

Milchtaich, Igal

Working Paper
Poleyquilibrium

Working Paper, No. 2015-06

Provided in Cooperation with:
Department of Economics, Bar-Ilan University

Suggested Citation: Milchtaich, Igal (2015) : Poleyquilibrium, Working Paper, No. 2015-06, Bar-Ilan University, Department of Economics, Ramat-Gan

This Version is available at:
<https://hdl.handle.net/10419/130536>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Polyequilibrium

Igal Milchtaich, *Bar-Ilan University*

June 2015

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 to every profile of non-excluded strategies of the other players at least as well as the first strategy does.

A polyequilibrium's description of the outcome of the game may be more or less specific, depending on the number and the identities of the non-excluded strategy profiles.

A particular property of the outcome is said to hold in a polyequilibrium if it holds for all non-excluded profiles. Such a property does not necessarily hold in any Nash equilibrium in the game. In this sense, the generalization proposed in this work extends the set of justifiable predictions concerning a game's results.

Keywords. Polyequilibrium, Polystrategy, Coarsening of Nash equilibrium, Subgame perfection, Bayesian 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 him 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 i . 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 any excluded strategy would not make him 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 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. Specifically, in a curb set, each excluded strategy must *not* to be a best response

to any belief about the other players' strategies that assigns positive probability only to strategy profiles without 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., probability distributions over strategy profiles. 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 may 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 any 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 .

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 4.1.

The condition defining polyequilibrium ostensibly allows any strategy of any player – 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 for a particular polyequilibrium. That is, a property is said to hold only if all strategy profiles in the polyequilibrium have it. In particular, a player's strategy is said to be a polyequilibrium strategy if there is a polyequilibrium where the player's other strategies are all excluded. This means that, for example, a strategy in a finite game that is played with positive probability in some mixed-strategy equilibrium – and is thus automatically rationalizable – is not necessarily a polyequilibrium strategy (because the support of the player's mixed strategy also includes other pure strategies). However, a polyequilibrium may be able to specify certain aspects of the game's outcome without singling out a unique strategy profile. It may specify, for example, that a particular player does *not* take a particular action, that he receives a positive payoff, that the total payoff is

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.

less than 2, 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.

As an illustration of the difference between (Nash) equilibrium and polyequilibrium, consider the following two statements about a particular 2×2 game: " T is an equilibrium strategy for player 1" and " T is a polyequilibrium strategy for player 1". Both statements do not mention player 2, which suggests that player 1 is the main subject of interest. The first statement means that it is possible to assign player 2 a strategy that is a best response to 1's T and vice versa. It just does not spell out whether that strategy is L or R (or whether there are two equilibria, one with L and the other with R). The second statement, by contrast, does not necessarily refer to a specific strategy of player 2 even implicitly. The polyequilibrium it refers to may specify that 2 plays L or that he plays R (equivalently, does not play L), but it may also leave 2's strategy unspecified. There is of course nothing unfamiliar about justifying a player's choice of strategy without considering the other players' play. Choosing T is justifiable if it is a dominant strategy. Thus, the notion of dominant strategy essentially also comes under the umbrella of polyequilibrium. A strictly *dominated* strategy, by contract, is never a polyequilibrium strategy.

2 Definitions and Basic Facts

A (strategic) game Γ is specified by a (possibly infinite) 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 (1) holds for all $x'' \in X$. If, in addition, at least some of the inequalities are strict or all of them are so, then x'_i (*weakly*) *dominates* or *strictly dominates* x_i , respectively, relative to X . Strategy x'_i 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 of player i does. A *never-best-response* strategy relative to a set of strategy profiles X is a strategy that is not a best response to any strategy profile in X . Such a strategy is not necessarily dominated relative to X or vice versa. However, a strategy that is *strictly*

dominated (by some other strategy) relative to X is clearly also a never-best-response strategy relative to it. In each of the above expressions involving the phrase “relative to (a set of strategy profiles) X ” the latter may be dropped, and in this case, it is understood that the expression refers to the entire collection of strategy profiles, that is, $X = S$.

A strategy *profile* x' (i) responds to a strategy profile x'' at least as well as strategy profile x does if (1) holds for every player i , (ii) responds to a set of strategy profiles X at least as well as x does if the previous condition holds for all $x'' \in X$, and (iii) 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. A game has the *best-response existence property* if there is a best response to every strategy profile. Clearly, all finite games and all mixed extensions of finite games have this property.

For a player i in a game Γ , a *polystrategy* is any nonempty set of strategies, $\emptyset \neq X_i \subseteq S_i$. Player i 's entire strategy set S_i is referred to as his *trivial polystrategy*. A polystrategy that is a singleton, $\{x_i\}$, may be identified with the strategy x_i . A *polystrategy profile* X is an assignment of a particular polystrategy X_i to each player i or, equivalently, a nonempty rectangular subset of S . If the subset is a singleton, $\{x\}$, 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 strategy profiles X with $X' \subseteq X \subseteq X''$.

Definition 1. A polystrategy profile X is a *polyequilibrium* if for every strategy profile $x \notin X$ there is some $x' \in X$ that responds to X at least as well as x does. X is a *strict polyequilibrium* if for every player i and strategy $x_i \notin X_i$ there is some $x'_i \in X_i$ that strictly dominates x_i relative to X . A polystrategy profile X is a *simple polyequilibrium* if some strategy profile $x \in X$ is a best response to X . A *minimal polyequilibrium* is one that does not contain another polyequilibrium.

The five facts below easily follow from the definitions.

Fact 1. Every game has at least one strict polyequilibrium, namely, the *trivial polyequilibrium* S , which consists of all strategy profiles.

Fact 2. A polystrategy profile that is a singleton, $\{x\}$, is a polyequilibrium or a strict polyequilibrium if and only if its single element x is a (Nash) equilibrium or a strict equilibrium, respectively.

Fact 3. A sufficient condition for a polystrategy profile X to be a polyequilibrium is that all 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 Γ^X is independent of his own strategy, that is, $u_i(x) = u_i(x | x'_i)$ for all i , $x \in X$ and $x'_i \in X_i$.²

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 dominated strategies.

Fact 4. For strategy profiles X' and X'' with $X' \subseteq X''$, $[X', X'']$ is a *polyequilibria interval* (that is, all its elements are polyequilibria) 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.

Fact 5. For a polyequilibrium X in a game with the best-response existence property, and for any strategy profile $x^1 \in X$, some $x^2 \in X$ is a best response to x^1 . Successive use of this fact yields a *best-response sequence* x^1, x^2, x^3, \dots where each entry is an element of X that is a best response to its immediate predecessor.

Fact 6. 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.

Note that the above definitions and facts do not involve randomization or beliefs, and in particular, ‘strategy’ and ‘equilibrium’ always mean pure strategy and pure-strategy equilibrium.³ Moreover, the concepts introduced here are all *ordinal* in the sense that they are invariant to arbitrary increasing transformations of the players’ payoff functions and, in fact, could be alternatively formulated in terms of preferences over strategy profiles rather than payoffs.

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 $\phi_i: S_i \rightarrow S_i$.⁴ Any family 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 x and i ,

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

The inequality means that it is optimal for player i to use the recommended substitute $\phi_i(x_i)$ to 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 7. A polystrategy profile X is a polyequilibrium if and only if it is the image of some rational substitution function ϕ (that is, $\phi(S) = X$).

³ This aspect of the setup reflects the view that the class of allowable strategies is part of the game’s specification. Allowing the use of mixed strategies in a game Γ would turn it into another game, namely, the *mixed extension* Γ^* . Since the polyequilibrium concept does not *introduce* randomization, a polyequilibrium in Γ is not the same thing as a polyequilibrium in Γ^* . See Section 4.

⁴ 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.

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 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 be specified either explicitly or implicitly, as a particular property or consequence of strategy profiles (for example, "player 1's payoff is higher than 2's payoff"). In the second case, R is the collection of all strategy profiles that have 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, for a player i , a strategy x_i is a *polyequilibrium strategy* if there is some polyequilibrium X with $X_i = \{x_i\}$, and a real number v_i is a *polyequilibrium payoff* if there is some polyequilibrium X with $u_i(x) = v_i$ for all $x \in X$. A generalization of polyequilibrium payoff 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 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 be used also for special kinds of polyequilibria, and in particular for equilibria. Every equilibrium result is also a polyequilibrium result but not conversely. There are also some logical differences between the two concepts. For example, 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.

Example 1. In the finite (that is, pure-strategy) version of matching pennies, the only polyequilibrium is the trivial one. Therefore, there is no polyequilibrium where player 1 plays Heads and no polyequilibrium where he does not play Heads (equivalently, plays Tails).

A game will be said to be *plain* if every polyequilibrium in it includes at least one equilibrium, equivalently, if all polyequilibrium results are also equilibrium results. In a plain game, a polyequilibrium where all strategy profiles have a particular property exists if and only if

some equilibrium has the property. In a non-plain game, some results that do not hold in any equilibrium do hold in a polyequilibrium. The next three examples present such games.

Example 2. For both players in the finite game

$$\begin{array}{c} T \\ M \\ B \end{array} \begin{array}{ccc} L & C & R \\ \left(\begin{array}{ccc} 1,1 & 0,0 & 0,0 \\ 0,0 & 2,3 & 3,2 \\ 0,0 & 3,2 & 2,3 \end{array} \right) \end{array}$$

(Basu and Weibull 1991) the only equilibrium payoff is 1, but $[2,3]$ is also a polyequilibrium payoff interval. The former corresponds to the game's unique (pure-strategy) equilibrium $\{T, L\}$ and the latter to the polyequilibrium $\{M, B\} \times \{C, R\}$. If a third player were added to the game, whose payoff is the average payoff of the original two players and who does not take any meaningful action, then only 1 would be an equilibrium payoff for that player but 2.5 would be a polyequilibrium payoff. The latter is obviously also a *mixed*-equilibrium payoff for all players (see Proposition 3 in Section 4).

