

Agnostic Sequential Rationality

Igal Milchtaich*

November 2018

Agnostic sequential equilibrium is a refinement of sequential equilibrium that does not force on the players a single, arbitrary belief system. In fact, a strategy profile in an extensive-form game with perfect recall is an ASE if and only if it is a sequential equilibrium with *every* fully consistent belief system. However, ASE is not defined in terms of full consistency (which uses perturbations of strategies) but is based on a simpler, local concept of *strong consistency* between strategy profiles and off-equilibrium beliefs, which is applicable to a large class of dynamic games, including games with a continuum of actions. ASE is generalized by the set-valued solution concept of agnostic sequential polyequilibrium, which allows leaving the players' actions unspecified in some (possibly, many) information sets.

Keywords: Agnostic sequential equilibrium, agnostic sequential polyequilibrium, strong consistency, perfect Bayesian equilibrium.

1 Introduction

In dynamic games with perfect information, the idea that the players' choices of strategies should be rational also off-equilibrium, that is, at decision nodes that are not actually reached, is captured by the notion of subgame perfect equilibrium. This refinement of Nash equilibrium excludes, in particular, non-credible threats: actions that are detrimental to the actor, and a rational player would therefore not carry out. In games with imperfect information, an action taken at a particular information set U may be beneficial or detrimental depending on the node at which it is taken. However, by definition of an information set, the player does not know which of the nodes in U was reached. Therefore, he can assess the wisdom of choosing a particular action only with respect to particular beliefs about the history of play, which are expressed by a particular probability distribution over the nodes in U . The problem, of course, is that the prior probability of every sequence of actions leading to an off-equilibrium information set is zero, so that the beliefs in question are, in a sense, conditional on an event (namely, reaching U) that should not have happened if the players adhered to their equilibrium strategies.

The standard solution to this problem is to provide the beliefs as part of the solution concept. Thus, a solution is a pair (μ, x) , called an *assessment*, where x is a strategy profile and μ is a *belief system* that specifies a probability distribution over the nodes in each information set U . The probability assigned to a set of nodes $V \subseteq U$ is denoted $\mu(V)$ (and $\mu(U) = 1$). This approach is used by the sequential equilibrium solution concept (Kreps and Wilson 1982) as well as by the more general, but somewhat nebulous, notion of perfect

* Department of Economics, Bar-Ilan University, Ramat Gan 5290002, Israel
igal.milchtaich@biu.ac.il <http://faculty.biu.ac.il/~milchti>




Figure 1. Non-agnostic sequential equilibria (black lines). **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 strategy P is justified only if $1/3 \leq p \leq 2/3$. **c** The choices of players 2 and 3 are justified only if $p \leq 1/3 \leq q$. (In this game, the players have identical payoffs.)

Bayesian equilibrium, which originated with Fudenberg and Tirole (1991). (The terminology may be naturally extended by referring to a *strategy profile* x as a sequential or perfect Bayesian equilibrium when there is *some* belief system μ such that the assessment (μ, x) is so.) These and similar solution concepts require the strategies in x to be *sequentially rational* under μ , in the sense that the strategy x_i of each player i is optimal in the continuation game starting at each of the player’s information sets U , given the other players’ strategies and i ’s beliefs in U about the events that preceded his arrival there, which are specified by μ . They also require these hypotheses, or conjectures, about the history of play to be reasonable. 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 when off-equilibrium information sets are concerned, which means that beliefs are usually effectively *chosen* much like strategies are chosen. However, unlike the choice of strategies, which needs to be justified in the sense of satisfying sequential rationality, the beliefs may be largely arbitrary. This is a divergence from the common view that equilibrium represents a self-enforcing convention, as the beliefs in off-equilibrium information sets are neither self-enforcing nor can they be described as a convention, because these sets are actually never reached in equilibrium.

Consider the equilibrium shown in Figure 1a. This is a sequential (and perfect Bayesian) equilibrium, because player 2’s choice of r rather than l is justified by *some* belief, namely, that if player 1 switched from his equilibrium strategy of *Out* to another strategy, it was R rather than L . However, all alternative beliefs of player 2 at his (off-equilibrium) information set U are arguably as reasonable as this one. But if these beliefs, in particular, a belief that 1 played L , are not excluded, then player 2’s choice of l also cannot be excluded, which calls

into question player 1's choice of *Out*. A similar problem does not occur in sequential equilibria where U is reached with positive probability, such as (L, l) or the one where both players' choices of actions are uniformly distributed.

The concept of agnostic sequential equilibrium proposed in this paper differs from sequential and perfect Bayesian equilibrium in that it does not allow specifying actions in off-equilibrium information sets that are justifiable only under particular arbitrary beliefs about the history of play. The choice of such actions (for example, r in Figure 1a) entails the exclusion of one or more alternative actions that under different but equally reasonable beliefs are actually better. The concept is based on the idea that an action may be excluded only if some other action in the same information set U gives the acting player the same or higher payoff under *all* beliefs in U that are consistent with the strategy profile. Pinning down the appropriate meaning of consistency in this context is key for implementing this idea. This is done in the next section.

Unless where otherwise stated, the discussion below assumes that the dynamic game under consideration is an *extensive form game*, that is, one that can be described by a finite game tree, possibly with chance nodes (where, without loss of generality, all outcomes are assumed to have positive probability). It also assumes perfect recall. Throughout this paper, 'strategy' may refer to either *pure* strategy, which prescribes a single action at each of the player's information sets, or *behavior* strategy, which prescribes a probability distribution over actions at each information set. The former is viewed as a special case of the latter, in which all probabilities are either 0 or 1. All the results in the paper hold for both kinds of strategies and any reference to (an unqualified) 'strategy' may be understood in either way. Another notational convention is that actions at different information sets are always tagged differently, so that an action's name uniquely identifies the information set at which it is playable.

