Conditional Probability Perfect Bayesian Equilibrium

Anna Merotto* Alexander Wolitzky†
MIT MIT

March 3, 2026

Abstract

We propose a definition of perfect Bayesian equilibrium in extensive-form games that is consistent with a common conditional probability system on terminal nodes: *conditional probability perfect Bayesian equilibrium* (CPPBE). In a CPPBE, players (i) are sequentially rational, (ii) apply Bayes' rule both on and off path, and (iii) are free to form any common belief after zero-probability events. In particular, players may believe that past trembles are correlated. CPPBE refines subgame perfect equilibrium and weak perfect Bayesian equilibrium, but does not require "no-signaling-what-you-don't-know." We show that consistency with a conditional probability system is equivalent to consistency with *ranked probabilities*, which express both probabilities and lexicographic degrees of disbelief. This result provides a simple alternative characterization of CPPBE that is directly expressed by conditions on decision nodes and does not require explicitly calculating a conditional probability system.

Keywords: Perfect Bayesian equilibrium, conditional probability system, lexicographic probability, ranked probability, extensive-form games

JEL Codes: C70, C72

*MIT Department of Economics; email: anna_mer@mit.edu
†MIT Department of Economics; email: wolitzky@mit.edu
We thank Pierpaolo Battigalli, Emiliano Catonini, Laura Doval, Drew Fudenberg, Takuo Sugaya, and the participants at the MIT Theory lunch for helpful comments.

1 Introduction

Game theorists have proposed several different notions of perfect Bayesian equilibrium (PBE) in extensive-form games. These all agree in requiring sequential rationality: players choose optimally at every information set, given their beliefs. They also all agree that beliefs should be updated by Bayes' rule along the equilibrium path of play. But they involve very different consistency conditions on beliefs at off-path information sets.

Kreps and Wilson (1982) impose a strong consistency condition, which requires players to attribute zero-probability events to mistakes—or *trembles*—that are independent across information sets. In their sequential equilibrium concept, beliefs are derived by Bayes' rule from a sequence of positive, vanishing tremble probabilities. While sequential equilibrium is extremely influential and widely applied, the requirement that beliefs are derived from a sequence of independent trembles can make this solution concept both unduly restrictive (as the motivation for independence is unclear) and hard to work with (as it can be difficult to determine whether such a sequence exists) (Fudenberg and Tirole, 1991; Battigalli, 1996; Watson, 2025b).[1] Nonetheless, some consistency condition on off-path beliefs is necessary: the weakest version of PBE (*weak PBE*), which places no restriction on off-path beliefs, has little bite beyond Nash equilibrium, and in particular does not imply subgame perfection (Myerson, 1991; Mas-Colell et al., 1995).

The current paper proposes that a general, tractable notion of perfect Bayesian equilibrium requires only that the players' beliefs are consistent with the existence of a common conditional probability system (CPS): an array of conditional probabilities on terminal nodes linked by a consistency condition. We call the resulting solution concept *conditional probability perfect Bayesian equilibrium* (CPPBE). In a CPPBE, players update their beliefs using Bayes' rule after positive-probability events, both on- and off-path, but can form any belief after zero-probability events. In particular, players may believe that past trembles are correlated. We suggest that this solution concept offers a useful balance of restrictiveness (implying subgame perfection and weak PBE), flexibility regarding off-path beliefs (relaxing the independence requirement of sequential equilibrium), and tractability (admitting a simple characterization, discussed below).

In addition to our main definition of CPPBE based on CPS's, we provide two equivalent definitions which do not directly involve CPS's. First, a CPPBE is a sequentially rational assessment satisfying a consistency condition in the spirit of Kreps and Wilson (1982), but allowing correlated trembles. (Hence, CPPBE is more permissive than sequential

[1] Recent papers on constructing sequential equilibria include Dilmé (2024) and Watson (2025a).

equilibrium.) Second, a CPPBE is a sequentially rational assessment where beliefs are derived from *ranked probabilities* on nodes, where a ranked probability (r_x, p_x) specifies both the discrete *degree of disbelief* $r_x \in \mathbb{N}_{\geq 0}$ that node x will be reached and the probability $p_x \in (0,1]$ of reaching node x conditional on excluding theories of play with lower degrees of disbelief. The latter definition provides a simpler and more tractable characterization of CPPBE, as a ranked probability just consists of a pair of numbers (r_x, p_x) for each node. Establishing the validity of this characterization is our main result.

Allowing correlated trembles is appropriate in settings where players may not view others' deviations as independent. For instance, upon arriving at an off-path information set, players may entertain the possibility of unmodeled communication between their opponents, or unmodeled correlated shocks.[2] Consider, for example, the following game:

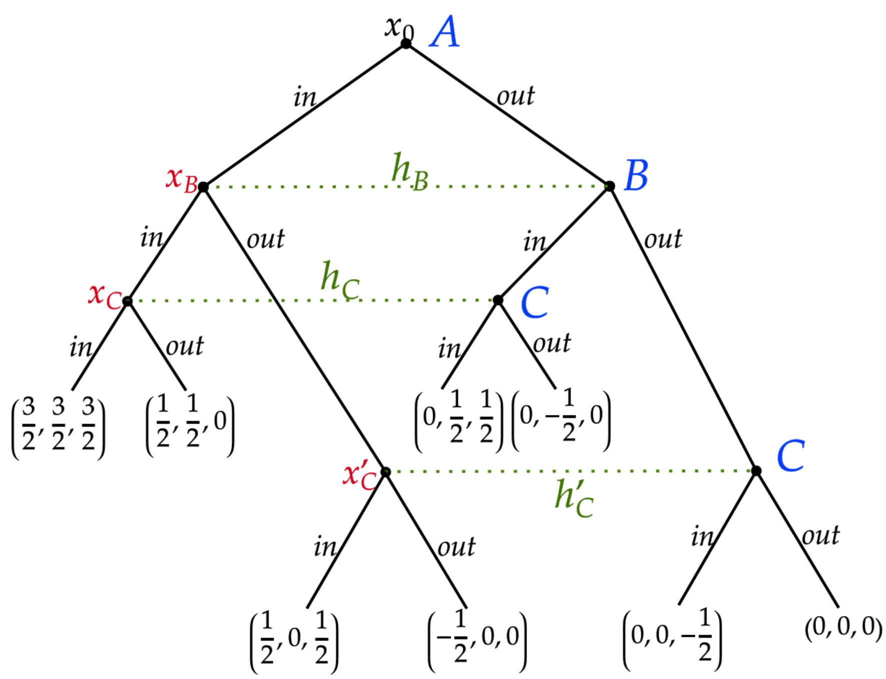

Figure 1: Investment game.

Here, three firms ($I = \{A, B, C\}$) must choose between investing ($a_i = in$) or not ($a_i = out$). Each firm gets a profit $\sum_{i \in I} \mathbf{1}_{a_i = in} - 3/2$ from investing, and 0 from not investing. The timing is that firms A and B (the "upstream market") simultaneously choose whether to invest or not, while C observes B's choice before acting. Suppose that in equilibrium all three firms invest, but C observes B not investing. In any sequential equilibrium, C

[2] If players attribute deviations to unmodeled shocks, one may ask whether these shocks can also affect other players' *future* play. In CPPBE, they cannot: players always believe their opponents will play their equilibrium strategies in the future. In contrast, Fudenberg et al. (1988) study equilibria where they can. The resulting solution concept is very permissive, and in particular does not refine subgame perfection.

continues to believe that A invested, and thus C chooses to invest as well. However, in a CPPBE, C can believe that if B deviated by not investing, then A may have done the same. For instance, C may suspect that a surprise shock has hit the upstream market, affecting both A's and B's decision. In this case, C may not invest after seeing B deviate.

Several related solution concepts lie in between weak PBE and sequential equilibrium, including *almost PBE* (Mailath, 2025), *plain PBE* (Watson, 2025b), and PBE satisfying *no-signaling-what-you-don't-know* (*NSWYDK*) (Fudenberg and Tirole, 1991) and *strategic independence* (Battigalli, 1996). In addition, Doval and Ely (2020) formulate a notion of PBE consistent with a CPS over strategy profiles (rather than terminal nodes), which we show is equivalent to CPPBE. Most of these concepts are ordered by strict inclusion, as shown in Figure 2.[3] The restrictions imposed by these solution concepts rule out different kinds of correlation in trembles. For example, NSWYDK implies that player i cannot view player j's trembles as correlated with information that neither i nor j hold, while strategic independence further implies that player i cannot view player j's trembles as correlated with information that only i holds.[4]

sequential equilibrium

$\subset$ strategically independent, sequentially rational assessment

$\subset$ PBE satisfying *NSWYDK*

$\subset$ PBE consistent with a CPS on S $=$ CPPBE

$\subset$ almost PBE

$\subset \underbrace{\left\{ \begin{matrix} \text{SPE} \\ \text{weak PBE} \end{matrix} \right\}}_{\text{non-nested}}$

Figure 2: Equilibria based on backward induction ordered by strict inclusion.

[3] The exception is plain PBE, which is non-nested with CPPBE, as shown in Appendix B. Of the strict inclusions shown in Figure 2, the first and second are shown by Battigalli (1996), the third is shown by our Theorem 2 and Observation 1, the fourth is shown by Doval and Ely (2020), and the last is shown by Mailath (2025). We also show that CPPBE and almost PBE coincide in multistage games with observed actions (Proposition 2)

[4] For example, modify the game in Figure 1 by assuming that B has three actions: not invest ($a_B = out$), low investment ($a_B = l$), and high investment ($a_B = h$), and suppose that in equilibrium $a_A = a_B = out$. Suppose also that A and B move simultaneously snd that C observes A's action and whether B has invested or not, but not the level of B's investment. Then, if B deviates, NSWYDK allows C to believe that B's action was correlated with A's, even if B has not observed it, while strategic independence does not. (See Battigalli (1996) for a similar example.)

The various solution concepts in Figure 2 can be more or less suited to different applications, depending on the tractability of applying them and the reasonableness of the implied belief restrictions. However, we believe CPPBE is a natural benchmark, for a couple reasons. First, it imposes no belief restrictions beyond consistency with a CPS, which in turn is implied by the PBE refinements most used in applications to date, such as NSWYDK and sequential equilibrium. Second, CPPBE is intuitively the weakest possible equilibrium concept consistent with sequential rationality and backward induction where Bayes' rule is always applied after positive-probability events.[5]

Sugaya and Wolitzky (2021) previously introduced CPPBE in multi-stage games with mediated communication, as a formalization of the non-cooperative solution concept implicit in Myerson (1986). However, their definition is complicated by issues specific to mediated games, such as the need to specify the set of possible messages from the mediator (the "mediation range") as part of the equilibrium. Sugaya and Wolitzky (2021) establish a revelation principle for CPPBE, but do not analyze general properties of CPPBE or relate it to other PBE notions in unmediated extensive-form games.

Our characterization of CPPBE in terms of ranked probabilities is reminiscent of prior results that establish equivalences between CPS's and lexicographic probability systems (Blume et al., 1991a,b; Hammond, 1994; Halpern, 2010). However, in contrast to our characterization, these results connect CPS's and lexicographic probability systems defined on the same domain, and they do not characterize when an assessment is consistent with a conditional or lexicographic probability system. We are not aware of prior work that uses our formulation of ranked probabilities, although it is easy to see that there is an equivalence between ranked probabilities and lexicographic conditional probability systems as defined by Blume et al. (1991a). Our formulation can also be viewed as combining a ranking function expressing degrees of disbelief in the spirit of Spohn (1988) together with conditional probabilities.

