Perfect Bayesian Equilibrium: Consistency Conditions for Practical Definitions

Joel Watson*

December 2024

Abstract

This paper develops and compares consistency conditions for belief-based solution concepts of noncooperative games, particularly for perfect Bayesian equilibrium (PBE). The paper promotes a condition called *plain consistency*, which requires independent updating across dimensions of the strategy space conditional on "conjunctive" events. For practicality, plain consistency constrains updating only on pairs of consecutive information sets for a player, and it can be imposed narrowly. The concept is defined for infinite games, strengthening the foundation of PBE in settings where sequential equilibrium is not defined. Implications of plain consistency are worked out, including the relation to subgame perfection, a version of "no signaling what you don't know," and the structure of beliefs in a wide class of sender/designer games. A key element of the approach taken herein is to express a player's belief at an information set as a probability distribution over strategy profiles.

1 Introduction

Solution concepts for sequential-move games are based on the notion of *sequential rationality*: that a players' strategy is optimal not just in the ex ante sense (before the game is played) but at all of her information sets. Trembling-hand perfect equilibrium (Selten 1975) and sequential equilibrium (Kreps and Wilson 1982) formulate sequential rationality by examining sequences of fully mixed behavior strategies that converge to the candidate equilibrium strategy profile. For sequential equilibrium, such sequences amount to a consistency condition on the players' updated beliefs at all information sets. Practitioners rarely utilize these solution concepts, however, because (i) constructing sequences of fully mixed strategies with the desired properties can be a cumbersome process for complex games; (ii) in some applications, more permissive concepts that allow for a greater range of beliefs may be desired; and (iii) many applications have infinite action spaces, a setting in which trembling-hand perfect equilibrium and sequential equilibrium are not defined.[1]

*UC San Diego; http://econ.ucsd.edu/∼jwatson/. This paper subsumes Watson (2017). The author thanks the NSF for financial support (SES-1227527) and the following people for their input: Nageeb Ali, Pierpaolo Battigalli, Jesse Bull, Jack Fanning, Simone Galperti, Keri Hu, David Miller, Phil Reny, Karl Schlag, Joel Sobel, seminar participants at Oxford, and several anonymous referees.

[1]Myerson and Reny (2020) discusses technical problems with extending sequential equilibrium to infinite games and develops an equilibrium concept for a class of infinite-action multi-stage games based on nets of strategy perturbations.

Rather, practitioners gravitate toward the perfect Bayesian equilibrium (PBE) concept, which is usually described in the same way as is sequential equilibrium: an *assessment* consisting of the equilibrium behavior-strategy profile and a system of beliefs as probability distributions over nodes in information sets. However, PBE puts structure on the beliefs with consistency conditions that are formulated without reference to strategy trembles. The practitioners' natural objective for consistency conditions is a suitable form of independence and proper conditional-probability updating (Bayes' rule), so as to deliver implications along the lines of what sequential equilibrium implies, such as subgame perfection. Formal definitions of PBE require a precise account of these consistency conditions.

Fudenberg and Tirole (1991) provides the focal PBE definition for a class of games known as *finite multi-stage games with observed actions and independent types*. Their definition assumes belief updating is "reasonable" in that (i) the updated belief in each stage t regarding the type of each player i depends on only the public history through stage $t-2$ and player i's action in stage $t-1$, (ii) the corresponding Bayes'-rule equation for the belief about player i's type is satisfied where well-defined, and (iii) beliefs exhibit independence regarding the types of different players. These consistency conditions are implied by sequential equilibrium, due to the supporting sequence of behavior strategies, although stronger conditions are required to reach equivalence on equilibrium outcomes.

Unfortunately, modern applications of PBE are increasingly outside the class of finite multi-stage games with observed actions and independent types. To develop a definition of PBE for general games, one option is to utilize Battigalli's (1996) *independence property* for conditional probability systems. Such a definition has not been adopted for applications because it has not been formulated in a practicable way, and it has not offered a practical advantage over sequential equilibrium because verifying the independence condition appears tantamount to finding a suitable sequence of fully mixed behavior strategies in a sequential-equilibrium construction.[2] Further, infinite games have not been addressed.

This paper endeavors to support wider application of PBE, as well as extensive-form rationalizability, by developing a simple belief-consistency condition called *plain consistency* that can be flexibly used in applications, and by promoting a strategy-based account of beliefs that simplifies the description of sequential rationality and belief updating. Plain consistency constrains only how an individual player updates her belief on a given pair of consecutive information sets (from one information set to the next one that arises for the same player). It says that, if the player's belief at the first information set exhibits independence between a given strategy component L and its complement, and if the marginal on L would reach the second information set with positive probability, then the belief at the second information set also exhibits independence and the marginal on L accords to the conditional-probability formula. The formal, general definition makes the same statement more broadly by allowing for conditioning on a set of strategies that is a product set relative to L.

This paper provides a definition of plain consistency for finite games, a generalized definition for infinite games, and a stronger version for beliefs derived from conditional-probability

[2]Since Fudenberg and Tirole's (1991) and Battigalli's (1996) framework clings to an assessment-based account of strategies and beliefs, such a PBE definition would amount to the following: Postulate a conditional probability system over strategy profiles, ensure that it has the independence property, calculate an implied strategy profile and a conditional probability system over terminal nodes, derive from it the assessments at information sets, and verify sequential rationality.

systems (CPS). Plain consistency is compared to other belief-consistency notions in the literature. Alternative versions of perfect Bayesian equilibrium are defined by combining consistency notions with sequential rationality and common initial beliefs. The paper proves that a moderate form of plain consistency implies subgame perfection. For multi-stage games with observed actions and independent types, the moderate form implies an element of Fudenberg and Tirole's (1991) reasonableness condition, and the CPS version implies play equivalent to a sequential equilibrium. The paper derives the implications of plain consistency for a wide class of sender/designer games described in the following subsection, allowing for infinite action spaces and infinite numbers of information sets.

1.1 Applied motivation

Although applications of PBE are widespread, a measure of ambiguity persists across the applied literature regarding the meaning of "imposing Bayes' rule where applicable." In some cases, PBE is the stated solution concept but it is not formally defined, leaving one to guess whether the reported analysis relies on implicit consistency assumptions or is an arbitrary selection of some weak version of PBE. In other cases, the modeling exercise alludes to a noncooperative game, but analysis is provided for only part of the game, and therefore a PBE is not fully constructed.

Examples abound from the large literature on *sender/designer games*, including models in categories known as information design, persuasion, mechanism design, cheap talk, and hard evidence. In these models, the following occur in order: There is a move of nature; player 1 may or may not observe some aspect of nature's move; player 1 makes a choice such as to design a signaling technology, send a message, or disclose evidence; and then other players make an observation (receive the message, see the signal, etc.) and take actions.

Consider the simplest example in the category of information design. At the beginning of the game, nature chooses whether the underlying state is 0 or 1, with equal probability. Player 1 does not observe nature's move. Player 1 then selects an experiment e, which is a mapping from the set of states $\{0, 1\}$ to probability distributions over the same space (randomization performed by nature). The outcome of the experiment can be interpreted as the recommended action. Player 2 observes e and the experiment outcome (not the state directly) and then chooses either action 0 or action 1. Suppose that player 2's payoff is 0 if she chooses action 0, regardless of the state. If she chooses action 1 then player 2's payoff is 1 in state 1 and it is -2 in state 0. Player 1 strictly prefers that player 2 take action 1 in both states.

The typical analysis of such a game involves the following steps: For any experiment, the analyst calculates player 2's optimal response to every experiment outcome, as though the experiment is exogenously given along with prior distribution of the state. This calculation results in an expression of the mixed action that the experiment implements as a function of the state. Then, the analyst picks the experiment e^* that maximizes the sender's expected payoff, implicitly claiming that this experiment would be chosen in every PBE of the game. The foregoing description is the approach taken by the majority of papers in the literature, including the most prominent such as Kamenica and Gentzkow (2011), the earliest entries such as Watson (1996), and the most recent such as Bergemann, Heumann, and Morris (2023).[3]

[3]Additional recent examples are noted in Section 5 to illustrate scope.

In our example, e^* specifies outcome 1 for sure in state 1, and in state 0 it puts probability $1/2$ on outcome 1 and probability $1/2$ on 0. In the event that player 2 observes experiment outcome 1, her posterior belief puts probability $2/3$ on state being 1, making her indifferent between the two actions, and she chooses action 1. This maximizes the probability that player 2 chooses action 1, over all experiments, and so e^* is optimal for player 1.

But notice that the typical analysis in the literature is not an analysis of the full game. To conduct such an analysis, one must describe player 2's beliefs and rational behavior following every *surprise choice* of player 1—that is, where player 1 deviates by selecting other than the equilibrium experiment—and for whatever is the experiment outcome. In such an off-path contingency, it is feasible for player 2 to update her belief about the state in a way that is related not only to the realized signal but also to the surprising choice of player 1. Player 2's updated belief about nature is, in fact, arbitrary without the analyst making a consistency assumption for such an off-path contingency.

In fact, in our example, there is a "weak PBE" in which, on the equilibrium path, player 1 selects not e^* but rather the most informative experiment, which outputs 1 for sure in state 1, and 0 for sure in state 0.[4] If player 1 were to deviate by selecting, say, e^* and the experiment outputs 1, then player 2 believes that the state is surely 0 and chooses action 0, to player 1's detriment. Note that in this equilibrium, player 1's deviation to e^* signals to player 2 something about the state that player 1 has not observed.

In addition to the weak PBE just described, there are many others, including one in which e^* is chosen in equilibrium. The broad set of weak PBE support a range of outcomes, equilibrium payoffs, and information transmission. In contrast, in a finite version of the game in which player 1 selects from a finite set of experiments, sequential equilibrium predicts player 1 will pick an available experiment close to e^*.[5] The beliefs and behavior of more informative weak PBE, such as the fully informative described in the previous paragraph, are not part of any sequential equilibrium.

One may or may not want to rule out the outcome in which player 1 chooses the most informative experiment. In either case, making the distinction for the specified infinite game requires articulating a defensible consistency condition that rules out such an equilibrium. This is not a multistage game with observed actions, so Fudenberg and Tirole's (1991) PBE definition described above does not apply, even holding aside that the game is infinite. One might argue that we could invalidate the fully informative outcome by examining weak PBE for a different, equivalent extensive form in which nature selects the experiment outcome before player 2's move and the actual state after player 2's move. It is true that the alternative game form does not have a weak PBE of the fully informative kind, because weak PBE assumes that the strategy profile nails down beliefs regarding future actions. However, this trick does not work in general and is normatively deficient.[6]

[4]A weak PBE (or weak sequential equilibrium, Myerson 1991) is an assessment that satisfies sequential rationality and Bayes' rule on the equilibrium path (discussed in Mas-Colell, Whinston, and Green 1995).

[5]It is the experiment that, under the prior belief about the state, maximizes the probability player 2 will strictly prefer action 1.

[6]For example, consider a four-player game in which nature chooses state θ_1 that players 1 and 3 care about, and nature chooses state θ_2 that players 2 and 4 care about. Players 1 and 3 interact as in our simple information-design example, as do players 2 and 4, except that player 1 observes θ_2 and player 2 observes θ_1. In every extensive-form depiction of this game, either nature's choice of θ_1 occurs before player 4's information sets, or

Thus, many of the entries in the literature do not acknowledge the existence of the full range of equilibria described above or rule out such equilibria with unstated belief-consistency conditions absent from the literature on solution concepts. The consideration becomes more subtle and complicated if, in relation to our example, player 1 received a private signal of nature's move. I suggest that it would be helpful to put the literature on firmer ground by defining an appropriate consistency condition that is amenable to analysis of the full game.[7]

A similar critique applies to some models of mechanism design, most notably in the context of an informed principal. In Myerson's (1983) seminal modeling exercise, nature chooses a vector of types for the strategic players, and then player 1 privately observes her type and selects a mechanism that governs communication between all of the players (a mediation system). The other players are informed of the selected mechanism and privately observe their types. Although Myerson (1983) claims to study sequential equilibrium, updated beliefs about types cannot be derived from a sequence of behavior strategies due to the set of mediation systems being uncountably infinite. Rather, Myerson assumes these beliefs take a particular form that is essentially structural consistency (Kreps and Wilson 1982).

I show that, for sender/designer games, plain consistency implies structural consistency of beliefs about the moves of nature and the sender/designer player. An implication is to rule out the unintuitive beliefs described above for our information-design example, so that every *plain* PBE has player 1 choosing e^* on the equilibrium path, just as sequential equilibrium predicts in the finite version of the game. The result also provides justification for Myerson's assumption for informed-principal games.[8] More generally, it provides a foundation for PBE definitions that rule out equilibria where, in off-path contingencies, the sender/designer signals something about nature's move that the sender/designer did not observe.

1.2 Technical overview and layout

For intuition and to describe plain consistency in the abstract, consider a finite two-dimensional state space $S = S_L \times S_{-L}$, where L refers to one dimension and $-L$ the other. Let p be a Bayesian decision maker's prior belief about the state $s = (s_L, s_{-L}) \in S$, and suppose p has full support. Imagine that the decision maker then learns that $s \in E$ for some event $E \subset S$. Let p' denote the updated belief, which from the conditional-probability formula is given by $p'(X) = p(X \mid E) = p(X \cap E)/p(E)$ for every $X \subset S$.

Consider the special case in which two conditions are satisfied: First, the prior belief has s_L and s_{-L} independently distributed. Second, the conditioning event is a product set $E = E_L \times E_{-L}$; in words, it is the *conjunctive event* "$s_L \in E_L$ AND $s_{-L} \in E_{-L}$." Then clearly the posterior belief can be written as the product of the updated marginal distribution of s_L, conditioned on E_L, and the updated marginal distribution of s_{-L}, conditioned on E_{-L}. That

nature's choice of θ_2 occurs before player 3's information sets. Further, apart from whether this trick works, a solution concept that predicts different outcomes for strategically equivalent games is undesirable because it obscures logic underlying beliefs and behavior.

[7] Incidentally, Krähmer (2021) is an example in which the full game is analyzed and where one of the main results relies on an equilibrium construction similar to the unintuitive one described here.

[8] Myerson (1982) and the subsequent literature have another, unrelated problem, discussed in Section 5.4: Rationality is not evaluated for some information sets, and therefore equilibrium constructions do not satisfy the sequential-rationality conditions required for sequential equilibrium or any perfect equilibrium concept.

is, for any set $X = X_L \times X_{-L} \subset S$, we have $p'(X) = p_L(X_L \,|\, E_L) \cdot p_{-L}(X_{-L} \,|\, E_{-L})$, where p_L is the marginal of p on dimension L and p_{-L} is the marginal on dimension $-L$. Updating on the L dimension can be performed separately from updating on the $-L$ dimension.

Plain consistency extends this idea to settings in which the prior distribution does not have full support, in particular where $p(E) = 0$. If the conditional-probability formula cannot be applied for X overall but is applicable for updating on the L dimension, then plain consistency requires that it be employed to calculate the posterior marginal distribution of s_L, independent of how s_{-L} is updated. Thus, if $p_L(E_L) > 0$ and yet $p_{-L}(E_{-L}) = 0$, plain consistency still requires $p'_L(X_L) = p_L(X_L \,|\, E_L)$ for every $X_L \subset S_L$, and it requires p' to exhibit independence between s_L and s_{-L}.

In the context of a game, s_L and s_{-L} represent complementary components of the strategy profile, and E represents information a player receives about the strategy profile by virtue of arriving at an information set. Probably distribution p gives the player's belief at her previous information set, and p' is the updated belief at the current information set.

Plain consistency is formally defined in a more general way that allows for conditioning on any product set $Y = Y_L \times Y_{-L}$. Specifically, suppose that $p(Y) > 0$ and the conditional probability measure $p(\cdot \,|\, Y)$ has s_L and s_{-L} independently distributed. Consider any $s_{-L} \in S_{-L}$ such that $Y_L \times \{s_{-L}\} \subset E$ and $p'(Y_L \times \{s_{-L}\}) > 0$. Then the general form of plain consistency requires $p'(X_L \times \{s_{-L}\})/p'(Y_L \times \{s_{-L}\})$ to equal $p(X_L \times Y_{-L})/p(Y_L \times Y_{-L})$ for every $X_L \subset E_L \cap Y_L$. Plain consistency is implied by Battigalli's independence condition for beliefs given by a conditional-probability system, which is further implied by Kreps and Wilson's (1982) full consistency condition.

A key element of the approach taken herein is to describe a player's belief and plan at an information set as a probability distribution over the strategy profile, which I call an *appraisal*. It captures the player's conjecture about both how the information set was reached and what will happen from this point in the game.[9] I specify that every player has a (possibly artificial) information set representing the beginning of the game. Thus, a player starts the game with an appraisal and then updates it from one information set to the next as play proceeds.

Describing the players' contingent beliefs and plans as appraisals rather than as an assessment entails no loss, because every assessment uniquely translates into an appraisal system. Appraisals offer a few advantages. First, appraisals provide for a simpler expression of belief-consistency conditions and continuation values.[10] Second, appraisals allow for a direct description of independence regarding strategy components.[11] Third, appraisals neatly

[9]Expressing beliefs as probability distributions over the strategy space is common for defining rationalizability, as in the original definition of Pearce 1984 and Battigalli's (1997) restatement, and it utilized for some equilibrium theory, such as Battigalli's (1996) exploration of independence underlying equilibrium notions and Govindan and Wilson's (2009) study of forward induction. I refine the approach by incorporating into the appraisal the player's own strategy and by incorporating additional structure.

[10]For an assessment (μ, π), where π is the strategy profile and μ gives the probabilities over nodes at information sets, calculating the continuation value from an information set entails averaging over nodes according to μ and assuming the other players' future behavior is according to π. With an appraisal system, the appraisal at an information set is plugged into the payoff function.

[11]To accomplish this with assessments, one would have to account for how, at a given information set, each node x describes the information sets reached, and actions taken, on the path to x. But such a structure is equivalent to keeping track of the set of strategy profiles that reach a given information set, as appraisals facilitate. Further, even if we imagine adding this structure and then using the equilibrium strategy profile for "future"

separate beliefs from plans and allow for consistency conditions to be built from the ground up, whereas an assessment, by using the strategy profile to determine future behavior, already embeds an independence assumption: that a player's belief about future behavior is treated independently of any observation about past moves. I submit that it is better for clarity and broad application to have a theoretical foundation in which the belief system is expressed in the abstract, and then one can impose consistency conditions.[12]

The next section lays out the basic notation and definitions. Section 3 presents benchmark consistency conditions, including full consistency, the independence condition for conditional-probability systems, and the notion of minimal consistency sufficient for the one-deviation property of rationality. Section 4 develops the notion of plain consistency and explores some of its implications. Section 5 studies implications of plain consistency for sender/designer games. Section 6 discusses rationalizability, offers further technical notes, and adds to the discussion of related literature. The Appendices contain additional definitions related to behavior strategies and the one-deviation property, and all of the proofs.