Example 3. Bilateral trade. A buyer has to offer a price p to the owner of an item whose worth of 1 to the buyer and 0 to the owner, and the latter has to decide whether to sell at that price. The strategy "accept any price greater than zero" is a dominant strategy for the owner, 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's strategy should be to offer "as little as possible", or "an ϵ ", is incompatible with the definition of equilibrium. However, this idea is compatible with polyequilibrium. For any $0 < \epsilon \leq 1$, the owner's strategy of accepting any positive price and the buyer's polystrategy $0 < p \leq \epsilon$ ("offer a positive price not higher than ϵ ") together constitute a polyequilibrium.

Example 4. In a symmetric two-player game, each player must choose a positive integer y and his payoff is

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

where z is the integer chosen by the other player. To receive his highest payoff of z , a player must choose $y = 2z$. However, such a choice prevents the other player from receiving his highest payoff of y (which would require $z = 2y$), and therefore an equilibrium does not exist. In fact, as the total 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 positive integer L 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.

$\begin{array}{c} T \\ M \\ B \end{array} \begin{array}{ccc} L & C & R \\ \left(\begin{array}{ccc} -1,2 & 0,0 & 2,-1 \\ 0,0 & 0,0 & 0,0 \\ 2,-1 & 0,0 & -1,2 \end{array} \right) \end{array}$	$\begin{array}{c} T \\ M \\ B \end{array} \begin{array}{cc} L & R \\ \left(\begin{array}{cc} 3,-3 & 0,0 \\ 1,-1 & 1,-1 \\ 0,0 & 3,-3 \end{array} \right) \end{array}$	$\begin{array}{c} T \\ M \\ B \end{array} \begin{array}{cc} L & R \\ \left(\begin{array}{cc} -1,-3 & 1,-4 \\ 1,3 & -1,4 \\ 0,3 & 0,4 \end{array} \right) \end{array}$
a	b	c

Figure 1. Three finite games.

4 Finite Games and Their Mixed Extensions

The game in Example 2 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 the exception rather than the rule. In particular, that game is special in that its polyequilibria are all strict, and as can easily be seen, a strict polyequilibrium in a finite game is always a curb set. This observation aligns with the fact that the concept of curb set is meant as a generalization of strict, rather than Nash, equilibrium (Basu and Weibull 1991).

Example 5. The finite game in Figure 1a has two mixed-strategy equilibria, two polyequilibria and two curb sets. None of the pairs coincides with another. The first, pure-strategy, equilibrium (M, C) is a polyequilibrium but it is not a curb set. The support $\{T, B\} \times \{L, R\}$ of the second mixed-strategy equilibrium is not a polyequilibrium but it is a curb set. The trivial polyequilibrium is (trivially) also a curb set.

Example 5 shows that, in general, a curb set in a finite game may not even contain a polyequilibrium and vice versa, and the support of a mixed-strategy equilibrium may not contain a polyequilibrium. However, the opposite of the last assertion is false: a polyequilibrium always contains the support of some mixed-strategy equilibrium (although it does not necessarily coincide with one).

Proposition 1. In a finite game, every polyequilibrium includes (as a subset) the support of some mixed-strategy equilibrium.

This result is an immediate corollary of Proposition 2 below, which identifies a connection between polyequilibria in a finite n -player game Γ and in its mixed extension Γ^* . By definition, an unqualified 'strategy' in Γ or in Γ^* is a pure or mixed strategy in Γ , respectively. As the collection S_i of all pure strategies for a player i in Γ may be viewed as a subset of the collection S_i^* of all mixed strategies, Γ may be viewed as a subgame of Γ^* , with the same symbol u_i denoting the payoff function of a player i in both games.

A polystrategy X_i for a player i in Γ^* is a nonempty subset of S_i^* , that is, a collection of mixed strategies in Γ , and the same is true for its convex hull $\text{conv } X_i$. In the special case of a *pure* polystrategy, $X_i \subseteq S_i$, $\text{conv } X_i$ consists of all mixed strategies of player i whose supports are subsets of X_i . Note that, for a polystrategy profile $X = X_1 \times X_2 \times \dots \times X_n$ in Γ^* or (as a special case) in Γ ,

$$\text{conv } X = \text{conv } X_1 \times \text{conv } X_2 \times \dots \times \text{conv } X_n$$

is a polystrategy profile in Γ^* .

Proposition 2. 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^* with convex polystrategies) includes at least one equilibrium (which is a mixed-strategy equilibrium in Γ).

Proof. Consider any polyequilibrium X in Γ or in 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 some polytope (that is, the convex hull of a finite number of strategies in Γ^*) that includes each of these (finitely many) strategies x'_i . A (mixed) strategy in S_i^* is a convex combination of elements in S_i . It follows, by the linearity of u_i in each player's 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 (again 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 $\text{conv } X$, this proves that both polystrategy profiles are polyequilibria in Γ^* (see Fact 4). 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 6, every such equilibrium is also an equilibrium in Γ^* . ■

By the first part of Fact 6 and Proposition 2, the following implications hold for every polystrategy profile X in a finite game Γ (equivalently, every pure-polystrategy profile in the mixed extension Γ^*):

X is a polyequilibrium in $\Gamma^* \Rightarrow X$ is a polyequilibrium in $\Gamma \Rightarrow \text{conv } X$ is a polyequilibrium in Γ^* .

As the next example shows, neither of the implications is an equivalence: the converse of the first implication may be false, and the same holds for the second implication. In either of the games in the following example, a counterexample for the former is obtained by taking $X = S$, the trivial polyequilibrium in Γ , and for the latter, by choosing $X = \{T, B\} \times \{L, R\}$.

Example 6. The mixed extension Γ^* of the game Γ in Figure 1b has two disjoint polyequilibria intervals, which together include all its polyequilibria. The first interval is $[\{x\}, X']$, where $x = (1/2 T + 1/2 B, 1/2 L + 1/2 R)$ is the unique equilibrium in Γ^* and

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

and the second interval is $[X'', S^*]$, where

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

and S^* is the trivial polyequilibrium in Γ^* . The polyequilibria in the first interval are all simple and those in the second one are strict. None of the polyequilibria is a pure-polystrategy polyequilibrium. The same is not true for the mixed extension of the game in Figure 1a, which has the equilibrium (M, C) . Its other polyequilibria are the elements of the disjoint polyequilibria intervals

$$[\{x\}, \{(pT + (1 - 2p)M + pB, qL + (1 - 2q)C + qR) \mid 0 \leq p, q \leq 1/2\}]$$

and

$$[\{(pT + (1 - p)B, qL + (1 - q)R) \mid 0 \leq p, q \leq 1\}, S^*],$$

where x is the equilibrium $(1/2 T + 1/2 B, 1/2 L + 1/2 R)$ and S^* is the trivial polyequilibrium.

The above examples and results underline the importance of clearly distinguishing between a finite game Γ and its mixed extension Γ^* . In particular, they show that the identity between the (pure-strategy) equilibria in Γ and those equilibria in Γ^* that only involve pure strategies does not extend to polyequilibria.⁵ The set of polyequilibria in Γ does not generally coincide with but is rather a superset of the (possibly, empty) set of pure-polystrategy polyequilibria in Γ^* . At the same time, there is a simple one-to-one correspondence between the former and a different set of polyequilibria in Γ^* : a polyequilibrium X in Γ is matched with its convex hull $\text{conv } X$.

The set of all *mixed*-strategy equilibria in Γ (which includes the pure-strategy ones) coincides with the set of all equilibria in Γ^* . As the next example demonstrates, this set may not be connected in a simple way with the set of polyequilibria in Γ^* . Specifically, the example shows that the inclusion indicated by Proposition 2 does not always hold for polyequilibria with non-convex polystrategies.

Example 7. *A polyequilibrium that does not include an equilibrium.* In the mixed extension of the game in Figure 1c, a strategy profile is an equilibrium if and only if players 1 and 2 play T and L , respectively, with probability $1/2$. However, regardless of player 2's strategy, any strategy $pT + p'M + (1 - p - p')B$ of player 1 yields him the same payoff as the corresponding non-equilibrium strategy

$$f(t)T + (f(t) - t)M + (1 - 2f(t) + t)B, \quad (2)$$

where $t = p - p'$ and

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

Therefore, the following polystrategy profile $X = (X_1, X_2)$ is a polyequilibrium: X_1 consists of all strategies of player 1 *except* the equilibrium ones and X_2 consists of all strategies of player 2. A smaller, minimal polyequilibrium is obtained by including in X_1 only the strategies of the form (2), with $-1 \leq t \leq 1$ and f given by (3).⁶ It is easy to check that, in this polyequilibrium, player 2's payoff satisfies $|u_2| \geq 1$. His unique equilibrium payoff, by contrast, is 0.

It would not be possible to find a similar example where player 2's *payoff* (rather than the absolute value of the payoff) in some polyequilibrium X is greater than 0, while his unique equilibrium payoff is 0. This fact is a consequence of the second part of the following proposition.

⁵ 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 players *could* use mixed strategies if they wanted. Both alternatives are legitimate, but there is no escaping choosing one.