The rest of the paper is organized as follows. Section 2 introduces CPPBE, its properties, and its relation to sequential equilibrium, weak PBE, and subgame perfection. Section 3 presents our main result, which provides a practical characterization of CPPBE based on ranked probabilities rather than CPS's. Section 4 shows that is equivalent to define CPPBE via a CPS over either terminal nodes or strategy profiles. Finally, Section 5 relates

[5] This observation can be formalized, building on Catonini and Penta (2025). In particular, it can be shown that their "backwards rationalizability" concept together with a common prior is equivalent to a correlated equilibrium version of CPPBE. We have opted not to include this result in the current paper because the connection is somewhat tenuous, as Catonini and Penta focus on games with observable actions and we do not allow correlating devices.

CPPBE to NSWYDK and almost PBE, while the relationship with plain PBE is explained in Appendix B.

2 Equilibrium Concept

2.1 Model

We consider extensive-form games of perfect recall with a finite number of players ($i \in I$) and a finite number of nodes ($x \in X$).[6] We do not explicitly introduce Nature, but it can be easily included as an additional player with singleton information sets.

Let $\succ$ and $\succeq$ denote respectively the strict and weak successor relation on X, and $\succ_0$ the direct successor relation. As usual, $Z \subseteq X$ is the set of terminal nodes, $\emptyset$ the initial history, and H the partition of $X \setminus Z$ representing the information set of the agent active at each node. Let $\mathcal{A}(h)$ denote the set of available actions at information set h, and fix payoffs $(u_i : Z \to \mathbb{R})_{i \in I}$.

Strategies and beliefs are defined as usual. For each agent i, a **behavior strategy** β_i is a function mapping each h where i is active into a distribution over $\mathcal{A}(h)$. Call $\mathcal{B}_i$ the set of all β_i, and let $\mathcal{B} := \times_{i \in I} \mathcal{B}_i$. We write $\beta(\cdot|h)$ to represent $\beta_i(\cdot|h)$ when i is active at h. Similarly, a pure **strategy** s_i maps each h where i is active into an element of $\mathcal{A}(h)$; S_i is the set of all s_i; and $S := \times_{i \in I} S_i$.

A **system of beliefs** is a function μ_0 mapping each h into a distribution over its nodes. We mostly follow the standard approach of assuming that all agents' beliefs come from the same system μ_0. However, this homogeneous beliefs assumption plays only a minor role in our analysis (in particular, it simplifies notation), and we explain how to relax it in Section 2.4.

For any node x, we call $h(x)$ the information set containing x, $S(x)$ the set of strategy profiles consistent with reaching x on path, and $Z(x)$ the set of terminal nodes that succeed x (with the convention that $Z(x) = \{x\}$ if $x \in Z$). With a slight abuse of notation, for any $Y \subseteq X$, we call $S(Y) := \bigcup_{x \in Y} S(x)$ and $Z(Y) := \bigcup_{x \in Y} Z(x)$. Moreover, we call $P^{\beta}(\cdot|x)$ the probability induced by β on X.[7]

[6] The definition of CPPBE extends straightforwardly to infinite games, but equilibrium existence becomes a subtle issue, as for any refinement of subgame perfection (Harris et al., 1995).

[7] Given $y, y' \in X$ with $y \prec_0 y'$, let $a_{y,y'} \in \mathcal{A}(h(y))$ be the action leading from y to y'. For any sequence $x \prec_0 x_1 \prec_0 \ldots \prec_0 x_m \prec_0 x'$ in X, define $P^{\beta}(x'|x) := \beta(a_{x,x_1}|h(x)) \ldots \beta(a_{x_m,x'}|h(x_m))$.

2.2 Solution Concept

In any PBE concept, the first requirement for an assessment (β, μ_0) to be an equilibrium is sequential rationality. For any node x, let $V^\beta(x) = (V_i^\beta(x))_{i\in I}$ be the vector of continuation values when β is played from x onward. An assessment (β, μ_0) is **sequentially rational** at information set h with active player i if

$$\sum_{x\in h} \mu_0(x|h)\, V_i^\beta(x) \geq \sum_{x\in h} \mu_0(x|h)\, V_i^{(\beta_i', \beta_{-i})}(x) \qquad \text{for all } \beta_i' \in \mathcal{B}_i.$$

Then, (β, μ_0) is sequentially rational if it is sequentially rational at any $h \in H$.

The second equilibrium requirement is a form of consistency between the strategy profile β and the belief system μ_0. Our consistency notion refers to conditional probability systems (CPS's) (Rényi, 1955; Myerson, 1986). Given a finite set Ω, a CPS on Ω is a function $\mu : 2^\Omega \times 2^\Omega \setminus \{\emptyset\} \to [0,1]$ such that, for any $A, B, C \subseteq \Omega$ with $C \neq \emptyset$, we have (i) $\mu(C|C) = \mu(\Omega|C) = 1$; (ii) if $A \cap B = \emptyset$ then $\mu(A \cup B|C) = \mu(A|C) + \mu(B|C)$; and (iii) if $A \subseteq B \subseteq C$ and $B \neq \emptyset$, then $\mu(A|B) \cdot \mu(B|C) = \mu(A|C)$. Condition (iii) is often called the chain rule.

Our consistency notion is as follows.

Definition 1. *An assessment (β, μ_0) is **consistent with a CPS μ on** Z if*

(1) for any $h \in H$ and $x \in h$,

$$\mu_0(x|h) = \mu(Z(x)|Z(h)), \qquad \text{and}$$

(2) for any $x, x' \in X$ with $x \prec x'$,

$$\mu_0(x|h(x)) > 0 \implies \mu(Z(x')|Z(x)) = P^\beta(x'|x).$$

In words, point (1) requires that the belief system μ_0 is consistent with the CPS μ, and point (2) requires that μ is consistent with the strategy profile β, starting from any node with positive conditional probability.

Our solution concept is defined as follows.

Definition 2. *An assessment (β, μ_0) is a **conditional probability perfect Bayesian equilibrium (CPPBE)** if it is sequentially rational and consistent with a CPS on Z.*

The idea of requiring an assessment to be consistent with a CPS on Z—as in point (1) of

Definition 1—is due to Fudenberg and Tirole (1991). Following Myerson (1986), Sugaya and Wolitzky (2021) introduce CPPBE in the context of multi-stage games with mediated communication, which requires a more intricate definition.

For an illustration of CPPBE, consider the investment game in Figure 1. Sequential rationality implies that any CPPBE (β, μ_0) has $\beta(in|h_C) = 1$, since investing is strictly dominant for C once B has invested, and $\beta(in|\emptyset) = \beta(in|h_B) = 1$, by backward induction. By the definition, this implies $\mu_0(x_B|h_B) = \mu_0(x_C|h_C) = 1$. Sequential rationality also implies that $\beta(in|h'_C) = 1$ if and only if $\mu_0(x'_C|h'_C) \geq 1/2$, since after seeing B deviate, C invests only if she assigns probability at least 1/2 to A having invested. However, CPPBE places no restriction on C's beliefs at h'_C.[8] Thus, in a CPPBE, firm C can form any off-path belief, including one that justifies not investing after B deviates.

2.3 Definition via Trembles and Basic Properties

CPPBE can be equivalently defined using Myerson's (1986) characterization of CPS's. Myerson shows that, for any finite set Ω, a function $\mu : 2^\Omega \times 2^\Omega \setminus \{\emptyset\} \to [0,1]$ is a CPS if and only if there is a sequence of full-support distributions $(\eta_j) \in \Delta(\Omega)$ such that, for any $A, B \subseteq \Omega$ with $B \neq \emptyset$, $\mu(A|B) = \lim_j \eta_j(A|B)$. This implies the following characterization, which clarifies how CPPBE differs from sequential equilibrium.

Proposition 1. *A sequentially rational assessment (β, μ_0) is a sequential equilibrium if and only if there exists a sequence $(\beta_j, \eta_j)_{j \in \mathbb{N}}$ of full-support behavior strategies β_j and full-support distributions η_j over Z such that, $\forall h \in H$, $x \in h$, $x' \succ x$,*

(1) $\lim_{j \downarrow 0} \beta_j = \beta$*;*

(2) $\lim_{j \downarrow 0} \eta_j(Z(x)|Z(h)) = \mu_0(x|h)$*;*

(3) $\mu_0(x|h) > 0 \implies \lim_{j \downarrow 0} \eta_j(Z(x')|Z(x)) = P^\beta(x'|x)$*; and*

(4) $\forall j \in \mathbb{N}$, $\eta_j(Z(x)|Z(h)) = P^{\beta_j}(x)/P^{\beta_j}(h)$.

Moreover, a sequentially rational assessment (β, μ_0) is a CPPBE if and only if there exists such a sequence that satisfies (1), (2), and (3) (but not necessarily (4)).

Proposition 1 is a direct consequence of Definition 2 and Myerson's (1986) characterization. It shows that the difference between sequential equilibrium and CPPBE is that, in a sequential equilibrium, the converging distributions (η_j) are all required to be consistent with the behavior strategies (β_j), and thus assign independent probabilities to trembles at

[8] Moreover, for any $\mu_0(x'_C|h'_C)$, Theorem 1 assures that there exist a CPS μ on Z consistent with it.

different information sets. In contrast, in a CPPBE, the (η_j) must be consistent only with the limit behavior strategy β, and thus allow correlated trembles.

We can also mention some basic properties of CPPBE.

Relation to subgame perfection: In every finite extensive-form game of perfect recall, CPPBE is refined by sequential equilibrium (by Proposition 1) and refines subgame perfect equilibrium. The latter result follows because every CPPBE is an almost PBE (as proved in Section 5.2), and every almost PBE is subgame perfect (as proved by Mailath (2025)). For example, in the investment game, the unique sequential equilibrium prescribes that $\mu_0(x'_C|h'_C) = 1$ and all players always invest, while all subgame perfect equilibria (and all CPPBE) prescribe that all players always invest on path, but that C can play anything off path.

Existence and upper hemi-continuity: In every finite extensive-form game of perfect recall, a sequential equilibrium always exists, and thus a CPPBE. CPPBE is also upper hemi-continuous with respect to payoffs.[9]

2.4 Heterogeneous CPPBE

To conclude this section, we note that CPPBE can be extended to let players hold heterogeneous off-path beliefs. Besides generality, this extension is useful for relating CPPBE to prior PBE concepts that allow heterogeneous beliefs (see Appendix B).

We define a CPPBE with heterogeneous beliefs as follows.

Definition 3. *A sequentially rational assessment (β, μ_0) is a **heterogeneous CPPBE** if there exists a profile of CPS's $(\mu_i)_{i\in I}$ on Z such that*

(1) for any $h \in H$ where agent i is active, and any $x \in h$,

$$\mu_0(x|h) = \mu_i(Z(x)|Z(h)), \qquad \textit{and}$$

(2) for any $x, x' \in X$ with $x \prec x'$, for any agent $i \in I$,

$$\mu_i(Z(x)|Z(h(x))) > 0 \implies \mu_i(Z(x')|Z(x)) = P^{\beta}(x'|x).$$

Thus, in a heterogeneous CPPBE, each player's belief is derived from a personal CPS that is consistent with everyone's equilibrium strategies, but these CPS's can differ across

[9] The proof follows from standard arguments.

players.

To preview, while we focus on (homogeneous) CPPBE for simplicity, our main results easily extend to heterogeneous CPPBE. In particular, in the same vein as Theorem 1, heterogeneous CPPBE can be characterized in terms of heterogeneous ranked probabilities.

3 Characterization via Ranked Probabilities

PBE refinements are often criticized for their lack of analytical tractability. In particular, for refinements based on CPS's, a key difficulty is showing that a consistent CPS exists (Battigalli, 1996; Watson, 2025b). To do so, one possibility is applying Myerson's (1986) characterization, which requires finding a converging sequence of totally mixed probabilities whose limit point is consistent with the candidate equilibrium. Another possibility is guessing a CPS and then verifying that the chain rule is satisfied conditional on any subset of the relevant space (typically, Z or S). Often, neither of these approaches is straightforward to apply.