2 Fundamental Concepts

This section and the next contain the fundamental definitions for finite n-player games of perfect recall, represented in the extensive form. Straightforward modifications allow for the extension to infinite games in Section 4.3.

Denote by $N \equiv \{0, 1, \ldots, n\}$ the set of players, which includes nature (chance) as the nonstrategic player 0, as well as strategic players $1, 2, \ldots, n$. Let $N_+ \equiv N \setminus \{0\}$ denote the set of strategic players. In some applications, to sort out items that players obtain information about, it may be convenient to specify that there are several nature players or that nature does not have perfect recall (an example appears in Section 4.2) but it is not necessary to do so. In the expressions to follow, "$\subset$" denotes (weak) "subset," not "proper subset."

2.1 Situations, strategies, payoffs, and information sets

Let H be the set of distinct *situations* in the game where individual players must make choices. Although it is customary to call these *information sets* in an extensive-form representation, I use the word "situation" for these decision points and "information set" for the information that players have at these points in the game, which will be put in terms of strategy profiles. As is standard, situations are distinctly labeled so that the players' individual sets of situations are disjoint. Let $H_+ \equiv \cup_{i \in N_+} H_i$ be the set of situations for the strategic players.

Denote by S the space of pure strategy profiles, including nature's strategy. A strategy profile $s \in S$ maps H to the space of actions and, for each situation, specifies an action that is

information sets when calculating expected payoffs, another complication arises: At a given information set, the other information sets in the game generally cannot be classified as coming either before or after the current one (see Kreps and Ramey 1987). These complications presumably led Fudenberg and Tirole (1991) and Battigalli (1996) to describe equilibrium in general games as a combination of assessments and conditional probability systems on terminal nodes.

[12]Also, assessments are not suitable for rationalizability theory and other solution concepts in which the players may have different beliefs about the future. Plain consistency, by the way, implies that the players' beliefs about the future are unchanged as the game progresses.

feasible at this situation. A strategy profile is equivalently expressed as a tuple $s = (s_h)_{h \in H}$ where, for every $h \in H$, s_h is a feasible action at situation h. For every $i \in N_+$, let $u_i : S \to \mathbb{R}$ be player i's payoff function.

For any subset of situations $L \subset H$, let s_L denote the restriction of s to the subdomain L. That is, s_L gives the actions that s specifies for the situations in L. For any $L \subset H$, define $-L \equiv H \setminus L$. We can then write $s = (s_L, s_{-L})$. For $X \subset S$, define $X_L \equiv \{s_L \mid s \in X\}$. In the case of $L = \{h\}$ for a single $h \in H$, we simplify notation by dropping the brackets; for instance, we write X_h instead of $X_{\{h\}}$. Likewise, let subscript "i" stand in for "H_i," so that s_i means the same thing as s_{H_i}, and similarly for $-i$ and H_{-i}. Note that S_h is the set of actions available at situation h, and s_h is the action chosen at h by the player on the move.

It will be useful to express information sets in the strategic form—that is, in terms of subsets of strategy profiles.

Definition: *For every $h \in H$, denote by $S(h)$ the set of strategy profiles that reach situation h, in that an action at h is chosen on the path of this strategy profile (in the extensive form representation, the path includes a node in h). Call $S(h)$ the* ***information set (in the strategic form) at h****.*

The information set at h represents what the player on the move has observed about past play, which is that the actual strategy profile must be in $S(h)$. Note that, for any $L \subset H$, $S(h)_L$ is the set of action profiles for the situations in L that are consistent with reaching h.[13]

Definition: *For $L \subset H$, say that a set* ***$X \subset S$ is an L-product set*** *if $X = X_L \times X_{-L}$.*

Perfect recall implies that the information sets for an individual player have a product structure. To be precise, for every player $i \in N_+$, $h \in H_i$, and $L \subset H_i$, it is the case that $S(h)$ is an L-product set.[14]

Let us assume that at every situation at least two actions are available, so that the player on the move has a nontrivial choice to make. This assumption, along with perfect recall, implies that the information sets for an individual player represent the precedence relation of situations. If $h \neq h'$, then either h *is a successor of* h', in which case $S(h) \subset S(h')$ and $S(h) \neq S(h')$; or h *is a predecessor of* h', in which case $S(h') \subset S(h)$ and $S(h) \neq S(h')$; or neither, in which case $S(h) \cap S(h') = \emptyset$. If h' is a successor of h then every strategy profile that reaches h' also reaches h.

Call $h' \in H_i$ *an immediate successor of* $h \in H_i$ for player i if h' is a successor of h and there is no other situation for player i between the two; that is, there is no $h'' \in H_i$ such that h'' is a successor of h and h' is a successor of h''. In this case, let us also call h and h' (in this order) *consecutive situations*.

[13] Expressing information sets as subsets of strategy profiles is standard for defining extensive-form rationalizability (Pearce 1984) and conditional dominance (Shimoji and Watson 1998). It has also been utilized for equilibrium theories, such as formulating solution concepts on the basis of "normal-form information sets" by Mailath, Samuelson, and Swinkels (1993) and analysis of forward induction by Govindan and Wilson (2009).

[14] For any player $i \in N_+$ and $h, h' \in H_i$, it must be that either $S(h) = S_{h'} \times S(h)_{-h'}$ or $S(h)_{h'}$ is a singleton and so $S(h) = S(h)_{h'} \times S(h)_{-h'}$. This is because, if the first equation does not hold, there exist strategy profiles s and s' such that $s_{-h'} = s'_{-h'}$, $s \in S(h)$, and $s' \notin S(h)$, which means that at h player i has already made the choice at h', and so by perfect recall it must be that $S(h)_{h'}$ is a singleton.

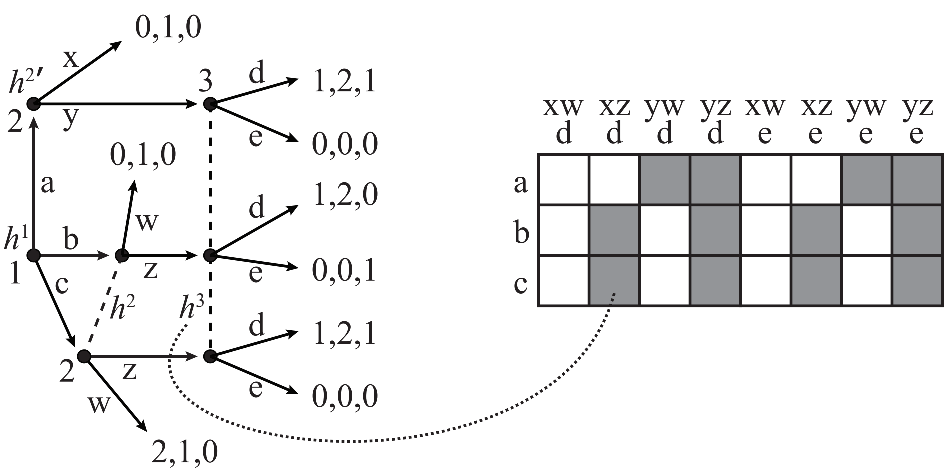

Figure 1: Example to illustrate theoretical components.

To review some of the definitions just described, consider the three-player game shown in Figure 1, where the situations are labelled h^1, h^2, $h^{2\prime}$, and h^3. The space of strategy profiles is $S = \{\mathrm{a,b,c}\}\times\{\mathrm{x,y}\}\times\{\mathrm{w,z}\}\times\{\mathrm{d,e}\}$, which is depicted by the table on the right side of the picture. The rows of the table are the actions feasible at player 1's situation h^1, whereas the columns are the profiles of actions for the other players' situations. For player 3's situation h^3, we have $S_{h^3} = \{\mathrm{d,e}\}$ and $S_{-h^3} = \{\mathrm{a,b,c}\}\times\{\mathrm{x,y}\}\times\{\mathrm{w,z}\}$. The subset of strategy profiles that are consistent with reaching h^3 is

$$S(h^3) = \{(\mathrm{a,y,w}), (\mathrm{a,y,z}), (\mathrm{b,x,z}), (\mathrm{b,y,z}), (\mathrm{c,x,z}), (\mathrm{c,y,z})\}\times\{\mathrm{d,e}\}.$$

This set corresponds to the shaded region of the table in the figure. Clearly $S(h^3)$ is a product set relative to h^3 but it is not a product set relative to the situation of player 1.

2.2 Strategy distributions and independence

Let ΔS denote the space of probability distributions over S, which we call the mixed strategy profiles. For any subset $T \subset S$, let us take "ΔT" to mean the subset of ΔS with support in T. Here are some standard definitions for reference:

Definition: *Given a distribution $p \in \Delta S$ and a set of situations $L \subset H$, the **marginal distribution on dimension L**, denoted by $p_L \in \Delta S_L$, is defined by $p_L(X_L) \equiv p(X_L \times S_{-L})$ for every $X_L \subset S_L$. We also call this the L-marginal distribution.*

Definition: *For a distribution $p \in \Delta S$ and an event $Y \subset S$ for which $p(Y) > 0$, the **distribution conditional on Y**, denoted by $p(\,\cdot\,|\,Y)$, is given by $p(X\,|\,Y) \equiv p(X \cap Y)/p(Y)$ for every $X \subset S$. We also denote the conditional distribution by $(p\,|\,Y)$.*

Occasionally we need to refer to the L-marginal of a conditional distribution—that is, first conditioning on a set Y to get $p(\,\cdot\,|\,Y)$, and then calculating the marginal on dimension L. I shall denote the resulting marginal distribution by $p(\,\cdot\,|\,Y)_L$ or $(p\,|\,Y)_L$. The latter notation will be used where we need the marginal evaluated on a particular set; for instance, for the distribution p conditional on Y, $(p\,|\,Y)_L(X_L)$ is the marginal probability of $X_L \subset S_L$. Likewise, we will sometimes need to calculate the conditional distribution of a given L-marginal. That is, for any $X_L, Y_L \subset S_L$ satisfying $p_L(Y_L) > 0$, we have $p_L(X_L \mid Y_L) \equiv p_L(X_L)/p_L(Y_L)$.

The next definition identifies whether a mixture of strategy profiles, conditional on some event, treats a specific set of situations $L \subset H$ independently of the rest, meaning that it can be expressed as the product of the marginal distributions on L and $-L$.

Definition: *Given a distribution $p \in \Delta S$, a set of situations $L \subset H$, and an L-product set $Y \subset S$, say that a distribution **$p \in \Delta S$ exhibits L-independence on Y** if $p(Y) > 0$ and, for every L-product set $X \subset Y$, $p(X \mid Y) = (p \mid Y)_L(X_L) \cdot (p \mid Y)_{-L}(X_{-L})$. Call $\sigma \in \Delta S$ a **behavior strategy profile** if it exhibits h-independence on S for every $h \in H$.*

This definition of a behavior strategy is equivalent to specifying a distribution over actions at every situation, independent across situations, which is the traditional way of describing a behavior strategy in the literature. In fact, if σ is a behavior strategy profile as defined above, then $\sigma(s) = \prod_{h \in H} \sigma_h(s_h)$ for every $s \in S$, and σ exhibits L-independence on S for every $L \subset H$. Technical details are provided in the Appendix. I use this description of behavior strategies because it puts strategy profiles and beliefs in the same space.

2.3 Beliefs as appraisal systems

It is convenient to express the players' beliefs at their various situations as probability distributions over the strategy space, which I call *appraisals*. To model belief updating, we think of the players as having artificial situations that refer to "before the game begins." To be precise, for each player $i \in N_+$, we define the *initial situation* $\underline{h}_i$ to have $S(\underline{h}_i) = S$, and we label these differently across players so as not to duplicate the situations in the game, H. For each player i, let $\underline{H}_i \equiv H_i \cup \{\underline{h}_i\}$ denote the extended set of situations. Let $\underline{H}_+ \equiv \cup_{i \in N_+} \underline{H}_i$. In the obvious way, we include the initial situations for the ordering that relates consecutive situations for each player. Specifically, for every situation $h \in H_i$, if h has no predecessors in H_i then $(\underline{h}_i, h)$ is defined to be a pair of consecutive situations.

Definition: *For any strategic player i and $h \in \underline{H}_i$, call a distribution $p^h \in \Delta S$ **an appraisal at h** if $p^h \in \Delta S(h)$ and if p^h exhibits h'-independence on S for every $h' \in H_i$. An **appraisal system** is a collection of appraisals, one for each situation of the strategic players, written $P = \left(p^h\right)_{h \in \underline{H}_+}$.*

The condition $p^h \in \Delta S(h)$ means that the appraisal at h puts probability one on reaching h. The independence condition means that, at any situation $h \in \underline{H}_+$, player i views his own strategy as independent of the other players' strategy profile, and player i's strategy is represented as a behavior strategy. Thus, one can always write the appraisal p^h as (σ_i, θ_{-i}), where $\theta_{-i} = p^h_{-i}$ is a distribution over the strategy profile of the other players (player i's belief), and $\sigma_i = p^h_i$ is player i's behavior strategy. The following condition means that a player's plan of behavior for as yet unreached situations does not change as the game progresses.

Definition: *Say that an appraisal system **P has the consistent-plan property** if for every player $i \in N_+$ and for every pair $h, h' \in H_i$ such that h' is not a predecessor of h, it is the case that $p^h_{h'} = p^{h'}_{h'}$.*

2.4 Sequential best response

We can test whether an appraisal at h specifies rational behavior for the player on the move, meaning that the actions assigned positive probability at situation h maximize the player's expected payoff.

Definition: *For a given situation $h \in H_+$ and two appraisals p^h and $\hat{p}^h$, say that **$\hat{\mathbf{p}}^{\boldsymbol{h}}$ is an h-deviation from $\mathbf{p}^{\boldsymbol{h}}$** if $p^h_{-h} = \hat{p}^h_{-h}$ (they are identical on all other situations).*

Definition: *For a strategic player i and a situation $h \in H_i$, say that an appraisal **$\mathbf{p}^{\boldsymbol{h}}$ is rational at h** if $u_i(p^h) \geq u_i(\hat{p}^h)$ for every h-deviation $\hat{p}^h$. Say that an **appraisal system $\boldsymbol{P} = \left(\boldsymbol{p^h}\right)_{\boldsymbol{h} \in \underline{\boldsymbol{H}}_+}$ is sequentially rational** if p^h is rational at h, for every $h \in H_+$.*

Here sequential rationality is defined in terms of single deviations, meaning that we evaluate player i's rationality at a given situation $h \in H_i$ by looking only at alternative choices at h rather than alternatives that would also adjust player i's behavior at other situations that may be reached in the continuation of the game.[15] The familiar *one-deviation property* (equivalence between single-deviation optimality and continuation-strategy-deviation optimality) holds under the condition of minimal consistency for belief updating, which is defined in the next section (see the Appendix for details).

For an example of an appraisal system and sequential best response, consider again the game shown in Figure 1. Suppose that the players' initial appraisals coincide on the distribution that puts probability one on the strategy profile $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$. That is, for each player i, $p^{\underline{h}_i}(\{(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})\}) = 1$. Suppose that the appraisal is the same at player 2's situation h^2, which itself is reached by $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$. Note that player 2's other situation $h^{2\prime}$ is not reached by $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$, so player 2 cannot have the same appraisal there. Let the appraisal at $h^{2\prime}$ put probability one on $(\mathrm{a}, \mathrm{x}, \mathrm{w}, \mathrm{e})$. Finally, suppose that at player 3's situation h^3, player 3's appraisal puts probability one on $(\mathrm{b}, \mathrm{x}, \mathrm{z}, \mathrm{e})$. It is easy to confirm that the appraisal system just described is sequentially rational. At each situation, the player on the move cannot gain by switching from the action prescribed by her appraisal to a different action.

2.5 Generic equilibrium definition

With beliefs represented as an appraisal system, it is straightforward to write the definition of sequential equilibrium and versions of perfect Bayesian equilibrium. Each such definition combines sequential rationality with some sort of consistency condition on the players' appraisal system, and also requires the equilibrium condition that, at the beginning of the game, the players' appraisals coincide on the equilibrium strategy profile.

Definition: *Consider a behavior strategy profile $\sigma \in \Delta S$. Say that an **appraisal system $\boldsymbol{P} = \left(\boldsymbol{p^h}\right)_{\boldsymbol{h} \in \underline{\boldsymbol{H}}_+}$ is rooted in $\boldsymbol{\sigma}$** if, for every player i, $p^{\underline{h}_i} = \sigma$.*

Here is the generic equilibrium definition for any given (yet to be specified) consistency requirement on the appraisal system:

[15] Note that since all strategy profiles in the support of p^h and $\hat{p}^h$ reach h, the expected payoffs shown in the rationality definition are conditional on reaching situation h.

Definition: *Taking nature's strategy $(\sigma_h)_{h \in H_0}$ as given, and take as given a consistency concept for beliefs that we refer to as ⊛-consistency. We say that **an appraisal system P is a ⊛-perfect Bayesian (or sequential) equilibrium** if there is a behavior strategy profile σ^* such that P is rooted in σ^*, P is ⊛-consistent, and P is sequentially rational. In this case, we say that σ^* is the equilibrium strategy profile.*

3 Benchmark Consistency Conditions on Beliefs

This section briefly reviews three consistency notions that will serve as benchmarks. The first is full consistency, standard in the literature as the basis for the definition of sequential equilibrium, here conveniently reformulated as a condition on appraisal systems. The next refers to beliefs constructed from a general conditional-probability system on the strategy space that satisfies Battigalli's (1996) independence condition, which I call strong consistency. The third is a much weaker requirement that I call minimal consistency. Minimal consistency operates on only the appraisals at consecutive situations and is sufficient for the one-deviation property of sequential rationality.

3.1 Full consistency and cps/independence

To define the strongest form of belief consistency, *full consistency*, assume that nature's strategy is exogenously given as part of the game representation and is in the form of a behavior strategy with full support.

