A General, Practicable Definition of Perfect Bayesian Equilibrium

Joel Watson*

February 2017

Abstract

This paper develops a general definition of *perfect Bayesian equilibrium (PBE)* for extensive-form games. It is based on a new consistency condition for the players' beliefs, called *plain consistency*, that requires proper conditional-probability updating on independent dimensions of the strategy space. The condition is familiar and convenient for applications because it constrains only how a player's belief is updated on consecutive information sets. The PBE concept is defined for infinite games, implies subgame perfection, and captures the notion of "no signaling what you don't know." 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

Standard solution concepts for dynamic games are based on the notion of *sequential rationality*, which requires players to maximize their expected payoffs not just in the ex ante sense (strategy selection before the game is played) but at all contingencies within the game where they are called upon to take actions. *Trembling-hand perfect equilibrium* (Selten 1975) and *sequential equilibrium* (Kreps and Wilson 1982) ensure that the rationality test is applied to all information sets in an extensive-form game, because these concepts are defined relative to convergent sequences of fully mixed behavior strategies.

Trembling-hand perfect equilibrium and sequential equilibrium aren't always the best choice for applications, for the following reasons. First, constructing sequences of fully mixed strategies with the desired properties can be difficult in complex games. Second, for some applications, more permissive concepts—allowing for a greater range of beliefs at information sets—may be desired. Third, while many applications are conveniently formulated with infinite action spaces, trembling-hand perfect equilibrium and sequential equilibrium are not

*UC San Diego; http://econ.ucsd.edu/∼jwatson/. The author thanks the NSF for financial support (SES-1227527) and the following people for their insightful input: Nageeb Ali, Pierpaolo Battigalli, Jesse Bull, Jack Fanning, Simone Galperti, Keri Hu, David Miller, Phil Reny, Karl Schlag, Joel Sobel, and two anonymous referees. This paper owes a great deal to Pierpaolo Battigalli and his work identifying independence conditions central to major solution concepts.

defined for such games.[1] All three of these factors are relevant for many modern applications, including some models of private contracting on networks, dynamic contracting, sequential evidence production, and repeated games on networks.

Practitioners therefore often turn to the *perfect Bayesian equilibrium* (PBE) concept, which is usually described in the same way as is sequential equilibrium—a behavior strategy profile and a system of *assessments* that give the players' beliefs at information sets as probability distributions over nodes—but puts structure on the assessments with consistency conditions that are formulated without reference to strategy trembles. At the heart of PBE is the idea that the players' beliefs should be consistent with proper conditional probability updating (Bayes' rule) where applicable. To state a formal definition of PBE, one must make "where applicable" precise.

Fudenberg and Tirole (1991) provide the leading formal definition of PBE in the literature, but this definition applies only to the class of "finite multi-period games with observed actions and independent types" whereas applications of PBE are increasingly outside this class. A more general definition of PBE based on Battigalli's (1996) *independence property* for conditional probability systems has not been utilized in applications because it lacks a practicable formulation and because verifying the independence condition appears tantamount to finding a suitable sequence of fully mixed behavior strategies in a sequential-equilibrium construction.[2] Further, an infinite-game extension has not been worked out.

Although applications of "perfect Bayesian equilibrium" are widespread in the literature, a measure of ambiguity persists regarding the technical conditions that practitioners are actually utilizing in individual modeling exercises. In some articles, PBE is the stated solution concept but there is no reference to a formal definition. Thus, for other than finite multi-period games with observed actions and independent types, it is not always clear what researchers have in mind. There is a range of possible consistency assumptions, and the assumptions matter in applications. Questions linger about what "Bayes' rule where applicable" should mean.

For the analysis of complex games, researchers sometimes retreat to the concept of *weak PBE* because of its simple structure and flexibility, despite that it was advanced as a pedagogical stepping stone.[3] Weak PBE imposes no constraints on beliefs off the equilibrium path. It does not imply subgame perfection.

This paper endeavors to support wider application of PBE by providing a general definition of perfect Bayesian equilibrium that meets several goals. First, it constrains only how individual players update beliefs on *consecutive information* sets—that is, from one information set to the next one that arises for the same player—thus lending itself to straightforward application in a way familiar to practitioners. Second, it applies to all finite games as well as to infinite games with the appropriate measurability structure. Third, it is a refinement

[1]Myerson and Reny (2015) discuss technical problems with extending sequential equilibrium to infinite games and propose a new concept called *open sequential equilibrium*.

[2]On the unwieldy point, working toward a PBE definition using Fudenberg and Tirole's (1991) and Battigalli's (1996) framework 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 conjecture system over terminal nodes, derive from it the assessments at information sets, and verify sequential rationality.

[3]See Myerson (1991) and Mas-Colell, Whinston, and Green (1995). Some stronger definitions that imply subgame perfection are discussed in Section 5.

of subgame perfection. And fourth, it captures the notion of "no signaling what you don't know," implying Fudenberg and Tirole's (1991) *reasonableness* condition for multi-period games with observed actions and independent types.[4]

The PBE definition proposed herein is based on a new consistency notion called *plain consistency*, which emulates some of the structure on beliefs inherent in sequential equilibrium. Central to the definitions of trembling-hand perfect equilibrium and sequential equilibrium is the assumption that choices made at different information sets are independent of one another, as represented by behavior strategies. This assumption puts structure on the players' beliefs, including at off-path information sets. Likewise, plain consistency imposes some independence in the operation of conditional-probability updating. For ease of use, the condition is limited to updating on consecutive information sets, but it is also strong enough to deliver the other desired properties stated above.[5]

To get the basic idea in the abstract, consider a setting in which values $x \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$ and $y \in \{\mathrm{d}, \mathrm{e}, \mathrm{f}\}$ will be realized. Suppose the prior belief of a Bayesian decision maker has x and y independently distributed, with positive probability on all outcomes. Imagine that the decision maker then learns that $(x, y) \in E = \{\mathrm{a}, \mathrm{b}\} \times \{\mathrm{d}, \mathrm{e}\}$. Note that E, as a product set, is a *conjunctive event*: "$x \in \{\mathrm{a}, \mathrm{b}\}$ AND $y \in \{\mathrm{d}, \mathrm{e}\}$." Because the prior satisfies independence and E is conjunctive, updating about x and y can be done separately. The decision maker's posterior belief will be given by the product of the conditional marginal probabilities,

$$\mathrm{Prob}\,[(x, y) \mid \{\mathrm{a}, \mathrm{b}\} \times \{\mathrm{d}, \mathrm{e}\}] = \mathrm{Prob}\,[x \mid \{\mathrm{a}, \mathrm{b}\}] \cdot \mathrm{Prob}\,[y \mid \{\mathrm{d}, \mathrm{e}\}], \tag{1}$$

with the marginal conditional probabilities defined by the conditional-probability formula.

The key idea is that we want the conditional probabilities on the right side of Equation 1 to be defined, and the equation to hold, *even* if the prior belief puts zero probability on $x \in \{\mathrm{a}, \mathrm{b}\}$ or on $y \in \{\mathrm{d}, \mathrm{e}\}$. For instance, if the prior satisfies $\mathrm{Prob}\,[\{a, b\}] > 0$ then the marginal conditional probability for x should be defined by the conditional-probability formula, so that $\mathrm{Prob}\,[\{a\} \mid \{a, b\}] = \mathrm{Prob}\,[\{a\}] / \mathrm{Prob}\,[\{a, b\}]$. If $\mathrm{Prob}\,[\{a, b\}] = 0$ then the marginal conditional probability for x is arbitrary but we still require it to be defined, and we still require Equation 1 to hold. It is easy to verify that the conditional properties would be defined and Equation 1 would always be satisfied if the decision maker's prior and posterior beliefs were given by the limit of a sequence of fully mixed joint distributions with x and y independent.

Shifting our attention back to games, components x and y in the abstract example now represent different dimensions of the strategy profile, and E now represents information a player receives about the strategy profile by virtue of arriving at an information set. The independence condition is imposed on updating from this player's previous information set if E is a conjunctive event.

A key element of the approach taken herein is to describe a player's belief 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. Each player is assumed to have a conjecture system that maps the player's information sets into appraisals. I specify that every player has a (possibly artificial)

[4]The idea is that a player i would not use the observation of a surprise choice by a player j to change i's belief about a move of some other player k (which could be nature) that player j did not observe.

[5]Battigalli's (1996) independence condition would be the strongest condition along these lines.

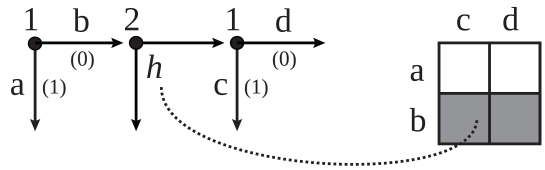

Figure 1: Independence across information sets of the same player.

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.

The following examples illustrate the building blocks of the theory. Consider first the game fragment shown in Figure 1 above. Just part of the extensive-form is pictured; the rest is inconsequential to the discussion here. Also pictured is a table indicating the possible combinations of actions for player 1 at his two information sets pictured in the tree. Suppose that backward induction identifies strategy ac for player 1 so that, at the beginning of the game, player 2 believes that player 1 will select action a at his first information set and would select action c at his second information set. This is indicated in the figure by the probability 1 on actions a and c, and probability 0 on actions b and d.