⁶ A rational substitution function whose image is this polyequilibrium is essentially defined by f .

Proposition 3. For each 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 i 's mixed-equilibrium payoffs in Γ .

Proof. Example 2 shows that both inclusions may be strict. The first inclusion and the first equality are trivial, and the second ones are special cases of the proposition's second part. 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 2, 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 3 shows that, in the mixed extension Γ^* of a finite game Γ , the only polyequilibrium payoffs are the equilibrium ones. It follows from the second part of the next theorem that for *generic* Γ a stronger proposition holds: *all* the polyequilibrium results in Γ^* are also equilibrium results.

Theorem 1. If a polyequilibrium X in the mixed extension Γ^* of a finite game Γ does not include any equilibrium, then its convex hull $\text{conv } X$ includes infinitely many equilibria. Therefore, if Γ^* has only finitely many equilibria, then it is plain: 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 2. Suppose then that $A = \{x^1, \dots, x^L\}$, with $L \geq 1$. For each $1 \leq l \leq L$, let $\bar{x}^l \in X$ (necessarily, $\bar{x}^l \neq x^l$) be a strategy profile that responds to X , and hence 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 2, except that it is also required that $\{x_i^1, \dots, x_i^L, \bar{x}_i^1, \dots, \bar{x}_i^L\} \subseteq X'_i$. As shown there, 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 $1 \leq l \leq L$

$$u_i(x \mid \bar{x}_i^l) \geq u_i(x \mid x_i^l), \quad x \in X', \quad (4)$$

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 $g_i: X'_i \rightarrow X'_i$ such that

$$u_i(x \mid g_i(x_i)) \geq u_i(x), \quad x \in X' \quad (5)$$

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

$$g_i(x_i) \neq x_i^l, \quad x_i \in X'_i. \quad (6)$$

The meaning of (5) is that changing player i 's strategy from any x_i to $g_i(x_i)$ cannot decrease his payoff in Γ' .

The function is defined as $g_i = g_i^L \circ \dots \circ g_i^1$, the successive composition of L functions $g_i^l: X_i^l \rightarrow X_i^l$ ($l = 1, \dots, L$), which are constructed as follows. If $x_i^l = \bar{x}_i^l$, g_i^l is the identity function, $g_i^l(x_i^l) = x_i^l$. If $x_i^l \neq \bar{x}_i^l$, it is defined by

$$g_i^l(x_i^l) = x_i^l + \alpha_i^l \phi_i^l(x_i^l)(\bar{x}_i^l - x_i^l), \quad (7)$$

where $0 < \alpha_i^l < 1$ is a constant that is chosen as described below and

$$\phi_i^l(x_i^l) = \max\{\alpha \geq 0 \mid x_i^l + \alpha(\bar{x}_i^l - x_i^l) \in X_i^l\}. \quad (8)$$

It is not difficult to see that the function $\phi_i^l: X_i^l \rightarrow [0, \infty)$ defined by (8) is continuous and satisfies

$$\phi_i^l(x_i^l) = \phi_i^l(x_i^l) + \alpha \quad (9)$$

for every $x_i^l, x_i^l \in X_i^l$ and α such that $x_i^l + \alpha(\bar{x}_i^l - x_i^l) = x_i^l$. In particular, for any x_i^l such that $g_i^l(x_i^l) = x_i^l$ (hence, $x_i^l + \alpha_i^l \phi_i^l(x_i^l)(\bar{x}_i^l - x_i^l) + (\bar{x}_i^l - x_i^l) = \bar{x}_i^l$), $\phi_i^l(x_i^l) = \phi_i^l(\bar{x}_i^l) + \alpha_i^l \phi_i^l(x_i^l) + 1$, and therefore

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

where M is the maximum of the function ϕ_i^l . To guarantee that, in fact,

$$g_i^l(x_i^l) \neq x_i^l, \quad x_i^l \in X_i^l, \quad (10)$$

it is required that α_i^l be chosen *greater* than $1 - 1/M$. An additional requirement is that α_i^l be sufficiently close to 1 (but smaller than it) to make the inequality $(1 - \alpha_i^l)M < \phi_i^l(x_i^l)$ hold for all $l' \neq l$ satisfying $\phi_i^l(x_i^l) > 0$. This requirement guarantees that, for every $l' \neq l$ and strategy x_i^l ,

$$g_i^l(x_i^l) = x_i^l \Rightarrow x_i^l = x_i^l. \quad (11)$$

This is because, if $g_i^l(x_i^l) = x_i^l$, then by (9) $\phi_i^l(x_i^l) = \phi_i^l(x_i^l) + \alpha_i^l \phi_i^l(x_i^l)$, so that $\phi_i^l(x_i^l) = (1 - \alpha_i^l)\phi_i^l(x_i^l) \leq (1 - \alpha_i^l)M$, which by the above requirement implies that $\phi_i^l(x_i^l) = 0$, and therefore also $\phi_i^l(x_i^l) = 0$, so that $x_i^l = g_i^l(x_i^l) = x_i^l$.

It follows from (4) and (7) that each of the functions g_i^l satisfies a condition similar to (5) (trivially so if g_i^l is the identity). Since by definition $g_i(x_i) = g_i^L(\dots(g_i^1(x_i))\dots)$, (5) itself clearly also holds. It remains to prove that (6) holds for every l satisfying $x_i^l \neq \bar{x}_i^l$. Suppose that this is not so: for some l' with $x_i^{l'} \neq \bar{x}_i^{l'}$ and strategy x_i^l , $g_i^{l'}(\dots(g_i^1(x_i^l))\dots) = x_i^{l'}$. Necessarily, $l' \neq L$, since an equality would violate (10) for $l = L$. Therefore, by (11) (again with $l = L$), $g_i^{L-1}(\dots(g_i^1(x_i^l))\dots) = x_i^{l'}$. A repeated use of the same argument now shows that $l' \neq l$ also 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 $\bar{\Gamma}$, which 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 , which implies that the set $B_i(x)$ of the player's best response strategies to any strategy profile x is a nonempty convex subset of his strategy set X_i^l . 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 $x^l \neq \bar{x}^l$, there is some i such that $x_i^l \neq \bar{x}_i^l$. By (6), $g_i(\bar{x}_i) \neq x_i^l$, which proves that $x^* = g(\bar{x}) \neq x^l$. Thus, $x^* \notin A$. For every player i and strategy $x_i \in X_i^l$,

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

where the first inequality follows from (5), the second inequality holds since \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' . ■

4.1 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*.⁷ He showed that, in a game as above, the intersection of any number of generalized saddle points is also a generalized saddle point, so that the intersection of all of them, called the *saddle*, is the game's unique minimal strict polyequilibrium. For weak generalized saddle points, this is not so. Two of them may have a nonempty intersection that does not even include a weak generalized saddle point, 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 for 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.

Proposition 4. A two-player zero-sum game is plain if and only if it has an equilibrium.

Proof. Suppose that a two-player zero-sum game Γ has an equilibrium, with payoffs v (which is the *value* of the game) and $-v$ to players 1 and 2 respectively, but some polyequilibrium $X = (X_1, X_2)$ does not include any equilibrium. The first assumption means that both players have optimal strategies, and the second one implies that at least one player's polystrategy in X , say X_1 , does not include any such strategy. Since any of the optimal strategies yields player 1 a payoff of at least v against every strategy in X_2 , there must be some $x_1 \in X_1$ that does the same. Since x_1 is not optimal, there is some strategy x_2 of player 2 against which x_1 yields less than v . However, this conclusion contradicts the definition of polyequilibrium, since it implies that no strategy in X_2 responds to x_1 at least as well as x_2 does. The

⁷ Duggan and Le Breton (1996) use the last term in a somewhat different sense.

contradiction proves that a polyequilibrium X as above cannot in fact exist in a game that has an equilibrium. ■

A finite two-player zero-sum game Γ may or may not have a (pure-strategy) equilibrium. (The former holds, for example, if Γ can be presented as an extensive form game with perfect information. The latter holds for the game in Figure 1b.) However, the mixed extension of such a game, that is, a *matrix game*, always has an equilibrium. It therefore follows from Proposition 4 that, in matrix games, the sets of equilibrium and polyequilibrium results always coincide (which is not true for their non-zero-sum counterparts, as Example 7 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 the mixed extension of the game in Figure 1b, have larger, richer sets of polyequilibria (see Example 6).

5 Successive Elimination of Strategies

In a polyequilibrium, the exclusion of strategies is given an after-the-fact justification. That is, each excluded strategy does not do better than some 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 relations between successive elimination and polyequilibria are explored below. Usefully, the elimination process itself can be naturally described in terms of polystrategies.

Definition 2. *Successive elimination* of (weakly) dominated or strictly dominated strategies means a nonincreasing finite sequence of polystrategy profiles $S = X^0 \supseteq X^1 \supseteq \dots \supseteq X^L$, 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 dominates x_i or strictly dominates it, 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} . A polystrategy profile is said to be obtained by successive elimination of one of these kinds if it is the last entry X^L in some sequence as above.

Clearly, every 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 following two propositions extend these observations to successive elimination.

Proposition 5. A polystrategy profile that is obtained by successive elimination of dominated strategies is a polyequilibrium.

