

Conditional Probability Perfect Bayesian Equilibrium

Anna Merotto*

MIT

Alexander Wolitzky[†]

MIT

September 1, 2026

Abstract

We study a notion 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).² 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

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

² We do not claim that this is a new solution concept: as we discuss, we show an equivalence between our definition and that of Doval and Ely (2020), and the literature contains many other similar concepts.

assessment satisfying a consistency condition in the spirit of Kreps and Wilson (1982), but allowing correlated trembles. (Hence, CPPBE is more permissive than sequential 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 that this information set was reached due to earlier unmodeled communication between their opponents, or an earlier unmodeled correlated shock.³ 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

³ 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.

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 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, 2019), *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 2.⁴ 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.⁵

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 of 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, in Appendix D, we establish a sense in which a version of CPPBE is the weakest equilibrium concept for dynamic Bayesian games consistent with sequential rationality and backward induction where Bayes’ rule is always applied after positive-probability events: specifically, we show that CPPBE is equivalent to the “backwards rationalizability” concept of Catonini and Penta (2026), to-

⁴The exception is plain PBE, which is non-nested with CPPBE, as shown in Appendix C. 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 (2019).

⁵For 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 and 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.




Figure 2: Equilibrium refinements in extensive-form games, ordered by strict inclusion.

gether with a common and independent prior.⁶ Of course, by suggesting this benchmark, we do not mean to discourage the use of either stronger or weaker refinements in applications where they prove useful.

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. In general, games with communication—either with a neutral mediator as in Myerson (1986) and Sugaya and Wolitzky (2021), or without one as in cheap talk games—are ones where the correlated trembles allowed by CPPBE seem especially natural, as it seems reasonable that players may entertain the possibility of unmodeled communication channels in addition to those explicitly incorporated in the extensive form. Moreover, the tractability of CPPBE in games with mediated communication (i.e., the revelation principle) is one motivation for studying it more generally.

Our characterization of CPPBE in terms of ranked probabilities builds on prior results that establish equivalences between CPS’s and lexicographic conditional probability systems (LCPS’s) (Blume et al., 1991a,b; Hammond, 1994; Halpern, 2010). However, in contrast to our characterization, these results connect CPS’s and LCPS’s defined on the same

⁶ This result holds in multistage games with observed actions, the class of games considered by Catonini and Penta (2026).

domain (whereas we connect CPS's on terminal nodes and ranked probabilities on all nodes), 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 LCPS's as defined by Blume et al. (1991a) (so our main contribution is not introducing ranked probabilities, but using them to characterize CPPBE). 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 it is equivalent to define CPPBE via a CPS over either terminal nodes or strategy profiles. The appendix explains CPPBE's relation to NSWYDK and to plain PBE, and provides an epistemic foundation.

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 ($x \in X$).⁷ We do not explicitly introduce Nature, but it can be included as an additional player with a fixed strategy.⁸

Let $>$ and $\geq$ denote respectively the strict and weak successor relation on X , and $>_0$ the direct successor relation. As usual, $Z \subseteq X$ is the set of terminal nodes, h_0 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 $(u_i : Z \rightarrow \mathbb{R})_{i \in I}$.

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 $S := \times_{i \in I} S_i$.

⁷The definition of CPPBE extends straightforwardly to infinite games, but equilibrium existence becomes a subtle issue, as with any refinement of subgame perfection (Harris et al., 1995).

⁸We return to this point in Section 2.3.

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 Appendix C.

For any node x , we let $h(x)$ be 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 let $S(Y) := \bigcup_{x \in Y} S(x)$ and $Z(Y) := \bigcup_{x \in Y} Z(x)$. We let $P^\beta(\cdot|x)$ the probability induced by β on X .⁹

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

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

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 $\mu : 2^\Omega \times 2^\Omega \setminus \{\emptyset\} \rightarrow [0, 1]$ 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 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 $x \in h$,

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

⁹Given $y, y' \in X$ with $y <_0 y'$, let $a_{y, y'} \in \mathcal{A}(h(y))$ be the action leading from y to y' . Fix any $x, x' \in X$. If $x < x'$, there exists a sequence $x <_0 x_1 <_0 \dots <_0 x_m <_0 x'$ in X . In this case, fix $P^\beta(x'|x) := \beta(a_{x, x_1}|h(x)) \dots \beta(a_{x_m, x'}|h(x_m))$. If $x = x'$, let $P^\beta(x|x) := 1$; otherwise, let $P^\beta(x'|x) := 0$.

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

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

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 beliefs to be consistent with a CPS on Z —as in point (1) of Definition 1—is due to Fudenberg and Tirole (1991).¹⁰ Following Myerson (1986), Sugaya and Wolitzky (2021) introduced CPPBE in the context of multi-stage games with mediated communication, which requires a more intricate definition. Doval and Ely (2020) define a notion of an assessment being consistent with a CPS on S , which we show is equivalent to our definition in Section 4.

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 $\beta(in|h_0) = \beta(in|h_B) = 1$, by backward induction. By the definition, this implies that $\mu_0(x_B|h_B) = \mu_0(x_C|h_C) = 1$. Sequential rationality also implies that $\beta(in|h'_C) = 1$ 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 .¹¹ 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 $\mu : 2^\Omega \times 2^\Omega \setminus \{\emptyset\} \rightarrow [0, 1]$ 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.

¹⁰ Most of Fudenberg and Tirole's paper focuses on beliefs about players' types in multistage games with observed actions, but their Section 6 considers beliefs in general extensive-form games.

¹¹ Moreover, for any $\mu_0(x'_C|h'_C)$, Theorem 1 ensures that there exist a CPS μ on Z consistent with it.

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) $\lim_j \beta_j = \beta$;
- (2) $\lim_j \eta_j(Z(x)|Z(h)) = \mu_0(x|h)$;
- (3) $\mu_0(x|h) > 0 \implies P^\beta(x'|x) = \lim_j \eta_j(Z(x')|Z(x))$; and
- (4) $\forall j \in \mathbb{N}, P^{\beta_j}(x|h_0)/P^{\beta_j}(h|h_0) = \eta_j(Z(x)|Z(h))$.

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.¹² It shows that the difference between sequential equilibrium and CPPBE is that, in a sequential equilibrium, each distribution η_j is required to be consistent with the behavior strategy β_j , which assigns independent probabilities to trembles at 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.

Proposition 1 also suggests a way to restrict *which* players' trembles can be correlated, rather than allowing correlation across the whole population as in CPPBE. Fix a partition $(B_k)_{k=1}^K$ of I into "blocks," interpreted as groups of players who share an unobserved communication channel or a common shock, and consider sequences η_j that are required to factor across blocks, $\eta_j = \eta_j^{(1)} \otimes \dots \otimes \eta_j^{(K)}$, where each $\eta_j^{(k)}$ is a full-support distribution over the projection of Z onto B_k 's own histories. Within a block, trembles remain entirely unrestricted and may be correlated in arbitrary ways, as in CPPBE; across blocks, however, the product structure precludes correlation, since no shared randomization device links the blocks at any stage of the approximating sequence, and hence in the resulting CPS. The resulting solution concept implies CPPBE (where $K = 1$) and can capture some relevant considerations. In particular, in games with Nature, one can require that Nature's moves cannot be correlated with any player's trembles.¹³

We can show basic properties of CPPBE.

¹² The first part of Proposition 1 also follows by Lemma 3.1 of Battigalli (1996); the equivalence between CPS's and sequential equilibrium is also studied by McLennan (1989a,b).

¹³ The extension suggested in this paragraph is somewhat in the spirit of Lipnowski and Sadler (2019), who introduce a solution concept where only some players are required to have correct conjectures about each other's strategies.

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 follows from Theorem 2 and Doval and Ely (2020)), and every almost PBE is subgame perfect (by Mailath (2019)). 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, by a standard argument.

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

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

or equivalently

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

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.¹⁴ 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 on 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 $2^{|Z|} - 2^{|Z|}$ probabilities $\mu(A|B)$, one for each pair of subsets of terminal nodes $A, B \subseteq Z$, $B \neq \emptyset$.

Definition 3. 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 $x < x'$ we have $r_x \leq r_{x'}$ and, if $\mu_0(x|h(x)) > 0$, then

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

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 3 is a formalization of the notion of “Bayes’ rule where possible.”¹⁵ 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

¹⁴Note that $\mu_0^{r,p}$ is invariant to positive rescalings of p : that is, beliefs depend only on the relative probabilities under p .

¹⁵We thank Laura Doval for suggesting this interpretation.

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. Definition 3 resolves this ambiguity by requiring that ranked probabilities propagate along paths of play in the sense of condition (2).¹⁶

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

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. Since a vector of ranked probabilities is a simpler object than a CPS, this characterization simplifies the task of checking whether a given assessment is a CPPBE (as we will see).

As mentioned in the Introduction, it is well-known that, in general, there is an equivalence between the sets of CPS's and **lexicographic conditional probability systems** (LCPS's; i.e., ordered sequences of probability measures with disjoint support) on a given state space (Blume et al., 1991a,b; Hammond, 1994; Halpern, 2010). Theorem 1 differs from these results in that it establishes an equivalence between two notions of consistent assessments (consistency with a CPS on Z and with a vector of ranked probabilities), rather than two notions of probability, and, moreover, the CPS's and ranked probabilities in the theorem have different domains (Z and X).