How should player 2 update her belief in the event that her information set h is reached? Backward-induction reasoning provides an answer: Player 2 has observed that player 1 selected b at his first information set—which is a surprise given player 2's initial belief—but this does not cause player 2 to change her belief that player 1 would select action c at his second information set.

We can dissect the logic as follows. At the beginning of the game, player 2 initially treats the actions at player 1's information sets as independent. Further, from the structure of the game, arriving at h provides player 2 with information that we can describe as a conjunctive event: "Player 1's action at his first information set is in $\{\mathrm{b}\}$ AND player 1's action at his second (as yet unreached) information set is in $\{\mathrm{c}, \mathrm{d}\}$." Thus, the combination of player 1's actions that are consistent with information set h being reached (the shaded region of the table) forms a product set, $\{\mathrm{b}\} \times \{\mathrm{c}, \mathrm{d}\}$. And in these circumstances, player 2 updates her belief about the action at player 1's second information set based on only what she has learned from the structure of the game about the action at player 1's second information set—which is, of course, nothing. Thus, she maintains the belief that player 1 would select c. Importantly, this logic applies even though reaching h is a surprise for player 2 in that her initial belief puts zero probability on h being reached.

Now let us apply the same logic to a player's belief about the actions of two other players, which also illustrates the idea of "no signaling what you don't know." Consider the game fragment shown in Figure 2. At the beginning of the game, player 3 believes that player 1 will select b for sure and that player 2 will choose c, d, and e with probabilities 0.2, 0.2, and 0.6 respectively. At information set h, player 3 has learned that player 1 chose a and that player 2 did not select c. That is, the set of action profiles that reach h is the product set $\{\mathrm{a}\} \times \{\mathrm{d}, \mathrm{e}\}$, and so the information at h is a conjunctive event. Player 3 then updates her belief about player 2's action on the basis of only what she has observed about player 2's action. Player 3 does not use the surprise regarding player 1's choice as an excuse to take liberties with the

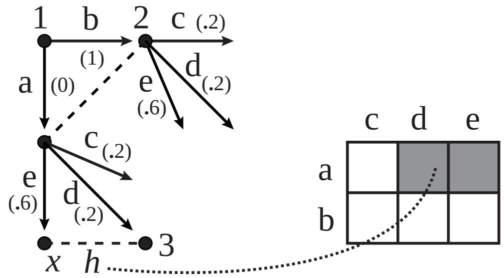

Figure 2: Independence and signaling.

probabilities of c, d, and e. Thus, player 3's updated belief puts probability 0.25 on d and probability 0.75 on e, as Bayes' rule requires for the marginal distribution over player 2's action.

The next example, shown in Figure 3 below, demonstrates that the logic of independence is embedded in the concept of subgame perfection and that weak PBE does not imply subgame perfection. The table in the figure shows the strategy profiles, with rows representing player 2's actions and columns representing profiles of actions taken at the information sets of players 1 and 3. The shaded region denotes the strategy profiles that reach player 3's information set, a conjunctive event. This game has a single subgame-perfect equilibrium, $(\mathrm{w}, \mathrm{a}, \mathrm{x})$. In the proper subgame, player 3 selects x in response to player 2's equilibrium action a. There is also a weak PBE in which $(\mathrm{z}, \mathrm{a}, \mathrm{y})$ is the strategy profile and, at information set h, player 3 believes that player 2 selected action b. Weak PBE allows for this belief because it imposes no restrictions on how beliefs are updated off the equilibrium path. Player 3 changes his belief about player 2 based on player 1's surprise selection of action w, contrary to the independence notion. Clearly, $(\mathrm{z}, \mathrm{a}, \mathrm{y})$ is not a subgame-perfect equilibrium.

In the formal definition of plain consistency, the foregoing logic is applied to the extent possible for belief updating on consecutive information sets. For instance, consider a player i who is on the move at information set h in a game, and let L denote any subset of the other players' information sets (not necessarily a proper subset and possibly including information sets of nature). At h, player i will have a belief about the strategies that the other players are using, including the actions they take at the information sets in L. Suppose that this belief exhibits independence between the behavior at the information sets in L and the behavior

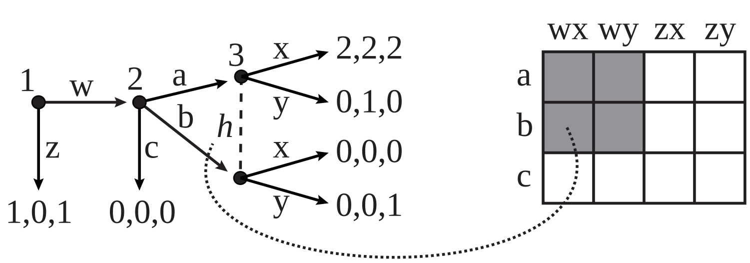

Figure 3: Weak PBE and subgame perfection.

at other information sets, in that player i views behavior in L as uncorrelated with behavior outside of L.

Suppose that the next information set encountered by player i is h', and suppose that the set of strategy profiles consistent with reaching h' is a product set with respect to L. Then plain consistency requires that player i's belief at h' must be the product of marginals with respect to L, and the marginal distribution on the information sets in L should be consistent with the conditional-probability formula restricted to these information sets. That is, player i's updated belief about actions taken at information sets in L is conditioned on only what player i has observed about these particular actions.

The definition of plain consistency goes a bit further by applying the same idea to subsets of strategy profiles. That is, we can impose the same logic on any subset Z of strategy profiles on which player i puts positive probability at both information sets h and h'. If Z is a product set with respect to L and the subset that reaches h' is also a product set, then what player i learns by arriving at h' allows for the L dimension to be separated from its complement. Assuming that player i's belief at h, when restricted to the set Z, exhibits independence between the behavior at the information sets in L and the behavior at other information sets, then his belief at h' ought to exhibit similar independence and updating should obey Bayes' rule on each dimension (as applicable).

Plain consistency puts no other restrictions on how players update their beliefs. The PBE definition combines plain consistency with the assumptions that the players' beliefs at the beginning of the game are concentrated on the the actual strategy profile and that each player's strategy is sequentially rational.

Before launching into the definitions, let me elaborate on why it is helpful to express beliefs in terms of appraisals rather than assessments. To describe whether beliefs exhibit independence regarding actions taken at different information sets, one must keep track of these various actions as separate components, as a strategy profile does. To accomplish this with assessments, one must put structure on the nodes at every information set so that each node x describes the actions taken on the path to x. But such a structure is equivalent to keeping track of the strategy profile, at least its restriction to information sets that came before the current information set. Further, even if we imagine adding this structure and then using the equilibrium strategy profile for "future" information sets when calculating expected payoffs, another complication arises: At a given information set, the other information sets in the game generally cannot be neatly classified as coming either before or after the current one.[6] Thus, I suggest that the most straightforward approach is to focus on strategy profiles and account for beliefs as appraisals.

The next section lays out the basic notation and definitions. Section 3 develops the notion of plain consistency and the equilibrium concept. Section 4 provides details on how to apply the PBE definition to infinite games and extends the definition of plain consistency accordingly. Section 5 compares plain PBE with other equilibrium definitions. The Appendix contains additional definitions related to the one-deviation property and a proof.

[6]See Kreps and Ramey (1987) for an example in which the player on the move does not know whether a particular information set for another player was already reached. 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.

2 Basic Concepts

Information Sets, Strategies, and Payoffs

Consider any extensive-form game of perfect recall, with n players and nature taking the role of "player 0." For convenience, in this and the next section the definitions and results are put in a form that applies to finite games; Section 4 extends the definitions to games with infinite strategy spaces. The definitions also apply to other dynamic representations (substitute "personal history" or "contingency" for "information set").

Define $N \equiv \{0, 1, \dots, n\}$ and let $N_+ \equiv N \setminus \{0\}$ denote the set of strategic players. Let H be the set of information sets. It is partitioned into sets $H_0, H_1, \dots, H_n$, where H_i denotes the set of information sets for player i. (As is standard, information sets are distinctly labeled so that the players' individual sets of information sets are disjoint.) Let $H_+ \equiv \cup_{i \in N_+} H_i$ be the set of information sets for the strategic players. Denote by S the space of pure strategy profiles, including nature's strategy. 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. Note that I use the symbol "$\subset$" to mean "subset," not "proper subset." Finally, let $u : S \to \mathbb{R}^n$ be the payoff function and extend it to the space of mixed strategies by the usual expected payoff calculation.

We will be dealing essentially with the *agent form* of the game, in that beliefs and choice are analyzed at individual information sets. A key element is that a player's choice at one information set is independent of his choice at another information set. For example, if player i has two different information sets, h and h', then we think of player i as being separated into two agents whose names are the information sets themselves: Agent h takes the action at information set h, agent h' takes the action at information set h', and these are independent choices. The agents in H_i all share the payoff function u_i.