Proof. With the notation of Definition 2, consider any player i and $1 \leq l \leq L$. If a strategy $x'_i \in X_i^l$ dominates a strategy $x_i \in X_i^{l-1} \setminus X_i^l$ relative to X^{l-1} , then in particular the first strategy responds to X^L at least as well as the second one does. The last relation between the two strategies is clearly transitive, which implies that for every strategy $x_i \notin X_i^l$ there is a strategy in X_i^l that responds to X^L at least as well as x_i does. ■

Proposition 6. In a game with the best-response existence property, a polystrategy profile 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. In a game with the best-response existence property, consider successive elimination of never-best-response strategies, that is, a finite sequence $S = X^0 \supseteq X^1 \supseteq \dots \supseteq X^L$ as in Definition 2. Consider also any polyequilibrium X , and a best-response sequence x^1, x^2, x^3, \dots of elements of X as in Fact 5. It is easy to see, by induction, 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 indicates that x^{l+1}, x^{l+2}, \dots are best responses to strategy profiles in X^{l-1} .) In particular, $X \cap X^L \neq \emptyset$. ■

Example 8. *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 $100, 99, \dots, 3$, simultaneously for both players. It follows, by Proposition 6, 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 in which either player does *not* choose 2, and therefore also no polyequilibrium where the total payoff is greater than 4. It follows from the second part of Theorem 1 that the same is true also for the mixed extension of the game.

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 so see that, in a game with the best-response existence property, such a singleton must be an equilibrium. However, if the obtained polystrategy profile X^L 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 in Figure 1b, elimination of the never-best-response strategy M gives the polystrategy profile $\{T, B\} \times \{L, R\}$, which is a curb set and the support of the game's unique mixed-strategy equilibrium but is not a polyequilibrium (because neither T nor B responds to *both* L and R at least as well as M does). The only polyequilibrium the game is the trivial one.⁸

In Proposition 6, 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 a 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.

⁸ It is noteworthy that, even though strategy M is strictly dominated by a *mixed* strategy and is therefore not rationalizable, *not* playing it is not a polyequilibrium result in the (finite) game.




Figure 2. The centipede game (with an odd number of decision nodes m).

Example 9. 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 $(y + z) \bmod 2$. 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, a polyequilibrium is obtained. Thus, each of player 2's two strategies is a polyequilibrium strategy even though it is eliminated during successive elimination of strictly dominated strategies.

6 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. Depending on the kind of strategies the players may use, 'strategy' means either *pure* strategy or *behavior* strategy. That is, at each of the player's information sets it prescribes either an action or a probability distribution over actions, respectively. A polystrategy X_i *excludes* 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 may be viewed as 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 10. *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 2). The odd- and even-numbered nodes are controlled by player 1 and 2, respectively. At each node,




Figure 3. 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.

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 node k 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, each 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 node m). Therefore, a polystrategy profile is any subset of $\{1, 2, \dots, m + 1\}$ (specifically, 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 5, 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 on 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, which means that the latter may be legitimately excluded. Thus, a 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 (corresponding to $l = 2$) is the game’s unique equilibrium. It is thus a plain game.

As the next example shows, there are also perfect-information extensive form games that are not plain. Moreover, there are such games where a player’s payoffs in some polyequilibrium are higher than his unique equilibrium payoff.

Example 11. *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 may object, and in this case, no one receives anything. Assuming that only pure strategies can 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 3a. 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 3b is a variant of the first one, and can be described as involving an additional, inconsequential coin toss. The polystrategy X_i shown for each player i includes three strategies, Black, Gray and Thick, which prescribe choosing the actions indicated by lines with these properties in all decision nodes. 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 one 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 a total payoff of either 1 or 1.5, whereas in every (pure-strategy) equilibrium in the game the total payoff is 2. If a third player, whose only role is to receive the money if the others do not get 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, 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. This fact is actually an immediate corollary of Proposition 3.

6.1 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 extend to strategy profiles, polystrategies and polystrategy profiles.

Definition 3. A polystrategy profile X in a dynamic game G is a *weak subgame perfect polyequilibrium* if, in every subgame of G , the induced polystrategy profile is a polyequilibrium. The last condition may be expressed more explicitly as follows: for every strategy profile $x \notin X$ and every subgame G' there is some $x' \in X$ that in 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 $x \notin X$ 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 4. (Another example is Figure 3, where the polyequilibrium in a is subgame perfect

⁹ A subgame by definition includes either none or all of the decision nodes in every 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).




Figure 4. 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 game, where only player 1's payoffs are shown, $X = \{LR', LL'\} \times \{r, l\}$ is a weak SPP. It is not a SPP, since 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).

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.

A polystrategy profile that is a singleton is a subgame perfect polyequilibrium if and only if its single element is a subgame perfect equilibrium. In general, however, a subgame perfect polyequilibrium may not include a subgame perfect, or even any, equilibrium (see Figure 3a). Two exceptions to this general rule are presented by the next two propositions.

Proposition 7. In a perfect-information extensive form game, every weak subgame perfect polyequilibrium X where each player's polystrategy is rectangular includes a subgame perfect equilibrium.

Proof. Consider the collection \mathcal{X} of all weak subgame perfect polyequilibria that are subsets of X and consist of rectangular polystrategies. (By assumption, $X \in \mathcal{X}$.) 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 nodes following v , only one action or distribution over actions is not excluded, but (ii) this is not so for v itself. Expand the set of excluded actions or distributions over actions at v to include all but a single, optimal one. This modification of X' gives a new polystrategy profile, $X'' \subseteq X'$, which is clearly also an element of \mathcal{X} , a contradiction to the minimality assumption concerning X' . ■

Proposition 8. If a perfect-information extensive form game G has a unique subgame perfect equilibrium x , then every subgame perfect polyequilibrium in G includes x .

Proof. Suppose that $x \notin X$ for some subgame perfect polyequilibrium X . To show that this assumption leads to a contradiction, consider a strategy profile $x' \in X$ that in all subgames responds to X at least as well as x does. There is some subgame G' in which 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, player i 's payoff in G' under x' must be the same as under x , which implies that both strategy profiles induce subgame perfect equilibria in G' . This conclusion clearly contradicts the assumption that G has a unique subgame perfect equilibrium. ■

In the above analysis of the centipede game (Figure 2), 'strategy' actually refers to a number of *equivalent* 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 any one (or more) of them may be arbitrarily 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 10 (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 moreover 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 in *all* subgames at least as well as x_1 does. 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 $2 \leq l \leq 4$. However, for all $2 \leq l \leq m + 1$, X includes the game's unique subgame perfect equilibrium, which corresponds to $l = 2$ (cf. Propositions 7 and 8).¹⁰

A significant 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 is allowed to ignore all but a relatively small number of relevant information sets.

Example 12. Sequential competition (Milchtaich, Glazer and Hassin 2015, Example 4).

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

¹⁰ 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.

first. In the first case, the consumer's payoff is $1 - x - p$, where x is the time he has been in the market and p is the price he paid for the good, and in the second case, the payoff is zero. Thus, the longer a consumer waits until he buys the good, the lower is his valuation of it. The two sellers produce the good at zero cost. Seller 1 arrives at time 1 and seller 2 arrives at time 1.1. An arriving seller announces a price p for the good, sells the demanded quantity q , and immediately leaves. His payoff is thus 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 , by changing the demand seller 2 faces. By definition, a strategy for seller 2 prescribes a price p_2 for every conceivable demand. Thus, it has to consider all possible partitions of the time interval $[0,1]$ that ends with seller 1's arrival into two subsets: the arrival times of the consumers who bought from that seller and of those who did not. 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, except for in the extreme-right case, seller 1 does not take advantage of this by setting his monopoly price of 0.5. The reason is the credible threat implicit in the consumers' strategy, which is the following. 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 the resulting 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 $p_1 > p^*$, only those who arrived in $[2(p_1 - 0.1), 1]$ do so (or no one does, if $p_1 > 0.6$). In the first case, seller 1's profit is $p_1(1 - p_1)$, and it thus attains its maximum, which is between 0.24 and 0.25, at $p_1 = p^*$. In the second case, the profit is only $p_1(1 - 2(p_1 - 0.1))$ (or 0), which is less than 0.16. For player 2, the only relevant decision nodes are those corresponding to the two kinds of time intervals indicated above. In those corresponding to $p_1 \leq p^* (\leq 0.5)$, player 2's profit-maximizing price p_2 is 0.9, and in those corresponding to $p_1 > p^*$, p_2 is the minimum between $p_1 - 0.1$ and 0.5. In the second case, the difference between p_2 and p_1 is precisely that required to compensate 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' strategy is credible, and it is what makes this strategy, together with seller 1's strategy of selling at price p^* and seller's 2 rectangular polystrategy just described, a subgame perfect polyequilibrium.

Note that the polyequilibrium in Example 12 has a well-defined *path* as it prescribes a unique action at every information set that may actually be reached. In this, it differs from every non-equilibrium polyequilibrium in the centipede game, which includes strategy profiles that specify several different paths in the game tree.

