Discussion Paper Series – CRC TR 224

Discussion Paper No. 748
Project B 02

Perfect Bayesian Equilibria and Cross-Pair Independence from Common Actions

Francesc Dilmé[1]

May 2026

[1] University of Bonn. Email: fdilme@uni-bonn.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Perfect Bayesian Equilibria and Cross-Pair Independence from Common Actions

Francesc Dilmé*

May 8, 2026

Abstract

This paper develops a simple notion of perfect Bayesian equilibrium for arbitrary finite extensive-form games with perfect recall. Its key ingredient is *cross-pair independence from common actions*, a belief restriction that compares the likelihoods of two pairs of histories, possibly in different information sets, after canceling actions common to both pairs. The condition extends *Bayes' rule whenever possible* and *no signaling what you don't know* beyond sequentially ordered information sets. In multi-stage games with observable actions and independent types, it implies the belief-reasonableness requirements of Fudenberg and Tirole (1991). Extending the same cancellation logic to arbitrary multisets is equivalent to consistency. The construction therefore provides a tractable, assessment-based notion of reasonableness without non-standard probabilities, plausibility orders, or auxiliary belief-revision structures.

Since Fudenberg and Tirole (1991) introduced perfect Bayesian equilibrium, it has become one of the most widely used equilibrium concepts in applied and theoretical work. Broadly speaking, an assessment is a perfect Bayesian equilibrium if it is sequentially rational and reasonable. Sequential rationality requires that, at each information set, the moving player's continuation strategy be optimal given her beliefs and the other players' strategies. Reasonableness concerns players' beliefs. One such requirement is that beliefs be updated according to Bayes' rule whenever the relevant conditioning event has positive probability, including at information sets reached only after deviations. Another crucial requirement rules out that actions signal information the acting player does not possess. Reasonableness also requires that different types of a player have the same beliefs about information that is independent of that player's type. These restrictions are automatically satisfied at on-path information sets and in fully consistent assessments (Kreps and Wilson, 1982), that is, assessments approximated by sequences of fully mixed strategy profiles.

*University of Bonn. fdilme@uni-bonn.de. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B02) and under Germany's Excellence Strategy (EXC 2126/1-390838866 and EXC 2047-390685813).

The original definition in Fudenberg and Tirole (1991) is tailored to multi-stage games with observable actions, independent types, and a single initial move by nature. For more general extensive-form environments, Fudenberg and Tirole also introduce generally reasonable extended assessments, which rely on a conditional probability system over terminal histories and include a version of no signaling what you don't know. Other generalizations of perfect Bayesian equilibrium likewise use auxiliary objects such as plausibility orders, AGM belief revision, or common conditional probability systems. While these approaches are often conceptually powerful, they can be difficult to apply directly in theoretical or applied work.

This paper proposes a simple assessment-based definition of perfect Bayesian equilibrium using an intuitive requirement, termed *cross-pair independence from common actions*, which compares the relative likelihoods of two pairs of histories in possibly different information sets. We show that this condition generalizes "Bayes' rule whenever possible" and "no signaling what you don't know" for arbitrary finite extensive-form games and, in multi-stage games with observable actions and independent types, implies the belief-reasonableness requirements of Fudenberg and Tirole (1991). We also identify the gap between perfect Bayesian equilibrium and sequential equilibrium: requiring independence from common actions across arbitrary multisets of histories, rather than just two pairs of histories, is equivalent to consistency.

To illustrate the basic idea, consider the game in panel (a) of Figure 1, where payoffs and player 4's moves are omitted because they are irrelevant for the discussion. In any assessment in which all actions are played with positive probability, standard Bayes' rule implies that $\alpha = \beta$, because player 2's strategy is the same at both nodes in her information set. Now consider an assessment in which T_1 and M_1 are played with probability zero, while T_2 is played with positive probability. Player 2's information set is now off path, but it is still natural to require that $\alpha = \beta$ whenever $\gamma > 0$. Bayes' rule whenever possible imposes exactly that: the relative likelihood of histories (T_1, T_2) and (M_1, T_2) at player 4's information set should equal the relative likelihood of histories T_1 and M_1 at player 2's information set.

If instead T_2 is played with probability zero, then player 4's information set is no longer reached with positive probability after player 2's information set, so Bayes' rule whenever possible is silent. Even so, if player 4 assigns positive probability to the event that play passed through player 2's information set (that is, if $\gamma > 0$), then player 2's off-path action should not convey information about whether player 1 previously chose T_1 or M_1, because player 2 did not know that fact when she moved. This is the role of the no-signaling-what-you-don't-know condition. In this game, it says that player 2's deviation should not reveal information unknown to player 2, so the relative likelihood of histories (T_1, T_2) and (M_1, T_2) at player 4's information set should again equal the relative likelihood of histories T_1 and M_1 at player 2's information set.

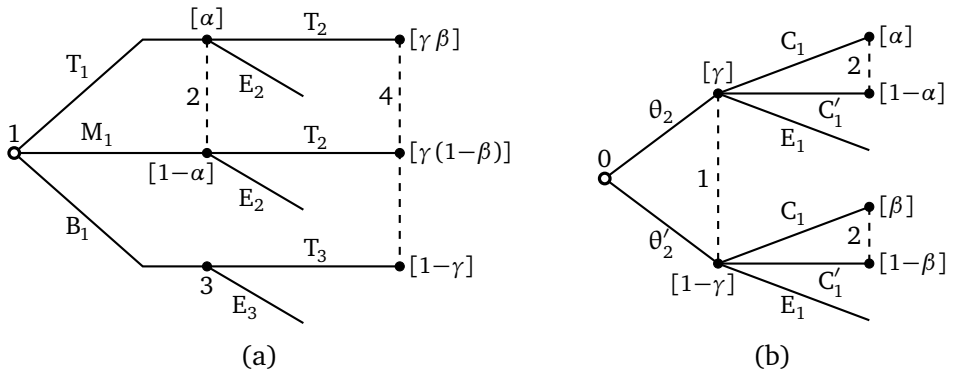


Figure 1

We define *simple perfect Bayesian equilibrium* as an assessment that is sequentially rational and satisfies Bayes' rule whenever possible and no signaling what you don't know. This concept is stronger than weak perfect Bayesian equilibrium (which only requires on-path Bayes' rule) and similarly straightforward to verify: it requires only a small number of local restrictions on strategies and beliefs at and between relevant information sets. This matters in applications because checking whether an assessment satisfies the full set of equilibrium conditions is often difficult.

Bayes' rule whenever possible and no signaling what you don't know only permit comparisons of the likelihoods of pairs of histories across information sets that are sequentially ordered; otherwise they are silent. This limitation prevents them from imposing some reasonable restrictions. For example, different types of a player who observe the same behavior of another player should hold the same beliefs about that player's unobserved action whenever the other player's information is independent of the observing player's type. To see this, consider the game in panel (b) of Figure 1. There, player 1 does not observe nature's move, which is interpreted as choosing player 2's type. Player 1 can exit or choose one of two continuation actions. Player 2 moves only if player 1 does not exit, and she observes nature's move but not player 1's continuation action. Even in an assessment in which player 1 exits with probability one, it is natural to require $\alpha = \beta$: player 2's belief about which action player 1 chose when deviating should not depend on player 2's type. Yet neither Bayes' rule whenever possible nor no signaling what you don't know imposes any belief restriction in this case.

This motivates the definition of *cross-pair independence (from common actions)*, which extends the logic behind Bayes' rule whenever possible and no signaling what you don't know beyond sequentially ordered information sets. Broadly speaking, the condition requires that if histories h and $\hat{h}$ lie in one information set and histories h' and $\hat{h}'$ lie in another, then comparing the product of the likelihoods of h and h' with that of $\hat{h}$ and $\hat{h}'$ should depend only on the actions that are not common to both sides. In panel (b) of Figure 1, this condition implies $\alpha = \beta$, because the actions

not common to histories (θ_2, C_1) and (θ_2, C_1') are the same as those not common to histories (θ_2', C_1) and (θ_2', C_1'). We show that cross-pair independence is stronger than Bayes' rule whenever possible and no signaling what you don't know, and that it implies the belief-reasonableness requirements in multi-stage games with observable actions and independent types. For that reason, we define a *perfect Bayesian equilibrium* as an assessment that is sequentially rational and satisfies cross-pair independence.

Finally, we study a stronger condition, simply termed *independence from common actions*, which requires that ratios formed from multisets of histories with matching information-set counts depend only on their non-common actions. This condition is equivalent to consistency. Nonetheless, it imposes infinitely many conditions, so verifying that a given assessment satisfies it is often difficult. We illustrate, however, that independence from common actions still provides an intuitive and powerful tool for ruling out consistency of assessments in practice.

Literature review. This paper studies belief restrictions for finite extensive-form games with perfect recall. Its benchmark is the relation between perfect Bayesian equilibrium and sequential equilibrium. Kreps and Wilson (1982) define consistency by requiring beliefs to be limits of Bayes' rule beliefs from fully mixed strategy profiles. This disciplines off-path beliefs through tremble sequences. Fudenberg and Tirole (1991) define perfect Bayesian equilibrium for finite multi-stage games with observable actions and independent types, combining sequential rationality with reasonableness restrictions such as common beliefs, belief independence, Bayes' rule whenever possible, and no signaling what you don't know. The question here is whether comparable discipline can be imposed directly on an assessment (σ, μ) in arbitrary finite extensive-form games, and how the resulting restrictions compare with consistency.

The closest related work falls into two strands. The first develops general perfect Bayesian equilibrium notions by enriching the belief system: Watson (2025) uses belief revision over strategies; Bonanno (2013, 2016) use plausibility orders and AGM-style consistency; and Merotto and Wolitzky (2026) uses a common conditional probability system on terminal nodes. These papers seek to retain some discipline of consistency without imposing sequential equilibrium, but do so with auxiliary belief-revision or conditional-probability objects rather than restrictions stated directly on assessments. A second strand characterizes consistency, including Kreps and Ramey (1987), Perea y Monsuwe et al. (1997), Kohlberg and Reny (1997), Pimienta (2014), and Dilmé (2023). The closest conceptual lineage within that strand is the conditional- and relative-probability approach of Myerson (1986), McLennan (1989), Battigalli (1996), and Kohlberg and Reny (1997); the difference is that the independence restriction is stated directly on histories after canceling common actions.

The paper's contribution is to provide a single cancellation principle that organizes the gap between perfect Bayesian equilibrium and sequential equilibrium. Bayes' rule is one implication of that principle when the conditioning event has positive probability, and Bayes' rule whenever possible and no signaling what you don't know are its local off-path implications when information sets are sequentially ordered. The new condition, cross-pair independence, applies the same logic pairwise across histories in arbitrary information sets by canceling actions common to both sides of the comparison. Our notion of perfect Bayesian equilibrium is then stated directly on assessments and capable of capturing cross-information-set restrictions missed by the local conditions. In multi-stage games with observable actions and independent types, Appendix B relates this condition to Fudenberg and Tirole's (1991) reasonableness requirements; with the additional cross-action strengthening introduced there, the formulations are equivalent. Requiring the cancellation logic for arbitrary multisets of histories rather than pairs then yields consistency itself.

The rest of the paper is organized as follows. Section 1 provides the basic definitions. Section 2 introduces Bayes' rule whenever possible and no signaling what you don't know, as well as simple perfect Bayesian equilibrium. Section 3 defines and analyzes cross-pair independence and independence from common actions. Appendix A provides the omitted proofs, and Appendix B compares the cross-pair independence condition with Fudenberg and Tirole's (1991) concept of reasonableness in multi-stage games with observable actions and independent types.

1 Basic definitions

1.1 Games in extensive form

We work with finite extensive-form games with perfect recall.

A (finite) *game* is a tuple $G := \langle A, H, \mathcal{I}, N, \iota, \pi, u \rangle$ with the following components:

1. A finite set of *actions* A.
2. A finite set of *histories* H. A history is a finite sequence of actions $h \equiv (h_j)_{j=1}^{|h|}$, where $|h|$ denotes the length of h. The set H is prefix-closed: if $h \in H$ and $|h| > 0$, then its predecessor $(h_j)_{j=1}^{|h|-1}$ also belongs to H. In particular, the empty history $\emptyset = (h_j)_{j=1}^{0}$ belongs to H. For histories $h, \hat{h} \in H$, we write $(h, \hat{h})$ for their concatenation whenever it belongs to H; in particular, (h, a) denotes the history obtained by appending action a to h. We write $h \preceq \hat{h}$ if h is a prefix of $\hat{h}$, $h \prec \hat{h}$ if $h \preceq \hat{h}$ and $h \neq \hat{h}$, and $\hat{h} \succ h$ if $h \prec \hat{h}$. We write $Z \subset H$ for the set of terminal histories, that is, those with no feasible continuation.
3. An *information partition* $\mathcal{I}$ of $H \setminus Z$, together with a partition $\{A^I \mid I \in \mathcal{I}\}$ of A, such that for every $I \in \mathcal{I}$, every $a \in A^I$, and every $h \in H$, we have $h \in I$ if and only if $(h, a) \in H$. The elements

of $\mathcal{I}$ are the *information sets*.[1]

4. A finite set of players N, with $0 \notin N$.
5. A *player assignment* $\iota : \mathcal{I} \to N \cup \{0\}$, which assigns each information set either to a player or to nature (represented by 0), where $|I| = 1$ whenever $\iota(I) = 0$. We assume perfect recall.[2]
6. A *strategy of nature* $\pi : \bigcup_{I \in \iota^{-1}(\{0\})} A^I \to (0,1]$ such that $\sum_{a \in A^I} \pi(a) = 1$ for each $I \in \iota^{-1}(\{0\})$.
7. For each player $i \in N$, a payoff function $u_i : Z \to \mathbb{R}$.

For each action $a \in A$, let I^a denote the unique information set at which a is available. A *behavior strategy* for player i is a map $\sigma_i : \bigcup_{I \in \iota^{-1}(\{i\})} A^I \to [0,1]$ such that $\sum_{a \in A^I} \sigma_i(a) = 1$ for every $I \in \iota^{-1}(\{i\})$. We write Σ_i for the set of player-i behavior strategies. A *strategy profile* is a collection $(\sigma_i)_{i \in N}$ together with nature's fixed strategy π; throughout the paper we identify a profile with the induced map $\sigma : A \to [0,1]$ satisfying $\sum_{a \in A^I} \sigma(a) = 1$ for all $I \in \mathcal{I}$ and $\sigma(a) = \pi(a)$ for all nature actions. We write Σ for the set of all such profiles. A strategy profile $\sigma \in \Sigma$ is *fully mixed* if $\sigma(a) > 0$ for every $a \in A$.

1.2 Histories as sets

We will often identify a history with the set of actions it contains. The next observation shows that, under our standing assumptions of perfect recall and unique availability of each action, this abuse of notation is harmless.

Lemma 1.1. *No two distinct histories have the same set of actions.*

Proof of Lemma 1.1. We first show that no information set can be visited twice along a history. Suppose there are two distinct prefixes $p \prec q$ of a single history that belong to the same information set I. Let c be the first action taken after p on the path from p to q. By the definition of p and q, we have $(p,c) \preceq q$. Apply perfect recall with $I = I' = I$, $h = q$, and $\hat{h} = p$. We obtain some $\tilde{p} \in I$ such that $(\tilde{p},c) \preceq p$, implying $\tilde{p} \prec p$. Repeating the same reasoning produces an infinite descending chain of histories in a finite game tree, which is impossible. Therefore an information set cannot be visited twice along any single history. Because each action is available at a unique information set, it follows immediately that no action can occur twice along a history.

---

[1] Without loss of generality, each action is assumed to be available at a unique information set; one can always relabel actions to ensure this.

[2] For all $I, I' \in \mathcal{I}$ with $\iota(I) = \iota(I')$ and all $h, \hat{h} \in I$, if $(h', a) \preceq h$ for some $h' \in I'$ and $a \in A$, then $(\hat{h}', a) \preceq \hat{h}$ for some $\hat{h}' \in I'$. This condition is imposed also when $\iota(I) = 0$, treating nature as player 0 for this formal purpose.