Note that a strategy profile $s \in S$ maps H to the space of actions and, for each information set, specifies an action that is feasible at this information set. For any subset of information sets $L \subset H$, let s_L denote the restriction of s to the subdomain L. That is, s_L gives the profile of actions that strategy s specifies for the information sets in L. For any $L \subset H$, define $-L \equiv H \setminus L$. Note that 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; so, for instance, we write X_h and s_h instead of $X_{\{h\}}$ and $s_{\{h\}}$. Note that S_h is the set of actions available at information set h. Also, for a given player i, the subscript "i" refers to the information sets H_i. For example, s_i means the same thing as s_{H_i}. Likewise, "$-i$" refers to H_{-i}. Thus, subscripts "i" and "$-i$" have their usual meaning of identifying the strategies of player i and the other players.

Definition 1: *For a given set $L \subset H$, say that a set* ***$X \subset S$ is a product set (relative to L)*** *if $X = X_L \times X_{-L}$.*

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

Definition 2: *Given $L \subset H$ and a product set $Y = Y_L \times Y_{-L} \subset S$, say that a distribution* ***$p \in \Delta S$ exhibits independence on Y relative to L*** *if for every product set $X = X_L \times X_{-L} \subset Y$, we have $p(X)p(Y) = p(X_L \times Y_{-L}) \cdot p(Y_L \times X_{-L})$. In the case of $L = \{h\}$ for a single $h \in H$, then let us drop the brackets and say "... relative to h." Say that* ***p exhibits complete independence*** *if, for every $h \in H$, p exhibits independence relative to h on S.*

Independence on Y means that the distribution conditional on Y, which is given by the standard conditional probability formula, exhibits independence across L and $-L$. That is, the conditional probability of a product set $X = X_L \times X_{-L} \subset Y$ is the product of the conditional marginal probabilities:

$$\begin{aligned}\text{Prob}\,[X \mid Y] = \frac{p(X)}{p(Y)} &= \frac{p(X_L \times Y_{-L})}{p(Y)} \cdot \frac{p(X_{-L} \times Y_L)}{p(Y)} \\ &= \text{Prob}\,[X_L \times Y_{-L} \mid Y] \cdot \text{Prob}\,[X_{-L} \times Y_L \mid Y].\end{aligned}$$

In the expression above, the second equality is due to independence on Y relative to L. Note that independence relative to L on S implies the same on every product set $Y \subset S$.

Let $\underline{\Delta} S$ be the set of mixed strategy profiles that exhibit complete independence. Note that a mixture in $\underline{\Delta} S$ is equivalent to a *behavior strategy* profile: It specifies, for every information set h, a probability distribution over the actions available at h, and the specification is independent across information sets. We typically write such a mixture as $\sigma = (\sigma_h)_{h \in H}$, where σ_h denotes the mixed action choice at information set h. Nature's mixed strategy is taken as exogenous and we assume that it exhibits independence relative to all of nature's information sets.

It will be useful to think about information sets in terms of subsets of strategy profiles. For each $h \in H$ and $s \in S$, let us say that s *reaches* h if the path of strategy profile s includes a node in h. Denote by $S(h)$ the set of strategy profiles that reach h. Note that, for any $L \subset H$, $S(h)_L$ is the set of action profiles for the information sets in L that are consistent with h being reached.[7]

Because of perfect recall, the information sets for an individual player have a particular product structure and precedence relation. For every pair of information sets $h, h' \in H_i$ for player i, it is the case that $S(h)$ is a product set relative to h'. Further, for $h, h' \in H_i$ with $h \neq h'$, either h *is a successor of* h', in which case $S(h) \subset S(h')$; or h *is a predecessor of* h', in which case $S(h') \subset S(h)$; or neither, in which case $S(h) \cap S(h') = \emptyset$. If h' is a successor of h then every path through h' also passes through h. We 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 information set for player i between the two; that is, there is no $g \in H_i$ such that g is a successor of h and h' is a successor of g.

[7] Expressing extensive-form information sets as subsets of strategy profiles is standard. Mailath, Samuelson, and Swinkels (1993) formulate solution concepts on the basis of "normal form information sets," where there is no reference to an extensive form, and Shimoji and Watson (1998) take such "restrictions" as given (whether or not they are derived from extensive-form information sets). Note that, here, I am taking the conventional approach of examining standard extensive-form information sets but simply represent them as subsets of the strategy space.

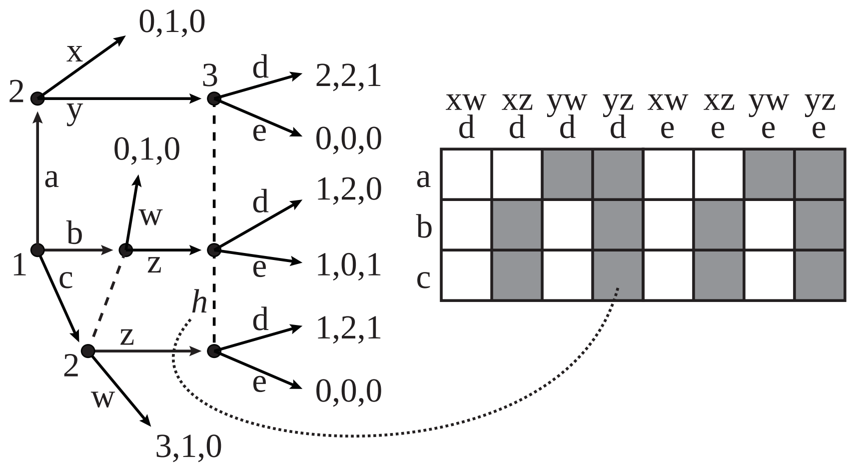

Figure 4: Example to illustrate theoretical components.

An Example

To review some of the definitions just described, consider the three-player game shown in Figure 4. The space of strategy profiles is $S = \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{x}, \mathrm{y}\} \times \{\mathrm{w}, \mathrm{z}\} \times \{\mathrm{d}, \mathrm{e}\}$, which is depicted by the table on the right side of the picture. Note that the rows of the table are the actions feasible at player 1's information set, whereas the columns are the profiles of actions for the other players' information sets. For the information set h identified in the picture (player 3's information set), we have

$$S_h = \{\mathrm{d}, \mathrm{e}\} \quad \text{and} \quad S_{-h} = \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{x}, \mathrm{y}\} \times \{\mathrm{w}, \mathrm{z}\}.$$

The subset of strategy profiles that are consistent with reaching h is

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

This set corresponds to the shaded region of the table in the figure. Clearly $S(h)$ is a product set relative to h but it is not a product set relative to the information set of player 1. That is, letting h' denote player 1's information set, we have

$$S(h)_{h'} = \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \quad \text{and} \quad S(h)_{-h'} = \{(\mathrm{x}, \mathrm{z}), (\mathrm{y}, \mathrm{w}), (\mathrm{y}, \mathrm{z})\} \times \{\mathrm{d}, \mathrm{e}\},$$

and $S(h)$ does not equal the Cartesian product of $S(h)_{h'}$ and $S(h)_{-h'}$.

Appraisal Systems

We must consider the beliefs that the strategic players hold at their various information sets. It will be useful to think of each player as having an information set that refers to "before the game begins." For this purpose, define *initial information sets* $\underline{h}_1, \underline{h}_2, \ldots, \underline{h}_n$ and extended sets $\underline{H}_1, \underline{H}_2, \ldots, \underline{H}_n$ for the strategic players as follows. For each strategic player i:

- If there exists $\hat{h} \in H_i$ such that $S(\hat{h}) = S$ then define $\underline{h}_i \equiv \hat{h}$ and let $\underline{H}_i \equiv H_i$.
- Otherwise let $\underline{h}_i$ be defined as an artificial information set with the property $S(\underline{h}_i) \equiv S$ and let $\underline{H}_i \equiv H_i \cup \{\underline{h}_i\}$.

Assume the players' artificial information sets are distinctly labelled so that $\underline{H}_1, \underline{H}_2, \ldots, \underline{H}_n$ are disjoint sets, and let $\underline{H}_+ \equiv \cup_{i \in N_+} \underline{H}_i$.

Definition 3: *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 independence on S relative to g for every $g \in H_i$. An* ***appraisal system*** *is a collection of appraisals, one for each information set of the strategic players, written $P = \left(p^h\right)_{h \in \underline{H}_+}$.*

An appraisal contains two things: The marginal on S_i gives player i's own strategy and the marginal on S_{-i} gives player i's belief about the strategy profile of the other players. In terms of player-agents, an appraisal at information set h describes agent h's belief about the other agents' behavior (this is the marginal on S_{-h}) as well as agent h's planned behavior (the marginal on S_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 information set $h \in \underline{H}_+$, player i views his strategy as independent of the other players' strategy profile, and player i's strategy is represented as a behavior strategy. Conditions on the relation between appraisals at different information sets are developed in the following sections.

Sequential Best Responses

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

Definition 4: *For a given information set $h \in H_+$ and two appraisals p^h and $\hat{p}^h$, say that* ***$\hat{p}^h$ is an h-deviation from p^h*** *if p^h and $\hat{p}^h$ are identical on all other information sets; that is, $p^h(X_{-h} \times S_h) = \hat{p}^h(X_{-h} \times S_h)$, for all $X_{-h} \subset S_{-h}$.*

Definition 5: *For a strategic player i and an information set $h \in H_i$, say that an appraisal* ***p^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 $\boldsymbol{P} = \left(\boldsymbol{p^h}\right)_{\boldsymbol{h \in \underline{H}_+}}$ is sequentially rational*** *if p^h is rational at h, for every $h \in H_+$.*

Here sequential rationality is defined in terms of what are commonly called "one-shot deviations," meaning that we evaluate player i's rationality at a given information set $h \in H_i$ by looking just at alternative choices at h rather than alternatives that would also adjust player i's behavior at other information sets that may be reached in the continuation of the game.[8] The familiar *one-deviation principle*—equivalence between single-deviation optimality and strategy-deviation optimality—holds here, assuming that player i's appraisal system has the property I call "minimal consistency." See the Appendix for more details. Minimal consistency is implied by the plain consistency condition defined in the next section.