6.2 Bayesian Perfection

In the context of dynamic games with imperfect information, the subgame perfection requirement is arguably too weak. A number of refinements of Nash equilibrium take a different approach and require each player's action or distribution over actions at each information set to be optimal with respect to a suitable system of beliefs about the history of play. A *belief system* μ specifies a probability distribution over the nodes in each information set U of every player i . The probability assigned to a set of nodes $V \subseteq U$ is denoted $\mu(V)$ (with $\mu(U) = 1$). A minimal consistency requirement is that this probability coincides with that derived from the players' actual strategy profile x by Bayes' rule if possible, that is,

$$\mu(V) = \frac{\Pr_x(V)}{\Pr_x(U)} \text{ if } \Pr_x(U) \neq 0, \quad (12)$$

where $\Pr_x(V)$ is the probability that, under x , one of the nodes in V is reached (and similarly for $\Pr_x(U)$). A belief system satisfying this requirement is said to be *weakly consistent* with the strategy profile x . Under *perfect recall*, which is assumed below, weak consistency entails that the beliefs concerning an information set U of a player i do not change if only i 's strategy changes, provided that the probability that U is reached is positive for both the old and the new strategy. The reason there is no change is that, since player i has a perfect recall of his actions, all nodes in U are preceded by the same sequence of i 's actions, and therefore the relative probabilities that they are reached only depend on the probabilities of the other players' actions. Thus, for $V \subseteq U$ and all strategies x'_i and x''_i of player i ,

$$\frac{\Pr_{x|x'_i}(V)}{\Pr_{x|x'_i}(U)} = \frac{\Pr_{x|x''_i}(V)}{\Pr_{x|x''_i}(U)} \text{ if } \Pr_{x|x'_i}(U), \Pr_{x|x''_i}(U) \neq 0,$$

where the subscripts refer to the strategy profiles obtained by replacing i 's strategy in x with the indicated strategies.

An existing solution concept that may be viewed as a beliefs-based refinement of the polyequilibrium concept is *essentially perfect Bayesian equilibrium* (Blume and Heidhues 2006). An EPBE may be described as a simple polyequilibrium X that is obtained from a specified strategy profile x by declaring certain information sets *irrelevant*. The collection of irrelevant information sets must (i) have zero probability of being reached when x is played and (ii) include every information set that follows any irrelevant one. X is then defined as the collection of all strategy profiles that agree with x (in the sense of specifying the same action or distribution over actions) in each of the *relevant* (that is, not irrelevant) information sets. By requirement (i), any belief system μ that is weakly consistent with x is also weakly consistent with every other strategy profile in X . Requirement (ii) implies that the polystrategy X_i of each player i is a special kind of rectangular polystrategy. The definition of EPBE is completed by the additional requirement that there is some weakly consistent belief system μ as above (which may be considered part of the EPBE) such that, for every relevant information set U of every player i , strategy x_i is a best response to X in the continuation game starting at U with the distribution over U 's nodes specified by μ . This requirement implies that, for each player i , strategy x_i is a best response to X also in the whole game, so that X is a simple polyequilibrium.

Essentially perfect Bayesian equilibrium may be regarded a solution to technical problems associated with the perfect Bayesian equilibrium solution concept (Blume and Heidhues 2006). Alternatively, as indicated, it may be viewed as one way of incorporating optimality with respect to a weakly consistent belief system into the polyequilibrium concept. However, EPBS has certain features that, in the general context of polyequilibria in dynamic games, are quite special and restrictive. The first special, albeit arguably desirable, aspect is that an EPBS is required to specify a unique outcome: all strategy profiles in the polyequilibrium give the same distribution over terminal nodes. The second restrictive aspect is that, at each information set, either the acting player's choice is completely specified or it is left completely unspecified, and in the second case, the same must hold also for each of the following information sets. As Blume and Heidhues (2006) point out, one consequence of this requirement is that, in the (perfect information) centipede game (Example 10), the only EPBE is the subgame perfect equilibrium.

The limitations of the EPBS concept reflect a fundamental tension between the underlying principles of Bayesian perfection and polyequilibrium, which concerns the way the players' beliefs at the various information sets relate to the actual strategy profile. In a perfect Bayesian equilibrium and a number of related solution concepts, beliefs are reasonable *hypotheses*, or conjectures, about the events that preceded the arrival at an information set. In particular, a choice of action at an off-path information set may be justified by a theory as to why the set has been reached even though it should not have been reached according to the strategy profile. Different solution concepts have different definitions of "reasonable" hypotheses and different requirements as to how they should reflect the players' actual strategies. However, they all leave at least some leeway, which means that beliefs are often effectively *chosen* much like strategies are chosen. It is hard to reconcile this element of choice with the spirit of an excluding solution concept such as polyequilibrium.

There is, however, a simple, natural way to incorporate beliefs in a manner consistent with exclusion rather than choice. In a polyequilibrium X , the exclusion of a strategy of a player i requires the existence of an alternative strategy that responds to every strategy profile $x \in X$ at least as well as the first strategy does. This condition may be extended by replacing "every $x \in X$ " with "every $x \in X$ and every belief system that is consistent with x ". Thus, at each of player i 's information sets U , the alternative strategy should be an adequate substitute to the excluded one for every strategy profile $x \in X$ and every distribution over U 's nodes that is consistent with it. If under x the probability that U is reached is positive, the consistency requirement singles out a unique distribution there. Otherwise, there may be many consistent distributions, each reflecting a different plausible hypothesis as to why U has been reached. As detailed below, the idea that all these hypotheses need to be considered leads to the notion of *perfect Bayesian polyequilibrium* (PBP).

Perfect Bayesian polyequilibrium is not a generalization of perfect Bayesian equilibrium but a fundamentally different, stronger concept. The difference is illustrated by Figure 5a. The equilibrium shown is perfect Bayesian, since Player 2's choice of r rather than l is justified by *some* belief, namely, that if player 1 switched from his strategy *Out* to some alternative strategy, it would be R rather than L . However, *any* belief of player 2 at his off-equilibrium information set is consistent with *Out*. Therefore, in a PBP, which requires the exclusion of a




Figure 5. Perfect Bayesian equilibria (black lines) that are not (singleton) perfect Bayesian polyequilibria. **a** Player 2's choice of r is justified by the belief $p = 0$ but not by $p = 1$. **b** Player 2's choice of P is justified only for $1/3 \leq p \leq 2/3$. (In this game, players 3's payoffs are identical to those of 1, and are not indicated.)

strategy to be justified under all consistent beliefs, l cannot be excluded, and consequently L and R cannot be both excluded. Therefore, there is no PBP where player 1 plays *Out*, so that this perfect Bayesian equilibrium result and player 2's maximal payoff of 2 are not PBP results.

The last example highlights a basic difference between perfect Bayesian polyequilibrium and other solution concepts such as perfect Bayesian equilibrium, essentially perfect Bayesian equilibrium and sequential equilibrium. The difference is that a PBP is not allowed to specify actions that are justifiable only under particular *arbitrary* beliefs about the history of play. The choice of such actions (for example, r in Figure 5a) entails the exclusion of one or more alternative actions that under different but equally reasonable beliefs are actually better. The disallowance of such exclusions may make it impossible to single out *any* action. Such an outcome is of course entirely in tune with the central idea of polyequilibrium, which is that the players' strategies may be only partly specified.

Example 13. Only player 1 knows whether or not it will rain tomorrow – two possibilities that a priori are equally likely. The value of this information for player 2, who would prefer taking an umbrella only if necessary, is \$1. Player 1 gives him a take-it-or-leave-it offer to buy the information. His payoff is the asking price p if the offer is accepted and 0 otherwise. For simplicity, only pure strategies are allowed. Every $0 \leq v_1 \leq 1$ is a perfect Bayesian equilibrium payoff for player 1. It is obtained in an equilibrium where $p = v_1$ and player 2 is willing to pay this price but would reject the offer, and take an umbrella, if player 1 asked any other price. This reaction is justified by a belief that a price different from v_1 indicates rain. By contrast, the only perfect Bayesian polyequilibrium payoff for player 1 is 1. To see this, consider a perfect Bayesian polyequilibrium X where the payoff is $v_1 < 1$. For every $v_1 < p < 1$, player 2's polystrategy X_2 excludes acceptance of this price, for if it included a strategy that prescribes acceptance, it would not be possible to exclude player 1's strategy of asking p . Consider some such p that is also different from the prices specified by some strategy $x_1 \in X_1$. (In principle, player 1's asking price may depend on whether or not it will rain tomorrow.) If player 1 uses x_1 , player 2's information set where he is asked to pay p is not reached. No action there does as well as the excluded one of acceptance under *all*

beliefs, or even only under the two extreme ones: that it will surely rain tomorrow and that it will not rain. For example, with no rain, rejecting the offer and taking an umbrella is a worse option than paying p and acting according to the provided information. This contradiction leaves $v_1 = 1$ as the only perfect Bayesian polyequilibrium payoff. It is obtained in the polyequilibrium where player 1's strategy is $p = 1$ and player 2's polystrategy is acceptance of this price and rejection of any higher one. The reaction to prices lower than 1 is left unspecified.

Solution concepts that involve a single, specific belief system μ usually impose on it certain *internal consistency* requirements, which express the idea that beliefs at different off-path information sets should not only reflect the players' strategy profile but also represent a coherent hypothesis about their deviation from it. In particular, beliefs at an information set that follows another information set of the same player should be derived from the beliefs at the latter whenever possible. This requirement is formally expressed by the *preconsistency* condition (Hendon et al. 1996, Perea 2002), which is based on Fudenberg and Tirole's (1991) notion of reasonable assessment. Internal consistency between beliefs at information sets belonging to different players is guaranteed by the stronger *full consistency* condition, which is the central pillar of Kreps and Wilson's (1982) sequential equilibrium solution concept.¹¹ Perfect Bayesian polyequilibrium does not specify a single belief system, which renders the whole internal consistency idea moot. This brings about a considerable simplification, since consistency is narrowed down to the "local" requirement that beliefs at a particular information set U of a particular player i are reconcilable with a specified strategy profile x .