To connect Theorem 1 and these prior results explicitly, we prove the theorem in two steps: establishing a bijection between CPS's on Z and ("reduced") ranked probabilities on X (which recapitulates prior results), and verifying that this bijection preserves consistency (which is novel).¹⁷

We start by introducing some notation. For any ranked probability (r, p) , define the **projection** of (r, p) onto Z as the pair $(\tilde{r}, \tilde{p}) : 2^Z \setminus \{\emptyset\} \rightarrow \mathbb{N}_{\geq 0} \times \mathbb{R}_{++}$ given by

$$\tilde{r}(E) := \min_{z \in E} r_z, \quad \tilde{p}(E) := \sum_{z \in E: r_z = \tilde{r}(E)} p_z \quad \text{for all non-empty } E \subseteq Z.$$

¹⁶ 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 .

¹⁷ A prior version of this paper contained a slightly shorter and more direct proof, which left the relation to CPS–LCPS equivalence implicit.

Let $\mathcal{P}$ the set of projections of all possible ranked probabilities on X . We say that a projection $(\tilde{r}, \tilde{p})$ is **reduced** if ranks are consecutive starting from zero, so $\{\tilde{r}(z) : z \in Z\} = \{0, \dots, \tilde{r}\}$ for some $\tilde{r} \in \mathbb{N}_{\geq 0}$, and the normalization $\sum_{z: \tilde{r}(z)=k} \tilde{p}(z) = 1$ holds for any $k \in \{0, \dots, \tilde{r}\}$. A reduced projection is thus exactly a LCPS on Z (Blume et al., 1991a): the level sets $\{z : \tilde{r}(z) = k\}$ are the disjoint supports of the probability distributions given by restricting $\tilde{p}$ to each level set.

In the first part of the proof, we introduce a bijection Ψ from CPS's on Z to reduced projections, with inverse Φ . First, for any CPS μ on Z , we define the **null operator** $K : 2^Z \setminus \{\emptyset\} \rightarrow 2^Z$ by

$$K(Y) := \{z \in Y : \mu(z|Y) = 0\} \quad \text{for all non-empty } Y \subseteq Z.$$

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 z , there is a unique number m_z such that $\mu(z|K^{m_z}(Z)) > 0$.

We then define the map Ψ from CPS's on Z to reduced projections by $\Psi(\mu) = (\tilde{r}, \tilde{p})$, where

$$\tilde{r}(E) = \min_{z \in E} m_z \geq 0, \quad \tilde{p}(E) = \mu(E|K^{\tilde{r}(E)}(Z)), \quad \text{for all non-empty } E \subseteq Z.$$

Conversely, we define the map Φ from reduced projections $(\tilde{r}, \tilde{p})$ to CPS's on Z by

$$\Phi(\tilde{r}, \tilde{p})(A|B) = \frac{\tilde{p}(A \cap B) \mathbf{1}\{\tilde{r}(A \cap B) = \tilde{r}(B)\}}{\tilde{p}(B)} \quad \text{for all } A \subseteq Z, B \subseteq Z, B \neq \{\emptyset\},$$

with $\Phi(\tilde{r}, \tilde{p})(A|B) = 0$ if $A \cap B = \emptyset$. In the appendix, we verify that Ψ and Φ are mutual inverses. This adapts the standard CPS–LCPS equivalence to our setting.

The second part of the proof verifies that this bijection preserves consistency. To do so, we let Γ map rprojections to the ranked probability given by $r_x = \tilde{r}(Z(x))$ and $p_x = \tilde{p}(Z(x))$ for each node x , and let Λ map a ranked probability to its projection. These functions, composed with Ψ and Φ respectively, define a mapping $\hat{\Psi}$ from CPS's on Z to ranked probabilities on X , and a mapping $\hat{\Phi}$ from ranked probabilities to CPS's. In particular, for the “only if” direction, this lets us select a ranked probability—among those with the same projection—that preserves consistency with the initial assessment. We then show that, if an assessment (β, μ_0) is consistent with a CPS μ on Z , then it is also consistent with the ranked probability $\hat{\Psi}(\mu)$; and, conversely, if it is consistent with a ranked probability (r, p) , then it is also consistent with the CPS $\hat{\Phi}(r, p)$.

The example in Figure 3 illustrates that different ranked probabilities can be consistent with the same strategy profile, 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 observe 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 $\mu_0(x_D|h_D) = 1/3$.

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 3), $2p_{w_C} = 9p_{z_C}$. Therefore, if D never observes 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.¹⁹

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.²⁰

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.²¹ Second, CPS's on Z are defined on observable events, while CPS's on S are not (as noted by Battigalli, 1996).²² On the other hand, a CPS on S explicitly describes off-path strategies, which may be conceptually appealing and may suggest approaches to selecting among CPPBE. The two formulations are thus complementary.

We denote a generic CPS on S by 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.

¹⁹ 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.

²⁰ 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.

²¹ This is roughly analogous to the computational advantage of behavior strategies over mixed strategies (Kuhn, 1953).

²² 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 .

Definition 4 (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 $x \in h$,

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

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

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

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 $\nu((l, a, d)|S(x)) = \beta(a|x)\beta(d|y) = 2/9$.




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

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

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 there is a unique CPS μ on Z consistent with (β, μ_0) . The class of CPS's on S that are equivalent to μ is

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

However, consistency of ν with (β, μ_0) imposes some additional constraints, such as $\nu((l, a.d)|S(x)) = 2/9$ (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 $\nu(\cdot|S(x))$, 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, a CPPBE where β is played pins down a CPS μ on Z up to a continuum (the CPS's differ only on the terminal nodes following x) and each μ is defined on the set of terminal nodes of cardinality $|Z|=6$. However, for any such μ , there are infinitely many CPS's ν on S consistent with both μ and a PBE where β is played, 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 Conclusion

This paper has studied 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.

Appendix A: Proofs

Proof of Theorem 1

Preliminaries and CPS–LCPS Equivalence:

We first establish a property of the null operator K defined in the text.

Claim 1. *For any $z \in Z$, there is a unique $m_z \in \mathbb{N}_{\geq 0}$ such that $\mu(z|K^{m_z}(Z)) > 0$.*

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

Recall also the definition of reduced projection $(\tilde{r}, \tilde{p}) \in \mathcal{P}$ of a ranked probability introduced in the main text. For any fixed set of terminal nodes, let $\mathcal{R}$ be the set of all reduced projections and $\mathcal{C}$ the set of all CPS's on Z . Recall also the maps Ψ and Φ defined in the main text.

We rely on the following lemma that restates, in our notation, results of Blume et al. (1991a) and Hammond (1994) and shows some properties of Ψ and Φ .

Lemma 1. *$\Psi : \mathcal{C} \rightarrow \mathcal{R}$ is a bijection from CPS's on Z onto reduced projections $(\tilde{r}, \tilde{p})$, with inverse $\Phi : \mathcal{R} \rightarrow \mathcal{C}$.*

Proof. Claim 1 implies that Ψ is a well-defined function, while Φ is a function by definition. First, we show that, for any μ , $\Psi(\mu) = (\tilde{r}, \tilde{p})$ is a reduced projection. To see this, define a ranked probability (r, p) as follows. For any node $x \notin Z$, let $r_x \in \mathbb{N}_{\geq 0}$ and $p_x \in (0, 1]$ arbitrarily, with $r_{x_0} = 0$, $p_{x_0} = 1$ for the initial node. For $z \in Z$, let $r_z = m_z$ as defined by Claim 1 and $p_z = \mu(z|K^{r_z}(Z))$. By construction, for any $E \subseteq Z$,

$$\begin{aligned} \tilde{r}(E) &= \min_{z \in E} m_z = \min_{z \in E} r_z, \\ \tilde{p}(E) &= \mu(E|K^{\tilde{r}(E)}(Z)) = \sum_{z \in E} \mu(z|K^{\tilde{r}(E)}(Z)) = \sum_{z \in E: r_z = \tilde{r}(E)} \mu(z|K^{r_z}(Z)) = \sum_{z \in E: r_z = \tilde{r}(E)} p_z. \end{aligned}$$

To see that the other conditions of reduced projections are satisfied, note that since Z is finite, there is a largest rank $\bar{r}$ where $K^{\bar{r}+1}(Z) = \emptyset$. For each $k \in \{0, \dots, \bar{r}\}$, the set of terminal nodes assigned rank k is $K^k(Z) \setminus K^{k+1}(Z)$, which is nonempty because the sets $(K^m(Z))_m$ are

ordered by strict inclusion (so the ranks are consecutive starting from zero), and satisfies

$$\sum_{z: \tilde{r}(z)=k} \tilde{p}(z) = \sum_{z: m_z=k} \mu(z | K^k(Z)) = 1.$$

Thus, $\Psi(\mu)$ is a reduced projection.

Next, we show that, for any reduced projection $(\tilde{r}, \tilde{p})$, $\Phi(\tilde{r}, \tilde{p}) = \mu$ is a CPS on Z . The only non-trivial part is the chain rule. Fix $A \subseteq B \subseteq C \subseteq Z$, $B \neq \emptyset$. Note that $\tilde{r}(B) \geq \tilde{r}(C)$. If $\tilde{r}(B) > \tilde{r}(C)$, then $\mu(A|C) = \mu(B|C) = 0$, and the chain rule trivially holds. Thus, suppose that $\tilde{r}(B) = \tilde{r}(C)$, which implies that $\mu(B|C) > 0$. Then, by definition,

$$\mu(A|B) \cdot \mu(B|C) = \frac{\tilde{p}(A)\mathbf{1}\{\tilde{r}(A) = \tilde{r}(B)\}}{\tilde{p}(B)} \cdot \frac{\tilde{p}(B)}{\tilde{p}(C)} = \frac{\tilde{p}(A)\mathbf{1}\{\tilde{r}(A) = \tilde{r}(C)\}}{\tilde{p}(C)} = \mu(A|C),$$

implying the chain rule.