Definition: *Call a behavior strategy $\sigma \in \Delta S$ **fully mixed** if $\sigma(s) > 0$ for every $s \in S$, and σ_0 is the strategy for nature given in the representation of the game.*

A fully mixed behavior strategy reaches every situation $h \in H$ with positive probability, implying a well-defined probability distribution $\sigma(\cdot \mid S(h))$ calculated from the conditional-probability formula. Full consistency requires that the players' beliefs are given by the limit of such conditional-probability updating, in relation to a sequence of fully mixed behavior strategies that converges to the equilibrium strategy profile. The strong form of independence is due to the use of behavior strategies in the construction.

Definition: *Say that **an appraisal system P is fully consistent** if there is a behavior strategy $\sigma \in \Delta S$ such that σ_0 is the given strategy of nature, P is rooted in σ, and the following condition holds: There is a sequence of fully mixed behavior strategies $\{\sigma^k\}$ that converges to σ and has the property that, for every $h \in H_+$, the sequence of conditional distributions $\{\sigma^k(\cdot \mid S(h))\}$ converges to p^h.*

Sequential equilibrium, due to Kreps and Wilson (1984), is the concept we get by requiring full consistency in Definition 11, although Kreps and Wilson put beliefs in terms of assessments rather than appraisals. As for consistency conditions that are not founded on strategy perturbations, the strongest would assume that beliefs are derived from a conditional-probability system (Rényi 1955, Myerson 1986) over the space of strategy profiles, and that it satisfies Battigalli's (1996) proposed independence condition, defined as follows.

Definition: *A **conditional probability system on S** is a mapping $\zeta\colon 2^S \times (2^S \setminus \{\emptyset\}) \to [0,1]$ such that $\zeta(\,\cdot \mid E) \in \Delta E$ for every nonempty $E \subset S$, and for sets $X \subset Y \subset Z \subset S$ with Z nonempty, the chain rule holds: $\zeta(X \mid Z) = \zeta(X \mid Y) \cdot \zeta(Y \mid Z)$.*

Because $\zeta(\,\cdot \mid E)$ is a probability distribution for which the marginals can be defined, we will sometimes denote this distribution overall by $(\zeta \mid E)$. Then the L-marginal is denoted by $(\zeta \mid E)_L$ and we write $(\zeta \mid E)_L(X_L)$ as the marginal evaluated at $X_L \subset S_L$.

Definition: *A conditional probability system ζ, defined on S, is said to **have the independence property** if for every $L \subset H$ and every nonempty L-product set $Y \subset S$, it is the case that $(\zeta \mid Y_L \times Y_{-L})_L = (\zeta \mid Y_L \times S_{-L})_L$.*

Battigalli (1996) shows that the independence property is weaker than full consistency, is equivalent to full consistency for games with "observable deviators," and, in equilibrium, implies subgame perfection. Following Fudenberg and Tirole (1991), Battigalli's (1996) equilibrium analysis is put in terms of assessments and conditional probability systems on terminal nodes, from which assessments and strategies are derived. The extra layers are not needed here, as we are expressing beliefs in terms of appraisals. Requiring the independence property leads to the most stringent notion of consistency for perfect Bayesian equilibrium.

Definition: *Say that **appraisal system P is derived from conditional-probability system ζ** if for every $h \in \underline{H}_+$, $p^h = (\zeta \mid S(h))$.*

Definition: *Say **an appraisal system is strongly consistent** if it is derived from a conditional-probability system that has the independence property.*

3.2 Minimal consistency

The definition of minimal consistency, like strong consistency, makes no reference to strategy perturbations. Unlike strong consistency, it also makes no reference to a conditional-probability system. It simply requires players to use the conditional-probability formula where applicable to update beliefs for all consecutive situations.

Definition: *An appraisal system P is called **minimally consistent** if it has the consistent-plan property and, for every strategic player i and consecutive situations h and h' in H_i satisfying $p^h_{-i}(S(h')_{-i}) > 0$, it is the case that $p^{h'}_{-i} = p^h_{-i}(\,\cdot \mid S(h')_{-i})$.*

Under minimal consistency, each player i maintains her continuation strategy and, regarding her belief about the other players' strategies, updates from one situation to the next using a conditional probability calculation where applicable. In the event that the conditional probability expression is not well defined (the denominator is zero), player i develops a new belief about the others that is consistent with reaching the current situation. Minimal consistency has the same implications for outcomes as what Perea (2000) calls "updating consistency."[16]

[16] Perea's (2000) updating consistency is weaker in the sense of allowing a player, at a situation that she previously put positive probability on reaching, to change her beliefs about behavior at situations that cannot be on the play path in the past or future.

4 Plain Consistency

My proposed consistency condition has been developed to achieve the following desiderata. First, given the difficulty of constructing fully consistent and strongly consistent belief systems, the consistency condition should practicable for use in common applications. Helpful in this regard is that it regulate only updating over consecutive situations, as with minimal consistency. Second, the condition should be equipped to narrow the scope of perfect Bayesian equilibrium by implying subgame perfection, ruling out unreasonable signaling, and such. Third, it should be flexible in the sense that the analyst can impose the condition narrowly or more broadly as suitable for applications. Fourth, the condition should generalize for use in games with infinite action spaces.

4.1 Definition of plain consistency

Here is the logic for the definition of plain consistency. Consider a strategic player i, consecutive situations $h, h' \in \underline{H}_i$, a set $L \subset H$, and an L-product set $Y \subset S(h)$. Suppose that p^h exhibits L-independence on Y. This means that at h, player i puts positive probability on Y and believes that behavior on dimensions L and $-L$ are independent conditional on Y.

Now suppose that at h' player i puts positive probability on $Y_L \times \{s_{-L}\}$, for some strategy component $s_{-L} \in Y_{-L}$ for which $Y_L \times \{s_{-L}\} \subset S(h')$. Player i's updated belief at h' is, of course, based on what has been learned at situation h' following h. But notice that, if s_{-L} is the actual strategy on the $-L$ dimension, then at h' nothing has been learned about the L dimension to differentiate between elements of Y_L, since $Y_L \times \{s_{-L}\} \subset S(h')$. Thus, at h' player i's marginal belief about s_L conditional on $Y_L \times \{s_{-L}\}$ should be the same as it was at h conditional on Y.

In other words, if on Y player i's prior belief exhibits independence between behavior on dimension L and behavior on dimension $-L$, then to the extent that information about the two dimensions is conjunctive, player i can update on the two dimensions separately. In this case, plain consistency requires that if the conditional-probability formula applies to one of the dimensions, then it should characterize the updated belief on this dimension.

Definition: *Consider any appraisal system P, any $L \subset H$ and L-product set $Y \subset S$, and consecutive situations h and h' for any strategic player, such that $Y \subset S(h)$ and p^h exhibits L-independence on Y. Say that* ***the update from $\mathbf{h}$ to $\mathbf{h'}$ is $\mathbf{L}$-plainly consistent on $\mathbf{Y}$*** *if*

$$(p^{h'} \mid Y_L \times \{s_{-L}\})_L = (p^h \mid Y_L \times Y_{-L})_L$$

for every $s_{-L} \in Y_{-L}$ for which $Y_L \times \{s_{-L}\} \subset S(h')$ and $p^{h'}(Y_L \times \{s_{-L}\}) > 0$.

Recall that the definition of p^h exhibiting L-independence on Y includes $p^h(Y) > 0$. Note that L-plain consistency does not constrain the updated belief about dimension $-L$, unless the conditional-probability formula applies there as well (reversing L and $-L$).

To illustrate how plain consistency can be utilized in applications, consider again the game pictured in Figure 1, where h^3 is player 3's situation and h^2 is the situation at which player 2 chooses between actions w and z. The table on the right is the space of strategy profiles, separated in dimensions h^1 and $-h^1$, and the shaded region is $S(h^3)$. Note that $S(h^3)$ is not an h^1-product set, but the subset $Y = \{\mathrm{b}, \mathrm{c}\} \times \{(\mathrm{x}, \mathrm{z}, \mathrm{d}), (\mathrm{y}, \mathrm{z}, \mathrm{d}), (\mathrm{x}, \mathrm{z}, \mathrm{e}), (\mathrm{y}, \mathrm{z}, \mathrm{e})\}$ is an h^1-product set on which we can impose plain consistency.

For instance, suppose that player 3's appraisal at $\underline{h}_3$ puts probability 1 on strategy profile $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$, which happens to be a Nash equilibrium and also a weak PBE strategy profile. In the weak PBE, player 3's appraisal at h^3 puts high probability on b and z having been chosen by players 1 and 2. However, such a belief is ruled out by h^1-plain consistency on Y for the update from $\underline{h}_3$ to h^3. Since player 3's appraisal at h^3 specifies, on the h^1 margin, positive probability on $\{\mathrm{b}, \mathrm{c}\}$, the h^1-marginal conditioned on this subset must match the marginal probability at $\underline{h}_3$. That is, the surprise of arriving at situation h^3 requires player 3 to change her belief about player 2's choice between w and z, but it provides no knowledge about player 1's choice between b and c. Therefore, plain consistency requires player 3's belief about player 1's choice between b and c to be the same as it was at the beginning of the game, namely that c was chosen. With this application of plain consistency, there is no PBE in which c and w are played with positive probability.

While I would regard as easy the application of plain consistency described above, it may be challenging to evaluate a large number of Y sets in more complicated games. One may want to limit attention to a smaller number of key cases. A particularly straightforward utilization of plain consistency is to impose the condition only for consecutive situations having the following property:

Definition: *Say that* ***situation*** $\boldsymbol{h \in \underline{H}}$ ***is L-conjunctive*** *if $S(h)$ is an L-product set.*

In this case, at h, the active player's knowledge about the strategy profile is expressed as a conjunction combining knowledge about dimension L and knowledge about dimension $-L$. If consecutive situations $h, h' \in \underline{H}$ are L-conjunctive and p^h treats the two dimensions as independent, then updating on the two dimensions can be handled separately by considering the conditioning set $Y = S(h')_L \times S(h)_{-L}$ in the plain-consistency definition.

Definition: *Call an appraisal system* ***m-plainly consistent*** *(m for moderate) if, for every $L \subset H$ and every L-conjunctive pair of consecutive situations $h, h' \in \underline{H}$ such that p^h exhibits L-independence on $S(h')_L \times S(h)_{-L}$, the update from h to h' is L-plainly consistent on $S(h')_L \times S(h)_{-L}$.*

Under m-plain consistency, nothing further is implied by looking at the belief conditional on any strict subsets of $S(h')_L \times S(h)_{-L}$. The consistency condition comes down to applying the conditional-probability formula separately on each dimension as applicable. As a byproduct, independence is preserved if the conditional-probability formula applies to at least one dimension.

Lemma 1: *An appraisal system is m-plainly consistent if and only if for every $L \subset H$ and every L-conjunctive pair of consecutive situations $h, h' \in \underline{H}$ for which p^h exhibits L-independence on $S(h')_L \times S(h)_{-L}$, it is the case that $p^{h'}$ exhibits L-independence on S, and $p_L^{h'} = p_L^h(\cdot \mid S(h')_L)$.*

Checking m-plain consistency should be relatively easy for many games. In some applications, one may get enough mileage from imposing the condition on a small number of L sets and situations. A broader use of plain consistency would look for instances in which, for some L-product set Y, $S(h) \cap Y$ is an L-product set, $S(h')_L = Y_L$, and $S(h') \cap Y$ is an L-product set. Further broadening the usage would include instances in which $S(h') \cap Y$ is not an L-product set. The concept is flexible in that plain consistency can be imposed in as many or few places in a game as the analyst views as useful and where the premise of the definition holds.

For beliefs derived from a conditional-probability system, moderate plain consistency can be strengthened to constrain probabilities in relation to all conditioning events on a component of the strategy space, as evaluated on consecutive situations. The following definition imposes the condition on both dimensions L and $-L$, if the belief at the former situation regarding s_L puts positive probability on the latter situation being reached.

Definition: *Call an appraisal system* ***cps/m-plainly consistent*** *if it is derived from a conditional-probability system ζ with the following property: For every $L \subset H$ and L-conjunctive consecutive situations $h, h' \in \underline{H}$, if $\zeta(\cdot \mid S(h))$ exhibits L-independence on S and either $(\zeta \mid S(h))_L(S(h')_L) > 0$ or $(\zeta \mid S(h))_{-L}(S(h')_{-L}) > 0$, then $\zeta(\cdot \mid S(h'))$ exhibits L-independence and $(\zeta \mid Y_L \times S(h')_{-L})_L = (\zeta \mid Y_L \times S(h)_{-L})_L$ for all $Y_L \subset S(h')_L$.*

4.2 Implications of moderate plain consistency

This subsection describes some of the implications of plain consistency, notably the simple application of moderate plain consistency.

Minimal consistency

Moderate plain consistency implies minimal consistency, and therefore the one-deviation property holds for sequential rationality.

Lemma 2: *If an appraisal system is m-plainly consistent then it is minimally consistent.*

The reason is that, for every strategic player i and every situation $h \in \underline{H}_i$, $S(h)$ is an $\underline{H}_i$-conjunctive set. Therefore, moderate plain consistency requires player i's belief about the others' strategy profile s_{-i} to be updated using the conditional-probability formula as applicable (separate from player i's own past and future planned actions).

Subgame perfection

Moderate plain consistency clearly implies that players do not change their beliefs about, or plans for, behavior at situations that have not yet been reached in the game. In the equilibrium context, this implies subgame perfection.

Theorem 1: *If P is an m-plain perfect Bayesian equilibrium of a finite game, then the equilibrium strategy profile σ is a subgame-perfect Nash equilibrium.*

A component of the proof of this theorem is the following observation:

Lemma 3: *Given a game, consider any subgame and let J be the set of situations in the subgame. Then every $h \in \underline{H}_+$ is J-conjunctive.*

To see why Theorem 1 is true, consider any subgame and let J be the set of situations in the subgame. At the initial situations, the players have the same appraisal, which is the equilibrium strategy profile, and this strategy profile specifies behavior in the subgame. Lemmas 1 and 3 imply that, as players update beliefs on a path to the subgame, their beliefs about play in the subgame do not change, for nothing has been learned about play in the subgame before it is reached. The proof then shows that rationality at whatever situation a player first acts in the subgame implies that the continuation strategy in the subgame is a Nash equilibrium.

Reasonable updating in a standard class of multistage games

On multi-stage games with observed actions and independent types, plain consistency produces some of the implications of Fudenberg and Tirole's (1991) reasonableness condition, which I refer to as FT-reasonableness, described in the Introduction. Moderate plain consistency suffices if we model nature as separated into agents who simultaneously choose the types of the strategic players.

To be precise, there are n strategic players and correspondingly n agents of nature called $01, 02, \ldots, 0n$. The game begins at stage 0, when, simultaneously, each agent $0i$ of nature, at a situation denoted by h^{0i}, chooses the type of strategic player i, denoted by θ_i. Types are drawn independently according to a given probability distribution with full support, and each strategic player i observes her own type θ_i but nothing else. Then in each of stages $1, 2, \ldots, T$, the strategic players simultaneously choose actions, where the sets of feasible actions may depend on previous actions but not on the types. At the end of each stage, the players observe the vector of actions chosen. Payoffs depend on the type vector and sequence of action profiles.

By imposing that the $-i$ players update their belief about θ_i from one stage to the next as a function of only player i's action in the former stage (even upon observing an unexpected action profile), FT-reasonableness makes sense because none of the other players have any direct information about player i's type. Moderate plain consistency has a similar implication, as follows: For any strategic player i, let $L^i \equiv \{h^{0i}\} \cup H_i$. For consecutive situations (h, h') of any strategic player and for any other strategic player i, if p^h exhibits L^i-independence and $p^h_{L^i}(S(h')_{L^i}) > 0$ then FT-reasonableness is implied for the update about dimension L^i of the strategy profile. This follows from Lemma 1 and the next result.

Lemma 4: *Take as given a multi-stage game with observed actions and independent types (as specified above). Then every* $h \in \underline{H}_+$ *is* L^i*-conjunctive.*

Fudenberg and Tirole (1991) provide a result showing that, in multi-stage games with observed actions and independent types, "reasonable" updating is equivalent to full consistency. Reasonableness combines the FT condition with the assumption that the conditional-probability formula is used where applicable to update the belief about the type of each player i (conditioned on this player's action), that beliefs maintain independence across player-types, and that beliefs are derived from a conditional-probability system. The next result is based on weaker conditions that are simpler to state.[17]

[17]Fudenberg and Tirole's (1991) analysis requires developing an "extended belief system" that includes a separate conditional probability system on the type space for every situation.

Theorem 2: *Suppose that P is a perfect Bayesian equilibrium of a finite multi-stage game with observed actions and independent types (as specified above), under cps/m-plain consistency. Then there is a sequential equilibrium that has exactly the same outcome distribution.*

I claim that Theorem 2 can be extended to multistage games in which we model nature as choosing the entire vector of types at one situation, by employing plain consistency more broadly than in the moderate sense. The application discussed in Section 5 is in the same direction. Further discussion of reasonableness and implications of independence may be found in Section 6.

4.3 Generalization for infinite games

An advantage of the plain-consistency concept is that it can be applied to infinite games, with the appropriate modifications to deal with distributions on infinite spaces. A general version of the condition can be stated for any application in which the assumed measurable spaces have an appropriate product structure and appraisals can be "disintegrated" between the two dimensions under consideration.

Definition: *Say that a given set $L \subset H$ satisfies the* ***tensor-product condition*** *if S, S_L, and S_{-L} are endowed with sigma-algebras $\mathcal{B}$, $\mathcal{B}_L$, and $\mathcal{B}_{-L}$ respectively, such that $\mathcal{B}$ is the tensor-product sigma-algebra in relation to $\mathcal{B}_L$ and $\mathcal{B}_{-L}$. Say that a given distribution $p \in \Delta S$ satisfies the* ***L-disintegration condition*** *if there exists a function $\kappa_L \colon S_L \times S_{-L} \to \mathbb{R} \cup \{\infty\}$, called the L-disintegration probability kernel, such that*

(a) for every $X_L \in \mathcal{B}_L$, $\kappa_L(X_L, \cdot\,)$ is a $\mathcal{B}_{-L}$-measurable function,

(b) for every $s_{-L} \in S_{-L}$, $\kappa_L(\,\cdot\,, s_{-L})$ is a probability measure on $(S_L, \mathcal{B}_L)$, and

(c) $p(X_L \times X_{-L}) = \int_{X_{-L}} \kappa_L(X_L, s_{-L}) p_{-L}(ds_{-L})$ for all $X_L \in \mathcal{B}_L$ and $X_{-L} \in \mathcal{B}_{-L}$.