[8]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 information set h.

Example Continued

For an example of an appraisal system and sequential best response, consider again the game shown in Figure 4. 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 lower information set, which itself is reached by $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$. Note that player 2's upper information set 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 player 2's upper information set put probability one on $(\mathrm{a}, \mathrm{x}, \mathrm{w}, \mathrm{e})$. Finally, suppose that at player 3's information set h shown in the figure, 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 information set, the player on the move cannot gain by switching from the action prescribed by her appraisal to a different action.

3 Perfect Bayesian Equilibrium

The concept of perfect Bayesian equilibrium combines sequential rationality with requirements concerning the players' appraisals, most notably regarding how updating occurs on consecutive information sets. We start by assuming that appraisals exhibit independence on product sets, in recognition of the fact that player-agents make their choices independently:

Definition 6: *Say that an* ***appraisal system $\boldsymbol{P}$ exhibits product-independence*** *if for every $h \in \underline{H}_+$, every $L \subset H$, and every product set $Y = Y_L \times Y_{-L} \subset S(h)$, appraisal p^h exhibits independence on Y relative to L.*

Plain Consistency

This takes us to the key definition. Consider a strategic player i, two information sets $h, h' \in \underline{H}_i$ such that h' is an immediate successor of h, and a set of information sets $L \subset H$. Let us call sets $Y \subset S$ and $Z \subset S$ *comparable* if $Z = Z_L \times Z_{-L} \subset S(h)$ and $Y = Y_L \times Y_{-L} = Z \cap S(h')$. That is, Y and Z are product sets, all strategy profiles in Z reach h, and Y is the subset of these strategies that reach h'. Figure 5 provides a graphical depiction.

Consider how the appraisal at h is updated to form the appraisal at h', in particular regarding the implied distributions conditional on Z. Note that, within Z, arriving at h' yields a conjunctive event Y. Thus, conditional on Z, updating on the L dimension should be separate from updating on the $-L$ dimension. Further, to calculate the updated marginals conditional on Z, Bayes' rule should be used on each dimension if applicable. Product independence already captures that, conditional on Z, the updated appraisal is the product of conditional marginals.

Consider updating on the L dimension and take any set $X_L \subset Y_L$. Note that, for the appraisal at h, the conditional probability formula defines the marginal probability of X_L conditional on Z to be $p^h(X_L \times Z_{-L})/p^h(Z)$, assuming the denominator is nonzero. The equivalent expression for the appraisal at h' is $p^{h'}(X_L \times Z_{-L})/p^{h'}(Z)$, which is equal to

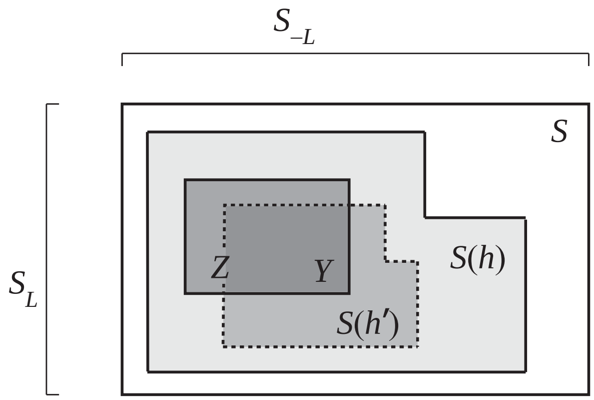

Figure 5: Illustration of the requirements for plain consistency.

$p^{h'}(X_L{\times}Y_{-L})/p^{h'}(Y)$ because $Y = Z \cap S(h')$. Bayes' rule applied to the conditional marginal probabilities on the L dimension thus requires

$$\frac{p^{h'}(X_L{\times}Y_{-L})}{p^{h'}(Y)} = \frac{p^h(X_L{\times}Z_{-L})}{p^h(Z)} \Bigg/ \frac{p^h(Y_L{\times}Z_{-L})}{p^h(Z)},$$

assuming the denominators $p^{h'}(Y)$ and $p^h(Y_L \times Z_{-L})$ are strictly positive. This expression simplifies to

$$\frac{p^{h'}(X_L{\times}Y_{-L})}{p^{h'}(Y_L{\times}Y_{-L})} = \frac{p^h(X_L{\times}Z_{-L})}{p^h(Y_L{\times}Z_{-L})}. \tag{2}$$

Another way to describe this condition is that updating preserves probability ratios at the margin, for all applicable dimensions of the strategy space.

Definition 7: *Say that* ***an appraisal system P is plainly consistent*** *if it exhibits product-independence and the following property holds for every player $i \in N_+$ and every pair of information sets $h, h' \in \underline{H}_i$ such that h' is an immediate successor of h: For every set $L \subset H$ and any sets Y and Z that are comparable and satisfy $p^h(Y_L{\times}Z_{-L}) > 0$ and $p^{h'}(Y_L{\times}Y_{-L}) > 0$, Equation 2 holds for every $X_L \subset Y_L$.*

Plain consistency has several implications. First, if $p^h(S(h')_{-i}{\times}S_i) > 0$ then plain consistency is equivalent to proper conditional-probability updating (Bayes' rule) on the space S_{-i} of the other players' strategy profile. Second, player i does not change his belief about the actions of other players at information sets that could not have been reached yet in the game. That is, if at h' player i knows that some information set $h'' \in H_{-i}$ has not been reached, player i's belief at h' about the action that would be chosen at h'' is the same as player i's belief at h, which by recursion is the same as player i's belief at the beginning of the game. Similarly, for any information set $h'' \in H_i$ that has not yet been reached at h', player i's plan at h' about the action to take at h'' is the same as player i's plan at the beginning of the game.[9] Fourth, plain consistency implies Fudenberg and Tirole's (1991) reasonableness condition for multi-period games with observed actions and independent types.[10]

[9]The first and third implications form the definition of "minimal consistency" described in the Appendix.

[10]This result is most easily seen by assuming that Nature chooses the types of the strategic players at distinct

Equilibrium Definition

The next definition means that, at the beginning of the game, the players' appraisals coincide on a given strategy profile.

Definition 8: *Consider a mixed (behavior) strategy profile $\sigma \in \underline{\Delta}S$. Say that an* ***appraisal system*** $\boldsymbol{P} = \left(\boldsymbol{p^h}\right)_{\boldsymbol{h} \in \underline{\boldsymbol{H}}_+}$ ***conforms to*** $\boldsymbol{\sigma}$ *if, for every player i, $p^{\underline{h}_i} = \sigma$.*

With the essential ingredients in place, the general definition of perfect Bayesian equilibrium is stated thus:

Definition 9: *Taking nature's strategy $(\sigma_h)_{h \in H_0}$ as given, say that* ***an appraisal system*** $\boldsymbol{P}$ ***is a perfect Bayesian equilibrium*** *and say that* ***behavior strategy profile*** $\boldsymbol{\sigma} \in \underline{\boldsymbol{\Delta}}\boldsymbol{S}$ ***is a PBE strategy profile*** *if P conforms to σ, P is plainly consistent, and P is sequentially rational.*

This version of PBE is straightforward to apply because it allows the analyst to examine the players individually and, for each player, one must look only at how beliefs evolve between consecutive information sets. It also implies subgame perfection.

Theorem: *If P is perfect Bayesian equilibrium of a finite game then the PBE strategy profile σ is a subgame perfect Nash equilibrium.*

Examples Continued

Consider the game shown in Figure 3 in the Introduction. As noted earlier, strategy profile $(\mathrm{z}, \mathrm{a}, \mathrm{y})$ is a weak PBE strategy profile but it is not subgame perfect and is not a PBE strategy profile. If it were a PBE strategy profile, then player 3's appraisal at the beginning of the game would put probability one on $(\mathrm{z}, \mathrm{a}, \mathrm{y})$. But at player 3's information set h, he has learned only that player 1 deviated, so plain consistency requires that player 3 continue to believe that player 2 selects action a, and then y would not be optimal.