We present a simpler characterization of CPPBE in terms of ranked probabilities rather than CPS's. We show that ranked probabilities always generate beliefs that are consistent with some CPS, and that, conversely, any beliefs that are consistent with a CPS are generated by ranked probabilities.

Formally, a **ranked probability** associates a rank $r_x \in \mathbb{N}_{\geq 0}$ and a positive probability $p_x \in (0,1]$ to each node x, where $r_{x_0} = 0$ and $p_{x_0} = 1$. For any ranked probability vector $(r,p) \in (\mathbb{N}_{\geq 0} \times (0,1])^{|X|}$, we define an associated system of beliefs $\mu_0^{r,p}$ according to

$$\mu_0^{r,p}(x|h) := \lim_{\varepsilon \downarrow 0} \frac{p_x \, \varepsilon^{r_x}}{\sum_{y \in h} p_y \, \varepsilon^{r_y}} \qquad \text{for all } h \in H \text{ and } x \in h,$$

or equivalently

$$\mu_0^{r,p}(x|h) := \frac{p_x \mathbf{1}\{r_x = \min_{x' \in h} r_{x'}\}}{\sum_{y \in h} p_y \mathbf{1}\{r_y = \min_{x' \in h} r_{x'}\}} \qquad \text{for all } h \in H \text{ and } x \in h.$$

Intuitively, r_x is the lexicographic "degree of disbelief" that node x will be reached (Spohn, 1988), and p_x is the probability of reaching node x relative to other nodes with the same degree of disbelief. (Note that $\mu_0^{r,p}$ is invariant to positive linear transformations of p—and more generally to positive linear transformations of $(p_x)_{x \in h}$ for any $h \in H$, fixing the rest of the vector p—so only the relative probabilities captured by p matter for beliefs.) An interpretation is that there is an ordered set of theories of how the game will be played, where, whenever the players observe an event with zero probability according to the current theory, they move to the next one. With this interpretation, r_x is the index of the theory that gives positive probability to reaching node x, and p_x is the probability of reaching node x under this theory.

A vector of ranked probabilities is a simpler object than a CPS on Z, as it is defined on individual decision nodes rather than sets of nodes. In particular, a vector (r,p) consists of $|X|$ ranks $r_x \in \mathbb{N}_{\geq 0}$ and $|X|$ probabilities $p_x \in (0,1]$, one for each node, while a CPS on Z consists of $2^{2|Z|-1}$ probabilities $\mu(A|B)$, one for each pair of subsets of terminal nodes $A \subseteq Z$, $B \subseteq Z \setminus \{\emptyset\}$.

Definition 4. *An assessment* (β, μ_0) *is* ***consistent with ranked probabilities*** (r,p) *if*

(1) $\mu_0 = \mu_0^{r,p}$,

(2) for any $x, x' \in X$ *with* $x \prec x'$ *we have* $r_x \leq r_{x'}$ *and, if* $\mu_0(x|h(x)) > 0$, *then*

$$\begin{aligned} &P^\beta(x'|x) > 0 \implies r_x = r_{x'} \quad \textit{and} \quad p_x P^\beta(x'|x) = p_{x'}, \qquad \textit{and} \\ &P^\beta(x'|x) = 0 \implies r_x < r_{x'}. \end{aligned}$$

Here, the first condition says that the belief system μ is generated by the ranked probabilities (r,p); and the second condition says that the ranked probabilities are consistent with the behavior strategy β, in that the probabilities p are updated in accordance with β where possible, and the ranks r reflect deviations from β.

Definition 4 is a formalization of the notion of "Bayes' rule where possible."[10] Informally, this notion states that beliefs can be arbitrary immediately following a deviation, but Bayes' rule must apply at subsequent information sets. However, this informal statement can be ambiguous, for example when there is no clear ordering of deviations (as in Figure 3) or when the first information set after a deviation is a singleton (as in Figure 8 later in the paper.) Definition 4 resolves this ambiguity by requiring that ranked probabilities propagate along paths of play in the sense of condition (2).[11]

Our characterization of CPPBE in terms of ranked probabilities is as follows.

[10] We thank Laura Doval for suggesting this interpretation.

[11] For instance, an important subtlety is that condition (2) applies even if $P^\beta(x'|x) > 0$ but $\mu_0(x'|h(x')) = 0$, which can occur if $h(x')$ is also on-path from nodes with lower rank than x.

Theorem 1. *An assessment (β, μ_0) is consistent with a CPS on Z if and only if it is consistent with a vector of ranked probabilities.*

Consequently, an assessment (β, μ_0) is a CPPBE if and only if it is sequentially rational and is consistent with a vector of ranked probabilities.

Theorem 1 clarifies what it means for an assessment to be consistent with a CPS: namely, this holds if and only if beliefs are consistent with a vector of ranked probabilities. As we will see, this characterization simplifies the task of checking whether a given assessment is a CPPBE. Intuitively, this is because a vector of ranked probabilities is a much simpler object than a CPS.

We sketch the proof of Theorem 1. For the "only if" direction, given an assessment (β, μ_0) consistent with a CPS μ on Z, we define the **null operator** $K : 2^X \setminus \{\emptyset\} \to 2^X$ by

$$K(Y) := \left\{x \in Y : \mu(\{x\}|Y) = 0\right\} \qquad \text{for all } Y \subseteq X,\ Y \neq \emptyset.$$

We also define the iterates of this operator by letting K^0 be the identity function and letting $K^m(Y) := K(K^{m-1}(Y))$ for any $m \geq 1$. In the appendix, we verify that, for any node x, there is a unique number m_x such that $\mu(\{x\}|K^{m_x}(X)) > 0$. We then define the ranked probabilities

$$r_x = m_x \quad \text{and} \quad p_x = \mu(\{x\}|K^{r_x}(X)), \qquad \text{for all } x \in X.$$

Finally, we show that the hypothesis that (β, μ_0) is consistent with μ and μ satisfies the chain rule implies that (β, μ_0) is also consistent with (r, p).

For the "if" direction, given an assessment (β, μ_0) consistent with ranked probabilities (r, p), we define the function $\mu : 2^Z \times 2^Z \setminus \{\emptyset\} \to [0,1]$ by

$$\mu(A|B) = \frac{\sum_{z \in A \cap B} p_z \mathbf{1}\{r_z = \min_{z' \in B} r_{z'}\}}{\sum_{z \in B} p_z \mathbf{1}\{r_z = \min_{z' \in B} r_{z'}\}} \qquad \text{for all } A \subseteq Z,\ B \subseteq Z \setminus \{\emptyset\}.$$

We then show that the hypothesis that (β, μ_0) is consistent with (r, p) implies that μ is a CPS and that it is consistent with (β, μ_0).

The example in Figure 3 illustrates that different ranked probabilities can be consistent with the same assessment, and shows that a node's rank can depend both on the "number of deviations" required to reach it and on the probability of these deviations. Suppose agents play the behavior strategy $\beta = out$, indicated in blue. Then, if (β, μ_0) is consistent with (r, p), we must have $r_{z_1} = 0$ and $r_x > 0$ for all $x \neq x_0, z_1$, because x_0 and z_1 are the only on-path nodes under β. Moreover, we must also have $r_{x_C} > r_{x_B}$, because $P^\beta(x_C|x_B) = 0$. However, the nodes x_C and y_C can be ranked in multiple ways: for example, we can have either $r_{x_B} = r_{y_C} = 1$ and $r_{x_C} = 2$, or $r_{x_B} = 1$ and $r_{x_C} = r_{y_C} = 2$. In the first ranking, the agents' "second theory of play" is that A mixes between l or r—so, at h_C, agent C has observed that the first (on-path) theory was wrong and has moved to the second one, and believes that y_C has been reached. Alternatively, in the second ranking, the agents' second theory is that A plays l and B plays out, and the third theory is that either A plays l and B plays c, or A plays r. Now, at h_C, agent C has observed that the first two theories were both wrong and has moved to the third one, under which either x_C or y_C may have been reached. Note that in the second ranking the nodes x_C and y_C have the same rank, even though reaching x_C requires more deviations than reaching y_C.

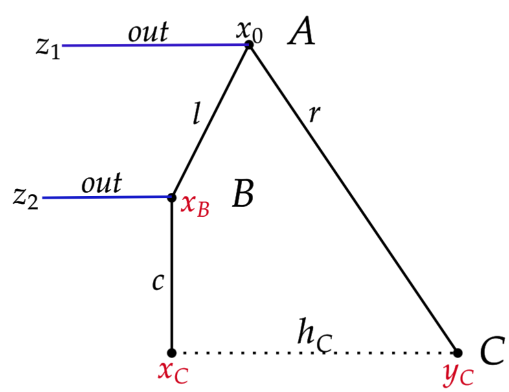

Figure 3: An example where r_x can take different values.

We illustrate the utility of Theorem 1 using the examples in Figures 4 and 5. In Figure 4, four players play a strategy profile β where A plays l, B plays out, and C has only one feasible action. We assume that B and D cannot observe A's choice, and C can only observe

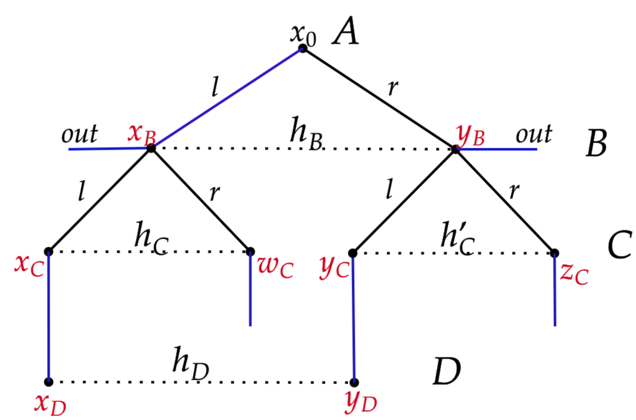

Figure 4: An example where checking if an assessment is a CPPBE is non-trivial.

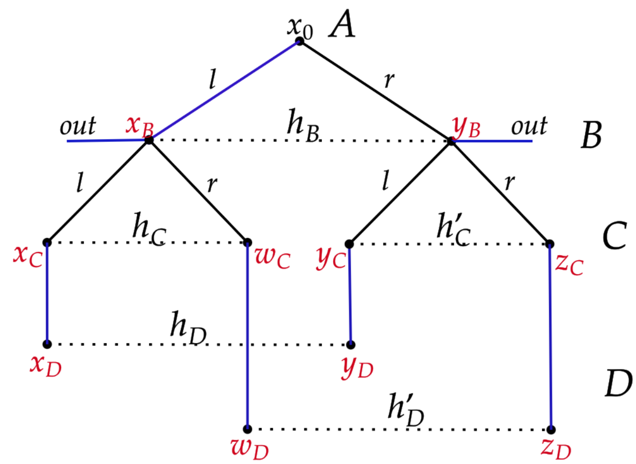

Figure 5: An extension of the example.

whether B chooses out or not. For simplicity, we do not represent payoffs or D's actions; however, we will assume that β is sequentially rational.

Suppose B's beliefs are $\mu_0(x_B|h_B) = 1$, and consider three candidate belief systems for players C and D: (i) $\mu_0(x_C|h_C) = \mu_0(y_C|h'_C) = 1/4$ and $\mu_0(x_D|h_D) = 1$; (ii) $\mu_0(x_C|h_C) = 1/4$, $\mu_0(y_C|h'_C) = 3/4$, and $\mu_0(x_D|h_D) = 1$; and (iii) $\mu_0(x_C|h_C) = 1/4$, $\mu_0(y_C|h'_C) = 3/4$, and $\mu_0(x_D|h_D) = 1/3$.