We interpret $\kappa_L(X_L, s_{-L})$ as the probability of X_L conditional on s_{-L}, which in the case of finite S is obviously well defined so long as $p_{-L}(s_{-L}) > 0$. Depending on the application, one could assume the L-disintegration condition directly, or use an existence result for the spaces in the representation of the game.[18]

For an L-disintegration probability kernel κ_L, a point $s_{-L} \in S_{-L}$, and a set $Y_L \in \mathcal{B}_L$ for which $\kappa_L(Y_L, s_{-L}) > 0$, define

$$\kappa_L(X_L, s_{-L} \,|\, Y_L) \equiv \frac{\kappa_L(X_L \cap Y_L, s_{-L})}{\kappa_L(Y_L, s_{-L})} \tag{1}$$

for every $X_L \in \mathcal{B}_L$, interpreted as the kernel probability of X_L conditional on $Y_L \times \{s_{-L}\}$.

The general definition of plain consistency follows.

[18]Baccelli, Blaszczyszyn, and Karray (2024, Section 14.D) prove that a sufficient condition for a measure on a product space to have a disintegration probability kernel, and thus to satisfy the L-disintegration condition, is that the output space, $(S_L, \mathcal{B}_L)$ above, is Polish.

Definition: *[Generization] Take as given any $L \subset H$ that satisfies the tensor-product condition, an appraisal system P, and consecutive situations h and h' for any strategic player, such that $p^{h'}$ satisfies the L-disintegration condition. Consider any L-product set $Y \subset S$, such that $Y \subset S(h)$ and p^h exhibits L-independence on Y. Say that* ***the update from h to h′ is L-plainly consistent on Y*** *if*

$$\kappa_L(\,\cdot\,, s_{-L}\,|Y_L) = (p^h\,|Y_L \times Y_{-L})_L$$

for every $s_{-L} \in Y_{-L}$ for which $Y_L \times \{s_{-L}\} \subset S(h')$ and $\kappa_L(Y_L, s_{-L}) > 0$, where κ_L is the L-disintegration probability kernel associated with $p^{h'}$.

Recall that $p(Y) > 0$ is required for p^h to exhibit L-independence on Y. This definition of plain consistency can be further generalized to allow $p(Y) = 0$ by extending the definition of $p^h(\,\cdot\,|\,\cdot\,)$ to allow such a conditioning set, and in the definition of "independence on," replace $p(Y) > 0$ with the requirement that $p^h(\,\cdot\,|Y)$ is well defined.[19] The arbitrary extension of the conditional probability formula would have to suit the application.

5 Application: Sender/Designer Games

This section explores the implications of plain consistency for a prominent class of games in which one player communicates or arranges to provide information to the others about a move of nature.

5.1 Description of the class of games

Let us use the name *sender/designer game* for any game in which n strategic players and nature interact in the following three stages:

1. Nature chooses the *state* $\theta \in \Theta$ according to a distribution $\sigma_0 \in \Delta\Theta$ with full support.
2. Player 1 privately observes $\phi = f(\theta) \in \Phi$ and then chooses $e \in E$.
3. For each $i \in N_+$, player i privately observes $\lambda_i = g_i(e, \theta) \in \Lambda_i$, and then simultaneously the players choose actions.

Sets Θ, Φ, E, and Λ_i for $i \in N_+$, along with functions $f\colon \Theta \to \Phi$ and $g_i\colon E \times \Theta \to \Lambda_i$ for $i \in N_+$, are given and commonly known. The game may end after the third stage, or it may continue in additional stages.

We shall characterize the implications of plain consistency for each player i's belief about θ upon observing λ_i in the third stage. But first I note some of the standard game-theoretic models in the sender/designer class.

[19] For instance, if $S_L = [0, 3]$, p^h_L is the uniform distribution, and $Y = \{1, 2\} \times S_{-L}$, then it would be appropriate to define $p^h(\,\cdot\,|Y)_L$ as the distribution that puts probability $1/2$ on each of 1 and 2. Extending the conditional probability in this way essentially defines a partial conditional-probability system with a small collection of conditioning sets.

Cheap-talk (Crawford and Sobel (1982)). Player 1 is the sender and the other players are receivers. Set Θ is typically taken to be an interval subset of a Euclidean space, Φ is a partition of Θ or homeomorphic to such a partition, $f(\theta)$ is the partition element containing θ, E is a message space, $\Lambda_i = E$ for every $i \in N_+$, and g_i is defined by $g_i(e, \theta) = e$. That is, the sender observes the partition element containing the true state and then sends a message to the receiver or receivers, who jointly observe only the message. In the standard model, the sender observes the state perfectly, which means Φ is the finest partition, and the single receiver (player 2) takes an action in the third stage. The game ends with the receiver's action choice. Payoffs are a function of the receiver's action and the state.

Hard evidence/persuasion.[20] Player 1 observes something about the state and may possess hard evidence, in the form of documents or other items whose existence (and thus availability for disclosure) depends on the state. State space Θ typically is taken to be an interval subset of a Euclidean space, and Φ is a partition of Θ. Set E is the set of possible evidentiary actions (subsets of documents to disclose); it is assumed to contain the empty action $\emptyset$ meaning no disclosure. For every $\theta \in \Theta$, a given subset $E^\theta \subset E$ defines the evidentiary actions that are feasible in state θ, where $E^\theta = E^{\theta'}$ for any pair θ and θ' that are in the same element of partition Φ. For every $\theta \in \Theta$, $f(\theta)$ is the partition element containing θ. Further, for every $i \in N_+$ we define $\Lambda_i \equiv E$, and we define g_i as follows: For every $\theta \in \Theta$ and $e \in E$, let $g_i(e, \theta) = e$ if $e \in E^\theta$ and let $g_i(e, \theta) = \emptyset$ if $e \notin E^\theta$.

In words, player 1 observes the partition element containing the true state, knows the corresponding set of feasible evidentiary actions, and then picks an evidentiary action (disclosure choice). If the disclosure is feasible then this is what the other players observe; otherwise they observe the empty disclosure. The game ends with the third-round action choices. Payoffs are a function of these actions, the state, and also the evidentiary action in the case of costly evidence.

Information design.[21] Nature's role is to choose a payoff-relevant value $\omega \in \Omega$ according to a given distribution $\sigma_0^\omega \in \Delta\Omega$ with full support, and to conduct a signaling experiment chosen by player 1. Let Z be the set of possible signal vectors of the form $(z_1, z_2, \ldots, z_n)$, where z_i denotes the signal privately observed by player i. An experiment is a function $e\colon \Omega \to \Delta Z$ that gives a distribution of signal vectors contingent on ω. Take as given a set E of feasible experiments. Assume that for every experiment e, nature's feasible draws are given by $Z^e \equiv \operatorname{supp} e$, which will imply that nature's strategy has full support. Let $B \equiv \prod_{e \in E} Z^e$ and let $\beta = (\beta^e)_{e \in E}$ denote a generic element, where $\beta^e = z^e$.

We can think of nature as choosing ω and privately conducting all experiments in the first stage. Player 1 obtains information about ω and then picks an experiment. The other players observe the identity of the experiment, and each player i privately observes her realized signal. Thus, $\Theta \equiv \Omega \times B$, and σ_0 is the product of σ_0^ω and every $e \in E$. Further, Φ is a partition of Ω, and for every $\theta = (\omega, \beta) \in \Theta$, $f(\theta)$ is the partition element containing ω. Finally, for every $i \in N_+$ we define $\Lambda_i \equiv E \times Z_i$, and we define g_i as follows: For every $\theta = (\omega, \beta) \in \Theta$ and

[20] In the literature on strategic settings with hard evidence, key conceptual elements were developed in various entries, including Green and Laffont (1986), Milgrom and Roberts (1986), Shin (1991), Lipman and Seppi (1995), and Bull and Watson (2004, 2007). The description here allows for statistical evidence (Bull and Watson 2019)).

[21] The description here allows for interim information design (Koessler and Skreta 2022).

$e \in E$, let $g_i(e, \theta) = (e, \beta_i^e)$. The players simultaneously choose stage-three actions, after which the game ends. Payoffs are a function of these actions and ω.

Mechanism design by an informed principal (Myerson (1983)). In this setting, nature chooses a payoff-relevant vector of types $\omega = (\omega_1, \omega_2, \ldots, \omega_n)$ from a set Ω according to a given distribution $\sigma_0^\omega \in \Delta\Omega$ with full support, and nature also serves as a mediation system for communication between the strategic players and an external enforcer, which is a mechanism chosen by player 1.

As in the previous setting, we can think of nature conducting the random draws for every mechanism in stage one. In stage two, player 1 observes her type ω_1 and selects a mechanism. In stage three, all players observe the identity of the chosen mechanism, each player $i \in N_+$ observes their type ω_i, and the players privately report their types to the mechanism. In stage four, the mechanism sends each strategic player $i \in N_+$ a private message, which is an element of a set D_i and is interpreted as the suggested action to take in stage five. Concurrently, the mechanism privately directs an external enforcer on what public action to take. Then in stage five, the strategic players simultaneously choose productive actions, where D_i is the set of feasible actions for player $i \in N_+$, and the external enforcer takes the public action from the set D_0 that the mechanism directed. The game ends after stage five, and payoffs are a function of ω and the productive-action profile $d = (d_0, d_1, \ldots, d_n) \in D \equiv \prod_{i \in N} D_i$.

Let $\overline{\Omega} \equiv \prod_{i \in N_+} \Omega_i$, where Ω_i is the set of feasible types for player $i \in N_+$, and note that Ω can be a strict subset of $\overline{\Omega}$ if types are correlated. The set of mechanisms is given by $E = \{e \colon \overline{\Omega} \to \Delta D\}$. As in the previous setting, assume that for every mechanism e and type-report vector $\hat{\omega}$, nature's set of feasible recommendation vectors is $Z^e(\hat{\omega}) \equiv \operatorname{supp} e(\hat{\omega})$, which will imply that nature's strategy has full support. Let $B \equiv \prod_{e \in E, \hat{\omega} \in \overline{\Omega}} Z^e(\hat{\omega})$ and, for a generic element β, let $\beta^e(\hat{\omega})$ denote the vector of recommended actions under mechanism e and report vector $\hat{\omega}$.

Thus, the structure of the first three stages is described as follows: $\Theta \equiv \Omega \times B$, and σ_0 is the product of σ_0^ω and $e(\hat{\omega})$ for every $e \in E$ and $\hat{\omega} \in \overline{\Omega}$. Further, $\Phi = \Omega_1$ and, for every $\theta = (\omega, \beta) \in \Theta$, we set $f(\theta) = \omega_1$. Finally, for every $i \in N_+$ we define $\Lambda_i \equiv E \times \Omega_i$, and we set $g_i(e, \theta) = (e, \omega_i)$ for every $\theta = (\omega, \beta) \in \Theta$. The players' observations and feasible actions in stage four are similarly described. Payoffs are a function of the type vector and productive actions.

5.2 Characterization of stage-three beliefs about the state

To keep notation simple, let λ_i denote the situation in stage three in which player i has observed λ_i, and let ϕ denote the situation in stage two where player 1 has observed $\phi \in \Phi$. I will thus refer to Φ as the set of situations for player 1 in stage two, s_Φ denotes player 1's strategy on these situations, and $-\Phi$ refers to the situations other than Φ. Also, let $M \equiv H \backslash (H_0 \cup \Phi)$ denote the set of situations in stage three and later. For any $\lambda_i \in \Lambda_i$ and $s_\Phi \in S_\Phi$, let

$$Y_0(\lambda_i, s_\Phi) \equiv \{\, \theta \mid \lambda_i = g_i(s_{f(\theta)}, \theta)\} \tag{2}$$

be the set of states for which (θ, s_Φ) reaches situation λ_i. Note that, because at situation λ_i, no other situations in M have yet been reached, $S(\lambda_i)$ is an M-product set and $S(\lambda_i) = S_M \times S_{-M}(\lambda_i)$.

Let us explore the implications of plain consistency for the updated belief about θ. If any specific λ_i would arise on the equilibrium path, then, under minimal consistency, the conditional-probability formula would apply to calculate player i's belief at situation λ_i about θ and about player 1's strategy component s_Φ. If λ_i is not on the equilibrium path, then anything goes under minimal consistency, and player i's belief at λ_i about θ is arbitrary. Plain consistency, however, has sharp implications for the belief about the state.

The next definition describes the regularity conditions on the appraisal system needed to impose plain consistency on the belief at λ_i. The second definition is the application of plain consistency at such a situation for dimensions M and $-\Phi$; the former ensures that beliefs about actions taken before stage three are treated independently of beliefs about actions in stage three and beyond, and the latter ensures that updating about player 1's actions in stage two are isolated from all other actions, with respect to conjunctive events.

Definition: *Call a sender/designer game and appraisal system P **amenable** if the following conditions hold: First, the tensor-product condition is satisfied by H_0, H_Φ, and M. Second, $Y_0(\lambda_i, s_\Phi)$ is measurable and $p^{\underline{h}_i}(\,\cdot\,|Y_0(\lambda_i, s_\Phi) \times S_M \times \{s_\Phi\})$ is well-defined, for every $i \in N_+$, $\lambda_i \in \Lambda_i$, and $s_\Phi \in S_\Phi$. Third, for every $i \in N_+$ and $\lambda_i \in \Lambda_i$, (a) $p^{\underline{h}_i}$ exhibits H_0-independence, Φ-independence, and M-independence on S, and (b) p^{λ_i} satisfies the L-disintegration condition for $L = M$ and $L = -\Phi$.*

Definition: *Given an amenable sender/designer game and appraisal system P, call the appraisal system **pc-compliant** if, for every $i \in N_+$ and $\lambda_i \in \Lambda_i$, the update from $\underline{h}_i$ to λ_i is both M-plainly consistent on S and $-\Phi$-plainly consistent on $Y_0(\lambda_i, s_\Phi) \times S_{-0}$ for every $s_\Phi \in S_\Phi$.*

To understand the next result, note that $p_\Phi^{\lambda_i}$ denotes player i's belief at situation λ_i about player 1's strategy component s_Φ (the Φ margin of p^{λ_i}) and $p_0^{\lambda_i}$ is player i's belief about nature's strategy θ.

Theorem 3: *Consider any amenable sender/designer game and appraisal system P, and let $Y_0(\lambda_i, s_\Phi)$ be defined by Equation (2). Then P is pc-compliant only if, for every player $i \in N_+$ and $\lambda_i \in \Lambda_i$,*

(a) p^{λ_i} *exhibits M-independence on S and satisfies* $p_M^{\lambda_i} = p_M^{\underline{h}_i}$, *and*

(b) $p_0^{\lambda_i} = \displaystyle\int_{S_\Phi} p_0^{\underline{h}_i}(\,\cdot\,|Y_0(\lambda_i, s_\Phi))\, p_\Phi^{\lambda_i}(ds_\Phi)$.

Likewise, the converse holds (P satisfying conditions (a) and (b) implies pc-compliance) for appropriately defined M- and $-\Phi$-disintegration kernels.

Theorem 3's characterization of beliefs regarding nature's move is simple to state in words: After observing λ_i, player i's belief about the state θ can be expressed by a two-stage distribution: First player 1's strategy component s_Φ is drawn according to distribution $p_\Phi^{\lambda_i}$, which is player i's updated belief about s_Φ at situation λ_i. Then θ is drawn according to the conditional distribution $p_0^{h_i}(\,\cdot\,|Y_0(\lambda_i, s_\Phi))$, which player i's initial belief about nature's strategy, conditioned on the subset of Θ that reaches λ_i for the given strategy component of player 1.

For example, if λ_i is a surprise event (assigned zero probability by player i's belief at the beginning of the game), then player i's explanation is that player 1 must have deviated and nature did not deviate. The deviation of player 1 cannot be used as a excuse for player i to update the belief about the state, except insofar as it reflects player 1's own information about the state. Updating does not allow player 1 to signal something that player 1 did not observe.

In fact, the updated belief can also be written as the conditional-probability update of a new prior belief that is identical to $p^{\underline{h}_i}$ except that player 1's second-stage behavior is given by some $\hat{\sigma}_\Phi$ that reaches situation λ_i. This is a form of *structural consistency*.[22] For instance, consider the setting in which Θ, Φ, E and Λ are finite, and let distributions $\sigma_0 \in \Delta S_0$, $\sigma_\Phi \in \Delta S_\Phi$, and $\sigma_M \in \Delta S_M$ be such that $p^{\underline{h}_i}(s) = \sigma_0(s_0)\sigma_\Phi(s_\Phi)\sigma_M(s_M)$ for every $s \in S$. Define $\hat{\sigma}_\Phi \in \Delta S_\Phi$ by

$$\hat{\sigma}_\Phi(s_\Phi) = \frac{p_\Phi^{\lambda_i}(s_\Phi)\eta}{\sigma_0(Y_0(\lambda_i, s_\Phi))}$$

for every s_Φ for which $Y_0(\lambda_i, s\Phi) \neq \emptyset$, and otherwise set $\hat{\sigma}_\Phi(s_\Phi) = 0$, where

$$\eta \equiv \left(\sum_{\{s_\Phi | Y_0(\lambda_i, s\Phi) \neq \emptyset\}} \frac{p_\Phi^{\lambda_i}(s_\Phi)}{\sigma_0(Y_0(\lambda_i, s_\Phi))} \right)^{-1}.$$

Define $\hat{\sigma} \in \Delta S$ by $\hat{\sigma}(s) = \sigma_0(s_0)\hat{\sigma}_\Phi(s_\Phi)\sigma_M(s_M)$. It is not difficult to verify that, under pc-compliance, $p^{\lambda_i}(s) = \hat{\sigma}(s|S(\lambda_i))$ for every $s \in S(\lambda_i)$. One can also verify that Theorem 3's characterization of beliefs regarding nature's move is exactly what full consistency implies.

The upshot is that, for finite sender/designer games, plain consistency has the same implications for stage-three beliefs as does full consistency but has the advantage of being a regularity condition on updating rather than being derived from a sequence of fully mixed behavior strategies. Further, with plain consistency, the characterization of stage-three beliefs generalizes to the class of infinite sender/designer games, with appropriate assumptions on the space of probability distributions.