If under x the probability that U is reached is zero, the weak consistency condition (Eq. (12)) does not specify any beliefs there. Yet player i may actually *know* a great deal about the history of play at that information set. Specifically, he knows that at another information set the acting player j took a particular action a if *all* nodes in U are preceded by (that information set and) action a . (By the perfect-recall assumption, this is so in particular for each of player i 's own past actions.) This means that in the set \mathcal{A} of all actions a as above the probabilities specified by x are irrelevant; the effective probabilities are 1.¹² Therefore, it remains for player i to only speculate about the other players' behavior at information sets that do not involve actions in \mathcal{A} . The simplest hypothesis is that the players adhere to x there. In other words, the hypothesis effectively replaces x with the strategy profile $x^{\mathcal{A}}$ obtained from it by specifying that every action in \mathcal{A} is taken with probability 1. If under $x^{\mathcal{A}}$ the probability that U is reached is positive, this hypothesis yields by Bayes' rule a unique probability distribution on the nodes in U , which arguably represents the unique beliefs

¹¹ In an extensive form game, a belief system μ is said to be fully consistent with a strategy profile x if they are, respectively, the limits of some sequences $(\mu^n)_{n=1}^{\infty}$ and $(x^n)_{n=1}^{\infty}$ such that each x^n is a *completely mixed* strategy profile, that is, one assigning positive probability for every action in every information set, and μ^n is the unique belief system that is weakly consistent with x^n .

¹² Probability 1 here means that the action was actually taken, not that the acting player j *meant* it to be played for sure. For an action a that precedes all the nodes in U , there is no effective difference between 1 and any other positive probability. Therefore, if the probability that player j 's strategy assigns to a is not 0, it may be left unchanged. Note that the discussion here and the definitions below are relevant both when the players are assumed to use only pure strategies and when all behavior strategies are allowed.

there that are consistent with x . However, if even under $x^{\mathcal{A}}$ the probability that U is reached is zero, then reaching it indicates that at least one action $a \notin \mathcal{A}$ was taken even though the probability assigned to a by x in the relevant information set is 0. As a generalization of the above reasoning, consider a set of actions \mathcal{A}' that is *minimally sufficient* for reaching U under x , in the sense that U is reached with positive probability under $x^{\mathcal{A}'}$ but not under $x^{\mathcal{B}}$, for all $\mathcal{B} \subsetneq \mathcal{A}'$ (which clearly implies that x assigns probability 0 to every action in \mathcal{A}'). Each minimally sufficient set \mathcal{A}' specifies by Bayes' rule a probability distribution on U , which may be viewed as consistent with the players' strategy profile x if U was reached.¹³ Note that the set of all such distribution is finite if the information set itself is finite. This is because a set of actions that is minimally sufficient for reaching a finite U necessarily consists of all actions that precede some node $v \in U$ and are assigned probability 0 by x .

A deviation from x that involves a non-minimally sufficient set of actions represents a non-parsimonious hypothesis as to why an information set U was reached; it assumes more than it has to. Moreover, such a hypothesis may have the troubling aspect that it implies a *future* deviation from x . This can happen if some players have information sets that include both nodes that precede U and nodes that follow it, as in Figure 5b, which is taken from Kreps and Ramey (1987). In this game, the players' order of moves is random – either player 1 or player 3 moves before player 2 moves – and is unknown to them. The only (Nash) equilibrium outcome is that players 1 and 3 both play Q for sure, so that player 2's information set is not reached. The choice of Q reflects the fact that there is no equilibrium strategy for player 2 that specifies playing either L or R with probability $2/3$ or greater. In particular, neither action is played with probability 1, which means that any equilibrium strategy for player 2 can be justified only by beliefs that assign positive probability to both nodes in his information set. Obtaining such probabilities require that both player 1 and player 3 deviate from their equilibrium strategies by playing P with positive probability. However, such simultaneous deviations are inconsistent with an assumption that the player acting *after* player 2 will be using his equilibrium strategy. Thus, this example shows that a non-parsimonious hypothesis about the past may project onto the future. With a parsimonious hypothesis about the deviations from x that led to an information set U of a player i being reached, this cannot happen. If a set of actions \mathcal{A} is minimally sufficient for reaching U under x , then the probability that some information set where x and $x^{\mathcal{A}}$ disagree is reached *after* U is reached is zero, both under $x^{\mathcal{A}}$ and under any strategy profile that differs from it only in player i 's strategy. This is because any history that has positive probability under such a strategy profile and reaches U must, by the minimal-sufficiency assumption, first go through *all* the actions in \mathcal{A} , which by the perfect-recall assumption means that it cannot reach any of the corresponding information sets again after U . A parsimonious hypothesis is thus *structurally consistent* with x in the sense of Kreps and Ramey (1987, p. 1338), which is a stronger requirement than Kreps and Wilson's (1982) notion of structural consistency of beliefs (for which the strategy profile is irrelevant). This is not the case for the beliefs that justify player 2's equilibrium strategy in the above example, which are only *convex structurally consistent*. That is, these beliefs can be obtained only as

¹³ It follows immediately from the minimality of \mathcal{A}' that the same distribution on U would result also if the actions in \mathcal{A}' were assigned arbitrary positive probabilities, that are not necessarily 1.

convex combinations of the two structurally consistent ones, which are those that assign probability 1 to one of player 2's nodes and reflect a hypothesized deviation by only one, particular other player.

The fundamental reason why this problematic aspect of perfect Bayesian and sequential equilibria is not shared by perfect Bayesian polyequilibrium lies in the logical difference between justifying the choice of a particular strategy and justifying its exclusion. As indicated, justifying the exclusion of a strategy of a player i under a strategy profile x involves examining, for each of the player's information sets U , all distributions over U 's nodes that are consistent with x . If the examination of two or more such distributions reveals that the excluded strategy indeed does not do better than the specified alternative strategy, then this is automatically the case also for all convex combinations of these distributions, which therefore do not need to be explicitly examined. By contrast, the *choice* of a strategy may be justified only by the beliefs expressed by one, particular convex combination. Hence the difference.

The above discussion leads to the following formal definition of PBP, where consistent beliefs are expressed as appropriately modified strategy profiles.¹⁴ The definition uses the following notation. For an information set U of a player i , $u_i(x \mid_U x'_i)$ denotes the player's payoff under a strategy profile that differs from x only in that, at U and all the information sets that follow it, player i plays according to strategy x'_i .

Definition 4. A polystrategy profile X in a dynamic game with perfect recall is a *perfect Bayesian polyequilibrium* if for every strategy profile $x \notin X$ there is some $x' \in X$ such that, for every player i and information set U of that player, the inequality

$$u_i(x''^{\mathcal{A}} \mid_U x'_i) \geq u_i(x''^{\mathcal{A}} \mid_U x_i) \quad (13)$$

holds for every $x'' \in X$ and every set of actions \mathcal{A} that is minimally sufficient for reaching U under x'' .¹⁵

The next proposition shows that PBP may be considered a generalization of subgame perfect polyequilibrium. As discussed above, it is not a generalization of perfect Bayesian equilibrium, essentially perfect Bayesian equilibrium or sequential equilibrium. The exact sense in which the last solution concept is weaker than singleton PBP is spelled out by Theorem 2 below.

¹⁴ Replacing this notion of consistency, which is the one presented above, with any weaker notion, and in particular with weak consistency, would result in a *stronger* definition. That is, the set of PBPs would be smaller. The opposite is true for any stronger notion of consistency.

¹⁵ If U is reached with positive probability under x'' , then the only minimally sufficient set is $\mathcal{A} = \emptyset$, for which $x''^{\mathcal{A}} = x''$. Technically, the definition is meaningful even if minimally sufficient sets do not exist, which can happen with a continuum of outcomes to chance moves. However, this case would be addressed more satisfactorily if sure past outcomes were assigned probability 1, like players' actions. A more comprehensive modification of the definition will need to deal with conditional probability distributions on zero-probability events, which is probability doable but technically more challenging.

Proposition 9. In an extensive form game, every perfect Bayesian polyequilibrium is also subgame perfect. With complete information, the converse holds too.

Proof. For given PBP X and strategy profile x , let $x' \in X$ be as in Definition 4, and consider any $x'' \in X$, player i and subgame G' . Denote player i 's payoff function in G' by u'_i . It has to be shown that, in G' , strategy x'_i responds to x'' at least as well as x_i does. Clearly, a sufficient condition for this is that $u'_i(x'' \mid_U x'_i) \geq u'_i(x'' \mid_U x_i)$ for every information set U of player i in G' that is not preceded by any other such information set and is reached with positive probability in G' when the strategy profile is x'' . The inequality is equivalent to (13), where \mathcal{A} is the set of all actions preceding the root of G' that are assigned probability 0 by x'' . This completes the proof of the first assertion in the proposition. The second assertion is clear from the fact that, in a complete-information game, every information set includes only one node, which is the root of a subgame. ■

Theorem 2. A strategy profile x in an extensive form game is a (singleton) perfect Bayesian polyequilibrium if and only if (x, μ) is a sequential equilibrium for *every* belief system μ that is fully consistent with x .

The proof of the theorem is based on the following lemma, which is of independent interest. The lemma connects three collections of beliefs at a particular information set: those that arise from fully consistent assessments, those arising from minimally sufficient sets of actions, and those that are convex combinations of the latter. It shows that the first collection includes the second collection and is included in the third one. (It can furthermore be shown that, in general, both inclusions are strict.) This result strengthens and extends the Proposition in Kreps and Ramey (1987), which asserts that every fully consistent assessment satisfies convex structural consistency.

