperfect information. It differs from equilibrium in specifying strategies that players do *not* choose and by requiring an after-the-fact justification for the exclusion of these strategies rather than the retainment of the non-excluded ones. Specifically, for each excluded strategy of each player there must be a non-excluded one that responds at least as well as the first strategy does to every profile of non-excluded strategies of the other players. A particular result (e.g., Pareto efficiency of the payoffs) is said to hold in a polyequilibrium if it holds for all non-excluded profiles. As such a result does not necessarily hold in any Nash equilibrium in the game, the generalization proposed in this work extends the set of justifiable predictions concerning a game's results. *JEL Classification*: C72.

Keywords: Polystrategy, Polyequilibrium, Coarsening of Nash equilibrium, Subgame perfection.

1 Introduction

A Nash equilibrium is a self-enforcing strategy profile. Each player i is assigned a strategy x_i that is an optimal choice for i if all the other players choose the strategies assigned to them. Viewed from a different perspective, a Nash equilibrium excludes all but a single strategy for each player. The exclusion is justified in that, if none of the other players chooses an excluded strategy, player i also has no incentive to do so; choosing an excluded strategy cannot make the player better off in comparison with choosing the unique non-excluded one.

The first, conventional view of Nash equilibrium generalizes to *rationalizability* (Bernheim 1984, Pearce 1984). A rationalizable strategy is a best response to some *belief* about each of the other players' play that assigns positive probability only to strategies that are themselves rationalizable. Thus, unlike Nash equilibrium, the self-referring rationalizability condition potentially involves a set of strategies for each player rather than a single strategy. The same is true for the related solution concept of *curb set* (for "closed under rational behavior"; Basu and Weibull 1991). However, whereas rationalizability provides justification for the inclusion of the strategies in a player's set, a curb set can be described as requiring justification for the exclusion of the strategies outside it, similarly to the above alternative view of Nash equilibrium.

[§] *Games and Economic Behavior* 113 (2019), 339–355.

* Department of Economics, Bar-Ilan University, Ramat Gan 5290002, Israel
igal.milchtaich@biu.ac.il /

Specifically, in a curb set, every excluded strategy is not a best response to any belief about each of the other players' play that assigns positive probability only to non-excluded strategies.¹

Polyequilibrium is similar to curb set in being an "excluding" set-valued solution concept but differs from it and from rationalizability in not involving beliefs, i.e., product probability distributions on other players' strategies. Furthermore, like pure-strategy Nash equilibrium, polyequilibrium is a purely ordinal concept, invariant to arbitrary player-specific increasing transformations of the payoff functions. It requires that, for each excluded strategy of each player, there is a non-excluded one that yields the same or higher payoff against *every* profile of non-excluded strategies. Note that this requirement is weaker than requiring the excluded strategies to be weakly dominated, because (a) it only considers strategies of the other players that are not themselves excluded and (b) it allows for selection, that is, choosing among equally good strategies.

In short, polyequilibrium can be described as a self-enforcing subgame. A subgame by definition restricts each player i to a designated set of allowable strategies, and polyequilibrium requires the restriction to be self-enforcing in the sense that every strategy x'_i outside the player's designated set has an adequate substitute within it: an allowable strategy x''_i that responds at least as well as x'_i does to every profile of allowable strategies for the other players. Note that this requirement is a substantially stronger kind of self-enforcement than that employed by another set-valued generalization of Nash equilibrium, the *Nash retract* (Kalai and Samet 1984). The latter's definition changes the order of logical quantifiers and only requires that, against any *given* profile of allowable strategies for the other players, every strategy x'_i has an adequate substitute x''_i among the player's allowable strategies.

Polyequilibrium and the corresponding notion of self-enforcement are essentially a straightforward generalization of Shapley's (1964) notion of *generalized saddle point* in the context of finite two-player zero-sum games. More precisely, generalized saddle point is a special case of strict polyequilibrium (see Section 2) and its weak version is a special case of polyequilibrium. See Section 5.

The condition defining polyequilibrium ostensibly allows any strategy of any player to be included – justification is only required for the excluded strategies. This lenience in the definition is counterbalanced by a unanimity requirement when it comes to stating that a particular property of the game's outcome holds in a particular polyequilibrium X . The property is said to hold only if *all* strategy profiles in X possess it; if only some of them do so, then the polyequilibrium is mute about whether or not the outcome of the game may be expected to have the property. This way, a polyequilibrium may be able to specify, or predict, certain results without singling out a unique strategy profile. It may specify, for example, that a particular player takes or does not take a particular action, that the payoffs are positive, that the outcome

¹ Another set-valued solution concept is *strategic stability* (Kohlberg and Mertens 1986). However, the solution in this case includes only strategy profiles that are themselves Nash equilibria.

is socially efficient, and so on. The polyequilibrium concept thus represents a somewhat different philosophy than Nash equilibrium and certain other solution concepts that are designed to be completely specific, or at least as specific as possible, about the players' play. A polyequilibrium does not have to satisfy any set requirements in terms of its predictive power. Indeed, the collection of all strategy profiles in a game is a polyequilibrium (which immediately settles the question of existence). A polyequilibrium is, however, as good as the predictions it makes. Crucially, which predictions are "good", or interesting, is not determined by any objective measure of interest but is entirely context-dependent and, ultimately, subjective.

The *raison d'être* of the polyequilibrium solution concept is that it supports a larger set of justifiable predications about the game's outcome than Nash equilibrium does. As indicated, an outcome, or result, corresponds to a set of strategy profiles that share a common property. The result holds in an equilibrium x or a polyequilibrium X if the latter is an element or a subset of that set, respectively.

Section 3 shows that the concept of polyequilibrium result is indeed a genuine extension of equilibrium result. Even very simple games, with finite, countably infinite or a continuum of strategies, may have an interesting polyequilibrium result that is not an equilibrium result, either because the result itself is easily justifiable but no single strategy profile possessing it is so (Example 1), or because of what are arguably merely technical reasons, like non-existence of best-response strategies (Example 2). On the other hand, there are games in which the collections of equilibrium and polyequilibrium results coincide. This property, which means that every polyequilibrium in the game includes at least one equilibrium, is dubbed PE-equivalence. A major question then is which kinds of games have this property.

Example 1 already indicates that PE-equivalence often does not hold for finite games (where only pure strategies are allowed). However, for the mixed extensions of finite games (which result from allowing mixed strategies), Section 4 shows that the answer is more nuanced. While PE-equivalence does not hold for all such games (Example 5), it does hold for the mixed extensions of *generic* finite games, specifically, those that have only finitely many mixed-strategy equilibria (Theorem 1). For zero-some games, this is not so. As Section 5 shows, all zero-sum games, whether finite or otherwise, that have at least one equilibrium satisfy PE-equivalence (Theorem 2).

An important kind of polyequilibrium result is polyequilibrium strategy, which is any strategy that a particular player uses in all strategy profiles in some polyequilibrium X . Thus, X justifies the use of the strategy without necessarily pinning down the other players' strategies. This, of course, is not unfamiliar: the use of any dominant strategy can be similarly justified. Section 6 shows that, more generally, if only one of a player's strategies survives a number of rounds of successive elimination of weakly dominated strategies, then it is a polyequilibrium strategy (Proposition 6). On the other hand, in a game with the best-response existence property (where a player can always best respond to whatever the others are doing), any strategy that is

eliminated during successive elimination of strictly dominated strategies is not a polyequilibrium strategy (Proposition 7).

The polyequilibrium concept is particularly natural in the context of dynamic games, where it provides a sound justification for specifying the players' actions in only some of their decision nodes or information sets. Section 7 shows that here, too, the set of polyequilibrium results may extend beyond the equilibrium results. For example, in a perfect-information extensive form game, a player may receive a positive payoff in some polyequilibrium but zero in every equilibrium (Example 9). However, if the game has a unique subgame perfect equilibrium, then the only subgame perfect polyequilibrium results are those holding in the subgame perfect equilibrium (Theorem 3).

2 Definitions and Basic Facts

A (strategic) game Γ is specified by a set of players and, for each player i , a nonempty set of strategies S_i and a payoff function u_i that determines i 's payoff for each strategy profile $x \in S \stackrel{\text{def}}{=} \prod_j S_j$. The game is *finite* if so are its set of players and each player's strategy set.

A strategy x_i'' of player i responds to a strategy profile x *at least as well* as strategy x_i' does if

$$u_i(x \mid x_i'') \geq u_i(x \mid x_i'), \quad (1)$$

where the argument on each side of the inequality is the strategy profile obtained from x by replacing i 's strategy x_i with the indicated one. Strategy x_i'' responds to a *set* of strategy profiles X at least as well as x_i' does if inequality (1) holds for all $x \in X$. If, in addition, at least one of the inequalities is strict or all of them are so, then x_i'' *weakly* or *strictly dominates* x_i' , respectively, relative to X . (The phrase "relative to X " may be dropped if the reference is to the entire set of strategy profiles, that is, $X = S$.) Strategy x_i'' is a *best response* to a strategy profile or to a set of strategy profiles if it responds to it at least as well as every other strategy of player i does. A strategy *profile* x'' responds to a strategy profile x at least as well as strategy profile x' does if (1) holds for all i , and responds to a set of strategy profiles X at least as well as x' does if the previous condition holds for all $x \in X$. A strategy profile is a best response to a strategy profile or a set of strategy profiles if it responds to it at least as well as every other strategy profile does.

For a player i in a game Γ , a *polystrategy* is any nonempty set of strategies, $\emptyset \neq X_i \subseteq S_i$. A polystrategy that is a singleton, $\{x_i\}$, may be identified with the strategy x_i . Player i 's entire strategy set S_i is referred to as the *trivial polystrategy*. A *polystrategy profile* X is a Cartesian product of polystrategies, one polystrategy X_i for each player i . In other words, it is a nonempty rectangular subset of S . If the subset is a singleton, $\{x\}$, then it may be identified with its single strategy profile x . Every polystrategy profile X defines a *subgame* of Γ , denoted Γ^X , in which the players are as in Γ but each player i can only choose among the strategies in X_i and his payoff function is the restriction of u_i to X . For polystrategy profiles X' and X'' with $X' \subseteq X''$, the *interval* $[X', X'']$ is the collection of all polystrategy profiles X with $X' \subseteq X \subseteq X''$.