5.3 Possible usefulness for applications

As noted in the Introduction, many entries in the literature on sender/designer games are ambiguous on belief consistency, by either (a) not formally defining the perfect Bayesian equilibrium concept to be utilized, or (b) not reporting analysis of a fully specified noncooperative game. For instance, looking at the ten articles with sender/designer elements that were published in the "top-five" general-interest economics journals in 2022-2023, I put one in the first category and seven in the second category.[23] The intent is not to hyper-criticize

[22]While it characterizes the updated beliefs in the application here, in general finite games, structural consistency is not equivalent to full consistency (Kreps and Ramey (1987)).

[23]In the first category I classified Kattwinkel and Knoepfle (2023), and in the second Au and Whitmeyer (2023), Bowen et al. (2023), de Clippel and Zhang (2022), Dworczak and Pavan (2022), Lipnowski, et al.

the literature. I merely wish to suggest there is room to clarify the foundations of applied theoretical analysis. One of my own articles, Watson (1996), is in the second category.[24]

Plain consistency may thus be useful for clarifying, and in some cases provide a foundation for, various modeling exercises in the sender/designer literature. A reasonable standard for application of noncooperative game theory is to fully describe the game to be studied as well as the solution concept to be utilized. For modeling exercises where it is convenient to present a partial analysis of the game and to focus on specific outcomes that emerge from the partial analysis, a good starting point would be to specify a solution concept that justifies the focus. In numerous examples, defining perfect Bayesian equilibrium based on plain consistency would complete the game-theoretic exercise and rule out outcomes that were previously dispensed with implicitly. In particular, Theorem 3 provides a foundation for the structural-consistency assumption that Myerson (1983), Koessler and Skreta (2023), and a few others make directly on belief systems. In any case, one might differentiate between possible predictions based on plain consistency and other consistency notions.

5.4 Comments on equilibrium existence

As for any solution concept, the issue of equilibrium existence arises with respect to versions of perfect Bayesian equilibrium for infinite games, where the existence theorems of trembling-hand perfect and sequential equilibrium do not apply. Existence is typically not a hurdle for some kinds of sender/designer games, such as cheap-talk settings in which a babbling equilibrium exists under mild assumptions, and similarly information-design settings with an uninformed designer. Existence is a more difficult proposition for other sender/designer games, such as general signaling games and informed-designer/principal games.

Presumably existence can be dealt with for narrow classes of games, and it comes down to familiar considerations and technical barriers (continuity, compactness, monotonicity, and so on). Whereas imposing plain consistency or any other consistency restriction may advance theoretical goals in an application, limiting the scope may grease the wheels of an existence argument by allowing for a wide range of beliefs where it may be useful in a proof.

Without getting too deep into the weeds, here is an example of an existence argument for a class of signaling games utilizing a method that focuses attention on a finite subset of an infinite space. There are n strategic players (a finite number). At stage one of the game, nature chooses a vector of types $\theta = (\theta_0, \theta_1, \theta_2, \ldots, \theta_n)$ according to a probability distribution μ that has full support on a finite product space $\Theta = \Theta_0 \times \Theta_1 \times \cdots \times \Theta_n$. Each player i will later choose an action $a_i \in A_i$, where A_i is a compact subset of the reals, we let $A \equiv A_1 \times A_2 \times \cdots \times A_n$, and the payoffs are given by a function $w\colon \Theta \times A \to \mathbb{R}^n$. In stage

(2022), McClellan (2022), and Perego and Yuksel (2022). Separately, Koessler and Skreta (2023) provide an equilibrium definition in line with Myerson (1983) that directly assumes structural consistency for stage-three beliefs.

[24]My own interest in consistency conditions for belief systems began early on, including in research to produce Shimoji and Watson (1998) and Watson (1998). Further, Pierpaolo Battigalli pointed out that I had implicitly assumed a cross-player consistency condition in Watson (1993), which led to the corrective Battigalli and Watson (1997). Research leading to the present paper began in 2011 when I aimed to develop a strong perfect Bayesian equilibrium concept for models of decentralized contractual networks, although the first resulting paper, Watson (2024), utilized sequential equilibrium for a complicated equilibrium construction.

two, player 1 observes θ_1 and chooses a_1, which everyone else observes. At stage three, each player $i \in \{2, 3, \ldots, n\}$ observes θ_i, and these players simultaneously choose their actions $a_2, \ldots, a_n$, ending the game.

Such a sender/designer game is a multi-stage game with observed actions, but it is generally not finite and does not necessarily have independent types. Let us adopt pc-compliance as the belief-consistency restriction for perfect Bayesian equilibrium. Deferring discussion of the assumptions on A and w that will be needed, consider the following method of attack for an existence proof.

Note that any action a_1 and distribution $\mu' \in \Delta\Theta$ (a posterior about nature's move) define a stage-three continuation game that is a Bayesian game in which nature chooses θ according to distribution μ', players 2 through n observe their types, and they simultaneously select their actions. Suppose that this continuation game has a (Bayesian) Nash equilibrium, let us arbitrarily select such an equilibrium, and let $w_1^*(a_1, \mu', \theta_1)$ denote the equilibrium expected payoff of player 1's type θ_1. We'll need an assumption implying that w_1^* is continuous in μ', and that guarantees the existence of a distribution $\underline{\mu} \in \Delta\Theta$ such that

$$w_1^*(a_1, \mu', \theta_1) \geq w_1^*(a_1, \underline{\mu}, \theta_1) \tag{3}$$

for every $\mu' \in \Delta\Theta$ that satisfies belief consistency, $a_1 \in A_1$, and $\theta_1 \in \Theta_1$. We'll also need $w_1^*(a_1, \underline{\mu}, \theta_1)$ to have a maximum by choice of a_1.

A constructive algorithm could then run as follows. First, for each $\theta_1 \in \theta_1$, find $a_1^{\theta_1}$ to maximize $w^*(a_1, \underline{\mu}, \theta_1)$, and define $\hat{A}_1 \equiv \{a_1^{\theta_1} \mid \theta_1 \in \Theta_1\}$. Our perfect Bayesian equilibrium will specify that the types of player 1 put zero probability on actions outside of $\hat{A}_1$, and if such an action is chosen then players 2 through n hold posterior belief $\underline{\mu}$ and their play in the stage-three continuation is that which leads to payoff w_1^* for the types of player 1. No type θ_1 would want to deviate to such an action, since $a_1^{\theta_1}$ would yield a weakly higher payoff given the definition of $a_1^{\theta_1}$ and Inequality 3.

Second, look at an artificial game between player 1's types whereby they choose mixed actions in $\Delta^\varepsilon \hat{A}_1$ and receive payoffs given by w_1^*, where $\Delta^\varepsilon \hat{A}_1$ denotes the (compact) subset of $\Delta \hat{A}_1$ in which each action gets a probability of at least ε. Using standard methods, we can find an equilibrium σ_1^ε of this artificial game. Third, by letting ε approach 0, we can construct a sequence $(\sigma_1^{1/k})_k$ that converges to some σ_1^* and from which we derive (using the conditional-probability formula) the posterior about nature's move that players 2 through n hold contingent on observing any action $a_1 \in \hat{A}_1$, as in a full-consistency construction. Coupled with posterior $\underline{\mu}$ for every other action, we have specified all stage-three beliefs. Player 1's equilibrium strategy is σ_1^* and sequential rationality holds by construction. More details and sufficient conditions may be found in Appendix A.4.

The beliefs of players 2 through n after observing any $a_1 \in \hat{A}_1$ conform to Theorem 3 by construction because they have the same formulation as with full consistency. Thus, the key constraint imposed by plain consistency is that $\underline{\mu}$, the off-equilibrium-path posterior about nature's move, also has the form required by Theorem 3. Weaker consistency conditions would loosen this requirement, with minimal and weak consistency imposing no constraint.

Moving to informed-designer games and mechanism design by an informed principal, it is worth noting that even existing equilibrium-existence results in these settings entail complications. I have already noted that plain consistency justifies the structure of stage-three beliefs that are directly assumed in some theoretical exercises such as Myerson (1983). However, using plain consistency, or any other consistency condition, to define perfect Bayesian equilibrium would not transform their equilibrium constructions into full-blown perfect Bayesian equilibria. For example, Myerson's (1983) analysis of incentives in stage five are put in terms of the expected payoffs from the start of stage four, calculated using the beliefs from stage three, as in a Bayes-Nash equilibrium calculation. This analysis does not imply sequential rationality for all situations that could be reached in stage five, because some of these situations may have zero probability given stage-three beliefs.[25] In other words, the equilibrium notion that Myerson (1983) uses does not require sequential rationality and hence is technically not in the "perfect" category. The same is true for other entries in the literature that followed, including Koessler and Skreta (2023).

6 More Definitions and Discussion

This section describes how to incorporate plain consistency, or any other consistency condition, into the concept of extensive-form rationalizability, presents a result on the relation between the various consistency concepts defined earlier, and provides further comments on related literature.

6.1 Extensive-form rationalizability

Pearce's (1984) notion of extensive-form rationaliability can be generalized to allow for any consistency condition on appraisals, as follows.[26] For any strategic player i, let $P_i = \left(p^h\right)_{h \in \underline{H}_i}$ denote player i's appraisal system (the appraisals at player i's situations), and let us refer to a given consistency condition applied to P_i as $\circledast$-consistency. Note that the definitions of "rooted in" and the consistent-plan property apply to P_i by restricting to these appraisals.

Definition: *Call $X \subset S$ a* ***player-product set*** *if it is an H_i-product set for every $i \in N_+$. Then, given a player $i \in N_+$ and player i's appraisal system P_i, say that* ***P_i conforms to X*** *if for every $h \in \underline{H}_i$ satisfying $S(h)_{-i} \cap X_{-i} \neq \emptyset$, it is the case that $p^h_{-i}(S(h)_{-i} \cap X_{-i}) = 1$.*

Given a player-product set X and a pure strategy $s_i \in S_i$, denote by $\mathcal{P}_i(X, s_i)$ the set of player i's appraisal systems that are $\circledast$-consistent, rooted in (s_i, σ_{-i}) for any behavior strategy profile $\sigma_{-i} \in \Delta X_{-i}$, have the consistent-plan property, and conform to X. If desired for the

[25] Here is an illustration: Suppose that player 1 deviates in stage two to selects a mechanism that would suggest action $\hat{d}_2$ to player 2 in the event that player 1 reports type $\hat{\theta}_1$ and otherwise would suggest another action. Suppose that in the equilibrium construction, upon observing this selected mechanism, player 2 puts zero probability on player 1's type $\hat{\theta}_1$. Then Myerson's incentive-compatibility condition does not constrain what player 2 believes or should do in stage five if the mechanism recommends $\hat{d}_2$ in stage four, although this is a situation in the game and it would be reached if type $\hat{\theta}_1$ of player 1 chose the deviation mechanism and then reported truthfully.

[26] The description here is a generalization of Shimoji and Watson's (1998) replacement-based definition and makes sense presuming at least minimal consistency of appraisals is assumed.

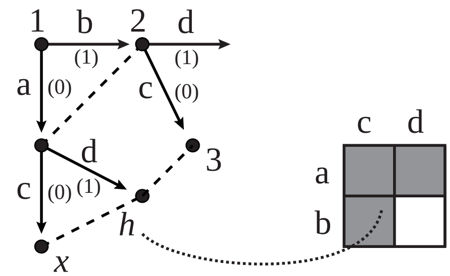

Figure 2: Example of belief regarding double deviation.

application, one can also take nature's strategy as given (a common prior). Define

$$B_i(X) \equiv \{s_i \in X_i \mid \text{there exists } P_i \in \mathcal{P}_i(X, s_i) \text{ such that}$$
$$p^h \text{ is rational for every } h \in H_i \text{ for which } S(h)_{-i} \cap X_{-i} \neq \emptyset\},$$

and let $B(X) \equiv \prod_{i \in N_+} B_i(X)$. Define operator B^k inductively by $B^0 \equiv B$ and $B^k(X) \equiv B(B^{k-1}(X))$ for $k = 1, 2, \ldots$.

Definition: *Take as given a consistency concept for beliefs, relative to which operator B is defined. If appropriate for the application, also take nature's strategy $(\sigma_h)_{h \in H_0}$ as given. The set of* ***rationalizable strategy profiles*** *for the strategic players is defined as the largest fixed point of B, which, for finite n-player games, is equal to $\cap_{k=0}^{\infty} B^k(S_1 \times S_2 \times \cdots \times S_n)$.*

Following Pearce (1984), rationality is imposed at only situations reached by X_{-i}. Under minimal consistency, this is Pearce's definition of rationalizability. The definition under plain consistency has already been featured in Bull and Watson's (2019) analysis of statistical evidence.

6.2 Additional notes and comparisons

The next example shows that plain consistency has some implications for disjunctive knowledge ("or" statements) and simultaneous surprise events. Consider the game fragment shown in Figure 2, where the game ends following actions d and b or after player 3's move at situation h. Suppose player 3's belief at the beginning of the game is as indicated in the game tree (that player 1 will select b and player 2 will select d with certainty) so that player 3 would be surprised to arrive at situation h. Because $S(h)_{-3}$ does not have a product structure, m-plain consistency does not address the update from $\underline{h}_3$ to h.

However, applying plain consistency for the update from $\underline{h}_3$ to h on S implies that p^h must put zero probability on both a and c having been chosen by players 1 and 2, which is the same implication of full consistency (zero probability on node x). The reason is that H_1-plain consistency on $S_1 \times S_{-1}$ implies that if $p_2^h(\{\mathrm{c}\}) > 0$ then p^h must put zero probability on both a and c being played. Likewise, letting H_2-plain consistency on $S_2 \times S_{-2}$ implies that if $p_1^h(\{\mathrm{a}\}) > 0$ then p^h must put zero probability on both a and c being played. This example shows that plain consistency is not equivalent to structural consistency, because the latter does not have such an implication.

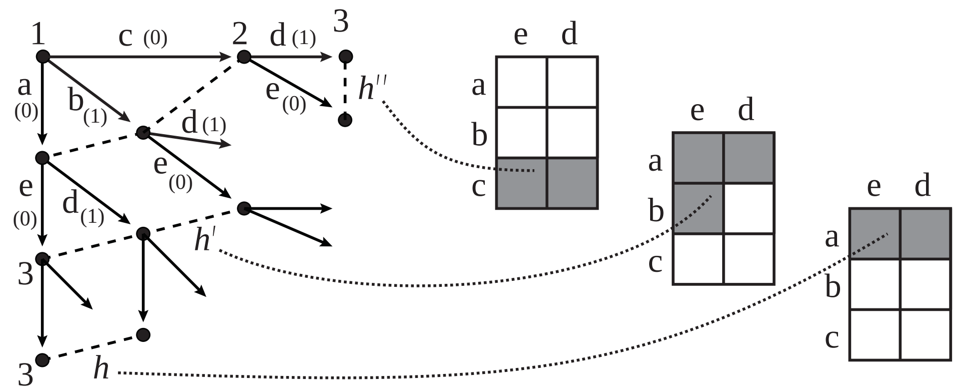

Figure 3: Example of path dependence.

A similar restriction can be derived for signaling games with two senders who jointly observe nature's choice of the state and simultaneously send messages to a receiver, with no hard evidence (all actions are available to every type). Following an off-equilibrium-path message profile that does not identify the deviator, plain consistency implies that the receivers must believe exactly one of the senders deviated and, conditional on which one, the posterior must satisfy structural consistency using the non-deviating player's equilibrium strategy. Any public unilateral deviation by one of the senders (not a case of one type pretending to be another type) will be identified by the receiving players, and in such a contingency the posterior of the receiving players must satisfy structural consistency using the non-deviating player's equilibrium strategy. Likewise, consider a asymmetric-information setting with a product type space and nature choosing the senders' types according to a distribution with full support, also with no hard evidence. Then any unilateral public deviation by one of the senders will be identified by the receiving players, and a similar structural consistency condition must hold.

Plain consistency does not require different players with the same information to have the same beliefs. Further, path dependence can arise in the beliefs of a single player, as the example shown in Figure 3 demonstrates. Suppose player 3's belief at the beginning of the game is as indicated in the game tree: that player 1 will select b and player 2 will select d with certainty. Plain consistency allows player 3 at situation h' to believe with certainty that actions b and e were chosen by the other players, and at h to have the same belief. However, plain consistency requires player 3's belief at h'' to put probability one on d. Full consistency and strong consistency require the appraisals at h and h'' to be the same.

Turning to the relation between consistency conditions, here is terminology to refer to the widest application of plain consistency:

Definition: *Call an appraisal system **f-plainly consistent** (f for fully) if it satisfies the plain-consistency condition for every applicable $L \subset H$, pair of consecutive situations, and L-product set Y.*

Recall that, by assuming any of the various consistency conditions in the equilibrium context, we obtain different versions of perfect Bayesian equilibrium (with sequential equilibrium as one version). The next result confirms what one would expect regarding the relation between the various consistency concepts discussed herein.

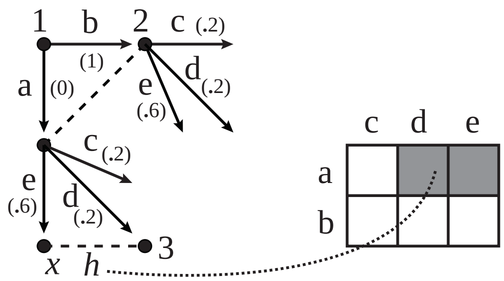

Figure 4: Another example.

Theorem 4: *As defined for finite games, consistency definitions are nested from strongest to weakest as follows: full, strong, f-plain, m-plain, and minimal.*

Other versions of perfect Bayesian equilibrium that have appeared in the literature include concepts defined by Bonanno (2013), González-Díaz and Meléndez-Jiménez (2014), Mailath (2016), and Doval and Ely (2020). Their consistency conditions apply to assessments and are engineered to strengthen weak PBE enough to ensure subgame perfection. Bonanno's (2013) definition has a plausibility-logic foundation ("AGM-consistency") that provides an intuitive lexic support restriction.[27] González-Díaz and Meléndez-Jiménez (2014) take the more direct approach of generalizing the notion of a subgame to "regular quasi-subtree," which they define as an situation and the nodes that follow, with the property that in this set each player knows he or she is in this set. González-Díaz and Meléndez-Jiménez require weak PBE on every regular quasi-subtree, which implies Nash equilibrium on every subgame as a special case. Mailath's (2016) definition of "almost PBE" is similar, imposing proper conditional-probability updating of assessments from one situation to the next (generally for a different player) in the case in which the latter situation follows with positive probability given the assessment at the former. Doval and Ely (2020) strengthen Mailath's definition by requiring beliefs to be derived from a conditional-probability system on the strategy space.

