To appear in the
International Journal of Game Theory
DOI: 10.1007/s00182-015-0506-6

AGM-consistency and perfect Bayesian equilibrium. Part II: from PBE to sequential equilibrium.

Giacomo Bonanno*
Department of Economics, University of California, Davis, CA 95616-8578, USA
gfbonanno@ucdavis.edu

January 2014. Revised August 2014. Revised February 2015.

Abstract

In [6] a general notion of perfect Bayesian equilibrium (PBE) for extensive-form games was introduced and shown to be intermediate between subgame-perfect equilibrium and sequential equilibrium. Besides sequential rationality, the ingredients of the proposed notion are (1) the existence of a plausibility order on the set of histories that rationalizes the given assessment and (2) the notion of Bayesian consistency relative to the plausibility order. We show that a cardinal property of the plausibility order and a strengthening of the notion of Bayesian consistency provide necessary and sufficient conditions for a PBE to be a sequential equilibrium.

Keywords: plausibility order, belief revision, Bayesian updating, sequential equilibrium, consistency.

1 Introduction

In [6] a solution concept for extensive-form games was introduced, called perfect Bayesian equilibrium (PBE), and shown to be a strict refinement of subgame-perfect equilibrium ([16]); it was also shown that, in turn, the notion of sequential equilibrium ([10]) is a strict refinement of PBE. Besides sequential rationality, the ingredients of the definition of perfect Bayesian equilibrium are (1) the qualitative notions of plausibility order and AGM-consistency and (2) the notion of Bayesian consistency relative to the plausibility order.[1] In this paper we continue the study of PBE by providing necessary and sufficient conditions for a PBE to be a sequential equilibrium. There are two such conditions. One is the notion of choice-measurability of the plausibility order, which was shown in [6] to be implied by sequential equilibrium. Choice measurability requires that the plausibility order $\precsim$ on the set of histories H that rationalizes the given assessment have a *cardinal* representation, in the sense that there is an integer-valued function $F : H \to \mathbb{N}$ such that (1) $F(h) \leq F(h')$ if and only if $h \precsim h'$ (this is the ordinal part) and (2) if a is an action available at h and h' (where h and

*I am grateful to an Associate Editor and two anonymous reviewers for helpful comments and suggestions.

[1] As shown in [5], these notions can be derived from the primitive concept of a player's epistemic state, which encodes the player's initial beliefs and her disposition to revise those beliefs upon receiving (possibly unexpected) information. The existence of a plausibility order that rationalizes the epistemic state of each player guarantees that the belief revision policy of each player satisfies the so-called AGM axioms for rational belief revision, which were introduced in [1].

h' are two histories in the same information set) then $F(h) - F(h') = F(ha) - F(h'a)$ (this is the cardinal part). A cardinal representation F can be interpreted as measuring the "plausibility distance" between histories and this distance is required to be preserved by the addition of a common action. Choice measurability imposes constraints on the supports of the belief system. The second condition concerns the distribution of probabilities over those supports and is a strengthening of the notion of Bayesian consistency which is part of the definition of PBE; we call the stronger condition "uniform Bayesian consistency". Both notions (choice measurability and uniform Bayesian consistency) are strictly related to existing notions in the literature, as detailed in Section 4. Although the characterization of sequential equilibrium provided in this paper is strictly related to earlier characterizations, it offers a novel understanding of sequential equilibrium in terms of its relationship to the notion of perfect Bayesian equilibrium.

The paper is organized as follows. The next section reviews the definition of PBE. In Section 3 it is shown that choice measurability and uniform Bayesian consistency are necessary and sufficient for a PBE to be a sequential equilibrium. Section 4 discusses related literature and Section 5 concludes. The proofs are given in the Appendix.

2 AGM-consistency and perfect Bayesian equilibrium

In this section we recall the notion of perfect Bayesian equilibrium introduced in [6]; we employ the same notation, which makes use of the history-based definition of extensive-form game (see [11]). As in [6], we restrict attention to *finite* extensive-form games with *perfect recall*.

A total pre-order on a set H is a binary relation $\precsim$ which is complete ($\forall h, h' \in H$, either $h \precsim h'$ or $h' \precsim h$) and transitive ($\forall h, h', h'' \in H$, if $h \precsim h'$ and $h' \precsim h''$ then $h \precsim h''$). We write $h \sim h'$ as a short-hand for "$h \precsim h'$ and $h' \precsim h$" and $h \prec h'$ as a short-hand for "$h \precsim h'$ and $h' \not\precsim h$".

Definition 1 *Given an extensive form, a* plausibility order *is a total pre-order $\precsim$ on the finite set of histories H that satisfies the following properties: $\forall h \in D$ (D is the set of decision histories),*

$PL1$. $h \precsim ha$, $\forall a \in A(h)$ *($A(h)$ is the set of actions available at h),*

$PL2$. *(i)* $\exists a \in A(h)$ *such that* $h \sim ha$,
(ii) $\forall a \in A(h)$, *if* $h \sim ha$ *then,* $\forall h' \in I(h)$, $h' \sim h'a$,
($I(h)$ is the information set that contains h),

$PL3$. *if history h is assigned to chance, then* $h \sim ha$, $\forall a \in A(h)$.

The interpretation of $h \precsim h'$ is that history h is *at least as plausible* as history h'.[2] Property $PL1$ says that adding an action to a decision history h cannot yield a more plausible history than h itself. Property $PL2$ says that at every decision history h there is at least one action a which is "plausibility preserving" in the sense that adding a to h yields a history which is as plausible as h; furthermore, any such action a performs the same role with any other history that belongs to the same information set. Property $PL3$ says that all the actions at a history assigned to chance are plausibility preserving.

[2] As in [6] we use the notation $h \precsim h'$ rather than the, perhaps more natural, notation $h \succsim h'$, for two reasons: (1) it is the standard notation in the extensive literature that deals with AGM belief revision (for a recent survey of this literature see the special issue of the *Journal of Philosophical Logic*, Vol. 40 (2), April 2011) and (2) when representing $\precsim$ numerically it is convenient to assign lower values to more plausible histories.

An *assessment* is a pair (σ, μ) where σ is a behavior strategy profile and μ is a system of beliefs.[3]

Definition 2 *Fix an extensive-form. An assessment* (σ, μ) *is* AGM-consistent *if there exists a plausibility order* $\precsim$ *on the set of histories* H *such that:*

(i) the actions that are assigned positive probability by σ *are precisely the plausibility-preserving actions:* $\forall h \in D, \forall a \in A(h)$,

$$\sigma(a) > 0 \textit{ if and only if } h \sim ha, \qquad (P1)$$

(ii) the histories that are assigned positive probability by μ *are precisely those that are most plausible within the corresponding information set:* $\forall h \in D$,

$$\mu(h) > 0 \textit{ if and only if } h \precsim h', \forall h' \in I(h). \qquad (P2)$$

If $\precsim$ *satisfies properties* $P1$ *and* $P2$ *with respect to* (σ, μ)*, we say that* $\precsim$ rationalizes (σ, μ).