Finally, we show that, for any $\mu \in \mathcal{C}$, $\Phi(\Psi(\mu)) = \mu$, and that, for any $(\tilde{r}, \tilde{p}) \in \mathcal{R}$, $\Psi(\Phi(\tilde{r}, \tilde{p})) = (\tilde{r}, \tilde{p})$. This shows that Ψ and Φ are bijections and mutual inverses.

For the first part, fix a CPS μ , let $(\tilde{r}, \tilde{p}) = \Psi(\mu)$ (which we have shown is a reduced projection), and fix $z \in B \subseteq Z$ such that $\mu(z|B) > 0$. By the chain rule, this implies that $\mu(z|K^{\tilde{r}(B)}(Z)) > 0$, and thus $\tilde{r}(z) = \tilde{r}(B)$ by the definition of K . Thus,

$$\Phi(\Psi(\mu))(z|B) = \frac{\tilde{p}(z)\mathbf{1}\{\tilde{r}(z) = \tilde{r}(B)\}}{\tilde{p}(B)} = \frac{\mu(z|K^{\tilde{r}(z)}(Z))}{\mu(B|K^{\tilde{r}(B)}(Z))} = \frac{\mu(z|K^{\tilde{r}(z)}(Z))}{\mu(B|K^{\tilde{r}(z)}(Z))} = \mu(z|B),$$

where the first equality follows from definition, the second and third from the steps shown above, and the last from the chain rule. If $z \notin B$, then by definition $\Phi(\Psi(\mu))(z|B) = 0 = \mu(z|B)$. By finite additivity, this reasoning extends to $\Phi(\Psi(\mu))(A|B)$ for $A \subseteq Z$. Thus, $\Phi(\Psi(\mu)) = \mu$.

For the second part, fix a reduced projection $(\tilde{r}, \tilde{p})$ with maximum rank $\bar{r}$, and let $\mu = \Phi(\tilde{r}, \tilde{p})$ (which we have shown is a CPS). For any $m \geq 0$, let $L_m := \{z \in Z : \tilde{r}(z) \geq m\}$. We first show that $K^m(Z) = L_m$ for any $m \geq 0$. We proceed by induction. For $m = 0$, $K^0(Z) = Z = L_0$. Suppose then that $K^m(Z) = L_m$ for some $m \leq \bar{r}$. Since ranks are consecutive, there exists $z \in Z$ with $\tilde{r}(z) = m$, and thus $\tilde{r}(L_m) = m$. Hence, for any $z \in L_m$,

$$\mu(z|L_m) = \frac{\tilde{p}(z)\mathbf{1}\{\tilde{r}(z) = \tilde{r}(L_m)\}}{\tilde{p}(L_m)} = \frac{\tilde{p}(z)\mathbf{1}\{\tilde{r}(z) = m\}}{\tilde{p}(L_m)},$$

which, since $\tilde{p}(z) > 0$ and $\tilde{p}(L_m) > 0$, is zero if and only if $\tilde{r}(z) > m$. By the definition of K ,

$$K^{m+1}(Z) = K(L_m) = \{z \in L_m : \tilde{r}(z) > m\} = L_{m+1},$$

completing the induction. In particular, $K^{\tilde{r}+1}(Z) = \emptyset$.

It follows that, for any $z \in Z$ and $m \geq 0$, we have $z \in K^m(Z) \setminus K^{m+1}(Z)$ if and only if $\tilde{r}(z) = m$. By Claim 1, this implies

$$\Psi(\mu)_1(z) = m_z = \tilde{r}(z).$$

Moreover, taking $m = \tilde{r}(z)$ in the display above and using the normalization $\tilde{p}(L_{\tilde{r}(z)}) = \sum_{z': \tilde{r}(z') = \tilde{r}(z)} \tilde{p}(z') = 1$,

$$\Psi(\mu)_2(z) = \mu(z | K^{\tilde{r}(z)}(Z)) = \frac{\tilde{p}(z)}{\tilde{p}(L_{\tilde{r}(z)})} = \tilde{p}(z).$$

Finally, fix any $E \subseteq Z$, $E \neq \emptyset$. Since both $(\tilde{r}, \tilde{p})$ and $\Psi(\mu)$ are projections,

$$\begin{aligned} \Psi(\mu)_1(E) &= \min_{z \in E} \Psi(\mu)_1(z) = \min_{z \in E} \tilde{r}(z) = \tilde{r}(E), \\ \Psi(\mu)_2(E) &= \sum_{\substack{z \in E: \\ \Psi(\mu)_1(z) = \Psi(\mu)_1(E)}} \Psi(\mu)_2(z) = \sum_{\substack{z \in E: \\ \tilde{r}(z) = \tilde{r}(E)}} \tilde{p}(z) = \tilde{p}(E). \end{aligned}$$

Thus, $\Psi(\Phi(\tilde{r}, \tilde{p})) = (\tilde{r}, \tilde{p})$. □

Next, define $\Gamma : (\tilde{r}, \tilde{p}) \mapsto (r, p)$ such that, for any $(\tilde{r}, \tilde{p}) \in \mathcal{P}$ and $x \in X$,

$$(\Gamma(\tilde{r}, \tilde{p}))_x = \left(\min_{z \in Z(x)} \tilde{r}(z), \sum_{z \in Z(x): \tilde{r}(z) = \min_{z' \in Z(x)} \tilde{r}(z')} \tilde{p}(z) \right).$$

Note that the resulting ranked probability has a reduced projection. We also let $\Lambda : (r, p) \mapsto (\tilde{r}, \tilde{p})$ denote the map from each ranked probability into its projection. Thus, $\Lambda \circ \Gamma$ is the identity on reduced projections.

"Only if" Direction:

Let (β, μ_0) be an assessment consistent with a CPS μ on Z . We show that it is also consistent with the ranked probability $(r, p) = \Gamma(\Psi(\mu))$. Recall that we use the notation $\mu(A|B) = \mu(Z(A)|Z(B))$ for any $A, B \subseteq X$, $B \neq \emptyset$, when this will not cause confusion.

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 $Z(h) \subseteq K^{r_h}(Z)$

and $\mu(h|K^{r_h}(Z)) > 0$. Moreover, for any $x \in h$, we have

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

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 $\mu(\{x\}|K^{r_h}(Z)) = 0$, which implies $r_x > r_h$ and $\mu_0^{r,p}(x|h) = 0$. If instead $\mu_0(x|h) > 0$, then $\mu(\{x\}|K^{r_h}(Z)) > 0$ and so $r_x = r_h$. Thus,

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

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 3. Fix any $x, x' \in X$ with $x' \succ_0 x$. Since $Z(x') \subseteq Z(x)$ and $r_x = \min_{z \in Z(x)} r_z$, this implies $\min_{z \in Z(x')} r_z \geq \min_{z \in Z(x)} r_z$, and hence $r_{x'} \geq r_x$. Moreover, suppose $\mu_0(x|h) > 0$. By definition of Γ and Ψ , $r_x = \min_{z \in Z(x)} \tilde{r}(z) = \min_{z \in Z(x)} m_z$. By definition of K and since $\mu_0(x|h) > 0$, for any $z \in Z(x)$, $m_z = r_x$ if and only if z is reached with positive probability from x . Thus, if $P^\beta(x'|x) = 0$, this implies that, for any $z' \in Z(x')$, $m_{z'} > r_x$, implying $r_{x'} > r_x$. Conversely, if $P^\beta(x'|x) > 0$, then for any $z' \in Z(x')$, $m_{z'} = r_x$ if and only if z' is reached with positive probability from x' . Thus, $r_x = r_{x'}$. Finally, letting $Q(x)$, $Q(x') \subseteq Z$ respectively the sets of terminal nodes in $Z(x)$ and $Z(x')$ reached with positive probability from x and x' ,

$$p_{x'} = \sum_{z' \in Q(x')} p(z') = \mu(Q(x')|K^{r_{x'}}(Z)) = \mu(Q(x')|Q(x))\mu(Q(x)|K^{r_x}(Z)) = P^\beta(x'|x)p_x.$$

where the first equality follows from definition of Γ and Ψ , and the third one by the chain rule applied to $Q(x') \subseteq Q(x)$.

"If" Direction:

Let (β, μ_0) be an assessment consistent with a ranked probability (r, p) on X . We show that it is also consistent with the CPS $\mu = \Phi(\Lambda(r, p))$.

We start with a preliminary result, which shows that (r, p) agrees with its reduction $\Gamma(\Lambda(r, p))$ at every node assigned positive probability.

Claim 2. *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) $\tilde{r}(Z(x)) = r_x$ and $\tilde{p}(Z(x)) = p_x$.

Proof. For any $z \in Z(x)$, Definition 3 implies that if $P^\beta(z|x) > 0$ then $r_z = r_x$, while if $P^\beta(z|x) = 0$ then $r_z > r_x$, proving point (1). Moreover, there is at least one $z \in Z(x)$ with $P^\beta(z|x) > 0$, and thus $\tilde{r}(Z(x)) = \min_{z \in Z(x)} r_z = r_x$. Finally, using point (1) and the fact that consistency with (r, p) implies $p_z = P^\beta(z|x)p_x$ for any such z ,

$$\tilde{p}(Z(x)) = \sum_{\substack{z \in Z(x): \\ r_z = r_x}} p_z = \sum_{\substack{z \in Z(x): \\ P^\beta(z|x) > 0}} p_z = p_x \sum_{\substack{z \in Z(x): \\ P^\beta(z|x) > 0}} P^\beta(z|x) = p_x. \quad \square$$

We are now ready to establish consistency.

Claim 3. For any $h \in H$ and $x \in h$, $\mu_0(x|h) = \mu(Z(x)|Z(h))$.

Proof. To ease notation, let $\rho := \tilde{r}(Z(h))$. 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 2 implies $\min_{y \in h} r_y = \rho$.