To express this logic formally, let L be the singleton consisting of player 2's information set, let $Z = \{\mathrm{w}, \mathrm{z}\} \times \{\mathrm{a}, \mathrm{b}\} \times \{\mathrm{x}, \mathrm{y}\}$, and define $Y \equiv S(h) \cap Z$. Observe that Z is a product set relative to L and

$$Y = \{\mathrm{w}\} \times \{\mathrm{a}, \mathrm{b}\} \times \{\mathrm{x}, \mathrm{y}\} = S(h)$$

is also a product set relative to L. Note that $Y_L = Z_L = \{\mathrm{a}, \mathrm{b}\}$, so $Y_L \times Z_{-L} = Z$. Note as well that h is an immediate successor of $\underline{h}_3$ for player 3, $p^{\underline{h}_3}(Y_L \times Z_{-L}) > 0$, and $p^h(Y) = 1 > 0$. Plain consistency thus implies that

$$\frac{p^h\left(\{\mathrm{a}\} \times \{(\mathrm{w}, \mathrm{x}), (\mathrm{w}, \mathrm{y})\}\right)}{p^h(Y)} = \frac{p^{\underline{h}_3}(\{\mathrm{a}\} \times \{(\mathrm{w}, \mathrm{x}), (\mathrm{w}, \mathrm{y}), (\mathrm{z}, \mathrm{x}), (\mathrm{z}, \mathrm{y})\})}{p^{\underline{h}_3}(Z)}.$$

information sets. That is, for each strategic player $i \in N_+$, let h^{0i} be the information set of Nature where player i's type is selected, and assume that $h^{0i} \neq h^{0j}$ for distinct strategic players i and j. Then, for a given player $i \in N_+$, let $L = H_i \cup \{h^{0i}\}$ and apply plain consistency with $Z = S(h')_L \times S(h)_{-L}$. Note that, in addition to the reasonableness condition, Fudenberg and Tirole assume that any two players always have the same belief about a third player's type. If we restrict attention to the class of multi-period games with observed actions and uncorrelated types, then requiring this condition along with plain consistency would produce Fudenberg and Tirole's PBE definition for this class of games.

The denominators equal one and the numerator on the right also equals one, so the numerator on the left equals one. That is, at h player 3 must believe that player 2 selected action a, and therefore action y is not optimal.

The only PBE in this example is the appraisal system that, at every information set, puts probability one on profile $(\mathrm{w}, \mathrm{a}, \mathrm{x})$. This strategy profile reaches all information sets, so plain consistency is trivially satisfied, and the appraisal system is clearly sequentially rational. Strategy profile $(\mathrm{w}, \mathrm{a}, \mathrm{x})$ is also the only subgame-perfect Nash equilibrium of the game.

Next consider the example shown in Figure 4 and the appraisal system described at the end of Section 2. The players' initial appraisals all put probability one on $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$, and player 2's lower information set has the same appraisal. At player 2's upper information set, the appraisal puts probability one on $(\mathrm{a}, \mathrm{x}, \mathrm{w}, \mathrm{e})$. Finally, at player 3's information set h shown in the figure, the appraisal puts probability one on $(\mathrm{b}, \mathrm{x}, \mathrm{z}, \mathrm{e})$. As noted before, this appraisal system is sequentially rational. Further, the beliefs trivially satisfy Bayes' rule updating on the equilibrium path and each player's belief about the others' behavior at future information sets coincides with the strategy profile, so the conditions for a weak PBE are satisfied and thus $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$ is a weak PBE profile.

However, the appraisal system just described is not plainly consistent. In particular, at information set h player 3 is supposed to believe that player 1 selected action b rather than action c, but player 3 has observed nothing from the structure of the game that could allow him to distinguish between actions b and c. Conditional on b or c having been played, player 3 has learned something about the choice at player 2's lower information set (that z was played) but he has learned nothing about the choice of b versus c. At h player 3 may think that action a was played, but he cannot change from his initial belief that c would be played to now believe it is b.

To demonstrate this formally, let $Z = \{\mathrm{b}, \mathrm{c}\} \times \{\mathrm{x}, \mathrm{y}\} \times \{\mathrm{w}, \mathrm{z}\} \times \{\mathrm{d}, \mathrm{e}\}$ and define $Y \equiv S(h) \cap Z$. Observe that Z is a product set relative to player 1's information set and

$$Y = \{\mathrm{b}, \mathrm{c}\} \times \{\mathrm{x}, \mathrm{y}\} \times \{\mathrm{z}\} \times \{\mathrm{d}, \mathrm{e}\}$$

is also a product set relative to player 1's information set. Note that h is an immediate successor of $\underline{h}_3$ for player 3, and $p^h(Y) = 1 > 0$. Plain consistency thus implies that

$$\frac{p^h\left(\{\mathrm{c}\} \times \{\mathrm{x}, \mathrm{y}\} \times \{\mathrm{z}\} \times \{\mathrm{d}, \mathrm{e}\}\right)}{p^h(Y)} = \frac{p^{\underline{h}_3}(\{\mathrm{c}\} \times \{\mathrm{x}, \mathrm{y}\} \times \{\mathrm{w}, \mathrm{z}\} \times \{\mathrm{d}, \mathrm{e}\})}{p^{\underline{h}_3}(Z)}.$$

The denominators equal one and the numerator on the right also equals one, so the numerator on the left equals one. That is, at h, conditional on player 3 putting positive probability on the subset of actions $\{\mathrm{b}, \mathrm{c}\}$, player 3 must put all of the probability weight on action c, as he did at the beginning of the game.

The foregoing discussion shows that, in this game, there is no PBE in which actions c and w are played for sure. There is a PBE in which the appraisals at all information sets put probability one on strategy profile $(\mathrm{a}, \mathrm{y}, \mathrm{z}, \mathrm{d})$, except at player 2's lower information set where the appraisal puts probability one on $(\mathrm{b}, \mathrm{y}, \mathrm{z}, \mathrm{d})$. To summarize, of this game's four pure-strategy Nash equilibria $(\mathrm{a}, \mathrm{y}, \mathrm{z}, \mathrm{d})$, $(\mathrm{c}, \mathrm{x}, \mathrm{z}, \mathrm{d})$, $(\mathrm{c}, \mathrm{y}, \mathrm{w}, \mathrm{e})$, and $(\mathrm{c}, \mathrm{x}, \mathrm{w}, \mathrm{e})$, the first and last are weak PBE profiles, but only the first is a PBE profile.[11]

[11] There is a continuum of similar PBE that differ on the appraisal at player 2's lower information set with

4 Extension to Infinite Games

This section develops an infinite-game extension of plain consistency, described in a way that applies to any game. The condition has force wherever the game and appraisals have the appropriate structure. There are two parts. First, we add the "measurable" qualifier to the sets considered in Definitions 2 and 6, where measurability is defined for whatever distributions are considered in the application of interest.

The second part deals with probability densities. To build intuition, consider any finite game. Fix L and consider an indexed collection $\{(X_L(r), Y(r), Z(r)) \mid r \in R\}$, where R is a finite set. Suppose that, for each $r \in R$, it is the case that $Y(r)$ and $Z(r)$ are comparable and $X_L(r) \subset Y_L(r)$.[12] Then for each r, plain consistency requires

$$p^{h'}(X_L(r) \times Y_{-L}(r)) = \frac{p^h(X_L(r) \times Z_{-L}(r))}{p^h(Y_L(r) \times Z_{-L}(r))} \cdot p^{h'}(Y_L(r) \times Y_{-L}(r)), \tag{3}$$

assuming that $p^h(Y_L(r) \times Z_{-L}(r)) > 0$.[13] Clearly, we can sum both sides over r and the resulting equation is an implication of plain consistency.

The extended notion of plain consistency considers r on a continuum, allows for p^h to be partly characterized by a density function, and replaces the summation with an integral. Further, assuming the sets $\{Z(r)\}$ are disjoint, we can replace the integral on the left side with the probability of the union of $X_L(r) \times Y_{-L}(r)$ over r.

Additional notation helps to make this formal. Suppose that $r \in R = [0, \overline{r}]$, where $\overline{r} \geq 0$. Define

$$\mathbb{XZ}(r) \equiv \bigcup_{r' \in [0,r]} X_L(r') \times Z_{-L}(r') \quad \text{and} \quad \mathbb{YZ}(r) \equiv \bigcup_{r' \in [0,r]} Y_L(r') \times Z_{-L}(r'),$$

and define $\mathbb{XY}(r)$, and $\mathbb{YY}(r)$ similarly. Suppose that for all r, $Y(r)$ and $Z(r)$ are comparable relative to L, h, and h'. Suppose that $X_L(r) \subset Y_L(r)$ for all r. Assume that the sets $\{Z(r)\}$ are disjoint, the sets $\{\mathbb{XZ}(r)\}$ and $\{\mathbb{YZ}(r)\}$ are p^h-measurable, and $p^h(\mathbb{YZ}(\overline{r})) > 0$.

Suppose that on $\{\mathbb{YZ}(r)\}$ and $\{\mathbb{XZ}(r)\}$, distribution p^h is characterized by a density function except where $p^h(\mathbb{YZ}(r))$ is discontinuous, and suppose that the density for $\{\mathbb{YZ}(r)\}$ is everywhere strictly positive. For values of r where the density at $\mathbb{YZ}(r)$ exists, define $q(r)$ to be the ratio of the density of p^h at $\mathbb{XZ}(r)$ to the density of p^h at $\mathbb{YZ}(r)$; and for values of r where $p^h(\mathbb{YZ}(r))$ is discontinuous, define $q(r)$ to be the ratio $p^h(X_L(r) \times Z_{-L}(r))/p^h(Y_L(r) \times Z_{-L}(r))$. Then extended plain consistency requires that the sets $\{\mathbb{XY}(r)\}$ and $\{\mathbb{YY}(r)\}$ are $p^{h'}$-measurable and

$$p^{h'}(\mathbb{XY}(\overline{r})) = \int_0^{\overline{r}} q(r) dp^{h'}(\mathbb{YY}(r)). \tag{4}$$

regard to the probabilities of actions b and c. There are also PBE in which player 3 randomizes. For instance, player 1 selects a with probability $1/4$, b with probability $1/2$, and c with probability $1/4$. Player 2 selects y and z for sure, and player 3 mixes between d and e with equal probabilities. Because all information sets are reached with positive probability, plain consistency is easy to confirm.