The PBE definitions developed in all of the entries cited above do not address conditional-probability updating on separate components of the strategy space, and in this critical sense are weaker than plain PBE. For example, in the game shown in Figure 4, with initial beliefs given by the indicated probabilities of actions, these definitions do not restrict the belief of player 3 at situation h regarding player 2's actions d and e, whereas plain consistency nails it down to probability $1/4$ on d and probability $3/4$ on e. Thus, these other definitions do not constrain beliefs along the lines of reasonableness, and they do not constrain stage-three beliefs in sender/designer games. On the other hand, their definitions all assume a greater level of conformity for the assessments of the various players, so in this sense they impose stronger requirements than does plain PBE.[28] Further, none of these alternative concepts in

[27]Bonanno requires Bayes' updating in the sense of minimal consistency (see the Appendix here), although it is characterized in terms of plausibility of nodes and a requirement that information-set assessments are derived by calculating the product of action probabilities on the feasible paths. In a given subgame, where plausibility of various nodes are anchored by the plausibility level of the subgame's initial node, the conditions imply that assessments at on-path situations (conditional on reaching the subgame) are consistent with the strategy profile, which implies subgame perfection.

[28]Weibell (1992) proposes another definition that sits between subgame perfection and perfect equilibrium; it

the literature are defined for infinite games.

7 Conclusion

This paper has endeavored to expand the toolbox of consistency conditions for the belief-based solution concepts, to support applied-theoretical analysis and to clarify the foundations for specific applications. It also promotes modeling beliefs as probability distributions over the space of strategy profiles, which is needed to represent complementary dimensions of the strategy space and also simplifies some of the definitions and logical steps. Whether or not the paper succeeds in these objectives, I hope that it helps theorists classify alternative assumptions for solution concepts, provides a pedagogical benefit, and encourages further research on consistency conditions and implications for classes of games beyond those studied herein.

A Appendix

This appendix contains additional technical material and proofs of the results. All subsections pertain to finite games except for Subsection A.4.

A.1 Behavior strategies in mixed-strategy space

Recall that I defined behavior strategy profiles as mixed strategies that satisfy h-independence on S for every $h \in H$. Thus, the set of behavior strategy profiles is a subset of ΔS (which otherwise allows for correlation across situations and players). In the prior literature, a behavior strategy profile is more commonly defined as a distribution $\pi_h \in \Delta S_h$ for every $h \in H$, with independence across situations. That is, π_h is a distribution over the actions available to the player on the move at h, and so the behavior strategy profile is an element of $\prod_{h \in H} \Delta S_h$, which is not a subset of ΔS.

The two ways of defining behavior strategies are, in fact, equivalent. Let $\underline{\Delta} S$ denote the set of mixed strategies that satisfy h-independence on S for every $h \in H$.

Lemma 5: *For any finite game and any distribution $\sigma \in \underline{\Delta} S$, it is the case that $\sigma(s) = \prod_{h \in H} \sigma_h(s_h)$ for every $s \in S$.*

Proof: Consider the following claim: For any given $\sigma \in \underline{\Delta} S$, $L \subset H$, and $h' \in -L$, if $\sigma(s) = \sigma_{-L}(s_{-L}) \cdot \prod_{h \in L} \sigma_h(s_h)$ for every $s \in S$, then $\sigma(s) = \sigma_{-L'}(s_{-L'}) \cdot \prod_{h \in L'} \sigma_h(s_h)$ for every $s \in S$, where $L' = L \cup \{h'\}$. This claim is proved as follows. From the premise of the claim, we have for every $s \in S$,

$$\sigma(s) = \sigma_{-L}(s_{-L}) \cdot \prod_{h \in L} \sigma_h(s_h).$$

---

is based on strategy trembles. Weibull examines beliefs as probability distributions over strategy profiles, as I advocate here.

For any $h' \in -L$, fixing any $s_{-h'}$ and summing over $s'_{h'} \in S_{h'}$, we obtain

$$\sum_{s'_{h'} \in S_{h'}} \sigma(s_{-h'}, s'_{h'}) = \sum_{s'_{h'} \in S_{h'}} \sigma_{-L}(s_{-L'}, s'_{h'}) \cdot \prod_{h \in L} \sigma_h(s_h).$$

The left side is equal to $\sigma_{-h'}(s_{-h'})$ and the summation term on the right is $\sigma_{-L'}(s_{-L'})$. Multiplying both sides by $\sigma_{h'}(s_{h'})$ yields:

$$\sigma_{h'}(s_{h'})\sigma_{-h'}(s_{-h'}) = \sigma_{-L'}(s_{-L'}) \cdot \prod_{h \in L'} \sigma_h(s_h).$$

Using the fact that σ exhibits h'-independence on S, we see that the left side is $\sigma(s)$, which gives the expression in the conclusion of the claim.

For any $s \in S$, starting with an arbitrary $h \in H$, for which we have $\sigma(s) = \sigma_h(s_h)\sigma_{-h}(s_{-h})$ by h-independence on S, we then apply the claim in iterative fashion over all other situations in H, producing the expression $\sigma(s) = \prod_{h \in H} \sigma_h(s_h)$. ■

Lemma 6: *For any finite game, there exists a bijection η: $\prod_{h \in H} \Delta S_h \to \underline{\Delta} S$ such that, for every $\pi \in \prod_{h \in H} \Delta S_h$, it is the case that π and $\sigma = \eta(\pi)$ induce the same probability distribution over play paths (terminal nodes) in the game.*

Proof: Bijection η is defined as follows. For every $\pi = (\pi_h)_{h \in H} \in \prod_{h \in H} \Delta S_h$, let $\sigma = \eta(\pi)$ be given by $\sigma(s) = \prod_{h \in H} \pi_h(s_h)$. Clearly $\sigma \in \underline{\Delta} S$ and $\sigma_h = \pi_h$ for every $h \in H$, by construction. Further, Lemma 5 impies that mixed strategy profiles in $\underline{\Delta} S$ are uniquely defined by their marginal distributions across $h \in H$. Therefore, η is a bijection.

To verify that σ and π induce the same distribution over play paths, consider any feasible path of actions $\rho = (a^1, a^2, \ldots, a^K)$ in the game, which is a terminal node in the extensive-form representation. Let $(h^1, h^2, \ldots, h^K)$ be the (uniquely identified) corresponding situations where these actions are taken; that is, a^k is chosen at h^k, for $k \in \{1, 2, \ldots, K\}$. If play is described by π, then the probability that path ρ occurs is $\prod_{k \in \{1,2,\ldots,K\}} \pi_{h^k}(a^k)$. If play is described by σ, then the probability that ρ occurs is $\sum_{s \in T^\rho} \sigma(s)$ where $T^\rho \equiv \{s \in S \mid s_{h^k} = a^k \text{ for } k \in \{1, 2, \ldots, K\}\}$. Let $H^\rho \equiv \{h^1, h^2, \ldots, h^K\}$. From the definition of T^ρ and that σ satisfies h-independence on S for every $h \in H$, we have that $\sum_{s \in T^\rho} \sigma(s) = \prod_{k \in \{1,2,\ldots,K\}} \sigma_{h^k}(a^k) \cdot \sum\{\sigma_{-H^\rho}(s_{-H^\rho}) \mid s_{-H^\rho} \in S_{-H^\rho}\}$. Of course, the marginal distribution σ_{-H^ρ} satisfies $\{\sigma_{-H^\rho}(s_{-H^\rho}) \mid s_{-H^\rho} \in S_{-H^\rho}\} = \sigma_{-H^\rho}(S_{-H^\rho}) = 1$. Recalling that $\sigma_h = \pi_h$, we conclude that $\sum_{s \in T^\rho} \sigma(s) = \prod_{k \in \{1,2,\ldots,K\}} \pi_{h^k}(a^k)$. ■

A.2 One-deviation property

For a given game, consider a situation $h \in H_i$ where player i is on the move. Let J^h be the subset of H_i consisting of h and all of its successors. Then the *continuation strategy* from h for player i that appraisal p^h prescribes is the marginal distribution on J^h, that is $p^h_{J^h}$.

Definition: *For a given situation $h \in H_+$ and two appraisals p^h and $\hat{p}^h$, say that* ***$\hat{p}^h$ is a continuation-strategy h-deviation from p^h*** *if $p^h_{-J^h} = \hat{p}^h_{-J^h}$ (they are identical on $H \setminus J^h$).*

In other words, for player i's situation h, a continuation-strategy h-deviation alters only the actions played at player i's situations in the continuation of the game (that is, at h and successor situations for player i).

Definition: *Call an appraisal system P* ***continuation-strategy sequentially rational*** *if for every player $i \in N_+$, every $h \in H_i$, and every continuation-strategy h-deviation $\hat{p}^h$, it is the case that $u_i(p^h) \geq u_i(\hat{p}^h)$.*

Theorem 5: [One-Deviation Property] *Given any minimally consistent appraisal system P for a finite game, P is sequentially rational if and only if it is continuation-strategy sequentially rational.*

Theorem 5 is essentially a mixed-strategy, n-player version of one direction of Perea's (2000) Theorem 3.1 (see also Hendon, Jacobsen, and Sloth 1996). I include a proof here for completeness and because of the slight generalization and variation on the logical method.

Proof: Because h-deviations are a subset of continuation-strategy h-deviations, continuation-strategy sequential rationality implies sequential rationality. To obtain the opposite implication, presume that appraisal system P is minimally consistent and sequentially rational, and yet there is a player $i \in N_+$, a situation $h' \in H_i$, and continuation-strategy h'-deviation $\hat{p}^{h'}$, such that $u_i(p^{h'}) < u_i(\hat{p}^{h'})$. We will find a contradiction.

The set of situations $\underline{H}_i$ can be partitioned into sets $(H_i^\beta)_{\beta=0}^{\overline{b}}$ for some integer $\overline{b}$, where $H_i^0 = \{\underline{h}_i\}$ and for each $\beta \in \{1, 2, \ldots, \overline{b}\}$ and every $h \in H_i^\beta$ there exists a situation in $H_i^{\beta-1}$ that is an immediate predecessor of h. Let b be the largest integer for which there exists $h \in H_i^b$ satisfying $\hat{p}_h^{h'} \neq p_h^{h'}$.

Consider every situation $h \in H_i^b$. By presumption, P is rational at h, which means that $u_i(p^h) \geq u_i(\hat{p}^h)$ for every h-deviation $\hat{p}^h$. Let appraisal q be defined so that

$$q_{H_i^b} = p_{H_i^b}^{h'} \text{ and } q_{-H_i^b} = \hat{p}_{-H_i^b}^{h'}.$$

That is, q is a continuation-strategy h'-deviation from $p^{h'}$ that deviates according to $\hat{p}^{h'}$ except for the last stage at which $\hat{p}^{h'}$ differs from $p^{h'}$. This is well defined because $p^{h'}$ and $\hat{p}^{h'}$, as appraisals for player i, satisfy h-independence on S, for every $h \in H_i$.

I claim that $u_i(\hat{p}^{h'}) \leq u_i(q)$. To see why this is the case, let us compare these appraisals. First note that the sets $S(h)$, for $h \in H_i^b$, are disjoint. Let

$$\hat{H}_i^b \equiv \{h \in H_i^b \mid \hat{p}^{h'}(S(h)) > 0\}$$

and let $Y \equiv \{S(h) \mid h \in \hat{H}_i^b\}$. By the definition of conditional probability, there must exist probability distributions $g \in \Delta \hat{H}_i^b$ and $r \in \Delta S$ such that $r(Y) = 0$ and

$$\hat{p}^{h'} = \hat{p}^{h'}(Y) \sum_{h \in \hat{H}_i^b} \hat{p}^{h'}(\,\cdot\,|S(h))g(h) + \left(1 - \hat{p}^{h'}(Y)\right) r. \tag{4}$$

Likewise, because q differs from $\hat{p}^{h'}$ only regarding the play at situations in H_i^b, we have

$$q = \hat{p}^{h'}(Y) \sum_{h \in \hat{H}_i^b} q(\,\cdot\,|S(h))g(h) + \left(1 - \hat{p}^{h'}(Y)\right) r, \tag{5}$$

for the same functions g and r.

We know that $S(h)$ is an H_i-product set for every $h \in H_i^b$. Because q and $\hat{p}^{h'}$ are appraisals for player i, they exhibit H_i-independence on S, which further implies that $\hat{p}^{h'}(\,\cdot\,|S(h))$ and $q(\,\cdot\,|S(h))$ exhibit H_i-independence on S. Further, we have $\hat{p}^{h'}(\,\cdot\,|S(h))_{-i} = q(\,\cdot\,|S(h))_{-i}$ by construction of q.

Applying the minimal-consistency condition in iterated fashion from h' to the situations in $\hat{H}_i^b$, and using the law of iterated expectation, we have

$$\hat{p}^{h'}(\,\cdot\,|S(h))_{-i} = \hat{p}^{h'}(\,\cdot\,|S(h)_{-i} \times S_i)_{-i} = p^h_{-i}$$

for every $h \in \hat{H}_i^b$. This implies that $q(\,\cdot\,|S(h))$ and p^h can differ only on situations in H_i that cannot be reached conditional on h having been reached, which implies that $u_i(q(\,\cdot\,|S(h))) = u_i(p^h)$. Although $\hat{p}^{h'}(\,\cdot\,|S(h))$ may not be an h-deviation from p^h, there exists an h-deviation that gives the same expected payoff. Because P is sequentially rational, we conclude that $u_i(q(\,\cdot\,|S(h))) \geq u_i(\hat{p}^{h'}(\,\cdot\,|S(h)))$. From equations 4 and 5, we obtain $u_i(q) \geq u_i(\hat{p}^{h'})$.

To recap, relative to $p^{h'}$, we found a continuation-strategy h'-deviation q that delivers a weakly higher expected payoff to player i as does $\hat{p}^{h'}$. Further, although $\hat{p}^{h'}$ differs from $p^{h'}$ at situations in stage b of the game, q differs from $p^{h'}$ through at most stage $b-1$. Replacing $\hat{p}^{h'}$ with q and repeating the argument, we eventually have $q = p^{h'}$, which contradicts that there is a continuation-strategy h'-deviation that strictly increases player i's expected payoff. ■

A.3 Implications of plain consistency

Results presented in Section 4.2 are restated here along with proofs.

Lemma 1: An appraisal system is m-plainly consistent if and only if for every $L \subset H$ and every L-conjunctive pair of consecutive situations $h, h' \in \underline{H}$ for which p^h exhibits L-independence on $S(h')_L \times S(h)_{-L}$, it is the case that $p^{h'}$ exhibits L-independence on S, and $p^{h'}_L = p^h_L(\,\cdot\,|S(h')_L)$.

Proof: The full implications of m-plain consistency are obtained by plugging $Y = S(h')_L \times S(h)_{-L}$ into the equation defining plain consistency, which then states that

$$(p^{h'}\,|S(h')_L \times \{s_{-L}\})_L = (p^h\,|S(h')_L \times S(h)_L)_L$$

for every $s_{-L} \in S(h)_{-L}$ for which $S(h')_L \times \{s_{-L}\} \subset S(h')$ and $p^{h'}(S(h')_L \times \{s_{-L}\}) > 0$. Because $S(h')$ is an L-product set, $s_{-L} \in S(h)_{-L}$ implies $S(h')_L \times \{s_{-L}\} \subset S(h')$. Therefore $(p^{h'}\,|S(h')_L \times \{s_{-L}\})_L$ is the same distribution for every $s_{-L} \in S(h')_{-L}$, proving L-independence for $p^{h'}$, and it is identical to $(p^{h'}\,|S(h')_L \times S(h')_{-L}\})_L$, which is $p^{h'}_L$ because the support is contained in $S(h')$. Note also that $(p^h\,|S(h')_L \times S(h)_{-L}\})_L = p^h_L(\,\cdot\,|S(h')_L)$, which gives the claimed expression. Reversing the steps proves the opposite relation. ■

Lemma 2: If an appraisal system is m-plain consistent then it is minimally consistent.

Proof: This is because, for any player $i \in N_+$ and pair of consecutive situations $h, h' \in \underline{H}_i$, the following are true: First, letting $\overline{L}(h')$ denote the set of situations consisting of h' and its successors for player i, we have that $S(h)$ and $S(h')$ are $\overline{L}(h')$-product sets satisfying $S(h)_{\overline{L}(h')} = S(h')_{\overline{L}(h')}$, and moderate plain consistency implies $p^{h'}_{\overline{L}(h)} = p^{h}_{\overline{L}(h)}$. Second, $S(h)$ and $S(h')$ are H_i-product sets due to perfect recall, and setting $L = -H_i$ gives the updating condition of minimal consistency (which references $-H_i$ as $-i$). ■

Lemma 3: Given a game, consider any subgame and let J be the set of situations in the subgame. Then every $h \in \underline{H}$ is J-conjunctive.

Proof: The initial node of the subgame is a singleton situation, denoted $\hat{h}$. Let M be the set of situations for all players on the unique path to $\hat{h}$. For each $h \in M$, let $\hat{a}_h$ be the unique action taken at situation h on the path to $\hat{h}$, and write $\hat{a}_M = (\hat{a}_h)_{h \in M}$. Let J be the set of situations in the subgame, and let $M' \equiv H \setminus (J \cup M)$.

For every strategic player i, we have $S(\underline{h}_i) = S = S_J \times S_{-J}$ and so trivially $\underline{h}_i$ is J-conjunctive. For every $h \in M \cup M'$, no situation in J is on any path to h, and therefore $S(h) = S_J \times S(h)_{-J}$. For every $h \in J$, the player on the move has observed $\hat{a}_M$ and there is no situation in M' is on any path to h, and therefore $S(h) = S(h)_J \times \{\hat{a}_M\} \times S_{M'}$. ■

Theorem 1: If P is an m-plain perfect Bayesian equilibrium of a finite game, then the equilibrium strategy profile σ is a subgame-perfect Nash equilibrium.

Proof: Let P be an m-plain perfect Bayesian equilibrium for a given finite game and let σ denote the equilibrium strategy profile. Consider any subgame containing a situation for any given strategic player i, and let L, $\hat{h}$, M, M' $\hat{a}_h$ for $h \in M$, and $\hat{a}_M$ be defined as in the proof of Lemma 3. We must show that σ_L is a Nash equilibrium in the subgame—that is, for every strategic player i, σ_{L_i} is a best response to $\sigma_{L_{-i}}$.

