



Australian School of Business Working Paper

Never Stand Still

Australian School of Business

Australian School of Business Research Paper No. 2012 ECON 02

Weakly-Bayesian and Consistent Assessments

Carlos Pimienta

This paper can be downloaded without charge from
The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1984332>

WEAKLY-BAYESIAN AND CONSISTENT ASSESSMENTS*

CARLOS PIMIANTA[†]

ABSTRACT. A *weakly-Bayesian assessment* is computed applying Bayes rule at positive probability information sets. We characterize the set of extensive-forms for which the sets of weakly-Bayesian and consistent assessments coincide. In doing so we disentangle the different restrictions imposed by consistency across information sets. We apply this knowledge to strengthen weakly-Bayesian assessments and to derive conditions for equivalence with consistency that can be useful in economic applications.

1. INTRODUCTION

A sequential equilibrium (Kreps and Wilson, 1982) is a sequentially rational *consistent assessment*. The notion of consistency incorporated in the definition of sequential equilibrium provides a way of selecting beliefs at zero probability information sets. Loosely speaking, consistent beliefs must admit an explanation consisting of “small trembles” made to reach those information sets.

There is a broad theoretical literature dealing with sequential equilibrium. This partly stems from the apparently ad-hoc procedure whereby consistency selects beliefs, which urged an effort to understand better the notion of consistency and its game theoretical implications. Battigalli (1996), Kohlberg and Reny (1997) and Swinkels (1993) show that consistency is related to the game theoretical principle of strategic independence. If different players choose their strategies independently then their assessments must be consistent.¹

[†] SCHOOL OF ECONOMICS, THE UNIVERSITY OF NEW SOUTH WALES, SYDNEY 2052, AUSTRALIA. PHONE: +61 2 9385 3358, FAX: +61 2 9313 6337.

E-mail address: c.pimienta@unsw.edu.au.

First version: November 23, 2009. *This version:* October 15, 2011.

JEL classification. C62, C72, D80.

Key words and phrases. Backwards induction, Weakly-Bayesian assessments, Consistent assessments, Sequential equilibrium, Extensive-forms.

* Preliminary discussions with Eric Maskin are gratefully acknowledged. This paper has benefited from numerous discussions with Cristian Litan. I thank Francesco De Sinopoli, Jeff Kline, Stephen Morris, Randy Silvers, Jeroen Swinkels, Bill Schworm and participants at the 5th Pan-Pacific Conference in Game Theory for helpful comments that greatly improved the presentation and content of the paper. This research has been economically supported by UNSW ASBRG 2009. The usual disclaimer applies.

¹ In fact they show that consistency is equivalent to a “strong” version of independence. See Kohlberg and Reny (1997) and Swinkels (1993) for different interpretations of this result.




FIGURE 1.

A number of papers also offer different characterizations of consistency and/or show that, under certain conditions, sequential equilibrium is equivalent to weaker equilibrium concepts. Fudenberg and Tirole (1991) define perfect Bayesian equilibrium imposing some intuitive restrictions on beliefs and show its equivalence to sequential equilibrium in multi-period games with observed actions. Perea y Monsuwé et al. (1997) provide an algebraic characterization of consistency without making use of trembles. Litan and Pimienta (2008) find the maximal class of extensive-forms such that sequential equilibrium and subgame perfection coincide in equilibrium strategies and equilibrium outcomes.

In this paper we look at those instances where consistency places no restrictions at zero probability information sets. This is the case of the *extensive-form* of Figure 1. If Player *I* moves *Out* then any belief at Player *II*'s information set is consistent as it can be justified by an appropriate sequence of trembles. A similar argument holds in the extensive-form of Figure 2. If players *I* and *II* play according to (r_1, r_2) then arbitrary beliefs at Player *III*'s information set are consistent. To generalize these ideas, we work with *weakly-Bayesian* assessments.² A weakly-Bayesian assessment imposes the only requirement that beliefs at positive probability information sets be computed from the strategy profile using Bayes' rule. Clearly, every consistent assessment is a weakly-Bayesian assessment and in general, not every weakly-Bayesian assessment is consistent. We characterize the set of extensive-forms such that every weakly-Bayesian assessments is consistent. Both in Figure 1 and Figure 2 the set of weakly-Bayesian and consistent assessments coincide.

It is not difficult to come up with examples of extensive-forms for which some weakly-Bayesian assessment is not consistent. The weakly-Bayesian assessment $(Out, l_2, r, \mu(x_2) = 1)$ is not consistent in the extensive-form of Figure 3. Consistent beliefs should place probability zero at the central node of Player *III*'s information set given that in a sequential equilibrium "correlation in defections are (partially) ruled out" (Kreps and Wilson, 1982, p. 875). That is, if Player *I* defects, it does

² Mas-Colell et al. (1995, Definition 9.C.3), among others, define *weak perfect Bayesian equilibrium*. In a weakly-Bayesian assessment we simply drop the sequentially rationality requirement from that definition. We do this to focus on belief formation and to facilitates the comparison with sequential equilibrium.




FIGURE 2.

not make a defection of Player *II* more likely. Figure 5 contains another example. Kreps and Wilson (1982, p. 876) explain how the consistency criterion invokes the “common knowledge” principle for beliefs. Hence, any assessment where Player *I* moves *Out* and players *II* and *III* assess different relative probabilities over their left-hand and right-hand nodes is not consistent.

This paper identifies the relevant characteristics shared by the extensive forms in figures 3, 5, and any other one where not every weakly-Bayesian assessment is consistent. Furthermore, we provide a characterization of the whole set of extensive forms where consistency imposes restrictions at zero probability information sets. While doing so, we disentangle the different restrictions imposed by consistency across information sets. This is useful not only to know under which conditions applying weakly-Bayesian assessments is not enough, but also to determine to which extent a concept that is more demanding than weakly-Bayesian closes the gap with respect to consistency.

Thus, as an application of the results, we also work with *preconsistent assessments*. A preconsistent assessment requires each player updates her beliefs even at zero probability parts of the extensive-form. Similarly to weakly-Bayesian assessments, preconsistent assessments are easy to compute and commonly used in economic applications. The results derived in this paper help describe a general class of extensive-forms where every restriction imposed by consistency is captured by preconsistent assessments. In addition, this allows us to arrive at the same




FIGURE 3.

results as Fudenberg and Tirole (1991) for multi-period games with observable actions where players have at most two types.

From a theoretical viewpoint, this paper can also help understand better how consistency brings about restriction in beliefs. While in some cases we may already have a very good understanding about how consistent beliefs are shaped, as it happens for instance when one information set comes after another like in Figure 5, in some other cases this relation may be more obscure or, at least, difficult to identify by arguments that are not context specific (see Figure 6). For this reason, a unifying explanation of the restrictions on beliefs entailed by consistency that is based solely on the characteristics of extensive-forms can be of theoretical interest.

In the next section we introduce the basic notation of extensive-form games and important definitions. Section 3 contains the results about equivalence between weakly-Bayesian and consistent assessments illustrated by a series of examples. Proofs of these results are offered in Section 4. Section 5 elaborates on the relationship between sequentially rational weakly-Bayesian assessments and sequential equilibria. To conclude, we apply what we have learned in the previous sections in Section 6. First, we strengthen weakly-Bayesian assessments defining preconsistent assessments. Then we derive a new result that establishes equivalence between this stronger concept and consistency. Throughout this section, we also describe several games where consistent beliefs can be computed by finding weakly-Bayesian or preconsistent assessments.

2. BASIC NOTATION AND DEFINITIONS

We start by describing notation and terminology for finite extensive-form games with perfect recall. For a full mathematical description of extensive-form games the reader is referred to Kreps and Wilson (1982). In what follows, for every two sets E and F , we use $E \subset F$ allowing for equality. As usual, $E \setminus F$ represents the set of elements in E that do not belong to F .

An *extensive-form* is a tuple $\Gamma = (\mathcal{N}, X, \prec, P, \mathcal{H}, \mathcal{A}, \lambda)$. The set of players is $\mathcal{N} = \{1, \dots, N\}$ and players are indexed by $n = 1, \dots, N$.

The finite set of nodes X is partially ordered by \prec . It contains a distinguished minimal element $x_o \in X$ called the root to the extensive form. The subset of final nodes is $Z \subset X$. The set $X \setminus Z$ is partitioned by the *player partition* $P = (P_0, P_1, \dots, P_N)$, where P_n represents the set of nodes where player n has to move (P_0 corresponds to the set of nodes where Nature moves).

The *information partition* $\mathcal{H} = (H_1, \dots, H_N)$ contains the information structure of the extensive form, where for each n , the collection H_n partitions P_n into *information sets* $h \in H_n$. An element $h \in H_n$ represents the set of nodes that player n cannot distinguish when she has to move at h . The information set that contains node x is denoted as $h(x)$. Furthermore, $H = \bigcup_n H_n$.

The set of *actions* in the extensive form is $\mathcal{A} = (A_0, A)$. For variety's sake, we use the terms *action*, *choice* and *move* interchangeably throughout the paper. The set of choices available to players is $A = \bigcup_{h \in H} A(h)$ where $A(h)$ represents the set of choices available at the information set h . The set of moves available to Nature is $A_0 = \bigcup_{p \in P_0} A_0(p)$ where $A_0(p)$ is the set of moves available to Nature at p . It will be convenient to write $A_n = \bigcup_{h \in H_n} A(h)$ for the set of choices available to player n across all information sets. Actions at different information sets are always labelled differently, that is, $A(h) \cap A(h') \neq \emptyset$ whenever $h \neq h'$. Furthermore, if choice a is taken at node x then the next node that follows is denoted (x, a) .

The vector λ contains the probability distributions over the moves of Nature by specifying for each $c \in A_0$ a number $\lambda(c) \in (0, 1)$ in such a way that $\sum_{c \in A_0(p)} \lambda(c) = 1$ for all $p \in P_0$. If Nature does not move at the root of the extensive-form, that is $x_o \notin P_0$, then we say that x_o is the unique initial node. If $x_o \in P_0$ then, with slight abuse of terminology, we say that each node that is only preceded by moves of Nature is an *initial* node. We explicitly allow that Nature moves at any other part of the extensive-form.

An extensive-form game is obtained from an extensive-form by specifying for each player n a *Bernoullian utility function* $u_n : Z \rightarrow \mathbb{R}$. Our characterizations are based on properties of the extensive-form.

In order to work with the space of extensive-forms we need to introduce concepts that summarize parts of their structure. Given any node x , there is a unique collection of choices (including those of Nature) that from the root of the extensive-form lead to that node. That set of choices is called *path to node* x and it is denoted by $\mathcal{P}(x)$. The subset of the path to node x made of by actions of players only is $\mathcal{P}_A(x) = \mathcal{P}(x) \setminus A_0$. Any path can be totally ordered by $<$. A *carrier* is any subset of choices C that satisfies $(C \cap A(h)) \neq \emptyset$ for all $h \in H$. That is, a carrier contains at least one action of each information set. The usage of the term ‘‘carrier’’ for a set of choices satisfying these properties is justified below.

We only consider extensive-forms with perfect recall. Whenever a player moves she remembers all the choices that she has taken in the past as well as the information that she knew before. In symbols, for any two nodes $x, x' \in h$ in an information set $h \in H_n$ that belongs to player n the inclusion $c \in (\mathcal{P}_A(x) \cap A_n)$ implies $c \in \mathcal{P}_A(x')$. Perfect recall implies that H_n is totally ordered by $<$.