An assessment (σ, μ) is sequentially rational if, for every player i and every information set I of hers, player i's expected payoff - given the strategy profile σ and her beliefs at I (as specified by μ) - cannot be increased by unilaterally changing her choice at I and possibly at information sets of hers that follow I.[4]

In conjunction with sequential rationality, the notion of AGM-consistency is sufficient to eliminate some subgame-perfect equilibria as "implausible". Consider, for example, the extensive game of Figure 1 and the pure-strategy profile $\sigma = (c, d, f)$ (highlighted by double edges), which constitutes a Nash equilibrium of the game (and also a subgame-perfect equilibrium since there are no proper subgames). Can σ be part of a sequentially rational AGM-consistent assessment (σ, μ)? Since, for Player 3, choice f can be rationally chosen only if the player assigns (sufficiently high) positive probability to history be, sequential rationality requires that $\mu(be) > 0$; however, any such assessment is *not* AGM-consistent. In fact, if there were a plausibility order $\precsim$ that satisfied Definition 2, then, by $P1$, $b \sim bd$ (since $\sigma(d) = 1 > 0$) and $b \prec be$ (since $\sigma(e) = 0$)[5] and, by $P2$, $be \precsim bd$

[3] A behavior strategy profile is a list of probability distributions, one for every information set, over the actions available at that information set. A system of beliefs is a collection of probability distributions, one for every information set, over the histories in that information set.

[4] The precise definition is as follows. Let Z denote the set of terminal histories and, for every player i, let $U_i : Z \to \mathbb{R}$ be player i's von Neumann-Morgenstern utility function. Given a decision history h, let $Z(h)$ be the set of terminal histories that have h as a prefix. Let $\mathbb{P}_{h,\sigma}$ be the probability distribution over $Z(h)$ induced by the strategy profile σ, starting from history h (that is, if z is a terminal history and $z = ha_1...a_m$ then $\mathbb{P}_{h,\sigma}(z) = \prod_{j=1}^{m} \sigma(a_j)$). Let I be an information set of player i and let $u_i(I|\sigma,\mu) = \sum\limits_{h \in I} \mu(h) \sum\limits_{z \in Z(h)} \mathbb{P}_{h,\sigma}(z) U_i(z)$ be player i's expected utility at I if σ is played, given her beliefs at I (as specified by μ). We say that player i's strategy σ_i is *sequentially rational at* I if $u_i(I|(\sigma_i, \sigma_{-i}), \mu) \geq u_i(I|(\tau_i, \sigma_{-i}), \mu)$ for every strategy τ_i of player i (where σ_{-i} denotes the strategy profile of the players other than i). An assessment (σ, μ) is *sequentially rational* if, for every player i and for every information set I of player i, σ_i is sequentially rational at I. Note that there are two definitions of sequential rationality: the *weakly local* one - which is the one adopted here - according to which at an information set a player can contemplate changing her choice not only there but possibly also at subsequent information sets of hers, and a *strictly local* one, according to which at an information set a player contemplates changing her choice only there. If the definition of perfect Bayesian equilibrium (Definition 5 below) is modified by using the strictly local definition of sequential rationality, then an extra condition needs to be added, namely the "pre-consistency" condition identified in [8] and [14] as being necessary and sufficient for the equivalence of the two notions. For simplicity we have chosen the weakly local definition.

[5] By $PL1$ of Definition 1, $b \precsim be$ and, by $P1$ of Definition 2, it is not the case that $b \sim be$ because e is not assigned positive probability by σ. Thus $b \prec be$.

(since - by hypothesis - μ assigns positive probability to be). By transitivity of $\precsim$, from $b \sim bd$ and $b \prec be$ it follows that $bd \prec be$, yielding a contradiction.

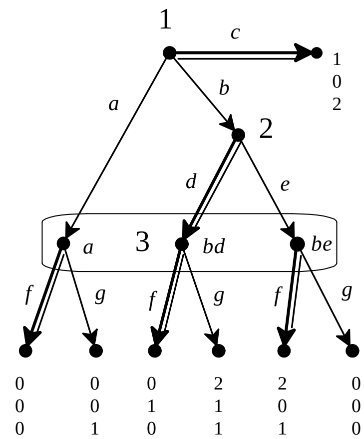

The Nash equilibrium $\sigma = (c, d, f)$ cannot be part
of a sequentially rational AGM-consistent assessment.
Figure 1

On the other hand, the Nash equilibrium $\sigma' = (b, d, g)$ together with $\mu'(bd) = 1$ forms a sequentially rational, AGM-consistent assessment: it can be rationalized by several plausibility orders, for instance the following ($\emptyset$ denotes the null history, that is, the root of the tree):

$$\begin{pmatrix} \emptyset, b, bd, bdg & \text{most plausible} \\ a, c, be, ag, beg & \\ af, bdf, bef & \text{least plausible} \end{pmatrix} \tag{1}$$

where each row represents an equivalence class. We use the following convention in representing a total pre-order: if the row to which history h belongs is above the row to which h' belongs, then $h \prec h'$ (h is more plausible than h') and if h and h' belong to the same row then $h \sim h'$ (h is as plausible as h').

Definition 3 *Fix an extensive form. Let $\precsim$ be a plausibility order that rationalizes the assessment (σ, μ). We say that (σ, μ) is* Bayesian relative to $\precsim$ *if for every equivalence class E of $\precsim$ that contains some decision history h with $\mu(h) > 0$ (that is, $E \cap D^+_\mu \neq \varnothing$, where $D^+_\mu = \{h \in D : \mu(h) > 0\}$) there exists a probability density function $\nu_E : H \to [0, 1]$ (recall that H is a finite set) such that:*

- $B1$. *$Supp(\nu_E) = E \cap D^+_\mu$.*
- $B2$. *If $h, h' \in E \cap D^+_\mu$ and $h' = ha_1...a_m$ (that is, h is a prefix of h') then $\nu_E(h') = \nu_E(h) \times \sigma(a_1) \times \ ... \ \times \ \sigma(a_m)$.*
- $B3$. *If $h \in E \cap D^+_\mu$, then, $\forall h' \in I(h)$, $\mu(h') = \nu_E(h' \mid I(h)) \overset{def}{=} \frac{\nu_E(h')}{\sum\limits_{h'' \in I(h)} \nu_E(h'')}$.*

Property $B1$ requires that $\nu_E(h) > 0$ if and only if $h \in E$ and $\mu(h) > 0$. Property $B2$ requires ν_E to be consistent with the strategy profile σ in the sense that if $h, h' \in E$, $\mu(h) > 0$, $\mu(h') > 0$

and $h' = ha_1...a_m$ then the probability that ν_E assigns to h' is equal to the probability that ν_E assigns to h multiplied by the probabilities (according to σ) of the actions that lead from h to h'.[6] Property $B3$ requires the system of beliefs μ to satisfy Bayes' rule in the sense that if $h \in E$ and $\mu(h) > 0$ (so that E is the equivalence class of the most plausible elements of $I(h)$) then for every history $h' \in I(h)$, $\mu(h')$ (the probability assigned to h' by μ) coincides with the probability of h' conditional on $I(h)$ using the probability measure ν_E.