Definition 1. A polystrategy profile X is a *polyequilibrium* if for every player i and strategy $x'_i \notin X_i$ there is some $x''_i \in X_i$ that responds to X at least as well as x'_i does, and it is a *strict polyequilibrium* if it satisfies the stronger requirement that x''_i strictly dominates x'_i relative to X . A polyequilibrium X is *simple* if there is some strategy profile that is a best response to X .

The polyequilibria and strict polyequilibria in a game are partially ordered by inclusion. The largest one is the *trivial polyequilibrium*, which includes all strategy profiles. A polyequilibrium or strict polyequilibrium is *minimal* if it does not contain any other polyequilibrium or strict polyequilibrium, respectively. A polyequilibrium is *small* or *large* if there is no polyequilibrium that can be obtained from X by the deletion or addition, respectively, of any single strategy of a single player.

The five facts below easily follow from the definitions.

Fact 1. A polystrategy profile that is a singleton, $\{x\}$, is a polyequilibrium or strict polyequilibrium if and only if its single element x is a (Nash) equilibrium or strict equilibrium, respectively.

Fact 2. A polystrategy profile X is a simple polyequilibrium if and only if some $x' \in X$ is a best response to X . Such a strategy profile x' is necessarily an equilibrium. Thus, simple polyequilibrium is a particularly simple generalization of equilibrium.

Fact 3. A sufficient condition for a polystrategy profile X in a game Γ to be a polyequilibrium is that all the strategy profiles in X are equilibria. However, this condition is not necessary. A polyequilibrium X satisfies it if and only if each player's payoff in the subgame Γ^X is independent of his own strategy, that is, $u_i(x) = u_i(x \mid x'_i)$ for all i , $x \in X$ and $x'_i \in X_i$.²

Fact 4. A sufficient condition for a polystrategy profile X in a subgame Γ' of a game Γ to be a polyequilibrium in Γ' is that X is a polyequilibrium in Γ . If the subgame is of the form $\Gamma' = \Gamma^{X'}$, where X' is a polyequilibrium in Γ , then this condition is also necessary.

Fact 5. For polystrategy profiles X' and X'' with $X' \subseteq X''$, $[X', X'']$ is an *interval of polyequilibria* (that is, all its elements are so) if and only if for every strategy profile x' there is some $x'' \in X'$ that responds to X'' at least as well as x' does. In the special case $X' = X''$, this fact reduces to a compact alternative description of polyequilibrium.

Note that the definitions and facts above are all ordinal in the sense that they are invariant to arbitrary increasing transformations of the players' payoff functions. None of them requires cardinal utilities or, fundamentally, *any* utilities, for everything could be alternatively formulated in terms of preferences over strategy profiles rather than payoffs. Correspondingly, the

² An example of such a polyequilibrium is the *dominance solution* of a dominance solvable finite game (Moulin 1979), which is obtained by the successive elimination of all weakly dominated strategies.

polyequilibrium concept does not involve randomization or beliefs and is thus a generalization of *pure-strategy* Nash equilibrium.

It could be argued that, with cardinal utilities, the definition of polyequilibrium should be extended to include also polystrategy profiles X where for each excluded strategy x'_i of a player i there is a *mixed* strategy that responds to X at least as well as x'_i does and whose support is included in X_i . At least, it may seem justifiable to exclude strategies that are strictly dominated by such a mixed strategy. However, an implicit assumption in this assertion is that the player is actually able to play the mixed strategy. But in this case, it (or an equivalent strategy) should have been included in the strategy set S_i , for otherwise the latter is mislabeled as it is not a complete specification of the player's possible choices. Put differently, allowing mixed strategies in a game Γ effectively turns it into another game, namely, the mixed extension Γ^* , and in this case, the relevant polyequilibria are those in Γ^* . The connections between the set of polyequilibria in a finite game Γ and in its mixed extension Γ^* are studied in Section 4.

2.1 Strategy Substitution

A polyequilibrium may be alternatively described in terms of strategy substitution. Whereas an equilibrium prescribes one, specific strategy for each player, a polyequilibrium may be viewed as a prescription of a suitable *substitute* for each of the player's strategies.

A prescription of substitute strategies for a player i is expressed by a function $\phi_i: S_i \rightarrow S_i$.³ Any profile of such functions, one for each player, defines a *substitution function* $\phi: S \rightarrow S$ by $(\phi(x))_i = \phi_i(x_i)$ for all i . A substitution function ϕ is *rational* if, for all i and x ,

$$u_i(\phi(x)) \geq u_i(\phi(x) \mid x_i).$$

The inequality means that it is acceptable for player i to use the recommended substitute $\phi_i(x_i)$ instead of strategy x_i if all the other players also follow their recommendations. This formulation differs from Definition 1 in that it combines the specification of the players' polystrategies with the justification for them. Specifically, the logical relation between the two concepts is as follows.

Fact 6. A polystrategy profile X is a polyequilibrium if and only if it is the image of some rational substitution function ϕ (that is, $\phi(S) = X$).

A number of continuity results immediately follow from the definition. For example, in a game where the players' strategy sets are topological spaces and their payoff functions are continuous (with respect to the product topology), any substitution function that is the pointwise limit of rational substitution functions is also rational.

³ It may be natural to require the function ϕ_i to be idempotent, which means that any strategy that is some other strategy's substitute is also its own substitute. Adding this requirement would not affect any of the assertions below.

3 Polyequilibrium Results

Polyequilibrium is a predictive solution concept, not a prescriptive or normative one. As a polyequilibrium generally includes multiple strategy profiles, there is no general sense in which it may be “played”. Instead, a polyequilibrium predicts certain outcomes, or results, of the players’ choice of strategies.

Formally, a *result* R in a game Γ is any set of strategy profiles. Its *negation* $\sim R$ is the complementary set $S \setminus R$. A result R *holds* in a polystrategy profile X if $X \subseteq R$, and it is a *polyequilibrium result* if it holds in some polyequilibrium in Γ . A result may also be specified implicitly, as a particular property or consequence of strategy profiles (for example, “player 1’s payoff is higher than 2’s payoff”). In this case, R is the collection of all strategy profiles with the specified property, so that the result holds in a polyequilibrium X if and only if all strategy profiles $x \in X$ have the property. In particular, a real number v_i is a *polyequilibrium payoff* for a player i if there is some polyequilibrium X with $u_i(x) = v_i$ for all $x \in X$, and a strategy x_i is a *polyequilibrium strategy* if there is some polyequilibrium X with $X_i = \{x_i\}$. A generalization of the first concept is *polyequilibrium payoff interval* for player i , which is any convex set of real numbers E such that “ i ’s payoff lies in E ” is a polyequilibrium result, that is, $u_i(X) \subseteq E$ for some polyequilibrium X . Another generalization is *limit polyequilibrium payoff*, which is any extended real number v_i (that is, a real number, ∞ or $-\infty$) such that every convex neighborhood of v_i is a polyequilibrium payoff interval for player i . Similar definitions may be applied to payoff vectors.

The concept of result may also be applied to special kinds of polyequilibria, and in particular to equilibria. Every equilibrium result is also a polyequilibrium result but not conversely. There is also a logical difference between the two concepts. For any result $R \neq S$, the proposition “ R holds in every polyequilibrium” is false (because the result does not hold in the trivial polyequilibrium) but the proposition “there does not exist a polyequilibrium where $\sim R$ holds” may or may not be false. Thus, the two propositions are not logically equivalent, even though they would be so if ‘polyequilibrium’ were replaced by ‘equilibrium’. The reason, of course, is the possibility that, in a polyequilibrium X , both R and its negation do not hold. For R that is the collection of all strategy profiles with a particular property, this is so if and only if some, but not all, strategy profiles in X have the property.

Definition 2. A game satisfies *PE-equivalence* if every polyequilibrium result is also an equilibrium result, equivalently, if every polyequilibrium in the game includes at least one equilibrium.

If a game satisfies PE-equivalence, then a polyequilibrium where all strategy profiles have a particular property exists if and only if some equilibrium has the property. In a game that does not satisfy PE-equivalence, some result that does not hold in any equilibrium holds in a polyequilibrium. The next three examples present such games.

Example 1. In the finite game

	L	C	R
T	$(1,1)$	$(0,0)$	$(0,0)$
M	$(0,0)$	$(2,3)$	$(3,2)$
B	$(0,0)$	$(3,2)$	$(2,3)$

(Basu and Weibull 1991) the only equilibrium payoff is 1 but $[2,3]$ is also a polyequilibrium payoff interval for each player (and “both players receive at least 2” is a polyequilibrium result). The former holds in the game’s unique (pure-strategy) equilibrium (T, L) and the latter in the polyequilibrium $\{M, B\} \times \{C, R\}$. If a third player were added to the game, who does not take any meaningful action and whose payoff is the average payoff of the original two players, then only 1 would be an equilibrium payoff for that player but $5/2$ would be a polyequilibrium payoff. The latter is obviously also a *mixed*-equilibrium payoff for all players (see Proposition 5 below).

Example 2. Bilateral trade. A buyer has to offer a price p to the owner of an item whose worth is 1 to the buyer and 0 to the seller, and the seller has to decide whether to sell at that price. The sensible strategy “accept any price greater than zero” is a weakly dominant strategy for the seller, yet it is not an equilibrium strategy because the buyer does not have a best response to it: offering any $p > 0$ is less profitable than offering, say, half that price. Thus, the intuitive idea that the buyer should offer “as little as possible”, or “an ϵ ”, is incompatible with the definition of equilibrium. However, the idea is compatible with polyequilibrium. Indeed, for any $0 < \epsilon \leq 1$, the buyer’s polystrategy $0 < p \leq \epsilon$ (“offer a positive price not higher than ϵ ”) and the seller’s strategy of accepting any positive price together constitute a polyequilibrium, where the result is that the item is sold at a positive price not higher than ϵ .

Example 3. In a symmetric two-player game, each player chooses a positive integer y and gets the payoff

$$z - \left| 1 - \frac{2z}{y} \right|,$$

where z is the number chosen by the other player. To receive his maximum payoff of z , a player must choose $y = 2z$. However, such a choice prevents the other player from receiving his maximum payoff of y (which would require $z = 2y$), and therefore an equilibrium does not exist. In fact, as the aggregate payoff is easily seen to be at most $y + z - 2$, even an ϵ -equilibrium does not exist, for all $0 < \epsilon < 1$. However, the nonexistence of equilibrium arguably does not reflect a significant misalignment of interests. In particular, if the players alternately escalate their “bids” by doubling that of their rival, both payoffs spiral upwards. This observation is reflected in the fact that for every $L \geq 2$ the (symmetric) polystrategy profile where both players’ polystrategy is $y > L$ (“choose a number greater than L ”) is a strict polyequilibrium, in which both of them receive at least L . Thus, infinity is a limit polyequilibrium payoff.