If $\hat{h} = \underline{h}_j$ for some strategic player j, then the subgame is the entire game, player i's strategy in the subgame is σ_i, and we know from rationality at $\hat{h}$ and the one-deviation property that σ_i is a best response to $p^{\hat{h}}_{-i} = \sigma_{-i}$. Otherwise, there is a unique sequence of situations $\{h^1, h^2, \ldots, h^K\} \subset H_i \setminus L_i$, such that:

- $h^1 = \underline{h}_i$;
- for each $k = 1, 2, \ldots, K-1$, h^k and h^{k+1} are consecutive situations for player i; and
- there is a situation $h \in L_i$ such that h^K and h are consecutive situations for player i.

Let $\underline{L}_i = \{h \in L_i \mid h^K$ and h that are consecutive situations for player $i.\}$.

To evaluate payoffs in the subgame, let $\hat{\sigma}$ be the behavior strategy constructed as follows. For each situation h on the path to $\hat{h}$, $\hat{\sigma}_h$ assigns probability one to action $\hat{a}_h$; otherwise, $\hat{\sigma}_h = \sigma_h$. Thus, $\hat{\sigma}_L = \sigma_L$, and $\hat{\sigma}$ reaches the subgame with probability one. Player i's payoff in the subgame (that is, conditional on the subgame being reached) under strategy profile σ is $u_i(\hat{\sigma})$.

Presume that σ_{L_i} is not a best response to $\sigma_{L_{-i}}$ in the subgame, and we will find a contradiction. Under the presumption, there exists a behavior strategy $\hat{\sigma}' \in \Delta S$ satisfying

$\hat{\sigma}'_{-L_i} = \hat{\sigma}_{-L_i}$ and $u_i(\hat{\sigma}') > u_i(\hat{\sigma})$. In words, $\hat{\sigma}'$ represents player i unilaterally deviating from $\hat{\sigma}$ in the sugbame, to her strict benefit.

Because $\underline{L}_i$ is the set of situations in which player i first acts in the subgame, there must be a situation $\tilde{h} \in \underline{L}_i$ such that $\hat{\sigma}(S(\tilde{h})) = \hat{\sigma}'(S(\tilde{h})) > 0$ and player i's continuation value from $\tilde{h}$ under $\hat{\sigma}'$ is strictly higher than under $\hat{\sigma}$. To write these continuation values, define probability distributions $\tilde{\sigma}$ and $\tilde{\sigma}'$ with support $S(\tilde{h})$ by $\tilde{\sigma} = \hat{\sigma}(\cdot \mid S(\tilde{h}))$ and $\tilde{\sigma}' = \hat{\sigma}'(\cdot \mid S(\tilde{h}))$. Player i's continuation value from $\tilde{h}$ under $\hat{\sigma}'$ is given by $u_i(\tilde{\sigma}')$, her continuation value from $\tilde{h}$ under $\hat{\sigma}$ is given by $u_i(\tilde{\sigma})$, and we have $u_i(\tilde{\sigma}') > u_i(\tilde{\sigma})$. By construction, $\tilde{\sigma}$ and $\tilde{\sigma}'$ are appraisals at $\tilde{h}$, and $\tilde{\sigma}'$ is a continuation-strategy $\tilde{h}$-deviation from $\tilde{\sigma}$.

Let us establish the relation between $p^{\tilde{h}}$ and $\tilde{\sigma}$. First we show that $\hat{\sigma}_L = p_L^{h^K}$. This follows from m-plain consistency using induction, Lemma 1, and Lemma 3. Start with $k = 1$, recall from Lemma 3 that h^1 and h^2 are L-conjunctive, and note that p^{h^1} exhibits L-independence on $S(h^2)_L \times S(h^1)_{-L} = S_L \times S(h^1)_{-L}$. Lemma 1 implies that $p_L^{h^2} = p_L^{h^1}(\cdot \mid S_L) = p_L^{h^1} = \sigma_L = \hat{\sigma}_L$ and p^{h^2} exhibits L-independence. Repeating the argument for the pairs (h^2, h^3), (h^3, h^4), and so on yields $p_L^{h^K} = \sigma_L = \hat{\sigma}_L$ and p^{h^K} exhibits L-independence.

Because h^K and $\tilde{h}$ are L-conjunctive, we can apply m-plain consistency once more, for h^K and $\tilde{h}$ are consecutive situations for player i and p^{h^K} exhibits L-independence. Lemma 1 implies that $p_L^{\tilde{h}} = p_L^{h^K}(\,\cdot\,|S(\tilde{h})_L) = \hat{\sigma}_L(\,\cdot\,|S(\tilde{h})_L)$ and $p^{\tilde{h}}$ exhibits L-independence. Because $\hat{\sigma}$ is a behavior strategy, it exhibits L-independence, and thus $\hat{\sigma}_L(\,\cdot\,|S(\tilde{h})_L) = \hat{\sigma}_L(\,\cdot\,|S(\tilde{h})) = \tilde{\sigma}_L$.

We have shown that $p^{\tilde{h}}$ and $\tilde{\sigma}$ are both appraisals at $\tilde{h}$ and satisfy $p_L^{\tilde{h}} = \tilde{\sigma}_L$. Define distribution $\overline{\sigma} \in \Delta S$ to exhibit L-independence, with $\overline{\sigma}_L \equiv \tilde{\sigma}_L$ and $\overline{\sigma}_{-L} \equiv p_{-L}^{\tilde{h}}$. Then $\overline{\sigma}$ is a continuation-strategy $\tilde{h}$-deviation from $p^{\tilde{h}}$. Because payoffs in the subgame depend only on s_L, we have $u_i(\overline{\sigma}) = u_i(\tilde{\sigma}') > u_i(\tilde{\sigma}) = u_i(p^{\tilde{h}})$. Thus, P is not continuation-strategy sequentially rational. Using Lemma 2 and Theorem 5, we conclude that P is not sequentially rational, which contradicts that P is an m-plain perfect Bayesian equilibrium. ■

Lemma 4: Take as given a multi-stage game with observed actions and independent types (as specified above). Then every $h \in \underline{H}_+$ is L^i-conjunctive.

Proof: Note that a situation in any stage K for any strategic player i is a personal history combining θ_i with the publicly observed sequence of actions profiles $(a^1, a^2, \ldots, a^{K-1})$ through the previous stage (the sequence being null if K is the first stage). Player i's strategy specifies the action to be taken in each such situation. Consider any situation h that entails a given sequence of action profiles $(a^1, a^2, \ldots, a^{K-1})$ for some K. The set $S(h)_{L^i}$ depends only on $(a^1, a^2, \ldots, a^{K-1})$. To be precise, $\hat{\theta}_i \hat{s}_i \in S(h)_{L^i}$ if and only if, for every $k = 1, \ldots, K-1$, it is the case that $\hat{s}_{h^k} = a^k$, where h^k is defined as the personal history combining $\hat{\theta}_i$ with action sequence $(a^1, a^2, \ldots, a^{k-1})$. Thus, for any strategy profile $s = \theta_1 \cdots \theta_n s_1 \cdots s_n$ such that $\theta_i s_i \in S(h)_{L^i}$ for every $i \in N_+$, we have $s \in S(h)$. Hence h is L^i-conjunctive for every $i \in N_+$. ■

Theorem 2: Suppose that P is a perfect Bayesian equilibrium of a finite multi-stage game with observed actions and independent types (as specified above), under cps/m-plain consistency. Then there is a sequential equilibrium that has exactly the same outcome distribution.

Proof: For any given finite multi-stage game with observed actions and independent types, let B be the set of feasible partial sequences of publicly observed action profiles from stage 1 through any stage prior to T. That is, each $\beta \in B$ can be written as $\beta = (a^1, a^2, \ldots, a^t)$, where $t < T$ is the length of the sequence and a^τ is the action profile played in stage τ. Assume B includes the null sequence at the beginning of the game.

For any $\beta \in B$, let $Z(\beta)$ denote the subset of S that would result in β being played through the stage equal to the length of β. Further, for any player $i \in N_+$ and type θ_i, let $Z^i(\beta, \theta_i)$ denote the subset of $Z(\beta)$ in which $s_{0i} = \theta_i$. Note that (β, θ_i) is a personal history for player i and therefore is exactly that situation $h \in H_+$ for which $S(h) = Z^i(\beta, \theta_i)$.

Take as given a perfect Bayesian equilibrium P under cps/m-plain consistency. Recall Lemma 4 showing that every situation in $\underline{H}_+$ is L^j-conjunctive for every $j \in N_+$. For any strategic player i and consecutive situations $h, h' \in \underline{H}_i$, if $\zeta(S(h')_{L^j} \times S(h)_{-L^j} \mid S(h)) = 0$ for at most one $j \in N_+$, then from cps/m-plain consistency we have

$$\zeta(X_{L^j} \times S(h')_{-L^j} \mid Y_{L^j} \times S(h')_{-L^j}) = \zeta(X_{L^j} \times S(h)_{-L^j} \mid Y_{L^j} \times S(h)_{-L^j})$$

for every $j \in N_+$ and $X_{L^j} \subset Y_{L^j} \subset S(h')_{L^j}$, and further $\zeta(\cdot \mid S(h'))$ exhibits L^j-independence for every $j \in N_+$.

Note that the condition $\zeta(S(h')_{L^j} \times S(h)_{-L^j} \mid S(h)) = 0$ means that at h' player i observed a deviation from $\zeta(\cdot \mid S(h))$ by player j in the previous stage of the game. If $j \neq i$ then this was a public deviation, meaning that player j took an action in the previous stage that player i did not expect, given player i's appraisal at h. (No type of player j that player i assigned positive probability would choose this action, given player i belief about player j's strategy.) If $j = i$ then the deviation was by player i herself, and it may have been public from the other players' perspective or private (whereby one type deviates to take the action that another type was supposed to take, where both types were assigned positive probability by the other players).

Clearly, if $j \neq i$ and $\zeta(S(h')_{L^j} \times S(h)_{-L^j} \mid S(h)) > 0$, so that player i did not observe a public deviation by player j, then by conditional-probability updating, player i believes that no private deviation was taken by player j in the previous stage.

Consider any situation $\hat{h}$ of player i in some stage t of the game, and let $(h^0, h^1, h^2, \ldots, h^t)$ be the sequence of consecutive situations for player i from the beginning of the game to $\hat{h}$, where $h^0 = \underline{h}_i$ and $h^t = \hat{h}$. Let us say that $\hat{h}$ *entails at most one deviation per stage* if for every $\tau \in \{1, 2, \ldots, t\}$, $\zeta(S(h^\tau)_{L^j} \times S(h^{\tau-1})_{-L^j} \mid S(h^{\tau-1})) = 0$ for at most one player $j \in N_+$. For such a situation $\hat{h}$, iterative application of cps/m-plain consistency implies that

$$\zeta(X_{L^j} \times S(\hat{h})_{-L^j} \mid Y_{L^j} \times S(\hat{h})_{-L^j}) = \zeta(X_{L^j} \times S_{-L^j} \mid Y_{L^j} \times S_{-L^j})$$

for every $j \in N_+$ and $X_{L^j} \subset Y_{L^j} \subset S(\hat{h})_{L^j}$, and $\zeta(\cdot \mid S(\hat{h}))$ exhibits L^j-independence for every $j \in N_+$. Setting $Y_{L^j} = S(\hat{h})_{L^j}$ yields the following relation of marginal distributions:

$$\zeta(\cdot \mid S(\hat{h}))_{L^j} = \zeta(\cdot \mid S(\hat{h})_{L^j} \times S_{-L^j})_{L^j}.$$

Using the same steps as in the proof of Lemma 5, we then obtain:

$$\zeta(s \,|\, S(\hat{h})) = \prod_{j \in N_+} (\zeta \,|\, S(\hat{h}))_{L^j}(s_{L^j}) = \prod_{j \in N_+} \zeta(\{s_{L^j}\} \times S_{-L^j} \,|\, S(\hat{h})_{L^j} \times S_{-L^j})$$

for every $s \in S$.

Let J denote the subset of $\underline{H}_+$ that entail at most one deviation per stage. Let P and σ^* denote the given perfect Bayesian equilibrium and equilibrium strategy profile. Because $\sigma^* = p^{\underline{h}_i}$ for every $i \in N_+$, we have $\sigma^* = \zeta(\cdot \,|\, S)$. Since ζ is a conditional probability system, Myerson's (1986) Theorem 1 guarantees that there is a sequence of fully mixed strategy profiles $\{\hat{\sigma}^k\}$ such that $\hat{\sigma}^k \to \sigma^*$, $\hat{\sigma}^k$ exhibits L^j-independence for every $j \in N_+$, and $p^h(s) = \zeta(s \,|\, S(h)) = \lim_{k\to\infty} \hat{\sigma}^k(s)/\hat{\sigma}^k(S(h))$ for every $h \in J$ and $s \in S(h)$. The first equality is from the definition of P deriving from ζ. The second equality is achieved by using Myerson's theorem separately on each dimension L^j, for the conditional probability system defined on S_{L^j} given by $\zeta(\cdot \times S_{-L^j} \mid \cdot \times S_{-L^j})$, to construct a sequence $\{\hat{\sigma}^k_{L^j}\}$, and taking the product over $j \in N_+$.

Because player j observes $\theta_0 j$ before taking actions, we can use Kuhn's Theorem regarding behavior strategies to replace $\hat{\sigma}^k_{L^j}$ with an equivalent behavior strategy profile $\sigma^k_{L^j}$, and take the product to get a behavior strategy profile σ^k that is equivalent to $\hat{\sigma}^k$ in the sense of giving the same distribution over terminal nodes, including in the limit and for any single deviation (the same h-deviation applied to both σ^k and $\hat{\sigma}^k$). In fact, because σ^*_0 has full support, without loss we can assume that $\sigma^k_0 = \sigma^*_0$ for all k. (To verify, examine the ratio of the constructed $\sigma^k(s)$ and the modified version with $\sigma_0(s) = \sigma^*_0$ and note that the latter is bounded away from zero.) We can also assume that $\sigma^k(s)/\sigma^k(S(h))$ converges for every $h \in \underline{H}_+$ and $s \in S(h)$ because there must be a subsequence with this property (due to J and S being finite).

Next let us define an appraisal system $\tilde{P}$ by $\tilde{p}^h(s) \equiv \sigma^k(s)/\sigma^k(S(h))$ for each $h \in \underline{H}_+$ and every $s \in S(h)$. By construction, $\tilde{P}$ inherits the following properties of P, by construction. First, for every $h \in J$, p^h exhibits L^j-independence for all $j \in N_+$. Second, $\tilde{p}^h$ is rational at h, for every $h \in J$. Further, the conditions for rationality at $h \in J$ do not depend on $\tilde{p}^h_{-J}$ because such dependence would require both the player on the move at h and another player to simultaneously deviate at h. This means that any appraisal system $\hat{P}$ satisfying $\hat{p}^h_J = \tilde{p}^h_J$ for every $h \in J$ has the property that $\hat{p}^h$ is rational at h, for every $h \in J$. And by construction, $\hat{p}^{\underline{h}_i}$ induces the same outcome distribution as does σ^*.

Third, every $h \in J \cap \underline{H}_i$ of player i entails at most one deviation per stage. That is, for any consecutive situations $h^{\tau-1}$ and h^τ in this sequence, $\tilde{p}^{h^{\tau-1}}(S(h^\tau)_{L^j} \times S(h^\tau)_{-L^j} \,|\, S(h^{\tau-1})) = 0$ for no more than one $j \in N_+$. Further, if $j \neq i$ and $\tilde{p}^{h^{\tau-1}}(S(h^\tau)_{L^j} \times S(h^\tau)_{-L^j} \,|\, S(h^{\tau-1})) > 0$, so that player i did not observe a public deviation by player j in the stage of h^τ, then player i believes that no private deviation was taken by player j in the previous stage. This means that at h, player i thinks that all situations reached so far in the game, including for the other players, are all in J. An implication is that, for any $s \in S$ for which $s_J \in \operatorname{supp} \tilde{p}^h_J$, it is the case that $s \in S(h)$. This implies that $\{s \in S(h) | s_J \in \operatorname{supp} \tilde{p}^h_J\}$ is a J-product set.

Let us define a partial appraisal $Q = \left(q^h\right)_{h \in J}$ by $q^h = \tilde{p}^h_J$ for every $h \in J$, and note that these probability distributions are over S_J. By construction, Q satisfies Watson's (2023) definition of "J-partial sequential equilibrium." The "rectangular margin-support condition" is also satisfied because we have shown that $\{s \in S(h) | s_J \in \operatorname{supp} \tilde{p}^h_J\}$ is a J-product set for every $h \in J$. Watson's (2023) Theorem then establishes the existence of the claimed sequential equilibrium. ■

A.4 Implications for sender/designer games

Below is a restatement of the result presented in Section 5 and a proof, followed by more details on the equilibrium existence example sketched in the main text.

Theorem 3: Consider any amenable sender/designer game and appraisal system P, and let $Y_0(\lambda_i, s_\Phi)$ be defined by Equation (2). Then P is pc-compliant only if, for every player $i \in N_+$ and $\lambda_i \in \Lambda_i$,

(a) p^{λ_i} exhibits M-independence on S and satisfies $p^{\lambda_i}_M = p^{h_i}_M$, and

(b) $p^{\lambda_i}_0 = \displaystyle\int_{S_\Phi} p^{h_i}_0(\,\cdot\,|Y_0(\lambda_i, s_\Phi))\, p^{\lambda_i}_\Phi(ds_\Phi)$.

Likewise, the converse holds (P satisfying conditions (a) and (b) implies pc-compliance) for appropriately defined M- and $-\Phi$-disintegration kernels.