It is straightforward to verify that assessment (i) (i.e., β together with belief system (i)) is a sequential equilibrium. However, assessments (ii) and (iii) are not are sequential equilibria, because sequential equilibrium requires that $\mu_0(x_C|h_C) = \mu_0(y_C|h'_C)$. In contrast, Theorem 1 can be used to verify that these assessments are both CPPBE.[12]

First, consider assessment (ii). Fix a vector (r,p) such that $0 = r_{x_0} = r_{x_B} < r_{y_B} = r_{x_C} = r_{w_C} = r_{x_D} < r_{y_C} = r_{z_C} = r_{y_D}$, $p_{x_0} = p_{x_B} = 1$, $p_{w_C} = 3p_{x_C} = 3p_{x_D}$, and $3p_{z_C} = p_{y_C} = p_{y_D}$. Then assessment (ii) is consistent with (r,p), which implies that it is a CPPBE by Theorem 1. Similarly, assessment (iii) is consistent with any vector (r,p) satisfying $0 = r_{x_0} = r_{x_B} < r_{y_B} = r_{x_C} = r_{w_C} = r_{x_D} = r_{y_C} = r_{z_C} = r_{y_D}$, $p_{x_0} = p_{x_B} = 1$, and $6p_{x_C} = 6p_{x_D} = 2p_{w_C} = 3p_{y_C} = 3p_{y_D} = 9p_{z_C}$.

In the example in Figure 5, we modify the example in Figure 4 by assuming that D always moves after C. Consider extending assessment (iii) above with either of two candidate beliefs for D at h'_D: $\mu_0(w_D|h'_D) = 9/11$ and $\mu_0(w_D|h'_D) = 5/6$. At first glance, it is not obvious which of these beliefs (if any) is consistent with CPPBE. However, Theorem 1 implies that only the first one is. To see this, note that in any vector (r,p) that is consistent with either belief (in the sense of Definition 4), $2p_{w_C} = 9p_{z_C}$. Therefore, if D never observes

[12] Moreover, (ii), but not (iii), satisfies no-signaling-what-you-don't-know, as we will discuss.

A's choice—as in Figure 5—and B plays r, the only possible equilibrium belief for D is $\mu_0(w_D|h'_D) = 9/11$. Intuitively, immediately after B's deviation, player D must form some belief about A's action. Since all information sets are on path after B's deviation, the existence of a CPS implies that beliefs at h_D and h'_D must be consistent with the same conjecture about A's move.[13]

4 CPS's on Z vs. S

Several authors have developed PBE concepts that require consistency with a CPS on the set of strategy profiles, rather than the set of terminal nodes (Battigalli, 1996; Kohlberg and Reny, 1997; Doval and Ely, 2020; Watson, 2025b). The closest such paper is due to Doval and Ely (2020), who define a PBE as an assessment consistent with a CPS on S, arguing that this captures the idea of "Bayes' rule where possible." In this section, we show that Doval and Ely's solution concept is equivalent to ours, in the sense that for any PBE consistent with a CPS on S there exists an equivalent CPPBE, and vice versa.[14]

Despite this equivalence, CPS's on Z have a couple advantages over those on S. First, a CPS on Z is a reduced-form object relative to a CPS on S: any CPS on S induces a unique CPS on Z, but many CPS's on S induce the same CPS on Z. This relationship makes CPS's on Z simpler and easier to compute.[15] Second, CPS's on Z are defined on observable events (as noted by Battigalli (1996)).[16]

We denote a generic CPS on S with ν, and write $h' \not\prec x$ if there is no $x' \in h'$ such that $x' \prec x$. Recall that $S(x)$ is the set of strategy profiles consistent with reaching node x on path.

Definition 5 (Doval and Ely, 2020). *An assessment (β, μ_0) is a **PBE consistent with a CPS on S** if it is sequentially rational and there is a CPS ν on S such that*

(1) for any $h \in H$ and $x \in h$,

$$\mu_0(x|h) = \nu(S(x)|S(h)),$$

[13] This example suggests that, given a strategy profile β, Theorem 1 can also be used to check whether there exists a belief system μ_0 such that (β, μ_0) is a CPPBE. To do so, one must find a vector (r, p) such that the induced beliefs satisfy sequential rationality at each information set, which amounts to solving a system of linear inequalities for each non-singleton information set.

[14] Consistent with this equivalence, it is possible to construct ranked probabilities starting from a CPS on S in a similar way as we do for a CPS on Z in the proof of Theorem 1.

[15] This is roughly analogous to the computational advantage of behavior strategies over mixed strategies (Kuhn, 1953).

[16] At the same time, given a candidate assessment, Theorem 2 will imply that it is equivalent to verify its consistency with a CPS on S or Z.

(2) for any $x \in X$, $h \in H$ with $x \in h$,

$$\mu_0(x|h) > 0 \implies \forall s \in S, \quad \nu(s|S(x)) = \begin{cases} \prod_{h' \not< x} \beta(s(h')|h') & \text{if } s \in S(x), \\ 0 & \text{otherwise.} \end{cases}$$

As in Definition 1, point (1) requires that the belief system μ_0 is consistent with the CPS ν, and point (2) requires that the strategy profile β and ν are consistent, in that the relative weights under ν of the strategies $s \in S(x)$ are determined by their β-probabilities at those nodes where their behaviors differ (that is, at nodes that do not precede x).

For an example of what point (2) implies, consider the tree in Figure 6a. Fix β such that $\beta(r|x_0) = 1$ and $\beta(a|x) = \beta(e|y) = \frac{1}{3}$. Note that if node x is reached then the histories that have not yet been played are $\{x\}$ and $\{y\}$. So, in any PBE consistent with a CPS ν on S in which β is played, point (2) requires that $\nu((l, a.d)|S(x)) = \beta(a|x)\beta(d|y) = \frac{2}{9}$.

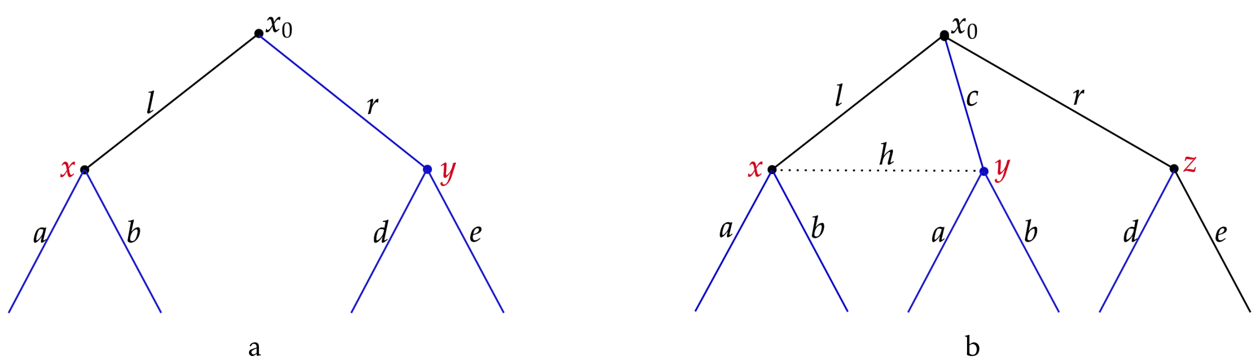

Figure 6: Differences between CPPBE and PBE consistent with a CPS on S.

We call two CPS's equivalent if they induce the same conditional probabilities on Z.

Definition 6. *Let μ be a CPS on Z and ν be a CPS on S. We say ν and μ are **equivalent** if*

$$\mu(A|B) = \nu(S(A)|S(B)) \qquad \text{for all } \forall A, B \subseteq Z,\ B \neq \emptyset.$$

This equivalence relation partitions the set of CPS's on S into equivalence classes: each CPS on Z is equivalent to a range of CPS's on S, but each CPS on S is equivalent to a unique CPS on Z.

The following theorem establishes the equivalence of CPPBE and PBE consistent with a CPS on S.

Theorem 2. *(1) Fix a PBE (β, μ_0) consistent with a CPS ν on S. Let μ be the CPS on Z that is equivalent to ν. Then, (β, μ_0) is a CPPBE consistent with μ.*

(2) Fix a CPPBE (β, μ_0) consistent with a CPS μ on Z. Then, there exists a CPS ν on S such that ν is equivalent to μ and (β, μ_0) is a PBE consistent with ν.

Part (1) is straightforward: since the CPS μ on Z that is equivalent to ν is unique, the proof only requires verifying that (β, μ_0) is consistent with μ. Part (2) requires more work, because we must first construct a CPS ν on S that is equivalent to μ. To do so, we rely on Myerson's (1986) characterization to fix a sequence of probabilities on Z converging to the CPS μ. We then use these distributions on Z to define a converging sequence of distributions on S, whose limit point ν is shown to be equivalent to μ and consistent with the assessment (β, μ_0).

Note that part (2) is a partial converse to part (1), as it applies to some CPS ν that is equivalent to μ, but not necessarily to every such CPS. To see this, consider the examples in Figure 6a and Figure 6b. In Figure 6a, let $\beta(r|x_0) = 1$ and $\beta(a|x) = \beta(e|y) = \frac{1}{3}$, and suppose that (β, μ_0) is a CPPBE. Since there is only one off-path move, it is easy to see that that there is a unique CPS μ on Z consistent with (β, μ_0). However, the class of CPS's on S that are equivalent to μ is

$$\Big\{\nu \text{ CPS on } S : \forall a_1 \in \{a,b\}, a_2 \in \{d,e\},\ \nu((l,a_1.a_2)|S) = \nu((l,a_1.a_2)|S(y)) = \nu((r,a_1.a_2)|S(x)) = 0,$$
$$\nu((r,a.d)|S) + \nu((r,b.d)|S) = \nu((r,a.d)|S(y)) + \nu((r,b.d)|S(y)) = \frac{2}{3},$$
$$\text{and} \quad \nu((l,a.d)|S(x)) + \nu((l,a.e)|S(x)) = \frac{1}{3}\Big\}.$$

However, consistency of ν with (β, μ_0) imposes some additional constraints, such as $\nu((l,a.d)|S(x)) = \frac{2}{9}$ (as shown above). Indeed, in this example, there are infinitely many CPS's on S equivalent to μ, but a unique one that can be part of a PBE consistent with a CPS on S where β is played.

In Figure 6b, let $\beta(c|x_0) = \beta(d|z) = 1$, $\beta(a|h) = \frac{1}{2}$, and suppose that (β, μ_0) is a CPPBE consistent with a CPS μ on Z. In any PBE (β, μ_0) consistent with a CPS on S, $\mu_0(x|h) = 0$. Thus, consistency of (β, μ_0) with a CPS ν on S does not impose any restrictions on $\nu(.|S(x))$, while equivalence with μ only requires that $\nu((l,a.d)|S(x)) + \nu((l,a.e)|S(x)) = \mu(Z(l,a)|Z(x))$, and similarly for $Z(l,b)$. Thus, there is a continuum of CPS's ν on S that are equivalent to μ and consistent with a PBE where β is played.

Figure 6b also illustrates the computational advantage of CPPBE over PBE consistent with a CPS on S. In this example, there is a unique CPS μ on Z consistent with a CPPBE where β is played, and μ must be defined on the set of terminal nodes of cardinality $|Z|= 6$. However, there are infinitely many CPS's ν on S consistent with a PBE where β, and ν must be defined on the set of strategy profiles of cardinality $|S|= 12$. Moreover, all of these CPS's on S lead to the same assessment (β, μ) as the CPPBE.

5 Relation to Similar Solution Concepts

In this section, we relate CPPBE to the most similar solution concepts: no-signaling-what-you-don't-know and almost PBE.