4 Finite Games and Their Mixed Extensions

The game in Example 1 has three polyequilibria, which coincide with the supports of its three mixed-strategy equilibria (one of which is pure) as well as with the game's three curb sets. However, such coincidences are not the rule. For example, the finite game

$$\begin{array}{c} T \\ M \\ B \end{array} \begin{array}{ccc} L & C & R \\ \left(\begin{array}{ccc} -1,2 & 0,0 & 2,-1 \\ 0,0 & 0,0 & 0,0 \\ 2,-1 & 0,0 & -1,2 \end{array} \right) \end{array} \quad (2)$$

has two supports of mixed-strategy equilibria, two polyequilibria and two curb sets, but none of the pairs coincides with another. The first, pure-strategy, equilibrium (M, C) is a polyequilibrium but is not a curb set. The support $\{T, B\} \times \{L, R\}$ of the second mixed-strategy equilibrium is a curb set but is not a polyequilibrium. The trivial polyequilibrium is (trivially) also a curb set.

By definition (Basu and Weibull 1991), a polystrategy profile X in a finite game is a curb set if and only if, for every *mixed*-strategy profile whose support is included in X , every (pure) strategy x_i of a player i that is a best response to the profile is included in X_i . The last example shows that a curb set may not even contain a polyequilibrium (as a subset), and vice versa. It follows from the next proposition that a *strict* polyequilibrium is a curb set,⁴ so in a finite game without payoff ties, all polyequilibria are curb sets. However, the reverse inclusion does not generally hold regardless of ties. The reason for this is the fundamentally stronger nature of the polyequilibrium condition. For X to be a polyequilibrium, it is not sufficient that each polystrategy X_i includes every best-response strategy to every mixed-strategy profile as above, or even to every similar correlated strategy, that is, a probability distribution on X .⁵ For example, $\{T, B\} \times \{L, R\}$ is a curb set in the game (2), because if the column player, for example, only mixes between L and R , then T is a best response for the row player or B is so. But the set is not a polyequilibrium, because neither T nor B responds to *both* L and R at least as well as M does.

Proposition 1. A necessary condition for a polystrategy profile X in a finite game to be a polyequilibrium is that, for every probability distribution on X and every player i , some strategy that is a best response to the distribution is included in X_i . A necessary condition for X to be a strict polyequilibrium is that *every* such best-response strategy is in X_i . The second, stronger condition is *sufficient* for X to be a curb set but it is not sufficient for it to be a polyequilibrium.

Proof. Consider a probability distribution on a polyequilibrium X , and some strategy x'_i of a player i that is a best response to it. Let $x''_i \in X_i$ be a strategy that responds to X at least as well

⁴ This result aligns with the fact that the concept of curb set is meant as a generalization of strict, rather than Nash, equilibrium. Note, however, that this meaning of curb set is not universally accepted. See van Damme (2002, p. 1568).

⁵ A strategy x'_i is a best response to a *distribution* if player i 's expected payoff $\mathbb{E}u_i(x | x'_i)$, where x is a random element of X with that distribution, cannot be increased by replacing x'_i with any other strategy.

as x'_i does, so that inequality (1) holds for all $x \in X$. Replacing x on both sides of the inequality with a random element x of X with the given distribution and taking expectations preserves the inequality, which shows that x''_i too is a best response to the distribution. Moreover, if x''_i strictly dominates x'_i relative to X , then it actually affords player i a higher expected payoff against the distribution than x'_i does. However, the conclusion contradicts the assumption that the latter is a best-response strategy. Therefore, if X is a strict polyequilibrium, then x'_i must lie in X_i . The last part of the proposition is proved above. ■

Proposition 2. Every polyequilibrium in a finite game contains (although it does not necessarily coincide with) the support of a mixed-strategy equilibrium, but not every such support contains a polyequilibrium.

A (pure) strategy that is played with positive probability in a mixed-strategy equilibrium is rationalizable (Bernheim 1984, Pearce 1984). Therefore, an immediate corollary of Proposition 2 is that, in every polyequilibrium, each player's polystrategy includes at least one rationalizable strategy. In particular, every polyequilibrium strategy is rationalizable. Put differently, if a strategy x_i of a player i is *not* rationalizable, then playing it is not a polyequilibrium result. However, *not* playing x_i may also not be a polyequilibrium result. For example, in the finite game

$$\begin{array}{c} T \\ M \\ B \end{array} \begin{array}{cc} L & R \\ \left(\begin{array}{cc} 3, -3 & 0, 0 \\ 1, -1 & 1, -1 \\ 0, 0 & 3, -3 \end{array} \right) \end{array} \quad (3)$$

strategy M is strictly dominated by a *mixed* strategy and is therefore not rationalizable, but there is no polyequilibrium where player 1's polystrategy does not include M , as the only polyequilibrium in the game is the trivial one. Thus, the polystrategy profile $\{T, B\} \times \{L, R\}$, which represents all strategies in the game that are rationalizable, and is also the unique minimal curb set and the support of the unique mixed-strategy equilibrium, does not even contain a polyequilibrium.

The negative assertions in Proposition 2 are demonstrated by game (3). The positive result is an immediate corollary of Proposition 4 below, which identifies a particular connection between polyequilibria in a finite game Γ and in its mixed extension Γ^* .

By definition, an unqualified 'strategy' in Γ or Γ^* is a pure or mixed strategy, respectively, in Γ , and similarly for 'equilibrium'. As the collection S_i of all (pure) strategies for a player i in Γ may be viewed as a subset of the player's strategy set S_i^* in Γ^* , and similarly for the collections S and S^* of all strategy profiles, Γ may be viewed as a subgame of Γ^* , with the same symbol u_i denoting the payoff function of player i in both games. This (standard) view identifies the set of equilibria in Γ with a subset of that in Γ^* , namely, the pure-strategy equilibria.

The relation between the polyequilibria in the two games is more complex. A polystrategy X_i for a player i in Γ^* is any nonempty subset of S_i^* , that is, a collection of mixed strategies in Γ . It is a *pure* polystrategy if $X_i \subseteq S_i$. By the first part of Fact 4, every *pure-polystrategy* polyequilibrium

in Γ^* (that is, one in which all the polystrategies are pure) is a polyequilibrium also in Γ . However, as the following example and proposition show, the converse is false. In particular, S , the trivial polyequilibrium in Γ , is usually not a polyequilibrium in Γ^* . Thus, the identity between the equilibria in Γ and those equilibria in Γ^* that only involve pure strategies does not extend to polyequilibria. This difference between equilibrium and polyequilibrium means that, for the latter, even if the particular strategies examined are all pure, it still matters whether or not mixed strategies are admissible alternatives, that is, whether they *could* be used.

Example 4. The mixed extension of the game (3) has two small polyequilibria, which are the equilibrium $x' = (1/2 T + 1/2 B, 1/2 L + 1/2 R)$ and

$$X' = \{ (pT + (1 - p)B, qL + (1 - q)R) \mid 0 \leq p, q \leq 1 \},$$

and two large polyequilibria, which are

$$X'' = \{ (pT + (1 - 2p)M + pB, 1/2 L + 1/2 R) \mid 0 \leq p \leq 1/2 \}$$

and the trivial polyequilibrium S^* . Its entire set of polyequilibria is the union of two disjoint intervals: the interval of simple polyequilibria $[\{x'\}, X'']$ and the interval of strict polyequilibria $[X', S^*]$. Thus, there is no pure-polystrategy polyequilibrium.

Proposition 3. A polyequilibrium in a finite game Γ is a polyequilibrium also in the mixed extension Γ^* if and only if it is simple. In particular, the trivial polyequilibrium in Γ is a polyequilibrium also in Γ^* if and only if every player has a (pure) strategy in Γ that is a best response to all strategy profiles.

The proof of the proposition uses the following result.

Lemma 1. In the mixed extension of a finite game, every finite polyequilibrium is simple.

Proof. For a finite polyequilibrium X , there exists for each player i a strategy $x'_i \in X_i$ that is not weakly dominated relative to X by any other strategy in X_i . Thus, every strategy $x_i \in X_i$ is of one of two kinds: either (i) $u_i(\bar{x} \mid x_i) < u_i(\bar{x} \mid x'_i)$ for some $\bar{x} \in X$, or (ii) $u_i(\bar{x} \mid x_i) = u_i(\bar{x} \mid x'_i)$ for all $\bar{x} \in X$. Again by the finiteness of X , there is some $0 < \epsilon < 1$ such that for every strategy $x_i \in X_i$ that is of the first kind there is some $\bar{x} \in X$ such that $u_i(\bar{x} \mid x_i) < u_i(\bar{x} \mid (1 - \epsilon)x'_i + \epsilon\tilde{x}_i)$ for all $\tilde{x}_i \in S_i^*$. Consider any strategy $\tilde{x}_i \in S_i^*$. Since X is a polyequilibrium, there is some $x_i \in X_i$ that responds to X at least as well as $(1 - \epsilon)x'_i + \epsilon\tilde{x}_i$ does. By definition of ϵ , x_i cannot be of the first kind above, and it is therefore of the second kind. It follows that x'_i too responds to X at least as well as $(1 - \epsilon)x'_i + \epsilon\tilde{x}_i$ does, and therefore at least as well as \tilde{x}_i does. The conclusion shows that each player i has a strategy (namely, x'_i) that is a best response to the polyequilibrium X , which means that the latter is simple. ■

Proof of Proposition 3. If a polyequilibrium X in Γ is simple, then by Fact 2 every player i has a strategy $x'_i \in X_i$ that responds to X at least as well every (pure) strategy $x_i \in S_i$ does. By the linearity of u_i in player i 's own strategy, the same is true also for every (mixed strategy) $x_i \in S_i^*$,

which proves that X is a (simple) polyequilibrium also in Γ^* . Conversely, if a polyequilibrium X in Γ is a polyequilibrium also in Γ^* , then by Lemma 1 and Fact 2 some $x' \in X$ is a best response to X in Γ^* , and therefore also in Γ , which proves that X is simple. The second part of the proposition is a special case of the first part. ■

As indicated, the second part of Proposition 3 shows that the set of polyequilibria in Γ does not generally coincide with, but is rather a (usually, proper) superset of, the (possibly, empty) set of pure-polystrategy polyequilibria in Γ^* . There is, however, a simple, natural one-to-one correspondence between the former and another set of polyequilibria in Γ^* . This correspondence, which is indicated by the next proposition, matches each polyequilibrium in Γ with its *convex hull*.