How should one interpret the probability $\nu_E(h)$ of Definition 3? This issue was not discussed in [6].[7] First of all, it should be noted that one *cannot* interpret $\nu_E(h)$ as the "probability that history h is *reached* by the actual play of the game". To see this, consider the game of Figure 1 and the assessment $\sigma' = (b, d, g)$, $\mu'(bd) = 1$, which is rationalized by the plausibility order (1). Let $E = \{\emptyset, b, bd, bdg\}$ be the top equivalence class of that order, so that $E \cap D^+_\mu = \{\emptyset, b, bd\}$. Then there is only one function ν_E that satisfies the properties of Definition 3, namely $\nu_E(\emptyset) = \nu_E(b) = \nu_E(bd) = \frac{1}{3}$. In particular, $\nu_E(\emptyset) = \frac{1}{3}$ which is at odds with the fact that the play of the game "reaches" history $\emptyset$ (the root of the tree) for sure, that is, with probability 1. However, $\nu_E(\emptyset) = \frac{1}{3}$ does have a meaningful interpretation as suggested in [2, pp.114-115]. Define two random variables $\mathbf{d}$ and $\mathbf{t}$, where $\mathbf{d}$ is the *decision* node at which the play is and $\mathbf{t}$ is the current time; assume further that each move takes one unit of time (the initial time being set to 0). In the example of Figure 1, conditional on the actual play of the game belonging to $E = \{\emptyset, b, bd, bdg\}$, $\mathbf{d}$ takes the values $\emptyset$, b, and bd, while the possible values of $\mathbf{t}$ can be taken to be 0, 1 and 2. Then, conditional on $\sigma' = (b, d, g)$, $P(\mathbf{d} = \emptyset \mid \mathbf{t} = 0) = 1$, while for $i \in \{1, 2\}$, $P(\mathbf{d} = \emptyset \mid \mathbf{t} = i) = 0$; similarly, $P(\mathbf{d} = b \mid \mathbf{t} = 1) = P(\mathbf{d} = bd \mid \mathbf{t} = 2) = 1$ and, for $i \in \{0, 2\}$ and $j \in \{0, 1\}$, $P(\mathbf{d} = b \mid \mathbf{t} = i) = P(\mathbf{d} = bd \mid \mathbf{t} = j) = 0$. Letting it be equally probable that the current time is 0, 1 or 2, that is, $P(\mathbf{t} = 0) = P(\mathbf{t} = 1) = P(\mathbf{t} = 2) = \frac{1}{3}$, and defining $\nu_E(\emptyset)$ as $P(\mathbf{d} = \emptyset \textit{ and } \mathbf{t} = 0)$ we get that $\nu_E(\emptyset) = P(\mathbf{d} = \emptyset \mid \mathbf{t} = 0) \times P(\mathbf{t} = 0) = \frac{1}{3}$. Note that $P(\mathbf{d} = \emptyset \textit{ and } \mathbf{t} = 0) = P(\mathbf{d} = \emptyset)$, where $P(\mathbf{d} = \emptyset) = \sum_{i=0}^{2} [P(\mathbf{d} = \emptyset \mid \mathbf{t} = i) \times P(\mathbf{t} = i)]$ and thus one can equivalently define $\nu_E(\emptyset)$ as $P(\mathbf{d} = \emptyset \textit{ and } \mathbf{t} = 0)$ or as $P(\mathbf{d} = \emptyset)$. Similarly, $\nu_E(b) \overset{def}{=} P(\mathbf{d} = b \textit{ and } \mathbf{t} = 1) = \frac{1}{3}$ and $\nu_E(bd) \overset{def}{=} P(\mathbf{d} = bd \textit{ and } \mathbf{t} = 2) = \frac{1}{3}$ (note, again, that $P(\mathbf{d} = b \textit{ and } \mathbf{t} = 1) = P(\mathbf{d} = b)$ and $P(\mathbf{d} = bd \textit{ and } \mathbf{t} = 2) = P(\mathbf{d} = bd)$).

In general, when conditioning on an equivalence class E which is not the top equivalence class, the values $\{\nu_E(h)\}_{h \in E \cap D^+_\mu}$ – while maintaining the same interpretation – incorporate also the probabilities of deviations from the most plausible play(s). To see this, consider the extensive form of Figure 2 and the assessment $\sigma(f) = 1$, $\sigma(L) = \sigma(R) = \frac{1}{2}$, $\sigma(\ell) = \frac{1}{5}$, $\sigma(r) = \frac{4}{5}$, $\sigma(A) = 1$, $\sigma(D) = 1$, $\mu(a) = 0$, $\mu(b) = \frac{1}{3}$, $\mu(c) = \frac{2}{3}$, $\mu(d) = \frac{3}{4}$, $\mu(e) = \frac{1}{4}$, which is rationalized by the following plausibility order:

$$\begin{pmatrix} \emptyset, f \\ b, c, d, e, bL, bR, cL, cR, d\ell, dr, e\ell, er, bLA, erD \\ a, aL, aR, bLB, erC \end{pmatrix} \tag{2}$$

[6] Note that if $h, h' \in E$ and $h' = ha_1...a_m$, then $\sigma(a_j) > 0$, for all $j = 1, ..., m$. In fact, since $h' \sim h$, every action a_j is plausibility preserving and therefore, by Property $P1$ of Definition 2, $\sigma(a_j) > 0$.

[7] I am grateful to a reviewer for raising this question.

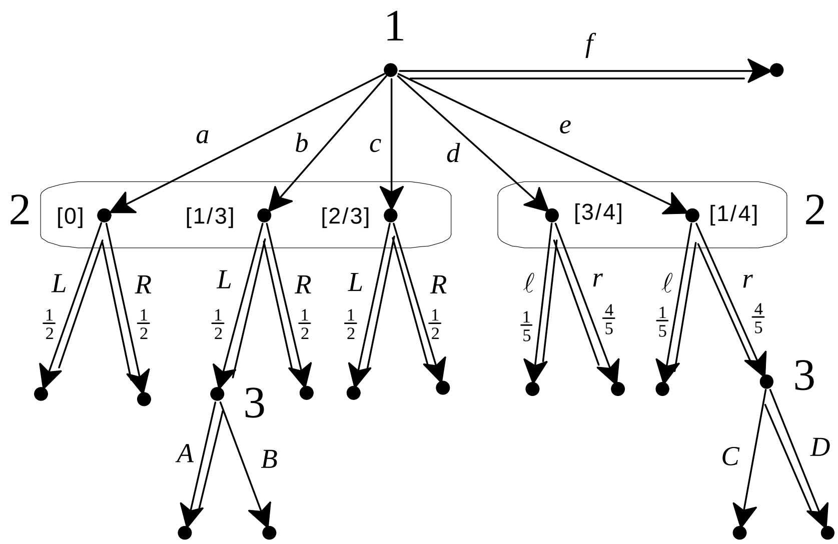

An extensive form and an assessment

Figure 2

Let E be the middle equivalence class of (2): $E = \{b, c, d, e, bL, bR, cL, cR, d\ell, dr, e\ell, er, bLA, erD\}$, so that $E \cap D^+_\mu = \{b, c, d, e, bL, er\}$. There is an infinite number of probability density functions $\nu_E : H \to [0, 1]$ that satisfy the properties of Definition 3: the degree of freedom is given by the relative likelihood of information set $\{a, b, c\}$ to information set $\{d, e\}$, conditional on $\mathbf{t} = 1$. Let $P(\{a, b, c\} \mid \mathbf{t} = 1) = \alpha \in (0, 1)$. Then, conditioning on *decision and terminal* histories in E and on times 1 and 2, and taking $P(\mathbf{t} = 1) = P(\mathbf{t} = 2) = \frac{1}{2}$, one gets the following joint probability distribution (only the positive values are shown), call it $\hat{\nu}_E(\cdot)$:

$$\left(\begin{array}{ccccc} & b & c & d & e \\ t = 1 & \frac{1}{6}\alpha & \frac{1}{3}\alpha & \frac{3}{8}(1-\alpha) & \frac{1}{8}(1-\alpha) \end{array} \right)$$

$$\left(\begin{array}{ccccccccc} & bL & bR & cL & cR & d\ell & dr & e\ell & er \\ t = 2 & \frac{1}{12}\alpha & \frac{1}{12}\alpha & \frac{1}{6}\alpha & \frac{1}{6}\alpha & \frac{3}{40}(1-\alpha) & \frac{3}{10}(1-\alpha) & \frac{1}{40}(1-\alpha) & \frac{1}{10}(1-\alpha) \end{array} \right)$$