2 Consistent Beliefs

A minimal consistency requirement for an assessment (μ, x) is that the probabilities specified by the belief system μ coincide with those derived from the strategy profile x using Bayes' rule (in other words, they are the conditional probabilities) whenever possible. That is, for every information set U ,

$$\mu(V) = \frac{\mathbb{P}_x(V)}{\mathbb{P}_x(U)} \text{ if } \mathbb{P}_x(U) > 0, \quad V \subseteq U, \quad (1)$$

where $\mathbb{P}_x(V)$ is the probability that, under x , one of the nodes in the subset V is reached. This requirement is referred to as *weak consistency* of the assessment (μ, x) (or of the belief system μ with the strategy profile x). A standard strengthening of (1), which may be dubbed *weak+ consistency* in the information set U , is the requirement that a similar condition holds with x replaced with the strategy profile $x \mid x'_i$ obtained from it by replacing x_i with x'_i , for every strategy x'_i of the player i acting at U . The additional requirement has a bite only if the information set U is *reachable* for i under x , in the sense that $\mathbb{P}_{x \mid x'_i}(U) > 0$ for some strategy x'_i . The reason the requirement is sensible 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 the nodes are reached only depend on the probabilities of the other players' actions. Thus, for all strategies x_i' and x_i'' of player i ,

$$\frac{\mathbb{P}_{x_i'}(V)}{\mathbb{P}_{x_i'}(U)} = \frac{\mathbb{P}_{x_i''}(V)}{\mathbb{P}_{x_i''}(U)} \quad \text{if } \mathbb{P}_{x_i'}(U), \mathbb{P}_{x_i''}(U) > 0, \quad V \subseteq U. \quad (2)$$

Solution concepts that involve a single, specific belief system μ usually impose on it also certain *internal consistency* requirements, which express the idea that beliefs at different off-equilibrium 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 condition of *full consistency* of the assessment (μ, x) (or of the belief system μ with the strategy profile x). The condition requires the assessment to be the (pointwise) limit of some sequence of weakly consistent assessments $(\mu^k, x^k)_{k=1}^{\infty}$ where each x^k is a *completely mixed* strategy profile, in the sense that it assigns positive probability to every action at every information set (which entails that μ^k is uniquely determined by the weak consistency requirement). A *sequential equilibrium* (Kreps and Wilson 1982) is a fully consistent assessment (μ, x) such that x is sequentially rational under μ .

Agnostic sequential equilibrium does not specify a single belief system, which renders the whole internal consistency requirement moot. This brings about a considerable simplification, since consistency is narrowed down to the "local" condition that beliefs at each information set U are reconcilable with the strategy profile x . If, under x , the probability that U is reached is positive, this local consistency requirement is simply the weak consistency condition expressed by (1). However, if the probability is zero, then weak consistency does not specify any beliefs at U . Nevertheless, the player i acting at U may actually *know* a great deal about the history of play there. In particular, 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 for the set \mathcal{A} of all actions a as above for which the probability specified by x is 0, the specification was evidently not followed; the effective probability is 1. Therefore, it only remains for player i to speculate about the other players' behavior at information sets that do not involve actions in \mathcal{A} . The simplest hypothesis is that they 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 1.¹ If, under $x^{\mathcal{A}}$, the probability $\mathbb{P}_{x^{\mathcal{A}}}(U)$ that U is reached is positive, this hypothesis yields a unique probability distribution over the

¹ Probability 1 here means that the action was actually taken, not that the acting player j *meant* it to be played for sure. Note that, technically, it is immaterial whether the probability assigned to any action in \mathcal{A} is 1 or any other positive number; both possibilities yield the same beliefs in U .

nodes in U , which arguably represents the only beliefs at that information set that are consistent with x . This probability distribution $\mu_U^{\mathcal{A}}$ is given by

$$\mu_U^{\mathcal{A}}(V) = \frac{\mathbb{P}_{x^{\mathcal{A}}}(V)}{\mathbb{P}_{x^{\mathcal{A}}}(U)}, \quad V \subseteq U. \quad (3)$$

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 is 0. A natural approach in this case is to enlarge the set \mathcal{A} by adding to it one or more such precluded actions a , in such a way that \mathcal{A} becomes *minimally sufficient* for reaching the information set 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 $\mathcal{B} \subsetneq \mathcal{A}$. Any probability distribution $\mu_U^{\mathcal{A}}$ over the nodes of U that corresponds (by way of (3)) to a minimally sufficient set \mathcal{A} will be said to be *strongly consistent* with the strategy profile x . Note that if x is a pure-strategy profile and there are no chance moves in the game then such a distribution is necessarily *degenerate*: it assigns probability 1 to some node in U .