Now suppose two distinct histories h and $\hat{h}$ contain exactly the same set of actions. Let j be the first position at which they differ and set

$$p=(h_1,\dots,h_{j-1})=(\hat{h}_1,\dots,\hat{h}_{j-1}),\qquad a:=h_j.$$

Because actions cannot repeat along a history, the action a does not appear in the common prefix p. Since a belongs to both h and $\hat{h}$, it must occur somewhere in $\hat{h}$ after p. In particular, there exists $q\succ p$ on the path $\hat{h}$ such that the action a is available after both p and q. Since each action is available at a unique information set I^a, it follows that p and q both lie in I^a. But then the history $\hat{h}$ visits the same information set I^a twice, contradicting the preliminary argument. This contradiction shows that two distinct histories cannot contain the same set of actions. □

For any action set A' (for instance, a history or part of a history), define

$$\sigma(A'):=\prod_{a\in A'}\sigma(a),$$

with the convention $\sigma(\emptyset)=1$. In particular, $\sigma(h)$ denotes the probability with which history h is reached under σ. For histories $h,\hat{h}\in H$, the term $\sigma(h\setminus\hat{h})$ denotes the product of the probabilities of the actions in h that are not in $\hat{h}$.

Finally, for an information set $I\in\mathcal{I}$, let

$$\sigma(I):=\sum_{h\in I}\sigma(h).$$

1.3 Assessments

We now recall the standard definitions of belief systems, Bayes' rule, and consistency (Kreps and Wilson, 1982). For every nonterminal history $h\in H\setminus Z$, let I^h denote the unique information set containing h.

A *belief system* is a map $\mu:H\setminus Z\to[0,1]$ such that, for each information set $I\in\mathcal{I}$, we have $\sum_{h\in I}\mu(h)=1$. Thus, for $h\in I$, the number $\mu(h)$ is interpreted as the probability assigned to history h by the player moving at I, that is, $\iota(I)$. For $\hat{H}\subset I$, we write $\mu(\hat{H}):=\sum_{\hat{h}\in\hat{H}}\mu(\hat{h})$. An *assessment* is a pair (σ,μ) consisting of a strategy profile and a belief system.

Definition 1.1. An assessment (σ,μ) *satisfies Bayes' rule* if, for every information set I with $\sigma(I)>0$, we have

$$\mu(h)=\frac{\sigma(h)}{\sigma(I)}\qquad\text{for all } h\in I.$$

Definition 1.2. An assessment (σ,μ) is *consistent* if there exists a fully mixed sequence of strategy profiles (σ_n) such that $\sigma_n(a)\to\sigma(a)$ for all $a\in A$ and

$$\frac{\sigma_n(h)}{\sigma_n(I)}\to\mu(h)\qquad\text{for all } I\in\mathcal{I} \text{ and } h\in I.$$

Bayes' rule—sometimes called *weak consistency*—is often viewed as a mild requirement. By contrast, consistency—sometimes called *strong consistency*—is usually regarded as demanding, and its definition through convergent fully mixed sequences is often cumbersome to use. One of the aims of the paper is to identify intermediate restrictions on assessments that remain easy to verify.

2 Simple perfect Bayesian equilibria

This section generalizes two of Fudenberg and Tirole's (1991) reasonableness requirements on assessments—Bayes' rule whenever possible and no signaling what you don't know—to arbitrary extensive-form games.

2.1 Definition of "satisfying Bayes' rule whenever possible"

A natural strengthening of Bayes' rule requires that whenever a player assigns positive probability to a preceding information set, she should use Bayes' rule to compare the relative likelihoods of the continuation histories compatible with that event. We next formulate this idea for arbitrary finite extensive-form games.

Fix an information set $I\in\mathcal{I}$, a set of histories $H'\subset H$, and an assessment (σ,μ). We say that H' *precedes* $h\in I$ if some history in H' is a prefix of h. In the cases used below, this prefix is unique; denote it by $h_{H'}$ and write $h\backslash h_{H'}$ for the continuation actions after that prefix. In particular, if $H'=I'$ is an information set, uniqueness follows from Lemma 1.1. We say that H' *strictly precedes* h if, in addition, $h\backslash h_{H'}\neq\emptyset$. We write $I_{\succeq H'}$ (resp. $I_{\succ H'}$) for the set of histories in I that are (resp. strictly) preceded by H'. We say that $h\in I$ is *on the path of play after* I' if $h\in I_{\succ I'}$ and $\mu(h_{I'})\sigma(h\backslash h_{I'})>0$. The information set I is *on the path of play after* I' if at least one history in I is on the path of play after I'.

Definition 2.1. An assessment (σ,μ) *satisfies Bayes' rule whenever possible* if, for all $I,I'\in\mathcal{I}$ such that I is on the path of play after I' and $\mu(I_{\succ I'})>0$, we have

$$\frac{\mu(h)}{\mu(I_{\succ I'})}=\frac{\mu(h_{I'})\sigma(h\backslash h_{I'})}{\sum_{\hat{h}\in I_{\succ I'}}\mu(\hat{h}_{I'})\sigma(\hat{h}\backslash\hat{h}_{I'})}\qquad\text{for all } h\in I_{\succ I'}. \tag{2.1}$$

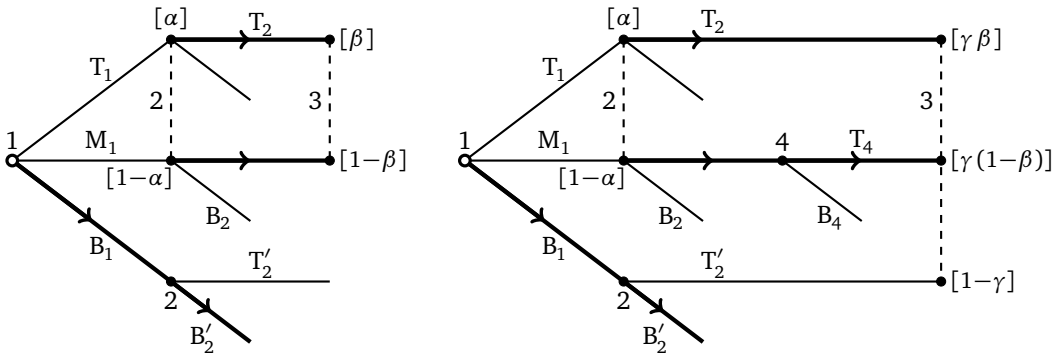


Figure 2

The intuition for Bayes' rule whenever possible is straightforward. If player $\iota(I)$ assigns positive probability to the event that play passed through I' and I can be reached from I' with positive probability, then the relative probabilities of the histories in $I_{\succ I'}$ should be obtained by standard Bayesian updating from beliefs at I'. The condition is also easy to verify. For each information set I, identify the information sets I' such that $\mu(I_{\succ I'})>0$ and I is on the path of play after I', and verify that (2.1) holds.

The examples in Figure 2 illustrate the restriction imposed by Bayes' rule whenever possible. In panel (a), player 3's information set is off path, so Bayes' rule does not directly determine β. However, that information set is on the path of play after player 2's upper information set. Hence Bayes' rule whenever possible implies $\alpha=\beta$. Indeed,

$$\beta=\mu(\mathrm{T}_1,\mathrm{T}_2)=\frac{\mu(\mathrm{T}_1)\sigma(\mathrm{T}_2)}{\mu(\mathrm{T}_1)\sigma(\mathrm{T}_2)+\mu(\mathrm{M}_1)\sigma(\mathrm{T}_2)}=\mu(\mathrm{T}_1)=\alpha. \tag{2.2}$$

In panel (b), player 3's information set is again off path, and it is no longer the case that player 3's information set succeeds player 2's information set: not every history in player 3's information set extends a history in player 2's information set. If $\gamma=0$, player 3 assigns probability zero to player 2's upper information set, and Bayes' rule whenever possible imposes no restriction on β. If instead $\gamma>0$, then conditional on player 3's information set being reached, player 3 assigns positive probability to player 2's upper information set. In that case Bayes' rule whenever possible again yields $\alpha=\beta$:

$$\beta=\frac{\mu(\mathrm{T}_1,\mathrm{T}_2)}{\mu(\mathrm{T}_1,\mathrm{T}_2)+\mu(\mathrm{M}_1,\mathrm{T}_2,\mathrm{T}_4)}=\frac{\mu(\mathrm{T}_1)\sigma(\mathrm{T}_2)}{\mu(\mathrm{T}_1)\sigma(\mathrm{T}_2)+\mu(\mathrm{M}_1)\sigma(\mathrm{T}_2)\sigma(\mathrm{T}_4)}=\mu(\mathrm{T}_1)=\alpha, \tag{2.3}$$

where we used that $\sigma(\mathrm{T}_4)=1$.

A useful characterization of the Bayes-rule-whenever-possible condition is the following.

Lemma 2.1. *An assessment (σ,μ) satisfies Bayes' rule whenever possible if and only if, for all histories h and $\hat{h}$ belonging to $I_{\succ I'}$ for some information sets $I, I' \in \mathcal{I}$, we have*

$$\frac{\mu(\hat{h})\mu(h_{I'})}{\mu(h)\mu(\hat{h}_{I'})} = \frac{\sigma(\hat{h}\backslash\hat{h}_{I'})}{\sigma(h\backslash h_{I'})} \tag{2.4}$$

whenever neither side is the indeterminate form $0/0$.[3]

Equation (2.4) can be interpreted by observing that its left-hand side can be rewritten as

$$\frac{\mu(\hat{h})/\mu(h)}{\mu(\hat{h}_{I'})/\mu(h_{I'})} \qquad \text{and} \qquad \frac{\mu(\hat{h})/\mu(\hat{h}_{I'})}{\mu(h)/\mu(h_{I'})}.$$

Using the first expression, (2.4) says that the relative likelihood of $\hat{h}$ and h at I is equal to their relative likelihood at I', adjusted by the relative probability of the continuations from I' to $\hat{h}$ and h. Using the second expression, (2.4) says that the ratio of the two updated likelihoods, $\mu(\hat{h})/\mu(\hat{h}_{I'})$ and $\mu(h)/\mu(h_{I'})$, should equal the ratio of the probabilities of the corresponding continuation histories, $\sigma(\hat{h}\backslash\hat{h}_{I'})$ and $\sigma(h\backslash h_{I'})$.

2.2 Definition of "no signaling what you don't know"

Consider the game and assessment in Figure 3(a). The game tree is the same as in Figure 2(a), but player 2 now chooses T_2 with probability zero. Because player 3's information set is no longer reached with positive probability after player 2's information set, Bayes' rule whenever possible imposes no restriction on α and β. Nonetheless, it is natural to require $\alpha=\beta$: player 2 does not know whether player 1 chose T_1 or M_1, so player 2's deviation should not change the relative likelihood of those two possibilities. This motivates the following condition, where we use $I'\times\{a\}$ to denote the set of histories of I' concatenated with $a\in A^{I'}$, that is, $\{(h',a)\,|\,h'\in I'\}$. Note that if h is preceded by $I'\times\{a\}$, then there is a unique $h'\in I'$ such that $(h',a)\preceq h$; we denote it by $h_{I'}$.

Definition 2.2. An assessment (σ,μ) *satisfies no signaling what you don't know* if, for all $I, I'\in\mathcal{I}$ and $a\in A^{I'}$ such that $\mu(I_{\succeq I'\times\{a\}})>0$, we have

$$\frac{\mu(h)}{\mu(I_{\succeq I'\times\{a\}})} = \frac{\mu(h_{I'})\sigma(h\backslash(h_{I'}\cup\{a\}))}{\sum_{\hat{h}\in I_{\succeq I'\times\{a\}}}\mu(\hat{h}_{I'})\sigma(\hat{h}\backslash(\hat{h}_{I'}\cup\{a\}))} \qquad \text{for all } h\in I_{\succeq I'\times\{a\}}, \tag{2.5}$$

whenever the denominator on the right-hand side is positive.

[3]Note that Lemma 2.1 is equivalent to requiring that $\mu(\hat{h})\mu(h_{I'})\sigma(h\backslash h_{I'})=\mu(h)\mu(\hat{h}_{I'})\sigma(\hat{h}\backslash\hat{h}_{I'})$ for all histories h and $\hat{h}$ belonging to $I_{\succ I'}$ for some information sets $I, I'\in\mathcal{I}$. We find the ratio formulation in (2.4) more intuitive.

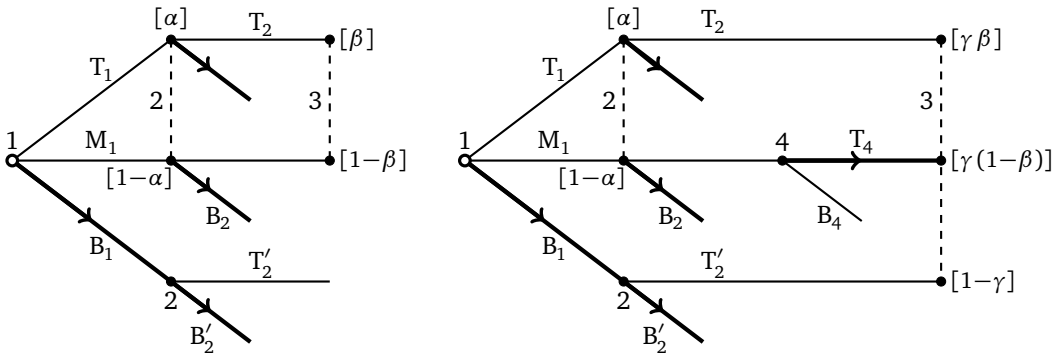


Figure 3

Thus, when several histories in the same information set share the same action at a preceding information set, that action should not affect their relative likelihoods. The point is precisely that the acting player could not have used this action to signal information she did not know.

In Figure 3(a), no signaling what you don't know imposes no restriction on player 2's upper information set, because the relevant denominator is zero, and no restriction on player 2's lower information set, because it contains only one history. At player 3's information set, however, (2.5) gives

$$\beta = \frac{\mu(T_1, T_2)}{\mu(T_1, T_2) + \mu(M_1, T_2)} = \frac{\mu(T_1)}{\mu(T_1) + \mu(M_1)} = \alpha.$$

Hence the assessment satisfies no signaling what you don't know if and only if $\alpha = \beta$. Relative to (2.2), the common factor $\sigma(T_2)$ has simply been canceled from the start.

Figure 3(b) gives the analogous comparison for the game in Figure 2(b) when player 2 now chooses T_2 with probability zero. If $\gamma = 0$, the conditioning event has probability zero and no signaling what you don't know imposes no restriction on β. If $\gamma > 0$, then

$$\beta = \frac{\mu(T_1, T_2)}{\mu(T_1, T_2) + \mu(M_1, T_2, T_4)} = \frac{\mu(T_1)}{\mu(T_1) + \mu(M_1)\sigma(T_4)} = \mu(T_1) = \alpha,$$

where we again use that $\sigma(T_4) = 1$. This is exactly the intended "no signaling what you don't know" logic: once player 3 assigns positive probability to player 2's upper information set, player 2's off-path choice of T_2 should not alter player 3's relative beliefs about T_1 and M_1.

The characterization of Bayes' rule whenever possible in Lemma 2.1 extends naturally to no signaling what you don't know as follows.

Lemma 2.2. *An assessment (σ,μ) satisfies no signaling what you don't know if and only if, for all $I,I'\in\mathcal{I}$, all $a\in A^{I'}$, and all $h,\hat{h}\in I_{\succeq I'\times\{a\}}$, we have*

$$\frac{\mu(\hat{h})\mu(h_{I'})}{\mu(h)\mu(\hat{h}_{I'})}=\frac{\sigma(\hat{h}\setminus(\hat{h}_{I'}\cup\{a\}))}{\sigma(h\setminus(h_{I'}\cup\{a\}))} \tag{2.6}$$

whenever neither side is the indeterminate form $0/0$.

2.3 Properties

We begin with the relative strength of the various belief restrictions.

Proposition 2.1. *Fix an assessment (σ,μ). Consider the following statements:*

(A) (σ,μ) is consistent.

(B) (σ,μ) satisfies no signaling what you don't know.

(C) (σ,μ) satisfies Bayes' rule whenever possible.

(D) (σ,μ) satisfies Bayes' rule.

Then $A\Rightarrow B$ and $A\Rightarrow C\Rightarrow D$.