Then $\nu_E(h) = \hat{\nu}_E(h \mid E \cap D^+_\mu) = \frac{\hat{\nu}_E(h)}{\hat{\nu}_E(b)+\hat{\nu}_E(c)+\hat{\nu}_E(d)+\hat{\nu}_E(e)+\hat{\nu}_E(bL)+\hat{\nu}_E(er)}$. For example, if one chooses $\alpha = \frac{1}{2}$, then $\nu_E(\cdot)$ is uniquely given by

$$\left(\begin{array}{ccccccc} & b & c & d & e & bL & er \\ t = 1 & \frac{10}{71} & \frac{20}{71} & \frac{45}{142} & \frac{15}{142} & 0 & 0 \\ t = 2 & 0 & 0 & 0 & 0 & \frac{5}{71} & \frac{6}{71} \end{array} \right)$$

which incorporates the following probabilistic judgements (given the information that the actual play of the game is at a decision history in E): (1) the probability that the actual play of the game is at history b is $\nu_E(b) = \frac{10}{71}$, (2) conditional on being at information set $\{a, b, c\}$ (and thus on $\mathbf{t} = 1$) the probabilities are $\mu(b) = \frac{\nu_E(b)}{\nu_E(b)+\nu_E(c)} = \frac{\frac{10}{71}}{\frac{10}{71}+\frac{20}{71}} = \frac{1}{3}$ and $\mu(c) = \frac{2}{3}$, (3) conditional on $\mathbf{t} = 1$

the probabilities are $\begin{pmatrix} b & c & d & e \\ \frac{1}{6} & \frac{2}{6} & \frac{3}{8} & \frac{1}{8} \end{pmatrix}$, (4) conditional on $\mathbf{t}=2$ the probabilities are $\begin{pmatrix} bL & er \\ \frac{5}{11} & \frac{6}{11} \end{pmatrix}$, etc. We summarize the above discussion in the following remark. The important difference between "being *at* history h" and "*reaching* h" was emphasized in [2, 7, 15].[8]

Remark 4 *The interpretation of $\nu_E(h)$ in Definition 3 is "the probability that the actual play of the game is currently* at *history h, conditional on it being in $E \cap D_\mu^+$", assuming that any two (relevant) dates are equally likely. The unique time $t(h)$ at which a decision history h is visited can be taken to be equal to the length of h, denoted by $\ell(h)$, which is defined recursively as follows: $\ell(\emptyset) = 0$ and $\ell(ha) = \ell(h) + 1$. As noted above, one can equivalently define $\nu_E(h)$ as "the probability that the actual play of the game is* at *history h* and *the time is $t(h)$" (conditional on $E \cap D_\mu^+$).*[9]

Definition 5 *An assessment (σ, μ) is a* perfect Bayesian equilibrium *if it is sequentially rational, it is rationalized by a plausibility order on the set of histories and is Bayesian relative to it.*

Remark 6 *It is proved in [6] that if (σ, μ) is a perfect Bayesian equilibrium then σ is a subgame-perfect equilibrium and that every sequential equilibrium is a perfect Bayesian equilibrium.*[10]

For the game illustrated in Figure 3, a perfect Bayesian equilibrium is given by $\sigma = (c, d, g)$, $\mu(a) = \mu(be) = 1$ (σ is highlighted by double edges). In fact (σ, μ) is sequentially rational and, furthermore, it is rationalized by the following plausibility order and is Bayesian relative to it (the trivial density functions on the equivalence classes that contain histories h with $\mu(h) > 0$ are written next to the order):

$$\begin{pmatrix} \emptyset, c \\ a, ad \\ b, bd \\ be, beg \\ ae, aeg \\ bef \\ aef \end{pmatrix}, \qquad \begin{pmatrix} \nu_{\{\emptyset,c\}}(\emptyset) = 1 \\ \nu_{\{a,ad\}}(a) = 1 \\ \nu_{\{be,beg\}}(be) = 1 \end{pmatrix} \tag{3}$$

The belief revision policy encoded in a perfect Bayesian equilibrium can be interpreted either as the epistemic state of an external observer[11] or as a belief revision policy which is shared by all the players. For example, the perfect Bayesian equilibrium $\sigma = (c, d, g)$ and $\mu(a) = \mu(be) = 1$ of the game of Figure 3 reflects the following belief revision policy: the initial beliefs are that Player 1 will play c; conditional on learning that Player 1 did not play c, the observer would become convinced

[8] The probability of reaching history h is the sum of the probabilities of the histories that have h as a prefix and is thus interpreted as the probability that the actual play of the game *is or was* at h. In the above example, conditional on $\mathbf{t} = 1$ or $\mathbf{t} = 2$ and considering decision and terminal histories in E, the probability of reaching b is given by $\hat{\nu}_E(\mathbf{d} = b \textit{ and } \mathbf{t} = 1) + \hat{\nu}_E(\mathbf{d} = bL \textit{ and } \mathbf{t} = 2) + \hat{\nu}_E(\mathbf{d} = bR \textit{ and } \mathbf{t} = 2) = \frac{1}{6}\alpha + \frac{1}{12}\alpha + \frac{1}{12}\alpha = \frac{1}{3}\alpha$ or, conditioning on $E \cap D_\mu^+$ and taking $\alpha = \frac{1}{2}$, by $\nu((\mathbf{d} = b \textit{ and } \mathbf{t} = 1) + \nu(\mathbf{d} = bL \textit{ and } \mathbf{t} = 2) = \frac{10}{71} + \frac{5}{71} = \frac{15}{71}$.

[9] The reason why we take the support of ν_E to be $E \cap D_\mu^+$, rather than E, is that terminal histories as well as decision histories h with $\mu(h) = 0$ are irrelevant for the notion of Bayesian consistency and ν_E so defined is a much simpler object (compare, for instance, the simpler function ν_E with the more extensive function $\hat{\nu}_E$ in the above example).

[10] The example of Figure 1 shows that PBE is a strict refinement of subgame-perfect equilibrium. It is shown in [6] that, in turn, sequential equilibrium is a strict refinement of PBE.

[11] For example, [9] adopt this interpretation. For a subjective interpretation of perfect Bayesian equilibrium see [5].

that Player 1 played a (that is, she would judge a to be strictly more plausible than b) and would expect Player 2 to play d; upon learning that Player 1 did not play c and Player 2 did not play d, the observer would become convinced that Player 1 played b and Player 2 played e, hence judging be to be strictly more plausible than ae, thereby reversing her earlier belief that a was strictly more plausible than b. Note that such a belief revision policy is consistent with the AGM rationality axioms introduced in [1]; however, it is ruled out by the stronger notion of sequential equilibrium.

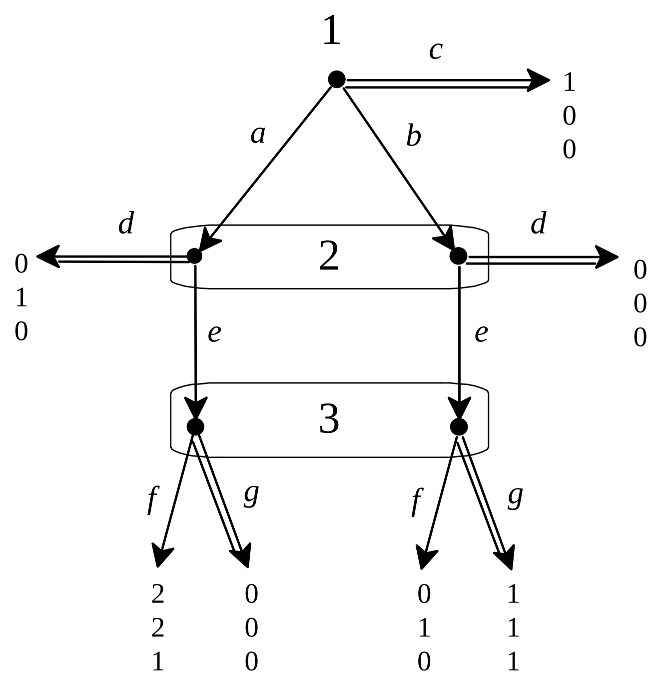

The assessments (1) $\sigma = (c, d, g)$ with $\mu(a) = \mu(be) = 1$ and
(2) $\sigma = (c, d, g)$ with $\mu(a) = \mu(be) = \frac{3}{4}$, $\mu(b) = \mu(ae) = \frac{1}{4}$
are both perfect Bayesian equilibria.
Figure 3