Proof: First I show that pc-compliance implies properties (a) and (b) stated in the theorem. Consider any player $i \in N_+$ and $\lambda_i \in \Lambda_i$. Recall that $S(\lambda_i) = S_M \times S_{-M}(\lambda_i)$. The update from $\underline{h}_i$ to λ_i being M-plainly consistent on S implies that $\kappa_M(\,\cdot, s_{-M} \,|\, S_M) = (p^{\underline{h}_i} \,|\, S)_M$ for every $s_{-M} \in S_{-M}$ for which $S_M \times \{s_{-M}\} \subset S(\lambda_i)$ and $\kappa_M(S_M, s_{-M}) > 0$, where κ_M is the M-disintegration probability kernel associated with p^{λ_i}. Since $\kappa_M(S_M, s_{-M}) = 1$ and $(p^{\underline{h}_i} \,|\, S)_M = p^{h_i}_M$, using Equation 1 we obtain $\kappa_M(\,\cdot\,, s_{-M}) = p^{h_i}_M$ for every $s_{-M} \in S_{-M}(\lambda_i)$. Therefore, from the definition of disintegration probability kernel, p^{λ_i} exhibits M-independence on S and satisfies $p^{\lambda_i}_M = p^{h_i}_M$, and hence condition (a) holds.

Next consider any player $i \in N_+$ and $\lambda_i \in \Lambda_i$, and let $\kappa_{-\Phi}$ be the $-\Phi$-disintegration probability kernel associated with p^{λ_i}. We can assume that $\kappa_{-\Phi}(Y_0(\lambda_i, s_\Phi) \times S_M, s_\Phi) = 1$ for every $s_\Phi \in S_\Phi$, because $\kappa_{-\Phi}(\,\cdot\,, s_\Phi)$ is a probability measure, κ_Φ is uniquely defined up to a measure-zero set, and if this condition were violated over a strictly positive measure of s_Φ then it would contradict that $p^{\lambda_i} \in \Delta S(\lambda_i)$.

Take any $X_0 \subset S_0$. From Equation 1 and because $\kappa_{-\Phi}(Y_0(\lambda_i, s_\Phi) \times S_M, s_\Phi) = 1$, we have

$$\begin{aligned}\kappa_{-\Phi}\left(X_0 \times S_M, s_\Phi \,|\, Y_0(\lambda_i, s_\Phi) \times S_M\right) &= \frac{\kappa_{-\Phi}\left((X_0 \times S_M) \cap (Y_0(\lambda_i, s_\Phi) \times S_M), s_\Phi\right)}{\kappa_{-\Phi}\left(Y_0(\lambda_i, s_\Phi) \times S_M, s_\Phi\right)} \\ &= \kappa_{-\Phi}(X_0 \times S_M, s_\Phi).\end{aligned}$$

By $-\Phi$-plain consistency on $Y_0(\lambda_i, s_\Phi) \times S_{-0}$, we know that

$$\kappa_{-\Phi}\left(X_0 \times S_M, s_\Phi \,|\, Y_0(\lambda_i, s_\Phi) \times S_M\right) = (p^{\underline{h}_i} \,|\, Y_0(\lambda_i, s_\Phi) \times S_{-0})_{-\Phi}(X_0 \times S_M).$$

By definition of $\kappa_{-\Phi}$,

$$p^{\lambda_i}_{-\Phi}(X_0 \times S_M) = \int_{S_\Phi} \kappa_{-\Phi}(X_0 \times S_M, s_\Phi)\, p^{\lambda_i}_\Phi(ds_\Phi).$$

Combining the last three equations yields

$$p^{\lambda_i}_{-\Phi}(X_0 \times S_M) = \int_{S_\Phi} (p^{\underline{h}_i} \,|\, Y_0(\lambda_i, s_\Phi) \times S_{-0})_{-\Phi}(X_0 \times S_M)\, p^{\lambda_i}_{\Phi}(ds_\Phi).$$

Note that $(p^{\underline{h}_i} \,|\, Y_0(\lambda_i, s_\Phi) \times S_{-0})_{-\Phi}(X_0 \times S_M) = p_0^{\underline{h}_i}(X_0 \,|\, Y_0(\lambda_i, s_\Phi))$, the H_0-marginal probability of X_0 conditional on $Y_0(\lambda_i, s_\Phi)$. Likewise, $p^{\lambda_i}_{-\Phi}(X_0 \times S_M) = p_0^{\lambda_i}(X_0)$. With these substitutions, we obtain

$$p_0^{\lambda_i}(X_0) = \int_{S_\Phi} p_0^{\underline{h}_i}(X_0 \,|\, Y_0(\lambda_i, s_\Phi))\, p^{\lambda_i}_{\Phi}(ds_\Phi),$$

which is condition (b) in the theorem.

To prove the converse result, consider any appraisal system P that satisfies conditions (a) and (b) for every player $i \in N_+$ and $\lambda_i \in \Lambda_i$. If, for some $i \in N_+$ and $\lambda_i \in \Lambda_i$, the given disintegration kernel $\kappa_{-\Phi}$ for p^{λ_i} violates $-\Phi$-plain consistency on $Y_0(\lambda_i, s_\Phi) \times S_{-0}$ for a positive-measure subset of S_Φ, then condition (b) would fail, and thus $-\Phi$-plain consistency can be violated only on a measure-zero subset. We can then alter the definition of $\kappa_{-\Phi}$ on the measure-zero subset, to everywhere conform to $-\Phi$-plain consistency. A similar adjustment can be done for κ_M. ■

Next, I present a few more details on the example of equilibrium existence sketched in Section 5. Picking up with the second part of the proposed algorithm, fix any small number $\varepsilon > 0$ and let $\Delta^\varepsilon \hat{A}_1$ be the (compact) subset of $\Delta \hat{A}_1$ in which each action gets a probability of at least ε. Imagine that for each $a_1 \in \hat{A}_1$, we have specified a distribution $\mu^{a_1} \in \Delta\Theta$ that will be the posterior belief of players 2 through n about nature's move, contingent on observing action a_1 chosen by player 1. Let us look at an artificial game between player 1's types whereby they choose mixed actions in $\Delta^\varepsilon \hat{A}_1$ and receive payoffs given by w_1^*. We then have a correspondence $\rho^\varepsilon\colon (\Delta\Theta)^{|\hat{A}_1|} \rightrightarrows (\Delta^\varepsilon \hat{A}_1)^{|\Theta_1|}$ that gives, for each vector of posteriors $(\mu^{a_1})_{a_1 \in \hat{A}_1}$, the set of optimal mixed-action vectors for player 1's types. Let function $\eta^\varepsilon\colon (\Delta^\varepsilon \hat{A}_1)^{|\Theta_1|} \to (\Delta\Theta)^{|\hat{A}_1|}$ give, for every mixed-action vector for player 1's types, the implied posteriors for the actions in $\hat{A}_1$, which are uniquely determined by Bayes rule. Let $\eta^\varepsilon(a_1, \sigma_1)$ denote the posterior conditional on action a_1 when player 1's types choose mixed action vector σ_1.

The next step is to find a fixed point σ_1^ε of the composition mapping $\rho^\varepsilon \circ \eta^\varepsilon$, which exists because η^ε is continuous and ρ^ε is nonempty, convex-valued, and upper hemi-continuous owing to the continuity of w_1^*. We can then construct a sequence $(\sigma_1^{1/k})_k$ that converges and has the property that $\eta^{1/k}(a_1, \sigma_1^\varepsilon)$ converges for every $a_1 \in \hat{A}_1$. Such a sequence exists from Kakutani's theorem since $(\Delta^\varepsilon \hat{A}_1)^{|\Theta_1|}$ is a compact subset of $\mathbb{R}^{|\Theta_1|-1}$. The limit of $\sigma_1^{1/k}$ is player 1's equilibrium strategy and the limit of $\eta^{1/k}(a_1, \sigma_1^\varepsilon)$ is the posterior about nature's move that players 2 through n hold contingent on observing action $a_1 \in \hat{A}_1$. The posterior contingent on observing any of the other actions is specified to be $\underline{\mu}$, as discussed above.

As noted in the text, the key constraint imposed by plain consistency in the foregoing argument is that $\underline{\mu}$, the off-equilibrium-path posterior about nature's move, also has the form required by Theorem 3. Also recall that, for the construction to work, we need assumptions guaranteeing that the stage-three continuation game has an equilibrium, that there is a selection of such equilibria that makes w_1^* continuous in μ', and there is a distribution $\underline{\mu}$ that has the form required by Theorem 3 and satisfies Inequality 3 for every $\mu' \in \Delta\Theta$ that satisfies belief consistency, $a_1 \in A_1$, and $\theta_1 \in \Theta_1$. For example, a supermodular setting with the following assumptions would suffice: A_i is an interval and Θ_i is a finite subset of the reals, for every player i. Types are positively related in that, for each $M \subset N_+$ and $-M = N \backslash M$, $(\mu \,|\, \Theta_M \times \{\theta_{-M}\})_M$ is increasing in θ_{-M} in the multivariate sense of first-order stochastic dominance (see, for instance, Østerdal 2010). Function w is supermodular in all arguments (Topkis 1979, Milgrom and Roberts 1990), continuous, and strictly quasiconcave; and w_1 is increasing in θ_{-1} and a_{-1}. We can select the maximal equilibrium of the stage-three continuation game to define w_1^*, and $\underline{\mu}$ is defined as the posterior conditional on the lowest type of player 1.

A.5 Proof of other result

Theorem 4: As defined for finite games, consistency definitions are nested from strongest to weakest as follows: full, strong, f-plain, m-plain, and minimal.

Proof: Battigalli (1996) proves that full consistency implies strong consistency. To show that strong consistency implies f-plain consistency, consider any strongly consistent appraisal system P, and let ζ be a conditional-probability system that has the independence property, from which P is derived. Take any $L \subset H$ and L-product set $Y \subset S$, and consecutive situations h and h' for any strategic player, such that $Y \subset S(h)$ and p^h exhibits L-independence on Y. Take any $s_{-L} \in Y_{-L}$ for which $Y_L \times \{s_{-L}\} \subset S(h')$ and $p^{h'}(Y_L \times \{s_{-L}\}) > 0$. Using that P is derived from ζ and that ζ has the independence property, we have

$$(p^{h'} \,|\, Y_L \times \{s_{-L}\})_L = (\zeta \,|\, Y_L \times \{s_{-L}\})_L = (\zeta \,|\, Y_L \times S_{-L})_L$$

and

$$(p^h \,|\, Y_L \times Y_{-L})_L = (\zeta \,|\, Y_L \times Y_{-L})_L = (\zeta \,|\, Y_L \times S_{-L})_L,$$

which implies $(p^{h'} \,|\, Y_L \times \{s_{-L}\})_L = (p^h \,|\, Y_L \times Y_{-L})_L$, as required for plain consistency. The remaining inclusion relations are implied by the moderate and full plain consistency definitions, and by Lemma 2. ■

References

Au, P.H. and M. Whitmeyer (2023): "Attraction Versus Persuasion: Information Provision in Search Markets," *Journal of Political Economy* 131 (1): 202–245.

Baccelli, François, Bartlomiej Blaszczyszyn, and Mohamed Karray (2024): *Random Measures, Point Processes, and Stochastic Geometry*, Inria, 2024, hal-02460214v2.

Battigalli, P. (1996): "Strategic Independence and Perfect Bayesian Equilibria," *Journal of Economic Theory* 70: 201–234.

Battigalli, P. (1997): “On Rationalizability in Extensive Form Games,” *Journal of Economic Theory* 74: 40–61.

Battigalli, P. and J. Watson (1997): “On ‘Reputation’ Refinements with Heterogeneous Beliefs,” *Econometrica* 65: 369–374.

Bergemann, D., T. Heumann, and S. Morris (2023): “Screening with Persuasion,” MIT working paper, September 2023.

Blume, L., A. Brandenburger, and E. Dekel (1991a): “Lexicographic Probabilities and Choice Under Uncertainty,” *Econometrica* 59 (1): 61–79.

Blume, L., A. Brandenburger, and E. Dekel (1991b): “Lexicographic Probabilities and Equilibrium Refinements,” *Econometrica* 59 (1): 81–98.

Bonanno, G. (2013): “AGM-Consistency and Perfect Bayesian Equilibrium, Part I: Definition and Properties,” *International Journal of Game Theory* 42 (3): 567–592

Bowen, T.R., D. Dmitriev, and S. Galperti (2023): “Learning from Shared News: When Abundant Information Leads to Belief Polarization,” *The Quarterly Journal of Economics* 138 (2): 955–1000.

Bull, J. and J. Watson (2017): “Statistical Evidence and the Problem of Robust Litigation,” unpublished draft, UC San Diego.

Coletti, G. and R. Scozzafava (2002): “Probabilistic Logic in a Coherent Setting,” *Trends in Logic*, Vol. 15, Kluwer, Dordrecht.

Cozman, F.G. (2013): “Independence for Full Conditional Probabilities: Structure, Factorization, Non-Uniqueness, and Bayesian Networks,” *International Journal of Approximate Reasoning*, 54: 1261–1278.

De Clippel, G. and X. Zhang (2022): “Non-Bayesian Persuasion,” *Journal of Political Economy* 130 (10): 2594–2642.

Doval, L. and J. Ely (2020): “Bayes’ Rule Where Possible,” working paper.

Dworczak, P. A. Pavan (2022): “Preparing for the Worst but Hoping for the Best: Robust (Bayesian) Persuasion,” *Econometrica* 90 (5): 2017–2051.

Fudenberg, D. and J. Tirole (1991): “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory* 53: 236–260.

González-Díaz, J. and M. A. Meléndez-Jiménez (2014): “On the Notion of Perfect Bayesian Equilibrium,” *TOP Journal of the Spanish Society of Statistics and Operations Research* 22 (1): 128–143.

Govindan, Srihari and Robert Wilson (2009): “On Forward Induction,” *Econometrica* 77 (1, January): 1–28.

Hammond, P.J. (1994): “Elementary Non-Archimedean Representations of Probability for Decision Theory and Games,” pp. 25–59 in P. Humphreys (Ed.), *Patrick Suppes: Scientific Philosopher*, vol. 1, Kluwer, Dordrecht.

Hendon, E., H. J. Jacobsen, and B. Sloth (1996): “The One-Deviation Principle for Sequential Rationality,” *Games and Economic Behavior* 12: 274-282.

Kamenica, E. and M. Gentzkow (2011): “Bayesian Persuasion,” *American Economic Review* 101 (6): 2590–2615.

Kattwinkel, D. and J. Knoepfle (2023): “Costless Information and Costly Verification: A Case for Transparency,” *Journal of Political Economy* 131 (2): 504–548.

Krauss, P. (1968): “Representation of Conditional Probability Measures on Boolean Algebras,” Acta Mathematica Hungarica 19 (3-4): 229–241.

Krähmer, D. (2021): “Information Design and Strategic Communication,” *American Economic Review: Insights* 3 (1): 51–66.

Kohlberg, E. and P. Reny (1997): “Independence on Relative Probability Spaces and Consistent Assessments in Game Trees,” *Journal of Economic Theory* 75 (2): 280–313.

Koessler, F. and V. Skreta (2023): “Informed Information Design,” *Journal of Political Economy* 131 (11, November): 3187–3232.

Kreps, D. and G. Ramey (1987): “Structural Consistency, Consistency, and Sequential Rationality,” *Econometrica* 55 (6): 1331–1348.

Kreps, D. and R. Wilson (1982): “Sequential Equilibrium,” *Econometrica* 50 (4): 863–894.

Lipnowski, E., D. Ravid, and D. Shishkin (2022): “Persuasion via Weak Institutions,” *Journal of Political Economy* 130 (10): 2705–2730.

Mailath, G., (2016): Class Notes for Penn Economics 703: Microeconomics II Modelling Strategic Behavior, unpublished draft dated January 18, Department of Economics, University of Pennsylvania.

Mailath, G., L. Samuelson, and J. Swinkels (1993): “Extensive Form Reasoning in Normal Form Games,” *Econometrica* 61 (March): 273–302.

Mas-Colell, A., M. D. Whinston, and J. Green (1995): *Microeconomic Theory*, New York: Oxford University Press.

McLennan, A. (1989): “Consistent Conditional Systems in Noncooperative Game Theory,” *International Journal of Game Theory* 18: 141–174.

McClellan, A. (2022): “Experimentation and Approval Mechanisms,” *Econometrica* 90 (5): 2215–2247.

Milgrom, P. and J. Roberts (1990): “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities,” *Econometrica* 58: 1255–77.

Myerson, R.B. (1983): “Mechanism Design by an Informed Principal,” *Econometrica* 51 (6): 1767–1797.

Myerson, R. B. (1986): “Games with Communication,” *Econometrica* 54 (2): 323–358.

Myerson, R. B. (1991): *Game Theory: Analysis of Conflict*, Cambridge, Massachusetts: Harvard University Press.

Myerson, R. B. and P. J. Reny (2020): “Perfect Conditional ε-Equilibria of Multi-Stage Games With Infinite Sets of Signals and Actions,” *Econometrica* 88 (2): 495–531.

Østerdal, L.P (2010): “The Mass Transfer Approach to Multivariate Discrete First Order Stochastic Dominance: Direct Proof and Implications,” *Journal of Mathematical Economics* 46 (6): 1222–1228.

Pearce, D. G. (1984): “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica* 52 (July): 1029–1050.

Perea, A. (2002): "A Note on the One-Deviation Property in Extensive-Form Games," *Games and Economic Behavior* 40: 322–338.

Perego, J. and S. Yuksel (2022): "Media Competition and Social Disagreement," *Econometrica* 90 (1): 223–265.

Rényi, A. (1955): "On a New Axiomatic Theory of Probability," *Acta Mathematica Hungarica* 6: 285–335.

Selten, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory* 4 (1): 25–55.

Shimoji, M. and J. Watson (1998): "Conditional Dominance, Rationalizability, and Game Forms," *Journal of Economic Theory* 83: 161–195.

Topkis, D. (1979): "Equilibrium Points in Non-Zero Sum n-Person Submodular Games," *SIAM Journal of Control and Optimization* 17: 773–87.

Watson, J. (1993): "A 'Reputation' Refinement without Equilibrium," *Econometrica* 61: 199–205.

Watson, J. (1996): "Information Transmission When the Informed Party is Confused," *Games and Economic Behavior* 12 (1): 143–161.

Watson, J. (1998): "Alternating-Offer Bargaining with Two-Sided Incomplete Information," *Review of Economic Studies* 65: 573–594.

Watson, J. (2017): "A General, Practicable Definition of Perfect Bayesian Equilibrium," unpublished draft, UC San Diego.

Watson, J. (2023): "Partially Constructed Sequential Equilibrium," working paper (November 2023).

Watson, J. (2024): "Contractual Chains," *Econometrica* 92 (5): 1735–1774.

Weibull, J. (1992): "On Self-Enforcement in Extensive-Form Games," *Games and Economic Behavior* 4: 450–462.