For an arbitrary set of choices of players $B \subset A$, we say that B *reaches* node x if $\mathcal{P}_A(x) \subset B$. Likewise, B *reaches* the information set h if $\mathcal{P}_A(x) \subset B$ for some node $x \in h$. Furthermore, we let $X^+(C)$ represent the set of decision nodes that are reached by the carrier C and let $X^0(C) = X \setminus X^+(C)$ represent its complement.

A *pure strategy* s_n of player n is a plan of action that specifies, for each information set $h \in H_n$, one choice $s_n(h) \in A(h)$. The set of player n 's pure strategies is S_n and the set of pure strategy profiles is $S = S_1 \times \cdots \times S_N$. The *carrier of*

a pure strategy profile s is $\mathcal{C}(s) = \bigcup_{n \in \mathcal{N}} \bigcup_{h \in H_n} s_n(h)$. We write $S(x)$ and $S(h)$ to denote the set of pure strategy profiles whose carriers reach, respectively, node x and information set h . The sets $S_n(x)$ and $S_{-n}(x)$ are the projections of $S(x)$ on S_n and $S_{-n} = \prod_{m \neq n} S_m$.

A behavioral strategy σ_n of player n specifies for every information set $h \in H_n$ a probability distribution $\sigma_n(\cdot | h)$ on $A(h)$. The probability that player n chooses $a \in A(h)$ is, therefore, $\sigma_n(a | h)$. A behavioral strategy profile σ specifies for every information set $h \in H$ a probability distribution on $A(h)$. The set of behavioral strategies of player n is denoted Σ_n and the set of behavioral strategy profiles $\Sigma = \Sigma_1 \times \cdots \times \Sigma_N$. The carrier of a behavioral strategy profile σ , denoted $\mathcal{C}(\sigma)$, is the union over all information sets h of the choices $a \in A$ that satisfy $\sigma(a | h) > 0$.

Every behavioral strategy profile σ induces, together with λ , a probability distribution $\mathbb{P}(\cdot | \sigma)$ on Z . Given an arbitrary subset of nodes $Y \subset X$ we let $Z(Y)$ denote the subset of final nodes $z \in Z$ that satisfy $y < z$ for some $y \in Y$. Furthermore, we write $\mathbb{P}(Y | \sigma)$ instead of $\mathbb{P}(Z(Y) | \sigma)$.

A system of beliefs μ specifies for every information set h a probability distribution $\mu(\cdot | h)$ over its nodes. An assessment is a behavioral strategy profile together with a system of beliefs (σ, μ) .

We now introduce our two objects of study.

Definition 1 (Consistent Assessments). The assessment (σ, μ) is *consistent* if it is the limit point of a sequence $\{(\sigma^t, \mu^t)\}_{t=0}^\infty$ such that, for all t , σ^t is completely mixed (i.e. $\sigma^t(a | h) > 0$ for all $h \in H$ and all $a \in A(h)$) and

$$\mu^t(x | h) = \frac{\mathbb{P}(x | \sigma^t)}{\mathbb{P}(h | \sigma^t)} \quad \text{for every } h \in H \text{ and every } x \in h.$$

Definition 2 (weakly-Bayesian Assessments). The assessment (σ, μ) is a *weakly-Bayesian assessment* if for every $h \subset X^+(\mathcal{C}(\sigma))$ and every $x \in h$

$$\mu(x | h) = \frac{\mathbb{P}(x | \sigma)}{\mathbb{P}(h | \sigma)}.$$

Of course, every consistent assessment is weakly-Bayesian but the converse is not true.

3. NON-CONSISTENT WEAKLY-BAYESIAN ASSESSMENTS

In this section we characterize the set of extensive-forms such that the set of consistent assessments is a strict subset of the set of weakly-Bayesian assessments. This is done in propositions 1 and 3. Theorem 1 will later assert that in the complement of the set laid out by the propositions the sets of weakly-Bayesian and consistent assessments coincide.

To provide a more clear intuition about the results we introduce relative probabilities over the set S of pure strategy profiles.³ A relative probability on S specifies the relative weight of each subset of pure strategy profiles with respect to any other subset. This includes subsets having prior probability equal to zero. A relative probability ρ on S must satisfy the following properties: for every subset $Q \subset S$ and all nonempty subsets $R, T \subset S$,

- (i) $\rho(Q, R) \in [0, \infty]$,
- (ii) $\rho(Q, Q) = 1$,
- (iii) $\rho(Q, T) + \rho(R, T) = \rho(Q \cup R, T)$ if $Q \cap R = \emptyset$, and
- (iv) $\rho(Q, T) = \rho(Q, R)\rho(R, T)$, whenever the product does not involve both 0 and ∞ .

Standard prior probabilities are therefore given by $\rho(\cdot, S)$.

Battigalli (1996) and Kohlberg and Reny (1997) show that every consistent assessment can be generated, in a way specified below, by a relative probability defined over the set of pure strategy profiles and satisfying a strong independence property. Strong independence implies weak independence and, for our purposes, the latter concept is restrictive enough.

The relative probability ρ defined on the set of pure strategies S is weakly independent if for every subset of players $M \subset \mathcal{N}$ and every two pairs of subsets of strategy profiles $Q_M, R_M \subset \prod_{n \in M} S_n$ and $Q_{-M}, R_{-M} \subset \prod_{n \in \mathcal{N} \setminus M} S_n$,⁴

$$\rho(Q_{-M} \times Q_M, Q_{-M} \times R_M) = \rho(R_{-M} \times Q_M, R_{-M} \times R_M). \quad (3.1)$$

A consistent assessment (σ, μ) can be generated by a relative probability ρ satisfying weak independence according to:^{5,6}

$$\sigma(a | h) = \rho(S(x, a), S(x)) \text{ for any } x \in h; \text{ and}, \quad (3.2)$$

$$\mu(x | h) = \rho(S(x), S(h)). \quad (3.3)$$

³ Relative probabilities are equivalent *conditional probability systems* (Myerson, 1986). In game theory conditional probability systems arise naturally from the need of specifying probabilities conditional on events that have prior probability zero. Among others, conditional probability systems have been studied by Battigalli (1996); Blume et al. (1991); Hammond (1994); Kohlberg and Reny (1997); McLennan (1989a,b); Myerson (1986); and Swinkels (1993).

⁴ Swinkels (1993) calls this condition *quasi-independence*. It is the (weak) notion of independence considered by Battigalli (1996). Swinkels uses the term *individual quasi-independence* for the analogous condition where M is substituted by just one player. The latter is the independence condition used by Kohlberg and Reny (1997). Quasi-independence implies individual quasi-independence but the opposite is not true. Swinkels (1993) offers an example, credited to Myerson, at that effect.

⁵ The converse is not true, i.e. not every weakly independent relative probability system generates a consistent assessment. See Kohlberg and Reny (1997) for an example.

⁶ Strictly speaking this only holds when the extensive-form does not contain moves of Nature. If it does, as we allow here, beliefs in Equation (3.3) need to be modified so that $\mu(x | h) = \rho(S(x), S(h)) \left(\prod_{c \in \mathcal{P}_A(x) \cap A_0} \lambda(c) \right) \left(\sum_{x' \in h} \prod_{c \in \mathcal{P}_A(x') \cap A_0} \lambda(c) \right)^{-1}$.

It can be shown that perfect recall and weak independence imply that (3.2) is well defined (i.e. it does not depend on the node $x \in h$ that is used).

We are going to derive a condition that implies restrictions on consistent beliefs at zero probability information sets. Recall that $\mathcal{P}(x)$ denotes the set of actions that form the path to node x and that $\mathcal{P}_A(x)$ is obtained from $\mathcal{P}(x)$ by removing the moves of Nature. Consider two nodes x and y (not necessarily in the same information set) and an action $a \in (\mathcal{P}_A(x) \cap \mathcal{P}_A(y))$. Let $\hat{\sigma}$ be a behavioral strategy profile that takes action a with probability zero and every other action in $\mathcal{P}_A(x)$ with positive probability. Even though $\hat{\sigma}$ does not reach either x or y , intuitively, node y cannot be *infinitely more likely* than node x . Let us now offer a more formal argument.

Let ρ be an independent relative probability defined on S that induces the consistent assessment $(\hat{\sigma}, \mu)$. We want to show that $\rho(S(x), S(y)) > 0$, that is, that (the set of strategy profiles that lead to) node y cannot be infinitely more likely than node x . Let x' be the unique node that satisfies $\mathcal{P}(x', a) \subset \mathcal{P}(x)$ and let y' be the unique node that satisfies $\mathcal{P}(y', a) \subset \mathcal{P}(y)$. Assume that action $a \in A(h')$ is available at the information set h' and that player n moves at h' . By property (iii) of relative probabilities we can find the value of $\rho(S(x), S(y))$ through

$$\rho(S(x), S(y)) = \rho(S(x), S(x', a)) \rho(S(x', a), S(y', a)) \rho(S(y', a), S(y)).$$

We obtain $\rho(S(x), S(x', a)) > 0$ because every choice in $\mathcal{P}(x)$ that follows a receives positive probability. It also holds that $\rho(S(y', a), S(y)) \geq 1$ because $S(y)$ is a subset of $S(y', a)$. In addition, from strategic independence it follows the equality $\rho(S(x', a), S(y', a)) = \rho(S(x', a'), S(y', a'))$ for every $a' \in A(h')$ because we are changing the same strategies of player n in both sides. Choose some $a' \in A(h')$ that satisfies $0 < \rho(S(x'), S(x', a')) \leq 1$. The value of $\rho(S(x', a'), S(y', a'))$ equals

$$\rho(S(x', a'), S(x')) \rho(S(x'), S(y')) \rho(S(y'), S(y', a')) = \rho(S(x'), S(y')),$$

where the last equality follows from the fact that weak independence makes (3.2) well defined. To conclude, $\rho(S(x'), S(y')) > 0$ because every choice in $\mathcal{P}(x)$ that precedes a receives positive probability. Therefore, $\rho(S(x), S(y)) > 0$.

Now, if x and y belong to the same information set h then $S(x)$ and $S(y)$ are both subsets of $S(h)$ and we obtain $\rho(S(y), S(h)) < 1$. Furthermore, if h is a zero probability information set according to $\hat{\sigma}$, i.e. $h \subset X^0(\mathcal{C}(\hat{\sigma}))$, then the last inequality and (3.3) imply $\mu(y | h) < 1$.

Consider again the extensive-form of the game in Figure 3. The leftmost node and the central node in Player III's information set have a common choice, action l_1 , in their respective paths. If players I and II play according to (Out, l_2) then (Out, l_2) is infinitely more likely than (Out, m) . We can use independence to conclude that the profile (l_1, l_2) is infinitely more likely than (l_1, m) , i.e. the leftmost




FIGURE 4.

node in Player *III*'s information set is infinitely more likely than the central node. Since Player *I* plays *Out*, Player *III*'s information set is reached with probability zero and we need to specify beliefs at her information set. It follows that consistent beliefs must assign probability zero to that central node. Of course, this restriction does not apply to weakly-Bayesian assessments.

We start our characterization with a sufficient condition for an extensive-form to admit weakly-Bayesian assessments that are not consistent. It corresponds to the set of properties suggested above and in our analysis of Figure 3.

Proposition 1. *Consider an extensive-form where we can find an information set h with two distinct nodes $x, y \in h$, a carrier C , and an action $c \in A$ such that:*

- (i) *the carrier C does not reach h ,*
- (ii) *$c \in (\mathcal{P}_A(x) \cap \mathcal{P}_A(y)) \setminus C$.*

Then the set of consistent assessments is strictly contained in the set of weakly-Bayesian assessments.