Another perfect Bayesian equilibrium of the game of Figure 3 is $\sigma = (c, d, g)$, $\mu(a) = \frac{3}{4}$, $\mu(b) = \frac{1}{4}$, $\mu(ae) = \frac{1}{4}$ and $\mu(be) = \frac{3}{4}$; it is rationalized by the following plausibility order and is Bayesian relative to it:

$$\begin{pmatrix} \emptyset, c \\ a, ad, b, bd \\ ae, aeg, be, beg \\ aef, bef \end{pmatrix}, \qquad \begin{pmatrix} \nu_{\{\emptyset,c\}}(\emptyset) = 1 \\ \nu_E(a) = \frac{3}{4},\ \nu_E(b) = \frac{1}{4} \\ \nu_F(ae) = \frac{1}{4},\ \nu_F(be) = \frac{3}{4} \end{pmatrix} \tag{4}$$

where $E = \{a, ad, b, bd\}$ and $F = \{ae, aeg, be, beg\}$. In this case the relative *plausibility* of a and b remains unchanged after observing e, since $a \sim b$ and $ae \sim be$; however, the relative *probability* does change, since a is judged to be three times as likely as b but ae is judged to be one third as likely as be. Such a revised judgment is also incompatible with the notion of sequential equilibrium.

We will show in the next section that sequential equilibrium can be characterized as a strenghtening PBE based on two properties: (1) a property of the plausibility order that constrains the supports of the belief system in a way that rules out the phenomenon highlighted by the first PBE of the game of Figure 3 and (2) a strengthening of the notion of Bayesian consistency, which imposes constraints on how the probabilities are distributed over those supports and thereby rules out the phenomenon highlighted by the second PBE of the game of Figure 3.

3 Perfect Bayesian equilibrium and sequential equilibrium.

Given a plausibility order $\precsim$ on the finite set of histories H, a function $F : H \to \mathbb{N}$ (where $\mathbb{N}$ denotes the set of non-negative integers) is said to be an *ordinal integer-valued representation* of $\precsim$ if, for every $h, h' \in H$,

$$F(h) \leq F(h') \text{ if and only if } h \precsim h'. \tag{5}$$

Since H is finite, the set of ordinal integer-valued representations is non-empty. Instead of an ordinal representation of the plausibility order $\precsim$ one could seek a *cardinal* representation which, besides (5), satisfies the following property: if h and h' belong to the same information set and $a \in A(h)$, then

$$F(h') - F(h) = F(h'a) - F(ha). \tag{CM}$$

If we think of F as measuring the "plausibility distance" between histories, then we can interpret CM as a distance-preserving condition: the plausibility distance between two histories in the same information set is preserved by the addition of the same action. The following definition is taken from [6].

Definition 7 *A plausibility order $\precsim$ on the set of histories H is* choice measurable *if it has at least one integer-valued representation that satisfies property CM.*

For example, the plausibility order (3) is not choice measurable, since any integer-valued representation F of it must be such that $F(b) - F(a) > 0$ and $F(be) - F(ae) < 0$. On the other hand, the plausibility order (4) *is* choice measurable, as shown by the following integer-valued representation: $F(\emptyset) = F(c) = 0$, $F(h) = 1$ for all $h \in \{a, ad, b, bd\}$, $F(h) = 2$ for all $h \in \{ae, aeg, be, beg\}$ and $F(aef) = F(bef) = 3$.

Choice measurability plays a crucial role in filling the gap between perfect Bayesian equilibrium and sequential equilibrium. First we recall the definition of sequential equilibrium. An assessment (σ, μ) is *KW-consistent* ('KW' stands for 'Kreps-Wilson') if there is an infinite sequence $\langle \sigma^1, ..., \sigma^m, ... \rangle$ of completely mixed strategy profiles such that, letting μ^m be the unique system of beliefs obtained from σ^m by applying Bayes' rule,[12] $\lim_{m\to\infty}(\sigma^m, \mu^m) = (\sigma, \mu)$. An assessment (σ, μ) is a *sequential equilibrium* if it is KW-consistent and sequentially rational.

It is shown in [6] that if (σ, μ) is a KW-consistent assessment then it is rationalized by a plausibility order that is choice measurable and that the notion of sequential equilibrium is a strict

[12] That is, for every $h \in D\backslash\{\emptyset\}$, $\mu^m(h) = \frac{\prod_{a \in A_h} \sigma^m(a)}{\sum_{h' \in I(h)} \prod_{a \in A_{h'}} \sigma^m(a)}$, where A_h is the set of actions that occur in history h. Since σ^m is completely mixed, $\sigma^m(a) > 0$ for every $a \in A$ and thus $\mu^m(h) > 0$ for all $h \in D\backslash\{\emptyset\}$.

refinement of perfect Bayesian equilibrium. We now show that choice measurability together with a strengthening of Definition 3, which we call uniform Bayesian consistency (Definition 9 below), is *necessary and sufficient* for a perfect Bayesian equilibrium to be a sequential equilibrium.

To motivate the next definition, let (σ, μ) be an assessment which is rationalized by a plausibility order $\precsim$. As before, let D_μ^+ be the set of decision histories to which μ assigns positive probability: $D_\mu^+ = \{h \in D : \mu(h) > 0\}$. Let $\mathcal{E}_\mu^+$ be the set of equivalence classes of $\precsim$ that have a non-empty intersection with D_μ^+. Clearly $\mathcal{E}_\mu^+$ is a non-empty, finite set. Suppose that (σ, μ) is Bayesian relative to $\precsim$ and fix a collection of probability density functions $\{\nu_E\}_{E\in\mathcal{E}_\mu^+}$ that satisfy the properties of Definition 3. We call a probability density function $\nu : D \to (0, 1]$ (recall that D is the set of all decision histories) a *full-support common prior* of $\{\nu_E\}_{E\in\mathcal{E}_\mu^+}$ if, for every $E \in \mathcal{E}_\mu^+$, $\nu_E(\cdot) = \nu(\cdot \mid E \cap D_\mu^+)$, that is, for all $h \in E \cap D_\mu^+$, $\nu_E(h) = \frac{\nu(h)}{\sum_{h'\in E\cap D_\mu^+} \nu(h')}$. Note that a full support common prior assigns positive probability to all decision histories, not only to those in D_μ^+. Since any two elements of $\mathcal{E}_\mu^+$ are mutually disjoint, a full support common prior always exists; indeed, there is an infinite number of them. To see this, let $\mathcal{E}_\mu^+ = \{E_1, ..., E_m\}$ $(m \geq 1)$ and choose arbitrary $\alpha_1, ..., \alpha_m \in (0, 1)$ and $\beta \in [0, 1)$ such that $\alpha_1 + ... + \alpha_m + \beta = 1$ and $\beta > 0$ if and only if $D\backslash D_\mu^+ \neq \varnothing$. Then $\nu(\cdot) = \nu^*(\cdot) + \nu^{**}(\cdot)$ is a full-support common prior of $\{\nu_{E_i}\}_{i\in\{1,...,m\}}$, where $\nu^* : D \to [0, 1]$ is given by $\nu^*(h) = \sum_{i=1}^m [\alpha_i \times \nu_{E_i}(h)]$ (so that $\nu^*(h) > 0$ if and only if $h \in D_\mu^+$) and $\nu^{**} : D \to [0, 1]$ is such that $\nu^{**}(h) > 0$ if and only if $h \in D\backslash D_\mu^+$ and $\sum_{h\in D\backslash D_\mu^+} \nu^{**}(h) = \beta$. Among the many full-support common priors there are some that satisfy further properties, as the following lemma, which is proved in the Appendix, shows.