Proof of Proposition 2.1. Assume (σ,μ) is consistent. Let σ_n be a fully mixed sequence witnessing consistency, and let μ_n denote the Bayes-rule beliefs generated by σ_n.

For $A\Rightarrow C$, fix $I,I'\in\mathcal{I}$ such that I is on the path of play after I' and $\mu(I_{\succ I'})>0$. Since (σ_n,μ_n) satisfies Bayes' rule,

$$\frac{\mu_n(h)}{\mu_n(I_{\succ I'})}=\frac{\mu_n(h_{I'})\sigma_n(h\setminus h_{I'})}{\sum_{\hat{h}\in I_{\succ I'}}\mu_n(\hat{h}_{I'})\sigma_n(\hat{h}\setminus\hat{h}_{I'})}.$$

The denominator on the right-hand side of (2.1) is strictly positive by the definition of "on the path of play after." The denominator in the last display therefore converges to that positive number, and taking limits yields (2.1). Hence (σ,μ) satisfies Bayes' rule whenever possible.

For $A\Rightarrow B$, fix $I,I'\in\mathcal{I}$ and $a\in A^{I'}$ such that $\mu(I_{\succeq I'\times\{a\}})>0$. To verify Definition 2.2, restrict attention to those instances for which the right-hand denominator in (2.5) is strictly positive. For $h\in I_{\succeq I'\times\{a\}}$, Bayes' rule for (σ_n,μ_n) gives

$$\frac{\mu_n(h)}{\mu_n(I_{\succeq I'\times\{a\}})}=\frac{\mu_n(h_{I'})\sigma_n(a)\sigma_n(h\setminus(h_{I'}\cup\{a\}))}{\sum_{\hat{h}\in I_{\succeq I'\times\{a\}}}\mu_n(\hat{h}_{I'})\sigma_n(a)\sigma_n(\hat{h}\setminus(\hat{h}_{I'}\cup\{a\}))}.$$

Since $\sigma_n(a)>0$, the common factor $\sigma_n(a)$ cancels for each n, before any limit is taken. Thus

$$\frac{\mu_n(h)}{\mu_n(I_{\succeq I'\times\{a\}})}=\frac{\mu_n(h_{I'})\sigma_n(h\setminus(h_{I'}\cup\{a\}))}{\sum_{\hat{h}\in I_{\succeq I'\times\{a\}}}\mu_n(\hat{h}_{I'})\sigma_n(\hat{h}\setminus(\hat{h}_{I'}\cup\{a\}))}.$$

The denominator in the last display converges to the positive denominator in (2.5), and $\mu_n(I_{\succeq I' \times \{a\}}) \to \mu(I_{\succeq I' \times \{a\}}) > 0$. Taking limits yields (2.5). Hence (σ, μ) satisfies no signaling what you don't know.

For $C \Rightarrow D$, let I be an information set with $\sigma(I) > 0$. If $I = \{\emptyset\}$, the claim is immediate. Otherwise I is on the path of play after the root information set $\{\emptyset\}$, and, by Lemma 1.1 and the definition of a game, $\{\emptyset\}$ is a singleton. Applying (2.1) with $I' = \{\emptyset\}$, we have

$$\mu(h) = \frac{\sigma(h)}{\sum_{\hat{h} \in I} \sigma(\hat{h})} = \frac{\sigma(h)}{\sigma(I)} \qquad \text{for all } h \in I.$$

Thus Bayes' rule holds at every information set with strictly positive reach probability. □

2.4 Simple perfect Bayesian equilibria

Fix an assessment (σ, μ). For each player $i \in N$ and information set $I \in \mathcal{I}$ with $\iota(I) \neq 0$, define the expected payoff of player i conditional on I by

$$u_i(\sigma, \mu \mid I) = \sum_{h \in I} \sum_{z \in Z^h} \mu(h) \sigma(z \setminus h) u_i(z), \tag{2.7}$$

where Z^h denotes the set of terminal histories that succeed h and $\sigma(z \setminus h)$ is the product of probabilities of the continuation actions after h.

Definition 2.3. An assessment (σ, μ) is *sequentially rational* if for all $I \in \mathcal{I}$ with $\iota(I) \neq 0$ we have

$$u_{\iota(I)}(\sigma, \mu \mid I) \geq u_{\iota(I)}(\sigma'_{\iota(I)}, \sigma_{-\iota(I)}, \mu \mid I) \qquad \text{for all } \sigma'_{\iota(I)} \in \Sigma_{\iota(I)}.$$

In words, sequential rationality requires that at every information set the prescribed continuation behavior of the moving player be optimal given her beliefs and the other players' strategies. Because payoffs are evaluated conditional on the information set, changes in the deviating player's behavior at information sets that cannot be reached afterward do not affect the comparison.

Definition 2.4. An assessment is a *simple perfect Bayesian equilibrium* if it is sequentially rational and satisfies Bayes' rule whenever possible and no signaling what you don't know.

Simple perfect Bayesian equilibrium is meant to impose only elementary and directly interpretable restrictions on beliefs. Players update by Bayes' rule whenever the relevant conditioning event has positive probability, and actions cannot signal information unavailable to the player who chose them. Because both restrictions are local, checking the definition in a finite game reduces to verifying sequential rationality and a relatively small collection of belief equalities of the form

(2.1) and (2.5). The concept therefore retains much of the discipline of perfect Bayesian reasoning while avoiding perturbing sequences and auxiliary belief-revision objects.

Subgame perfection. Following Selten (1965), subgame perfection has been regarded as a basic reasonableness condition for equilibrium concepts. It requires that the restrictions of the equilibrium objects to a subgame constitute an equilibrium of that subgame.

A *subgame* is a set of histories $H' \subset H$ such that: if $h \in H'$ and $h \preceq \hat{h} \in H$, then $\hat{h} \in H'$; if $I \cap H' \neq \emptyset$, then $I \subseteq H'$; and H' has a unique minimal history. The corresponding induced subgame is obtained by taking that minimal history as the root and retaining exactly the continuations and information sets contained in H'. If r is the minimal history, subgame histories are identified with suffixes k such that $(r,k) \in H'$. The restriction of an assessment to a subgame restricts strategies to actions in the subgame, inherits beliefs at non-root information sets through the map $k \mapsto (r,k)$, and assigns probability one to the induced root.

Proposition 2.2. *The restriction of a simple perfect Bayesian equilibrium to a subgame is a simple perfect Bayesian equilibrium of the induced subgame.*

Proposition 2.2 can be generalized to subforms, analogously to the subform-perfection result in Dilmé (2024) for sequential equilibria.

3 Independence from common actions

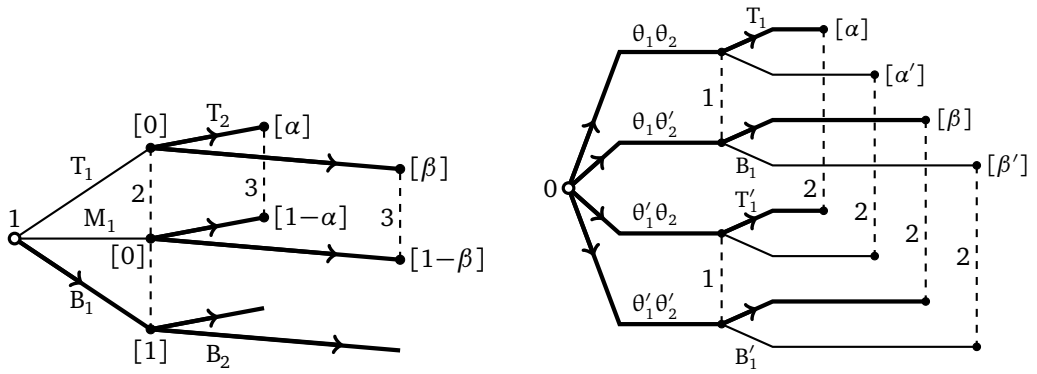
3.1 Cross-pair independence

As we have seen, Bayes' rule whenever possible and no signaling what you don't know are useful for disciplining beliefs across sequentially ordered information sets. They are much less effective across information sets that are not sequentially ordered, even when the same underlying intuition seems to apply. The next examples illustrate the point.

Consider the game in panel (a) of Figure 4. Player 2 does not observe player 1's action, and player 3 observes player 2's action when player 1 does not exit. Bayes' rule whenever possible cannot be applied because player 3's information sets are not on the path after the other players' information sets. Similarly, no signaling what you don't know is silent because the beliefs assigned to histories leading to each of player 3's information sets are zero. Still, for any fully mixed strategy profile σ_n we have

$$\frac{\alpha}{1-\alpha} \cdot \frac{1-\beta}{\beta} = \frac{\sigma_n(\mathrm{T}_1)\sigma_n(\mathrm{T}_2)}{\sigma_n(\mathrm{T}_1)\sigma_n(\mathrm{M}_2)} \cdot \frac{\sigma_n(\mathrm{B}_1)\sigma_n(\mathrm{M}_2)}{\sigma_n(\mathrm{B}_1)\sigma_n(\mathrm{T}_2)} = 1,$$

so every consistent assessment supporting the displayed strategy profile must satisfy $\alpha = \beta$. Intu-



(a) (b)

Figure 4

itively, once player 3 conditions on player 1's deviation, the inference about player 2's choice should not depend on information that player 1 herself did not have.

A similar logic applies to games with observable actions. Consider the game in panel (b) of Figure 4. In this game, nature first chooses the types of players 1 and 2. We assume that types are independent. Then player 1 chooses her action. Finally, player 2 observes whether player 1 chose top or bottom, but not player 1's type. Bayes' rule whenever possible implies that $\alpha = \beta$, that is, types θ_2 and θ_2' have the same beliefs about player 1's type if she chooses top, and these coincide with the prior. Neither Bayes' rule whenever possible nor no signaling what you don't know imposes any condition on beliefs following player 1's bottom action. Still, the logic behind no signaling what you don't know suggests that types θ_2 and θ_2' have the same beliefs about player 1's type also when she chooses bottom, that is, that $\alpha' = \beta'$. This relation is satisfied in any consistent assessment.

Cross-pair independence. Lemmas 2.1 and 2.2 shed light on a powerful approach to establishing reasonableness conditions: relating the relative likelihoods of two pairs of histories to the relative likelihoods of the actions that the pairs do not have in common. For example, the numerator on the right-hand side of (2.4) is the likelihood of the actions appearing in the histories in the numerator on the left-hand side but not in the histories in the denominator; the denominator is defined analogously.

We now define a property that extends the logic behind Bayes-rule-whenever-possible and no-signaling-what-you-don't-know beyond information sets that are necessarily sequentially ordered. To do that, we first briefly recall the definition of a multiset. A *multiset* is an unordered collection of elements that allows multiple instances of the same element. For instance, ordinary sets can be viewed as multisets without repeated elements. If $\hat{A}$ is a multiset of actions, we naturally define $\sigma(\hat{A}) = \prod_{a \in \hat{A}} \sigma(a)^{m_{\hat{A}}(a)}$, where $m_{\hat{A}}(a)$ is the multiplicity of action a, with the product equal to one

when $\hat{A}$ is empty. The multiset union, intersection, and difference of $\hat{A}$ and $\tilde{A}$ are denoted by $\hat{A} \sqcup \tilde{A}$, $\hat{A} \sqcap \tilde{A}$, and $\hat{A} \setminus \tilde{A}$, respectively.[4]

Definition 3.1. An assessment (σ, μ) *satisfies cross-pair independence (from common actions)* if, for any $I_1, I_2 \in \mathcal{I}$ and any two pairs of histories $h_1, \hat{h}_1 \in I_1$ and $h_2, \hat{h}_2 \in I_2$, we have

$$\frac{\mu(\hat{h}_1)\mu(\hat{h}_2)}{\mu(h_1)\mu(h_2)} = \frac{\sigma((\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2))}{\sigma((h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2))} \tag{3.1}$$

whenever neither side is an indeterminate form $0/0$.

The cross-pair independence condition allows one to compare the likelihoods of two pairs of histories in possibly different information sets. As we shall see, it is stronger than Bayes' rule whenever possible and no signaling what you don't know, but weaker than consistency. It also involves a larger (but finite) number of constraints. Note that by Lemma 1.1, no primitive action occurs twice along a single history, but repeated copies may arise when several histories are added as multisets.

The left-hand side of (3.1) is familiar: it contains the ratio of the products of the likelihoods of two pairs of histories, each pair belonging to the same information set. The left sides of (2.4) and (2.6) have the same form (there, information sets are sequentially ordered), and correspond to the quotient of the likelihood ratios of $\hat{h}_1$ and h_1, and h_2 and $\hat{h}_2$. On the right-hand side, note that $(\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2)$ contains the actions in $\hat{h}_1 \sqcup \hat{h}_2$ that are not common to those in $h_1 \sqcup h_2$. Similarly, $(h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2)$ contains the actions in $h_1 \sqcup h_2$ that are not common to those in $\hat{h}_1 \sqcup \hat{h}_2$. Cross-pair independence requires that the product of the likelihood ratios depends only on the quotient of the corresponding non-common actions.

Using the cross-pair independence condition is often straightforward. For example, consider again the game in panel (a) of Figure 4. Cross-pair independence implies that[5]

$$\frac{\alpha}{1-\alpha} \cdot \frac{1-\beta}{\beta} = \frac{\sigma(\{\mathrm{T}_1, \mathrm{T}_2, \mathrm{B}_1, \mathrm{B}_2\} \setminus \{\mathrm{T}_1, \mathrm{B}_2, \mathrm{B}_1, \mathrm{T}_2\})}{\sigma(\{\mathrm{T}_1, \mathrm{B}_2, \mathrm{B}_1, \mathrm{T}_2\} \setminus \{\mathrm{T}_1, \mathrm{T}_2, \mathrm{B}_1, \mathrm{B}_2\})} = \frac{\sigma(\emptyset)}{\sigma(\emptyset)} = 1,$$

hence $\alpha(1-\beta) = (1-\alpha)\beta$, so $\alpha = \beta$. For the game in panel (b) of Figure 4, we have

$$\frac{\alpha'}{1-\alpha'} \cdot \frac{1-\beta'}{\beta'} = \frac{\sigma(\{\theta_1\theta_2, \mathrm{B}_1, \theta_1'\theta_2', \mathrm{B}_1'\} \setminus \{\theta_1'\theta_2, \mathrm{B}_1', \theta_1\theta_2', \mathrm{B}_1\})}{\sigma(\{\theta_1'\theta_2, \mathrm{B}_1', \theta_1\theta_2', \mathrm{B}_1\} \setminus \{\theta_1\theta_2, \mathrm{B}_1, \theta_1'\theta_2', \mathrm{B}_1'\})} = \frac{\sigma(\theta_1\theta_2)\sigma(\theta_1'\theta_2')}{\sigma(\theta_1'\theta_2)\sigma(\theta_1\theta_2')} = 1,$$

[4] If, for example, $\hat{A} = \{a_1, a_1, a_1, a_2\}$ and $\tilde{A} = \{a_1, a_3\}$ (where a_1, a_2, and a_3 are different actions), then $\sigma(\hat{A}) = \sigma(a_1)^3\, \sigma(a_2)$, $\hat{A} \sqcup \tilde{A} = \{a_1, a_1, a_1, a_1, a_2, a_3\}$, $\hat{A} \sqcap \tilde{A} = \{a_1\}$, and $\hat{A} \setminus \tilde{A} = \{a_1, a_1, a_2\}$.

[5] We do not explicitly specify which histories in the game correspond to the histories in condition (3.1) when it is clear. In this example, $\hat{h}_1 = (\mathrm{T}_1, \mathrm{T}_2)$, $\hat{h}_2 = (\mathrm{B}_1, \mathrm{B}_2)$, $h_1 = (\mathrm{T}_1, \mathrm{B}_2)$, and $h_2 = (\mathrm{B}_1, \mathrm{T}_2)$.

where the last equality follows from type independence. Thus $\alpha' = \beta'$ in any assessment satisfying cross-pair independence.

Relative strength. The following result establishes that cross-pair independence is stronger than both Bayes' rule whenever possible and no signaling what you don't know. In Section 3.4, we show that cross-pair independence is weaker than consistency.