Condition (i) above is clear. The sets of weakly-Bayesian and consistent assessments may differ only if there exists some strategy profile that reaches some information set with probability zero.

Condition (ii) captures the relevant features in Figure 3. To understand better why we need an action $c \in (\mathcal{P}_A(x) \cap \mathcal{P}_A(y))$ observe that in figures 1 and 2, where every weakly-Bayesian assessment is consistent, we cannot find two nodes in the same information set that have a common action in their respective paths. To see why we need $c \notin C$, consider the extensive-form of Figure 4 and a behavioral strategy profile where choices *l* and *Out* are taken with probability one. There is no restriction on how Player *III* should form her consistent beliefs. (Every conceivable belief vector at that information set is the limit of a sequence of conditional probabilities generated by an appropriately chosen sequence of trembles.) In this case, for any system of beliefs μ , the resulting assessment can be associated to a well defined independent relative probability system on the set of pure strategy profiles—and every weakly-Bayesian assessment is consistent.

Figure 5 is another example where the set of consistent assessment is a strict subset of the set of weakly-Bayesian assessments. Consistency implies common




FIGURE 5.

knowledge of beliefs. This means that Player *II* and Player *III* must have the same belief over their left-hand and right-hand nodes and that, consequently, not every weakly-Bayesian assessment is consistent. In order to see this in terms of Proposition 1 note that action r_2 belongs to the path of the two nodes in Player *III*'s information set. In Figure 5, moreover, Player *II*'s information set does admit arbitrary beliefs, but once those are fixed beliefs at her second information set are determined. This suggests that we should explore further how the values assumed by consistent beliefs at two different information sets relate to each other.

We start this analysis studying signaling games, where it is well known that sequential equilibrium does not impose restrictions on beliefs.⁷ Once Nature moves, the sender observes her type and sends a signal to the receiver. If the receiver observes a signal which is sent in equilibrium, she applies Bayes' rule to derive her beliefs about the type of the sender. If a signal is not sent in equilibrium, the receiver has no restrictions whatsoever on how to form her beliefs about the type of the sender upon receiving that signal. Note well that every pair of nodes that belong to the same information set have completely different paths from each other (the same signal can be sent from two different information sets so they are, in fact, different actions).

Consider the extensive-form of a slightly modified signaling game in Figure 6. After the sender learns her type, and before she sends a signal, she can end the game. As it is the case in a standard signaling game, no pair of nodes that belong to the same information set have a common action in their respective paths. Nonetheless, there are players' actions that are common to the paths to nodes that belong to different information sets. That is, $a_1 \in (\mathcal{P}_A(x_1) \cap \mathcal{P}_A(y_2))$ and $a_2 \in (\mathcal{P}_A(x_2) \cap \mathcal{P}_A(y_1))$. Consider the behavioral strategy profile $(f_1, f_2, u_1, d_2, r_u, l_d)$ which assigns probability zero to a_1 and a_2 . As mentioned previously, if a strategy profile assigns positive probability to all the actions leading to a node but one,

⁷ See, for instance, Kohlberg (1990).




FIGURE 6.

which is also in the path to a second node, then the underlying independent relative probability must consider the (set of pure strategy profiles leading to the) first node as infinitely more likely than the second node. In terms of the example this means $\rho(S(x_1), S(y_2)) = \rho(S(x_2), S(y_1)) = \infty$ which in turn implies that we cannot have $\rho(S(y_1), S(x_1)) = \rho(S(y_2), S(x_2)) = \infty$ because otherwise a node must be infinitely more likely than itself. From this argument we obtain that a consistent assessment with a behavioral strategy profile as above cannot display beliefs such that $\mu(y_1 | h_1) = \mu(y_2 | h_2) = 1$. Obviously, this restriction does not apply to weakly-Bayesian assessments.

The relevant features of the previous example are generalized in the next proposition into a new sufficient condition on extensive-forms so that the set of consistent assessments is a strict subset of the set of weakly-Bayesian assessments.

Proposition 2. *Consider an extensive-form where we can find two distinct information sets h_1, h_2 , two distinct nodes $x_1, y_1 \in h_1$ in the first information set, two distinct nodes $x_2, y_2 \in h_2$ in the second, a carrier C , and two choices $c_1, c_2 \in A$ such that:*

- (i) *the carrier C reaches neither h_1 nor h_2 ; and*
- (ii) *$c_1 \in (\mathcal{P}_A(x_1) \cap \mathcal{P}_A(y_2)) \setminus C$ and $c_2 \in (\mathcal{P}_A(x_2) \cap \mathcal{P}_A(y_1)) \setminus C$.*

Then the set of consistent assessments is strictly contained in the set of weakly-Bayesian assessments.

Remark 1. • If $h_1 = h_2$ then we can apply Proposition 1. However, it is important that $x_1 \neq y_1$ and $x_2 \neq y_2$. Otherwise we cannot guarantee that some weakly-Bayesian assessments is not consistent.

- The existence of two different actions c_1 and c_2 in this proposition is not dispensable. Suppose that in the extensive-form of Figure 6 we delete action a_2 and replace the two consecutive information sets of Player I by a single information set where Player I has available the three choices f_2 ,




FIGURE 7.

u_2 and d_2 . With this modification every weakly-Bayesian assessment is consistent.

In our signaling game the two zero probability information sets do not come one after another as it occurs, for instance, in the extensive-form of Figure 5. We have seen that this extensive-form satisfies the conditions of Proposition 1, therefore, we already know that some weakly-Bayesian assessments are not consistent. We can now see that this extensive-form also satisfies the conditions of Proposition 2. Indeed, if Player *I* moves *Out*, then Player *II*'s information set and Player *III*'s information set receive probability zero. The left-hand node in *II*'s information set and the left-hand node in *III*'s information set have action l_2 in their respective paths. The analogous is true for the right-hand nodes and action r_2 . Nevertheless, note that if moving *Out* was not a possible action the resulting extensive-form would satisfy the conditions of Proposition 1, but not the conditions of Proposition 2.

Similar arguments to those in our previous two examples are also valid when three or more information sets are involved. Figure 7 illustrates this with three information sets. Suppose that players *I* and *II* play according to (Out, r_1, r_2, r_3) . We must specify beliefs at the three information sets of Player *III*. The following pairs of nodes have an action in their respective paths that the previous profile attaches probability zero to: (x_1, y_2) , (x_2, y_3) and (x_3, y_1) . Weak independence implies that, for each of these pairs, the pure strategy profile leading to the first node is infinitely more likely than the pure strategy profile leading to the second node. It follows that a system of beliefs such that $\mu(y_1 | h_1) = \mu(y_2 | h_2) = \mu(y_3 | h_3) = 1$ is not consistent. Again, this restriction does not apply to weakly-Bayesian assessments.

The extensive-form of Figure 7 is such that no two information sets receiving probability zero come one after another. Figure 8 contains an extensive-form where for some information sets this is the case. For reasons analogous to those discussed




FIGURE 8.

in the previous examples, the assessment $(Out, l_1, l_2, \mu(x_1 | h_1) = 1, \mu(x_2 | h_2) = 1, \mu(x_3 | h_3) = 1, \mu(x_4 | h_4) = 1)$ is not consistent.⁸

The general result, which subsumes Proposition 2, is the following:⁹

Proposition 3. *Consider an extensive-form where we can find $K \geq 2$ distinct information sets h_1, \dots, h_K , two distinct nodes $x_i, y_i \in h_i$ for each $i = 1, \dots, K$, a carrier C , and K distinct actions c_1, \dots, c_K such that:*

- (i) *for each $i = 1, \dots, K$ the carrier C does not reach h_i ; and*
- (ii) *for each $i = 1, \dots, K - 1$ we have $c_i \in (\mathcal{P}_A(x_i) \cap \mathcal{P}_A(y_{i+1})) \setminus C$, likewise, $c_K \in (\mathcal{P}_A(x_K) \cap \mathcal{P}_A(y_1)) \setminus C$.*

Then the set of consistent assessments is strictly contained in the set of weakly-Bayesian assessments.

Remark 2. • Proposition 2 corresponds to $K = 2$.

- It is embedded in the statement of the proposition that we have to find K information sets *and an order* of those information sets such that the conditions are true. The conditions will not typically hold for every possible order.
- Again, it is important that the $2K$ nodes be distinct. However, if some of the K information sets are not different it would only mean that the conditions are satisfied for an integer strictly smaller than K .
- If the conditions are satisfied for some value of K it does not follow that they are also satisfied for some integer smaller than K . See for instance,

⁸ One difference is that in this example $S(x_1)$ and $S(y_2)$ are of the same “order of magnitude”, i.e. neither set is infinitely more likely than the other because l_1 is chosen by the strategy profile with probability one. The same is true for $S(x_3)$ and $S(y_4)$.

⁹ Nonetheless, we have presented propositions 2 and 3 separately to facilitate the exposition of the results. On a technical note, the proof of the Proposition 3 is done by induction (see page 19), and Proposition 2 corresponds to the initial step.

Figure 7 where $K = 3$ and Figure 8 where $K = 4$. In both cases the conditions do not hold for smaller values of K .

Theorem 1 states that the conditions in propositions 1 and 3 are not only sufficient but also necessary.

Theorem 1. *Consider an extensive-form where some weakly-Bayesian assessment is not consistent. Then the extensive-form satisfies either the conditions listed in Proposition 1 or the conditions listed in Proposition 3.*

Intuitively, for a fixed behavioral strategy profile, if a zero probability choice is in the path to two different nodes then we are losing freedom to choose consistent beliefs because the same tremble is associated to two different nodes. The fact that the conditions in propositions 1 and 3 are not satisfied means that we always have enough freedom so that arbitrary beliefs are consistent.

Take a behavioral strategy profile σ and two nodes x and y . Suppose that there is at least one choice in the path to x that is not in the path to y and at least one choice in the path to y that is not in the path to x . Suppose further that none of these two choices is taken with positive probability under σ . Based only on this information about σ and the structure of the extensive-form we cannot derive a definite likelihood ordering between nodes x and y . Recurring to an argument based on trembles, we can fine-tune the trembles associated to the two actions pinpointed before to make one of the sets of trembles necessary to reach one of the nodes as likely as we want with respect to the other. If x and y are the only two nodes in the same information set h and this is reached with probability zero under σ then we can freely choose consistent beliefs at h .

If the conditions in Proposition 1 are not satisfied the above argument holds even if we have more than two nodes in h and for any σ . So consistent beliefs at h can be freely chosen. But this choice of beliefs could, in principle, constrain the set of consistent beliefs available for a second zero probability information set h' . This can happen when one of the choices whose tremble we fine-tuned before is in the path to some node in h' . If the conditions in Proposition 3 are not satisfied for $K = 2$ then for each of the remaining nodes in that information set we can again fine-tune some tremble associated to a choice that is only in the path to that node and none else in h' . The argument can be repeated again for a third zero probability information set h'' but if the conditions in Proposition 3 are not satisfied for $K = 3$ we do not have restrictions in choosing beliefs at h'' . Continuing in this fashion we can pick arbitrary beliefs at every zero probability information set.

4. PROOFS

The main tool used in the proofs of Propositions 1, 2 and 3 is Lemma A1 in Kreps and Wilson (1982). But before stating that lemma, a few concepts are necessary.