Strong consistency implies weak consistency. If an information set U is reached with positive probability under a strategy profile x , then the only minimally sufficient set for reaching U is $\mathcal{A} = \emptyset$, for which $x^{\mathcal{A}} = x$. Strong consistency moreover implies weak+ consistency. If U is reachable for the player i acting there, then there is a unique set of actions \mathcal{A} that are minimally sufficient for reaching U and the corresponding probability distribution $\mu_U^{\mathcal{A}}$ is weakly+ consistent. The set \mathcal{A} consists of all actions of player i that are assigned probability zero by x_i and precede (all the nodes in) U , so that $x^{\mathcal{A}}$ differs from x (if at all) only in the strategy of player i . In general, however, there may be more than one set of actions that are minimally sufficient for reaching an information set U , so strongly consistent beliefs are generally not unique. The beliefs corresponding to different sets represent alternative hypotheses about the actions that preceded the arrival at U . The next lemma (specifically, part (ii)) shows that these beliefs are mutually singular. It also establishes the following important relations between strongly and fully consistent beliefs in an information set U :

$$\mathfrak{B}^S \subseteq \mathfrak{B}^F \subseteq \text{conv } \mathfrak{B}^S, \quad (4)$$

where \mathfrak{B}^S is the set of all strongly consistent beliefs, $\text{conv } \mathfrak{B}^S$ is the convex hull of this set and \mathfrak{B}^F is the set of all beliefs arising from belief systems that are fully consistent with x . One corollary of (4) is that a sufficient condition for \mathfrak{B}^S and \mathfrak{B}^F to coincide is that (i) either set is a singleton or (ii) all elements of \mathfrak{B}^F are degenerate. Another corollary is that $\mathfrak{B}^F = \text{conv } \mathfrak{B}^S$ if and only if \mathfrak{B}^F is a convex set. In general, however, both inclusions in (4) may be proper.²

² Consider, for example, a game where player 5 acts after players 1 through 4 choose L or R , but he only knows how many of them chose each action. Suppose that x specifies that everyone chooses R , and consider the off-path information set U where player 5 is informed that only two of his predecessors actually did so. The set \mathfrak{B}^S consists of the six degenerate distributions: those that assign the value 1 to the probability p_{ij} that players i and j played L , for some $1 \leq i < j \leq 4$. Therefore, $\text{conv } \mathfrak{B}^S$ includes all beliefs in U . \mathfrak{B}^F differs from both sets, since (i) it includes the uniform

Lemma 1. Let \mathfrak{A} be the collection of all sets of actions that are minimally sufficient for reaching a given information set U under a given strategy profile x . For each $\mathcal{A} \in \mathfrak{A}$, let $\mu_U^{\mathcal{A}}$ be the probability distribution on U induced by $x^{\mathcal{A}}$, which is given by (3). Then:

- (i) Each $\mathcal{A} \in \mathfrak{A}$ consists of all actions that are assigned probability 0 by x and precede a particular node $u \in U$.
- (ii) The probability distributions $\{\mu_U^{\mathcal{A}}\}_{\mathcal{A} \in \mathfrak{A}}$ have pairwise disjoint supports.
- (iii) Each of them $\mu_U^{\mathcal{A}}$ coincides with the probability distribution on U specified by some belief system μ that is fully consistent with x , that is,

$$\mu_U^{\mathcal{A}}(V) = \mu(V), \quad V \subseteq U. \quad (5)$$

- (iv) 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 $\{\mu_U^{\mathcal{A}}\}_{\mathcal{A} \in \mathfrak{A}}$, that is,

$$\mu(V) = \sum_{\mathcal{A} \in \mathfrak{A}} \lambda^{\mathcal{A}} \mu_U^{\mathcal{A}}(V), \quad V \subseteq U \quad (6)$$

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

Proof. Consider any $\mathcal{A} \in \mathfrak{A}$, with cardinality $|\mathcal{A}| (\geq 0)$, and any node $u \in U$ with $\mathbb{P}_{x^{\mathcal{A}}}(\{u\}) > 0$.

Clearly, \mathcal{A} must include all the actions that precede u and are assigned probability 0 by x . By the minimal sufficiency condition, it cannot include any other actions. This proves (i). To prove (ii), it has to be shown that every $\mathcal{A}' \neq \mathcal{A}$ in \mathfrak{A} satisfies $\mathbb{P}_{x^{\mathcal{A}'}}(\{u\}) = 0$. For this, it suffices to note that the actions in $\mathcal{A} \setminus \mathcal{A}'$ are assigned probability zero by $x^{\mathcal{A}'}$.

To prove (iii), let z 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 :

$$x^\epsilon(a) = (1 - \epsilon - \epsilon^{|\mathcal{A}|+1})x(a) + \epsilon x^{\mathcal{A}}(a) + \epsilon^{|\mathcal{A}|+1}z(a), \quad (7)$$

where $x(a)$, $x^{\mathcal{A}}(a)$ and $z(a)$ are the probabilities specified by x , $x^{\mathcal{A}}$ and z . In this context, it is useful to refer also to possible outcomes of change moves as “actions”, whose (positive, by assumption) probabilities are fixed and are the same in all strategy profiles. The unique belief system μ^ϵ that is weakly consistent with the completely mixed strategy profile x^ϵ satisfies

$$\mu^\epsilon(V) = \frac{\mathbb{P}_{x^\epsilon}(V)}{\mathbb{P}_{x^\epsilon}(U)}, \quad V \subseteq U. \quad (8)$$

For every $v \in U$, $\mathbb{P}_{x^\epsilon}(\{v\}) = \prod_k x^\epsilon(a^k)$, where the a^k 's are all the actions preceding node v . In view of (7), this product can be expressed as a polynomial in ϵ , $\beta_0 + \beta_1\epsilon + \beta_2\epsilon^2 + \dots$.

distribution on the nodes in U , and (ii) all its elements satisfy $p_{12}p_{34} = p_{13}p_{24} = p_{14}p_{23}$, because similar equalities hold under every completely mixed strategy profile.