Proposition 3.1. *If an assessment satisfies cross-pair independence, then it satisfies both Bayes' rule whenever possible and no signaling what you don't know.*

Proof of Proposition 3.1. Let (σ, μ) satisfy cross-pair independence. We first prove Bayes' rule whenever possible. Let h and $\hat{h}$ belong to $I_{\succ I'}$ for some information sets $I, I' \in \mathcal{I}$, and suppose that neither side of (2.4) is the indeterminate form $0/0$. Put

$$C := h \setminus h_{I'}, \qquad \hat{C} := \hat{h} \setminus \hat{h}_{I'}.$$

Apply cross-pair independence to the two pairs $(h, \hat{h})$ in I and $(\hat{h}_{I'}, h_{I'})$ in I'. The left side is the left side of (2.4), while the right side is

$$\frac{\sigma((\hat{h} \sqcup h_{I'}) \setminus (h \sqcup \hat{h}_{I'}))}{\sigma((h \sqcup \hat{h}_{I'}) \setminus (\hat{h} \sqcup h_{I'}))} = \frac{\sigma(\hat{C} \setminus C)}{\sigma(C \setminus \hat{C})}.$$

If $\sigma(C) > 0$, then every action in $C \sqcap \hat{C}$ has positive probability, and therefore

$$\frac{\sigma(\hat{C} \setminus C)}{\sigma(C \setminus \hat{C})} = \frac{\sigma(\hat{C} \setminus C)\sigma(C \sqcap \hat{C})}{\sigma(C \setminus \hat{C})\sigma(C \sqcap \hat{C})} = \frac{\sigma(\hat{C})}{\sigma(C)}.$$

If $\sigma(C) = 0 < \sigma(\hat{C})$, the same multiset cancellation shows that $\sigma(\hat{C} \setminus C) > 0$ and $\sigma(C \setminus \hat{C}) = 0$, so both ratios are $+\infty$. The case $\sigma(\hat{C}) = 0 < \sigma(C)$ is symmetric and gives value 0 on both sides. If $\sigma(C) = \sigma(\hat{C}) = 0$, the right-hand side of (2.4) is the indeterminate form $0/0$ and the pair is outside the stated domain. Thus (2.4) holds whenever it is not an indeterminate form $0/0$. By Lemma 2.1, (σ, μ) satisfies Bayes' rule whenever possible.

The proof for no signaling what you don't know is similar, except that the common action at the preceding information set is also canceled. Fix $I, I' \in \mathcal{I}$, $a \in A^{I'}$, and $h, \hat{h} \in I_{\succeq I' \times \{a\}}$, and suppose that neither side of (2.6) is the indeterminate form $0/0$. Put

$$C := h \setminus (h_{I'} \cup \{a\}), \qquad \hat{C} := \hat{h} \setminus (\hat{h}_{I'} \cup \{a\}).$$

Applying cross-pair independence to $(h, \hat{h})$ and $(\hat{h}_{I'}, h_{I'})$ again gives the left side of (2.6). On the

right side, the prefixes and the common copy of a cancel, leaving

$$\frac{\sigma(\hat{C}\setminus C)}{\sigma(C\setminus\hat{C})}.$$

As above, this ratio is $\sigma(\hat{C})/\sigma(C)$ whenever the latter is not $0/0$, including the cases in which it is 0 or $+\infty$. Hence (2.6) holds. By Lemma 2.2, (σ,μ) satisfies no signaling what you don't know. □

Recall that no-signaling-what-you-don't-know implies $\alpha=\beta$ for the assessment depicted in panel (b) of Figure 3 when $\gamma>0$. By Proposition 3.1, the same requirement can also be obtained using cross-pair independence: (3.1) yields

$$\frac{\alpha}{1-\alpha}\cdot\frac{1-\beta}{\beta}=\frac{\alpha}{1-\alpha}\cdot\frac{\gamma(1-\beta)}{\gamma\beta}=\frac{\sigma(\{T_1,M_1,T_2,T_4\}\setminus\{M_1,T_1,T_2\})}{\sigma(\{M_1,T_1,T_2\}\setminus\{T_1,M_1,T_2,T_4\})}=\frac{\sigma(T_4)}{\sigma(\emptyset)}=1,$$

that is, $\alpha(1-\beta)=(1-\alpha)\beta$, so $\alpha=\beta$. Note that exactly the same argument applies for the assessment in panel (b) of Figure 2, where $\alpha=\beta$ follows from Bayes' rule whenever possible.

3.2 Perfect Bayesian equilibrium

As discussed in the introduction, Fudenberg and Tirole (1991) used assessments to define the concept of perfect Bayesian equilibrium for multi-stage games with observable actions and independent types. Most extensions of perfect Bayesian equilibrium to arbitrary extensive-form games are not defined in terms of assessments. We now provide a definition of perfect Bayesian equilibrium using cross-pair independence.

Fudenberg and Tirole's concept of reasonable beliefs imposes different conditions. These include common beliefs, belief independence, Bayes' rule whenever possible, and no signaling what you don't know. Nevertheless, these are tailored to multi-stage games with observable actions and independent types and are difficult to extend to general extensive-form games. In Appendix B we review their definitions. Importantly, we show that, for games with observable actions and independent types, cross-pair independence implies Fudenberg and Tirole's (1991) reasonableness conditions, and that, with an additional reasonableness condition, the two notions are equivalent. We therefore view the following definition as a natural extension of perfect Bayesian equilibrium to arbitrary extensive-form games.

Definition 3.2. An assessment is a *perfect Bayesian equilibrium* if it is sequentially rational and satisfies cross-pair independence.

From Proposition 3.1, it follows that every perfect Bayesian equilibrium is a simple perfect

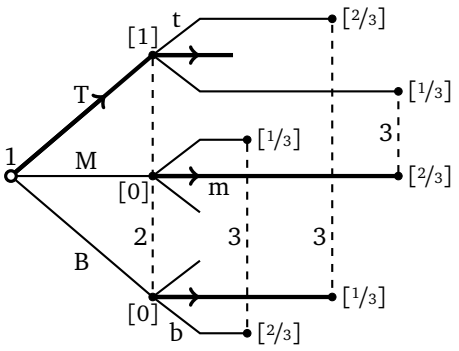


Figure 5

Bayesian equilibrium, but the converse need not hold in general. Moreover, perfect Bayesian equilibria are subgame perfect, as the following result establishes.

Proposition 3.2. *The restriction of a perfect Bayesian equilibrium to a subgame is a perfect Bayesian equilibrium of the induced subgame.*

3.3 Independence from common actions

As we shall now see, consistency is stronger than cross-pair independence. Intuitively, some constraints imposed by consistency require comparing three or more pairs of histories, while cross-pair independence only allows comparing two. In this section, we generalize cross-pair independence to an arbitrary number of pairs of histories that will turn out to be equivalent to consistency.

Consider Figure 5, which reproduces Figure 7 in Kohlberg and Reny (1997). Kohlberg and Reny show, using limits of mixed strategy profiles, that the depicted assessment is not consistent. Note that, nonetheless, this assessment satisfies cross-pair independence. Indeed, for any two pairs of player 3's histories where each pair belongs to some information set, the numerator and the denominator of the right-hand side of equation (3.1) are zero. For example, taking

$$h_1 = (\mathrm{T}, \mathrm{t}), \quad \hat{h}_1 = (\mathrm{B}, \mathrm{m}), \quad h_2 = (\mathrm{B}, \mathrm{b}), \quad \hat{h}_2 = (\mathrm{M}, \mathrm{t}),$$

we have $\sigma((\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2)) = \sigma(\{\mathrm{m}, \mathrm{M}\}) = 0$ and $\sigma((h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2)) = \sigma(\{\mathrm{b}, \mathrm{T}\}) = 0$. The same applies if one or both pairs are in player 2's information set. Hence, pairwise conditions are not sufficient to fully chracterize consistency.

We now generalize cross-pair independence by allowing arbitrary multisets of nonterminal histories.

Definition 3.3. An assessment (σ,μ) *satisfies independence from common actions* if, for all finite multisets of nonterminal histories $\hat{H}$ and $\tilde{H}$ satisfying

$$m_{\hat{H}}(I)=m_{\tilde{H}}(I) \qquad \text{for all } I\in\mathcal{I},$$

we have

$$\frac{\prod_{\hat{h}\in\hat{H}}\mu(\hat{h})}{\prod_{\tilde{h}\in\tilde{H}}\mu(\tilde{h})}=\frac{\sigma\big(\big(\bigsqcup_{\hat{h}\in\hat{H}}\hat{h}\big)\backslash\big(\bigsqcup_{\tilde{h}\in\tilde{H}}\tilde{h}\big)\big)}{\sigma\big(\big(\bigsqcup_{\tilde{h}\in\tilde{H}}\tilde{h}\big)\backslash\big(\bigsqcup_{\hat{h}\in\hat{H}}\hat{h}\big)\big)} \tag{3.2}$$

whenever neither side is an indeterminate form $0/0$. Products over empty multisets of histories are equal to one.

The condition says that relative likelihoods should depend only on the actions that remain after canceling the actions common to both sides. When all histories involved are on path, the condition follows immediately from Bayes' rule. Indeed, if each information set containing a history in $\hat{H}$ or $\tilde{H}$ is reached with positive probability, then

$$\frac{\prod_{\hat{h}\in\hat{H}}\mu(\hat{h})}{\prod_{\tilde{h}\in\tilde{H}}\mu(\tilde{h})}=\frac{\prod_{\hat{h}\in\hat{H}}\frac{\sigma(\hat{h})}{\sigma(I^{\hat{h}})}}{\prod_{\tilde{h}\in\tilde{H}}\frac{\sigma(\tilde{h})}{\sigma(I^{\tilde{h}})}}\cdot\frac{\prod_{I\in\mathcal{I}}\sigma(I)^{m_{\tilde{H}}(I)}}{\prod_{I\in\mathcal{I}}\sigma(I)^{m_{\hat{H}}(I)}}=\frac{\prod_{\hat{h}\in\hat{H}}\sigma(\hat{h})}{\prod_{\tilde{h}\in\tilde{H}}\sigma(\tilde{h})}\cdot\frac{\prod_{I\in\mathcal{I}}\sigma(I)^{m_{\tilde{H}}(I)}}{\prod_{I\in\mathcal{I}}\sigma(I)^{m_{\hat{H}}(I)}},$$

where I^h is the information set containing h. The information-set factors cancel by the requirement that $\hat{H}$ and $\tilde{H}$ have the same number of histories of each information set, and the remaining ratio is exactly the ratio of uncanceled action probabilities on the right-hand side of (3.2).

We now show that the assessment in Figure 5 does not satisfy independence-from-common-actions. To see this, define

$$\hat{H}=\{(\mathrm{T},\mathrm{t}),(\mathrm{M},\mathrm{m}),(\mathrm{B},\mathrm{b})\} \qquad \text{and} \qquad \tilde{H}=\{(\mathrm{M},\mathrm{b}),(\mathrm{B},\mathrm{t}),(\mathrm{T},\mathrm{m})\},$$

so the left side of (3.2) is

$$\frac{\mu(\mathrm{T},\mathrm{t})\mu(\mathrm{M},\mathrm{m})\mu(\mathrm{B},\mathrm{b})}{\mu(\mathrm{M},\mathrm{b})\mu(\mathrm{B},\mathrm{t})\mu(\mathrm{T},\mathrm{m})},$$

which is equal to 2^3, while the right side is

$$\frac{\sigma(\{\mathrm{T},\mathrm{t},\mathrm{M},\mathrm{m},\mathrm{B},\mathrm{b}\}\setminus\{\mathrm{M},\mathrm{b},\mathrm{B},\mathrm{t},\mathrm{T},\mathrm{m}\})}{\sigma(\{\mathrm{M},\mathrm{b},\mathrm{B},\mathrm{t},\mathrm{T},\mathrm{m}\}\setminus\{\mathrm{T},\mathrm{t},\mathrm{M},\mathrm{m},\mathrm{B},\mathrm{b}\})}=\frac{\sigma(\emptyset)}{\sigma(\emptyset)},$$

which is equal to 1. Thus both sides of (3.2) are well defined but have different values, so the assessment fails independence-from-common-actions.

3.4 Consistency

We now state the equivalence between independence from common actions and consistency.

Proposition 3.3. *An assessment satisfies the independence from common actions condition if and only if it is consistent.*

Proof of Proposition 3.3. We prove the two directions separately.

Consistency ⇒ independence from common actions. Let σ_n be a fully mixed sequence witnessing consistency, and define

$$\mu_n(h) := \frac{\sigma_n(h)}{\sigma_n(I)} \qquad (h \in I).$$

Take multisets $\hat{H}$ and $\tilde{H}$ satisfying the counting condition in Definition 3.3, and write

$$m_I := m_{\hat{H}}(I) = m_{\tilde{H}}(I).$$

For every n, Bayes' rule gives

$$\prod_{\hat{h}\in\hat{H}} \mu_n(\hat{h}) = \prod_{\hat{h}\in\hat{H}} \frac{\sigma_n(\hat{h})}{\sigma_n(I^{\hat{h}})}, \qquad \prod_{\tilde{h}\in\tilde{H}} \mu_n(\tilde{h}) = \prod_{\tilde{h}\in\tilde{H}} \frac{\sigma_n(\tilde{h})}{\sigma_n(I^{\tilde{h}})}.$$

As explained after Definition 3.3, the information-set denominators therefore cancel. Since histories are identified with multisets of actions,

$$\frac{\prod_{\hat{h}\in\hat{H}} \mu_n(\hat{h})}{\prod_{\tilde{h}\in\tilde{H}} \mu_n(\tilde{h})} = \frac{\sigma_n\big(\bigsqcup_{\hat{h}\in\hat{H}} \hat{h}\big)}{\sigma_n\big(\bigsqcup_{\tilde{h}\in\tilde{H}} \tilde{h}\big)} = \frac{\sigma_n\big(\big(\bigsqcup_{\hat{h}\in\hat{H}} \hat{h}\big) \setminus \big(\bigsqcup_{\tilde{h}\in\tilde{H}} \tilde{h}\big)\big)}{\sigma_n\big(\big(\bigsqcup_{\tilde{h}\in\tilde{H}} \tilde{h}\big) \setminus \big(\bigsqcup_{\hat{h}\in\hat{H}} \hat{h}\big)\big)}.$$

All cancellations above are made for the fully mixed profile σ_n, before taking limits. Passing to the limit in the extended nonnegative reals gives (3.2) whenever the limiting ratio is not an indeterminate form $0/0$. Thus (σ, μ) satisfies independence from common actions.

Independence from common actions ⇒ consistency. We use the following integer-row form of the Kohlberg and Reny (1997) characterization of consistency. The relevant part of their analysis concerns linear systems of equations in which the right-hand sides may take on infinite values.

Lemma 3.1 (Kohlberg–Reny)**.** *Fix an assessment* (σ, μ)*. Set* $\log 0 := -\infty$*. Consider the following system of equations in the variables* $\{x_a \mid a \in A\}$*:*

$$\sum_{a\in h^+\setminus h^-} x_a - \sum_{a\in h^-\setminus h^+} x_a = \log \mu(h^+) - \log \mu(h^-)$$

for every ordered pair $h^+, h^- \in H$ in the same information set whose beliefs are not both zero, together with the rows

$$x_a = \log \sigma(a)$$

for every action a with $\sigma(a) > 0$. A finite integer linear combination of right-hand sides is well defined if, after collecting terms, it does not contain both $+\infty$ and $-\infty$. The assessment (σ, μ) is consistent if and only if every integer linear combination of rows that eliminates all variables has right-hand side equal to zero whenever that right-hand side is well defined.

Consider the corresponding system of equations for $\{x_a \mid a \in A\}$:

$$\sum_{a \in h^+ \setminus h^-} x_a - \sum_{a \in h^- \setminus h^+} x_a = \log \mu(h^+) - \log \mu(h^-)$$

for every ordered pair $h^+, h^- \in H$ in the same information set with not both beliefs equal to zero, together with

$$x_a = \log \sigma(a)$$

for the positive-probability actions. Take an integer linear combination of rows that eliminates all variables and whose right-hand side is well defined. If a pair row (h^+, h^-) enters with coefficient $\lambda > 0$, put λ copies of h^+ into a multiset $\hat{H}$ and λ copies of h^- into a multiset $\tilde{H}$; if $\lambda < 0$, reverse the roles of h^+ and h^-. For the action rows, first combine equal actions to their net coefficients: put positive net copies of an action into $\hat{A}$ and negative net copies into $\tilde{A}$. After canceling common copies, we may assume $\hat{A} \sqcap \tilde{A} = \emptyset$. We also cancel common history copies from $\hat{H}$ and $\tilde{H}$; this preserves the information-set counts and the multiset equality below.