Given a system of beliefs μ , its *support* $\text{supp}(\mu)$ is the union over all information sets h of the decision nodes x for which $\mu(x | h) > 0$. We reserve the term *carrier* to talk about the set of choices that receive strictly positive probability for probabilities defined by strategy profiles, and the term *support* for the analogous concept for probabilities associated to beliefs. Moreover, we write $\mathcal{C}(\Sigma)$ for the set of all possible carriers.

A *labelling* is a function $L : A \rightarrow \mathbb{N}$ that maps each choice $c \in A$ into an integer number $L(c)$. For each labelling L there is an associated function $F_L : X \rightarrow \mathbb{N}$ defined by

$$F_L(x) = \sum_{c \in \mathcal{P}_A(x)} L(c).$$

Definition 3. Given a carrier C , the labelling L is said to be a C -labelling if we have $c \in C$ if and only if $L(c) = 0$.

A *basis* (C, Y) is a subset of $A \times X$. We say that the basis (C, Y) is consistent if there exists at least one consistent assessment (σ, μ) such that $\mathcal{C}(\sigma) = C$ and $\text{supp}(\mu) = Y$.

Lemma 1 (Kreps and Wilson (1982, Lemma A1)). *The basis (C, Y) is consistent if and only if there is a C -labelling L such that the following condition holds:*

$$x \in Y \text{ if and only if } x \text{ minimizes } F_L(\cdot) \text{ on } h(x).$$

For our purposes Lemma 1 implies the following. Take a behavioral strategy profile σ that does not reach the information set h and suppose that nodes x and y belong to h . If $F_L(x) \leq F_L(y)$ for every conceivable $\mathcal{C}(\sigma)$ -labelling L then a necessary condition for (σ, μ) to be consistent is that $\mu(y | h) \neq 1$. In order to prove Proposition 1 we show that, if the extensive-form meets the conditions given in the proposition, we can always find such a strategy profile or, more precisely, such a carrier.

For a behavioral strategy profile σ with carrier C , we are thus interested in inequalities of the form $F_L(x) \leq F_L(y)$ that remain true for every C -labelling L . It will be useful to write $F(x, C) \leq F(y, C)$ when this is the case. For instance, if node y comes after x we can readily conclude that $F(x, C) \leq F(y, C)$. The expression $F(x, C) \leq F(y, C)$ can be meaningfully read as “node y cannot be infinitely more likely than node x under any strategy profile with carrier C .” If x and y belong to the same zero probability information set this must be respected by consistent beliefs.

To avoid duplication of arguments we provide some results in the next lemma that are used continuously throughout the proofs.

Lemma 2. *Consider a carrier C , an information set $h \subset X^0(C)$, and two nodes x and y with a common action $c \in (\mathcal{P}_A(x) \cap \mathcal{P}_A(y)) \setminus C$. Let $C' = C \cup (\mathcal{P}_A(x) \setminus \{c\})$, the following holds:*

- (i) $0 < F(x, C') \leq F(y, C')$;
- (ii) *if $(h \cap X^+(C')) \neq \emptyset$ then there is at least one node $\hat{y} \in h$ that satisfies $\mathcal{P}_A(\hat{y}) \subset (C \cup \mathcal{P}_A(x))$; and moreover*
- (iii) $(\mathcal{P}_A(x) \cap \mathcal{P}_A(\hat{y})) \setminus C \neq \emptyset$, i.e. $0 < F(\hat{y}, C) < F(x, C)$.

Proof. Take a carrier C that does not contain some action $c \in (\mathcal{P}_A(x) \cap \mathcal{P}_A(y))$. Suppose that under that carrier $h \subset X^0(C)$. (Note that x and y belong to $X^0(C)$ because $c \notin C$ and that we do not necessarily assume $x, y \in h$.) We obtain a new carrier C' by adding to C all the choices in the path to x except for c , therefore, we also obtain that x and y belong to $X^0(C')$.

To prove part (i) take any C' -labelling L . Action c is the only element in $\mathcal{P}_A(x)$ that is not in C' . Hence, $F_L(x) = L(c) > 0$ and since c is also in $\mathcal{P}_A(y)$ we obtain $F_L(y) \geq L(c)$. That is, node y cannot be infinitely more likely than x under a strategy with carrier C' .

Let us turn to part (ii). Let \hat{y} be some node in h such that $\hat{y} \in X^+(C')$. We have $\mathcal{P}_A(\hat{y}) \subset C' = C \cup (\mathcal{P}_A(x) \setminus \{c\}) \subset (C \cup \mathcal{P}_A(x))$.

Part (iii) follows from $\hat{y} \in X^0(C)$ and $\hat{y} \in X^+(C')$, which means that the actions in the path to \hat{y} that are not in C are in the path to x . It also follows that $0 < F(\hat{y}, C) < F(x, C)$, where the second inequality is strict because $c \in \mathcal{P}_A(x)$ but $c \notin \mathcal{P}_A(\hat{y})$. \square

We can now prove the first proposition.

Proof of Proposition 1. Recall that $h(x)$ represents the information set that contains node x . The conditions listed in the proposition are equivalent to the following set being nonempty.

$$\Phi_1 = \left\{ (x, y, c, C) \in X^2 \times A \times \mathcal{C}(\Sigma) : y \in h(x), y \neq x, h(x) \subset X^0(C), \right. \\ \left. c \in (\mathcal{P}_A(x) \cap \mathcal{P}_A(y)) \setminus C \right\}.$$

Take an arbitrary element $(x, y, c, C) \in \Phi_1$ and construct the carrier

$$\hat{C} = C \cup (\mathcal{P}_A(x) \setminus \{c\}).$$

Using Lemma 2 (i) we obtain that $F(x, \hat{C}) \leq F(y, \hat{C})$. If $h(x) \subset X^0(\hat{C})$ then the desired result follows. If otherwise some $\hat{y} \in h(x)$ satisfies $\hat{y} \in X^+(\hat{C})$ then Lemma 2 (iii) implies $F(\hat{y}, C) < F(x, C)$. Since $h(x) \subset X^0(C)$ this concludes the proof. \square

The argument behind the proof of Proposition 2 is similar but slightly more involved because we have to deal with nodes in two different information sets. Given two information sets h_1 and h_2 , we want to use the structure of the extensive-form to find a carrier C and four nodes $x_1, y_1 \in h_1, x_2, y_2 \in h_2$ such that $F(x_1, C) \leq F(y_2, C)$ and $F(x_2, C) \leq F(y_1, C)$. If h_1 and h_2 are subsets of $X^0(C)$ then consistency implies that for every assessment (σ, μ) with $\mathcal{C}(\sigma) = C$ beliefs cannot be such that $\mu(y_1 | h_1) = \mu(y_2 | h_2) = 1$. These equalities would imply that for some C -labelling L it holds $F_L(y_1) < F_L(x_1)$ and $F_L(y_2) < F_L(x_2)$. But since we also have $F_L(x_1) \leq F_L(y_2)$ and $F_L(x_2) \leq F_L(y_1)$ we reach a contradiction. The main difficulty of the argument consists of showing that the information sets h_1 and h_2 are reached with probability zero. Lemma 2 will be of great help at this effect.

Proof of Proposition 2. We need to show that, if the conditions in Proposition 2 are satisfied, not every weakly consistent assessment is consistent. We already proved that this is the case whenever $\Phi_1 \neq \emptyset$, so it is enough that we prove the result for the case $\Phi_1 = \emptyset$. The conditions in the proposition imply that the following set is nonempty:

$$\begin{aligned} \Phi_2 = \{ & (x_1, y_1, x_2, y_2, c_1, c_2, C) \in X^4 \times A^2 \times \mathcal{C}(\Sigma) : h(x_1) \neq h(x_2), \\ & y_1 \in h(x_1), y_1 \neq x_1, y_2 \in h(x_2), y_2 \neq x_2, h(x_1) \subset X^0(C), h(x_2) \subset X^0(C), \\ & c_1 \in (\mathcal{P}_A(x_1) \cap \mathcal{P}_A(y_2)) \setminus C, c_2 \in (\mathcal{P}_A(x_2) \cap \mathcal{P}_A(y_1)) \setminus C \}. \end{aligned}$$

Take any $(x_1, y_1, x_2, y_2, c_1, c_2, C) \in \Phi_2$. Let $h_1 = h(x_1)$ and $h_2 = h(x_2)$. Construct the carrier

$$\hat{C} = C \cup (\mathcal{P}_A(x_1) \setminus \{c_1\}) \cup (\mathcal{P}_A(x_2) \setminus \{c_2\}).$$

From Lemma 2 (i) we obtain $F(x_1, \hat{C}) \leq F(y_2, \hat{C})$ and $F(x_2, \hat{C}) \leq F(y_1, \hat{C})$. If both h_1 and h_2 are contained in $X^0(\hat{C})$ then no consistent assessment (σ, μ) with $\mathcal{C}(\sigma) = \hat{C}$ satisfies $\mu(y_1 | h_1) = \mu(y_2 | h_2) = 1$. So we need to prove the result for the cases where either $(h_1 \cap X^+(\hat{C})) \neq \emptyset$ or $(h_2 \cap X^+(\hat{C})) \neq \emptyset$.

Thus, let us assume that $(h_2 \cap X^+(\hat{C})) \neq \emptyset$ (the other case is analogous) and construct the carrier

$$\hat{C}_1 = C \cup (\mathcal{P}_A(x_1) \setminus \{c_1\}).$$

Now we show that $h_1 \subset X^0(\hat{C}_1)$. Suppose to the contrary that $(h_1 \cap X^+(\hat{C}_1)) \neq \emptyset$. Lemma 2 implies that some node in h_1 (that is not x_1) and x_1 have a common action in their respective paths that is not contained in C . This in turn would imply that the set Φ_1 is not empty, but this is the case that we assumed at the beginning of the proof. We can conclude $h_1 \subset X^0(\hat{C}_1)$.

We also know that $c_2 \notin \hat{C}_1$, for otherwise c_2 would be a common choice between x_1 and y_1 . Furthermore, $c_1 \notin \hat{C}_1$ by construction. It follows that $x_2, y_2 \in X^0(\hat{C}_1)$. Since $(h_2 \cap X^+(\hat{C})) \neq \emptyset$ and we are assuming $\Phi_1 = \emptyset$ we know that there is a subset

of nodes $\hat{h}_2 \subset h_2$ that does not include x_2 nor y_2 and that satisfies $\hat{h}_2 \subset X^+(\hat{C}_1)$. From Lemma 2 (iii) we obtain $F(\hat{y}_2, C) < F(x_1, C)$ for each $\hat{y}_2 \in \hat{h}_2$.

Construct the new carrier

$$\hat{C}'_1 = C \cup (\mathcal{P}_A(y_1) \setminus \{c_2\}).$$

By Lemma 2 (i), $F(y_1, \hat{C}'_1) \leq F(x_2, \hat{C}'_1)$. Using the same arguments as before we can show that $h_1 \subset X^0(\hat{C}'_1)$ and $x_2, y_2 \in X^0(\hat{C}'_1)$. If $h_2 \subset X^0(\hat{C}'_1)$ then the desired result follows because $F(\hat{y}_2, C) < F(x_1, C)$ keeps holding when we change C for \hat{C}'_1 . That is, a consistent assessment does not satisfy $\mu(x_1 | h_1) = \mu(x_2 | h_2) = 1$.