For $l < |\mathcal{A}|$, the coefficient β_l is zero, because a positive coefficient would mean that $\mathbb{P}_{x^{\mathcal{B}}}(\{v\}) > 0$ for some $\mathcal{B} \subsetneq \mathcal{A}$, which contradicts the minimal-sufficiency assumption concerning \mathcal{A} . By a similar argument, $\beta_{|\mathcal{A}|} = \mathbb{P}_{x^{\mathcal{A}}}(\{v\})$. It follows that, for $V \subseteq U$, $(1/\epsilon^{|\mathcal{A}|}) \mathbb{P}_{x^\epsilon}(V) \rightarrow \mathbb{P}_{x^{\mathcal{A}}}(V)$ as $\epsilon \rightarrow 0$, which implies that the quotient in (8) converges to that in (3). Therefore, if $(\epsilon_k)_{k=1}^\infty$ is any sequence of positive numbers converging to 0 such that $(\mu^{\epsilon_k})_{k=1}^\infty$ converges to some limit μ , then that belief system, which is clearly fully consistent with x , satisfies (5). The existence of such a sequence follows from the obvious compactness of the set of all belief systems.

To prove (iv), consider any belief system μ that is fully consistent with x , and some sequence $(\mu^k, x^k)_{k=1}^\infty$ as in the definition of full consistency. For every k , $\mathbb{P}_{x^k}(\{u\}) = \prod_l x^k(a^l)$, where the a^l 's are all the actions (and outcomes of chance moves) preceding u . Therefore,

$$\frac{1}{\prod_{a \in \mathcal{A}} x^k(a)} \mathbb{P}_{x^k}(\{u\}) = \prod_{\substack{l \\ a^l \notin \mathcal{A}}} x^k(a^l) \xrightarrow{k \rightarrow \infty} \mathbb{P}_{x^{\mathcal{A}}}(\{u\}). \quad (9)$$

A result similar to (9) holds with u replaced by any other node $v \in U$ that is preceded by all the actions in \mathcal{A} . Therefore, for such v ,

$$\frac{\mu_U^{\mathcal{A}}(\{v\})}{\mu_U^{\mathcal{A}}(\{u\})} = \frac{\mathbb{P}_{x^{\mathcal{A}}}(\{v\})}{\mathbb{P}_{x^{\mathcal{A}}}(\{u\})} = \lim_{k \rightarrow \infty} \frac{\mathbb{P}_{x^k}(\{v\})}{\mathbb{P}_{x^k}(\{u\})} = \lim_{k \rightarrow \infty} \frac{\mu^k(\{v\})}{\mu^k(\{u\})} = \frac{\mu(\{v\})}{\mu(\{u\})}$$

if $\mu(\{u\}) > 0$, and if $\mu(\{u\}) = 0$, then also $\mu(\{v\}) = 0$. Thus, for $v \in U$ that is preceded by all the actions in \mathcal{A} , $\mu(\{v\}) = \lambda^{\mathcal{A}} \mu_U^{\mathcal{A}}(\{v\})$, where $\lambda^{\mathcal{A}} = \mu(\{u\})/\mu_U^{\mathcal{A}}(\{u\})$. For all $\mathcal{A}' \neq \mathcal{A}$ in \mathfrak{A} , $\mu_U^{\mathcal{A}'}(\{v\}) = 0$ (because $x^{\mathcal{A}'}$ assigns probability zero to the actions in $\mathcal{A} \setminus \mathcal{A}'$), so that the equality in (6) holds for $V = \{v\}$. To prove that the equality holds generally, it remains to note that every $v \in U$ is preceded by all the actions in *some* element of \mathfrak{A} , because the set of actions that precede v and are assigned probability zero by x necessarily has a subset that is minimally sufficient for reaching U under x . Setting $V = U$ in (6) proves that the coefficients sum up to 1. ■

2.1 Structural Consistency

Reaching an information set U may also be explainable by a deviation from the players' strategy profile x that involves a non-minimally sufficient set of actions. However, such an explanation represents a non-parsimonious hypothesis as to why U was reached; it assumes more than it has to. Moreover, the explanation may have the troubling aspect that it implies a *future* deviation from x . This can happen if some players' information sets include both nodes that precede U and nodes that follow it.

Example 1 (Kreps and Ramey 1987). In the game in Figure 1b, the players' order of moves is uncertain – player 2 moves after either player 1 or 3 moves – and is unknown to them. It is not difficult to see that 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, in any equilibrium, the probabilities that player 2's strategy assigns to playing L and R are not greater than $1/3$, and in particular, neither of them is 1. Such a strategy can be justified only by beliefs that attach positive probability to both nodes in player 2's

information set, and such beliefs are induced only by strategy profiles in which 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.

Example 1 shows that a non-parsimonious hypothesis about the past may project onto the future. With a parsimonious hypothesis about the deviations from a strategy profile that led to an information set being reached, this cannot happen.

Proposition 1. If a set of actions \mathcal{A} is minimally sufficient for reaching an information set U under a strategy profile x , then $x^{\mathcal{A}}$ has the property that, (i) under $x^{\mathcal{A}}$, the probability that U is reached is positive, and (ii) under any strategy profile that either coincides with $x^{\mathcal{A}}$ or differs from it only in the strategy of the player acting at U , the probability that some information set where x and $x^{\mathcal{A}}$ disagree is reached after U is reached is zero.