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 $\tilde{r}(Z(x)) \geq r_x$, and the preceding observations imply $r_x > \rho$. Thus $\tilde{r}(Z(x)) > \tilde{r}(Z(h))$ and

$$\mu(Z(x)|Z(h)) = 0 = \mu_0(x|h).$$

Now suppose that $\mu_0(x|h) > 0$. Since the sets $\{Z(y)\}_{y \in h}$ are disjoint and their union is $Z(h)$, and since $\tilde{r}(Z(y)) > \rho$ whenever $\mu_0(y|h) = 0$, point (2) of Claim 2 gives

$$\tilde{p}(Z(h)) = \sum_{y \in h} \sum_{\substack{z \in Z(y): \\ r_z = \rho}} p_z = \sum_{\substack{y \in h: \\ \mu_0(y|h) > 0}} \tilde{p}(Z(y)) = \sum_{\substack{y \in h: \\ \mu_0(y|h) > 0}} p_y.$$

Therefore, using point (2) of Claim 2 once more,

$$\mu(Z(x)|Z(h)) = \frac{\tilde{p}(Z(x))}{\tilde{p}(Z(h))} = \frac{p_x}{\sum_{y \in h: \mu_0(y|h) > 0} p_y} = \mu_0^{r,p}(x|h) = \mu_0(x|h). \quad \square$$

Claim 4. For any $x < x'$ with $\mu_0(x|h(x)) > 0$, we have $\mu(Z(x')|Z(x)) = P^\beta(x'|x)$.

Proof. By point (2) of Claim 2, $\tilde{r}(Z(x)) = r_x$ and $\tilde{p}(Z(x)) = p_x$.

If $P^\beta(x'|x) = 0$, then consistency with (r, p) implies $r_{x'} > r_x$, and hence $\tilde{r}(Z(x')) \geq r_{x'} > r_x = \tilde{r}(Z(x))$. Since $Z(x') \subseteq Z(x)$, it follows that

$$\mu(Z(x')|Z(x)) = 0 = P^\beta(x'|x).$$

Now suppose that $P^\beta(x'|x) > 0$, so that consistency with (r, p) gives $r_{x'} = r_x$ and $p_{x'} = P^\beta(x'|x)p_x$. Fix any $z \in Z(x')$. Since $P^\beta(z|x) = P^\beta(z|x')P^\beta(x'|x)$ and $P^\beta(x'|x) > 0$, we have $P^\beta(z|x) > 0$ if and only if $P^\beta(z|x') > 0$. Thus, by point (1) of Claim 2 applied to the pair $x < z$, we have $r_z = r_x = r_{x'}$ if and only if $P^\beta(z|x') > 0$, and in that case consistency with (r, p) gives $p_z = P^\beta(z|x)p_x = P^\beta(z|x')p_{x'}$. Hence

$$\tilde{p}(Z(x')) = \sum_{\substack{z \in Z(x'): \\ r_z = r_{x'}}} p_z = p_{x'} \sum_{\substack{z \in Z(x'): \\ P^\beta(z|x') > 0}} P^\beta(z|x') = p_{x'},$$

and $\tilde{r}(Z(x')) = r_{x'} = r_x = \tilde{r}(Z(x))$. Therefore

$$\mu(Z(x')|Z(x)) = \frac{\tilde{p}(Z(x'))}{\tilde{p}(Z(x))} = \frac{p_{x'}}{p_x} = P^\beta(x'|x). \quad \square$$

Proof of Theorem 2

Recall that we write $h < x'$ when there is $x \in h$ such that $x < x'$, and similarly for $h \not< x'$.

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 $x < x'$, 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 $h < x'$. Since $\mu_0(x | h(x)) > 0$, Doval and Ely's (2020) definition of PBE gives $\nu(s | S(x)) = \prod_{h \not< x} \beta(s(h) | h)$.

Summing this over all strategy profiles that reach x' gives

$$\mu(Z(x') | Z(x)) = \sum_{s \in S(x')} v(s | S(x)) = \sum_{s \in S(x')} \prod_{h \prec x} \beta(s(h) | h) = P^\beta(x' | x).$$

By the chain rule, for any $z \in Z(x')$, we have $\mu(\{z\} | Z(x)) = \mu(Z(x') | Z(x))\mu(\{z\} | Z(x'))$. Substituting, we have

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

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

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

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

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

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

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 . Additionally, Definition 4 imposes that agents should see as probabilistically independent off-path nodes not preceding z . Thus, agents' beliefs about play at nodes not preceding z should be consistent with independent trembles—or, equivalently, their beliefs can allow for correlated trembles only at nodes preceding z .

Finally, passing to a subsequence where each $\sigma_\varepsilon(s|B)$ converges if necessary, define ν by

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

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 $\nu(s|S(x)) = \prod_{h' \prec x} \beta(s(h')|h')$ 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

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

where the second equality follows because $S(x) = \cup_{z \in Z(x)} S(z)$, the sets $\{S(z)\}_{z \in Z(x)}$ are mutually 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 $z_s \succ x$.

Appendix B: Relation to *No-Signaling-What-You-Don't-Know*

In this section, we compare CPPBE to another closely related solution concept, No-signaling-what-you-don't-know (NSWYDK, Fudenberg and Tirole, 1991). NSWYDK 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.”²⁴

Definition 6 (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 $x <_0 x'$,*

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

(2) *for any $h \in H$, $x, y \in h$, $x <_0 x'$ and $y <_0 y'$ such that $a_{x,x'} = a_{y,y'}$,*²⁵

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

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 “tree-extended assessment.” Note that, if this assessment satisfies sequential rationality, it is already a stronger solution concept than CPPBE, since condition (1) is imposed also at nodes with zero conditional probability.

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.²⁶ We thus have:

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

²⁴ This generalizes Fudenberg and Tirole’s “perfect extended Bayesian equilibrium,” which is defined for multistage games with observed actions, to general games.

²⁵ Recall that if $x <_0 x'$ then $a_{x,x'} \in \mathcal{A}(h(x))$ denotes the action leading from x to x' .

²⁶ 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.

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 $\mu(Z(x_C)|Z(x_C) \cup Z(y_C)) = \mu(Z(x_B)|Z(x_B) \cup Z(y_B))$, 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 6 refers to the condition that $a_{x,x'} = a_{y,y'}$, which seems difficult to express with ranked probabilities.




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 do 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 from before—namely, $\mu_0(x_C|h_C) = \mu_0(y_C|h'_C)$.²⁷ Therefore, any characterization of NSWYDK based on ranked probabilities would require

²⁷ 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:

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

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).

Appendix C: Heterogeneous CPPBE and Plain PBE

Heterogeneous CPPBE

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. We define a CPPBE with heterogeneous beliefs as follows.

Definition 7. 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 $x \in h$,

$$\mu_0(x|h) = \mu_i(Z(x)|Z(h)), \quad \text{and}$$

(2) for any $x, x' \in X$ with $x < x'$, 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 players. 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. However, in general, heterogeneous CPPBE need not be subgame perfect.

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

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

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.²⁸ 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 $s(\underline{h}_i) = \text{wait}$ for every i not active at h_0 .

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

Definition 9 (Watson, 2025b). *For any agent i and history $h \in \underline{H}_i$, an **appraisal** p_h is a distribution on S such that*

- (1) $\text{supp } p_h \subseteq S(h)$;
- (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 $x < x'$, and for any such x, x' there does not exist $h'' \in \underline{H}_i$ and $x'' \in h''$ such that $x < x'' < x'$.

Definition 10 (Watson, 2025b). *The appraisal system P satisfies **plain consistency** if, for every agent i , every pair $h, h' \in \underline{H}_i$ of consecutive histories, every $L \subseteq \underline{H}$, and every L -product*

²⁸ Formally, $\underline{h}_i$ precedes any $h \in H_i$ and $S(\underline{h}_i) = S$. Moreover, $\mathcal{A}(\underline{h}_i) = \{\text{wait}\}$ for all i except j active at h_0 , and $\underline{h}_j = h_0$.

set $Y \subseteq S(h)$ such that p_h exhibits L -independence on Y ,

$$\frac{p_{h'}(B_L \times \{s_{-L}\})}{p_{h'}(Y_L \times \{s_{-L}\})} = \frac{p_h(B_L \times Y_{-L})}{p_h(Y)}$$

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

To get an intuition, suppose that $h, h' \neq h_0$, agent $j \neq i$ is active at h_0 , and i does not observe j 's choice. Fix $L = \{h_0\}$. Plain consistency implies that, if at h agent i believes that j 's trembles at h_0 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 h_0 are independent at h' .

The definition of plain PBE can be stated as follows.

Definition 11. An appraisal system $P = (p_h)_{h \in \underline{H}}$ is a **plain PBE** if there exists a behavior strategy β such that

- (1) for all i and s , $p_{h_i}(s) = \prod_{h \in H} \beta(s_h | h)$;
- (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 3.

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.²⁹

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.

²⁹Formally, let $\zeta : S \rightarrow Z$ be the outcome map, $\zeta(s)$ being the terminal node reached under s , and let $\zeta_{\#} p \in \Delta(Z)$ denote the pushforward of $p \in \Delta(S)$. Since $S(x) = \zeta^{-1}(Z(x))$, an appraisal at h satisfies $\zeta_{\#} p_h \in \Delta(Z(h))$, the same simplex as $\mu_i(\cdot | Z(h))$. A plain PBE $P = (p_h)_{h \in \underline{H}}$ consistent with β is equivalent to a heterogeneous CPPBE (β, μ_0) consistent with $(\mu_i)_i$ if $\zeta_{\#} p_h = \mu_i(\cdot | Z(h))$ for all i and $h \in H_i$, and $\zeta_{\#} p_{h_i} = \mu_i(\cdot | Z)$.