Hence, the next case we need to explore is when a subset of nodes $\hat{h}'_2 \subset h_2$ exists such that $\hat{h}'_2 \subset X^+(\hat{C}'_1)$. We construct a new carrier using the set of choices:

$$\tilde{B} = \left\{ \mathcal{P}_A(x_1) \setminus \left(\bigcup_{\hat{y}_2 \in \hat{h}_2} \mathcal{P}_A(\hat{y}_2) \right) \right\} \cup \left\{ \mathcal{P}_A(y_1) \setminus \left(\bigcup_{\hat{x}_2 \in \hat{h}'_2} \mathcal{P}_A(\hat{x}_2) \right) \right\}.$$

Let the new carrier be:

$$\tilde{C} = C \cup (\tilde{B} \setminus \{c_1, c_2\}).$$

We still obtain $h_1 \subset X^0(\tilde{C})$ from Lemma 2 (ii) and the assumption $\Phi_1 = \emptyset$. We obtain $h_2 \subset X^0(\tilde{C})$ by construction. The last carrier that we consider is:

$$C^* = \begin{cases} \tilde{C} \cup (\mathcal{P}_A(x_2) \setminus \{c_2\}) & \text{if } |\hat{h}_2| \leq |\hat{h}'_2|; \\ \tilde{C} \cup (\mathcal{P}_A(y_2) \setminus \{c_1\}) & \text{if } |\hat{h}_2| > |\hat{h}'_2|. \end{cases}$$

In either case, the information set h_2 is included in $X^0(C^*)$ because Φ_1 is empty. Regarding h_1 , if some node $\tilde{x}_1 \in h_1$ belongs to $X^+(C^*)$ then $F(\tilde{x}_1, \tilde{C}) \leq F(y_2, \tilde{C})$. Since we also have $F(\hat{x}_2, \tilde{C}) \leq F(y_1, \tilde{C})$ for every $\hat{x}_2 \in \hat{h}'_2$ the restriction on the values that consistent beliefs can take for strategy profiles with carrier equal to \tilde{C} results. Therefore, we only need to analyze $h_1 \subset X^0(C^*)$ together with $|\hat{h}_2| \leq |\hat{h}'_2|$ (given that the other case is similar). Let L be an arbitrary C^* -labelling. The following holds:

$$F_L(x_1) = \sum_{\hat{y}_2 \in \hat{h}_2} F_L(\hat{y}_2) + L(c_1),$$

$$F_L(y_1) = \sum_{\hat{x}_2 \in \hat{h}'_2} F_L(\hat{x}_2) + L(c_2),$$

$$F_L(y_2) \geq L(c_1),$$

$$F_L(x_2) = L(c_2).$$

To understand better the first (and the second) equality notice that, by construction, the only actions in the path to x_1 that do not belong to C^* are c_1 and those that we can also find in the path to some $\hat{y}_2 \in \hat{h}_2$.

Take now a consistent assessment (σ, μ) with $\mathcal{C}(\sigma) = C^*$. If $\mu(y_2 | h_2) > 0$ then $L(c_1) \leq L(c_2)$. If moreover $\mu(\hat{y}_2 | h_2) > 0$ for all $\hat{y}_2 \in \hat{h}'_2$ then $F_L(x_1) \leq F_L(y_1)$ because $|\hat{h}_2| \leq |\hat{h}'_2|$. This completes the proof as $\mu(y_1 | h_1)$ cannot take a strictly positive value. \square

The proof of Proposition 3 is similar. Given K information sets h_1, \dots, h_K the plan is to find two nodes x_k and y_k for each information set h_k so that for some carrier C we have $F(x_1, C) \leq F(y_2, C), \dots, F(x_K, C) \leq F(y_1, C)$. If every h_k is included in $X^0(C)$ then we cannot have $F_L(y_k) < F_L(x_k)$ for every k and some C -labelling L . This implies that if (σ, μ) is consistent and $\mathcal{C}(\sigma) = C$ then $\mu(x_k | h_k) \neq 1$ for at least one k .

Proof of Proposition 3. We need to show that when the conditions in the proposition hold not every weakly-Bayesian assessment is consistent. Therefore, as in the proof of the previous proposition it will be enough to prove it when $\Phi_1 = \emptyset$. Moreover, using the induction procedure, once the proposition has been proven for $K = 2$ we assume that the proposition has also been proven for every value $K' < K$ strictly larger than 2. This allows us to also assume henceforth $\Phi_{K'} = \emptyset$ for $2 \leq K' < K$.

The conditions in the proposition imply that the following set is nonempty:¹⁰

$$\begin{aligned} \Phi_K = \{ & (x_1, y_1, \dots, x_K, y_K, c_1, \dots, c_K, C) \in X^{2K} \times A^K \times \mathcal{C}(\Sigma) : \\ & h(x_i) \neq h(x_j) \text{ for all } i \neq j, \text{ and for all } i = 1, \dots, K, \\ & h(x_i) \subset X^0(C), y_i \in h(x_i), y_i \neq x_i, c_i \in (\mathcal{P}_A(x_i) \cap \mathcal{P}_A(y_{i+1})) \setminus C \}. \end{aligned}$$

Let $h_i = h(x_i)$. Take an arbitrary element of Φ_K and construct the carrier:

$$\hat{C} = C \cup \left(\bigcup_{i=1}^K (\mathcal{P}_A(x_i) \setminus \{c_i\}) \right).$$

From Lemma 2 (i) we obtain $F(x_i, \hat{C}) \leq F(y_{i+1}, \hat{C})$ for every $i = 1, \dots, K$. If for every i we also obtain $h_i \subset X^0(\hat{C})$ then a consistent assessment (σ, μ) with $\mathcal{C}(\sigma) = \hat{C}$ cannot satisfy $\prod_{i=1}^K \mu(y_i | h_i) = 1$. Thus, we need to prove the result when, for some i , the set h_i is not included in $X^0(\hat{C})$. Let I represent the collection of indexes i such that $(h_i \cap X^+(\hat{C})) \neq \emptyset$.

Take an $i \in I$. Let \hat{h}_i represent the subset of those $\hat{y}_i \in h_i$ that belong to $X^+(\hat{C})$. Nodes x_i and y_i do not belong to \hat{h}_i because c_i and c_{i-1} are not in \hat{C} . Moreover, every action $c \in (\mathcal{P}_A(\hat{y}_i) \setminus C)$ is contained in $\mathcal{P}_A(x_{i-1})$ and, therefore, $F(\hat{y}_i, C) \leq F(x_{i-1}, C)$ for all $\hat{y}_i \in \hat{h}_i$. In other words, no action $c \in (\mathcal{P}_A(\hat{y}_i) \setminus C)$ can be contained in some $\mathcal{P}_A(x_k)$ with $k \neq i-1$. This would imply that $\Phi_{K'} \neq \emptyset$ for some integer $K' < K$.¹¹

¹⁰ Throughout the proof, when the index i equals K , the index $i+1$ refers to 1. Likewise, if $i = 1$, the index $i-1$ refers to K .

¹¹ For instance, if action $c \in (\mathcal{P}_A(\hat{y}_i) \setminus C)$ belongs to $\mathcal{P}_A(x_{i-2})$ then Φ_{K-1} would be nonempty. One element of this set can be obtained from our initial choice from Φ_K by dropping the entries corresponding to the nodes at information set h_{i-1} and action c_{i-1} , and substituting node y_i by node \hat{y}_i and action c_{i-2} by action c .

Construct the carrier:

$$\hat{C}' = C \cup \left(\bigcup_{k \notin I} (\mathcal{P}_A(y_k) \setminus \{c_{k-1}\}) \right).$$

Lemma 2 (i) implies $F(y_k, \hat{C}') \leq F(x_{k-1}, \hat{C}')$ for every $k \notin I$. Furthermore, it also implies that we still have $F(\hat{y}_i, \hat{C}') \leq F(x_{i-1}, \hat{C}')$ for every $i \in I$ and every $\hat{y}_i \in \hat{h}_i$. If all the information sets h_1, \dots, h_K are contained in $X^0(\hat{C}')$ then we obtain that no consistent assessment whose strategy has carrier \hat{C}' can assign a belief equal to one to every decision node x_k . Therefore, we need to analyze what would happen otherwise. Let J represent the collection of indexes j such that $(h_j \cap X^+(\hat{C}')) \neq \emptyset$, furthermore, for each $j \in J$, let $\hat{h}'_j \subset h_j$ be the subset of nodes in h_j that belong to $X^+(\hat{C}')$. The assumption that $\Phi_{K'} = \emptyset$ whenever $K' < K$ indicates that if $j \in J$ then for every $\hat{x}_j \in \hat{h}'_j$ we obtain $(\mathcal{P}_A(\hat{x}_j) \setminus C) \subset \mathcal{P}_A(y_{j+1})$. With this in mind we use the set of choices

$$\begin{aligned} \tilde{B} = & \left\{ \bigcup_{i \in I} \left[\mathcal{P}_A(x_{i-1}) \setminus \left(\bigcup_{\hat{y}_i \in \hat{h}_i} \mathcal{P}_A(\hat{y}_i) \right) \right] \right\} \cup \left\{ \bigcup_{i \in I} \mathcal{P}_A(y_i) \right\} \cup \\ & \left\{ \bigcup_{j \in J} \left[\mathcal{P}_A(y_{j+1}) \setminus \left(\bigcup_{\hat{x}_j \in \hat{h}'_j} \mathcal{P}_A(\hat{x}_j) \right) \right] \right\} \cup \left\{ \bigcup_{j \in J} \mathcal{P}_A(x_j) \right\} \end{aligned}$$

to construct the new carrier

$$\tilde{C} = C \cup (\tilde{B} \setminus \{c_1, \dots, c_K\}). \quad (4.1)$$

We can assume that every information set h_1, \dots, h_K is contained in $X^0(\tilde{C})$. To see why note that if $(h_k \cap X^+(\tilde{C})) \neq \emptyset$ then there must be an action \tilde{c} in the second or in the fourth component of \tilde{C} that is in the path of some node \tilde{x}_k of h_k . Moreover, this node cannot be either x_k or y_k . Since $\Phi_{K'} = \emptyset$ for every $K' < K$ the action \tilde{c} must be contained in either $\mathcal{P}_A(y_{k+1})$ or $\mathcal{P}_A(x_{k-1})$. Consider that $\tilde{c} \in \mathcal{P}_A(y_{k+1})$ then by Lemma 2 (iii) we have that $F(\tilde{x}_k, \tilde{C}) \leq F(y_{k+1}, \tilde{C})$ and we only need to replace x_k by \tilde{x}_k and redefine I so that it does not include k . Analogously, suppose now that $\tilde{c} \in \mathcal{P}_A(x_{k-1})$. Lemma 2 (iii) implies $F(\tilde{x}_k, \tilde{C}) < F(x_{k-1}, \tilde{C})$. We now need to replace y_k by \tilde{x}_k and remove the index k from J .

The last carrier that we consider is:¹²

$$C^* = \begin{cases} \tilde{C} \cup \left(\bigcup_{k \notin (I-1) \cup (J+1)} (\mathcal{P}_A(x_k) \setminus \{c_k\}) \right) & \text{if } \frac{\prod_{i \in I} (|\hat{h}_i| + 1)}{\prod_{j \in J} (|\hat{h}'_j| + 1)} \geq 1, \\ \tilde{C} \cup \left(\bigcup_{k \notin (I-1) \cup (J+1)} (\mathcal{P}_A(y_k) \setminus \{c_{k-1}\}) \right) & \text{otherwise.} \end{cases}$$

¹² Note the sets $(I-1) = \{i : i+1 \in I\}$ and $(J+1) = \{j : j-1 \in J\}$.