Proof. Part (i) holds by definition. Consider any strategy profile as in (ii). Any path that has positive probability under it and reaches U must, by the minimal-sufficiency assumption, first go through *all* the actions in \mathcal{A} . By the perfect-recall assumption, the path cannot revisit any of the information sets where these actions are taken. ■

Kreps and Ramey (1987) call an assessment (μ, x) *structurally consistent*³ if the beliefs in every information set U are weakly+ consistent and coincide with those obtained using Bayes' rule (as in (1)) when x is replaced with some strategy profile x' that has the property specified for $x^{\mathcal{A}}$ in Proposition 1. The proposition thus shows that strongly consistent beliefs satisfy structural consistency. Kreps and Ramey (1987) call an assessment *convex structurally consistent* if the beliefs in every information set U are convex combinations of the structurally consistent ones (with weights that may depend on U). Their Proposition asserts that every fully consistent assessment (hence every sequential equilibrium) satisfies convex structural consistency. Proposition 1 and part (iv) of Lemma 1 ($\mathfrak{B}^F \subseteq \text{conv } \mathfrak{B}^S$) strengthen this result. They prove that, for a fully consistent assessment, the beliefs at every information set can in fact be presented as convex combinations of the strongly consistent beliefs.

The beliefs that justify a sequential equilibrium strategy for player 2 in Example 1 are not structurally consistent. They can be obtained only as convex combinations of the two structurally (and strongly) consistent beliefs, 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 example thus demonstrates the unavoidability of dealing with structurally inconsistent beliefs in a sequential (and perfect Bayesian) equilibrium. Agnostic sequential equilibrium, by contrast, involves only structurally consistent beliefs.

³ Note that theirs is a stronger concept than Kreps and Wilson's (1982) notion of structural consistency of *beliefs*, for which the strategy profile is irrelevant.

3 Agnostic Sequential Equilibrium

The discussion in the previous sections leads to the following formal definition, where strongly consistent beliefs are expressed as appropriately modified strategy profiles. 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 the strategy profile that differs from x only in that, at U and all the information sets that follow it, player i uses strategy x'_i .

Definition 1. A strategy profile x is an *agnostic sequential equilibrium (ASE)* if for every player i and strategy x'_i of that player the inequality

$$u_i(x^{\mathcal{A}}) \geq u_i(x^{\mathcal{A}} \mid_U x'_i) \quad (10)$$

holds for every information set U of player i and every set of actions \mathcal{A} that is minimally sufficient for reaching U under x .

If the notion of strong consistency of beliefs and a strategy profile implicit in Definition 1 were replaced with an even stronger, less inclusive kind of consistency, the result would be a *weaker* definition of agnostic sequential equilibrium. That is, the set of qualifying strategy profiles would expand or remain unchanged. The opposite is true for any weaker notion, such as weak or weak+ consistency. By the first inclusion in (4), $\mathfrak{B}^S \subseteq \mathfrak{B}^F$ ((iii) in Lemma 1), the notion of local consistency corresponding to full consistency is also weaker than strong consistency. However, replacing the latter with the former would actually not change the meaning of ASE. This conclusion follows from the second inclusion in (4), $\mathfrak{B}^F \subseteq \text{conv } \mathfrak{B}^S$ ((iv) in Lemma 1), because if each element in a particular set of beliefs at an information set U justifies the choice (or the exclusion) of a particular strategy there, then any convex combination of these beliefs automatically justifies it too. The conclusion shows that, fundamentally, sequential equilibrium and ASE do not differ because they employ different notions of consistency. The difference between the two solution concepts is wholly due to the logical difference between requiring sequential rationality with respect to *some* consistent beliefs or *all* such beliefs.

Theorem 1. A strategy profile x is an agnostic sequential equilibrium if and only if (μ, x) is a sequential equilibrium for *every* belief system μ that is fully consistent with x .

Proof. It follows from (iii) in Lemma 1 that if (μ, x) is a sequential equilibrium for every belief system μ that is fully consistent with x , then the condition in Definition 1 holds. It follows from (iv) in that lemma that if the second condition holds, then the first one holds. ■

Example 2. The sequential equilibrium in Figure 1c is supported by the fully consistent belief system where $p = q = 1/3$, which represents the common belief by players 2 and 3 that, if player 1 deviated from his equilibrium strategy of *Out*, it is twice as likely that he played *R* rather than *L*. However, any other common belief, which assigns other probabilities to *L* and *R*, is also fully consistent with the strategy profile in that figure. Therefore, the latter is not an agnostic sequential equilibrium. By contrast, a strategy profile where player 1's behavior strategy is changed to one that does actually specify choosing *L* with positive probability and *R* with twice that probability is an ASE, as for it the above beliefs constitute the unique fully consistent belief system (even though player 3's information set is not reached).

3.1 One-Deviation Property

Generally speaking, a solution concept for dynamic games satisfies *perfectness* if the strategy x_i prescribed to each player i is optimal not only in the whole game but also in the continuation game starting at each of the player's information sets U . The solution concept has the *one-deviation property* if it is sufficient to check deviations that disagree with x_i only at U (rather than there and/or at some later information sets), because non-existence of profitable deviations of this special kind automatically implies the same for any deviation. Sequential equilibrium and similar solution concepts have the one-deviation property (Hendon et al. 1996, Perea 2002). The same is true for agnostic sequential equilibrium.

Theorem 2. A strategy profile x is an agnostic sequential equilibrium if and only if for every player i and action a of that player, playable at an information set U , the inequality

$$u_i(x^{\mathcal{A}}) \geq u_i(x^{\mathcal{A} \cup \{a\}}) \quad (11)$$

holds for every set of actions \mathcal{A} that is minimally sufficient for reaching U under x .