Each remaining pair row contributes one history from a given information set to each side, so $\hat{H}$ and $\tilde{H}$ satisfy the counting condition in Definition 3.3. Eliminating all variables means exactly that

$$\Big(\bigsqcup_{\hat{h} \in \hat{H}} \hat{h}\Big) \sqcup \hat{A} = \Big(\bigsqcup_{\tilde{h} \in \tilde{H}} \tilde{h}\Big) \sqcup \tilde{A}.$$

Hence, because $\hat{A} \sqcap \tilde{A} = \emptyset$, we have

$$\bigsqcup_{\hat{h} \in \hat{H}} \hat{h} \setminus \bigsqcup_{\tilde{h} \in \tilde{H}} \tilde{h} = \tilde{A}, \qquad \bigsqcup_{\tilde{h} \in \tilde{H}} \tilde{h} \setminus \bigsqcup_{\hat{h} \in \hat{H}} \hat{h} = \hat{A}.$$

Because the corresponding linear combination of right-hand sides is well defined, the expression

$$\frac{\prod_{\hat{h} \in \hat{H}} \mu(\hat{h})}{\prod_{\tilde{h} \in \tilde{H}} \mu(\tilde{h})} \frac{\sigma(\hat{A})}{\sigma(\tilde{A})} \tag{3.3}$$

is well defined in the extended nonnegative reals. Every action in $\hat{A} \sqcup \tilde{A}$ has strictly positive probability, so $\sigma(\hat{A})$ and $\sigma(\tilde{A})$ are strictly positive. Therefore the belief ratio in (3.3) is not an indeterminate form $0/0$, and neither side of (3.2) is an indeterminate form $0/0$. Applying independence from common actions gives

$$\frac{\prod_{\hat{h} \in \hat{H}} \mu(\hat{h})}{\prod_{\tilde{h} \in \tilde{H}} \mu(\tilde{h})} = \frac{\sigma(\tilde{A})}{\sigma(\hat{A})}.$$

Thus expression (3.3) is equal to 1. Taking logarithms shows that the original linear combination of right-hand sides is well defined and equal to 0. By Lemma 3.1, (σ, μ) is consistent. □

As a corollary, every consistent assessment satisfies cross-pair independence, since (3.1) is the two-pair special case of (3.2).

A Omitted proofs

Proof of Lemma 2.1. Fix $I, I' \in \mathcal{I}$ and let $E := I_{\succ I'}$. For each $k \in E$ set

$$p_k := \mu(k_{I'}) \sigma\big(k \setminus k_{I'}\big).$$

Write $B := \mu(E)$ and $S := \sum_{r \in E} p_r$. The term S is the denominator in (2.1).

Only if direction. Suppose (σ, μ) satisfies Bayes' rule whenever possible. Note that if $B = 0$, then $\mu(k) = 0$ for every $k \in E$; if $S = 0$, then $p_k = 0$ for every $k \in E$. In that case the cross-product relation $\mu(\hat{k})\, p_k = \mu(k)\, p_{\hat{k}}$ reduces to $0 = 0$ for all $k, \hat{k} \in E$, so (2.4) is satisfied without imposing any additional conditions (recall footnote 3). Otherwise, Definition 2.1 says

$$\frac{\mu(k)}{B} = \frac{p_k}{S} \qquad (k \in E).$$

Cross-multiplying shows that $\mu(\hat{k})\, p_k = \mu(k)\, p_{\hat{k}}$ for all $\hat{k}, k \in E$, which is exactly (2.4) on its domain.

If direction. Conversely, assume (2.4) holds whenever neither side is $0/0$. Suppose that $B > 0$ and $S > 0$. Pick $r \in E$ with $\mu(r) > 0$ and $k \in E$ with $p_k > 0$. Since $p_k > 0$, the ratio in (2.4) is not $0/0$; hence $\mu(r)\, p_k = \mu(k)\, p_r$ implies $p_r > 0$ and $\mu(k) > 0$. Thus $\mu(\cdot)$ and $p(\cdot)$ have the same support on E. Fix k_0 in this common support. Applying (2.4) with $\hat{k} = k_0$ and summing over $k \in E$ gives

$$B\, p_{k_0} = \mu(k_0) S.$$

Because $p_{k_0} > 0$, dividing $B\,p_{k_0} = \mu(k_0)\,S$ by $B\,S$ and using $\mu(k)\,p_{k_0} = \mu(k_0)\,p_k$ for each $k \in E$ shows

$$\frac{\mu(k)}{B} = \frac{p_k}{S} \qquad (k \in E).$$

This is exactly (2.1), so (σ, μ) satisfies Bayes' rule whenever possible. □

Proof of Lemma 2.2. Fix $I, I' \in \mathcal{I}$ and $a \in A^{I'}$. Set

$$E := I_{\geq I' \times \{a\}}, \qquad p_k := \mu(k_{I'})\sigma\big(k \setminus \big(k_{I'} \cup \{a\}\big)\big) \quad (k \in E).$$

Let $B := \mu(E)$ and $S := \sum_{r \in E} p_r$, the denominator in (2.5).

If (σ, μ) satisfies no signaling what you don't know then, as in the proof of Lemma 2.1, either $B = 0$ or $S = 0$, or else

$$\frac{\mu(k)}{B} = \frac{p_k}{S} \qquad (k \in E)$$

by Definition 2.2, which implies the cross-product relation (2.6). Conversely, assume (2.6) holds on its domain and that $B > 0$ and $S > 0$. Picking r with $\mu(r) > 0$ and k with $p_k > 0$ shows, as in the previous proof, that the supports of $\mu(\cdot)$ and $p(\cdot)$ on E coincide. Fixing k_0 in this common support, summing $\mu(k)\,p_{k_0} = \mu(k_0)\,p_k$ over $k \in E$ yields $B\,p_{k_0} = \mu(k_0)\,S$ and therefore

$$\frac{\mu(k)}{B} = \frac{p_k}{S} \qquad (k \in E),$$

which is (2.5). Thus (σ, μ) satisfies no signaling what you don't know. □

Proof of Proposition 2.2. Let r be the unique minimal history of the subgame and let $\phi(k) := (r, k)$ be the embedding of the induced subgame into the original game. Let (σ', μ') be the restriction of a simple perfect Bayesian equilibrium (σ, μ) to this subgame. Let I be the original information set containing r. Since the subgame does not cut information sets, I is contained in the subgame. If $h \in I$ and $h \neq r$, then $h \prec r$ would contradict the minimality of r, while $r \prec h$ would make the history h visit I twice, contradicting Lemma 1.1. If h is incomparable with r, then the subgame has a minimal history on the path to h distinct from r, again a contradiction. Hence $I = \{r\}$. Thus the induced root has belief one, as does $\{r\}$ in the original game.

Sequential rationality transfers immediately. Any deviation inside the subgame can be extended to the whole game by leaving behavior unchanged outside the subgame. Because the subgame is closed under succession and preserves information sets, the continuation problem at every non-root information set of the subgame is exactly the corresponding continuation problem in the original game under ϕ. At the induced root, the same argument applies to the original singleton

information set $\{r\}$. Hence a profitable deviation in the subgame would give a profitable deviation in the full game.

The belief conditions transfer for the same reason. If I, I' are non-root information sets of the induced subgame and $a \in A^{I'}$, then

$$h \in I_{\succ I'} \quad \Longleftrightarrow \quad \phi(h) \in \phi(I)_{\succ \phi(I')},$$

and similarly for $I_{\succeq I' \times \{a\}}$. Moreover, if $h_{I'}$ is the predecessor of h in the subgame, then

$$\phi(h)_{\phi(I')} = \phi(h_{I'}), \qquad \sigma'(h \setminus h_{I'}) = \sigma(\phi(h) \setminus \phi(h_{I'})),$$

and $\mu'(h) = \mu(\phi(h))$. Hence every instance of (2.1) and (2.5) in the subgame is the corresponding instance in the original game. If the conditioning information set is the induced root, the same conclusion follows by comparing it with the original singleton information set $\{r\}$ and using the fact that both assign probability one to the root.

Thus (σ', μ') is sequentially rational and satisfies Bayes' rule whenever possible and no signaling what you don't know in the subgame, hence is a simple perfect Bayesian equilibrium of the subgame. □

Proof of Proposition 3.2. Let r be the unique minimal history of the subgame, and let $\phi(k) := (r, k)$ be the canonical embedding of induced-subgame histories into the original game. For a set J of induced-subgame histories, write $\phi(J) := \{\phi(k) : k \in J\}$. Let (σ', μ') be the restriction of the perfect Bayesian equilibrium (σ, μ). If r is terminal, the claim is vacuous, so suppose that r is nonterminal.

First, the original information set containing r is the singleton $\{r\}$. Let I be that information set. Since the subgame does not cut information sets, $I \subseteq H'$. If $h \in I \setminus \{r\}$, then $h \prec r$ contradicts the minimality of r, while $r \prec h$ makes the history h visit I twice, contradicting Lemma 1.1. If h is incomparable with r, the first prefix of h that belongs to H' is a minimal history of the subgame distinct from r, again a contradiction. Hence $I = \{r\}$, so $\mu(r) = 1 = \mu'(\emptyset)$. Thus $\phi(J)$ is an original information set for every information set J of the induced subgame, including the root.

Sequential rationality transfers exactly as in Proposition 2.2. Fix a player information set J of the induced subgame and a deviation τ_i' by its moving player. Extend τ_i' to the original game by using τ_i' on subgame actions and σ_i outside the subgame. Since the subgame is closed under successors and does not cut information sets, continuation terminal histories, beliefs, continuation action probabilities, and payoffs after J coincide with those after $\phi(J)$. Therefore a profitable deviation in the induced subgame would give a profitable deviation in the original game, contradicting sequential rationality of (σ, μ).

It remains to verify cross-pair independence. Fix information sets J_1, J_2 of the induced subgame and histories

$$k_\ell, \hat{k}_\ell \in J_\ell, \qquad \ell = 1, 2,$$

such that the induced-subgame instance of (3.1) is not excluded by the 0/0 convention. Then $\phi(k_\ell), \phi(\hat{k}_\ell) \in \phi(J_\ell)$, and

$$\mu'(k) = \mu(\phi(k))$$

for every induced-subgame information-set history k, including the root.

Let R be the multiset of actions in r. For every subgame history k, the multiset of actions in $\phi(k)$ is $R \sqcup k$. Hence, by multiplicity-wise cancellation,

$$(\phi(\hat{k}_1) \sqcup \phi(\hat{k}_2)) \setminus (\phi(k_1) \sqcup \phi(k_2)) = (\hat{k}_1 \sqcup \hat{k}_2) \setminus (k_1 \sqcup k_2),$$

and similarly

$$(\phi(k_1) \sqcup \phi(k_2)) \setminus (\phi(\hat{k}_1) \sqcup \phi(\hat{k}_2)) = (k_1 \sqcup k_2) \setminus (\hat{k}_1 \sqcup \hat{k}_2).$$

Since σ' is the restriction of σ, the strategy-probability ratio in the induced subgame is exactly the corresponding original-game ratio. The belief ratio is also identical by $\mu'(k) = \mu(\phi(k))$. Therefore the corresponding original-game instance of (3.1) is not excluded by the 0/0 convention. Applying cross-pair independence in the original game gives

$$\frac{\mu'(\hat{k}_1)\mu'(\hat{k}_2)}{\mu'(k_1)\mu'(k_2)} = \frac{\sigma'\big((\hat{k}_1 \sqcup \hat{k}_2) \setminus (k_1 \sqcup k_2)\big)}{\sigma'\big((k_1 \sqcup k_2) \setminus (\hat{k}_1 \sqcup \hat{k}_2)\big)}.$$

Thus the restricted assessment satisfies cross-pair independence in the induced subgame. Together with sequential rationality, this proves that (σ', μ') is a perfect Bayesian equilibrium of the induced subgame. □

B Belief reasonableness in Fudenberg and Tirole (1991)

In this section, we express the belief restrictions underlying the belief-reasonableness definition in Fudenberg and Tirole (1991), for multi-stage games with observable actions and independent types, in the notation of this paper. This translation is involved: the notation that is convenient for multi-stage games with observable actions and independent types is often difficult to reconcile with the notation used for general extensive-form games. Similarly, the conditions for reasonableness are difficult to extend to settings with unobservable actions or information sets that are not sequentially ordered. For this reason, the exact relationship between generalizations of perfect

Bayesian equilibrium and the notion in Fudenberg and Tirole (1991) is often unclear. Because of the simplicity of Definition 3.1, however, we are able to show that cross-pair independence implies reasonableness. Furthermore, if reasonableness is supplemented with a mild and intuitive condition, the two concepts are equivalent.

B.1 Multi-stage games with observable actions

Each player $i \in N = \{1, \dots, |N|\}$ has a type θ_i drawn from a finite set Θ_i. Let $\theta = (\theta_i)_{i \in N}$ denote a type profile, that is, an element of $\Theta := \prod_{i \in N} \Theta_i$. As usual, we write $\Theta_{-i} := \prod_{j \neq i} \Theta_j$ and $\Theta_{-(i,j)} := \prod_{\ell \notin \{i,j\}} \Theta_\ell$. We let $\Delta(X)$ be the simplex of probability distributions on a finite set X. For each player i, let $\pi_i \in \Delta(\Theta_i)$ denote the prior over types; under the standing convention on nature's move, $\pi_i(\theta_i) > 0$ for every $\theta_i \in \Theta_i$. Types are independent:

$$\pi(\theta) = \prod_{i \in N} \pi_i(\theta_i).$$

We assume that the type sets are pairwise disjoint, that is, $\Theta_i \cap \Theta_j = \emptyset$ for all $i \neq j$. Thus, for the multiset calculations in this section, the components of the initial chance move are treated as formal primitive chance factors and can be identified with the type labels themselves. Equivalently, an initial chance action selecting the whole type profile is decomposed into the independent formal factors $\{\theta_i : i \in N\}$, with $\sigma(\theta_i) = \pi_i(\theta_i)$. All bracketed strategic action labels introduced below are distinct from all type labels.

The game unfolds in stages. In each stage t, the players move simultaneously. In the underlying extensive-form representation we order the stage-t moves as $1, 2, \dots, |N|$, and player i does not observe the stage-t actions of players $1, \dots, i-1$ before choosing.

A *public history at stage t* is a tuple $a^t = (a_1, \dots, a_{t-1})$ where, for every $1 \leq \tau < t$, we have $a_\tau = (a_{\tau,1}, \dots, a_{\tau,|N|})$, $a^\tau := (a_1, \dots, a_{\tau-1})$, and $a^1 := \emptyset$. For every public history a^t, player i's set of feasible *public* actions is denoted by $A_i(a^t)$, and

$$A(a^t) := \prod_{i \in N} A_i(a^t).$$

The set of feasible public histories at stage t is therefore

$$H^t_{\text{pub}} := \left\{ a^t = (a_1, \dots, a_{t-1}) \,\middle|\, a_\tau \in A(a^\tau) \text{ for every } \tau = 1, \dots, t-1 \right\}.$$

To keep the section compatible with the standing convention that each primitive action is available at a unique information set, we distinguish public action labels from primitive extensive-

form actions. For each $\theta_i \in \Theta_i$ and each public history $a^t \in H^t_{\text{pub}}$, let

$$\widetilde{A}_{i,t}(\theta_i, a^t) := \left\{ [a_i]_{\theta_i, a^t} \mid a_i \in A_i(a^t) \right\}.$$

A history at player i's move in stage t is then identified with a tuple

$$(\theta, a^t, a_{t,1}, \ldots, a_{t,i-1}),$$

which abbreviates the primitive history

$$(\theta_1, \ldots, \theta_{|N|}, [a_{1,1}]_{\theta_1, a^1}, \ldots, [a_{1,|N|}]_{\theta_{|N|}, a^1}, \ldots, [a_{t,1}]_{\theta_1, a^t}, \ldots, [a_{t,i-1}]_{\theta_{i-1}, a^t}).$$

For every feasible pair (θ_i, a^t) with $\theta_i \in \Theta_i$ and $a^t \in H^t_{\text{pub}}$, let

$$I_{i,t}(\theta_i, a^t) = \left\{ (\theta, a^t, a_{t,1}, \ldots, a_{t,i-1}) \;\middle|\; \theta_{-i} \in \Theta_{-i},\ a_{t,j} \in A_j(a^t) \text{ for } j < i \right\}.$$

This is the set of tuples associated with player i's type θ_i and public history a^t; it is an actual stage-t information set exactly when stage t is reached after a^t. We therefore write

$$H^t_{\text{pub},+} \subseteq H^t_{\text{pub}}$$

for the set of *active public histories* at stage t, that is, the set of public histories after which stage t is actually reached in the induced extensive form. Equivalently, $a^t \in H^t_{\text{pub},+}$ if and only if $I_{i,t}(\theta_i, a^t)$ is a stage-t information set for every player i and type θ_i. Every non-singleton player information set in the induced extensive form is of this form. These active public histories are defined structurally in the induced extensive form and may be on or off the path of play under σ. If $a^t \in H^t_{\text{pub},+}$, then every prefix a^τ with $\tau \le t$ belongs to $H^\tau_{\text{pub},+}$.