[12] Recall that comparable means that $Z(r) = Z_L(r) \times Z_{-L}(r) \subset S(h)$ and $Y(r) = Y_L(r) \times Y_{-L}(r) \equiv Z(r) \cap S(h')$.

[13] This is just Equation 2 with the left-side denominator multiplied through. Note that Equation 3 is valid even if $p^{h'}(Y_L(r) \times Y_{-L}(r)) = 0$. Further, $Y(r) = \emptyset$ is allowed.

The basic definition of plain consistency in Section 3 follows from this expression in the special case of $\overline{r} = 0$.

Example – Evidence Disclosure

Consider a version of the evidence-disclosure game studied by Bull and Watson (2017). At the beginning of the game, nature chooses a vector (θ, ω, e), where $\theta \in \{0, 1\}$, $\omega \in [0, 1]$, and $e \in \{\text{doc}, \phi\}$. Nature's selection is made according to an atomless joint distribution, with density $f(\theta, \omega, e)$ defined everywhere and integrable. The value ω represents a private signal to player 1. The value e refers to a document that player 1 may possess and, if so, can be disclosed to player 2. Existence of the document is signified by $e = \text{doc}$, whereas $e = \phi$ means the document does not exist. Assume that $f(1, \omega, \text{doc}) + f(0, \omega, \text{doc}) > 0$ for all ω.

Player 1 privately observes ω and e. If $e = \text{doc}$, then player 1 decides whether to disclose the document. If $e = \phi$ then player 1 has no action. Player 2 observes only whether the document is disclosed. If the document is not disclosed, then player 2 cannot distinguish between the case in which the document does not exist and the case in which it exists but player 1 did not disclose it. Player 2 then chooses an action $a \in [0, 1]$. Player 1's payoff is a and player 2's payoff is $-(a - \theta)^2$.

Because e, ω, and θ are generally correlated, player 1's disclosure or nondisclosure of the document may provide information about θ to player 2. The document itself provides hard evidence, because it cannot be disclosed when it does not exist. The document can also serve as a "soft" signal of ω and θ, to the extent that player 1's disclosure choice depends on ω.

Regardless of nature's distribution, this game has a weak PBE in which no information is transmitted from player 1 to player 2 in equilibrium. Player 1's strategy is to never disclose the document. If the document is not disclosed then player 2 believes that (θ, ω, e) is distributed according to f and she selects $a = E[\theta]$. If the document is disclosed then player 2 believes that $\theta = 0$ and she chooses $a = 0$. Note that, essentially, disclosure of the document perfectly signals $\theta = 0$ to player 2.

This equilibrium could possibly make sense if there is a value $\hat{\omega} \in [0, 1]$ for which $f(0, \hat{\omega}, \text{doc}) > 1$ and $f(1, \hat{\omega}, \text{doc}) = 0$, because then there is a contingency in which player 1 learns from his private signal and the existence of the document that $\theta = 0$. Player 2 can believe that player 1 would err in implementing his strategy only if $\hat{\omega}$ was realized, and so a surprise disclosure would lead player 2 to believe that $\theta = 0$. But if there is no such value $\hat{\omega}$ then this equilibrium is unrealistic, because off-equilibrium-path disclosure would be signaling something that player 1 does not know. How can we rule out such an equilibrium in this case? The game is not finite, so definitions of sequential equilibrium and PBE from the prior literature do not apply. Additionally, this setting is not a multistage game with observed actions and independent types. But we can apply plain consistency and the PBE definition developed herein.

Let $h_2(\text{doc})$ denote the information set in which player 2 has observed disclosure of the document. We can apply the extended plain consistency condition to get an expression for player 2's belief about θ at this information set. Let $L = H_0 = \{\underline{h}_0\}$, so that $-L \equiv H_{-0}$ is the set of information sets for players 1 and 2. For every $\omega \in [0, 1]$, let S_1^ω be the set of strategies for player 1 that disclose in the event that ω is the private signal and the document exists. Note that the sets $\{S_1^\omega\}$ are distinct and not disjoint. We have

$$S(h_2(\mathrm{doc})) = \bigcup_{\substack{\theta\in\{0,1\}\\ \omega\in[0,1]}} \left[\{(\theta,\omega,\mathrm{doc})\}\times S_1^\omega \times S_2\right].$$

Define:

$$\begin{array}{ll} X_L(\omega) \equiv \{(1,\omega,\mathrm{doc})\} & Y_L(\omega) \equiv \{(\theta,\omega,\mathrm{doc}) \mid \theta \in \{0,1\}\} \\ Y_{-L}(\omega) \equiv S_1^\omega \times S_2 & Z(\omega) \equiv Y_L(\omega)\times S_1 \times S_2. \end{array}$$

Note that $S(h_2(\mathrm{doc}))$ is not a product set relative to L, but it is the union over ω of the product sets $Y_L(\omega)\times Y_{-L}(\omega)$, which we recall is written $\mathbb{YY}(1)$. Also, $\mathbb{XY}(1)$ is the subset of $S(h_2(\mathrm{doc}))$ in which $\theta = 1$.

We want an expression for $p^{h_2(\mathrm{doc})}(\mathbb{XY}(1))$, which is the probability that player 2 puts on $\theta = 1$ in the event that the document is disclosed. We can use plain consistency by relating player 2's appraisal at $h_2(\mathrm{doc})$ to her appraisal at $\underline{h}_2$, the beginning of the game.

The required conditions to apply plain consistency are satisfied. For all ω, $Y(\omega)$ and $Z(\omega)$ are comparable (relative to L, $\underline{h}_2$, and $h_2(\mathrm{doc})$) and $X_L(\omega) \subset Y_L(\omega)$. The sets $\{Z(\omega)\}$ are disjoint, the sets $\{\mathbb{XZ}(\omega)\}$ and $\{\mathbb{YZ}(\omega)\}$ are $p^{\underline{h}_2}$-measurable, and $p^{\underline{h}_2}(\mathbb{YZ}(1)) > 0$ by assumption. Observe that the density of $p^{\underline{h}_2}$ at $\mathbb{XZ}(\omega)$ is $f(1,\omega,\mathrm{doc})$ and the density of $p^{\underline{h}_2}$ at $\mathbb{YZ}(\omega)$ is $f(1,\omega,\mathrm{doc}) + f(0,\omega,\mathrm{doc})$. The ratio,

$$\frac{f(1,\omega,\mathrm{doc})}{f(1,\omega,\mathrm{doc}) + f(0,\omega,\mathrm{doc})} \equiv q(\omega),$$

is the probability of $\theta = 1$ conditional on the value of ω, using the standard conditional-probability formula for the density f.

Plain consistency requires that the sets $\{\mathbb{XY}(\omega)\}$ and $\{\mathbb{YY}(\omega)\}$ are $p^{h_2(\mathrm{doc})}$-measurable and

$$p^{h_2(\mathrm{doc})}(\mathbb{XY}(1)) = \int_0^1 q(\omega)dp^{h_2(\mathrm{doc})}(\mathbb{YY}(\omega)).$$

Since $S(h_2(\mathrm{doc})) = \mathbb{YY}(1)$, we know that $dp^{h_2(\mathrm{doc})}(\mathbb{YY}(\omega))$ integrates to 1. Thus, in the event of disclosure, the probability that player 2 puts on $\theta = 1$ must be a convex combination of the conditional probabilities $\{q(\omega)\}$. That is, essentially player 2 updates her belief about player 1's strategy (for which values of ω he discloses) but does not change her belief about nature's strategy. An implication is that if f is bounded away from zero, then the probability player 2 puts on $\theta = 1$ in the event of disclosure must be similarly bounded and the unrealistic equilibrium described above is not a PBE. Additional implications can be easily derived regarding the relation between the hard and soft aspects of the evidence.

Extended plain consistency also implies that, upon observing ω and e, player 1's belief about nature's selection of θ is properly derived from f. For example, let $L = \{\underline{h}_0\}$ as before and consider the information set $h_1(\omega)$ where player 1 observes ω and has the document. For every ω, define

$$\begin{array}{ll} X_L(\omega) \equiv \{(1,\omega,\mathrm{doc})\} & Y_L(\omega) \equiv \{(\theta,\omega,\mathrm{doc}) \mid \theta \in \{0,1\}\} \\ Y_{-L}(\omega) \equiv S_1 \times S_2 & Z(\omega) \equiv Y_L(\omega)\times S_1 \times S_2. \end{array}$$

The conditions for plain consistency are clearly satisfied. Note that, for any specific private signal value ω', $S(h_1(\omega') \cap Z(\omega)$ is nonempty only for $\omega = \omega'$. So plain consistency requires $p^{h_1(\omega)}(X_L(\omega) \times S_1 \times S_2) = q(\omega) p^{h_1(\omega)}(Y_L(\omega) \times S_1 \times S_2)$, where q is defined as before. Because $p^{h_1(\omega)}(Y_L(\omega) \times S_1 \times S_2) = 1$, player 1's belief about θ is given by the standard conditional-probability formula for densities.