Proof. The necessity of the condition is obvious, as (11) is a special case of (10), with x'_i that differs from x_i only in specifying that i takes the action a at U (with probability 1). To prove sufficiency, suppose that x is not an agnostic sequential equilibrium, so that $u_i(x^{\mathcal{A}} \mid_U x'_i) > u_i(x^{\mathcal{A}})$ for some i , U , x'_i and \mathcal{A} as in Definition 1. Without loss of generality, x'_i can be assumed to be a pure strategy, so the left-hand side of the above inequality can be written as $u_i(x^{\mathcal{A} \cup \mathcal{B}'})$, where the elements of the set \mathcal{B}' are actions of player i at either U or some later information set. Also without loss of generality, \mathcal{B}' is minimal, that is, $u_i(x^{\mathcal{A} \cup \mathcal{B}'}) \leq u_i(x^{\mathcal{A}})$ for all $\mathcal{B}' \subsetneq \mathcal{B}'$. In particular, for every $a \in \mathcal{B}'$ that is playable at an information set U' which does not precede any other information set where an element of \mathcal{B}' is playable, $u_i(x^{\mathcal{A} \cup \mathcal{B}'}) \leq u_i(x^{\mathcal{A}}) < u_i(x^{\mathcal{A} \cup \mathcal{B}'})$ holds for $\mathcal{B}' = \mathcal{B}' \setminus \{a\}$. The inequalities imply that U' is reached with positive probability under $x^{\mathcal{A} \cup \mathcal{B}'}$ and that deviating by playing a there increases player i 's payoff in the continuation game (with the beliefs in U' given by Bayes' rule under $x^{\mathcal{A} \cup \mathcal{B}'}$). The same is true with $\mathcal{A} \cup \mathcal{B}'$ replaced with any set $\mathcal{A}' \subseteq \mathcal{A}' \subseteq \mathcal{A} \cup \mathcal{B}'$ that is minimally sufficient for reaching U' under x (such a set necessarily exists, because U precedes U' and therefore any subset of $\mathcal{A} \cup \mathcal{B}'$ that is minimally sufficient for reaching U' must contain \mathcal{A}), as changing from $x^{\mathcal{A} \cup \mathcal{B}'}$ to $x^{\mathcal{A}'}$ does not change the beliefs in U' (see (2)). Thus, $u_i(x^{\mathcal{A}'}) < u_i(x^{\mathcal{A}' \cup \{a\}})$, which shows that the condition in the theorem does not hold. ■

3.2 Beyond Finite Trees

The formal setting of the above discussion is that of extensive-form games. However, the actual definition of agnostic sequential equilibrium is also applicable to perfect-recall dynamic games that cannot be described by a finite game tree, in particular, games with a continuum of actions. (If the characterization of ASE in Theorem 2 is taken as a *definition*, then the assumption of perfect recall can also be dispensed with.) In this, ASE is more similar to perfect Bayesian equilibrium than to sequential equilibrium.

Definition 1 is formally meaningful even if minimally sufficient sets do not necessarily exist, which can happen with a continuum of outcomes to chance moves. However, in games with a continuum of actions, agnostic sequential rationality may run into conceptual problems

similar to those plaguing more traditional forms of sequential rationality (Myerson and Reny 2015). These problems stem from the fact that reaching a zero-probability information set may actually be expected, and does not necessarily indicate a deviation.

Example 3 (Myerson and Reny 2015). Player 1 can play L or R . If he chooses L , then there is a chance move in which a number is drawn uniformly from the unit interval $[0,1]$, and if he chooses R , then the player himself has to choose a number in $[0,1]$. Player 2 is told the number s selected, but not whether it resulted from a chance move or was chosen by player 1. His possible actions are also L and R , and the payoffs are as in the battle of the sexes game. In particular, only (L, L) and (R, R) yield positive payoffs. In BoS, both strategy profiles are equilibria. Here, however, (R, R) must be played with probability 1 in any agnostic sequential equilibrium. To see this, consider for $s \in [0,1]$ the information set of player 2 where he is informed that s was selected. The probability that a chance move yields s is 0. Therefore, regardless of player 1's actual strategy, there is a unique set of actions that is minimally sufficient for reaching the information set, which corresponds to the belief that player 1 played R and then chose s . Player 2's should therefore play R regardless of s , and player 1's best response to this is also playing R . Therefore, the strategy profile where player 1 plays L and player 2 plays L regardless of s is not an ASE, even though the actions are best response to one another and, intuitively, as all outcomes of the chance move are equally likely, no value of s should raise a suspicion of deviation by player 1.

4 Beliefs Based on Strategic Reasoning

Strong consistency 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, *forward induction* arguments (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 form of the game tree, forward induction also involves examination of the strategic interests of the deviating players. For example, in the agnostic sequential equilibrium shown in Figure 2a, player 2's choice of r is supported by the unique strongly (and fully) consistent belief, which is that player 1 would follow his strategy in the proper 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 l , player 2 may reason that the most likely explanation for the deviation is that player 1 is aiming for the better equilibrium in the game, with payoffs 1, which means that he also played L rather than R . 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 2a 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 the one considered above, so that the same argument applies: if player 1 deviated once, then strategic considerations suggest he intends to deviate again.