Whenever $a^t \in H^t_{\text{pub},+}$ and $a_i \in A_i(a^t)$, we write $\sigma(a_i \mid \theta_i, a^t)$ for the probability of the primitive action $[a_i]_{\theta_i, a^t}$. Whenever $a^t \in H^t_{\text{pub},+}$ and

$$h = (\theta, a^t, a_{t,1}, \ldots, a_{t,i-1}) \in I_{i,t}(\theta_i, a^t),$$

let $\text{type}_j(h) := \theta_j$ for each $j \in N$, and write $\text{type}_{-i}(h) := (\text{type}_j(h))_{j \neq i}$. For $j \neq i$, define the marginal posterior

$$\mu(\theta_j \mid \theta_i, a^t) := \sum_{\substack{h \in I_{i,t}(\theta_i, a^t): \\ \text{type}_j(h) = \theta_j}} \mu(h) \quad \text{and} \quad \mu(\theta_{-i} \mid \theta_i, a^t) := \sum_{\substack{h \in I_{i,t}(\theta_i, a^t): \\ \text{type}_{-i}(h) = \theta_{-i}}} \mu(h).$$

B.2 Definition of belief reasonableness

We now state the belief restrictions used by Fudenberg and Tirole (1991). The next definition packages, in our notation, their common-belief, belief-independence, Bayes-updating, and "no signaling what you do not know" requirements, with the zero-probability strengthening used below. To keep the dependence on the public history explicit, we write the posterior of player i's type at an active public history $a^t \in H^t_{\text{pub},+}$ as

$$\mu_i(\cdot \mid a^t) \in \Delta(\Theta_i).$$

No such object is needed at terminal public histories, because there is then no continuation game.

Definition B.1. An assessment (σ, μ) for a multi-stage game with observable actions and independent types is *reasonable* if there exists a family of probability distributions

$$\mu_i(\cdot \mid a^t) \in \Delta(\Theta_i) \qquad (i \in N,\ a^t \in H^t_{\text{pub},+})$$

such that the following conditions hold.

(R1) *Common beliefs.* For every active public history $a^t \in H^t_{\text{pub},+}$, every player $i \in N$, every type $\theta_i \in \Theta_i$, and every history

$$h = (\theta, a^t, a_{t,1}, \dots, a_{t,i-1}) \in I_{i,t}(\theta_i, a^t),$$

we have

$$\mu(h) = \Big(\prod_{j \neq i} \mu_j(\theta_j \mid a^t)\Big) \prod_{j<i} \sigma(a_{t,j} \mid \theta_j, a^t). \tag{B.1}$$

(R2) *Bayes' rule whenever possible.* For every active public history $a^t \in H^t_{\text{pub},+}$, every player $i \in N$, and every public action profile $a_t \in A(a^t)$ such that the successor public history $a^{t+1} = (a^t, a_t)$ belongs to $H^{t+1}_{\text{pub},+}$, if

$$\sum_{\hat{\theta}_i \in \Theta_i} \mu_i(\hat{\theta}_i \mid a^t)\, \sigma(a_{t,i} \mid \hat{\theta}_i, a^t) > 0,$$

then

$$\mu_i(\theta_i \mid a^{t+1}) = \frac{\mu_i(\theta_i \mid a^t)\, \sigma(a_{t,i} \mid \theta_i, a^t)}{\sum_{\hat{\theta}_i \in \Theta_i} \mu_i(\hat{\theta}_i \mid a^t)\, \sigma(a_{t,i} \mid \hat{\theta}_i, a^t)} \qquad \text{for all } \theta_i \in \Theta_i. \tag{B.2}$$

Also, at the initial public history, $\mu_i(\cdot \mid \emptyset) = \pi_i$ for every $i \in N$.

(R3) *No signaling what you don't know.* For every active public history $a^t \in H^t_{\text{pub},+}$, every player $i \in N$, and every pair of public action profiles $a_t, \hat{a}_t \in A(a^t)$ with $a_{t,i} = \hat{a}_{t,i}$ such that the two

successor public histories $a^{t+1}=(a^t,a_t)$ and $\hat{a}^{t+1}=(a^t,\hat{a}_t)$ both belong to $H^{t+1}_{\text{pub},+}$, we have

$$\mu_i(\cdot\,|\,a^{t+1})=\mu_i(\cdot\,|\,\hat{a}^{t+1}).$$

Note that, whenever $a^t\in H^t_{\text{pub},+}$, summing (B.1) over the current-period actions of players $1,\ldots,i-1$ yields

$$\mu(\theta_{-i}\,|\,\theta_i,a^t)=\prod_{j\neq i}\mu_j(\theta_j\,|\,a^t),$$

and therefore

$$\mu(\theta_j\,|\,\theta_i,a^t)=\mu_j(\theta_j\,|\,a^t)\qquad\text{for every } j\neq i.$$

For our main result in this section, we strengthen reasonableness by adding the following condition.

Definition B.2. An assessment (σ,μ) for a multi-stage game with observable actions and independent types is *extra reasonable* if it is reasonable and the following condition holds.

(R4) *Cross-action consistency.* For every active public history $a^t\in H^t_{\text{pub},+}$, every player $i\in N$, and every pair of public action profiles $a_t,\hat{a}_t\in A(a^t)$ such that the two successor public histories $a^{t+1}:=(a^t,a_t)$ and $\hat{a}^{t+1}:=(a^t,\hat{a}_t)$ belong to $H^{t+1}_{\text{pub},+}$, we have, for all $\theta_i,\hat{\theta}_i\in\Theta_i$,

$$\frac{\mu_i(\theta_i\,|\,a^{t+1})}{\sigma(a_{t,i}\,|\,\theta_i,a^t)}\cdot\frac{\mu_i(\hat{\theta}_i\,|\,\hat{a}^{t+1})}{\sigma(\hat{a}_{t,i}\,|\,\hat{\theta}_i,a^t)}=\frac{\mu_i(\hat{\theta}_i\,|\,a^{t+1})}{\sigma(a_{t,i}\,|\,\hat{\theta}_i,a^t)}\cdot\frac{\mu_i(\theta_i\,|\,\hat{a}^{t+1})}{\sigma(\hat{a}_{t,i}\,|\,\theta_i,a^t)}\tag{B.3}$$

whenever neither side is the indeterminate form $0/0$.

We view (R4) as a mild and intuitive condition on beliefs. It follows from (R1)–(R3) except for assessments and active public histories at which, for some player i, at least two of i's types have zero public posterior probability and at least two of i's feasible current actions have zero likelihood under $\mu_i(\cdot\,|\,a^t)$ and σ. Hence, if every player with three or more types has at most two feasible actions at each information set, reasonableness and extra reasonableness coincide. Even in games where this exceptional configuration can arise, (R4) may have little additional bite once it is combined with sequential rationality.

To obtain some intuition for (R4), consider the game in Figure 6. For simplicity, some information sets and actions are omitted, but the example can be extended straightforwardly to a multi-stage game with observable actions and independent types. Consider the assessment depicted in the figure. After the first-period public action t, the public belief assigns probability zero to types θ_1 and $\hat{\theta}_1$, that is,

$$\mu_1(\theta_1\,|\,\mathrm{t})=\mu_1(\hat{\theta}_1\,|\,\mathrm{t})=0.$$

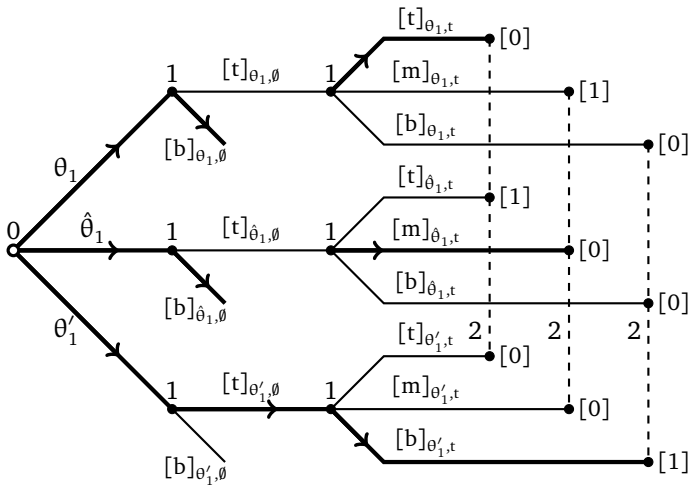


Figure 6

This assessment is reasonable (i.e., satisfies Definition B.1).[6] Nevertheless, the assessment is counterintuitive: in the second period, type θ_1 chooses t for sure and type $\hat{\theta}_1$ chooses m for sure, yet θ_1 is assigned lower probability after t than after m, while $\hat{\theta}_1$ is assigned lower probability after m than after t. Condition (R4) rules this out. Taking $a_{2,1} := \text{t}$, $\hat{a}_{2,1} := \text{m}$, $\theta_1 := \theta_1$, and $\hat{\theta}_1 := \hat{\theta}_1$, the left-hand side of (B.3) is 0, whereas the right-hand side is $+\infty$. More generally, (R4) requires the relative probabilities of player i's types after zero-likelihood actions to be disciplined by the same likelihood-ratio logic as Bayes' rule.[7]

B.3 Equivalence

Proposition B.1. *For any multi-stage game with observable actions and independent types, an assessment is extra reasonable if and only if it satisfies cross-pair independence.*

Proof of Proposition B.1. **Part 1: Extra reasonableness implies cross-pair independence.** Let (σ, μ) be extra reasonable, and fix $\{\mu_j(\cdot \mid a^t)\}$ satisfying (R1)–(R4).

If $|N| = 1$, then every information set in the induced extensive-form representation is a singleton. Hence, for every pair $h, \hat{h}$ in a common information set, $h = \hat{h}$, $\mu(h) = 1$, and both multiset differences in (3.1) are empty. Thus (3.1) is immediate. Hence assume $|N| \geq 2$.

---

[6] At this public history, player 1's second-period actions t and m have zero likelihood under $\mu_1(\cdot \mid \text{t})$ and σ. Thus (R2) does not pin down the posteriors following these actions; (R3) only identifies posteriors across public histories with the same current action of player 1.

[7] This requirement is similar to that imposed by cross-pair independence (note that (B.3) resembles (3.1) after reordering terms). Still, (R4) imposes cross-consistency across a narrow class of information sets (owned by the same player) and compares public histories instead of full histories.

For a player j, an active public history a^τ, and $\theta_j \in \Theta_j$, write

$$L_j(\theta_j \mid a^\tau) := \pi_j(\theta_j) \prod_{s=1}^{\tau-1} \sigma(a_{s,j} \mid \theta_j, a^s).$$

We first record the following consequence of (R2). For every active public history a^τ and every player j, either $L_j(\cdot \mid a^\tau) \equiv 0$ or there is a scalar $c_j(a^\tau) > 0$ such that

$$\mu_j(\cdot \mid a^\tau) = c_j(a^\tau) L_j(\cdot \mid a^\tau).$$

Consequently,

$$\mu_j(\theta_j \mid a^\tau) L_j(\hat{\theta}_j \mid a^\tau) = \mu_j(\hat{\theta}_j \mid a^\tau) L_j(\theta_j \mid a^\tau) \tag{B.4}$$

for all $\theta_j, \hat{\theta}_j \in \Theta_j$.

The proof is by induction on τ. At $\tau = 1$ this follows from (R2) and the full-support prior. Suppose it holds at a^τ, and let $a^{\tau+1} = (a^\tau, a_\tau)$ with $a_j := a_{\tau,j}$. If $L_j(\cdot \mid a^\tau) \equiv 0$, then $L_j(\cdot \mid a^{\tau+1}) \equiv 0$. Otherwise $\mu_j(\cdot \mid a^\tau) = c L_j(\cdot \mid a^\tau)$ for some $c > 0$. If

$$D := \sum_{\eta \in \Theta_j} \mu_j(\eta \mid a^\tau) \sigma(a_j \mid \eta, a^\tau)$$

is positive, (R2) gives

$$\mu_j(\theta_j \mid a^{\tau+1}) = \frac{c}{D} L_j(\theta_j \mid a^{\tau+1}).$$

If $D = 0$, then all terms in the preceding sum are zero and hence $L_j(\cdot \mid a^{\tau+1}) \equiv 0$. This proves the claim.

For $a_j \in A_j(a^t) \cup \{\varnothing\}$, define

$$\sigma_j(a_j \mid \theta_j, a^t) := \begin{cases} \sigma(a_j \mid \theta_j, a^t), & a_j \neq \varnothing, \\ 1, & a_j = \varnothing, \end{cases}$$

and

$$B_j(\theta_j, a^t, a_j) := \{\theta_j\} \sqcup \left\{ [a_{s,j}]_{\theta_j, a^s} : 1 \le s < t \right\} \sqcup \begin{cases} \{[a_j]_{\theta_j, a^t}\}, & a_j \neq \varnothing, \\ \emptyset, & a_j = \varnothing. \end{cases}$$

Also put

$$W_j(\theta_j, a^t, a_j) := \mu_j(\theta_j \mid a^t) \sigma_j(a_j \mid \theta_j, a^t).$$

A pair

$$x=(\theta_j,a^t,a_j),\qquad \hat{x}=(\hat{\theta}_j,a^t,\hat{a}_j)$$

is called admissible if either $a_j=\hat{a}_j=\varnothing$, or $a_j,\hat{a}_j\in A_j(a^t)$. For admissible triples write $B_j(x)$ and $W_j(x)$ in the evident way.

We shall use two elementary identities. First, every admissible pair satisfies

$$W_j(\hat{x})\,\sigma\big(B_j(x)\setminus B_j(\hat{x})\big)=W_j(x)\,\sigma\big(B_j(\hat{x})\setminus B_j(x)\big). \tag{B.5}$$

If $\theta_j=\hat{\theta}_j$, all past labels cancel and the assertion is immediate. If $\theta_j\neq\hat{\theta}_j$, the two type-specific blocks have no primitive label in common, and (B.5) is exactly (B.4), multiplied by the relevant current-action factors $\sigma_j(a_j\mid\theta_j,a^t)$ and $\sigma_j(\hat{a}_j\mid\hat{\theta}_j,a^t)$.