5.1 *No-Signaling-What-You-Don't-Know*

No-signaling-what-you-don't-know (NSWYDK, Fudenberg and Tirole, 1991) is a consistency condition that precludes an agent from viewing an opponent's past moves as being correlated with information that neither the opponent nor the player herself possesses. NSWYDK thus rules out some forms of correlation among agents' trembles, but it does not require full strategic independence (Battigalli, 1996). The appeal of NSWYDK varies in different settings. For instance, in the investment game in Figure 1, we may want to let firm C believe that the other firms' mistakes could be linked to a common shock, which is ruled out by NSWYDK. Moreover, in general games, checking whether an assessment is a CPPBE appears easier than checking whether it is a PBE satisfying NSWYDK due to Theorem 1, as discussed below.

In finite extensive-form games, NSWYDK requires beliefs to be consistent with a CPS on Z, along with some additional conditions that Fudenberg and Tirole call "general reasonableness".

Definition 7 (Fudenberg and Tirole, 1991)**.** *Fix an assessment (β, μ_0) and suppose μ_0 is consistent with a CPS μ on Z: that is, for any h and $x \in h$, $\mu_0(x|h) = \mu(Z(x)|Z(h))$. Then, (β, μ_0) is* ***generally reasonable*** *if*

(1) for any $x, x' \in X$ with $x \prec_0 x'$,

$$P^{\beta}(x'|x) = \mu(Z(x')|Z(x)),$$

(2) for any $h \in H$, $x, y \in h$, $x \prec_0 x'$ and $y \prec_0 y'$ such that $a_{x,x'} = a_{y,y'}$,[17]

$$\mu\Big(Z(x')\Big|Z(x') \cup Z(y')\Big) = \mu\Big(Z(x)\Big|Z(x) \cup Z(y)\Big).$$

An assessment whose belief system is consistent with a CPS on Z and satisfies condition (1) (but not necessarily (2)) coincides with what Battigalli (1996) terms a "sequentially rational tree-extended assessment." Note that this is already a stronger solution concept than CPPBE, since condition (1) is imposed also at nodes with zero conditional probability.[18]

A generally reasonable assessment is obtained by additionally imposing condition (2), which requires that belief updates immediately following an agent's deviation can depend only on the player's own information. In general, replacing the consistency condition in the definition of CPPBE with conditions (1) and (2) strengthens CPPBE to PBE satisfying NSWYDK.[19] We thus have:

Observation 1. *Any assessment that is sequentially rational and generally reasonable is a CPPBE.*

The converse is not true. For example, in the investment game in Figure 1, if invest is played on path, point (2) of general reasonableness implies that $\mu_0(x'_C|h'_C) = \mu_0(x_C|h_C)$. Intuitively, NSWYDK requires that firm C does not interpret B's deviation as correlated with A's choice, which neither C nor B can observe. Similarly, in the example in Figure 4, point (2) of general reasonableness implies that $\mu(Z(x_C)|Z(x_C) \cup Z(y_C)) = \mu_0(Z(x_B)|Z(x_B) \cup Z(y_B))$, which implies that assessment (iii) in Section 3 is not generally reasonable, while it can be shown that assessment (ii) is.

We now argue that checking if a given assessment is a CPPBE is easier than checking if it is a PBE satisfying NSWYDK. Theorem 1 provides a tractable characterization of CPPBE, which does not require explicitly constructing a CPS. However, we are not aware of an approach to checking whether an assessment is a PBE satisfying NSWYDK without constructing a CPS. The issue is that point (2) of Definition 7 refers to the condition that $a_{x,x'} = a_{y,y'}$, which seems difficult to express with ranked probabilities.

[17] Recall that if $x \prec_0 x'$ then $a_{x,x'} \in \mathcal{A}(h(x))$ denotes the action leading from x to x'.

[18] However, the difference between CPPBE and sequentially rational tree-extended assessment disappears in multistage games with observed actions.

[19] A similar comment applies to heterogeneous CPPBE. Fudenberg and Tirole (1991) focus on homogeneous beliefs, although they also introduce *extended assessments*, which use CPS's to guarantee that, for each $i \in I$, that all agents $j \neq i$ have a coherent belief system about i's type.

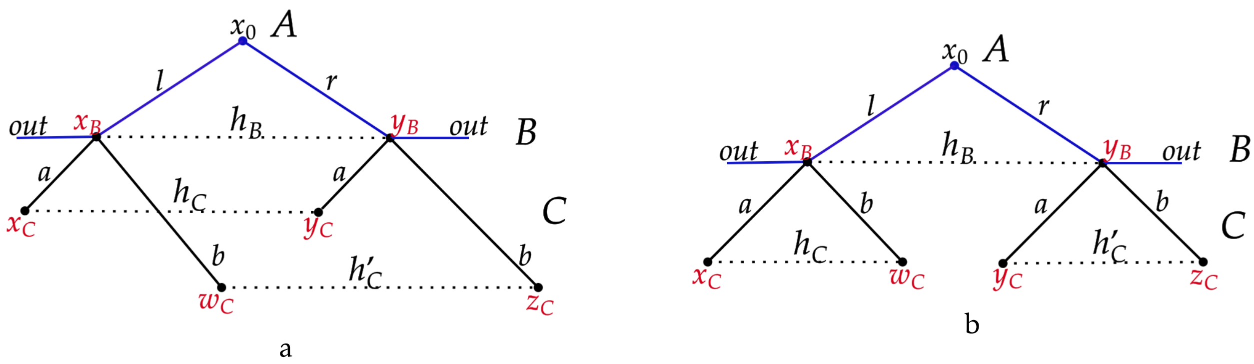

Figure 7: Implications of no-signaling-what-you-don't-know on beliefs.

To see this issue in greater detail, consider the examples in Figure 7. In these examples, there are three players. A randomizes with probabilities $\beta(l|x_0) = 2/3$, and B plays *out*. For simplicity, we don not represent payoffs or C's actions. In Figure 7a, NSWYDK imposes that $\mu_0(x_C|h_C) = \mu_0(w_C|h'_C) = 2/3$. Thus, any characterization of NSWYDK based on ranked probabilities would have to impose some conditions that depend on whether two nodes at the same information set are reached with the same action (as is the case for x_C and y_C in the example). In Figure 7b, NSWYDK has some implications for the beliefs at h_C and h'_C, but the requirements are different than before—namely, $\mu_0(x_C|h_C) = \mu_0(y_C|h'_C)$.[20] Therefore, any characterization of NSWYDK based on ranked probabilities would require some condition that depend on whether, for any information set (like h_C) and any node (like $x_C \in h_C$), there is any other node in a different information set (like $y_C \in h'_C$), that results from playing the same action (like a) at the preceding information set, and on whether this is true for any other node in the same information set (h_C).

5.2 *Almost PBE*

Among the closest solution concepts to CPPBE, one that is weaker is almost PBE (Mailath, 2025). Like CPPBE, almost PBE requires Bayesian updating on path as well as in off-path subgames, but allows arbitrary beliefs immediately following a deviation. However, unlike CPPBE, almost PBE does not require any consistency conditions on beliefs across unordered information sets. This is made precise in the definition and example below.

[20] Indeed, suppose there is a CPS μ on Z consistent with the candidate assessment. Then, $2/3 = \mu(Z(x_C)|Z(x_C, y_C)) = \mu(Z(w_C)|Z(w_C, z_C)) = \mu(Z(h_C)|Z(h_C, h'_C))$. Using the chain rule, we can write:

$$\mu(Z(x_C)|Z(x_C, y_C)) = \frac{\mu(Z(x_C)|Z(h_C))\,\mu(Z(h_C)|Z(h_C, h'_C))}{\mu(Z(x_C)|Z(h_C))\,\mu(Z(h_C)|Z(h_C, h'_C)) + \mu(Z(y_C)|Z(h'_C))\,\mu(Z(h'_C)|Z(h_C, h'_C))}.$$

This in turn implies $\mu(Z(x_C)|Z(h_C)) = \mu(Z(y_C)|Z(h'_C))$.

To define almost PBE, we first introduce some notation. For any $h, h' \in H$, we write $h \mathrel{\hat{\lesssim}} h'$ if there exist $x \in h$, $x' \in h'$ such that $x \preceq x'$. Whenever $h \mathrel{\hat{\lesssim}} h'$, define

$$h'_h := \left\{x' \in h' : \exists x \in h, x \preceq x'\right\}.$$

Moreover, given an assessment (β, μ_0), let $P^{\beta,\mu_0}(x'|h) = P^{\beta}(x'|x)\,\mu_0(x|h)$ if there exists $x \in h$ such that $x \preceq x'$, and $P^{\beta,\mu_0}(x'|h) = 0$ otherwise.[21]

Definition 8 (Mailath, 2025)**.** *An assessment (β, μ_0) is an **almost PBE** if it is sequentially rational and, for any $h \mathrel{\hat{\lesssim}} h'$ such that $\sum_{y \in h'_h} P^{\beta,\mu_0}(y|h) > 0$ and any $x' \in h'_h$,*

$$\mu_0(x'|h') = \frac{P^{\beta,\mu_0}(x'|h)}{\sum_{y \in h'_h} P^{\beta,\mu_0}(y|h)} \sum_{y \in h'_h} \mu_0(y|h').$$

The following result clarifies the relation between CPPBE and almost PBE. Recall that a **multistage game with observed actions** is an extensive-form game where (i) the information sets can be partitioned into finitely many stages, and each stage t is associated with a collection of initial nodes $W_t \subseteq X$ (the public histories at the start of stage t); (ii) for any node $w_t \in W_t$ and any stage-t information set h, either $w_t \preceq x$ for all $x \in h$ or $w_t \npreceq x$ for all $x \in h$; and (iii) for any node $w_t \in W_t$, each player is active for at most one stage-t information set that succeeds w_t.

Proposition 2. *(1) Any CPPBE is an almost PBE.*

(2) In any finite multistage game with observed actions, any almost PBE is a CPPBE.

We provide a direct proof in Appendix A. Part (1) can also be shown using Theorem 2 and the fact that every PBE consistent with a CPS on S is an almost PBE (which is the main result of Doval and Ely (2020)), while Part (2) is a new result, which we prove by applying the characterization from Theorem 1.

Intuitively, CPPBE requires the relative probabilities of different information sets that cannot be ordered by the $\hat{\lesssim}$ relation to be consistent with Bayes' rule, while almost PBE does not. However, the concepts coincide in multistage games with observed actions, as in these games the relative probabilities of such information sets can only be 0 or 1. The example below, based on Doval and Ely (2020), shows that this equivalence does not

[21] Note that P^{β,μ_0} is well-defined since perfect recall implies that, if a predecessor $x \in h$ of x' exists, then x is unique.

extend to all finite extensive-form games.[22]

Consider the example in Figure 8, with β as in blue. There are four players and, on path, players C and D do not play. However, if either A or B deviates, both C and D play, and we suppose that C randomizes at both x_C and y_C with uniform probability.

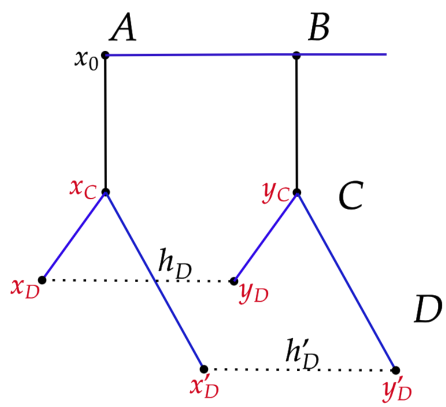

Figure 8: An almost PBE that is not a CPPBE.