For a polystrategy X_i of a player i in Γ^* , the convex hull of X_i , denoted $\text{conv } X_i$, is also a polystrategy in Γ^* . If X_i is pure (in other words, a polystrategy in Γ), $\text{conv } X_i$ consists of all mixed strategies whose supports are subsets of X_i . For a polystrategy profile $X = \prod_i X_i$ in Γ^* or (as a special case) in Γ , the convex hull of X is the polystrategy profile in Γ^* given by $\text{conv } X = \prod_i \text{conv } X_i$.

Proposition 4. For every polyequilibrium X in a finite game Γ or in its mixed extension Γ^* , $\text{conv } X$ is a polyequilibrium in Γ^* . This polyequilibrium (hence, every polyequilibrium in Γ^* that consists of convex polystrategies) includes at least one equilibrium (which is a mixed-strategy equilibrium in Γ).

Proof. Consider any polyequilibrium X in Γ or Γ^* . For each player i , select for each (pure) strategy $x'_i \in S_i$ some $x''_i \in X_i$ such that (1) holds for all $x \in X$, and let $X'_i \subseteq \text{conv } X_i$ be a polytope (that is, the convex hull of a finite number of strategies) that includes each of these (finitely many) strategies x''_i . Every (mixed) strategy in S_i^* is a convex combination of elements in S_i . It follows, by the linearity of u_i in player i 's own strategy, that for every $x'_i \in S_i^*$ there is some $x''_i \in X'_i$ such that (1) holds for all $x \in X$, and therefore, by the multilinearity of u_i , also for all $x \in \text{conv } X$. Since the polystrategy profile $X' = \prod_i X'_i$ is a subset of $\text{conv } X$, the last conclusion proves that these two polystrategy profiles are polyequilibria in Γ^* (see Fact 5). The subgame $\Gamma' = \Gamma^{X'}$ of Γ^* , where each player i is restricted to strategies in the polytope X'_i , is (identifiable with) the mixed extension of a finite game, and therefore has at least one equilibrium. By the second part of Fact 4, such an equilibrium is an equilibrium also in Γ^* . ■

As indicated, one corollary of Proposition 4 is Proposition 2. Another corollary is the following list of connections between the equilibrium and polyequilibrium payoffs in a finite game and in its mixed extension.

Proposition 5. For a player i in a finite game Γ , with mixed extension Γ^* , the following inclusions and equalities hold, and the inclusions may be strict:

$$\begin{aligned} \text{equilibrium payoffs in } \Gamma &\subseteq \text{polyequilibrium payoffs in } \Gamma \subseteq \text{mixed-equilibrium payoffs in } \Gamma \\ &= \text{equilibrium payoffs in } \Gamma^* = \text{polyequilibrium payoffs in } \Gamma^*. \end{aligned}$$

Moreover, in both Γ and Γ^* , every polyequilibrium payoff interval for player i includes at least one of the player's mixed-equilibrium payoffs in Γ .

Proof. Example 1 shows that both inclusions above may be strict. (In that finite game, $5/2$ is a polyequilibrium, but not an equilibrium, payoff for the additional, third player, and it is a mixed-equilibrium, but not a polyequilibrium, payoff for the original, row and column players, who get that expected payoff when they play $(1/2 M + 1/2 B, 1/2 C + 1/2 R)$.) The first inclusion and the first equality are trivial, and the second ones are special cases of the second part of the proposition. To prove the latter, consider a polyequilibrium payoff interval E for player i in either Γ or Γ^* and a corresponding polyequilibrium X , such that $u_i(X) \subseteq E$. By Proposition 4, there is some equilibrium payoff v_i for player i in Γ^* such that $v_i \in u_i(\text{conv } X) \subseteq \text{conv } u_i(X) \subseteq E$. ■

Proposition 5 shows that, in the mixed extension Γ^* of a finite game Γ , the only polyequilibrium payoffs are the equilibrium payoffs. It follows from the second part of Theorem 1 below that for a *generic* Γ a stronger proposition holds. Namely, all the polyequilibrium *results* in Γ^* are also equilibrium results. In other words, mixed extensions of finite games generically satisfy PE-equivalence. However, not *all* such games have this property. Specifically, as the next example shows, the inclusion indicated by the second part of Proposition 4 does not necessarily hold for polyequilibria with non-convex polystrategies.

Example 5. *A polyequilibrium that does not include an equilibrium.* Player 1 chooses whether to ask for \$1 or \$2. Player 2, who does not know the choice of player 1, has three options. He can do nothing and get \$1, or try to guess player 1's choice and get \$2 if he is correct. Player 1 gets the amount he asked for, unless player 2 guessed \$2, in which case player 1 *loses* the amount he asked for. Thus, the payoff matrix is

$$\begin{array}{c} \\ T \\ B \end{array} \begin{array}{ccc} L & C & R \\ \left(\begin{array}{ccc} 1,1 & 1,2 & -1,0 \\ 2,1 & 2,0 & -2,2 \end{array} \right) \end{array}.$$

In the mixed extension of the game, a strategy profile is an equilibrium if and only if players 1 and 2 play T and R , respectively, with probability 0.5. However, regardless of player 1's strategy, any strategy $(1 - p - p')L + pC + p'R$ of player 2 yields him the same payoff as the non-equilibrium strategy

$$(1 + t - 2f(t))L + f(t)C + (f(t) - t)R, \quad (4)$$

where $t = p - p'$ and the function f is defined by

$$f(t) = \begin{cases} (1 + t)/2, & -1 \leq t \leq -1/3 \\ 0, & -1/3 < t < 0 \\ t, & 0 \leq t \leq 1 \end{cases}. \quad (5)$$

Therefore, the following polystrategy profile $X = X_1 \times X_2$ is a polyequilibrium: X_1 consists of all strategies of player 1 and X_2 consists of all strategies of player 2 *except* the equilibrium ones.

A smaller, minimal polyequilibrium is obtained by including in X_2 only the strategies of the form (4), with $-1 \leq t \leq 1$ and $f(t)$ given by (5).⁶ It is easy to check that, in this polyequilibrium, player 1's payoff satisfies $|u_1| \geq 1/3$. His unique equilibrium payoff, by contrast, is 0.⁷

The game in Example 5 has infinitely many equilibria, but some polyequilibria do not include any of them. Somewhat paradoxically, with a *finite* set of equilibria, no polyequilibrium can be similarly shielded.

Theorem 1. If a polyequilibrium X in the mixed extension Γ^* of a finite game Γ does not include an equilibrium, then its convex hull $\text{conv } X$ must include infinitely many equilibria. Therefore, if Γ^* has only finitely many equilibria, then it satisfies PE-equivalence: every polyequilibrium includes at least one equilibrium.

Proof. It has to be shown that, for any polyequilibrium X in Γ^* and any finite set $A \subseteq \text{conv } X \setminus X$, the set $\text{conv } X \setminus A$ includes an equilibrium. If $A = \emptyset$, the inclusion follows from Proposition 4. Suppose then that $A = \{x^1, \dots, x^L\}$, with $L \geq 1$. For each $1 \leq l \leq L$, let \bar{x}^l be a strategy profile in X (hence, $\bar{x}^l \neq x^l$) that responds to X , and therefore also to $\text{conv } X$, at least as well as x^l does. For each player i , let $X'_i \subseteq \text{conv } X_i$ be a polytope as in the proof of Proposition 4, with the additional requirement that $\{x_i^1, \dots, x_i^L, \bar{x}_i^1, \dots, \bar{x}_i^L\} \subseteq X'_i$. As shown in that proof, every equilibrium x^* in the subgame Γ' of Γ^* obtained by restricting each player i to strategies in X'_i is an equilibrium also in Γ^* . Therefore, it suffices to show that some such equilibrium satisfies $x^* \notin A$. Note that, by construction, for every player i and $1 \leq l \leq L$

$$u_i(x \mid \bar{x}_i^l) \geq u_i(x \mid x_i^l), \quad x \in X', \quad (6)$$

where u_i is player i 's payoff function in Γ^* (and in Γ') and $X' = \prod_i X'_i$.

Claim 1. For each player i there is a continuous function $g_i: X'_i \rightarrow X'_i$ that satisfies the following two conditions:

$$u_i(x \mid g_i(x_i)) \geq u_i(x), \quad x \in X', \quad (7)$$

and, for every $1 \leq l \leq L$ with $\bar{x}_i^l \neq x_i^l$,

$$g_i(x_i) \neq x_i^l, \quad x_i \in X'_i. \quad (8)$$

The meaning of (7) is that changing player i 's strategy from any x_i to $g_i(x_i)$ cannot decrease his payoff in Γ' . The meaning of (8) is that (if $\bar{x}_i^l \neq x_i^l$) the image of g_i does not include x_i^l .

⁶ The function f essentially defines a rational substitution function for this polyequilibrium (see Section 2.1).

⁷ Note that, by the second part of Proposition 5, it would not be possible to find a similar example where player 1's *payoff* (rather than its absolute value) in some polyequilibrium X is greater than 0 while his unique equilibrium payoff is 0.