Lemma 1. For a strategy profile x and information set U in an extensive form game, let \mathfrak{A} be the (finite) collection of all sets of actions that are minimally sufficient for reaching U under x . For $\mathcal{A} \in \mathfrak{A}$, let $\mu^{\mathcal{A}}$ be the probability distribution on U induced by $x^{\mathcal{A}}$:

$$\mu^{\mathcal{A}}(V) = \frac{\Pr_{x^{\mathcal{A}}}(V)}{\Pr_{x^{\mathcal{A}}}(U)}, \quad V \subseteq U. \quad (14)$$

- (i) The probability distributions $\{\mu^{\mathcal{A}}\}_{\mathcal{A} \in \mathfrak{A}}$ have pairwise disjoint supports.
- (ii) Each of them $\mu^{\mathcal{A}}$ coincides with the probability distribution on U specified by some belief system μ that is fully consistent with x :

$$\mu^{\mathcal{A}}(V) = \mu(V), \quad V \subseteq U. \quad (15)$$

- (iii) For every belief system μ that is fully consistent with x , the probability distribution on U specified by μ is a convex combination of the distributions $\{\mu^{\mathcal{A}}\}_{\mathcal{A} \in \mathfrak{A}}$:

$$\mu(V) = \sum_{\mathcal{A} \in \mathfrak{A}} \lambda^{\mathcal{A}} \mu^{\mathcal{A}}(V), \quad V \subseteq U \quad (16)$$

for some (unique, by (i)) nonnegative coefficients $\{\lambda^{\mathcal{A}}\}_{\mathcal{A} \in \mathfrak{A}}$ that sum up to 1.

Proof. To prove (i), it has to be shown that $\Pr_{x^{\mathcal{A}}}(\{v\})\Pr_{x^{\mathcal{A}'}}(\{v\}) = 0$ for all $v \in U$ and $\mathcal{A}, \mathcal{A}' \in \mathfrak{A}$ with $\mathcal{A}' \neq \mathcal{A}$. If $\Pr_{x^{\mathcal{A}}}(\{v\}) > 0$, then the minimal sufficiency of the set of actions \mathcal{A} implies that all its elements precede node v . Because the actions in $\mathcal{A} \setminus \mathcal{A}'$ are assigned probability 0 by $x^{\mathcal{A}'}$, this implies that $\Pr_{x^{\mathcal{A}'}}(\{v\}) = 0$.

To prove (ii), consider any $\mathcal{A} \in \mathfrak{A}$, with cardinality $|\mathcal{A}| = L (\geq 0)$. Let \tilde{x} be some fixed completely mixed strategy profile and, for $0 < \epsilon < 1/2$, let x^ϵ be the strategy profile that, at each information set, assigns the following probability $x^\epsilon(a)$ to each action a :

$$x^\epsilon(a) = (1 - \epsilon - \epsilon^{L+1})x(a) + \epsilon x^{\mathcal{A}}(a) + \epsilon^{L+1}\tilde{x}(a), \quad (17)$$

where $x(a)$, $x^{\mathcal{A}}(a)$ and $\tilde{x}(a)$ are the probabilities specified by x , $x^{\mathcal{A}}$ and \tilde{x} . The unique belief system μ^ϵ that is weakly consistent with x^ϵ satisfies

$$\mu^\epsilon(V) = \frac{\Pr_{x^\epsilon}(V)}{\Pr_{x^\epsilon}(U)}, \quad V \subseteq U. \quad (18)$$

For any $v \in V$, $\Pr_{x^\epsilon}(\{v\}) = \prod_k x^\epsilon(a^k)$, where the a^k 's are all the actions preceding node v . Using (17), this product can be expressed as a polynomial in ϵ . For $l < L$, the coefficient of ϵ^l is 0, because a positive coefficient would mean that $\Pr_{x^{\mathcal{B}}}(\{v\}) > 0$ for some $\mathcal{B} \subsetneq \mathcal{A}$, which contradicts the minimal-sufficiency assumption concerning \mathcal{A} . By a similar argument, the coefficient of ϵ^L is $\Pr_{x^{\mathcal{A}}}(\{v\})$. It follows that, for $V \subseteq U$, $(1/\epsilon^L) \Pr_{x^\epsilon}(V) \rightarrow \Pr_{x^{\mathcal{A}}}(V)$ as $\epsilon \rightarrow 0$, which implies that the quotient in (18) converges to that in (14). Therefore, if $(\epsilon_n)_{n=1}^\infty$ is any sequence of positive numbers converging to 0 such that $(\mu^{\epsilon_n})_{n=1}^\infty$ converges to some limit μ , then that belief system, which is clearly fully consistent with x , satisfies (15). The existence of such a sequence follows from the obvious compactness of the set of all belief system.

To prove (iii), consider any belief system μ that is fully consistent with x , and some sequences $(\mu^n)_{n=1}^\infty$ and $(x^n)_{n=1}^\infty$ as in footnote 11. For every $\mathcal{A} \in \mathfrak{A}$, the minimal-sufficiency condition implies that there is some $u \in V$ that is preceded by all the actions in \mathcal{A} and satisfies $\Pr_{x^{\mathcal{A}}}(\{u\}) > 0$. For every $n \geq 1$, $\Pr_{x^n}(\{u\}) = \prod_k x^n(a^k)$, where the a^k 's are all the actions preceding u . Therefore,

$$\frac{1}{\prod_{a \in \mathcal{A}} x^n(a)} \Pr_{x^n}(\{u\}) = \prod_{\substack{k \\ a^k \notin \mathcal{A}}} x^n(a^k) \xrightarrow{n \rightarrow \infty} \Pr_{x^{\mathcal{A}}}(\{u\}). \quad (19)$$

A result similar to (19) holds with u replaced by any other node $v \in V$ that is preceded by all the actions in \mathcal{A} , which implies that

$$\frac{\mu^{\mathcal{A}}(\{v\})}{\mu^{\mathcal{A}}(\{u\})} = \frac{\Pr_{x^{\mathcal{A}}}(\{v\})}{\Pr_{x^{\mathcal{A}}}(\{u\})} = \lim_{n \rightarrow \infty} \frac{\Pr_{x^n}(\{v\})}{\Pr_{x^n}(\{u\})} = \lim_{n \rightarrow \infty} \frac{\mu^n(\{v\})}{\mu^n(\{u\})} = \frac{\mu(\{v\})}{\mu(\{u\})}$$

if $\mu(\{u\}) > 0$, and if $\mu(\{u\}) = 0$, then $\mu(\{v\}) = 0$. Therefore, for $v \in V$ that is preceded by all the actions in \mathcal{A} , $\mu(\{v\}) = \lambda^{\mathcal{A}} \mu^{\mathcal{A}}(\{v\})$, where $\lambda^{\mathcal{A}} = \mu(\{u\})/\mu^{\mathcal{A}}(\{u\})$. As shown in the first part of the proof, such v also satisfies $\mu^{\mathcal{A}'}(\{v\}) = 0$ for all $\mathcal{A}' \in \mathfrak{A}$ with $\mathcal{A}' \neq \mathcal{A}$, so that the equality in (16) holds for $V = \{v\}$. To prove that it holds generally, it remains to




Figure 6. The destructive and constructive potential of beliefs reflecting strategic considerations. **a** In the (singleton) polyequilibrium shown, player 2's action is justified by the unique beliefs consistent with player 1's strategy in the subgame. However, a deviation by player 1 from his strategy that leads to the subgame being reached may suggest an additional deviation there – an attempt to get positive payoff. **b** The Beer-Quiche game. The game has two pure-strategy perfect Bayesian equilibria (black and gray lines), only one of which (black) satisfies the intuitive criterion. Both equilibria are not perfect Bayesian polyequilibria. However, adopting the restriction on off-equilibrium beliefs underlying the intuitive criterion would make one of them (black) a PBP.

note that every $v \in V$ is preceded by all the actions in *some* $\mathcal{A} \in \mathfrak{A}$, since the set of actions preceding v that are assigned probability 0 by x necessarily has a subset that is minimally sufficient for reaching U under x . Setting $V = U$ in (16) proves that the coefficients sum up to 1. ■

Proof of Theorem 2. By definition, x is a PBP if and only if, for every player i and strategy y_i and information set U of that player, the inequality

$$u_i(x^{\mathcal{A}} \mid_U x_i) \geq u_i(x^{\mathcal{A}} \mid_U y_i)$$

holds for every set of actions \mathcal{A} that is minimally sufficient for reaching U under x . It follows from (ii) in Lemma 1 that a sufficient condition for this is that (x, μ) is a sequential equilibrium for every belief system μ that is fully consistent with x . Conversely, it follows from (iii) in the lemma that the last condition holds if x is a PBP. ■

6.3 Beliefs Based on Strategic Reasoning

The notion of consistency implicit in Definition 4 is based on the principle of parsimony: a particular deviation from the players' strategies is assumed to have occurred only if this assumption is needed for explaining why an off-path information set was reached. However, the simplest explanation is not always the most convincing one. In particular, a *forward induction* argument (Kohlberg and Mertens 1986) may lend credence to *inconsistent* beliefs. That is, a detected past deviation of another player from his strategy may hint at an additional, unobservable deviation. Unlike consistency, which is a notion based wholly on the game form, forward induction also involves an examination of the strategic interests of the deviating players. For example, in the perfect Bayesian equilibrium and (singleton) polyequilibrium shown in Figure 6a, player 2's choice of r is justified by a belief that player 1 would follow his strategy in the subgame and choose R there. However, if 2's information set is actually reached, which indicates that player 1 deviated from his strategy in the whole

game by playing ln , player 2 may reason that the most likely explanation for the deviation is that player 1 is aiming for the better equilibrium of the subgame, with payoffs 1, which means that he played L rather than R there. Such a belief makes l the better choice for player 2.