In this example, the definition of almost PBE does not impose any restriction on the probabilities that D assigns to A having deviated at h_D and h'_D, as neither of these information sets precedes each other according to $\hat{\lesssim}$. However, by the characterization in Theorem 1, in a CPPBE these probability must be the same, since we must have $r_{x_C} = r_{x_D} = r_{x'_D}$, $p_{x_C} = 2p_{x_D} = 2p_{x'_D}$, $r_{y_C} = r_{y_D} = r_{y'_D}$, and $p_{y_C} = 2p_{y_D} = 2p_{y'_D}$.

6 Conclusion

This paper has introduced a solution concept, Conditional Probability Perfect Bayesian Equilibrium (CPPBE), which requires sequential rationality and consistency with a common conditional probability system on terminal nodes, but nothing else. We have argued that this concept is appealing in settings where players may entertain the possibility that their opponents' deviations are correlated, for example as a consequence of unmodeled communication or common shocks. We have also shown that CPPBE is an especially tractable version of PBE, as it admits a characterization in terms of ranked probabilities, which are much simpler objects than conditional probability systems. We hope these features can make CPPBE a useful solution concept for future research on dynamic games.

[22] Sugaya and Wolitzky (2021) conjectured that that CPPBE and almost PBE coincide in all multistage games, but the example shows that this is only true with observed actions.

Appendix A: Proofs

Proof of Theorem 1

"Only if" direction:

Let (β, μ_0) be an assessment consistent with a CPS μ on Z. We first establish a property of the null operator K defined in the text.

Claim 1. *For any $Y \subseteq X$, $Y \neq \emptyset$ and $x \in Y$, there is a unique $m_x \in \mathbb{N}$ such that $\mu(x|K^{m_x}(Y)) > 0$.*

Proof. The definition of K implies that the sets $(K^m(Y))_m$ are ordered by set inclusion. Morever, these inclusions are strict, because if $K^{m+1}(Y) = K^m(Y)$ then $\mu(K^m(Y)|K^m(Y)) = 0$, contradicting the definition of μ. Since Y is finite, this implies that there is a unique $m_x \in \mathbb{N}$ such that $x \in K^{m_x}(Y) \setminus K^{m_x+1}(Y)$. The definition of K then implies that $\mu(\{x\}|K^{m_x}) > 0$. $\square$

Now define $r_x = m_x$ and $p_x = \mu(\{x\}|K^{r_x}(X))$, for all $x \in X$. We verify that (β, μ_0) is consistent with the vector (r, p).

We first show that $\mu_0^{r,p} = \mu_0$. For any $h \in H$, let $r_h := \min_{w \in h} r_w$, and note that $h \subseteq K^{r_h}(X)$ and $\mu(h|K^{r_h}(X)) > 0$. Moreover, for any $x \in h$, we have

$$\mu(\{x\}|K^{r_h}(X)) > 0 \iff \mu(\{x\}|h) > 0 \iff \mu_0(x|h) > 0,$$

where the first equivalence follows from the chain rule and the second from consistency of (β, μ_0) with μ. Thus, if $\mu_0(x|h) = 0$ then $\mu(\{x\}|K^{r_h}(X)) = 0$, which implies $r_x > r_h$ and $\mu_0^{r,p}(x|h) = 0$. If instead $\mu_0(x|h) > 0$, then $\mu(\{x\}|K^{r_h}(X)) > 0$ and so $r_x = r_h$. Thus,

$$\mu_0^{r,p}(x|h) = \frac{p_x}{\sum_{y \in h} p_y \mathbf{1}\{r_y = r_h\}} = \frac{\mu(\{x\}|K^{r_x}(X))}{\sum_{y \in h} \mu(\{y\}|K^{r_h}(X))} = \frac{\mu(\{x\}|K^{r_h}(X))}{\mu(h|K^{r_h}(X))} = \mu(\{x\}|h) = \mu_0(x|h),$$

where the first and second equalities are by definition, the third is by the preceding observations, the fourth is by the chain rule, and the last is by consistency.

It remains to show that (r, p) satisfies Condition (2) of Definition 4. For any $x, x' \in X$, if $x' \succ x$ then $Z(x') \subseteq Z(x)$. Thus, by the chain rule, for any m such that $x \in K^m(X)$, we have

$$\mu(\{x'\}|K^m(X)) = \mu(\{x'\}|\{x\})\mu(\{x\}|K^m(X)).$$

Hence, $r_{x'} \geq r_x$. Moreover, if $\mu_0(x|h(x)) > 0$, then consistency of μ with (β, μ_0) implies that if $P^\beta(x'|x) = 0$ then $r_{x'} > r_x$; while if $P^\beta(x'|x) > 0$ then $r_{x'} = r_x$ and $p_{x'} = P^\beta(x'|x)p_x$.

"If" direction:

Let (β, μ_0) be an assessment consistent with a vector of ranked probabilities (r, p). Define the function $\mu : 2^Z \times 2^Z \setminus \{\emptyset\} \to [0,1]$ by

$$\mu(A|B) = \frac{\sum_{z \in A \cap B} p_z \mathbf{1}\{r_z = \min_{z' \in B} r_{z'}\}}{\sum_{z \in B} p_z \mathbf{1}\{r_z = \min_{z' \in B} r_{z'}\}} \qquad \text{for all } A \subseteq Z,\ B \subseteq Z \setminus \{\emptyset\}.$$

We show that μ is a CPS consistent with (β, μ_0). That μ is a CPS is straightforward.

Claim 2. *The function μ is a CPS on Z.*

Proof. The only non-trivial part is the chain rule. Fix $A \subseteq B \subseteq C \subseteq Z$, $B \neq \emptyset$. Note that $\min_{z \in B} r_z \geq \min_{z \in C} r_z$. If $\min_{z \in B} r_z > \min_{z \in C} r_z$, then $\mu(A|C) = \mu(B|C) = 0$, and the chain rule trivially holds. Thus, suppose that $\min_{z \in B} r_z = \min_{z \in C} r_z$, which implies that $\mu(B|C) > 0$. To ease notation, let $\rho := \min_{z \in B} r_z$. Then, by definition,

$$\mu(A|B) \cdot \mu(B|C) = \frac{\sum_{z \in A} p_z \mathbf{1}\{r_z = \rho\}}{\sum_{z \in B} p_z \mathbf{1}\{r_z = \rho\}} \cdot \frac{\sum_{z \in B} p_z \mathbf{1}\{r_z = \rho\}}{\sum_{z \in C} p_z \mathbf{1}\{r_z = \rho\}} = \frac{\sum_{z \in A} p_z \mathbf{1}\{r_z = \rho\}}{\sum_{z \in C} p_z \mathbf{1}\{r_z = \rho\}} = \mu(A|C),$$

implying the chain rule. □

The most involved step is consistency of (β, μ_0) and μ. We start with a preliminary result.

Claim 3. *For any $h \in H$ and $x \in h$ such that $\mu_0(x|h) > 0$, we have*

(1) For any $z \in Z(x)$, $r_z = r_x$ if and only if $P^\beta(z|x) > 0$; and

(2) $\min_{z \in Z(x)} r_z = r_x$.

Proof. For any $z \in Z(x)$, Definition 4 implies that if $P^\beta(z|x) > 0$ then $r_z = r_x$, while if $P^\beta(z|x) = 0$ then $r_z > r_x$. Moreover, there is at least one $z \in Z(x)$ such that $P^\beta(z|x) > 0$. □

We are now ready to establish consistency.

Claim 4. *For any $h \in H$ and $x \in h$, $\mu_0(x|h) = \mu(Z(x)|Z(h))$.*

Proof. To ease notation, let $\rho := \min_{z\in Z(h)} r_z$. By consistency with (r,p), for any $x \in h$,

$$\mu_0(x|h) > 0 \iff r_x = \min_{y\in h} r_y.$$

Since there exists some $x' \in h$ with $\mu_0(x'|h) > 0$, point (2) of Claim 3 implies that

$$\min_{y\in h} r_y = \rho.$$

We consider in turn the cases where $\mu_0(x|h)$ equals zero and where it is positive. If $\mu_0(x|h) = 0$, then consistency with (r,p) implies that $\min_{z\in Z(x)} r_z \geq r_x$, and the preceding observations imply that $r_x > \min_{y\in h} r_y = \rho$. Thus,

$$\mu(Z(x)|Z(h)) = \frac{\sum_{z\in Z(x)} p_z \mathbf{1}\{r_z = \rho\}}{\sum_{z\in Z(h)} p_z \mathbf{1}\{r_z = \rho\}} = 0 = \mu_0(x|h).$$

Now suppose that $\mu_0(x|h) > 0$. By consistency with (r,p), for any $z \in Z(x)$ such that $P^\beta(z|x) > 0$, we have $p_z = P^\beta(z|x)p_x$. Moreover, by point (1) of Claim 3, for any $z \in Z(x)$, we have $r_z = \rho$ if and only if $P^\beta(z|x) > 0$. Thus,

$$\sum_{z\in Z(x)} p_z \mathbf{1}\{r_z = \rho\} \quad = \sum_{\substack{z\in Z(x):\\ P^\beta(z|x)>0}} p_z \quad = \quad p_x \sum_{\substack{z\in Z(x):\\ P^\beta(z|x)>0}} P^\beta(z|x) \quad = \quad p_x.$$

Since this holds for any $x \in h$ such that $\mu_0(x|h) > 0$, we have

$$\begin{aligned}\sum_{z\in Z(h)} p_z \mathbf{1}\{r_z = \rho\} \quad &= \quad \sum_{y\in h} \sum_{z\in Z(y)} p_z \mathbf{1}\{r_z = \rho\} = \sum_{\substack{y\in h:\\ \mu_0(y|h)>0}} \sum_{\substack{z\in Z(y):\\ P^\beta(z|y)>0}} p_z \\ &= \sum_{\substack{y\in h:\\ \mu_0(y|h)>0}} p_y \sum_{\substack{z\in Z(y):\\ P^\beta(z|y)>0}} P^\beta(z|y) = \sum_{\substack{y\in h:\\ \mu_0(y|h)>0}} p_y.\end{aligned}$$

Finally, we have

$$\mu(Z(x)|Z(h)) = \frac{\sum_{z\in Z(x)} p_z \mathbf{1}\{r_z = \rho\}}{\sum_{z\in Z(h)} p_z \mathbf{1}\{r_z = \rho\}} = \mu_0(x|h).$$

□

Claim 5. *For any $x \prec x'$ with $\mu_0(x|h(x)) > 0$, we have $\mu(Z(x')|Z(x)) = P^\beta(x'|x)$.*

Proof. By point (2) of Claim 3, we have

$$\min_{z \in Z(x)} r_z = r_x.$$

Moreover, consistency with (r, p) implies that $r_{x'} \geq r_x$, and point (1) of Claim 3 implies that this inequality holds with equality if and only if $P^\beta(x'|x) > 0$.

We consider in turn the cases where $P^\beta(x'|x)$ equals zero and where it is positive. If $P^\beta(x'|x) = 0$, then

$$\mu(Z(x')|Z(x)) = \frac{\sum_{z \in Z(x')} p_z \mathbf{1}\{r_z = \min_{z' \in Z(x)} r_{z'}\}}{\sum_{z \in Z(x)} p_z \mathbf{1}\{r_z = \min_{z' \in Z(x)} r_{z'}\}} = \frac{\sum_{z \in Z(x')} p_z \mathbf{1}\{r_z = r_x\}}{\sum_{z \in Z(x)} p_z \mathbf{1}\{r_z = r_x\}} = 0 = P^\beta(x'|x).$$

Now suppose that $P^\beta(x'|x) > 0$. Fix any $z \in Z(x')$ such that $P^\beta(z|x') > 0$. Since $P^\beta(z|x) = P^\beta(z|x') \cdot P^\beta(x'|x)$, we have $P^\beta(z|x) > 0$. Thus, the preceding observations and point (1) of Claim 3 imply that $r_z = r_x = r_{x'}$. We thus have

$$\begin{aligned}
\sum_{z \in Z(x')} p_z \mathbf{1}\{r_z = \min_{z' \in Z(x')} r_{z'}\} &= \sum_{z \in Z(x')} p_z \mathbf{1}\{r_z = r_x\} = \sum_{\substack{z \in Z(x'):\\ P^\beta(z|x')>0}} p_z \\
&= p_x \sum_{\substack{z \in Z(x'):\\ P^\beta(z|x')>0}} P^\beta(z|x) = p_x \sum_{\substack{z \in Z(x'):\\ P^\beta(z|x')>0}} P^\beta(z|x') P^\beta(x'|x) \\
&= P^\beta(x'|x) p_x \sum_{\substack{z \in Z(x'):\\ P^\beta(z|x')>0}} P^\beta(z|x') = P^\beta(x'|x) p_x,
\end{aligned}$$

where the third equality follows because consistency with (r, p) implies that, for any z such that $P^\beta(z|x') > 0$, we have $p_z = p_x P^\beta(z|x)$.