The function in Claim 1 is defined as $g_i = g_i^l \circ \dots \circ g_i^1$, the successive composition of L functions $g_i^l: X_i^l \rightarrow X_i^l$ ($l = 1, \dots, L$). These functions are defined by

$$g_i^l(x_i) = x_i + \alpha_i^l \phi_i^l(x_i)(\bar{x}_i^l - x_i^l), \quad (9)$$

where, if $\bar{x}_i^l \neq x_i^l$ (otherwise, the second term in (9) is zero), $0 < \alpha_i^l < 1$ is any constant that satisfies two requirements which are spelled out below and

$$\phi_i^l(x_i) = \max\{\alpha \geq 0 \mid x_i + \alpha(\bar{x}_i^l - x_i^l) \in X_i^l\}. \quad (10)$$

It is not difficult to see that the function $\phi_i^l: X_i^l \rightarrow [0, \infty)$ defined by (10) is continuous and satisfies

$$\phi_i^l(x_i) = \phi_i^l(x_i') + b \quad (11)$$

for every $x_i, x_i' \in X_i^l$ and real number b such that $x_i + b(\bar{x}_i^l - x_i^l) = x_i'$. In particular, $\phi_i^l(x_i') = \phi_i^l(\bar{x}_i^l) + 1$, which shows that the maximum M of the function ϕ_i^l satisfies $M \geq 1$. In addition, any strategy x_i satisfying $g_i^l(x_i) = x_i^l$ (hence, $x_i + (\alpha_i^l \phi_i^l(x_i) + 1)(\bar{x}_i^l - x_i^l) = \bar{x}_i^l$) also satisfies $\phi_i^l(x_i) = \phi_i^l(\bar{x}_i^l) + \alpha_i^l \phi_i^l(x_i) + 1$, and therefore

$$\alpha_i^l \leq 1 - 1/\phi_i^l(x_i) \leq 1 - 1/M.$$

The first requirement that the constant α_i^l has to satisfy is that it is, in fact, *greater* than $1 - 1/M$. This requirement guarantees that, if $\bar{x}_i^l \neq x_i^l$, then

$$g_i^l(x_i) \neq x_i^l, \quad x_i \in X_i^l. \quad (12)$$

The second requirement is that α_i^l is sufficiently close to (but smaller than) 1 to make the inequality $(1 - \alpha_i^l)M < \phi_i^l(x_i^{l'})$ hold for all $l' \neq l$ with $\phi_i^l(x_i^{l'}) > 0$. This requirement guarantees that for every $l' \neq l$

$$g_i^l(x_i) = x_i^{l'} \implies x_i = x_i^{l'}, \quad x_i \in X_i^l. \quad (13)$$

This is because, if $g_i^l(x_i) = x_i^{l'}$, then by (11) $\phi_i^l(x_i) = \phi_i^l(x_i^{l'}) + \alpha_i^l \phi_i^l(x_i)$, so that $\phi_i^l(x_i^{l'}) = (1 - \alpha_i^l)\phi_i^l(x_i) \leq (1 - \alpha_i^l)M$, which by the above requirement implies that $\phi_i^l(x_i^{l'}) = 0$, and therefore also $\phi_i^l(x_i) = 0$, so that $x_i = g_i^l(x_i) = x_i^{l'}$.

It follows from (6) and (9) that each of the functions g_i^l satisfies a condition similar to (7). Since by definition $g_i(x_i) = g_i^L(\dots(g_i^1(x_i))\dots)$, (7) itself clearly also holds. It remains to prove that (8) holds for every l with $\bar{x}_i^l \neq x_i^l$. Suppose that this is not so, which means that $g_i^L(\dots(g_i^1(x_i'))\dots) = x_i^{l'}$ for some strategy x_i' and some $1 \leq l' \leq L$ with $\bar{x}_i^{l'} \neq x_i^{l'}$. Necessarily, $l' \neq L$, since an equality would violate (12) for $l = L$. Therefore, by (13) (again with $l = L$), $g_i^{L-1}(\dots(g_i^1(x_i'))\dots) = x_i^{l'}$. A repeated use of the same argument now shows that the inequality $l' \neq l$ also holds for all $l < L$. This contradiction completes the proof of Claim 1.

Define a function $g: X' \rightarrow X'$ by $(g(x))_i = g_i(x_i)$ for all i . Construct a game $\bar{\Gamma}$ that has the same players and strategy sets as Γ' (but is not necessarily the mixed extension of a finite game) by assigning to each player i the payoff function \bar{u}_i defined by

$$\bar{u}_i(x) = u_i(g(x) \mid x_i).$$

The function \bar{u}_i is linear in player i 's own strategy x_i , and therefore the set $B_i(x)$ of best response strategies to any strategy profile x is a nonempty convex subset of the player's strategy set X'_i . The continuity of g and of (the multilinear function) u_i implies that the correspondence $x \mapsto B_i(x)$ has a closed graph. It therefore follows from Kakutani fixed-point theorem that $\bar{\Gamma}$ has some equilibrium \bar{x} . To complete the proof of the theorem, it remains to establish the following.

Claim 2. The strategy profile $x^* = g(\bar{x})$ satisfies $x^* \notin A$ and it is an equilibrium in Γ' .

Consider any $1 \leq l \leq L$. Since $\bar{x}^l \neq x^l$, there is some i such that $\bar{x}_i^l \neq x_i^l$. By (8), $g_i(\bar{x}_i) \neq x_i^l$, and therefore $x^* = g(\bar{x}) \neq x^l$. This proves that $x^* \notin A$. For every player i and strategy $x_i \in X'_i$,

$$\begin{aligned} u_i(x^*) &= u_i(g(\bar{x})) = u_i(g(\bar{x}) \mid g_i(\bar{x}_i)) \geq u_i(g(\bar{x}) \mid \bar{x}_i) \\ &= \bar{u}_i(\bar{x}) \geq \bar{u}_i(\bar{x} \mid x_i) = u_i(g(\bar{x}) \mid x_i) = u_i(x^* \mid x_i), \end{aligned}$$

where the first inequality follows from (7), the second inequality holds because \bar{x} is an equilibrium in $\bar{\Gamma}$, and all the equalities follow from the definitions. This proves that x^* is an equilibrium in Γ' . ■

5 Zero-Sum Games

Shapley (1964) called a strict polyequilibrium in a finite two-player zero-sum game a *generalized saddle point*, and called a polyequilibrium a *weak generalized saddle point*.⁸ He showed that, in a game of this kind, the intersection of any number of strict polyequilibria is also a strict polyequilibrium, so that the intersection of all of them, called the *saddle*, is the game's unique minimal strict polyequilibrium. A similar result does not hold for polyequilibria. Two of them may have a nonempty intersection that does not even contain a polyequilibrium, and every equilibrium is a minimal polyequilibrium. However, it follows as a conclusion from the next proposition that a *unique* equilibrium is necessarily also the game's unique minimal polyequilibrium. By Theorem 1, the same is true in every game that is the mixed extension of a finite game. However, the conclusion here and the proposition from which it follows concern any two-player zero-sum game: finite, the mixed extension of a finite game, or otherwise.

Theorem 2. A two-player zero-sum game satisfies PE-equivalence if and only if it has an equilibrium.

⁸ Duggan and Le Breton (1996) use the last term in a somewhat different sense.

Proof. It has to be shown that if the game has an equilibrium x' , with payoff vector $(v, -v)$, then every polyequilibrium X includes some equilibrium. In fact, a stronger conclusion holds: every strategy profile $x'' \in X$ that responds to X at least as well as x' does is an equilibrium. To prove this, it suffices to show that if player 1 plays x''_1 , then player 2 cannot get a higher payoff than $-v$, and if player 2 plays x''_2 , then player 1 cannot get a higher payoff than v . If there were, for example, a strategy x_1 of player 1 that yields a higher payoff than v against x''_2 , then this would be so also for every strategy in X_1 that responds to X at least as well as x_1 does. However, the existence of such a strategy in X_1 contradicts the assumption that strategy x''_2 responds to X at least as well as the equilibrium strategy x'_2 does, so that when player 2 plays x''_2 against any strategy in X_1 , he gets at least $-v$ and player 1 gets at most v . The contradiction proves that a strategy x_1 as above does not exist, so that x'' is indeed an equilibrium. ■

A finite two-player zero-sum game may or may not have a (pure-strategy) equilibrium. The latter holds, for example, for the 4×4 game where the row player's payoff matrix is

$$\begin{array}{c}
 T \\
 \\
 B
 \end{array}
 \begin{array}{c}
 L \\
 \\
 \\
 \\
 R
 \end{array}
 \begin{pmatrix}
 3 & 2 & 4 & 1 \\
 2 & 5 & -5 & 0 \\
 0 & -5 & 5 & 2 \\
 1 & 4 & 2 & 3
 \end{pmatrix}.$$

Its saddle (actually, the unique non-trivial polyequilibrium) is $\{T, B\} \times \{L, R\}$. Therefore, positive payoff for the row player is a polyequilibrium result in this game but not an equilibrium result.

The mixed extension of a finite two-player zero-sum game, that is, a *matrix game*, always has an equilibrium. It therefore follows from Theorem 2 that PE-equivalence holds for all matrix games: their sets of equilibrium and polyequilibrium results always coincide (which is not true for these games' non-zero-sum counterparts, the *bimatrix games*, as Example 5 demonstrates). The actual collections of equilibria and polyequilibria vary in their relative sizes. In rock-scissors-paper, the game's unique equilibrium is also its only non-trivial polyequilibrium. Other matrix games with a unique equilibrium, such as that in Example 4, have larger, richer sets of polyequilibria.

6 Successive Elimination of Strategies

In a polyequilibrium, the exclusion of strategies is given an after-the-fact justification. Each excluded strategy does not do better than a particular retained strategy of the same player against any profile of the other players' retained strategies. In this, exclusion differs from elimination of dominated strategies, which involves the stronger requirement that some alternative is better even relative to the collection of original strategy profiles. *Successive* elimination of dominated strategies blurs this distinction. The connections between successive elimination and polyequilibrium are explored below.

Successive elimination may involve weakly dominated, strictly dominated or never-best-response strategies. A player's strategy is a *never-best-response strategy* relative to a set of strategy profiles X (or, if X is the entire collection of strategy profiles S , simply "never-best-

response strategy”) if it is not a best response to any strategy profile in X . Such a strategy is not necessarily weakly dominated relative to X or vice versa. However, a strategy that is strictly dominated relative to X is clearly also a never-best-response strategy relative to it.

Definition 3. A polystrategy profile X is obtained by *successive elimination* of weakly or strictly dominated strategies if there is a nonincreasing finite sequence of polystrategy profiles $S = X^0 \supseteq X^1 \supseteq \dots \supseteq X^L = X$, with $L \geq 1$, such that for every $1 \leq l \leq L$, player i and (eliminated) strategy $x_i \in X_i^{l-1} \setminus X_i^l$ there is some $x'_i \in X_i^l$ that weakly or strictly dominates x_i , respectively, relative to X^{l-1} . Successive elimination of never-best-response strategies is defined similarly, except that the requirement concerning x_i is replaced by the requirement that it is a never-best-response strategy relative to X^{l-1} .