Lemma 8 *There always exists a full-support common prior ν that satisfies the following property: if $a \in A(h)$ and $ha \in D$, then (A) $\nu(ha) \leq \nu(h)$ and (B) if $\sigma(a) > 0$ then $\nu(ha) = \nu(h) \times \sigma(a)$.*

The following definition requires that, among the many full-support common priors, there be one that, besides the property of Lemma 8, satisfies the additional property that the relative likelihood of any two histories in the same information set be preserved by the addition of the same action.

Definition 9 *Fix an extensive form. Let (σ, μ) be an assessment which is rationalized by the plausibility order $\precsim$. We say that (σ, μ) is* uniformly Bayesian relative to $\precsim$ *if there exists a collection of probability density functions $\{\nu_E\}_{E\in\mathcal{E}_\mu^+}$ that satisfy the properties of Definition 3 and a full-support common prior $\nu : D \to (0, 1]$ of $\{\nu_E\}_{E\in\mathcal{E}_\mu^+}$ that satisfies the following properties.*

- $UB1$. *If $a \in A(h)$ and $ha \in D$, then (A) $\nu(ha) \leq \nu(h)$ and (B) if $\sigma(a) > 0$ then $\nu(ha) = \nu(h) \times \sigma(a)$.*
- $UB2$. *If $a \in A(h)$, h and h' belong to the same information set and $ha, h'a \in D$ then $\frac{\nu(h)}{\nu(h')} = \frac{\nu(ha)}{\nu(h'a)}$.*

We call such a function ν a uniform full-support common prior *of $\{\nu_E\}_{E\in\mathcal{E}_\mu^+}$.*

$UB1$ is the property of Lemma 8, which can always be satisfied by an appropriate choice of a full-support common prior. $UB2$ requires that the relative probability, according to the common prior ν, of any two histories that belong to the same information set remain unchanged by the addition of the same action.

Remark 10 *Choice measurability and uniform Bayesian consistency are independent properties. For example, the perfect Bayesian equilibrium $\sigma = (c, d, g)$ and $\mu(a) = \mu(be) = 1$ of the game of*

Figure 3 is such that any plausibility order that rationalizes it cannot be choice measurable[13] *and yet (σ,μ) is uniformly Bayesian relative to plausibility order (3) that rationalizes it.*[14] *On the other hand, the perfect Bayesian equilibrium $\sigma=(c,d,g)$, $\mu(a)=\mu(be)=\frac{3}{4}$, $\mu(b)=\mu(ae)=\frac{1}{4}$ of the game of Figure 3 is rationalized by the choice measurable plausibility order (4) (as shown above) but it cannot be uniformly Bayesian relative to any plausibility order that rationalizes it.*[15]

We can now state the main result of this paper, namely that choice measurability and uniform Bayesian consistency are necessary and sufficient for a perfect Bayesian equilibrium to be a sequential equilibrium. The proof, which exploits the characterization of sequential equilibrium provided in [12] (for an alternative and similar characterization see [17]), is given in the Appendix.

Proposition 11 *Fix an extensive game and an assessment (σ,μ). The following are equivalent:*

- (I) (σ,μ) *is a perfect Bayesian equilibrium which is rationalized by a choice measurable plausibility order and is uniformly Bayesian relative to it.*
- (II) (σ,μ) *is a sequential equilibrium.*

As an application of Proposition 11 consider the extensive game of Figure 4. Let (σ,μ) be an assessment with $\sigma(a)=\sigma(T)=\sigma(f)=\sigma(L)=1$ (highlighted by double edges; note that σ is a subgame-perfect equilibrium), $\mu(b)>0$ and $\mu(c)>0$. Then (σ,μ) can be rationalized by a choice-measurable plausibility order only if μ is such that[16]

$$\text{either } \mu(bB)=\mu(cBf)=0, \text{ or } \mu(bB)>0 \text{ and } \mu(cBf)>0. \tag{6}$$

If, besides from being rationalized by a choice-measurable plausibility order $\precsim$, (σ,μ) is also uniformly Bayesian relative to $\precsim$ (Definition 9), then[17]

[13]Because, by $P2$ of Definition 2, any such plausibility order $\precsim$ would have to satisfy $a\prec b$ and $be\prec ae$, so that any integer-valued representation F of it would be such that $F(b)-F(a)>0$ and $F(be)-F(ae)<0$.

[14]As can be seen by taking ν to be the uniform distribution over the set $D=\{\emptyset,a,b,ae,be\}$ ($UB1$ is clearly satisfied and $UB2$ is also satisfied, since $\frac{\nu(a)}{\nu(b)}=\frac{\frac{1}{5}}{\frac{1}{5}}=\frac{\nu(ae)}{\nu(be)}$).

[15]Because, by $P2$ Definition 2, any such plausibility order $\precsim$ would have to satisfy $a\sim b$ and $ae\sim be$, so that - letting E be the equivalence class $\{a,b,ad,bd\}$ and F the equivalence class $\{ae,be,aeg,beg\}$ (thus $E\cap D_\mu^+=\{a,b\}$ and $F\cap D_\mu^+=\{ae,be\}$) - if ν is any common prior then $\nu_E(a)=\frac{\nu(a)}{\nu(a)+\nu(b)}$, $\nu_E(b)=\frac{\nu(b)}{\nu(a)+\nu(b)}$. By $B3$ of Definition 3, $\mu(a)=\frac{\nu_E(a)}{\nu_E(a)+\nu_E(b)}$ and $\mu(b)=\frac{\nu_E(b)}{\nu_E(a)+\nu_E(b)}$. Thus $\frac{\nu(a)}{\nu(b)}=\frac{\nu_E(a)}{\nu_E(b)}=\frac{\mu(a)}{\mu(b)}=3$; similarly, $\frac{\nu(ae)}{\nu(be)}=\frac{\nu_F(ae)}{\nu_F(be)}=\frac{\mu(ae)}{\mu(be)}=\frac{1}{3}$, yielding a violation of $UB2$ of Definition 9.

[16]Proof. Let $\precsim$ be a choice measurable plausibility order that rationalizes (σ,μ) and let F be a cardinal representation of it. Since $\mu(b)>0$ and $\mu(c)>0$, by $P2$ of Definition 2, $b\sim c$ and thus $F(b)=F(c)$. By choice measurability, $F(b)-F(c)=F(bB)-F(cB)$ and thus $F(bB)=F(cB)$, so that $bB\sim cB$. Since $\sigma(f)>0$, by $P1$ of Definition 2, $cB\sim cBf$ and therefore, by transitivity of $\precsim$, $bB\sim cBf$. Hence if $\mu(bB)>0$ then, by $P2$ of Definition 2, $bB\in Min_{\precsim}\{bB,cBf,d\}$ (for any $S\subseteq H$, $Min_{\precsim}S$ is defined as $\{h\in S: h\precsim h', \forall h'\in S\}$) and thus $cBf\in Min_{\precsim}\{bB,cBf,d\}$ so that, by $P2$ of Definition 2, $\mu(cBf)>0$. The proof that if $\mu(cBf)>0$ then $\mu(bB)>0$ is similar.