5 Comparisons and Conclusion

Other versions of PBE that have appeared in the recent literature include concepts defined by Bonanno (2013), González-Díaz and Meléndez-Jiménez (2014), and Mailath (2016). Their technical 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.[14] 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 information set 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 information set to the next (generally for a different player) in the case in which the latter information set follows with positive probability given the assessment at the former.

The PBE definitions proposed by Bonanno (2013), González-Díaz and Meléndez-Jiménez (2014), and Mailath (2016) do not address conditional-probability updating on separate dimensions of the strategy space—as is needed for reasonableness—and thus in this sense are weaker than plain PBE. For example, in the game shown in Figure 2, these definitions do not restrict the belief of player 3 at information set h regarding player 2's actions d and e, whereas plain consistency nails it down to probability 0.25 on d and probability 0.75 on e. 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.[15]

Next consider the relation between plain PBE and sequential equilibrium. Plain consistency is implied by *full consistency* (Kreps and Wilson 1982), in which the beliefs at all information sets are derived from a sequence of fully mixed behavior strategies.[16] Thus, plain PBE

[14]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 information sets (conditional on reaching the subgame) are consistent with the strategy profile, which implies subgame perfection.

[15]Weibell (1992) proposes another definition that sits between subgame perfection and perfect equilibrium; it is based on strategy trembles. Weibull examines beliefs as probability distributions over strategy profiles, as I advocate here.

[16]Specifically, full consistency considers a convergent sequence of fully mixed behavior strategies (satisfying complete independence), $\{\sigma^k\}$. For each element k of the sequence, we define an appraisal system P^k using the standard conditional probability formula: $p^h(X) \equiv \sigma^k(X \cap S(h))/\sigma^k(S(h))$ for all $X \subset S$ and $h \in \underline{H}_+$. An appraisal system P is called *fully consistent* if it is the limit of such a sequence $\{P^k\}$. It is not difficult to

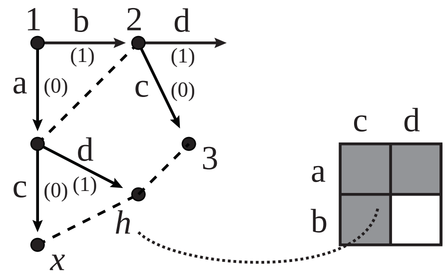

Figure 6: Beliefs allowed by plain consistency.

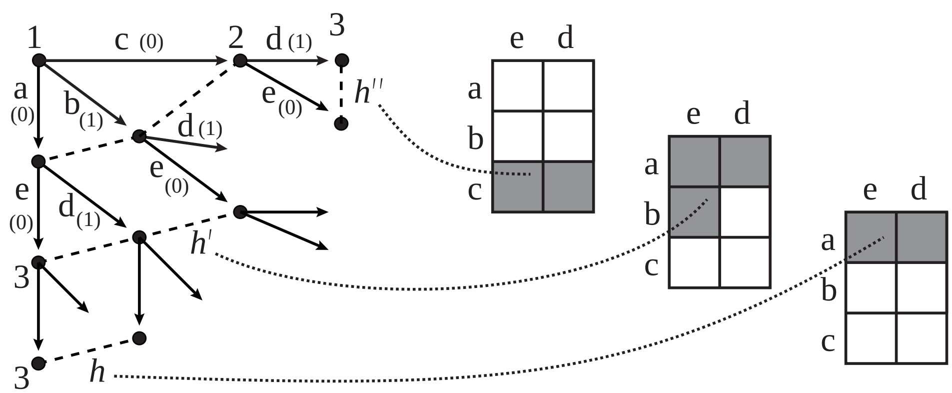

Figure 7: Path dependence.

is weaker than sequential equilibrium, and existence is ensured for finite games. Existence for classes of infinite games is beyond the scope of this note.

To see how plain PBE is weaker than sequential equilibrium, consider the game fragment shown in Figure 6. 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 information set h. Because $S(h)_{-3}$ does not have a product structure allowing player 1's and player 2's actions to be separated, the only requirements of plain consistency are that $p^h \in \Delta S(h)$ and that p^h treats player 3's own actions independently of s_{-3}. Full consistency, as well as Battigalli's independence condition, requires putting zero probability on the profile ac (node x).

The example shown in Figure 7 demonstrates that plain PBE allows for some path dependence in beliefs. 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 information set h' to have any probability distribution over $S(h')_{-3} = \{(\mathrm{a}, \mathrm{e}), (\mathrm{a}, \mathrm{d}), (\mathrm{b}, \mathrm{e})\}$, because his initial belief puts zero probability on reaching h' and $S(h')_{-3}$ does not have a product structure that allows for separating the actions of players 1 and 2. This means that player 3's beliefs about the relative likelihood of e and d may be

directly verify that a fully consistent appraisal system is plainly consistent. A two-step route to this conclusion is that Kohlberg and Reny (1997) show that full consistency implies Battigalli's (1996) independence property, and Watson (2017) shows that this implies plain consistency.

different at h and h'', whereas full consistency and Battigalli's independence condition require these to be the same.

One could strengthen plain PBE, and rule out the path dependence just described, by assuming that the condition of plain consistency must be applied between every information set and all of its successors, not just to immediate successors. The most stringent PBE concept, called "mutual PBE" in Watson (2017), requires that all of the players' appraisals be derived from a single conditional probability system that satisfies Battigalli's (1996) independence condition.[17] It is the closest PBE concept to sequential equilibrium, but it is still weaker than sequential equilibrium. Although it may have normative appeal, mutual PBE shares with sequential equilibrium some disadvantages relative to plain PBE. One is computational complexity, as it appears the most straightforward way to find a conditional probability system that satisfies the independence condition in a finite game is to construct a convergent sequence of fully mixed behavior strategies with desired properties, which means calculating a sequential equilibrium anyway. And there is the issue of how to analyze infinite games.[18] Another reason to focus on plain PBE is that, for some applications, it may not be realistic to impose many restrictions on appraisals across information sets of the same player and across players. By imposing conditions on only the relation between appraisals at each player's consecutive information sets, plain PBE is meant to capture how practitioners typically construct beliefs in applications and it affords flexibility to accommodate theories of belief revision.

A Appendix

This appendix contains some additional definitions related to the one-deviation property and a proof of the theorem presented in Section 3. Let us start by defining *minimal consistency* for a player's appraisals on consecutive information sets to mean that (i) the player maintains his plan of action for the current and unreached information sets and (ii) the player's belief about the other players' strategy profile is updated using Bayes' rule if possible.

Definition 10: *Say that* ***an appraisal system P is minimally consistent*** *if the following conditions hold for every strategic player i:*

1. *For every pair of information sets $g, h \in H_i$ such that g is not a predecessor of h (with $g = h$ allowed), $p^h(X_g \times S_{-g}) = p^{\underline{h}_i}(X_g \times S_{-g})$ for all $X_g \subset S_g$.*
2. *For every pair of information sets $h, h' \in \underline{H}_i$ such that h' is an immediate successor of h, if $p^h(S(h')_{-i} \times S_i) > 0$ then for all $X_{-i} \subset S(h')_{-i}$ we have:*

$$p^{h'}(X_{-i} \times S_i) = \frac{p^h(X_{-i} \times S_i)}{p^h(S(h')_{-i} \times S_i)}.$$

[17] It is without consequence to require that each player i's beliefs about the others is derived from a conditional probability system on S_{-i}. The restriction is in requiring all of the players' appraisals to be derived from the same conditional probability system, and the independence condition further restricts the beliefs.

[18] The computational issues lead to trade-offs in modeling choices. For instance, continuum action spaces offer a convenience for some applied settings, especially where calculus can be used to characterize optima. For these settings, a plain PBE construction on the infinite game may be preferred to examining sequential equilibrium on a finite version.

Under minimal consistency, each player i starts the game with an appraisal $p^{\underline{h}_i}$ and then, as the game progresses, he maintains his own strategy for the current and future information sets and, regarding his beliefs about the others, he updates from one information set to the next using a proper conditional probability calculation wherever possible. 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 information set.[19] Minimal consistency is clearly implied by plain consistency.

Consider an information set $h \in H_i$ where player i is on the move. Let K^h be the subset of H_i consisting of h and all of its successors. Then at information set $h \in H_i$, the *continuation strategy* for player i that appraisal p^h prescribes is simply the marginal distribution on K^h, that is $p^h(\cdot \times S_{-K^h})$.

Definition 11: *For a given information set $h \in H_+$ and two appraisals p^h and $\hat{p}^h$, say that* ***$\hat{\mathbf{p}}^{\mathbf{h}}$ is a continuation-strategy h-deviation from $\mathbf{p}^{\mathbf{h}}$*** *if p^h and $\hat{p}^h$ are identical on information sets $H \backslash K^h$; that is, $p^h(X_{-K^h} \times S_{K^h}) = \hat{p}^h(X_{-K^h} \times S_{K^h})$, for all $X_{-K^h} \subset S_{-K^h}$.*

In other words, for player i's information set h, a continuation-strategy h-deviation alters only the actions played at player i's information sets in the continuation of the game (that is, at h and successor information sets for player i). Standard arguments can be used to prove the following one-deviation result (see Hendon, Jacobsen, and Sloth 1996 and Perea 2002).[20]

Result: *[One-Deviation Property] Consider a minimally consistent appraisal system P for a finite game. Then P is sequentially rational if and only 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 $u_i(p^h) \geq u_i(\hat{p}^h)$.*

The theorem on subgame perfection is restated next, along with a proof.