Second, if

$$x_r=(\theta_{jr},a_r^{t_r},a_{jr}),\qquad \hat{x}_r=(\hat{\theta}_{jr},a_r^{t_r},\hat{a}_{jr})\qquad (r=1,\dots,m)$$

are admissible pairs, with $m=0,1,2$, then

$$\prod_{r=1}^{m}W_j(\hat{x}_r)\,\sigma\left(\left(\bigsqcup_{r=1}^{m}B_j(x_r)\right)\setminus\left(\bigsqcup_{r=1}^{m}B_j(\hat{x}_r)\right)\right)=\prod_{r=1}^{m}W_j(x_r)\,\sigma\left(\left(\bigsqcup_{r=1}^{m}B_j(\hat{x}_r)\right)\setminus\left(\bigsqcup_{r=1}^{m}B_j(x_r)\right)\right). \tag{B.6}$$

The cases $m=0$ and $m=1$ are immediate from (B.5). For $m=2$, write $X_r:=B_j(x_r)$ and $Y_r:=B_j(\hat{x}_r)$. Multiplying the two one-coordinate identities gives

$$\prod_{r=1}^{2}W_j(\hat{x}_r)\,\sigma\big((X_1\setminus Y_1)\sqcup(X_2\setminus Y_2)\big)=\prod_{r=1}^{2}W_j(x_r)\,\sigma\big((Y_1\setminus X_1)\sqcup(Y_2\setminus X_2)\big).$$

For each primitive label, the elementary equality

$$(x_1-y_1)_+ +(x_2-y_2)_+ -(x_1+x_2-y_1-y_2)_+ =(y_1-x_1)_+ +(y_2-x_2)_+ -(y_1+y_2-x_1-x_2)_+$$

applied to its multiplicities implies that there is a common residual multiset R such that

$$(X_1\setminus Y_1)\sqcup(X_2\setminus Y_2)=\big((X_1\sqcup X_2)\setminus(Y_1\sqcup Y_2)\big)\sqcup R$$

and

$$(Y_1\setminus X_1)\sqcup(Y_2\setminus X_2)=\big((Y_1\sqcup Y_2)\setminus(X_1\sqcup X_2)\big)\sqcup R.$$

All type labels in R have positive probability, so they may be canceled. Positive-probability strategic

labels in R may also be canceled. It remains to justify deletion of common zero-probability strategic labels.

Order such labels by public time and choose one of maximal public time, say $[\alpha_j]_{\xi,a^\tau}$. The portions of the two relevant type-specific blocks after this label contain no later zero-probability residual label, so the induction hypothesis, together with cancellation of positive labels, reduces the deletion of this copy to a one-step comparison at the common predecessor a^τ. Let

$$a^{\tau+1} := (a^\tau, \alpha), \qquad \hat{a}^{\tau+1} := (a^\tau, \hat{\alpha})$$

be the two successor public histories that appear in the two diagonal blocks, and let ξ, ζ be the two relevant types. If $\alpha_j = \hat{\alpha}_j$, the needed local equality is

$$\mu_j(\xi \mid a^{\tau+1})\mu_j(\zeta \mid \hat{a}^{\tau+1}) = \mu_j(\zeta \mid a^{\tau+1})\mu_j(\xi \mid \hat{a}^{\tau+1}),$$

which follows from (R3). If $\alpha_j \neq \hat{\alpha}_j$, the needed local equality is

$$\begin{aligned}\mu_j(\xi \mid a^{\tau+1})\mu_j(\zeta \mid \hat{a}^{\tau+1})\sigma(\alpha_j \mid \zeta, a^\tau)\sigma(\hat{\alpha}_j \mid \xi, a^\tau) \\ = \mu_j(\zeta \mid a^{\tau+1})\mu_j(\xi \mid \hat{a}^{\tau+1})\sigma(\alpha_j \mid \xi, a^\tau)\sigma(\hat{\alpha}_j \mid \zeta, a^\tau),\end{aligned}$$

which is (R4). Multiplying by the common prefix factors and by the already-established suffix identities removes one zero-probability copy from R without division by its probability. Since the public tree is finite, iteration proves (B.6).

Now take two pairs of histories

$$h_1, \hat{h}_1 \in I_1, \qquad h_2, \hat{h}_2 \in I_2.$$

If one of the two information sets is a singleton, then the corresponding two histories coincide, have belief one, and contribute the same primitive multiset to the numerator and denominator. We may therefore discard that pair. Hence suppose both information sets are non-singleton. For $r = 1, 2$, write

$$I_r = I_{i_r, t_r}(\theta^r_{i_r}, a^{t_r}_r)$$

and

$$h_r = (\theta^r, a^{t_r}_r, a^r_{t_r,1}, \ldots, a^r_{t_r,i_r-1}), \qquad \hat{h}_r = (\hat{\theta}^r, a^{t_r}_r, \hat{a}^r_{t_r,1}, \ldots, \hat{a}^r_{t_r,i_r-1}).$$

For each $j \neq i_r$, define

$$x_{jr} := \begin{cases} (\theta_j^r, a_r^{t_r}, a_{t_r,j}^r), & j < i_r, \\ (\theta_j^r, a_r^{t_r}, \varnothing), & j > i_r, \end{cases}$$

and define $\hat{x}_{jr}$ analogously from $\hat{h}_r$.

By (R1),

$$\mu(h_1)\mu(h_2) = \prod_{j \in N} \prod_{\{r : j \neq i_r\}} W_j(x_{jr}), \qquad \mu(\hat{h}_1)\mu(\hat{h}_2) = \prod_{j \in N} \prod_{\{r : j \neq i_r\}} W_j(\hat{x}_{jr}).$$

The omitted coordinates $j = i_r$ are identical within the pair $(h_r, \hat{h}_r)$ and therefore cancel from both multiset differences. Since primitive labels belonging to different players are disjoint, applying (B.6) to each player j, with $m = |\{r : j \neq i_r\}|$, gives

$$\mu(\hat{h}_1)\mu(\hat{h}_2)\sigma\big((h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2)\big) = \mu(h_1)\mu(h_2)\sigma\big((\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2)\big). \tag{B.7}$$

Set

$$A := \mu(\hat{h}_1)\mu(\hat{h}_2), \qquad B := \mu(h_1)\mu(h_2),$$

and

$$C := \sigma\big((\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2)\big), \qquad D := \sigma\big((h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2)\big).$$

Then (B.7) says $AD = BC$. For nonnegative A, B, C, D satisfying $AD = BC$, the extended ratios A/B and C/D are equal whenever neither ratio is the indeterminate form $0/0$: if $B, D > 0$, divide by BD; if $B = 0$, then $A > 0$ because A/B is not $0/0$, so $D = 0$, and since C/D is defined we have $C > 0$, hence both ratios are $+\infty$; the case $D = 0$ is symmetric. Therefore, whenever neither side of (3.1) is $0/0$,

$$\frac{\mu(\hat{h}_1)\mu(\hat{h}_2)}{\mu(h_1)\mu(h_2)} = \frac{\sigma\big((\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2)\big)}{\sigma\big((h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2)\big)}.$$

This is (3.1). Hence the assessment satisfies cross-pair independence.

Part 2: Cross-pair independence implies extra reasonableness. Let (σ, μ) satisfy cross-pair independence. We construct, for every player j and every active public history a^t, a distribution $\mu_j(\cdot \mid a^t)$ satisfying (R1)–(R4).

We shall use two elementary facts. First, consider any instance of (3.1) and write

$$A := \mu(\hat{h}_1)\mu(\hat{h}_2), \qquad B := \mu(h_1)\mu(h_2),$$

$$C := \sigma((\hat{h}_1 \sqcup \hat{h}_2) \setminus (h_1 \sqcup h_2)), \qquad D := \sigma((h_1 \sqcup h_2) \setminus (\hat{h}_1 \sqcup \hat{h}_2)).$$

If C/D is not the indeterminate form $0/0$, then

$$AD = BC. \tag{B.8}$$

Indeed, if $A = B = 0$, this is immediate. Otherwise the left-hand side of (3.1) is defined, and cross-pair independence gives $A/B = C/D$, which is equivalent to (B.8), including the cases in which the common value is 0 or $+\infty$.

Second, we use the following factorization lemma.

Lemma B.1. *Let $X_1, \dots, X_m$ be finite nonempty sets, let $X := \prod_{r=1}^{m} X_r$, and let $\mu \in \Delta(X)$. Suppose that, for all $x, y, u, v \in X$ satisfying*

$$\{x_r, u_r\} = \{y_r, v_r\} \qquad (r = 1, \dots, m)$$

as multisets, we have

$$\mu(x)\mu(u) = \mu(y)\mu(v).$$

Then there exist $\mu_r \in \Delta(X_r)$ such that

$$\mu(x_1, \dots, x_m) = \prod_{r=1}^{m} \mu_r(x_r) \qquad ((x_1, \dots, x_m) \in X).$$

Proof. Choose $a \in X$ with $\mu(a) > 0$. For each r and $\xi \in X_r$, let $a^{r \leftarrow \xi}$ be obtained from a by replacing a_r by ξ, and put

$$\lambda_r(\xi) := \frac{\mu(a^{r \leftarrow \xi})}{\mu(a)}.$$

We claim that

$$\mu(x) = \mu(a) \prod_{r=1}^{m} \lambda_r(x_r) \qquad (x \in X).$$

This is proved by induction on the number of coordinates in which x differs from a. The claim is immediate when this number is 0 or 1. If x differs from a in $k \geq 2$ coordinates, choose one such coordinate r and put $x' := x^{r \leftarrow a_r}$. Applying the hypothesis to $x, a, x', a^{r \leftarrow x_r}$ gives

$$\mu(x)\mu(a) = \mu(x')\mu(a^{r \leftarrow x_r}),$$

and the induction hypothesis applied to x' gives the claim. Finally set $c_r := \sum_{\xi \in X_r} \lambda_r(\xi) > 0$ and $\mu_r(\xi) := \lambda_r(\xi)/c_r$. Summing the displayed factorization over X gives $1 = \mu(a) \prod_r c_r$, and hence

$\mu(x) = \prod_r \mu_r(x_r)$. □

We shall also use the following observation: if $p, r \in \Delta(X)$ and

$$p(\hat{x})r(x) = p(x)r(\hat{x}) \qquad (x, \hat{x} \in X),$$

then $p = r$. Indeed, choosing x^* with $p(x^*) > 0$ shows first that $r(x^*) > 0$ and then that $r = \lambda p$; since both distributions sum to one, $\lambda = 1$.

We first dispose of the one-player case. If $|N| = 1$, every information set in the induced extensive-form representation is a singleton, so (R1) is automatic. Identify a public action profile with its unique component and set

$$\mu_1(\cdot \mid \emptyset) := \pi_1.$$

Proceed recursively over active public histories. Suppose $\mu_1(\cdot \mid a^t)$ is defined. For every $a \in A_1(a^t)$ with $(a^t, a) \in H^{t+1}_{\text{pub},+}$, put

$$D_a := \sum_{\eta \in \Theta_1} \mu_1(\eta \mid a^t)\sigma(a \mid \eta, a^t).$$

If $D_a > 0$, define the successor distribution by Bayes' rule:

$$\mu_1(\eta \mid a^t, a) = \frac{\mu_1(\eta \mid a^t)\sigma(a \mid \eta, a^t)}{D_a}.$$

For successors with $D_a = 0$, choose a full-support distribution $r_{a^t} \in \Delta(\Theta_1)$ and put

$$E_a := \sum_{\eta \in \Theta_1} r_{a^t}(\eta)\sigma(a \mid \eta, a^t).$$

If $E_a > 0$, set

$$\mu_1(\eta \mid a^t, a) = \frac{r_{a^t}(\eta)\sigma(a \mid \eta, a^t)}{E_a};$$

if $E_a = 0$, choose $\mu_1(\cdot \mid a^t, a)$ arbitrarily.

Then (R2) holds by construction, and (R3) is tautological. To verify (R4), fix successor actions $a, \hat{a}$ and types $\theta, \hat{\theta}$. If $D_a, D_{\hat{a}} > 0$, both successor distributions are Bayesian updates from $\mu_1(\cdot \mid a^t)$, so (B.3) holds. If $D_a = D_{\hat{a}} = 0$ and $E_a, E_{\hat{a}} > 0$, the same argument applies with r_{a^t} in place of $\mu_1(\cdot \mid a^t)$. If either $E_a = 0$ or $E_{\hat{a}} = 0$, the corresponding action has zero probability under every type, because r_{a^t} has full support; both sides of (B.3) are therefore zero. Finally, suppose, without loss of generality, that $D_a > 0$ and $D_{\hat{a}} = 0$. Whenever $\sigma(\hat{a} \mid \eta, a^t) > 0$, the equality $D_{\hat{a}} = 0$ implies $\mu_1(\eta \mid a^t) = 0$, and hence the Bayesian successor after a assigns probability zero to η. If instead $\sigma(\hat{a} \mid \eta, a^t) = 0$, the

relevant side of (B.3) already contains a zero action factor. Applying this to $\eta=\theta$ and $\eta=\hat{\theta}$ shows that both sides are zero. Thus (R4) holds in the one-player case. Henceforth assume $|N|\geq 2$.

Step 1: factorization inside each information set. Fix an active public history a^t, a player i, and a type $\theta_i\in\Theta_i$. Histories in $I_{i,t}(\theta_i,a^t)$ are indexed by

$$X_i(a^t):=\Big(\prod_{j<i}(\Theta_j\times A_j(a^t))\Big)\times\Big(\prod_{j>i}\Theta_j\Big).$$

Let $h_i(x)$ denote the corresponding history and define

$$\mu_i^{t,\theta_i}(x):=\mu(h_i(x)).$$

If $x,y,u,v\in X_i(a^t)$ satisfy

$$\{x_r,u_r\}=\{y_r,v_r\}\qquad\text{for every coordinate } r,$$

then

$$h_i(x)\sqcup h_i(u)=h_i(y)\sqcup h_i(v)$$

as primitive-action multisets. Hence the right-hand side of (3.1), applied to the ordered pairs $(h_i(x),h_i(y))$ and $(h_i(u),h_i(v))$, is 1. By (B.8),

$$\mu_i^{t,\theta_i}(x)\mu_i^{t,\theta_i}(u)=\mu_i^{t,\theta_i}(y)\mu_i^{t,\theta_i}(v).$$

Lemma B.1 therefore gives distributions

$$m_{ij}^{t,\theta_i}\in\Delta(\Theta_j\times A_j(a^t))\quad(j<i),\qquad \mu_{ij}^{t,\theta_i}\in\Delta(\Theta_j)\quad(j>i),$$

such that

$$\mu_i^{t,\theta_i}(x)=\prod_{j<i}m_{ij}^{t,\theta_i}(x_j)\prod_{j>i}\mu_{ij}^{t,\theta_i}(x_j). \tag{B.9}$$

Step 2: separation of earlier movers' current actions. Fix $j<i$ and $\theta_j\in\Theta_j$, and define

$$\mu_{ij}^{t,\theta_i}(\theta_j):=\sum_{a\in A_j(a^t)}m_{ij}^{t,\theta_i}(\theta_j,a).$$

We prove that

$$m_{ij}^{t,\theta_i}(\theta_j,a)=\mu_{ij}^{t,\theta_i}(\theta_j)\sigma(a\mid\theta_j,a^t)\qquad(a\in A_j(a^t)). \tag{B.10}$$

If $\mu_{ij}^{t,\theta_i}(\theta_j)=0$, this is immediate. Otherwise choose $\bar{a}$ with $m_{ij}^{t,\theta_i}(\theta_j,\bar{a})>0$, and fix all other coordinates at values with positive one-dimensional factors in (B.9). Let $x(a)$ be the resulting point when the jth coordinate is (θ_j,a). Then $\mu_i^{t,\theta_i}(x(a))=Km_{ij}^{t,\theta_i}(\theta_j,a)$ for some $K>0$.

First, $\sigma(\bar{a}\mid\theta_j,a^t)>0$. If not, choose $\tilde{a}$ with $\sigma(\tilde{a}\mid\theta_j,a^t)>0$ and apply (3.1) to the ordered pairs

$$(h_i(x(\bar{a})),h_i(x(\tilde{a})))\quad\text{and}\quad(h_i(x(\bar{a})),h_i(x(\bar{a}))).$$

The right-hand side is $\sigma(\tilde{a}\mid\theta_j,a^t)/\sigma(\bar{a}\mid\theta_j,a^t)=+\infty$, whereas the left-hand side has a positive denominator and finite numerator, a contradiction.

Now fix a. Applying the same comparison with a in place of $\tilde{a}$ gives

$$m_{ij}^{t,\theta_i}(\theta_j,a)\sigma(\bar{a}\mid\theta_j,a^t)=m_{ij}^{t,\theta_i}(\theta_j,\bar{a})\sigma(a\mid\theta_j,a^t),$$

where, if $\sigma(a\mid\theta_j,a^t)=0$, the equality follows because cross-pair independence forces $m_{ij}^{t,\theta_i}(\theta_j,a)=0$. Summing over a yields

$$m_{ij}^{t,\theta_i}(\theta_j,\bar{a})=\mu_{ij}^{t,\theta_i}(\theta_j)\sigma(\bar{a}\mid\theta_j,a^t),$$

and substituting back proves (B.10).