Suppose that $\prod_{i \in I} (\hat{h}_i + 1) \geq \prod_{j \in J} (\hat{h}'_j + 1)$. We now show that a consistent assessment (σ, μ) with $\mathcal{C}(\sigma) = C^*$ cannot satisfy at the same time all of the following:

- (i) for every $i \notin I$, $\mu(y_i | h_i) = 1$,
- (ii) for every $i \in I$, $\mu(y_i | h_i) > 0$ and $\mu(\hat{y}_i | h_i) > 0$ for all $\hat{y}_i \in \hat{h}_i$,
- (iii) for every $i \in I$, $\mu(y_i | h_i) + \sum_{\hat{y}_i \in \hat{h}_i} \mu(\hat{y}_i | h_i) = 1$.

If the consistent assessment (σ, μ) satisfies (i), (ii) and (iii) then we should be able to find a C^* -labelling, say L , as in Lemma 1. Given that the set \tilde{C} defined in (4.1) is contained in C^* we can write $F_L(x_{i-1}) = \sum_{\hat{y}_i \in \hat{h}_i} F_L(\hat{y}_i) + F_L(y_i)$ for every $i \in I$. Additionally, (ii) and (iii) above imply that $F_L(x_{i-1}) = (\hat{h}_i + 1)F_L(y_i)$ for every $i \in I$. A similar argument shows that for every $j \in J$ the equality $F_L(x_j) = (\hat{h}'_j + 1)^{-1}F_L(y_{j+1})$ also holds. The definition of C^* for the case that we are considering entails $F_L(x_k) \leq F_L(y_{k+1})$ whenever $k \notin (I - 1) \cup (J + 1)$. Finally, $F_L(y_k) < F_L(x_k)$ for every $k = 1, \dots, K$ given that we always have $\mu(x_k | h_k) = 0$ and $\mu(y_k | h_k) > 0$. We only have to put all these inequalities together to obtain:

$$\frac{\prod_{i \in I} (\hat{h}_i + 1)}{\prod_{j \in J} (\hat{h}'_j + 1)} < 1.$$

Which provides a contradiction. We can also conclude the proof because the case $\prod_{i \in I} (\hat{h}_i + 1) < \prod_{j \in J} (\hat{h}'_j + 1)$ is analogous. \square

The next step is to prove that the conditions given Propositions 1 and 3 are not only sufficient but also necessary. In order to prove this we need a characterization of consistent assessments.

Lemma 3 (Kreps and Wilson (1982, Lemma A2)). *Let (C, Y) be a consistent basis and let (σ, μ) satisfy $\mathcal{C}(\sigma) = C$ and $\text{supp}(\mu) = Y$. The assessment (σ, μ) is consistent if and only if there exists a function $\pi : A \cup A_0 \rightarrow (0, 1)$ such that $\pi(c) = \lambda(c)$ whenever $c \in A_0$, $\pi(c) = \sigma(c | h)$ whenever $c \in \mathcal{C}(\sigma)$, and for every $x \in X$ with $\mu(x | h) > 0$:*

$$\mu(x | h) = \frac{\prod_{c \in \mathcal{P}(x)} \pi(c)}{\sum_{\{x' \in h; \mu(x' | h) > 0\}} \left(\prod_{c \in \mathcal{P}(x')} \pi(c) \right)}. \quad (4.2)$$

Now we can turn to prove Theorem 1.

Proof of Theorem 1. Fix an extensive-form that satisfies neither the conditions of Proposition 1 nor the conditions of Proposition 3. Given any carrier C a consistent basis (C, Y) always exists (Lemma 1 gives a way of seeing this). Take a consistent assessment (σ, μ) with $\mathcal{C}(\sigma) = C$ and $\text{supp}(\mu) = Y$. Let L be the associated labelling and let π be a function such as the one in equation (4.2).

The collection of non-singleton information sets h that satisfy $h \subset X^0(C)$ is denoted H^0 . Take any information set $h \in H^0$. It is enough to prove that for every

μ' that only differs from μ at information set h , i.e. satisfies $\mu'(\cdot | h') = \mu(\cdot | h')$ for every $h' \neq h$, the assessment (σ, μ') is consistent.

First we show that if $\text{supp}(\mu') = \text{supp}(\mu)$ then (σ, μ') is consistent. Given the system of beliefs μ' we are going to construct a function π' such as the one in Lemma 3 that justifies it. Fix an arbitrary node x^* that belongs to h and Y , the support of both μ and μ' . Let $\pi'(c) = \pi(c)$ for every $c \in \mathcal{P}(x^*)$. For the rest of the nodes in h and Y we only modify the value taken by π' with respect to π for just one choice in its path. In symbols, for each $x \in ((h \cap Y) \setminus \{x^*\})$ choose any action $c_x \in (\mathcal{P}_A(x) \setminus C)$ and let $\pi'(c) = \pi(c)$ for every other $c \in (\mathcal{P}(x) \setminus \{c_x\})$. The value of $\pi'(c_x)$ is calculated in order to adjust the relative values of μ' with respect to μ appropriately:

$$\pi'(c_x) = \frac{\mu'(x | h) \mu(x^* | h)}{\mu'(x^* | h) \mu(x | h)} \pi(c_x). \quad (4.3)$$

However, for each x the choice c_x may also belong to the path to a node that is not in h . To keep track of those choices we let the set A^h consists of those actions whose value under π' has been assigned by (4.3). Likewise, the set Y^h consists of those nodes y that belong to some information set in $H^0 \setminus \{h\}$ and that have an action in their paths that belongs to A^h . A node that belongs to Y^h may contain in its path more than one choice in A^h but, by assumption, Y^h cannot contain two nodes that belong to the same information set.

For each $y^* \in Y^h$ we maintain $\pi'(c) = \pi(c)$ for every $c \in (\mathcal{P}(y^*) \setminus A^h)$. For the rest of the nodes $y \in (h(y^*) \setminus \{y^*\})$ that belong to Y , the support of μ , we choose any action $c_y \in (\mathcal{P}_A(y) \setminus C)$ and let $\pi'(c) = \pi(c)$ for every other action $c \in (\mathcal{P}(y) \setminus \{c_y\})$. We have to adjust the value of π' to maintain in the information set $h(y^*)$ the same beliefs as in μ . To do that we offset the changes made in 4.3 so that

$$\pi'(c_y) = \pi(c_y) \prod_{c \in \mathcal{P}_A(y^*) \cap A^h} \left(\frac{\pi(c)}{\pi'(c)} \right). \quad (4.4)$$

Again we can define the set of actions $A^{h(y^*)}$ whose value under π' has been defined by (4.4) and the set $Y^{h(y^*)}$ of nodes that belong to some information set in $H^0 \setminus h(y^*)$ and that satisfy $(\mathcal{P}_A(y) \cap A^{h(y^*)}) \neq \emptyset$. The set $A^{h(y^*)}$ does not contain two nodes from the same information set. Furthermore, since the conditions given in Proposition 3 are not met, it does not contain nodes in h and, for any $y^* \in A^h$, it does not contain any node in $h(y^*)$ either.

Since the set H^0 is finite, we can continue in the same fashion until all the actions in the paths to nodes in information sets that belong to H^0 are exhausted without redefining any value of π' . Finally, we have to set $\pi'(c) = \pi(c)$ for every unassigned c . One can check that the resulting π' satisfies equation (4.2) for the system of beliefs μ' .

Now we prove that for any $x^* \in h$ the bases $(C, Y \cup \{x^*\})$ and $(C, Y \setminus \{x^*\})$ are also consistent. We show it first for the basis $(C, Y \cup \{x^*\})$.

Let $Y' = Y \cup \{x^*\}$. As stated in Lemma 1 we are going to construct a C -labelling L' such that $x \in Y'$ if and only if x minimizes $F_{L'}(\cdot)$ over $h(x)$. Set $L'(c) = L(c)$ for every $c \in \mathcal{P}_A(x^*)$ and for the rest of the nodes $x \neq x^*$ in h take an arbitrary $c_x \in (\mathcal{P}_A(x) \setminus C)$ and let

$$L'(c_x) = L(c_x) + F_L(x^*) - F_L(x). \quad (4.5)$$

We fix $L'(c) = L(c)$ for every other action $c \in (\mathcal{P}_A(x) \setminus \{c_x\})$. That is, we are adjusting L' so that $F_{L'}(x) = F_{L'}(x^*)$ for every $x \in Y$.

We will assign the remaining values of L' recursively. For the same reasons as before, we know that no value is going to be redefined. Let A^h now be the set of those actions whose value under L' has been assigned in (4.5) and, similarly, let Y^h now be the set of those nodes that belong to some information set in $H^0 \setminus \{h\}$ and whose paths have an action in A^h . For each $y^* \in Y^h$ we fix $L'(c) = L(c)$ for every action $c \in \mathcal{P}_A(y^*)$ and for each $y \in (h(y^*) \setminus \{y^*\})$ select an arbitrary $c_y \in (\mathcal{P}_A(y) \setminus C)$. Let $L'(c) = L(c)$ for every $c \in (\mathcal{P}_A(y) \setminus \{c_y\})$ and

$$L'(c_y) = L(c_y) + \sum_{c \in \mathcal{P}_A(y^*) \cap A^h} (L'(c) - L(c)).$$

We can continue in the same fashion until we have exhausted all the actions in the paths to the nodes that belong to some information set in H^0 . In order to make L' completely defined let $L'(c) = L(c)$ for every action that remains unassigned. It is easy to check that the labelling L' satisfies the condition given in Lemma 1 for the basis (C, Y') .

To conclude it remains to show that for any $x^* \in h$ the basis $(C, Y \setminus \{x^*\})$ is also consistent. Take an arbitrary $c_{x^*} \in (\mathcal{P}_A(x^*) \setminus C)$ and let $L'(c_{x^*}) = L(c_{x^*}) + 1$. We fix $L'(c) = L(c)$ for every other action $c \in (\mathcal{P}_A(x^*) \setminus \{c_{x^*}\})$ in the path to x^* and also for every action $c \in \mathcal{P}_A(x)$ in the path to any other node $x \in h$ different from x^* . The next step is to assign the values of L' for those actions leading to nodes contained in each information set $h(y^*) \in H^0$ that satisfies $c_{x^*} \in \mathcal{P}_A(y^*)$. Since hereafter everything is analogous to the previous case we can conclude the proof. \square

5. SEQUENTIALLY RATIONAL WEAKLY-BAYESIAN ASSESSMENTS

In this section we consider extensive-form games and sequentially rational weakly-Bayesian assessments. Obviously, if for an extensive-form every weakly-Bayesian assessment is consistent then, for every payoff vector, every sequentially rational weakly-Bayesian assessment is a sequential equilibrium.¹³ Suppose that we are given an extensive-form where some weakly-Bayesian assessment is not consistent. We want to address whether we can always find payoffs so that in the

¹³ A sequentially rational weakly-Bayesian assessment is a *weak perfect Bayesian equilibrium* as defined by Mas-Colell et al. (1995, Definition 9.C.3).

resulting extensive-form game sequential equilibrium refines the set of sequentially rational weakly-Bayesian assessments.

We first introduce some additional notation needed to define sequential rationality. If players play according to the strategy profile σ the expected utility to player n is then given by the expression $U_n(\sigma) = \sum_{z \in Z} \mathbb{P}(z | \sigma) u_n(z)$. Let $\mathbb{P}_x(\cdot | \sigma)$ be the probability distribution generated on Z if players use the strategy profile σ and the game starts at the decision node x . (Note that $\mathbb{P}_x(\cdot | \sigma)$ is always well defined.) The expected utility to player n from the strategy profile σ at the information set h given the system of beliefs μ is equal to $U_n(\sigma | h, \mu) = \sum_{x \in h} \mu(x | h) \sum_{z \in Z} \mathbb{P}_x(z | \sigma) u_n(z)$.