[17]Proof. Suppose that $\mu(b)>0$, $\mu(c)>0$ (so that $b\sim c$) and $\mu(bB)>0$. Let ν be a full-support common prior that satisfies the properties of Definition 9. Then, by $UB2$, $\frac{\nu(c)}{\nu(b)}=\frac{\nu(cB)}{\nu(bB)}$ and, by $UB1$, since $\sigma(f)=1$, $\nu(cBf)=\nu(cB)\times\sigma(f)=\nu(cB)$. Let E be the equivalence class that contains b. Then $E\cap D_\mu^+=\{b,c\}$. Since $\nu_E(\cdot)=\nu(\cdot\mid E\cap D_\mu^+)$, by $B3$ of Defnition 3, $\mu(b)=\frac{\nu(b)}{\nu(b)+\nu(c)}$ and $\mu(c)=\frac{\nu(c)}{\nu(b)+\nu(c)}$, so that $\frac{\mu(c)}{\mu(b)}=\frac{\nu(c)}{\nu(b)}$. Let G be the equivalence class that contains bB. Then, since - by hypothesis - $\mu(bB)>0$, it follows from (6) that either

$$\mu(bB) > 0 \ \Rightarrow \frac{\mu(cBf)}{\mu(bB)} = \frac{\mu(c)}{\mu(b)}. \tag{7}$$

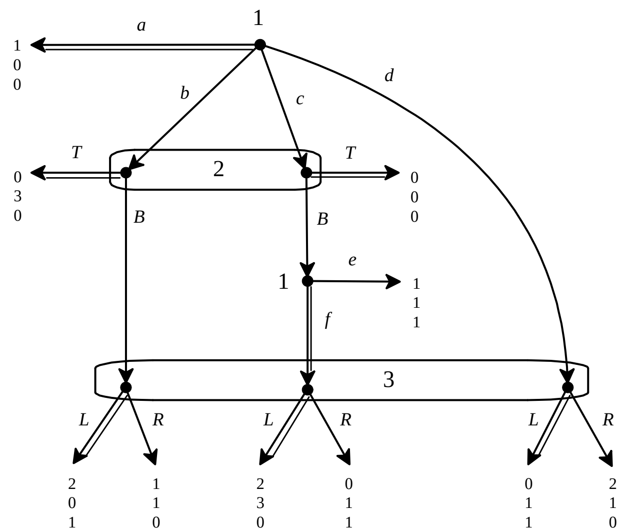

Implications of choice measurability and the uniform Bayesian property for assessments of the form $\sigma(a) = \sigma(T) = \sigma(f) = \sigma(L) = 1$, $\mu(b) > 0$, $\mu(c) > 0$.

Figure 4

Thus, for example, continuing to assume that $\sigma = ((a, f), T, L)$, the assessment $(\sigma, \tilde{\mu})$ with $\tilde{\mu}(b) = \frac{7}{10}$, $\tilde{\mu}(c) = \frac{3}{10}$, $\tilde{\mu}(bB) = \frac{7}{18}$, $\tilde{\mu}(cBF) = \frac{3}{18}$ and $\tilde{\mu}(d) = \frac{8}{18}$ is a sequential equilibrium,[18] while the assessment $(\sigma, \hat{\mu})$ with $\hat{\mu}(b) = \frac{7}{10}$, $\hat{\mu}(c) = \frac{3}{10}$, $\hat{\mu}(bB) = \hat{\mu}(cBf) = \hat{\mu}(d) = \frac{1}{3}$ is a perfect Bayesian equilibrium but not a sequential equilibrium.[19]

$G \cap D_\mu^+ = \{bB, cBf\}$ or $G \cap D_\mu^+ = \{bB, cBf, d\}$. Since $\nu_G(\cdot) = \nu(\cdot \mid G \cap D_\mu^+)$, by $B3$ of Defnition 3, in the former case $\mu(bB) = \frac{\nu(bB)}{\nu(bB)+\nu(cBf)}$ and $\mu(cBf) = \frac{\nu(cBf)}{\nu(bB)+\nu(cBf)}$ and in the latter case $\mu(bB) = \frac{\nu(bB)}{\nu(bB)+\nu(cBf)+\nu(d)}$ and $\mu(cBf) = \frac{\nu(cBf)}{\nu(bB)+\nu(cBf)+\nu(d)}$; thus in both cases $\frac{\mu(cBf)}{\mu(bB)} = \frac{\nu(cBf)}{\nu(bB)}$. Hence, since $\nu(cBf) = \nu(cB)$, $\frac{\mu(cBf)}{\mu(bB)} = \frac{\nu(cB)}{\nu(bB)}$ and, therefore, since - as shown above - $\frac{\nu(cB)}{\nu(bB)} = \frac{\nu(c)}{\nu(b)}$ and $\frac{\nu(c)}{\nu(b)} = \frac{\mu(c)}{\mu(b)}$, we have that $\frac{\mu(cBf)}{\mu(bB)} = \frac{\mu(c)}{\mu(b)}$.

[18] It follows from Proposition 11 and the fact that (σ, μ) is sequentially rational and rationalized by the following choice-measurable plausibility order: $\begin{pmatrix} \precsim \ : & F: \\ & \\ \emptyset, a & 0 \\ b, c, bT, cT & 1 \\ d, bB, cB, cBf, dL, bBL, cBfL & 2 \\ bBR, cBe, cBfR, dR & 3 \end{pmatrix}$ and is uniformly Bayesian relative to it: letting E_1, E_2 and E_3 be the top three equivalence classes, there is a unique collection of probability density functions that satify the properties of Definition 3, namely $\nu_{E_1}(\emptyset) = 1$, $\nu_{E_2}(b) = \frac{7}{10}$, $\nu_{E_2}(c) = \frac{3}{10}$, $\nu_{E_3}(d) = \frac{8}{21}$, $\nu_{E_3}(bB) = \frac{7}{21}$, $\nu_{E_3}(cB) = \nu_{E_3}(cBf) = \frac{3}{21}$; then the following is a full-support uniform common prior: $\nu(\emptyset) = \frac{9}{40}, \nu(b) = \nu(bB) = \frac{7}{40}$, $\nu(c) = \nu(cB) = \nu(cBf) = \frac{3}{40}$, $\nu(d) = \frac{8}{40}$.

[19] Both $(\sigma, \tilde{\mu})$ and $(\sigma, \hat{\mu})$ are sequentially rational and are rationalized by the choice measurable plausibility order given in Footnote 18; $(\sigma, \hat{\mu})$ is Bayesian relative to that plausibility order but cannot be uniformly Bayesian relative to any rationalizing order, because it fails to satisfy (7).

4 Related literature

The notion of perfect Bayesian equilibrium is built on two properties (besides sequential rationality):

(1) rationalizability of the assessment (σ, μ) by a plausibility order (Definition 2) and

(2) Bayesian consistency relative to the plausibility order (Definition 3).

The first property identifies the set of decision histories that can be assigned positive conditional probability by the system of beliefs, while the second property imposes constraints on how conditional probabilities can be distributed over that set in order to guarantee "Bayesian updating whenever possible".[20] The characterization of sequential equilibrium provided in this paper is built on a strengthening of each of those two properties:

(1′) the plausibility order must be choice measurable (Definition 7) and

(2′) the collection of conditional probability density functions identified by Bayesian consistency must be compatible with each other, in the sense that there exists a full-support common prior that preserves the relative likelihood of two decision histories in the same information set when a common action is added (we called this property uniform Bayesian consistency: Definition 9).

This two-part characterization is similar to a two-part characterization of KW-consistency provided in the literature ([10, 12, 17]).[21] This characterization is based on two functions defined on the set of actions A:

(i) the first function concerns the support of the assessment: it is expressed in terms of the labeling K in [10, p. 887], the function ε in [12, p. 241] and the function e in [17, p. 11];

(ii) the second function concerns the distribution of probabilities on that support: it is expressed in terms of the function ξ in [10, p. 888], the function $\bar{\sigma}$ in [12, p. 241] and the function c in [17, p. 11]; this function is used to define "pseudo" probabilities on decision histories.