Clearly, every weakly dominant strategy is a polyequilibrium strategy, whereas a strictly dominated strategy, or more generally a never-best-response one, is not a polyequilibrium strategy. The next two propositions extend these observations to successive elimination.

Proposition 6. A polystrategy profile X that is obtained by (any number of rounds of) successive elimination of (any number of) weakly dominated strategies is a polyequilibrium.

Proof. With the notation of Definition 3, each X^l ($1 \leq l \leq L$) is clearly a polyequilibrium in the subgame $\Gamma^{X^{l-1}}$ (where $\Gamma = \Gamma^{X^0}$ is the original game). Repeated use of the second part of Fact 4 completes the proof. ■

A game has the *best-response existence property* if for every strategy profile x there is a strategy profile that is a best response to x . Clearly, all finite games and all mixed extensions of finite games have this property.

Fact 7. For a polyequilibrium X in a game with the best-response existence property, and for every strategy profile $x^1 \in X$, some $x^2 \in X$ is a best response to x^1 . Repeated use of this fact yields a *best-response sequence* x^1, x^2, x^3, \dots where each entry, except the first one, is an element of the polyequilibrium that is a best response to its immediate predecessor.

Proposition 7. In a game with the best-response existence property, a polystrategy profile X that is obtained by successive elimination of never-best-response strategies shares at least one strategy profile with each of the game’s polyequilibria. Therefore, in such a game, any strategy that is eliminated during successive elimination of never-best-response (or, as a special case, strictly dominated) strategies is not a polyequilibrium strategy.

Proof. Consider a finite sequence $S = X^0 \supseteq X^1 \supseteq \dots \supseteq X^L = X$ as in (the second part of) Definition 3. Let X' be any polyequilibrium, and x^1, x^2, x^3, \dots a best-response sequence of elements of X' as in Fact 7. Using induction, it is easy to see that $\{x^{l+1}, x^{l+2}, \dots\} \subseteq X^l$ for all $0 \leq l \leq L$. (If $l = 0$, the inclusion is trivial, and if $l > 0$, it is implied by the inclusion for $l - 1$, since the latter shows that x^{l+1}, x^{l+2}, \dots are best responses to strategy profiles in X^{l-1} .) In particular, $X' \cap X \neq \emptyset$. ■

Example 6. *The Traveler's Dilemma* (Basu 1994). In this finite symmetric two-player game, the strategy sets are $S_1 = S_2 = \{2, 3, \dots, 100\}$. For a player choosing strategy y , the payoff is $y, y + 2$ or $z - 2$ if the other player's choice z is equal to, greater than or less than y , respectively. Clearly, the unique best response to any strategy z is $y = \max\{z - 1, 2\}$. Therefore, successive elimination of never-best-response strategies eliminates $100, 99, \dots, 3$, for both players. It follows, by Proposition 7, that the game's unique equilibrium $(2, 2)$ is included in each of its (489, as it turns out) polyequilibria. Thus, there is no polyequilibrium where either player's payoff is greater than 2. The same is true also for the mixed extension of the game. This follows from the second part of Proposition 5 and the fact that $(2, 2)$ is also the unique mixed-strategy equilibrium.

The last example is somewhat special in that the polystrategy profile obtained by successive elimination of never-best-response strategies is a singleton. It is not difficult to see that, in a game with the best-response existence property, such a singleton must be an equilibrium. However, if the polystrategy profile X obtained is not a singleton, then it is not necessarily a polyequilibrium (but only intersects every polyequilibrium in the game). For example, in the finite game (3), elimination of the never-best-response strategy M does not give a polyequilibrium.

In Proposition 7, the assumption that the game has the best-response existence property cannot be dropped. In fact, as the next example shows, in a game without this property, successive elimination (unlike one-time elimination) of strictly dominated strategies *may* eliminate a polyequilibrium strategy. Note that the same is not true for equilibrium strategies, which never get eliminated this way. Thus, equilibria and polyequilibria differ in this respect.

Example 7. The strategy set of player 1 consists of all integers and that of player 2 is $\{0, 1\}$. If they choose y and z , respectively, player 1 receives y and player 2 receives 0 if $y + z$ is even and 1 if it is odd. Consider the following two ways of successively eliminating strictly dominated strategies: first, either all odd or all even numbers are eliminated for player 1, and then strategy 0 or 1, respectively, is eliminated for player 2. In both cases (odd or even numbers), a polyequilibrium is obtained. Thus, each of player 2's two strategies is a polyequilibrium strategy, even though successive elimination of strictly dominated strategies may eliminate it.

7 Dynamic Games

The defining property of polystrategy is that a player's course of action may be only partially specified. In a dynamic context, this may mean that the specification is restricted to only some of the player's information sets.

As for strategic games, a polystrategy X_i of a player i in a dynamic game G with either perfect or imperfect information is any nonempty set of strategies. The meaning of 'strategy' here is viewed as part of the game's specification. The term may refer to *pure* strategies, which prescribe a single action at each of the player's information sets, or to *behavior* strategies, which




Figure 1. The centipede game (with an odd number of decision nodes m).

prescribe a probability distribution over actions at each information set. A polystrategy X_i is said to *exclude* a particular action or distribution over actions at a particular information set if none of the strategies in X_i prescribes it (in other words, if every strategy that does prescribe it is excluded). A polystrategy is *rectangular* if it includes *all* the strategies that do not prescribe excluded actions or distributions over actions at any of the player's information sets. A profile of rectangular polystrategies corresponds to a polystrategy profile in the *agent normal form* of the game.

The simplest kind of dynamic game is an (either perfect- or imperfect-information) *extensive form game*, which is one that can be described by a finite game tree, possibly with chance nodes. As in the case of general dynamic games, it still needs to be specified whether all behavior strategies or only pure strategies may be used. Any statement where this is not specified or can be understood from the context is to be interpreted as referring to both cases.

Example 8. *The centipede game* (Rosenthal 1981). In this extensive form game with perfect information, there are $m \geq 2$ decision nodes, numbered from 1 to m (see Figure 1). The odd- and even-numbered nodes are controlled by player 1 and 2, respectively. At each node, the controlling player has to choose between Stop and Continue, except that at node m only Stop can be chosen. The payoffs are determined by the first node k at which Stop was chosen. The player controlling that node receives 2^k and the other player receives $2^k/3$.

Consider the version of the centipede game where only pure strategies may be used. Effectively, a strategy is described by the index $1 \leq k \leq m + 1$ of the first node at which the player chooses Stop, with $k = m + 1$ standing for the strategy of never stopping (which is relevant only for the player not controlling the last node m). Therefore, a polystrategy profile is any subset of $\{1, 2, \dots, m + 1\}$ (specifying the collection of "first Stop" nodes) that includes at least one odd number and at least one even number. A necessary condition for such a subset to be a polyequilibrium is that it is of the form $\{1, 2, \dots, l\}$, for some $2 \leq l \leq m + 1$. This is because, by Fact 7, a polyequilibrium that includes any strategy $2 \leq k \leq m + 1$ must also include the unique best response to it, which is strategy $k - 1$. The above condition is also sufficient for polyequilibrium. This is because, for any $2 \leq l \leq m + 1$, the strategy of first stopping at l or $l - 1$ (depending on the player's identity and the evenness or oddness of l) responds to $\{1, 2, \dots, l\}$ at least as well as any strategy that prescribes a later stopping time does. Thus, a




Figure 2. Two polyequilibria that do not include an equilibrium, in complete-information games that begin with a chance move. **a** Each player’s polystrategy comprises two strategies, one indicated by black lines and the other by gray ones. **b** Each player’s polystrategy comprises three strategies: one indicated by black lines, one by gray lines, and one by thick lines of either color.

polystrategy profile is a non-trivial polyequilibrium if and only if the two players’ polystrategies are to stop no later than node l , for some (fixed, common) $2 \leq l \leq m$. The game therefore has m nested polyequilibria. The largest polyequilibrium is the trivial one, and the smallest (which is also the only small polyequilibrium) is the game’s unique equilibrium (corresponding to $l = 2$). Thus, the game satisfies PE-equivalence.

As the next example shows, there are also perfect-information extensive form games that do not satisfy PE-equivalence. Moreover, there are such games where a player has a unique equilibrium payoff but receives a higher payoff in a polyequilibrium.

Example 9. Semi-dictator games. Players 1 and 2 have \$2 to share. They flip a coin, and the winner can either dictate an equal split of the money or ask for the whole sum. If he chooses the latter, however, the other player is allowed to object, and in this case, no one gets anything. Assuming that only pure strategies may be used, each player has four strategies. However, since the only decision that affects a player’s own payoff is the one he makes if he wins the coin toss, any polystrategy that does not exclude any of the two possible decisions there is part of a polyequilibrium. Such a polyequilibrium X is shown in Figure 2a. Each player’s polystrategy includes two strategies, Black and Gray, which prescribe choosing the actions indicated by black and gray lines, respectively, in both decision nodes. It is easy to see that none of the four strategy profiles in X is an equilibrium. The game in Figure 2b is a variant of the first one, and can be described as involving an additional, payoff-irrelevant public coin toss. The polystrategy X_i shown for each player i includes three strategies, Black, Gray and Thick, which prescribe choosing the action indicated by a line with that property at each decision node. In particular, at the two decision nodes of player i that immediately follow the chance node C , his polystrategy prescribes three pairs of actions; the only pair missing is choosing the actions indicated by thin lines in both nodes. However, the latter yields player i the same (expected) payoff as Thick, as long as the other player j only uses strategies belonging to his polystrategy X_j . This proves that $X = (X_1, X_2)$ is a polyequilibrium. None of the nine strategy profiles in X is an equilibrium.

Moreover, the *outcome* of each of them, that is, the distribution it induces over the terminal nodes, is not an equilibrium outcome. Indeed, it is not difficult to check that every strategy profile in X yields the players an aggregate payoff of either 1 or $3/2$, whereas in every (pure-strategy) equilibrium in the game the aggregate payoff is 2. If a third player, whose only role is to get the money if the others do not receive it, were added to the game, that player's payoff would be greater than 0 in the polyequilibrium X but 0 in every equilibrium. Parenthetically, it is an immediate corollary of the second part of Proposition 5 that the last statement would not be true if behavior strategies were allowed, as a corresponding equilibrium yielding a positive payoff to the third player does exist. However, the polyequilibrium X yields such a payoff without requiring the players to randomize.