Definition 4. The assessment (σ, μ) is sequentially rational if at every information set h the strategy of the player moving at h , say player n , satisfies

$$U_n(\sigma_{-n}, \sigma_n | h, \mu) \geq U_n(\sigma_{-n}, \sigma'_n | h, \mu) \text{ for every } \sigma'_n \in \Sigma_n.$$

The next lemma asserts that if we can find weakly-Bayesian assessments that are not consistent then, for some payoffs, there are behavioral strategies that are part of sequentially rational weakly-Bayesian assessments that are not sequential equilibrium strategies. The proof of the theorem consists of constructing such a payoff vector.

Proposition 4. *Consider an extensive-form where the set of set consistent assessments is strictly contained in the set of weakly-Bayesian assessments. We can find a game with that extensive-form such that the set of sequential equilibrium strategies is a strict subset of the projection on Σ from the set of sequentially rational weakly-Bayesian assessments.*

Proof. Let K be such that $\Phi_K \neq \emptyset$ and either $\Phi_{K-1} = \emptyset$ or $K = 1$. Propositions 1 and 3 imply that we can find a carrier C and K information sets h_1, \dots, h_K that belong to $X^0(C)$ such that, for every consistent assessment (σ, μ) with $\mathcal{C}(\sigma) = C$, each h_i strictly contains a subset \hat{h}_i with $\prod_{i=1}^K (\sum_{y \in \hat{h}_i} \mu(y | h_i)) < 1$. That is, if (σ, μ) is a consistent assessment there must be at least one information set $h_i \in \{h_1, \dots, h_K\}$ with at least one node $x \in h_i \setminus \hat{h}_i$ that satisfies $\mu(x | h_i) > 0$.

For each $i = 1, \dots, K$ let c_i be an action available at h_i such that $\sigma(c_i | h_i) = 0$. (If at least one does not exist we only need to modify the carrier C appropriately.) Assign a payoff equal to zero to the player who moves at h_i at every ending node that follows some action in $A(h_i) \setminus \{c_i\}$. Also assign a payoff equal to zero to ending nodes that follow action c_i when taken at any node in \hat{h}_i . Assign a payoff equal to 1 to every player elsewhere. A weakly-Bayesian assessment (σ, μ') such that $\prod_{i=1}^K (\sum_{y \in \hat{h}_i} \mu'(y | h_i)) = 1$ is sequentially rational but not consistent. \square

A possible criticism to the relevance of Proposition 4 is that (as the proof takes advantage of) differences in strategies may only occur at parts of the extensive-form that are not reached by the strategy profile. In principle, we would like to




FIGURE 9.

show that if some weakly-Bayesian assessment is not consistent then, for some payoffs, sequential equilibrium selects only a strict subset from the set outcomes generated by sequentially rational weakly-Bayesian assessments. However this may not be possible.

Consider Figure 9. The five open circles are the initial nodes, each of them is selected with equal probability by Nature. Proposition 3 implies that in this extensive-form some weakly-Bayesian assessment is not consistent. Those assessments must attach probability zero to the two information sets of Player IV. That means that the actions l_1 , r_2 , l_3 and r_4 have to be taken with probability one which leaves, for instance, actions r_1 and l_2 as the two actions that Proposition 3 requires for $K = 2$. (This corresponds to the carrier C^* constructed in the proofs of propositions 2 and 3.) In this example, consistent beliefs can be arbitrary at the bottom information set of Player IV but they impose restrictions on the set of consistent beliefs at her top information set. Weakly-Bayesian beliefs can take, by definition, arbitrary values at both information sets. Consider now any game with that extensive-form. Whether or not actions l_1 , r_2 , l_3 and r_4 are sequentially rational does not depend on what is the behavior at the top information set of Player IV. The reason is that that information set can only be reached from zero probability nodes at positive probability information sets. This implies that if both information sets of Player IV are reached with probability zero the strategy part of a sequentially rational weakly-Bayesian assessment and a sequential equilibrium strategy may only differ in behavior at Player IV's top information set. However, behavior at that information set cannot affect the sequential equilibrium path.¹⁴

¹⁴ The same is true even if we consider concepts stronger than weakly-Bayesian such as *pre-consistent assessments* as defined in the next section. In the extensive-form of Figure 9, for any

6. APPLICATIONS: STRENGTHENING WEAKLY-BAYESIAN ASSESSMENTS

Weakly-Bayesian assessments offer a useful benchmark against which to compare consistent assessments. Not only did such a comparison allow us to find the maximal set of extensive forms where consistency never imposes restrictions at zero probability information sets, but also disentangled the different restrictions entailed by consistency. This in turn can help expand the set of extensive forms for which we can find the whole set of consistent assessments in a simple and intuitive way.

We already mentioned that consistent and weakly-Bayesian assessment coincide in standard signaling games (i.e. those with one sender, one receiver, different types of sender and only one type of receiver). We can also obtain equivalence within variations of this standard model of signaling games, such as the one in Figure 6 once we substitute both sets of consecutive binary choices of Player I by a single three way choice (although substituting just one is enough). Blackboard examples such as the ones in figures 1 and 3 also feature the equivalence. A common and easily observable feature of all these games is that nodes inside non-singleton information sets have completely independent paths.

Instances of long-horizon games such that every weakly-Bayesian assessment is consistent include those where every information set is always reached with positive probability, e.g. games where players always receive noisy signals about past moves of their opponents (formalized by a move of Nature after each move) and no player moves twice along any path of play.

In order to find a larger collection of extensive-forms where the computation of consistent assessments can be done in a simple manner we now strengthen weakly-Bayesian assessments. A natural way of doing so is to require that players also update their beliefs at parts of the extensive forms that are not reached by the strategy profile.

6.1. Subgame Consistency

In what follows we consider assessments (σ, μ) such that, for every node x that is either a node where nature moves $x \in P_0$ or a singleton information set $\{x\} \in H$:

$$Z(h') \subset Z(x) \text{ and } \mathbb{P}_x(h' | \sigma) > 0 \Rightarrow \mu(x' | h') = \frac{\mathbb{P}_x(x' | \sigma)}{\mathbb{P}_x(h' | \sigma)} \text{ for every } x' \in h'. \quad (6.1)$$

That is, if the play of the game must go through the decision node x before going through h' then beliefs at h' must be computed via Bayes rule taking node x as reference point. Every assessment that satisfies (6.1) is a weakly-Bayesian assessment and when coupled with sequential rationality induces a subgame perfect

assignment of payoffs to ending nodes, the sets of sequential equilibria and sequentially rational preconsistent assessments coincide even though not every preconsistent assessment is consistent.

equilibrium—and a perfect Bayesian equilibrium as defined by Ritzberger (2002, Definition 6.2)—in every game with that extensive-form.

An assessments satisfying (6.1) may be consistent even if the extensive-form satisfies the hypothesis of Proposition 1. The set of extensive-forms where this is the case is not difficult to characterize once we understand Proposition 1. Games played in stages where every time a new stage commences all previous uncertainties are resolved constitute a good example as long as every “stage extensive-form”, properly defined, does not satisfy Proposition 1 nor Proposition 3 (e.g. any sequence of simultaneous move games or any sequence of signaling games where every uncertainty is resolved after each stage).

6.2. Updating Consistency

The second minimal strengthening of weakly-Bayesian assessments that we consider is requiring that players always update their beliefs based on their own previous beliefs and on the strategy profile. To capture this idea we use *updating consistent* assessments as introduced by Hendon et al. (1996) and Perea (2002). It is important to notice that updating consistency requires that when a player updates her beliefs she takes into account, whenever it must have been the case, her own past deviations from the strategy profile. That is, every player recognizes that if she actually has to move at some of her information sets it must be because she did not previously preclude that information set from happening. Formally:

Definition 5. An assessment (σ, μ) is a *updating consistent* if for every player n , every two information sets $h, h' \in H_n$ satisfying $h < h'$, and any pure strategy $s_n \in S_n(h')$ of player n , $\sum_{x \in h} \mu(x | h) \mathbb{P}_x(h' | \sigma_{-n}, s_n) > 0$ implies

$$\mu(x' | h') = \frac{\sum_{x \in h} \mu(x | h) \mathbb{P}_x(x' | \sigma_{-n}, s_n)}{\sum_{x \in h} \mu(x | h) \mathbb{P}_x(h' | \sigma_{-n}, s_n)} \quad \text{for every } x' \in h'.^{15} \quad (6.2)$$

Hendon et al. (1996) and Perea (2002) show that updating consistency is sufficient and necessary for the one-shot deviation principle to hold (given any strategy profile, if a player cannot improve her payoff by changing just one action at one information set then she cannot improve by deviating to a new strategy). Weakly-Bayesian assessments that are updating consistent are called *preconsistent* assessments by Hendon et al. (1996).

As a quick illustration of preconsistent assessments we reconsider games where players only receive noisy signals about past moves of their opponents. In these games every updating consistent weakly-Bayesian assessment is consistent whether or not players move twice along the same path of play.

¹⁵ Note that the sum in the numerator has only one nonzero term, i.e. the unique node y in information set h that satisfies $\mathcal{P}(y) \subset \mathcal{P}(x')$.

As the name suggests, every preconsistent assessment is consistent. We analyze, however, the stronger concept obtained by selecting from the set of assessments that satisfy (6.1) those that are updating consistent. Not to introduce yet an additional concept we will call those preconsistent as well.

Definition 6. An assessment (σ, μ) is a *preconsistent* assessment if it satisfies both (6.1) and (6.2).

It is easy to see that not every preconsistent assessment is consistent. Note that players beliefs in a preconsistent assessment need not be common knowledge. Hence, it is not be surprising that for preconsistent and consistent assessments to coincide we need two different players not to share the same uncertainty. A general condition for equivalence between preconsistent and consistent assessments is given in the following theorem.

Theorem 2. *The set of consistent and preconsistent assessment coincide in every extensive form where the following two conditions hold:*

- (i) *if $h_0 \in H_n$ makes the extensive-form satisfy Proposition 1 with choice c then either c belongs to player n , i.e. $c \in A(h)$, $h \in H_n$; or there exists a node x with $\{x\} \in H$ (or $x \in P_0$) such that $Z(h_0) \subset Z(x)$ and $c \in \mathcal{P}_A(x)$;*
- (ii) *if h_1, \dots, h_K make the extensive-form satisfy Proposition 3 with choices c_1, \dots, c_K then $\{h_1, \dots, h_K\} \subset H_n$ for some $n \in \mathcal{N}$; if, moreover, $Z(h_i) \cap Z(h_j) = \emptyset$ for all $i \neq j$ then there exists an additional information set $h \in H_n$ with at most K nodes such that $Z(h_i) \subset Z(h)$ for all $i = 1, \dots, K$ and $(\{c_1, \dots, c_K\} \setminus A_n) \subset \bigcup_{x \in h} \mathcal{P}_A(x)$.*

(If h_1, \dots, h_K are as in Proposition 3 some selections of those K information sets may also make the extensive-form satisfy the hypothesis of the definition. Note that for the theorem to hold we need the conditions to be satisfied *for any* family of information sets as in Proposition 3.)

Sketch of the proof. The idea is to use the same reasoning as in the proof of Theorem 1. The main difficulty lies with how differently weakly-Bayesian and preconsistent assessments are determined. Given a basis (C, Y) we can identify the set $H^*(C, Y)$ of information sets where the value assumed by μ is not pinned down by (6.1) or (6.2). We need to show that at those information sets consistency does not impose any restrictions.