By a similar reasoning,

$$\begin{aligned}
\sum_{z \in Z(x)} p_z \mathbf{1}\{r_z = \min_{z' \in B} r_{z'}\} &= \sum_{z \in Z(x)} p_z \mathbf{1}\{r_z = r_x\} = \sum_{\substack{z \in Z(x):\\ P^\beta(z|x)>0}} p_z \\
&= p_x \sum_{\substack{z \in Z(x):\\ P^\beta(z|x)>0}} P^\beta(z|x) = p_x.
\end{aligned}$$

In sum, we have

$$\mu(Z(x')|Z(x)) = \frac{\sum_{z\in Z(x')} p_z \mathbf{1}\{r_z = \min_{z'\in B} r_{z'}\}}{\sum_{z\in Z(x)} p_z \mathbf{1}\{r_z = \min_{z'\in B} r_{z'}\}} = P^{\beta}(x'|x).$$

□

Proof of Theorem 2

Recall that we write $h \prec x'$ when there is $x \in h$ such that $x \prec x'$, and similarly for $h \nprec x'$.

Part (1):

Fix a PBE (β, μ_0) consistent with a CPS ν on S. Let μ be the CPS on Z equivalent to ν. We will show that (β, μ_0) is a CPPBE consistent with μ.

By definition, (β, μ_0) is sequentially rational. Moreover, by definition, for any h and $x \in h$,

$$\mu_0(x|h) = \nu(S(x)|S(h)) = \mu(Z(x)|Z(h)).$$

It remains to show that, for any $x \prec x'$, if $\mu_0(x|h(x)) > 0$, then $\mu(Z(x')|Z(x)) = P^{\beta}(x'|x)$. For such x, x', by definition,

$$\mu(Z(x')|Z(x)) = \sum_{z\in Z(x')} \mu(\{z\}|Z(x)).$$

Note that, for every s, s' such that $z_s = z_{s'} \in Z(x')$, we have $s(h) = s'(h)$ for all $h \prec x'$. Thus, using Doval and Ely's (2020) definition of PBE, if $z \in Z(x')$ then

$$\begin{aligned}\mu(\{z\}|Z(x)) &= \sum_{s:\ z_s=z} \nu(s|S(x)) = \sum_{s:\ z_s=z} \prod_{h\nprec x} \beta(s(h)|h) = \sum_{s:\ z_s=z} \Big(\prod_{h\nprec x,\ h\prec x'} \beta(s(h)|h)\Big)\Big(\prod_{h\nprec x'} \beta(s(h)|h)\Big) \\ &= P^{\beta}(x'|x) \sum_{s:\ z_s=z} \Big(\prod_{h\nprec x'} \beta(s(h)|h)\Big) = P^{\beta}(x'|x) \sum_{s:\ z_s=z} \nu(s|S(x')) = P^{\beta}(x'|x)\mu(\{z\}|Z(x')),\end{aligned}$$

and therefore

$$\mu(Z(x')|Z(x)) = \sum_{z\in Z(x')} P^{\beta}(x'|x)\mu(\{z\}|Z(x')) = P^{\beta}(x'|x).$$

Part (2):

Fix a CPPBE (β, μ_0) consistent with a CPS μ on Z. We construct a CPS ν on S such that (β, μ_0) is a PBE consistent with ν.

Since μ is a CPS on Z, there exists a sequence of totally mixed probabilities $(\eta_\varepsilon)_\varepsilon$ on Z such that

$$\mu(\{z\}|A) = \lim_{\varepsilon} \eta_\varepsilon(z|A) \qquad \text{for all } z \in Z, A \subseteq Z, A \neq \emptyset$$

Fix any sequence of totally mixed behavior strategies $(\beta_\varepsilon)_\varepsilon$ such that at each information set h, $\lim_\varepsilon \beta_\varepsilon(\cdot|h) = \beta(\cdot|h)$. Now define

$$\sigma_\varepsilon(s) := \eta_\varepsilon(z_s) \prod_{h \nless z_s} \beta_\varepsilon(s(h)|h) \qquad \text{for all } s \in S.$$

Note that $\sigma_\varepsilon > 0$ and, for each $z \in Z$, $\sigma_\varepsilon(S(z)) = \eta_\varepsilon(z)$. Thus, for each ε, σ_ε is a totally mixed probability function on S. In words, $\sigma_\varepsilon(s)$ weights the path of play induced by s according to η_ε, and coincides with β_ε at each history off this path.[23]

Finally, define ν by

$$\nu(s|B) = \lim_{\varepsilon} \sigma_\varepsilon(s|B) \qquad \text{for all } s \in S,\, B \subseteq S.$$

Then, ν is a CPS by Myerson (1986) and is equivalent to μ by construction. It thus remains only to show that for any h and $x \in h$ with $\mu_0(x|h) > 0$, we have $\nu(s|S(x)) = \prod_{h \nless x} \beta(s(h)|h)$ if $s \in S(x)$ and $\nu(s|x) = 0$ otherwise (the other parts of the definition of consistency of (β, μ_0) and ν following from the equivalence of ν and μ).

So, fix $h \in H$, $x \in h$ with $\mu_0(x|h) > 0$, and $s \in S$. By construction,

$$\nu(s|S(x)) = \lim_{\varepsilon} \frac{\sigma_\varepsilon(s \cap S(x))}{\sigma_\varepsilon(S(x))},$$

which equals zero if $s \notin S(x)$. If instead $s \in S(x)$, then

$$\begin{aligned}
\nu(s|S(x)) = \lim_{\varepsilon} \frac{\sigma_\varepsilon(s)}{\sigma_\varepsilon(S(x))} &= \lim_{\varepsilon} \frac{\eta_\varepsilon(z_s) \prod_{h' \nless z_s} \beta_\varepsilon(s(h')|h')}{\eta_\varepsilon(Z(x))} = \lim_{\varepsilon} \eta_\varepsilon(z_s|Z(x)) \lim_{\varepsilon} \prod_{h' \nless z_s} \beta_\varepsilon(s(h')|h') \\
&= \mu(\{z_s\}|Z(x)) \prod_{h' \nless z_s} \beta(s(h')|h') = P^\beta(z_s|x) \prod_{h' \nless z_s} \beta(s(h')|h') \\
&= \beta(s(h)|h) \prod_{h' > x} \beta(s(h')|h') \prod_{h' \nless z_s} \beta(s(h')|h') = \prod_{h' \nless x} \beta(s(h')|h'),
\end{aligned}$$

where the second equality follows because $S(x) = \cup_{z \in Z(x)} S(z)$, the sets $\{S(z)\}_{z \in Z(x)}$ are mu-

[23] Intuitively, backward induction implies that, at every terminal node z, the agents believe that the equilibrium would have been followed at all nodes not preceding z. Thus, their beliefs about play at these nodes should be consistent with independent trembles—or, equivalently, their beliefs can allow for correlated trembles only at nodes preceding z.

tually disjoint, and σ_ε is finitely additive for each ε; the third follows since both limits exist finite; the fifth follows because μ is consistent with (β, μ_0) and $\mu_0(x|h) > 0$; and the last two follow from properties of trees and $z_s \succ x$.

Proof of Proposition 2

Part (1):

Fix a CPPBE (β, μ_0) consistent with a CPS μ on Z. Fix h, h' such that $h \hat{\lesssim} h'$ and $\sum_{y \in h'_h} P^{\beta,\mu_0}(y|h) > 0$. For any $x' \in h'_h$, consistency with μ and the chain rule imply that

$$\mu_0(x'|h') = \mu(x'|h') = \mu(x'|h'_h)\,\mu(h'_h|h') = \mu(x'|h'_h) \sum_{y \in h'_h} \mu(y|h') = \mu(x'|h'_h) \sum_{y \in h'_h} \mu_0(y|h').$$

Moreover, for any $x' \in h'_h$, fix $x \in h$ such that $x \prec x'$. Consistency with μ implies that $\mu_0(x|h) = \mu(x|h)$ and, if $\mu_0(x|h) > 0$, then $\mu(x'|x) = P^\beta(x'|x)$. Thus, $P^{\beta,\mu_0}(x'|h) = \mu(x'|h)$.

Since $0 < \sum_{y \in h'_h} P^{\beta,\mu_0}(y|h) = \mu(h'_h|h)$, the chain rule the preceding observations imply that

$$\mu(x'|h'_h) = \frac{\mu(x'|h)}{\mu(h'_h|h)} = \frac{P^{\beta,\mu_0}(x'|h)}{\sum_{y \in h'_h} P^{\beta,\mu_0}(y|h)},$$

completing the proof.

Part (2):

Fix an almost PBE (β, μ_0). Note that (β, μ_0) is sequentially rational by assumption; we will show that there is a CPS on Z consistent with it by defining a vector (r, p) that satisfies the hypotheses of Theorem 1.

For any node $x \neq x_0$, let $\pi(x)$ be the direct predecessor of x. Let $r_{x_0} = 0$. For any $x \neq x_0$, if $P^\beta(x|\pi(x)) = 0$ then let $r_x = r_{\pi(x)} + 1$ and fix any $p_x \in (0,1]$. If instead $P^\beta(x|\pi(x)) > 0$, let $r_x = r_{\pi(x)}$ and $p_x = p_{\pi(x)} P^\beta(x|\pi(x))$. Note that the resulting vector (r, p) satisfies condition (2) of Definition 4. We will show that it also satisfies $\mu_0 = \mu_0^{r,p}$, and thus is consistent with (β, μ_0).

To this end, fix any $h \in H$ and $x \in h$. If h is a singleton, then trivially $\mu_0(\cdot|h) = \mu_0^{r,p}(\cdot|h)$. Otherwise, then there exists some $w \in X$ such that h belongs to the stage starting from w, which implies that $\{w\} \hat{\prec} h$ and $\sum_{y \in h} P^{\beta,\mu_0}(h|\{w\}) = 1$. By definition, this implies that $\mu_0(x|h) = P^\beta(x|w)$.

It thus remains to show that $\mu_0^{r,p}(x|h) = P^\beta(x|w)$. Note that, by construction, for any $y \in h$, we have $r_y = \min_{y' \in h} r_{y'}$ if and only if $P^\beta(y|\{w\}) > 0$. Thus,

$$\mu_0^{r,p}(x|h) = \frac{p_x \mathbf{1}\{r_x = \min_{y' \in h} r_{y'}\}}{\sum_{y \in h} p_y \mathbf{1}\{r_y = \min_{y' \in h} r_{y'}\}} = \frac{p_w P^\beta(x|w)}{p_w \sum_{y \in h} P^\beta(y|w)} = P^\beta(x|w).$$

Appendix B: Relation with Plain PBE

Watson (2025b) requires PBE to satisfy a consistency notion labeled plain consistency. Intuitively, plain consistency requires that, if a player's belief treats a component of the strategy set and its complement as independent, this independence must be preserved at any immediately succeeding information set that is reached with positive probability.