7.1 Subgame Perfection

A strategy x_i of a player i in a dynamic game G induces a strategy for i in each subgame of G . That strategy, which may also be denoted by x_i if the meaning is clear from the context, is obtained by restricting the original strategy to the information sets included in the subgame.⁹ These observation and notation convention naturally extend to strategy profiles, polystrategies and polystrategy profiles.

Definition 4. A polystrategy profile X in a dynamic game¹⁰ G is a *weak subgame perfect polyequilibrium* if, in every subgame of G , the induced polystrategy profile is a polyequilibrium. This condition may be expressed more explicitly as follows: for every strategy profile x' and every subgame G' there is some $x'' \in X$ that, in the subgame G' , responds to X at least as well as x' does. A polystrategy profile X is a *subgame perfect polyequilibrium* if it satisfies the following stronger condition: for every strategy profile x' there is some $x'' \in X$ that in *all* subgames of G responds to X at least as well as x' does.

The difference between subgame perfection and weak subgame perfection is illustrated by Figure 3. (Another example is Figure 2, where the polyequilibrium in a is subgame perfect whereas that in b is only a weak subgame perfect polyequilibrium.) Both properties are “hereditary” in the sense that a polyequilibrium with either property induces a polyequilibrium with that property in every subgame. For a polystrategy profile that is a singleton, $X = \{x\}$, the two properties are equivalent and hold if and only if x is a subgame perfect equilibrium. In general, however, a subgame perfect polyequilibrium may not include a subgame perfect or even *any* equilibrium. (The polyequilibrium in Figure 2a is an example of this.) Two exceptions to this general rule are given by the next theorem.

⁹ A subgame by definition includes either none or all of the nodes in each information set. Note that the meaning of ‘subgame’ in the present, dynamic context (Selten 1975) is different from that in the strategic context (Shapley 1964; see Section 2).

¹⁰ The definition is applicable to both perfect- and imperfect-information games. However, for the latter, subgame perfection is a relatively weak requirement, which does not fully capture the notion of sequential rationality. A refinement of subgame perfect polyequilibrium that accommodates that notion is introduced in Milchtaich (2018).




Figure 3. Subgame perfect polyequilibrium (SPP) and weak SPP. **a** In this one-player game, the singleton $\{RR'\}$ is a SPP and its complement $\{RL', LR', LL'\}$ is a weak SPP. **b** In this two-player game, where only player 1's payoffs are shown, $X = \{LR', LL'\} \times \{r, l\}$ is a weak SPP. It is not a SPP because neither of player 1's strategies responds to X at least as well as RL' does *both* in the whole game (where LR' does so) and in the subgame starting at the player's second decision node (where LL' does so).

Theorem 3. In a perfect-information extensive form game G :

- (1) Every weak subgame perfect polyequilibrium X where the players' polystrategies are rectangular includes a subgame perfect equilibrium.
- (2) If G has a unique subgame perfect equilibrium x' , then every subgame perfect polyequilibrium X includes x' .

Proof. (1) Consider the collection \mathcal{X} of all subsets of X that are weak subgame perfect polyequilibria and consist of rectangular polystrategies. This collection is not empty, as $X \in \mathcal{X}$. For each element of \mathcal{X} , count the number of decision nodes at which at least two actions or distributions over actions are *not* excluded, and consider some $X' \in \mathcal{X}$ for which this number is minimal. If the number is zero, then X' is a singleton, which implies that it is a subgame perfect equilibrium, and the proof is complete. It therefore suffices to assume that the number is greater than zero, and show that this assumption leads to a contradiction. The assumption implies that, for X' , there is some decision node v such that, (i) at each of the decision nodes following v , only one action or distribution over actions is not excluded, but (ii) this is not so for v itself. Let x' be a strategy profile such that, (i) the actions or distributions over actions specified by x' at the nodes following v are those singled out by X' , and (ii) the one specified at v best responds to them in the subgame G' starting at that node. Since X' is a weak subgame perfect polyequilibrium, it includes a strategy profile x'' which, in G' , responds to X' at least as well as x' does. Let X'' be the polystrategy profile obtained from X' by removing all strategy profiles that do not agree with x'' (in the sense of specifying the same action or distribution over actions) at the node v . It is not difficult to see that X'' is also an element of \mathcal{X} . However, this conclusion contradicts the minimality assumption concerning X' .

(2) Consider a strategy profile $x'' \in X$ that in all subgames of G responds to X at least as well as x' does. If $x'' \neq x'$, then there is some subgame G' where the strategy profiles induced by x' and by x'' differ only at the root, where they prescribe different actions or distribution over

actions to the acting player i . Since in G' strategy x_i'' responds to x'' (hence, to x') at least as well as x_i' does, and x' is a subgame perfect equilibrium, both x' and x'' induce subgame perfect equilibria in G' . This conclusion clearly contradicts the assumption that (the whole game) G has a unique subgame perfect equilibrium. The contradiction proves that $x'' = x'$. ■

In the above analysis of the centipede game, ‘strategy’ actually refers to a number of *equivalent* (pure) strategies, which specify the same first Stop node but differ in their prescription of actions at the player’s later decision nodes. This is so in general: polyequilibrium analysis never requires distinguishing between equivalent strategies, because the payoffs they yield are always identical and therefore all but (any) one of them may be excluded. However, since equivalent strategies may differ in the strategies they induce in *subgames*, the distinction between them may be important in the context of subgame perfection.

Example 8 (continued). As shown, a polystrategy profile in the centipede game is a polyequilibrium if and only if it has the form $\{1, 2, \dots, l\}$, for some $2 \leq l \leq m + 1$. For each such polyequilibrium, and for each $1 \leq k \leq l$, which represents all equivalent strategies whose first Stop is at node k , consider the representative strategy that prescribes Stop also at each of the player’s later decision nodes. Such a choice of representative strategies makes the polyequilibrium a weak subgame perfect polyequilibrium, since it is not difficult to check that the polystrategy profile it induces in every subgame is again of the general form indicated above. This weak subgame perfect polyequilibrium X is actually subgame perfect if $2 \leq l \leq 5$. However, for larger l this is not so. To see this, suppose that $l \geq 6$ and consider the strategy x_1' of player 1 that instructs him to continue only at his second decision node (no. 3). No strategy $x_1'' \in X_1$ responds to X at least as well as x_1' does in *all* subgames. This is because, to do so in the two subgames starting at player 1’s first and second decision nodes, x_1'' must specify the same actions there as x_1' . However, by construction, no strategy in X_1 does so. This proves that the weak subgame perfect polyequilibrium X is not subgame perfect if $l \geq 6$. Note that X consists of rectangular polystrategies if and only if $2 \leq l \leq 4$. Nevertheless, for *all* $2 \leq l \leq m + 1$, X includes the game’s unique subgame perfect equilibrium, which corresponds to $l = 2$ (cf. Theorem 3).¹¹

Dynamic games that are not extensive form ones do not always have a subgame perfect equilibrium. For such games, the concept of subgame perfect polyequilibrium may be particularly pertinent.

Example 10. “Almost perfect” information and a continuous action space (Harris et al. 1995, Myerson and Reny 2015). First, players 1 and 2 choose their actions simultaneously. Then, players 3 and 4, who are informed of their predecessors’ choices, do so. The set of actions for player 1 is the interval $[-1, 1]$ and that of each of the other players is the pair $\{-1, 1\}$. Denoting the action of player i by a_i , the four players’ payoffs are given by

¹¹ For an analysis of curb sets in the centipede game, see Pruzhansky (2003). If m is odd, the only curb set is the trivial polyequilibrium.

$$\begin{aligned}
u_1 &= 5(a_3 a_4 - 1) - |a_1| a_2 a_3 - a_1^2 & u_2 &= \frac{1}{2}(1 + 3a_2)a_3 \\
u_3 &= a_1 a_3 & u_4 &= a_1 a_4
\end{aligned}$$

It can be shown that, even with all behavior strategies allowed, this game does not have a subgame perfect equilibrium. Roughly, the reason is that, to make the first, second and third term in player 1's payoff function as large as possible, a_1 should be, respectively, (i) different from 0, so that the responses of players 3 and 4 will match, (ii) positive and negative with close-to-equal probabilities, so that player 2 will not be able to match player 3's action, and (iii) close to 0. No behavior strategy of player 1 (equivalently, mixed strategy, probability measure on $[-1, 1]$) optimally satisfies these three requirements. However, for any $0 < \epsilon \leq 1$, consider the polystrategy $X_1 = \{\pm a \mid 0 < a \leq \epsilon\}$, where $\pm a$ means playing a with probability 0.5 and $-a$ with probability 0.5. Thus, the smaller is ϵ , the tighter is the description of player 1's behavior. For the other players, the behavior is completely specified, as follows: X_2 includes only the single strategy ± 1 , and X_3 and X_4 include only the strategy that specifies choosing 1, -1 or ± 1 if player 1's action a_1 is positive, negative or zero, respectively. Clearly, the strategy of each of these three players is a best response to the polystrategy profile $X = X_1 \times X_2 \times X_3 \times X_4$, and this is so also in every subgame (that is, after players 1 and 2 play). To show that X is a subgame perfect polyequilibrium, it remains to show that for every mixed strategy x_1 of player 1 there is some $0 < a \leq \epsilon$ such that $\pm a$ responds to X at least as well as x_1 does. If player 1 uses x_1 , his payoff is $-5\mathbb{P}(a_1 = 0) - \mathbb{E}a_1^2$, where the expectation and probability are those specified by x_1 . This payoff is lower than the $-a^2$ player 1 would get from playing $\pm a$, for every a smaller than $\mathbb{E}|a_1|$ (or, if $\mathbb{E}|a_1| = 0$, for any $0 < a \leq \epsilon$).

Another advantage of subgame perfect polyequilibrium over subgame perfect equilibrium manifests itself in dynamic games with many information sets. Whereas a subgame perfect equilibrium must prescribe a carefully selected action at each information set, including those lying far away from the equilibrium path, a polyequilibrium may legitimately ignore all but a relatively small number of relevant information sets.