Recall that any path is totally ordered by $<$. For each player n and each information set $h \in (H^*(C, Y) \cap H_n)$ assign the choice $c_x = \max_c \{(\mathcal{P}_A(x) \setminus A_n) \setminus C\}$ to every node $x \in h$. It follows from (i) that no two nodes in the same information set are assigned the same choice. For each two different nodes $x, x' \in h$ actions c_x and $c_{x'}$ are different as well and satisfy $c_x \in (\mathcal{P}_A(x) \setminus C)$ and $c_{x'} \in (\mathcal{P}_A(x') \setminus C)$. We can choose trembles associated to c_x and $c_{x'}$ so as to make the relative likelihood between x and x' equal any value as we take the trembles to zero.

However, c_x may be the only choice in the path to some node in some other information set that is not in C . Thus, suppose that for two different information sets $h_1, h_2 \in H_n$ such that $h_1 < h_2$ there are nodes $x_1 \in h_1$ and $x_2 \in h_2$ for which $c_{x_1} = c_{x_2}$. Condition (i) implies that we must have $\mathcal{PP}(x_1) \subset \mathcal{PP}(x_2)$. Moreover, $\mu(x_1 | h_1) = 0$ whenever $h_2 \in H^*(C, Y)$. If $\mu(y_1 | h_1) > 0$ then every node y_2 in h_2 that follows h_1 is assigned a choice $c_{y_2} \notin \mathcal{P}_A(y_1)$, which in turn implies that the tremble associated to c_{y_2} cannot modify consistent beliefs at h_1 . Henceforth, we can justify as consistent arbitrary beliefs at h_2 when those are not given by (6.2). We can do so by adjusting trembles associated to choices that are either are not the path to any node in h_1 or, if they are, the corresponding node is assigned zero probability by the system of beliefs. In other words, under the condition of the theorem, preconsistent assessments will not fail to be consistent due to a “wrong” assignment of beliefs at h_2 as long as the value of beliefs up to information set h_1 agrees with consistency.

Consider now an indexed family of information sets $\{h_i\}_{i=1}^K \subset (H^*(C, Y) \cap H_n)$ such that $Z(h_i) \cap Z(h_j) = \emptyset$ for all $i \neq j$. Consider further the existence of an indexed family of pairs of nodes $\{(x_i, y_i)\}_{i=1}^K$ such that $x_i, y_i \in h_i$ and $c_{x_i} = c_{y_{i+1}}$ for every $i = 1, \dots, K-1$. By Proposition 3 consistency imposes restrictions on $\{h_i\}_{i=1}^K$ if we also have $c_{x_K} = c_{y_1}$. In such a case, by assumption, there must be an information set $h \in H_n$ with at most K nodes that precedes every information set in $\{h_i\}_{i=1}^K$. Moreover, some node, say $x \in h$ must belong to Y and $c_{x_i} \in \mathcal{P}_A(x)$ for some $i, 1 \leq i \leq K$. Since $h_i \in H^*(C, Y)$ the choice c_{x_i} is not the maximal element in $(\mathcal{P}_A(x) \setminus A_n) \setminus C$. This provides a contradiction and shows that consistency does to impose restrictions on $\{h_i\}_{i=1}^K$ beyond those imposed by preconsistency. \square

The two conditions in the last proposition guarantee that, whenever consistent beliefs cannot be arbitrarily chosen in a group of information sets controlled by player n they are preceded by another information set of player n , that she can take as reference point to apply Bayesian updating. In particular, the second part of condition (ii), by imposing restrictions in the number of nodes of that reference point, ensures that the relative probability between zero probability nodes does not need to be specified to compute consistent beliefs in the following information sets.

Before giving examples of extensive-forms where Theorem 2 is applicable it is convenient to analyze one where it is not so. Figure 10 contains one such an instance. The two last information sets of Player II satisfy Proposition 3 with $K = 2$ and associated choices l_1 and m_1 . They are preceded by another information set of Player II but it contains three nodes. This permits the left-hand node and middle node to not being ordered in terms of their likelihood by the beliefs at that information set thereby not imposing enough restrictions on preconsistent beliefs at the two last information sets.




FIGURE 10.

On the other hand, Theorem 2 can be applied to any game with one player who moves every other turn, who is the only player with nonsingleton information sets and whose information sets are totally ordered. It can also be applied to any two-player game, with or without moves of Nature, where Player II plays every other move, every information set of Player I is a singleton, and every information set of Player II has at most two nodes. In this last family of games, players do not have common uncertainties so we do not need to worry about their beliefs contradicting each other; player II moves sufficiently often, meaning that there is an explicit information set whose beliefs must be specified and updated via Bayes rule; and, moreover, the fact that information sets have at most two nodes implies that no two nodes are given zero probability by any system of beliefs in the same information set. A subset of this family is formed by *multi-period games with observed actions* as defined by Fudenberg and Levine (1983) and analyzed in Fudenberg and Tirole (1991) where Player I has two possible types and Player II can only be of one type. As we show in the next section, this family is enough to characterize consistency in multi-period games with observed actions with more than two players and with at most two types per player.

6.3. Multi-Period Games with Observed Actions

A *multi-period game with observed actions* is played in stages. In the first stage or period Nature chooses independently—we assume here so for simplicity—the type of each player and that information is only revealed to that player. At each following stage players move simultaneously and at the end of the period their moves are fully revealed. Therefore, the only uncertainty during the game concerns the initial move of Nature. A small warning is appropriate: unlike us, Fudenberg and Tirole (1991) assign the same label to choices of a player that are available at different information sets but that the opponents observe as identical.

We can recover Proposition 3.1 in Fudenberg and Tirole (1991) which characterizes the set of consistent assessments in every multi-period game with observed

actions where any player has at most two types. Fudenberg and Tirole (loc. cit.) do so requiring that beliefs about player n 's type depend only on player n 's actions and through a commonality requirement on beliefs. Instead of giving such conditions here we will take an alternative route by decomposing the whole game into pieces. Afterwards, we will be able to apply Theorem 2 to each of those pieces.

Indeed, given a multi-period game with observed actions, for each player n we can construct a derived two-player multi-period game. In this derived game, Nature moves first and chooses the type of player n with the same probability distribution as in the original game. After the initial move of Nature, player n has two singleton information sets (as many as possible types). Each information set has the same moves available as the corresponding first information set of player n in the original game. The next actor to move in this derived game is an *observer* of player n . The observer has as many information sets as previous moves were available in each of the singleton information sets of player n . Each of these information sets has two nodes (as many as different types of player n) and as many moves available as action profiles of player n 's opponents in the first period of the original game. The moves of the observer recreate, jointly with player n 's moves, the histories that can occur in the original game. So after each move of the observer (action profile of the opponents of player n), a singleton information set of player n follows again with the same actions available as in the stage of the original multi-period game that is preceded by the same history.

If we continue with this construction we obtain a two-player multi-period game where we can apply Theorem 2. Moreover, the set of consistent assessments of the original multi-period game can be characterized by the N different sets of consistent assessments of the different derived games. (By the same token we can also recover Proposition 3.2 in Fudenberg and Tirole (1991) which characterizes the set of consistent assessments in every multi-period game with two stages. In this case the derived game that we obtain satisfies the hypothesis of Theorem 1.)

Weak independence (3.1) is the reason why consistency in the N -player multi-period game can be characterized through consistency of the N different two-player games where one player plays the role of the observer. This is true in every multi-period game with observed actions no matter how many types or strategies (as long as the numbers are finite). The reason why preconsistency cannot be used to characterize consistency when a Player has three types or more is illustrated in Figure 11. A similar Figure can be found in Fudenberg and Tirole (1991, Figure 1) or in Osborne and Rubinstein (1994, Figure 235.1) to explain basically the same point. We can look at it here, however, from the point of view of Theorem 2 and see that it has the same relevant characteristic as Figure 10. The figure represents a period of a two-player multi-period game with observable actions where Player II (who can be thought as the *observer* of our previous construction) has one type




FIGURE 11.

and Player I has three possible types. After the history that precedes that period, Player II gives a belief equal to one to the rightmost node in her first information set. For simplicity, only one choice is available at that information set, after which is Player I 's turn to move. The starred choices are those that Player I takes with positive probability in the strategy profile. The information set receiving the l 's actions and the one receiving the m 's actions make the extensive form satisfy Proposition 3 with $K = 2$. The actions associated can be found in the paths to the two zero probability nodes on the first Player II 's information set. Note that that information set has three nodes and not two as required by Theorem 2. This allows that the leftmost and middle nodes of that information set not be given a likelihood ordering by the system of beliefs.

REFERENCES

- Battigalli, P., 1996. Strategic independence and perfect Bayesian equilibria. *J. Econ. Theory* 70 (1), 201–234.
- Blume, L., Brandenburger, A., Dekel, E., 1991. Lexicographic probabilities and choice under uncertainty. *Econometrica* 59 (1), 61–79.
- Fudenberg, D., Levine, D., 1983. Subgame-perfect equilibria of finite- and infinite-horizon games. *J. Econ. Theory* 31, 251–268.
- Fudenberg, D., Tirole, J., 1991. Perfect Bayesian equilibrium and sequential equilibrium. *J. Econ. Theory* 53, 236–260.
- Hammond, P., 1994. Elementary non-Archimedean representations of probability for decision theory and games. In: Patrick Suppes: Scientific Philosopher. Vol. 1, Probability and Probabilistic causality. Kluwer Academic, pp. 25–25.
- Hendon, E., Jacobsen, H., Sloth, B., 1996. The one-shot-deviation principle for sequential rationality. *Games Econ. Behav.* 12 (2), 274–282.
- Kohlberg, E., 1990. Refinement of Nash equilibrium: The main ideas. In: Ichiishi, T., Neyman, A., Tauman, Y. (Eds.), *Game Theory and Applications*. Academic Press, San Diego, pp. 3–45.

- Kohlberg, E., Reny, P., 1997. Independence on relative probability spaces and consistent assessments in game trees. *J. Econ. Theory* 75 (2), 280–313.
- Kreps, D., Wilson, R., 1982. Sequential equilibria. *Econometrica* 50, 863–894.
- Litan, C., Pimienta, C., 2008. Conditions for equivalence between sequentiality and subgame perfection. *Econ. Theory* 35 (3), 539–553.
- Mas-Colell, A., Whinston, M., Green, J., 1995. *Microeconomic theory*. Oxford University Press New York.
- McLennan, A., 1989a. Consistent conditional systems in noncooperative game theory. *Int. J. Game Theory* 18 (2), 141–174.
- McLennan, A., 1989b. The space of conditional systems is a ball. *Int. J. Game Theory* 18, 125–139.
- Myerson, R., 1986. Multistage games with communication. *Econometrica* 54 (2), 323–358.
- Osborne, M., Rubinstein, A., 1994. *A course in game theory*. The MIT press.
- Perea, A., 2002. A note on the one-deviation property in extensive form games. *Games Econ. Behav.* 40 (2), 322–338.
- Perea y Monsuwé, A., Jansen, M., Peters, H., 1997. Characterization of consistent assessments in extensive form games. *Games Econ. Behav.* 21 (1-2), 238–252.
- Ritzberger, K., 2002. *Foundations of non-cooperative game theory*. Oxford University Press, USA.
- Swinkels, J. M., Nov. 1993. Independence for conditional probability systems. Discussion Papers 1076, Northwestern University.