Theorem: *If P is perfect Bayesian equilibrium of a finite game then the PBE strategy profile σ is a subgame perfect Nash equilibrium.*

Proof: Let P be a PBE for a given finite game and let σ denote the PBE strategy profile. Consider any subgame and let G be the set of information sets in this subgame. For any player i, define $G_i = G \cap H_i$ and $G_{-i} = G \cap H_{-i}$. The initial node of the subgame is a singleton information set, denoted $\hat{h}$. We must show that σ_G is a Nash equilibrium in the subgame—that is, for every strategic player i, σ_{G_i} is a best response to $\sigma_{G_{-i}}$.

If $\hat{h} = \underline{h}_i$ 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 single-deviation property that σ_i is a best response

[19] To extend the definition of minimal consistency to infinite strategy spaces, the game must have enough structure so that sets of the form $S(h)$ are measurable, and we must restrict attention to measurable sets X_{-i}. The condition $p^h(S(h')_{-i} \times S_i) > 0$ becomes "$S(h')_{-i} \times S_i$ is in the support of p^h" and the conditional-probability expression should be put in density-function form where appropriate. Note that by requiring Bayes' rule on unreached paths (those not assigned positive probability by appraisals at the beginning of the game), minimal consistency is in one sense more stringent than the condition for weak PBE. However, weak PBE is more stringent in another respect: It captures the idea that "future play" is specified by the strategy profile, even following zero-probability events.

[20] Hendon, Jacobsen, and Sloth's (1996) notion of "preconsistency" is equivalent to the second condition of minimal consistency.

to $p^{\hat{h}}_{-i} = \sigma_{-i}$. Otherwise, there is a unique sequence of information sets $\{h^1, h^2, \ldots, h^M\} \subset H_i \backslash G_i$, such that:

- $h^1 = \underline{h}_i$;
- for each $m = 1, 2, \ldots, M-1$, h^{m+1} is an immediate successor of h^m; and
- there is an information set in G_i that is an immediate successor of h^M.

Let $\underline{G}_i$ be the subset of G_i that are immediate successors of h^M for player i.

Suppose that σ_{G_i} is not a best response to $\sigma_{G_{-i}}$ in the subgame, and we will find a contradiction. Observe first that, for every $h \in H_+$, $S(h)$ has a product structure relative to G. To see this, note that if $h \notin G$ then actions taken at information sets in G cannot influence whether h is reached and $S(h)_G = S_G$. Further, if $h \in G$ then the player on the move knows that the initial node of the subgame was reached, implying a unique sequence of actions on the path leading to the subgame and no implications for the other information sets outside of G.

To evaluate payoffs in the subgame, we can let $\hat{\sigma}$ be any element of ΔS that exhibits independence on S relative to G, exhibits independence on S relative to all of player i's information sets, puts probability one on the actions needed to reach $\hat{h}$ so that $\hat{\sigma}$ reaches the subgame with probability one, and otherwise coincides with σ (in particular, $\hat{\sigma}_G = \sigma_G$). Player i's payoff in the subgame (that is, conditional on the subgame being reached) under strategy profile σ is $u_i(\hat{\sigma})$.

On the presumption that σ_{G_i} is not a best response to $\sigma_{G_{-i}}$ in the subgame, there exists a mixture $\hat{\sigma}' \in \Delta S$ that exhibits independence on S relative to all of player i's information sets and that differs from $\hat{\sigma}$ only on G_i, such that $u_i(\hat{\sigma}') > u_i(\hat{\sigma})$. In words, player i gains by altering his behavior in the subgame. Note that player i's expected payoff in the subgame can be written as a weighted sum over payoffs from each of the information sets in $\underline{G}_i$. Further, changing his behavior from one such information set does not affect the expect payoff from any of the other information sets in $\underline{G}_i$. Therefore, for at least one information set $\tilde{h} \in \underline{G}_i$, player i can increase his expected payoff in the subgame by altering his behavior in the continuation from $\tilde{h}$. That is, we can assume that $\hat{\sigma}'$ differs from $\hat{\sigma}$ only on the information sets in the continuation of the game from $\tilde{h}$.

Note that distributions $\hat{\sigma}$ and $\hat{\sigma}'$ reach $\hat{h}$ for sure, the path from $\hat{h}$ to $\tilde{h}$ does not involve any actions of player i, and $\hat{\sigma}'_{G_{-i}} = \hat{\sigma}_{G_{-i}}$ (behavior of the other players in the subgame is the same for these distributions). These facts imply $\hat{\sigma}(S(\tilde{h})) = \hat{\sigma}'(S(\tilde{h}))$. Also, with both distributions, $\tilde{h}$ must be reached with positive probability, or otherwise player i's behavior from $\tilde{h}$ would not affect his expected payoff in the subgame. Thus, we have $\hat{\sigma}(S(\tilde{h})) = \hat{\sigma}'(S(\tilde{h})) > 0$. Using the fact that $\hat{\sigma}'$ differs from $\hat{\sigma}$ only from $\tilde{h}$, we can rewrite the inequality $u_i(\hat{\sigma}') > u_i(\hat{\sigma})$ by conditioning on $\tilde{h}$. Specifically, define probability distributions $\tilde{\sigma}$ and $\tilde{\sigma}'$ with support $S(\tilde{h})$ by

$$\tilde{\sigma}(X) = \frac{\hat{\sigma}(X)}{\hat{\sigma}(S(\tilde{h}))} \text{ and } \tilde{\sigma}'(X) = \frac{\hat{\sigma}'(X)}{\hat{\sigma}'(S(\tilde{h}))} \tag{5}$$

for all $X \subset S(\tilde{h})$. We then have $u_i(\tilde{\sigma}') > u_i(\tilde{\sigma})$.

The final step involves showing that $p^{\tilde{h}}_G = \tilde{\sigma}_G$. That is, regarding behavior in the subgame, player i's appraisal at $\tilde{h}$ coincides with $\tilde{\sigma}$. This follows from plain consistency using induction.

Start with $m = 1$ and note the following for $Z = S(h^2)_G \times S(h^1)_{-G}$ and $Y = S(h^2) = S(h^2)_G \times S(h^2)_{-G}$: By construction, h^2 is an immediate successor of h^1. By definition, $Z_G \subset Y_G$ (in fact, these sets are equal) and $S(h^2)$ is a product set relative to G. Also, $p^{h^2}(Y) = 1$ and $p^{h^1}(Z) = 1$. Using the definition of plain consistency, we conclude that $p_G^{h^2} = p_G^{h^1} = \sigma_G$. Repeating the argument establishes that $p_G^{h^M} = \sigma_G$. Note that, at each of the information sets in the sequence, the appraisal exhibits independence on S relative to G, due to plain consistency.

Let us conduct one more application of plain consistency, from h^M to immediate successor $\tilde{h}$, using $Z = S(\tilde{h})_G \times S(h^M)_{-G}$ and $Y = S(\tilde{h}) = S(\tilde{h})_G \times S(\tilde{h})_{-G}$. We obtain:

$$p^{\tilde{h}}(X_G \times S_{-G}) = \frac{p^{h^M}(X_G \times S_{-G})}{p^{h^M}(S(\tilde{h})_G \times S_{-G})} = \frac{\sigma_G(X_G)}{\sigma_G(S(\tilde{h})_G)} \tag{6}$$

for all $X_G \subset S(\tilde{h})_G$. Recall that $\hat{\sigma}$ and σ exhibit independence on S relative to G and satisfy $\hat{\sigma}_G = \sigma_G$. Recall as well that $\hat{\sigma}$ reaches the subgame for sure and that actions planned for unreached information sets outside of G are irrelevant for determining whether $\tilde{h}$ is reached. Using Equation 5, these facts imply that

$$\tilde{\sigma}_G(X_G) = \frac{\hat{\sigma}_G(X_G)}{\hat{\sigma}_G(S(\tilde{h})_G)} = \frac{\sigma_G(X_G)}{\sigma_G(S(\tilde{h})_G)}. \tag{7}$$

Combining Equations 6 and 7 yields $p_G^{\tilde{h}} = \tilde{\sigma}_G$.

We therefore can translate inequality $u_i(\tilde{\sigma}') > u_i(\tilde{\sigma})$ into $u_i(p^{\tilde{h}'}) > u_i(p^{\tilde{h}})$, where $\tilde{h}'$ is the continuation strategy $\tilde{h}$-deviation that plays according to $\hat{\sigma}'$ at player i's information sets starting from $\tilde{h}$. From the single-deviation property, this implies that P is not sequentially rational, which is the contradiction that completes the proof. ■

References

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

Blume, L., A. Brandenburger, and E. Dekel (1991a): "Lexicographic Probabilities and Choice Under Uncertainty," *Econometrica* 59 (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

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.

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.

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.

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

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.

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.

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.

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 (2015), "Sequential Equilibria of Multi-Stage Games with Infinite Sets of Types and Actions," manuscript, University of Chicago.

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

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.

Watson, J. (2017): "Mutual Perfect Bayesian Equilibrium," unpublished draft, UC San Diego.

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