Example 11. Sequential competition (Milchtaich et al. 2018). A market for a particular good has a continuum of consumers on one side and two sellers on the other side. The consumers arrive in a steady flow: in any time interval of length l , the total mass of arriving consumers is l . Each consumer demands a single unit of the good, and leaves the market after buying it or spending a unit of time in the market, whichever comes first. In the first case, the consumer's payoff is $1 - x - p$, where x is the time he spent in the market and p is the price he paid for the good, and in the second case, the payoff is zero. Thus, a consumer's valuation of the good decreases linearly with his "age" x . The two sellers produce the good at zero cost. Seller 1 arrives at time 1 and seller 2 arrives shortly afterwards, at 1.1. An arriving seller announces a price p for the good, sells the demanded quantity q , and immediately leaves. His payoff is the revenue pq .

For any price $0 < p_1 < 1$ that seller 1 sets, the total mass of the consumers who would receive positive payoff from buying at that price is $1 - p_1$. However, the seller's actual profit from setting price p_1 may be lower than his *monopoly profit* of $p_1(1 - p_1)$, because if consumers

expect seller 2's price p_2 to be significantly lower than p_1 , some of them may choose to wait. The wait may in turn affect p_2 , since it changes the demand seller 2 faces. Thus, a strategy for seller 2 has to prescribe a price p_2 for every p_1 and every possible total mass and age distribution for the consumers who are in the market at the time seller 2 arrives there. One may suspect that these myriad decision nodes are not all equally relevant. This is where the notions of polystrategy and (subgame perfect) polyequilibrium, which legitimize the consideration of only some decision nodes, come in handy.

As shown below, for every price $0.4 \leq p^* \leq 0.5$ there is a subgame perfect polyequilibrium with $p_1 = p^*$ and $p_2 = 0.9$. Since the second price is higher, none of seller 1's potential customers opts for waiting. Nevertheless, with $p^* < 0.5$, seller 1 does not take advantage of this by choosing his *monopoly price* of 0.5, which maximizes the monopoly profit. The reason is the credible threat implicit in the consumers' strategies, which are the following ones. Consumers always buy at any price p that leaves them with nonnegative payoff, except when it is seller 1's price and it is higher than p^* , in which case they buy only if their payoff is at least $p - 0.2$. Thus, if seller 1 sets a price $p_1 \leq p^*$, all the consumers who arrived in the time interval $[p_1, 1]$ buy from him, but if $p_1 > p^*$, then only those who arrived in $[2(p_1 - 0.1), 1]$ do so (or, if $p_1 > 0.6$, no one buys). In the first case, seller 1's profit is $p_1(1 - p_1)$, and it is therefore maximal at $p_1 = p^*$, where the profit is between 0.24 and 0.25. In the second case, the profit is only $p_1(1 - 2(p_1 - 0.1))$ (or 0), which is less than 0.16. Thus, $p_1 = p^*$ is the profit-maximizing price for seller 1. Seller 2's polystrategy specifies a profit-maximizing price p_2 only at decision nodes where the consumers' total mass and age distribution correspond to the price p_1 set by seller 1 (that is, they are as described above). For $p_1 \leq p^* (\leq 0.5)$, the price is $p_2 = 0.9$, and for $p_1 > p^*$, it is $p_2 = \min\{0.5, p_1 - 0.1\}$. In the second case, the difference between p_2 and p_1 compensates for the loss of payoff due to the waiting time of 0.1 and guarantee that both the consumers who bought from seller 1 (if there are any such consumers) and those who chose to wait acted optimally. This optimality is the sense in which the threat implicit in the consumers' strategies is credible, and it is what makes these strategies, together with seller 1's strategy of selling at price p^* and sellers 2's rectangular polystrategy just described, a subgame perfect polyequilibrium.

Note that the polyequilibrium in Example 11 has a well-defined *path* as it prescribes a unique action at every information set that is actually reached. In this, it differs from every non-singleton polyequilibrium in the centipede game, in which different strategy profiles may specify different paths in the game tree.

8 Discussion

The notion of polyequilibrium generalizes – in a rather simple, straightforward manner – the Nash equilibrium solution concept. Like the latter, it is universally applicable, in that it does not rely on any assumptions about the nature of the game (simultaneous, dynamic, complete or incomplete information, etc.), the cardinality of the player set (finite, countably infinite or a continuum; see Example 11), the structure of the strategy sets (linear structure, topological

properties), the functional form of the payoff functions (multilinear, continuous) or the interpretation of the payoffs (ordinal, von Neumann-Morgenstern utilities). Unlike *mixed-strategy* Nash equilibrium, polyequilibrium does not involve randomization. It is perfectly compatible with mixed strategies, but it does not introduce them. Thus, a finite game and its mixed extension are viewed as distinct games, with distinct – and often quite dissimilar (see Proposition 3) – sets of polyequilibria. The same applies to dynamic games with pure strategies and (the same) games with behavior strategies.

Polyequilibrium is similar to other set-valued solution concepts, like rationalizable strategies and curb sets, in that it identifies polystrategy profiles (that is, rectangular sets of strategy profiles) that can be described as self-enforcing. However, it differs from them quite fundamentally in that beliefs play no role. The first implication of this difference is that polyequilibrium does not require cardinal utilities and does not involve probabilities, so its applicability is not limited to only certain families of games. Second, the definition of polyequilibrium, which is based on strategy substitution, is sufficiently different to make the concept logically incomparable with rationalizable strategies and curb sets even in simple finite games (see Section 4). Third, and most profoundly, polyequilibrium is a strictly excluding solution concept. This means that the inclusion of a strategy in a player's polystrategy does not suggest that anyone believes or should believe that the use of that strategy is likely or even possible. Self-enforcement only concerns the *exclusion* of strategies from the players' polystrategies. Thus, the latter are not themselves interpreted as strategy recommendations or as the collections of all strategies that other players view as possible. By contrast, the concepts of rationalizable strategies and curb sets *are* based on an interpretation along these lines, from which they derive the requirement that the polystrategy profile must be a subset or a superset, respectively, of the set of all best responses to the beliefs that are supported in it (van Damme 2002, p. 1527). For polyequilibrium, neither inclusion is a defining property, but in games with the best-response existence property, every strategy profile in a polyequilibrium X has some best response that is also in X (Fact 7). Not all of them must be included, though, so this solution concept allows for selection among best responses.

As an illustration of the difference between an emphasis on inclusion or exclusion of strategies, consider the chain of nested polyequilibria that are obtained by successively eliminating weakly or strictly dominated strategies (Section 6). The temporary inclusion of a strategy does not constitute a contradiction to its later exclusion, because the inclusion is not a suggestion that the use of the strategy is possible. The first elimination may only establish that, say, "player i does *not* play some strategy T " is a polyequilibrium result, and a later stage may establish the same for "player i does not play T or M ". These results are perfectly harmonious. Indeed, any logical consequence of a polyequilibrium result is also so. Although it might seem natural to refer only to the last polystrategy profile in the chain as a polyequilibrium, or more radically, to reserve the term to *minimal* polyequilibria, this mathematical reflex is to be resisted. A larger polyequilibrium yields weaker results than a smaller polyequilibrium (which is included in it) does, because fewer properties are common to all its strategy profiles. However, if the larger

polyequilibrium already yields the result of interest, there is little point in chasing the absolute minimum or pondering whether it exists at all. Thus, multiplicity of equilibria is a feature, not a bug. This philosophy sets polyequilibrium apart from other set-valued solution concepts, which put a premium on either minimality or maximality.

The last point is particularly true for dynamic games (Section 7), where an insistence on minimality would actually be self-defeating. Part of the appeal of the polyequilibrium concept in such games, particularly those with many information sets, is the legitimization it provides to specifying the players' actions at only some of these sets, thus doing away with the need for pruning actions at information sets that are, in a logically consistent sense, unimportant. In particular, if all strategy profiles in a subgame perfect polyequilibrium give the same outcome (that is, the same distribution over terminal nodes), laboring to decrease its size would be counterproductive.

As the results in this paper demonstrate, the kinds of questions, results, directions of inquiry and conceptual issues raised by the polyequilibrium concept are quite different from those associated with other – more similar or less so – solution concepts. Moreover, the potential for extension is vast. For example, the concept may be applied to additional, important classes of games, such as games with incomplete information. It can be extended to encompass (suitably defined) correlated polyequilibria. And it may suggest a fresh look at other, established solution concepts, such as sequential equilibrium (Milchtaich 2018).

Acknowledgements

The author thanks a reviewer and an advisory editor for helpful comments and suggestions.

References

- Basu, K., 1994. The traveler's dilemma: paradoxes of rationality in game theory. *Am. Econ. Rev.* 84, 391–395.
- Basu, K., Weibull, J.W., 1991. Strategy subsets closed under rational behavior. *Econ. Lett.* 36, 141–146.
- Bernheim, D., 1984. Rationalizable strategic behavior. *Econometrica* 52, 1007–1028.
- Duggan, J., Le Breton, M., 1996. Dutta's minimal covering set and Shapley's saddles. *J. Econ. Theory* 70, 257–265.
- Harris, C., Reny, P., Robson, A., 1995. The existence of subgame-perfect equilibrium in continuous games with almost perfect information: a case for public randomization. *Econometrica* 63, 507–544.
- Kalai, E., Samet, D., 1984. Persistent equilibria in strategic games. *Int. J. Game Theory* 13, 129–144.

- Kohlberg, E., Mertens, J.-F., 1986. On the strategic stability of equilibria. *Econometrica* 54, 1003–1037.
- Milchtaich, I., 2018. [Agnostic sequential rationality](#). Mimeo.
- Milchtaich, I., Glazer, A., Hassin, R., 2018. [Time-differentiated monopolies](#). Mimeo.
- Moulin, H., 1979. Dominance solvable voting schemes. *Econometrica* 47, 1337–1351.
- Myerson, R.B., Reny, P.J., 2015. Sequential equilibria of multi-stage games with infinite sets of types and actions. Mimeo.
- Pearce, D.G., 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029–1050.
- Pruzhansky, V., 2003. On finding curb sets in extensive games. *Int. J. Game Theory* 32, 205–210.
- Rosenthal, R., 1981. Games of perfect information, predatory pricing and the chain store paradox. *J. Econ. Theory* 25, 92–100.
- Selten, R., 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory* 4, 25–55.
- Shapley, L.S., 1964. Some topics in two-person games. In: Dresher, M., Shapley, L.S., Tucker, A.W. (Eds.), *Advances in Game Theory* (Annals of Mathematics Studies 52). Princeton University Press, Princeton, pp. 1–28.
- van Damme, E., 2002. Strategic equilibrium. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory*, vol. 3. North-Holland, Amsterdam, pp. 1521–1596.