Watson's definition allows agents to hold heterogeneous beliefs, so we compare it with heterogeneous CPPBE. In particular, we show that the two solution concepts are non-nested. For the reader's convenience, we briefly review the relevant definitions here.

The definition of plain consistency is based on the notion of an *appraisal*: that is, for each information set, a distribution over complete strategy profiles S representing the beliefs and strategy of the active player at that information set. For each agent i, let H_i the subset of histories h where i is active. We also need to introduce, for each agent i, a fictitious information set $\underline{h}_i$ representing i's initial history before the game has started.[24] Let $\underline{H}_i := H_i \cup \underline{h}_i$, and $\underline{H} := \cup_i \underline{H}_i$. Note that every strategy s can be extended to have domain $\underline{H}$ by setting $s(\underline{h}_i) = wait$ for every i not active at $\emptyset$.

The formal definition of an appraisal involves an independence condition on product sets of strategy profiles. For any strategy profile s, we focus on the sequence-representation of s: that is, $s = (s(h))_{h \in \underline{H}} \in \times_{h \in \underline{H}} \mathcal{A}(h)$. For any $L \subseteq \underline{H}$, let $\pi_L(s)$ the L-projection of s: that is, $\pi_L(s) := (s(h))_{h \in L}$. Moreover, for any $B \subseteq S$, let $B_L := \{s_L : s \in B\} = \pi_L(B)$. Finally, let $-L := \underline{H} \setminus L$. We say that $B \subseteq S$ is an L-**product set** if $B = B_L \times B_{-L}$.

Definition 9 (Watson, 2025b)**.** *For an L-product set $C \subseteq S$, we say that $p \in \Delta(S)$* ***exhibits L-independence*** *on C if $p(C) > 0$ and, for every L-product set $B \subseteq C$, $p(B|C) = p(B_L \times C_{-L}|C) \cdot p(C_L \times B_{-L}|C)$.*

Definition 10 (Watson, 2025b)**.** *For any agent i and history $h \in \underline{H}_i$, an* ***appraisal*** *p_h is a distribution on S such that*

[24] Formally, $\underline{h}_i$ precedes any $h \in H_i$ and $S(\underline{h}_i) = S$. Moreover, $\mathcal{A}(\underline{h}_i) = \{wait\}$ for all i except j active at $\emptyset$, and $\underline{h}_j = \emptyset$.

(1) *supp* $p_h \subseteq S(h)$;

(2) *for every* $h' \in H_i$, p_h *exhibits* h'*-independence on* S.

An ***appraisal system*** *is a profile of appraisals* $P = (p_h)_{h \in \underline{H}}$.

We say that $h, h' \in \underline{H}_i$ are **consecutive** whenever there exist $x \in h$, $x' \in h'$ such that $x \prec x'$, and for any such x, x' there does not exist $h'' \in \underline{H}_i$ and $x'' \in h''$ such that $x \prec x'' \prec x'$.

Definition 11 (Watson, 2025b)**.** *Fix* $L \subseteq H$ *and an* L*-product set* $C \subseteq S$. *For some agent* i, *fix* $h, h' \in \underline{H}_i$ *consecutive and suppose that* $C \subseteq S(h)$ *and that* p_h *exhibits* L*-independence on* C. *Then, the appraisal system* P *satisfies* ***plain consistency*** *if*

$$\frac{p_{h'}(B_L \times \{s_{-L}\})}{p_{h'}(C_L \times \{s_{-L}\})} = \frac{p_h(B_L \times C_{-L})}{p_h(C)}$$

for every $B_L \subseteq C_L$, $s_{-L} \in C_{-L}$ *such that* $C_L \times \{s_{-L}\} \subseteq S(h')$ *and* $p_{h'}(C_L \times \{s_{-L}\}) > 0$.

To get an intuition, suppose that $h, h' \neq \emptyset$, agent $j \neq i$ is active at $\emptyset$, and i does not observe j's choice. Fix $L = \{\emptyset\}$. Plain consistency implies that, if at h agent i believes that j's trembles at $\emptyset$ are independent from trembles at any other information set, and if h' is on path given h, then i keeps believing that j's trembles at $\emptyset$ are independent at h'.

The definition of plain PBE can be stated as follows.

Definition 12. *An appraisal system* $P = (p_h)_{h \in H}$ *is a* ***plain PBE*** *if there exists a behavior strategy* β *such that*

(1) *for all* s, $p_\emptyset(s) = \prod_{h \in H} \beta(s_h | h)$;

(2) P *is sequentially rational;*

(3) P *satisfies plain consistency.*

Note that plain consistency needs to be checked at each pair of consecutive information sets and for each possible product set of strategy profiles consistent with these information sets. By contrast, checking whether an assessment is a CPPBE only requires checking the existence of a vector of ranked probabilities satisfying Definition 4.

We now show that plain PBE and CPPBE are non-nested. The intuition is that plain PBE rules out some correlation in agents' trembles, but it may not be consistent with agents using Bayes' rule whenever the conditioning event has positive probability. Below, we make this point precise with two examples. Intuitively, we say that a plain PBE and a CPPBE are equivalent if they induce the same play and the same beliefs at each information set.[25]

The first example is a plain PBE that is not equivalent to any heterogeneous CPPBE. Consider the game below played by A, B, and C. Let β be as in blue, with B randomizing with uniform probability. Assume that payoffs are such that β is sequentially rational.

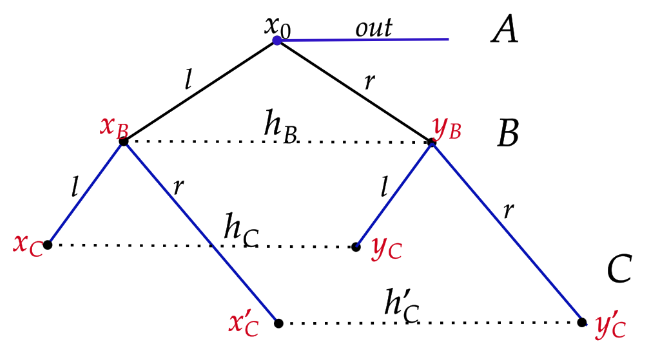

Figure 9: A plain PBE that is not equivalent to any heterogeneous CPPBE.

The definition of heterogeneous CPPBE implies that, if β is played in a heterogeneous CPPBE (β, μ_0), then $\mu_0(x_C|h_C) = \mu_0(x'_C|h'_C)$. That is, whatever belief C forms after A deviates, C subsequently updates following Bayes' rule on path. This is not true in a plain PBE, since plain consistency does not impose any restrictions on agent C's beliefs at h_C and h'_C. In particular, note that the only consecutive histories of agent C are $\underline{h}_C$, h_C and $\underline{h}_C$, h'_C. For Definition 11 to restrict the beliefs at h_C and h'_C, we need to fix L containing the initial history $\emptyset$. However, this implies that any $C \subseteq S$ satisfying $C_L \times S_{-L} \subseteq S(h_C) \cup S(h'_C)$ does not contain any strategy in which A plays *out*. At the same time, any C satisfying $p_{\underline{h}_C}(C) > 0$ needs to contain at least one of these strategies. Thus, plain consistency has no bite in pinning down C's beliefs at h_C and h'_C.

The second example is a heterogeneous CPPBE that is not equivalent to any plain PBE. It is based on Watson (2025b). There are again three players, A, B, and C, and the game is played as shown by the tree below. Fix β as in blue, with $\beta(b|h_B) = 0.6$, $\beta(c|h_B) = 0.2$, and payoffs such that β is sequentially rational.

[25] More precisely, a plain PBE $P = (p_h)_{h \in \underline{H}}$ consistent with a behavior strategy β is equivalent to a heterogeneous CPPBE (β, μ_0) consistent with CPS $(\mu_i)_i$ if, for all i and $h \in H_i$, $p_h = \mu_i(\cdot|h)$ and $p_{\underline{h}_i} = \mu_i(\cdot|\emptyset)$

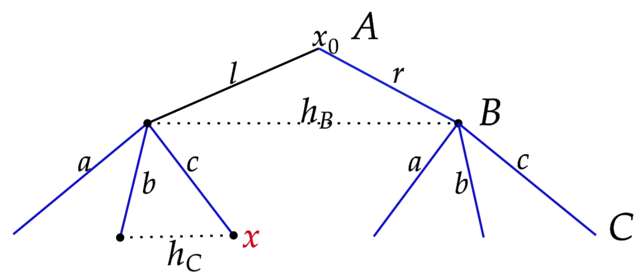

Figure 10: A heterogeneous CPPBE that is not equivalent to any plain PBE.

If β is played in a CPPBE (β, μ_0), then a straightforward application of the characterization in Theorem 1 implies that $\mu_0(x|h_C)$ can take any value. However, if (β, μ_0) is a plain PBE, then the only value $\mu_0(x|h_C)$ can take is $\frac{1}{4}$. This follows from plain consistency applied to h_B and h_C, which implies that C sees as independent A's trembles and B's action - that is, C believes that even if A deviates B plays following β.

References

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

Blume, L., A. Brandenburger, and E. Dekel (1991a): "Lexicographic probabilities and choice under uncertainty," *Econometrica: Journal of the Econometric Society*, 61–79.

——— (1991b): "Lexicographic probabilities and equilibrium refinements," *Econometrica: Journal of the Econometric Society*, 81–98.

Catonini, E. and A. Penta (2025): "Backward induction reasoning beyond backward induction," *American Economic Journal: Microeconomics*, forthcoming.

Dilmé, F. (2024): "A characterization of consistent assessments using power sequences of strategy profiles," *International Journal of Game Theory*, 53, 673–693.

Doval, L. and J. C. Ely (2020): "A note on Bayes' rule where possible," Tech. rep.

Fudenberg, D., D. M. Kreps, and D. K. Levine (1988): "On the robustness of equilibrium refinements," *Journal of Economic Theory*, 44, 354–380.

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

Halpern, J. Y. (2010): "Lexicographic probability, conditional probability, and nonstandard probability," *Games and Economic Behavior*, 68, 155–179.

Hammond, P. J. (1994): "Elementary non-Archimedean representations of probability for decision theory and games," in *Patrick Suppes: Scientific Philosopher: Volume 1. Probability and Probabilistic Causality*, Springer, 25–61.

Harris, C., P. Reny, and A. Robson (1995): "The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A case for public randomization," *Econometrica: Journal of the Econometric Society*, 507–544.

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 R. Wilson (1982): "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 863–894.

Kuhn, H. W. (1953): "Extensive games and the problem of information," *Contributions to the Theory of Games*, 2, 193–216.

Mailath, G. J. (2025): "Modeling Strategic Behavior," Available at url: "https://bpb-us-w2.wpmucdn.com".

Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995): *Microeconomic theory*, vol. 1, Oxford university press New York.

Myerson, R. B. (1986): "Multistage Games with Communication," *Econometrica: Journal of the Econometric Society*, 323–58.

——— (1991): *Game Theory: Analysis of Conflict*, Harvard university press.

Rényi, A. (1955): "On a new axiomatic theory of probability," *Acta Mathematica Hungarica*, 6, 285–335.

Spohn, W. (1988): "Ordinal conditional functions: A dynamic theory of epistemic states," in *Causation in decision, belief change, and statistics: Proceedings of the Irvine Conference on Probability and Causation*, Springer, 105–134.

Sugaya, T. and A. Wolitzky (2021): "The revelation principle in multistage games," *The Review of Economic Studies*, 88, 1503–1540.

Watson, J. (2025a): "Partially constructed sequential equilibrium," *Journal of Economic Theory*, 106111.

——— (2025b): "Perfect Bayesian Equilibrium: Consistency Conditions for Practical Definitions," Tech. rep., mimeo.