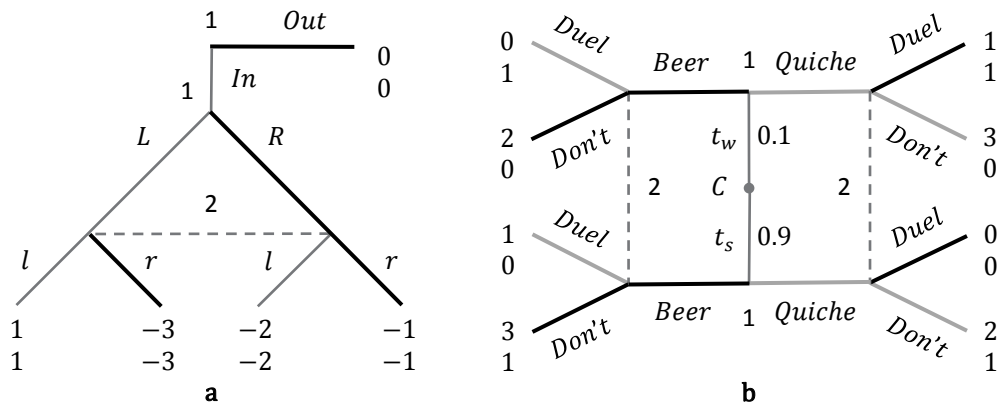


Figure 2. The destructive and constructive potential of beliefs reflecting strategic reasoning. **a** In the agnostic sequential equilibrium shown, player 2's action is supported by the unique beliefs consistent with player 1's strategy in the proper subgame. However, a deviation by player 1 from his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

A different potential outcome of strategic reasoning is selection among beliefs that *are* strongly consistent 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 sequential or perfect Bayesian equilibria, restriction to reasonable consistent beliefs may eliminate some of the equilibria. For example, it may eliminate all pooling equilibria in the Spence education model with two types of worker (Cho and Kreps 1987). The effect of a similar restriction on beliefs on agnostic sequential equilibria is in a sense the diametric opposite. In Spence's model, in particular, none of the pooling equilibria is an ASE to begin with. The reason is that the choice of any off-equilibrium education level cannot be excluded, because the requirement of strong consistency does not preclude a belief by the employer that such a choice indicates a high quality worker.⁴ Thus, as for sequential and perfect Bayesian 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 ASEs but can only *add* new ones (see the remark that follows Definition 1). Thus, the set of agnostic sequential equilibria, which is typically contained in the sets of perfect Bayesian and sequential equilibria (Theorem 1), 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 4. Consider the Beer-Quiche game shown in Figure 2b, where for simplicity only pure strategies are allowed. There are two (Nash) equilibria, which are both pooling. In one

⁴ More generally, in an agnostic sequential equilibrium a player would not know what to make of a signal that the sender's strategy never specifies: such a signal might mean anything. In this, ASE differs from perfect Bayesian equilibrium, which requires a ready interpretation for every physically possible signal.

equilibrium, Black (indicated by black lines), types t_w and t_s of player 1 both choose *Beer*, and in the other, Gray, 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, where both types of player 1 choose *Beer* and player's 2 would choose *Duel* only as a response to *Quiche*. This 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). The same argument also shows that a restriction to reasonable beliefs would make Black an agnostic sequential equilibrium. Thus, the logic underlying the intuitive criterion singles out the same equilibrium for both solution concepts, sequential and agnostic sequential equilibrium. This coincidence contrasts with the situation for the original, unmodified definitions, according to which both equilibria are sequential equilibria but neither of them is an ASE. 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.

5 Polyequilibrium

In an off-path information set, where there may be multiple strongly consistent beliefs, there may also be no single action that is the best choice under all beliefs. This fact raises the possibility that the action at such an information set may have to be left unspecified. Doing so leads to a set-valued solution concept. Instead of a single strategy profile x , the solution is a set of strategy profiles X .

An example of such a solution concept is *essentially perfect Bayesian equilibrium* (Blume and Heidhues 2006). In an incomplete-information game with perfect recall, an EPBE is a set X that is obtained from a specified strategy profile x by declaring certain information sets *relevant* and the remaining ones *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 information set of the same player. X is then defined as the collection of all strategy profiles that agree with x at each of relevant information sets. By requirement (i), any belief system μ that is weakly consistent with x is also weakly consistent with every other element of X . The definition of EPBE is completed by the 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 *all* strategy profiles in X in the continuation game starting at U with the distribution over U 's nodes specified by μ .

Essentially perfect Bayesian equilibrium extends perfect Bayesian equilibrium in being a set-valued solution concept. However, it is still based on a single, possibly arbitrary, belief system. Thus, it is not in the spirit of, and does not extend, agnostic sequential equilibrium. To extend the latter, a more general set-valued solution concept is needed. Such a concept is

polyequilibrium (Milchtaich 2018). (EPBE is a special kind of polyequilibrium, indeed of *simple* polyequilibrium.)

For a player i , with strategy set S_i and payoff function u_i , in either a simultaneous-move or a dynamic game, a *polystrategy* is any nonempty set of strategies, $\emptyset \neq X_i \subseteq S_i$. A *polystrategy profile* X is a Cartesian product of polystrategies, one polystrategy X_i for each player i . In other words, it is a nonempty rectangular set of strategy profiles. X is a *polyequilibrium* if for every player i and strategy $x'_i \notin X_i$ there is some $x''_i \in X_i$ that responds to X at least as well as x'_i does, in the sense that

$$u_i(x \mid x''_i) \geq u_i(x \mid x'_i), \quad x \in X.$$

Thus, the polyequilibrium condition is that for every strategy x'_i excluded by player i 's polystrategy X_i there is a non-excluded strategy x''_i that is an adequate substitute against *all* strategy profiles in X . In a dynamic context, the condition may be naturally strengthened by requiring x''_i to be an adequate substitute also in the continuation game starting at each of player i 's information sets U , for every strategy profile $x \in X$ and every probability distribution over the nodes in U that is strongly consistent with x . This idea leads to the following definition.