Past deviations may also be taken as indicators of intended *future* ones. This possibility is illustrated by the game that differs from that in Figure 6a only in the order of moves in the subgame: player 2 chooses his action before player 1 does so. As the two moves are actually effectively simultaneous, this game is essentially identical to that considered above, so that the same argument applies: if player 1 deviated once, then strategic considerations suggest he intends to deviate again.

A different potential outcome of strategic reasoning is ruling out certain beliefs that are *not* inconsistent with the players' strategies. For example, a belief that some player chose a strictly dominated action may be considered unreasonable if there is an alternative explanation for reaching an off-equilibrium information set that does not involve dominated actions. In games with multiple perfect Bayesian or sequential equilibria, restriction to reasonable consistent beliefs may eliminate some of the equilibria. For example, it may eliminate all pooling equilibria in Spence's signaling model with two types of worker (Cho and Kreps 1987). The effect of a similar restriction on beliefs on PBPs is in a sense the diametric opposite. In Spence's signaling model, in particular, none of the pooling equilibria is a perfect Bayesian polyequilibrium to begin with. The reason is that the choice of any non-equilibrium education level cannot be excluded, because consistency in the sense described in the previous subsection does not preclude a belief by the employer that such a choice indicates a high quality worker. Thus, as for equilibria, further restrictions on off-equilibrium beliefs, based on strategic considerations, may be warranted. However, as such restrictions effectively lead to a stronger notion of consistency, they do not eliminate PBPs but can only *add* new ones (see footnote 14). Thus, in particular, the set of singleton perfect Bayesian polyequilibria, which is typically contained in the sets of perfect Bayesian and sequential equilibria (Theorem 2), expands while the latter contract, which means that the gap between the corresponding sets of results may narrow. The following example illustrates this possibility.

Example 14. Consider the Beer-Quiche game shown in Figure 6b, where for simplicity only pure strategies are allowed. There are two (Nash) equilibria, which are both pooling. In one equilibrium (black lines), types t_w and t_s of player 1 both choose Beer, and in the other (gray lines), they choose Quiche. The second equilibrium is eliminated by the *intuitive criterion* (Cho and Kreps 1987). The criterion is based on a restriction of player 2's possible beliefs regarding player 1's type, which in particular precludes beliefs that, following a choice of Beer, attach a positive probability to t_w . The reason such beliefs are deemed unreasonable is that this type's equilibrium payoff of 3 is higher than anything he may get by choosing Beer. The same problem does not arise in the first equilibrium (black lines), where both types of player 1 choose Beer and player's 2 would choose Duel only as a response to Quiche. That response is justified by the unique reasonable belief following a choice of Quiche by player 1, which is that his type is t_w (because t_s would necessarily be harmed by such a choice). Put differently, the unreasonableness of all other beliefs at the off-equilibrium information set

justifies the exclusion of the alternative action (Don't Duel) there. The same argument also shows that a restriction to reasonable beliefs would make the equilibrium under consideration a PBP. Thus, the logic underlying the intuitive criterion singles out the same equilibrium for both solution concepts, perfect Bayesian equilibrium and perfect Bayesian polyequilibrium. This coincidence contrasts with the situation for the original, unmodified definitions, according to which both equilibria are perfect Bayesian equilibria but neither of them is a PBP. It is, however, a rather special outcome, which is due to the fact that the additional reasonableness requirement on off-equilibrium beliefs pins them down uniquely.

7 Incomplete Information Games

In a game with incomplete information, the players' payoffs are determined by both their actions and the *state (of the world)* ω , which is an element of some set Ω . Each player i is associated with a certain partition of Ω , the player's *information partition*, whose elements are i 's *information sets*. Each state ω thus lies in one, and only one, information set I of player i . The player knows the identity of I but, unless I is a singleton, does not know which of its element is the true state. A *strategy* for player i prescribes an action at each of his information sets. In other words, it is a mapping from Ω to i 's set of actions that is measurable with respect to the player's information partition. A strategy profile thus determines the action each player takes at each state. The players' payoffs in the game are their resulting expected payoffs, where the expectation is with respect to a specified probability measure on Ω , the players' *common prior*.

Example 15. *The electronic mail game* (Rubinstein 1989). This two-player game has a countably infinite set of states, $\Omega = \{\omega_0, \omega_1, \omega_2, \dots\}$. Players 1 and 2 must each choose one of two actions, A or B . Depending on whether the state is ω_0 or one of $\omega_1, \omega_2, \dots$, their payoffs are given by the payoff matrices

$$\begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} 1,1 & 0,-L \\ -L,0 & 0,0 \end{pmatrix} \end{array} \text{ or } \begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} 0,0 & 0,-L \\ -L,0 & 1,1 \end{pmatrix},$$

respectively, where $L > 1$ is a fixed parameter. The common prior assigns probability $1/2 < p < 1$ to ω_0 and probability $(1-p)\epsilon(1-\epsilon)^{l-1}$ to every other state ω_l ($l \geq 1$), with $0 < \epsilon < 1$. The information partition of player 1 is $\{\{\omega_0\}, \{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \dots\}$ and that of player 2 is $\{\{\omega_0, \omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4, \omega_5\}, \dots\}$. A strategy profile x can therefore be viewed as a sequence of actions a^0, a^1, a^2, \dots , where each entry $a^l \in \{A, B\}$ ($l = 0, 1, 2, \dots$) is the action taken at the information set where the highest-index state is ω_l . Depending on whether l is even or odd, that information set belongs to player 1 or 2, respectively.

Since $p > 1/2$, a strategy profile x may yield a payoff of p or higher to either player only if $a^0 = A$. Among all strategy profiles satisfying this condition, only one is an equilibrium, namely, that where $a^l = A$ for *all* l (Rubinstein 1989, Proposition 1). In the equilibrium, both players' payoff is precisely p . Thus, there is a no equilibrium where either player receives more than p .

The equilibrium just described is also the only *polyequilibrium* in the electronic mail game where some player i ($\in \{1,2\}$) receives at least p . To see this, consider any such polyequilibrium X . Necessarily, every strategy profile in X satisfies $a^0 = a^1 = A$. Consider any best-response sequence x^1, x^2, x^3, \dots of elements of X as in Fact 5. It is not difficult to show, by induction, that for all l the strategy profile x^l satisfies $a^0 = a^1 = \dots = a^l = A$. Consider now any strategy $x_i \in X_i$. If one or more of i 's actions in x_i were B , the limit $\lim_{l \rightarrow \infty} u_i(x^l | x_i)$ would be less than p , a contradiction to the assumption concerning the player's payoff in X . Therefore, X_i must be the singleton whose single element is the (equilibrium) strategy of player i where he chooses A in all information sets. This strategy is not a best response to any strategy of the other player j that prescribes choosing action B at any information set of that player other than the first one. Therefore, j 's polystrategy cannot include any such strategy, which means that the polyequilibrium X includes only the equilibrium strategy profile $a^0 = a^1 = \dots = A$.

References

- Basu, K. (1994). The Traveler's Dilemma: Paradoxes of rationality in game Theory. *American Economic Review* 84, 391–395.
- Basu, K. and Weibull, J. W. (1991). Strategy subsets closed under rational behavior. *Economics Letters* 36, 141–146.
- Bernheim, D. (1984). Rationalizable strategic behavior. *Econometrica* 52, 1007–1028.
- Blume, A. and Heidhues, P. (2006). Private monitoring in auctions. *Journal of Economic Theory* 131, 179–211.
- Cho, I-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics* 102, 179–221.
- Duggan, J. and Le Breton, M. (1996). Dutta's minimal covering set and Shapley's saddles. *Journal of Economic Theory* 70, 257–265.
- Fudenberg, D. and Tirole, J. (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory* 53, 236–260.
- Hendon, E., Jacobsen, J. and Sloth, B. (1996). The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior* 12, 274–282.
- Kalai, E. and Samet, D. (1984). Persistent equilibria in strategic games. *International Journal of Game Theory* 13, 129–144.
- Kohlberg, E. and Mertens, J.-F. (1986). On the strategic stability of equilibria. *Econometrica* 54, 1003–1037.
- Kreps, D. M. and Ramey, G. (1987). Structural consistency, consistency, and sequential rationality. *Econometrica* 55, 1331–1348.
- Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica* 50, 863–894.

- Milchtaich, I., Glazer, A. and Hassin, R. (2015). [Time-differentiated monopolies](#). Manuscript.
- Moulin, H. (1979). Dominance solvable voting schemes. *Econometrica* 47, 1337–1351.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029–1050.
- Perea, A. (2002). A note on the one-deviation property in extensive form games. *Games and Economic Behavior* 40, 322–338.
- Pruzhansky, V. (2003). On finding curb sets in extensive games. *International Journal of Game Theory* 32, 205–210.
- Rosenthal, R. (1981). Games of perfect information, predatory pricing and the chain store paradox. *Journal of Economic Theory* 25, 92–100.
- Rubinstein, A. (1989). The electronic mail game: Strategic behavior under “almost common knowledge”. *American Economic Review* 79, 385–391.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 25–55.
- Shapley, L. S. (1964). Some topics in two-person games. In: *Advances in Game Theory* (*Annals of Mathematics Studies* 52). M. Dresher, L. S. Shapley, and A. W. Tucker (Eds.), Princeton: Princeton University Press, 1–28.