Combining (B.9) and (B.10), every

$$h=(\theta,a^t,a_{t,1},\dots,a_{t,i-1})\in I_{i,t}(\theta_i,a^t)$$

satisfies

$$\mu(h)=\Big(\prod_{j<i}\sigma(a_{t,j}\mid\theta_j,a^t)\Big)\prod_{j\neq i}\mu_{ij}^{t,\theta_i}(\theta_j).\tag{B.11}$$

Step 3: the one-dimensional factors are common. Fix a player j and two information sets at the same active public history a^t, with observing players $i,i'\neq j$ and own types $\theta_i,\theta_{i'}$. We show that $\mu_{ij}^{t,\theta_i}=\mu_{i'j}^{t,\theta_{i'}}$.

Take $\theta_j,\hat{\theta}_j\in\Theta_j$. For each observer $\ell\in\{i,i'\}$ with $j<\ell$, choose actions $a_\ell,\hat{a}_\ell\in A_j(a^t)$ such that

$$\sigma(a_\ell\mid\theta_j,a^t)>0,\qquad\sigma(\hat{a}_\ell\mid\hat{\theta}_j,a^t)>0.$$

For observers ℓ with $j>\ell$, no current action of player j appears in $I_{\ell,t}(\theta_\ell,a^t)$; set the corresponding action factors equal to one. Compare, in $I_{i,t}(\theta_i,a^t)$, the ordered pair that changes player j's coordinate from θ_j to $\hat{\theta}_j$, and compare it with the reversed ordered pair in $I_{i',t}(\theta_{i'},a^t)$, keeping all other coordinates matched. The right-hand side of (3.1) is precisely the ratio of the positive current-action factors that appear in (B.11). Applying (B.8), substituting (B.11), summing over all

remaining coordinates, and canceling those positive action factors gives

$$\mu_{ij}^{t,\theta_i}(\hat{\theta}_j)\mu_{i'j}^{t,\theta_{i'}}(\theta_j)=\mu_{ij}^{t,\theta_i}(\theta_j)\mu_{i'j}^{t,\theta_{i'}}(\hat{\theta}_j).$$

By the elementary observation above, the two distributions are equal.

Since $|N|\geq 2$, for every j there exists at least one observer $i\neq j$. Define $\mu_j(\cdot\,|\,a^t)$ to be this common distribution. Then (B.11) becomes

$$\mu(h)=\Big(\prod_{j\neq i}\mu_j(\theta_j\,|\,a^t)\Big)\prod_{j<i}\sigma(a_{t,j}\,|\,\theta_j,a^t), \tag{B.12}$$

for every active public history a^t, every player i, every type θ_i, and every $h\in I_{i,t}(\theta_i,a^t)$. Thus (R1) holds.

Step 4: the initial prior. Fix $k\in N$ and choose $i\neq k$. Choose any $\bar{\theta}_i\in\Theta_i$. By Proposition 3.1, the maintained cross-pair independence hypothesis implies Bayes' rule whenever possible. The information set $I_{i,1}(\bar{\theta}_i,\emptyset)$ is on the path of play after the singleton root: all type priors have full support and, for each earlier mover, some action has positive probability. Hence Bayes' rule whenever possible gives, for

$$h=(\bar{\theta}_i,\theta_{-i},a_1,\ldots,a_{i-1})\in I_{i,1}(\bar{\theta}_i,\emptyset),$$

$$\mu(h)=\frac{\pi(\bar{\theta}_i,\theta_{-i})\prod_{j<i}\sigma(a_j\,|\,\theta_j,\emptyset)}{\sum_{\hat{\theta}_{-i},\hat{a}_1,\ldots,\hat{a}_{i-1}}\pi(\bar{\theta}_i,\hat{\theta}_{-i})\prod_{j<i}\sigma(\hat{a}_j\,|\,\hat{\theta}_j,\emptyset)}.$$

The denominator equals $\pi_i(\bar{\theta}_i)$. Summing over $a_1,\ldots,a_{i-1}$ yields

$$\mu(\theta_{-i}\,|\,\bar{\theta}_i,\emptyset)=\prod_{j\neq i}\pi_j(\theta_j).$$

By (B.12) at $t=1$,

$$\mu(\theta_{-i}\,|\,\bar{\theta}_i,\emptyset)=\prod_{j\neq i}\mu_j(\theta_j\,|\,\emptyset).$$

Marginalizing over all coordinates except k gives

$$\mu_k(\cdot\,|\,\emptyset)=\pi_k.$$

Step 5: Bayes' rule whenever possible. Fix an active public history a^t, a player j, and an action profile a_t such that

$$a^{t+1}:=(a^t,a_t)\in H^{t+1}_{\text{pub},+}.$$

Write

$$r^t(\eta) := \mu_j(\eta \mid a^t), \qquad r^{t+1}(\eta) := \mu_j(\eta \mid a^{t+1}), \qquad \sigma(\eta) := \sigma(a_{t,j} \mid \eta, a^t).$$

Assume $D := \sum_\eta r^t(\eta)\sigma(\eta) > 0$ and choose $\bar{\eta}$ with $r^t(\bar{\eta})\sigma(\bar{\eta}) > 0$.

If $j > 1$, fix $\theta_1 \in \Theta_1$. Compare the ordered pair in $I_{1,t+1}(\theta_1, a^{t+1})$

$$(\theta_1, \eta, z, a^{t+1}), \qquad (\theta_1, \bar{\eta}, z, a^{t+1})$$

with the reversed ordered pair in $I_{1,t}(\theta_1, a^t)$

$$(\theta_1, \bar{\eta}, \hat{z}, a^t), \qquad (\theta_1, \eta, \hat{z}, a^t).$$

The right-hand side of (3.1) is $\sigma(\bar{\eta})/\sigma(\eta)$, interpreted as $+\infty$ if $\sigma(\eta) = 0$. It is not $0/0$, so (B.8) applies. Summing over $z, \hat{z}$ and using (B.12) gives

$$r^{t+1}(\bar{\eta})r^t(\eta)\sigma(\eta) = r^{t+1}(\eta)r^t(\bar{\eta})\sigma(\bar{\eta}) \qquad (\eta \in \Theta_j).$$

If $j = 1$, fix $\theta_2 \in \Theta_2$. Because a^{t+1} is active, for every $\eta \in \Theta_1$ choose $b_\eta \in A_1(a^{t+1})$ with

$$\sigma(b_\eta \mid \eta, a^{t+1}) > 0.$$

Compare the ordered pair in $I_{2,t+1}(\theta_2, a^{t+1})$

$$((\eta, b_\eta, z), (\bar{\eta}, b_{\bar{\eta}}, z))$$

with the reversed ordered pair in $I_{2,t}(\theta_2, a^t)$

$$((\bar{\eta}, a_{t,1}, \hat{z}), (\eta, a_{t,1}, \hat{z})).$$

The stage-t action labels of player 1 cancel, and the right-hand side is the strictly positive ratio

$$\frac{\sigma(b_{\bar{\eta}} \mid \bar{\eta}, a^{t+1})}{\sigma(b_\eta \mid \eta, a^{t+1})}.$$

Applying (B.8), substituting (B.12), summing over $z, \hat{z}$, and canceling the positive auxiliary action probabilities gives the same identity:

$$r^{t+1}(\bar{\eta})r^t(\eta)\sigma(\eta) = r^{t+1}(\eta)r^t(\bar{\eta})\sigma(\bar{\eta}) \qquad (\eta \in \Theta_1).$$

Since $r^t(\bar{\eta})\sigma(\bar{\eta})>0$, the preceding identity implies $r^{t+1}(\bar{\eta})>0$ and hence

$$r^{t+1}(\eta)=c\,r^t(\eta)\sigma(\eta) \qquad (\eta\in\Theta_j)$$

for some $c>0$. Summing over η gives $c=1/D$. Therefore

$$\mu_j(\eta\mid a^{t+1})=\frac{\mu_j(\eta\mid a^t)\sigma(a_{t,j}\mid\eta,a^t)}{\sum_{\hat{\eta}}\mu_j(\hat{\eta}\mid a^t)\sigma(a_{t,j}\mid\hat{\eta},a^t)},$$

so (R2) holds.

Step 6: no signaling through other players' actions. Fix an active public history a^t, a player j, and two action profiles $a_t,\hat{a}_t\in A(a^t)$ such that $a_{t,j}=\hat{a}_{t,j}$ and

$$a^{t+1}:=(a^t,a_t), \qquad \hat{a}^{t+1}:=(a^t,\hat{a}_t)$$

are active successor public histories.

If $j>1$, compare, in the two player-1 successor information sets, the pair

$$(\theta_1,\theta_j,z,a^{t+1}), \qquad (\theta_1,\hat{\theta}_j,z,a^{t+1})$$

with the reversed pair

$$(\theta_1,\hat{\theta}_j,\hat{z},\hat{a}^{t+1}), \qquad (\theta_1,\theta_j,\hat{z},\hat{a}^{t+1}).$$

Because player j's stage-t public action is the same in the two successor histories, the right-hand side of (3.1) is 1. Summing the cross-product identity over $z,\hat{z}$ gives

$$\mu_j(\hat{\theta}_j\mid a^{t+1})\mu_j(\theta_j\mid\hat{a}^{t+1})=\mu_j(\theta_j\mid a^{t+1})\mu_j(\hat{\theta}_j\mid\hat{a}^{t+1}).$$

The elementary observation implies

$$\mu_j(\cdot\mid a^{t+1})=\mu_j(\cdot\mid\hat{a}^{t+1}).$$

If $j=1$, fix $\theta_2\in\Theta_2$ and choose

$$b_{\theta_1},b_{\hat{\theta}_1}\in A_1(a^{t+1}), \qquad \hat{b}_{\theta_1},\hat{b}_{\hat{\theta}_1}\in A_1(\hat{a}^{t+1})$$

with all four corresponding strategy probabilities strictly positive. Compare the ordered pair in

$I_{2,t+1}(\theta_2, a^{t+1})$

$$((\theta_1, b_{\theta_1}, z), (\hat{\theta}_1, b_{\hat{\theta}_1}, z))$$

with the reversed ordered pair in $I_{2,t+1}(\theta_2, \hat{a}^{t+1})$

$$((\hat{\theta}_1, \hat{b}_{\hat{\theta}_1}, \hat{z}), (\theta_1, \hat{b}_{\theta_1}, \hat{z})).$$

The common stage-t action of player 1 cancels. After applying (B.8), substituting (B.12), summing over $z, \hat{z}$, and canceling the four positive auxiliary action probabilities, we obtain the same cross-product identity for μ_1. Hence

$$\mu_1(\cdot \mid a^{t+1}) = \mu_1(\cdot \mid \hat{a}^{t+1}).$$

Thus (R3) holds.

Step 7: cross-action no signaling. Fix an active public history a^t, a player j, and two action profiles $a_t, \hat{a}_t \in A(a^t)$ such that

$$a^{t+1} := (a^t, a_t), \qquad \hat{a}^{t+1} := (a^t, \hat{a}_t)$$

are active successor public histories. We prove (B.3).

First suppose $j > 1$. Fix $\theta_1 \in \Theta_1$ and $\theta_j, \hat{\theta}_j \in \Theta_j$. Compare the ordered pair in $I_{1,t+1}(\theta_1, a^{t+1})$

$$(\theta_1, \theta_j, z, a^{t+1}), \qquad (\theta_1, \hat{\theta}_j, z, a^{t+1})$$

with the reversed ordered pair in $I_{1,t+1}(\theta_1, \hat{a}^{t+1})$

$$(\theta_1, \hat{\theta}_j, \hat{z}, \hat{a}^{t+1}), \qquad (\theta_1, \theta_j, \hat{z}, \hat{a}^{t+1}).$$

All primitive actions not belonging to player j cancel. Hence the right-hand side of (3.1) is

$$\frac{\sigma(a_{t,j} \mid \hat{\theta}_j, a^t)\sigma(\hat{a}_{t,j} \mid \theta_j, a^t)}{\sigma(a_{t,j} \mid \theta_j, a^t)\sigma(\hat{a}_{t,j} \mid \hat{\theta}_j, a^t)}.$$

If this ratio is $0/0$, then both action-products in (B.3) are zero and the desired equality is immediate. Otherwise, (B.8) applies. Substituting (B.12) and summing over $z, \hat{z}$ gives precisely (B.3).

Now suppose $j = 1$. Fix $\theta_2 \in \Theta_2$ and $\theta_1, \hat{\theta}_1 \in \Theta_1$. Choose

$$b_{\theta_1}, b_{\hat{\theta}_1} \in A_1(a^{t+1}), \qquad \hat{b}_{\theta_1}, \hat{b}_{\hat{\theta}_1} \in A_1(\hat{a}^{t+1})$$

with all four corresponding strategy probabilities strictly positive. Compare the ordered pair in $I_{2,t+1}(\theta_2, a^{t+1})$

$$((\theta_1, b_{\theta_1}, z), (\hat{\theta}_1, b_{\hat{\theta}_1}, z))$$

with the reversed ordered pair in $I_{2,t+1}(\theta_2, \hat{a}^{t+1})$

$$((\hat{\theta}_1, \hat{b}_{\hat{\theta}_1}, \hat{z}), (\theta_1, \hat{b}_{\theta_1}, \hat{z})).$$

The right-hand side of (3.1) is the product of the stage-t ratio in (B.3) and a strictly positive auxiliary stage-$(t+1)$ ratio. If the stage-t ratio is $0/0$, then both sides of (B.3) are zero. Otherwise the full right-hand side is not $0/0$, so (B.8) applies. Substituting (B.12), summing over $z, \hat{z}$, and canceling the four positive auxiliary action probabilities yields (B.3).

Thus (R4) holds. We have constructed a family $\{\mu_j(\cdot \mid a^t)\}$ satisfying (R1)–(R4), so the assessment is extra reasonable. □

References

BATTIGALLI, P. (1996): "Strategic independence and perfect Bayesian equilibria," *Journal of Economic Theory*, 70, 201–234.

BONANNO, G. (2013): "AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties," *International Journal of Game Theory*, 42, 567–592.

——— (2016): "AGM-consistency and perfect Bayesian equilibrium. Part II: from PBE to sequential equilibrium," *International Journal of Game Theory*, 45, 1071–1094.

DILMÉ, F. (2023): "A characterization of consistent assessments using power sequences of strategy profiles," *International Journal of Game Theory*, 1–21.

——— (2024): "Trembles by Nature, Stability, and Reputation," Mimeo.

FUDENBERG, D. AND J. TIROLE (1991): "Perfect Bayesian equilibrium and sequential equilibrium," *Journal of Economic Theory*, 53, 236–260.

KOHLBERG, E. AND P. J. RENY (1997): "Independence on relative probability spaces and consistent assessments in game trees," *Journal of Economic Theory*, 75, 280–313.

KREPS, D. M. AND G. RAMEY (1987): "Structural consistency, consistency, and sequential rationality," *Econometrica: Journal of the Econometric Society*, 1331–1348.

KREPS, D. M. AND R. WILSON (1982): "Sequential equilibria," *Econometrica*, 50, 863–894.

McLennan, A. (1989): "Consistent conditional systems in noncooperative game theory," *International Journal of Game Theory*, 18, 141–174.

Merotto, A. and A. Wolitzky (2026): "Conditional probability perfect Bayesian equilibrium," Mimeo.

Myerson, R. B. (1986): "Multistage games with communication," *Econometrica*, 323–358.

Perea y Monsuwe, A., M. Jansen, and H. Peters (1997): "Characterization of consistent assessments in extensive form games," *Games and Economic Behavior*, 21, 238–252.

Pimienta, C. (2014): "Bayesian and consistent assessments," *Economic Theory*, 55, 601–617.

Selten, R. (1965): "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit: Teil I: Bestimmung des dynamischen Preisgleichgewichts," *Zeitschrift für die gesamte Staatswissenschaft / Journal of Institutional and Theoretical Economics*, 121, 301–324.

Watson, J. (2025): "Perfect Bayesian Equilibrium: Consistency Conditions for Practical Definitions," *unpublished draft*.