Figure 8: 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 10 to restrict the beliefs at h_C and h'_C , we need to fix L containing the initial history h_0 . However, this implies that any $Y \subseteq S$ satisfying $Y_L \times S_{-L} \subseteq S(h_C) \cup S(h'_C)$ does not contain any strategy in which A plays *out*. At the same time, any Y satisfying $p_{\underline{h}_C}(Y) > 0$ needs to contain at least one of these strategies. Thus, plain consistency has no bite in pinning down Y '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.




Figure 9: 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 $\frac{1}{4}$. To see this, apply Definition 10 to agent C 's con-

secutive pair (h_C, h_C) , taking $L = \{h_B\}$ and $Y = \{b, c\} \times S_{-L}$. Note that $S(h_C)$ is an L -product set with $S(h_C)_L = \{b, c\}$: conditional on B playing b or c , reaching h_C does not depend on which of the two, so $Y_L \times \{s_{-L}\} \subseteq S(h_C)$ for every $s_{-L} \in Y_{-L}$ with $p_{h_C}(Y_L \times \{s_{-L}\}) > 0$. Moreover, p_{h_C} is the product distribution induced by β and hence exhibits L -independence on Y . Plain consistency then pins down C 's belief that B played c conditional on each such s_{-L} , and since the value is the same for all of them,

$$\mu_0(x|h_C) = \frac{p_{h_C}(\{c\} \times Y_{-L})}{p_{h_C}(Y)} = \frac{\beta(c|h_B)}{\beta(b|h_B) + \beta(c|h_B)} = \frac{0.2}{0.8} = \frac{1}{4}.$$

Appendix D: Epistemic Foundations

In the Introduction, we argued that CPPBE imposes sequential rationality, Bayesian updating, and nothing more. We now formalize this claim.

To do so, we first introduce *a posteriori CPPBE*, a version of CPPBE where Bayesian updating and sequential rationality are retained for each agent separately, while the assumptions that tie different agents' beliefs together are dropped.³⁰ This is the notion of CPPBE for which we provide an epistemic foundation. In Remark 1, we make the relation between the two concepts precise and show when CPPBE corresponds to a posteriori CPPBE.

We then recall the definition of *Backward Rationalizability* (BR), a solution concept introduced by Catonini and Penta (2026) to capture the logic of backward induction reasoning, and nothing more. Catonini and Penta also introduce *interim perfect equilibrium* (IPE), a solution concept for sequential games with incomplete information. In Proposition 3, we show that, for any fixed type space, a posteriori CPPBE and IPE induce the same strategies and the same beliefs about opponents' types at every history. This implies that BR characterizes the robust predictions of a posteriori CPPBE. Together with Remark 1, this also implies that, if a strategy is played in a CPPBE and if agents' beliefs satisfy a minimal agreement requirement, then the strategy is backward rationalizable. Conversely, if a strategy is backward rationalizable in the canonical type space and beliefs satisfy a notion equivalent to the sequential version of common prior, then the strategy is played in a CPPBE. In this sense, along with the standard description of agents' beliefs by a common belief system, CPPBE imposes sequential rationality, Bayesian updating, and

³⁰ This solution concept is obtained from CPPBE in the same spirit as Brandenburger and Dekel's (1987) *a posteriori* equilibrium.

nothing more.

D.1 A Posteriori CPPBE

We work in the setting of multistage games with observed actions and unobserved moves by Nature, as this is the setting studied by Catonini and Penta (2026). Following our baseline model, players are indexed by $i \in I$ and information sets by $h \in H$. Each player i has a payoff function $u_i : Z \times \Theta \rightarrow \mathbb{R}$ where $\Theta = \Theta_0 \times \dots \times \Theta_n$ is finite. Payoff types θ_i are private information of each agent i , while the state of nature θ_0 is unobserved. Let $G := \langle I, H, Z, \Theta_0, (\Theta_i, u_i)_{i \in I} \rangle$ be the resulting game with payoff uncertainty.

We introduce a type space $\mathcal{T} = (\Omega, \vartheta_0, (\pi_i, \tau_i, T_i, \vartheta_i)_{i \in I})$ to describe heterogeneous belief hierarchies. Let Ω be the finite uncertainty space and $\vartheta_0 : \Omega \rightarrow \Theta_0$ the map specifying the state of nature. For each agent i , π_i is i 's prior over Ω and T_i is the finite set of types. The map $\tau_i : \Omega \rightarrow T_i$ assigns a type to i and $\vartheta_i : T_i \rightarrow \Theta_i$ is an onto function specifying a payoff parameter for each type t_i .

When a player i observes type t_i , i deduces that the true underlying state must belong to the set $\tau_i^{-1}(t_i) \subseteq \Omega$, and assigns to events $E \subseteq \Omega$ a probability computed according to $\pi_i(E|\tau_i^{-1}(t_i))$.³¹ In particular, i 's initial belief about other agents' types and the state of nature is given by, for any t_{-i} and θ_0 , $\pi_i(\vartheta_0^{-1}(\theta_0) \cap (\cap_{j \neq i} \tau_j^{-1}(t_j)) | \tau_i^{-1}(t_i))$. To ease notation, we label this belief $\pi_i(\theta_0, t_{-i} | t_i)$. Let $G^{\mathcal{T}}$ be the game expanded with the type space.

In $G^{\mathcal{T}}$, we introduce two new objects to model agents' beliefs and actions throughout the game. For each agent i and type t_i , define a **Bayesian behavior strategy** as a function b_{t_i} mapping each h at which i is active into a distribution over $\mathcal{A}(h)$.³² For any agent i , let $b_i := (b_{t_i})_{t_i \in T_i}$.

Definition 12. For each type t_i , a **Bayesian CPS** $\hat{\mu}^{t_i}$ over $\Theta_0 \times T_{-i} \times Z$ is an array of conditional probabilities $(\hat{\mu}^{t_i}(\cdot | Z(h)))_{h \in H} \in \Delta(\Theta_0 \times T_{-i} \times Z)$ such that:

- (1) for any $h \in H$, $\hat{\mu}^{t_i}(\Theta_0 \times T_{-i} \times Z(h) | Z(h)) = 1$;
- (2) for any $h < h'$ and $E \subseteq \Theta_0 \times T_{-i} \times Z(h')$,

$$\hat{\mu}^{t_i}(E | Z(h)) = \hat{\mu}^{t_i}(E | Z(h')) \cdot \hat{\mu}^{t_i}(\Theta_0 \times T_{-i} \times Z(h') | Z(h)); \quad \text{and}$$

³¹ We assume that $\pi_i(\tau_i^{-1}(t_i)) > 0$ for any t_i , that is, i assigns positive probability to each of his types.

³² For notational convenience, we work with behavior strategies. Catonini and Penta define IPE with mixed strategies and observe that, for a fixed type space, there may be more IPE in mixed than in behavior strategies; they also note that Proposition 2 holds for IPE in behavior strategies as well.

(3) for any $\theta_0 \in \Theta_0$ and $t_{-i} \in T_{-i}$, $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z | Z) = \pi_i(\theta_0, t_{-i} | t_i)$.

From $\hat{\mu}^{t_i}$, we can define an equivalent Bayesian CPS over $\Theta_0 \times \Theta_{-i} \times Z$ by setting, for any h, θ_0, θ_{-i} , and z ,

$$\mu^{t_i}((\theta_0, \theta_{-i}, z) | Z(h)) := \hat{\mu}^{t_i}(\{\theta_0\} \times \mathfrak{S}_{-i}^{-1}(\theta_{-i}) \times \{z\} | Z(h)).$$

In a similar way, for type t_i , we can define a Bayesian CPS $\hat{\nu}^{t_i}$ over $\Theta_0 \times T_{-i} \times S_{-i}$ conditioning on $(S_{-i}(h))_{h \in H}$ and the corresponding Bayesian CPS ν^{t_i} over $\Theta_0 \times \Theta_{-i} \times S_{-i}$. Note that we specify a Bayesian CPS over the opponents' strategies, since this is the information about the continuation game that each agent uses when choosing an action.

We extend the definition of probability P^β over nodes to Bayesian behavior strategies. Fix two nodes $x <_0 x'$ and call $\iota(x)$ and $\iota(h)$ the agent active at x and h respectively. For any type profile $t = (t_j)_{j \in I}$ and strategy profile $b = (b_i)_{i \in I}$, let

$$P^{b,t}(x' | x) := b_{t_{\iota(h(x))}}(a_{x,x'} | h(x)).$$

In words, $P^{b,t}(x' | x)$ coincides with the probability that the agent active at $h(x)$ plays the action leading to x' given his behavior strategy and type. For any $h <_0 h'$, let also $P^{b,t}(h' | h) = P^{b,t}(x' | x)$ for any $x \in h, x' \in h'$, with $x <_0 x'$; it is well-defined since G is a multistage game with observed players' actions.

Definition 13. A tuple $(b_i, (\hat{\mu}^{t_i})_{t_i \in T_i})_{i \in I}$ is an **a posteriori CPPBE** of G^T if, for any i and t_i ,

- (1) for any s_i , if $b_{t_i}(s_i(h) | h) > 0$ for all h where i is active, then s_i is sequentially rational given μ^{t_i} ;
- (2) for any $\theta_0, t_{-i}, h <_0 h'$, and $\iota(h) = j$, if $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h) | Z(h)) > 0$ and $j \neq i$, then:

$$\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h') | Z(h)) = P^{b, (t_i, t_{-i})}(h' | h) \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h) | Z(h)),$$

while if $j = i$ then $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z | Z(h')) = \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z | Z(h))$.