Definition 2. A polystrategy profile X is an *agnostic sequential polyequilibrium (ASPE)* if for every player i and strategy $x'_i \notin X_i$ there is some $x''_i \in X_i$ such that the inequality

$$u_i(x^{\mathcal{A}} \mid_U x''_i) \geq u_i(x^{\mathcal{A}} \mid_U x'_i) \tag{12}$$

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

Polyequilibrium is an “excluding” solution concept. It requires justification for the exclusion of the strategies *not* in a player’s polystrategy rather than the inclusion of those in it. The (dynamic-game) refinement of polyequilibrium presented here extends this idea to beliefs. Non-consistent beliefs are (implicitly) excluded, but this should not be interpreted as an assertion that all the remaining, consistent ones are necessarily justifiable (see Section 4).

Agnostic sequential polyequilibrium is a generalization of agnostic sequential equilibrium. A strategy profile x is an ASE if and only if the singleton $\{x\}$ is an ASPE. Of these two solution concepts, the latter should be viewed as the principal one and the former as a mere appendage. For ASPE, unlike ASE, existence is not an issue. Indeed, the set of all strategy profile is an agnostic sequential polyequilibrium, called the *trivial* ASPE.

For a polyequilibrium that includes more than one strategy profile, the question of the predictive content of the polyequilibrium arises. The answer is the concept of *polyequilibrium result* (Milchtaich 2018). A *result* R is any set of strategy profiles in the game. It *holds* in a polystrategy profile X if $X \subseteq R$, and it is a *polyequilibrium result* if it holds in some polyequilibrium in the game. A result may also be specified implicitly, as a particular property or consequence of strategy profiles (for example, “player 1’s payoff is higher than 2’s payoff”). In this case, R is the collection of all strategy profiles that have the specified property, so that the result holds in a polyequilibrium X if and only if *all* strategy profiles in X

have the property. In particular, a real number v_i is a *polyequilibrium payoff* for a player i if there is some polyequilibrium X with $u_i(x) = v_i$ for all $x \in X$, and a strategy x_i is a *polyequilibrium strategy* if there is some polyequilibrium X with $X_i = \{x_i\}$. The concept of result may also be applied to special kinds of polyequilibria, and in particular to equilibria (or special kinds thereof) and to ASPEs.

Example 5. There is a \$2 sum in either the blue or the yellow box, and the two possibilities, B and Y , are equally likely. Player 2 has to choose one of the boxes, open it and take the money if it is there. Player 1 is not allowed to take the money, but he knows where it is hidden. He has to choose the price $0 \leq p \leq 1$ for which he offers to sell that information to player 2, who can pay the price and pick up the money, reject the offer and open the blue box, or reject it and open the yellow box. For simplicity, only pure strategies are allowed. Every v_1 between zero and \$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 will reject the offer and open the blue box if player 1 asks a different price. This reaction is supported by a belief that a price different from v_1 indicates that the money is in the blue box. By contrast, the only agnostic sequential polyequilibrium payoff for player 1 is \$1. To see this, suppose that there is an ASPE X where the payoff is $v_1 < 1$, and consider some strategy $x_1 \in X_1$ and some price $v_1 < p < 1$ that is different from the two prices that x_1 specifies for the two cases B and Y (which need not be the same). Player 2's polystrategy X_2 excludes acceptance of price p , for if it included a strategy prescribing acceptance, it would not be possible to exclude player 1's strategy of asking p . If player 1 uses x_1 , player 2's information set where he is asked to pay p is not reached. However, at that information set, there is no action of player 2 that under all strongly consistent beliefs does at least as well as the excluded action of acceptance. The alternative of rejecting the offer and opening the blue box, say, is worse than accepting the offer under the belief that, if player 1 asks p , then the money is in the yellow box. (This belief is strongly consistent with x_1 , and reflects an hypothesis that player 1 deviated by asking p only in case Y .) This conclusion contradicts the assumption that X is an agnostic sequential polyequilibrium. The contradiction leaves \$1 as the only possible ASPE payoff. This payoff is obtained in the ASPE where player 1 asks this price and player 2's polystrategy is to accept it. Thus, the reaction to any lower price is left unspecified.

References

- Blume, A., Heidhues, P. (2006). Private monitoring in auctions. *Journal of Economic Theory* 131, 179–211.
- Cho, I.-K., Kreps, D.M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics* 102, 179–221.
- Fudenberg, D., Tirole, J. (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory* 53, 236–260.
- Hendon, E., Jacobsen, J., Sloth, B. (1996). The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior* 12, 274–282.

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

Kreps, D.M., Ramey, G. (1987). Structural consistency, consistency, and sequential rationality. *Econometrica* 55, 1331–1348.

Kreps, D.M., Wilson, R. (1982). Sequential equilibria. *Econometrica* 50, 863–894.

Milchtaich, I. (2018). **Polyequilibrium**. *Games and Economic Behavior*
<https://doi.org/10.1016/j.geb.2018.09.013>

Myerson, R.B., Reny, P.J. (2015). Sequential equilibria of multi-stage games with infinite sets of types and actions. Mimeo.

Perea, A. (2002). A note on the one-deviation property in extensive form games. *Games and Economic Behavior* 40, 322–338.