In words, condition (1) states that every strategy a type plays with positive probability in equilibrium must be sequentially rational given that type's beliefs. Condition (2) requires that, whenever a player assigns positive probability to a profile of co-players' types and a state of nature, he should update his belief from there consistently with the strategy profile. This parallels condition (2) of Definition 1. However, because different players

may now update beliefs differently, condition (2) here also requires that every agent does not update his beliefs on co-players and the state of nature after observing his own moves.

The next remark clarifies the relation between a posteriori CPPBE and CPPBE. Intuitively, it provides conditions such that, in multistage games with observed actions and moves by Nature, a CPPBE corresponds to an a posteriori CPPBE, and vice versa.

The first implication relies on two conditions, which we label **common support** and **own-move invariance**. With Nature's moves, agents may disagree on beliefs in a CPPBE, since, even if all beliefs are described by the same CPS, different agents have different initial information, and thus use different conditioning events. If at some information set agents disagree on which nodes have positive probability, CPPBE consistency implies that all agents are free to hold any belief after nodes that the active agent deems impossible, while a posteriori CPPBE, IPE, and BR, apply Bayes' rule to update the belief of each agent that deemed these nodes possible. To rule out this discrepancy, we assume that the CPS satisfies *common support*: at any history, all players agree on which nodes have positive conditional probability.

Next, own-move invariance captures the fact that in solution concepts like IPE, BR, and a posteriori CPPBE, each agent's own moves are not new information. For IPE and BR, this holds because CPS's are defined on S_{-i} ; for a posteriori CPPBE, it is implied by the last condition in Definition 13. Therefore, these solution concepts prevent a player from changing beliefs following their own deviation, which in general is possible in a CPPBE. To rule this out, we assume that the CPS satisfies *own-move invariance*: each player's conditional beliefs about $\Theta_0 \times \Theta_{-i}$ coincide before and after each of their own moves.

For the second statement, from a posteriori CPPBE to CPPBE, the relevant sufficient condition is existence of a common CPS on the uncertainty space $\Theta \times Z$ from which all the Bayesian CPS are derived. This is the dynamic counterpart of a common prior, assuring that agents' beliefs agree not only at the initial history but also after any zero-probability event.

Let $\mathcal{T}^*$ denote the canonical type space where $\Omega = \Theta$, $T_i = \Theta_i$, ϑ_0 and τ_i denote the coordinate projections respectively on Θ_0 and T_i , ϑ_i denotes the identity map, and each prior π_i equals a common prior $\pi \in \Delta(\Omega)$. In games with unobserved moves by Nature, Definition 1 of CPPBE extends with behavior strategies $\beta_i(\cdot|h, \theta_i)$ conditioning on private information as well as on histories, and with CPS's μ on $\Theta \times Z$.³³

³³ The second statement of Remark 1 can be extended to any type space such that $b_{t_i} = b_{t'_i}$ for any types such that $\vartheta_i(t_i) = \vartheta_i(t'_i)$.

Remark 1. Fix a game G .

- (1) Let (β, μ_0) be a CPPBE consistent with a CPS μ that satisfies common support and own-move invariance, and let π be the marginal on Θ of $\mu(\cdot | \Theta \times Z)$. Then $((b_{\theta_i}, \hat{\mu}^{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ is an a posteriori CPPBE of G^{T^*} , where $b_{\theta_i} := \beta_i(\cdot | \cdot, \theta_i)$ and $\hat{\mu}^{\theta_i}(\cdot | Z(h)) := \mu(\cdot | \{\theta_i\} \times \Theta_{-i} \times Z(h))$ for every i , θ_i , and h .
- (2) Conversely, let $((b_{\theta_i}, \hat{\mu}^{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ be an a posteriori CPPBE of G^{T^*} whose belief systems admit a common CPS with common support: a CPS ω on $\Theta \times Z$ such that each $\hat{\mu}^{\theta_i}(\cdot | Z(h))$ is the conditional of ω given $\{\theta_i\} \times \Theta_{-i} \times Z(h)$. Then, (β, μ_0) , where $\beta_i(\cdot | h, \theta_i) := b_{\theta_i}(\cdot | h)$ for any θ_i and h and μ_0 is the initial element of ω , is a CPPBE of G consistent with ω .

D.2 Backward Rationalizability

For the reader's convenience, we first restate the definition of BR. Similarly to extensive-form rationalizability, BR is an iterative deletion procedure. For any $i \in I$ and $\theta_i \in \Theta_i$, let $BR_i^0(\theta_i) = S_i$. Then, define, for any $k > 0$,

$$BR_{-i}^{k-1} := \{(\theta_j, s_j)_{j \neq i} \in \Theta_{-i} \times S_{-i} : \forall j \neq i, s_j \in BR_j^{k-1}(\theta_j)\},$$

and let $s_i \in BR_i^k(\theta_i)$ if there is a CPS ν^i satisfying conditions (1) and (2) of Definition 12 over $\Theta_0 \times \Theta_{-i} \times S_{-i}$ and conditioning on $(S_{-i}(h))_{h \in H}$ such that:

- (1) s_i is sequentially rational given ν^i ;
- (2) for each $h \in H$ and $(\theta_{-i}, s_{-i}) \in \Theta_{-i} \times S_{-i}$, if $\nu^i(\Theta_0 \times \{(\theta_{-i}, s_{-i})\} | S(h)) > 0$, then there exists $s'_{-i} \in S_{-i}$ such that s'_{-i} and s_{-i} coincide on the equilibrium path starting from h and $(\theta_{-i}, s'_{-i}) \in BR_{-i}^{k-1}$.

The set of *backward rationalizable strategies* for type θ_i is $BR_i(\theta_i) = \bigcap_{k>0} BR_i^k(\theta_i)$. We let $BR_i := \{(\theta_i, s_i) \in \Theta_i \times S_i : s_i \in BR_i(\theta_i)\}$ and $BR := \times_{i \in I} BR_i$.

Intuitively, at each round of deletion, s_i survives for type θ_i if it is justified by a CPS that assigns probability one to opponents' strategies whose continuation is consistent with the previous round of deletion. Note in particular that the CPS need not be concentrated on strategies that themselves survive the previous round of deletion. This lets agents form any belief after a deviation about what actions led to the deviation.

Catonini and Penta (2026) define IPE as follows. For each agent i , let $p_i : T_i \rightarrow (\Delta(\Theta_0 \times T_{-i}))^H$ be the function mapping each of i 's types into an array of beliefs about the oppo-

nents' types for each history.³⁴

Definition 14. An assessment $(b_i, p_i)_{i \in I}$ is an *interim perfect equilibrium* if, for all i , θ_0 , and t_i ,

- (1) for any t_{-i} , $p_i(\theta_0, t_{-i} | t_i, h_0) = \pi_i(\theta_0, t_{-i} | t_i)$;
- (2) for any $h <_0 h'$, if $\sum_{(\theta'_0, t'_{-i})} P^{b, (t_i, t'_{-i})}(h' | h) p_i(\theta'_0, t'_{-i} | h, t_i) > 0$ then $p_i(\cdot | h', t_i)$ is derived from $p_i(\cdot | h, t_i)$ and b_{-i} by Bayes's rule, meaning that, for any (θ_0, t_{-i}) ,

$$p_i(\theta_0, t_{-i} | h', t_i) = \frac{P^{b, (t_i, t_{-i})}(h' | h) p_i(\theta_0, t_{-i} | h, t_i)}{\sum_{(\theta'_0, t'_{-i})} P^{b, (t_i, t'_{-i})}(h' | h) p_i(\theta'_0, t'_{-i} | h, t_i)};$$

- (3) for any s_i with $b_{t_i}(s_i(h) | h) > 0$ for all h where i is active, s_i is sequentially rational given the CPS ν^{t_i} over $\Theta_0 \times \Theta_{-i} \times S_{-i}$ induced by $p_i(\cdot | t_i)$ and b_{-i} .

Catonini and Penta (2026) prove the following result.

Proposition 2. Fix a game G . For each $i \in I$, $(\theta_i, s_i) \in BR_i$ if and only if there exists a type space $\mathcal{T}$, an IPE $(b_i, p_i)_{i \in I}$ of $G^{\mathcal{T}}$, and a type $t_i \in T_i$ such that $\vartheta_i(t_i) = \theta_i$ and $b_{t_i}(s_i(h) | h) > 0$ for any h where i is active.

We show that this result implies that BR also characterizes the robust predictions of a posteriori CPPBE: that is, the predictions that do not depend on exogenous assumptions about players' beliefs about co-players. This follows from the next proposition, which shows that, for every IPE, there is an a posteriori CPPBE where the players follow the same strategies and hold the same beliefs about opponents' types at each history, and vice versa.

Proposition 3. Fix a game G and a type space $\mathcal{T}$. For any IPE $(b_i, p_i)_{i \in I}$ there is an a posteriori CPPBE $(b_i, (\hat{\mu}^{t_i})_{t_i \in T_i})_{i \in I}$ such that, for every i , t_i , h , θ_0 , and t_{-i} ,

$$p_i(\{(\theta_0, t_{-i})\} | h, t_i) = \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z | Z(h)).$$

Moreover, the converse result holds for any a posteriori CPPBE $(b_i, (\hat{\mu}^{t_i})_{t_i \in T_i})_{i \in I}$.

Together, Proposition 2, Proposition 3, and Remark 1 imply that every strategy played in a

³⁴ IPE can be seen as a version of a posteriori PBE defined on strategy profiles, obtained from Doval and Ely's (2020) in the same spirit as a posteriori CPPBE.

CPPBE satisfying common support and own-move invariance is backward rationalizable. Conversely, if a strategy is backward rationalizable in the canonical type space and the a posteriori CPPBE where it is played admits a common CPS with common support, then the strategy is played in a CPPBE. We now prove Proposition 3.

Proof. From IPE to a posteriori CPPBE: Fix an IPE $(b_i, p_i)_{i \in I}$. For any i and t_i , we construct a Bayesian CPS $\hat{v}^{t_i}$ on $\Theta_0 \times T_{-i} \times S_{-i}$ recursively as follows. Recall that h_0 represents the initial history. For any θ_0, t_{-i} , and s_{-i} , let

$$\begin{aligned} \hat{v}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i} | S_{-i}(h_0)) &= \pi_i(\theta_0, t_{-i} | t_i), \\ \text{and } \hat{v}^{t_i}((\theta_0, t_{-i}, s_{-i}) | S_{-i}(h_0)) &= \hat{v}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i} | S_{-i}(h_0)) \prod_{j \neq i} \prod_{h: \iota(h)=j} b_{t_j}(s_j(h) | h). \end{aligned} \quad (1)$$

Then, suppose $h <_0 h'$. If $\hat{v}^{t_i}(\Theta_0 \times T_{-i} \times S_{-i}(h') | S_{-i}(h)) > 0$, for each $\{\theta_0, t_{-i}, s_{-i}\} \in \Theta_0 \times T_{-i} \times S_{-i}(h')$, we let

$$\hat{v}^{t_i}((\theta_0, t_{-i}, s_{-i}) | S_{-i}(h')) = \frac{\hat{v}^{t_i}((\theta_0, t_{-i}, s_{-i}) | S_{-i}(h))}{\hat{v}^{t_i}(\Theta_0 \times T_{-i} \times S_{-i}(h') | S_{-i}(h))}. \quad (2)$$

Suppose then that $\hat{v}^{t_i}(\Theta_0 \times T_{-i} \times S_{-i}(h') | S_{-i}(h)) = 0$. For each θ_0, t_{-i} , and s_{-i} , if $s_{-i} \in S_{-i}(h')$ we fix

$$\hat{v}^{t_i}((\theta_0, t_{-i}, s_{-i}) | S_{-i}(h')) = p_i(\theta_0, t_{-i} | h', t_i) \prod_{j \neq i} \prod_{h'' \prec h', \iota(h'')=j} b_{t_j}(s_j(h'') | h''). \quad (3)$$

If $s_{-i} \notin S_{-i}(h')$, we let $\hat{v}^{t_i}((\theta_0, t_{-i}, s_{-i}) | S_{-i}(h')) = 0$. In words, if t_i gives zero probability to reaching h' , i 's beliefs at h' are fixed to be consistent with $p_i(\cdot | h', t_i)$ and with co-players following the equilibrium from h' onward.

We now prove two results about the Bayesian CPS just constructed. The first shows that it is consistent with $P^{b,t}$ for each type profile t .

Lemma 2. Fix a Bayesian CPS $\hat{v}^{t_i}$ over $\Theta_0 \times T_{-i} \times S_{-i}$ constructed from a profile of Bayesian strategies and belief maps $(b_i, p_i)_{i \in I}$ as in Equations (1)–(3). For every $i, t_i, h <_0 h'$, and every (θ_0, t_{-i}) ,

$$\hat{v}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h') | S_{-i}(h)) = \begin{cases} P^{b, (t_i, t_{-i})}(h' | h) \hat{v}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h) | S_{-i}(h)) & \text{if } \iota(h) \neq i, \\ \hat{v}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h) | S_{-i}(h)) & \text{otherwise.} \end{cases}$$

Proof. Fix (θ_0, t_{-i}) and $h <_0 h'$. If $\iota(h) = i$ then $S_{-i}(h') = S_{-i}(h)$: the two conditioning events

coincide and the claim is immediate. Let then $\iota(h) \neq i$ and note that

$$S_{-i}(h') = \{s_{-i} \in S_{-i}(h) : s_j(h) = a_{h,h'}\},$$

where $a_{h,h'}$ is the action that leads from any node in h to h' (which is well-defined since the game is multistage with observed players' moves).

We show by induction on $<_0$ that, conditional on (θ_0, t_{-i}) , the measure $\hat{\nu}^{t_i}(\cdot | S_{-i}(h))$ is a product across players and each player's histories. At h_0 this holds by Equation (1). At all remaining histories it is inherited, because of the multistage structure of the game and since $S_{-i}(h)$ is a product set, hence the conditioning in Equation (2) operates coordinate by coordinate and, when Equation (3) applies, the measure is a product by construction. The probability of $\{s_j(h) = a_{h,h'}\}$ is therefore computed under the marginal of j alone, namely $b_{t_j}(a_{h,h'} | h)$. $\square$

The second lemma shows that, for any Bayesian CPS $\hat{\nu}^{t_i}$, there is a Bayesian CPS $\hat{\mu}^{t_i}$ over $\Theta_0 \times T_{-i} \times Z$ that induces the same beliefs on the continuation path at every history.

Lemma 3. *Fix a type t_i , and a Bayesian CPS $\hat{\nu}^{t_i}$ over $\Theta_0 \times T_{-i} \times S_{-i}$ with conditioning events $(S_{-i}(h))_{h \in H}$. There is a CPS $\hat{\mu}^{t_i}$ over $\Theta_0 \times T_{-i} \times Z$ conditioning on $(Z(h))_{h \in H}$ that is equivalent to $\hat{\nu}^{t_i}$ in the sense that, for any θ_0, t_{-i} , and h ,*

$$\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z | Z(h)) = \hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i} | S_{-i}(h)).$$

Proof. Fix a type t_i . For any strategy s_i and history h , let $\rho_h(s_i)$ be the strategy $s'_i \in S_i(h)$ such that $s'_i(\cdot | h'') = s_i(\cdot | h'')$ for any $h'' \not\prec h$. Let $\sigma^{t_i} = (\sigma^{t_i}(\cdot | h))_{h \in H}$ be the CPS on S_i , with conditioning events $(S_i(h))_{h \in H}$, obtained from b_{t_i} by the same recursion as Equation (2) and Equation (3): for any s_i , $\sigma^{t_i}(s_i | h_0) = \prod_{h: \iota(h)=i} b_{t_i}(s_i(h) | h)$ and, for $h <_0 h'$, $\sigma^{t_i}(\cdot | h')$ is obtained from $\sigma^{t_i}(\cdot | h)$ by conditioning on $S_i(h')$ when $\sigma^{t_i}(S_i(h') | h) > 0$, and is the image of $\sigma^{t_i}(\cdot | h)$ under $\rho_{h'}$ otherwise.

Note that the product $\hat{\nu}^{t_i}(\cdot | h) \otimes \sigma^{t_i}(\cdot | h)$ is a Bayesian CPS on $\Theta_0 \times T_{-i} \times S$ with conditioning events $(S(h))_{h \in H}$. Conditions (1) and (3) of Definition 12 are immediate, and condition (2) holds because $S(h') = S_i(h') \times S_{-i}(h')$, so the two chain rules multiply.

Given that $S(h) = \{s \in S : z_s \in Z(h)\}$, define $\hat{\mu}^{t_i}$ as the image of this product under the map $(\theta_0, t_{-i}, s) \mapsto (\theta_0, t_{-i}, z_s)$: for every $h \in H$ and (θ_0, t_{-i}, z) ,

$$\hat{\mu}^{t_i}((\theta_0, t_{-i}, z) | Z(h)) := \sum_{s \in S : z_s = z} \sigma^{t_i}(s_i | h) \hat{\nu}^{t_i}((\theta_0, t_{-i}, s_{-i}) | h).$$

Each condition of Definition 12 is inherited: for (1), $s \in S(h)$ implies $z_s \in Z(h)$; for (2), if $h \leq h'$ and $E \subseteq \Theta_0 \times T_{-i} \times Z(h')$ has preimage $\tilde{E}$, then $\tilde{E} \subseteq \Theta_0 \times T_{-i} \times S(h')$ and the chain rule for the product measure gives the chain rule for $\hat{\mu}^{t_i}$; for (3), the initial element is $\pi_i(\theta_0, t_{-i}|t_i)$ by Equation (1). $\square$

By construction, for every i, t_i, h, θ_0 , and t_{-i} , the new Bayesian CPS satisfies:

$$\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h)) = \hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}|S_{-i}(h)) = p_i(\theta_0, t_{-i}|h, t_i).$$

Indeed, at histories where Equation (2) applies, the marginal of $\hat{\nu}^{t_i}$ on $\Theta_0 \times T_{-i}$ of the Bayes update is $p_i(\cdot|h, t_i)$ by condition 2 of IPE, and at the remaining histories it is $p_i(\cdot|h, t_i)$ by Equation (3).

Finally, fix $h <_0 h'$ and (θ_0, t_{-i}) . By construction, $\hat{\mu}^{t_i}(\cdot|Z(h))$ is the image of $\hat{\nu}^{t_i}(\cdot|S_{-i}(h)) \otimes \sigma^{t_i}(\cdot|h)$ under $(\theta_0, t_{-i}, s) \mapsto (\theta_0, t_{-i}, z_s)$, and the preimage of $\{(\theta_0, t_{-i})\} \times Z(h')$ is $\{(\theta_0, t_{-i})\} \times S_i(h') \times S_{-i}(h')$. Hence, if $i(h) \neq i$, then $S_i(h') = S_i(h)$: the factor $\sigma^{t_i}(\cdot|h)$ is unaffected, so

$$\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h')|Z(h)) = \hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h')|S_{-i}(h)) = P^{b, (t_i, t_{-i})}(h'|h) \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h)|Z(h)),$$

where the second equality is Lemma 2 and the last one uses $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h)|Z(h)) = \hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h)|S_{-i}(h))$, given by Lemma 3. If instead $i(h) = i$, then $S_{-i}(h') = S_{-i}(h)$, so Lemma 3 gives $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h')) = \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h))$.

The last part is showing sequential rationality. By condition 3 of IPE, every s_i with $\prod_{h: i(h)=i} b_{t_i}(s_i(h)|h) > 0$ is sequentially rational given ν^{t_i} . By Lemma 2, conditional on each (θ_0, t_{-i}) , $\nu^{t_i}(\cdot|h)$ assigns probability $P^{b, (t_i, t_{-i})}(h'|h)$ to each co-player's move, so the expected continuation payoff it assigns at h to a strategy s'_i is $\sum_{(\theta_0, t_{-i})} p_i(\theta_0, t_{-i}|h, t_i) V_i^{(s'_i, b_{-i})}(\theta_0, t_{-i}, h)$, where we recall that $V_i^{(s'_i, b_{-i})}$ stands for i 's continuation value if the profile of strategies played is (s'_i, b_{-i}) . Since the weights equal $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h))$, this is the criterion of condition (1) of Definition 13, which therefore holds.

From a posteriori CPPBE to IPE: Fix an a posteriori CPPBE $(b_i, (\hat{\mu}^{t_i})_{t_i \in T_i})_{i \in I}$. For each i and t_i we define $p_i(\cdot|\cdot, t_i)$ and a Bayesian CPS $\hat{\nu}^{t_i}$ over $\Theta_0 \times T_{-i} \times S_{-i}$ recursively as follows.

Let $p_i(\theta_0, t_{-i}|h_0, t_i) := \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z)$ and let $\hat{\nu}^{t_i}(\cdot|S_{-i})$ be given by Equation (1). By condition (3) of the Bayesian CPS, $p_i(\cdot|h_0, t_i) = \pi_i(\cdot|t_i)$, so condition 1 of IPE holds.

Fix $h <_0 h'$ and suppose $\hat{\nu}^{t_i}(\cdot|S_{-i}(h))$ is defined.

- (a) If $\hat{\nu}^{t_i}(\Theta_0 \times T_{-i} \times S_{-i}(h')|S_{-i}(h)) > 0$, define $\hat{\nu}^{t_i}(\cdot|S_{-i}(h'))$ by Equation (2) and let $p_i(\cdot|h', t_i)$ be its marginal on $\Theta_0 \times T_{-i}$.

(b) Otherwise, let $p_i(\theta_0, t_{-i}|h', t_i) := \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h'))$ and define $\hat{\nu}^{t_i}(\cdot|S_{-i}(h'))$ by Equation (3).

In case (a) $p_i(\cdot|h', t_i)$ is derived by Bayes' rule from $p_i(\cdot|h, t_i)$ given b_{-i} ; in case (b) Bayes' rule does not apply. Hence condition 2 of IPE holds, and $\hat{\nu}^{t_i}$ is by construction the CPS induced by $p_i(\cdot|t_i)$ and b_{-i} .

Lemma 2 implies that the Bayesian CPS constructed in this way is consistent with the Bayesian strategies. Next, we show that the belief function p_i we have constructed induces the same belief at every history as the initial a posteriori CPPBE.

Lemma 4. *For every i, t_i, h, θ_0 , and t_{-i} ,*

$$p_i(\theta_0, t_{-i}|h, t_i) = \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h)).$$

Proof. We proceed by induction. At h_0 the two sides agree by definition. Let $h <_0 h'$ and assume the claim at h . In case (b) the claim at h' holds by definition, so let case (a) apply. If $\iota(h) = i$ then $S_{-i}(h') = S_{-i}(h)$, so $p_i(\cdot|h', t_i) = p_i(\cdot|h, t_i)$, and condition (2) of Definition 13 gives the same invariance for $\hat{\mu}^{t_i}$; the claim then follows from the induction hypothesis. Then, let $\iota(h) \neq i$. By condition Definition 12 and Definition 13, for any (θ_0, t_{-i}) ,

$$\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h')|Z(h)) = P^{b, (t_i, t_{-i})}(h'|h) \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h)).$$

On the other hand, by Lemma 2,

$$\begin{aligned} \hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h')|S_{-i}(h)) &= P^{b, (t_i, t_{-i})}(h'|h) \hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}|S_{-i}(h)) \\ &= P^{b, (t_i, t_{-i})}(h'|h) \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h)), \end{aligned}$$

where the last equality holds because the marginal of $\hat{\nu}^{t_i}(\cdot|S_{-i}(h))$ on $\Theta_0 \times T_{-i}$ is $p_i(\cdot|h, t_i)$ by construction, which equals the marginal of $\hat{\mu}^{t_i}$ by the induction hypothesis.

Since case (a) holds, summing over (θ_0, t_{-i}) gives $\sum_{(\theta'_0, t'_{-i})} P^{b, (t_i, t'_{-i})}(h'|h) \hat{\mu}^{t_i}(\{(\theta'_0, t'_{-i})\} \times Z|Z(h)) = \hat{\nu}^{t_i}(\Theta_0 \times T_{-i} \times S_{-i}(h')|S_{-i}(h)) > 0$, and in turn, by the previous steps, $\hat{\mu}^{t_i}(\Theta_0 \times T_{-i} \times Z(h')|Z(h)) > 0$. The chain rule of $\hat{\mu}^{t_i}$ then gives

$$\begin{aligned} \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h')) &= \frac{\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z(h')|Z(h))}{\hat{\mu}^{t_i}(\Theta_0 \times T_{-i} \times Z(h')|Z(h))} \\ &= \frac{P^{b, (t_i, t_{-i})}(h'|h) \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h))}{\sum_{(\theta'_0, t'_{-i})} P^{b, (t_i, t'_{-i})}(h'|h) \hat{\mu}^{t_i}(\{(\theta'_0, t'_{-i})\} \times Z|Z(h))}. \end{aligned}$$

By Equation (2), $p_i(\theta_0, t_{-i}|h', t_i)$ is the marginal of $\hat{\nu}^{t_i}(\cdot|S_{-i}(h'))$ on $\Theta_0 \times T_{-i}$, that is, $\hat{\nu}^{t_i}(\{(\theta_0, t_{-i})\} \times S_{-i}(h')|S_{-i}(h))$ divided by $\hat{\nu}^{t_i}(\Theta_0 \times T_{-i} \times S_{-i}(h')|S_{-i}(h))$. By the two equations above, the numerator and denominator coincide with those of the chain-rule expression, hence $p_i(\theta_0, t_{-i}|h', t_i) = \hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h'))$. □

Finally, by definition, sequential rationality given μ^{t_i} depends on μ^{t_i} only through the beliefs $\hat{\mu}^{t_i}(\{(\theta_0, t_{-i})\} \times Z|Z(h))$ over the nodes of i 's information set at h , with the continuation values being computed from the Bayesian strategy profile b . By Lemma 4 these beliefs coincide with $p_i(\cdot|h, t_i)$ at every h , and by Lemma 2 the CPS ν^{t_i} induced by $p_i(\cdot|t_i)$ and b_{-i} assigns, conditional on each (θ_0, t_{-i}) , probability $P^{b, (t_i, t_{-i})}(h'|h)$ to each move of a co-player j , so that the continuation values it induces are those of b_{-i} . Hence, for every s_i with $\prod_{h: I(h)=i} b_{t_i}(s_i(h)|h) > 0$, s_i is sequentially rational given ν^{t_i} , which is condition 3 of IPE. Together with conditions 1 and 2, verified above, $(b_i, p_i)_{i \in I}$ is an IPE, and Lemma 4 is the equality in Proposition 3. □

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*, 59, 61–79.
- (1991b): “Lexicographic probabilities and equilibrium refinements,” *Econometrica*, 59, 81–98.
- BRANDENBURGER, A. AND E. DEKEL (1987): “Rationalizability and correlated equilibria,” *Econometrica*, 55, 1391–1402.
- CATONINI, E. AND A. PENTA (2026): “Backward induction reasoning beyond backward induction,” *American Economic Journal: Microeconomics*, 18, 30–58.
- 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,” Working pa-

per, Columbia University and Northwestern University; latest version at <https://www.laura-doval.com/research>.

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*, Dordrecht: 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*, 63, 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*, 50, 863–894.

KUHN, H. W. (1953): “Extensive games and the problem of information,” in *Contributions to the Theory of Games, Volume II*, ed. by H. W. Kuhn and A. W. Tucker, Princeton, NJ: Princeton University Press, no. 28 in Annals of Mathematics Studies, 193–216.

LIPNOWSKI, E. AND E. SADLER (2019): “Peer-confirming equilibrium,” *Econometrica*, 87, 567–591.

MAILATH, G. J. (2019): *Modeling strategic behavior: A graduate introduction to game theory and mechanism design*, vol. 6 of *World Scientific Lecture Notes in Economics and Policy*, Singapore: World Scientific, updated PDF version (6 August 2026) freely available at <https://web.sas.upenn.edu/gmailath/books/modeling-strategic-behavior/>.

MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic theory*, New York: Oxford University Press.

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

——— (1989b): “The space of conditional systems is a ball,” *International Journal of Game*

Theory, 18, 125–139.

MYERSON, R. B. (1986): “Multistage games with communication,” *Econometrica*, 54, 323–358.

——— (1991): *Game theory: Analysis of conflict*, Cambridge, MA: Harvard University Press.

RÉNYI, A. (1955): “On a new axiomatic theory of probability,” *Acta Mathematica Academiae Scientiarum Hungaricae*, 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*, 231, 106111.

——— (2025b): “Perfect Bayesian equilibrium: Consistency conditions for practical definitions,” Working paper, University of California, San Diego; <https://econweb.ucsd.edu/~jwatson/>.