Our definition of choice measurability is a reformulation of (i) that shows it to be a strengthening of the notion of AGM consistency (Definition 2), while the notions of Bayesian consistency and uniform Bayesian consistency disentangle two properties implied by (ii): "Bayesian updating whenever possible" and "preservation of the relative likelihood of any two histories in the same information set after the addition of a common action" (the first property constitutes Bayesian consistency and the conjunction of both properties consitutes uniform Bayesian consistency).

The proof of Proposition 11 in the Appendix shows precisely how to translate the two-part characterization of sequential equilibrium given in [12] (a similar characterization is provided in [17]) into the characterization offered in this paper and *vice versa*. Here we comment on other relevant contributions.

In [18] the author shows the following (using our notation). A set of actions and histories $B \subseteq A \cup H$ is called a "basement" if it coincides with the support of at least one assessment, that is, if there is an assessment (σ, μ) such that $a \in B$ if and only if $\sigma(a) > 0$ and $h \in B$ if and only if $\mu(h) > 0$. Given a basement, one can construct a partial relation $\precsim$ on the set of histories H as

[20] By "Bayesian updating whenever possible" we mean the following: (1) when information causes no surprises, because the actual play of the game is consistent with the most plausible play(s) (that is, when information sets are reached that have positive prior probability), then beliefs should be updated using Bayes' rule and (2) when information is surprising (that is, when an information set is reached that had zero prior probability) then new beliefs can be formed in an arbitrary way, but from then on Bayes' rule should be used to update those new beliefs, whenever further information is received that is consistent with those beliefs.

[21] [4] provides an indirect proof of the fact that consistent assessments are determined by finitely many algebraic equations and inequalities.

follows: (1) if h and h' belong to the same information set then $h \sim h'$ if $h, h' \in B$ and $h \prec h'$ if $h \in B$ and $h' \notin B$, (2) if $a \in A(h)$ then $h \sim ha$ if $a \in B$ and $h \prec ha$ if $a \notin B$. A total pre-order $\precsim^*$ on H is said to be "additively representable" if there exists a function $\lambda : A \to \mathbb{R}$ such that $h \precsim^* h'$ if and only if $\sum_{a \in A_h} \lambda(a) \leq \sum_{a \in A_{h'}} \lambda(a)$ (A_h denotes the set of actions that appear in history h).

Proposition 12 *[18, Theorem 1 and Corollary 1, pp. 17 and 20] If the relation $\precsim$ derived from a basement B can be extended to a total pre-order $\precsim^*$ that has an additive representation, then there exists a KW-consistent assessment (σ, μ) whose support coincides with B. Conversely, given a KW-consistent assessment (σ, μ), the relation $\precsim$ derived from its basement can be extended to a total pre-order $\precsim^*$ that has an additive representation.*

There is a clear connection between Proposition 12 and Proposition 11. However, while Proposition 12 characterizes basements that are supported by a KW-consistent assessment, Proposition 11 focuses on a particular PBE (σ, μ) and on the conditions that are necessary and sufficient for (σ, μ) to be a sequential equilibrium. One of these conditions is choice measurability of the plausibility order that rationalizes (σ, μ), which is equivalent to the existence of an additive representation of a total pre-order that extends the plausibility relation obtained from the support of (σ, μ). The other condition is that (σ, μ) be uniformly Bayesian relative to the plausibility order (there is no counterpart to this property in Proposition 12, since it only deals with supports).[22]

The characterizations of sequential equilibrium provided in [10, 12, 17] and in Proposition 11 do not make any use of sequences and limits. A "limit free" characterization of sequential equilibrium is also provided in [9] using relative probability spaces (Ω, ρ) (which express the notion of an event being infinitely less likely than another) and random variables $s_i : \Omega \to S_i$ (where S_i is the set of pure strategies of player i), representing the beliefs of an external observer (who can assess the relative probabilities of any two strategy profiles, even those that have zero probability); the authors provide a characterization of KW-consistency in terms of the notion of strong independence for relative probability spaces (and, in turn, a characterization of strong independence in terms of weak independence and exchangeability).[23]

5 Conclusion

Besides sequential rationality, the notion of perfect Bayesian equilibrium introduced in [6] is based on two elements: (1) the qualitative notions of plausibility order and AGM-consistency and (2) the notion of Bayesian consistency relative to the plausibility order. In this paper we showed that by strengthening these two conditions one obtains a characterization of sequential equilibrium. The strengthening of the first condition is that the plausibility order that rationalizes the given assessment be choice measurable, that it, that there be a cardinal representation of it (which can be interpreted as measuring the plausibility distance between histories in a way that is preserved by the addition of a common action). The strengthening of the second condition imposes "uniform consistency" on the conditional probability density functions on the equivalence classes of the plausibility order, by requiring that there be a full-support common prior that preserves the relative probabilities of two decision histories in the same information set when a common action is added.

[22] At the 13^{th} SAET conference in July 2013 Streufert presented a characterization of KW-consistent assessments in terms of additive plausibility and a condition that he called "pseudo-Bayesianism" which is essentially a reformulation of one of the conditions given in [17, Theorem 2.1, p. 11].

[23] [3] shows that in games with observable deviators weak independence suffices for KW-consistency.

Although the characterization provided in this paper is strictly related to earlier characterizations, it offers a novel understanding of sequential equilibrium, expressed in terms of a strengthening of the notion of perfect Bayesian equilibrium.[24]

A Appendix: proofs

Proof of Lemma 8. We shall construct a full-support common prior that satisfies the properties of Lemma 8. Let $\mathcal{E}$ be the finite collection of equivalence classes of $\precsim$ and let $(E_1, ..., E_m)$ be the ordering of $\mathcal{E}$ according to decreasing plausibility, that is, $\forall h, h' \in H$, $\forall i, j \in \{1, ..., m\}$, if $h \in E_i$ and $h' \in E_j$ then $h \prec h'$ if and only if $i < j$. Let $\mathcal{E}^+ = \left\{E \in \mathcal{E} : E \cap D_\mu^+ \neq \varnothing\right\}$ and let $\mathcal{N} = \{\nu_E\}_{E \in \mathcal{E}^+}$ be an arbitrary collection of probability density functions that satisfy the properties of Definition 3. Fix an equivalence class E_i and define the function $f_{E_i} : H \to (0, 1]$ recursively as follows.
Step 0. For every $h \notin E_i$ set $f_{E_i}(h) = 0$.

Step 1. For every $h \in E_i \cap D_\mu^+$ set $f_{E_i}(h) = \nu_{E_i}(h)$, where $\nu_{E_i}(\cdot)$ is the relevant element of $\mathcal{N}$. Note that, by Property $B2$ of Definition 3, if $h, ha \in E_i \cap D_\mu^+$ then $f_{E_i}(ha) = f_{E_i}(h) \times \sigma(a)$.

Step 2. Let $h, ha \in E_i$; then (a) if $h \notin D_\mu^+$ and $ha \in D_\mu^+$, set $f_{E_i}(h) = \frac{f_{E_i}(ha)}{\sigma(a)}$ (note that, by $P1$ of Definition 2, $h \sim ha$ implies $\sigma(a) > 0$) and (b) if $h \in D_\mu^+$ and $ha \notin D_\mu^+$, set $f_{E_i}(ha) = f_{E_i}(h) \times \sigma(a)$. Note that, because of Property $B2$ of Definition 3, the values assigned under Step 2 cannot be inconsistent with the values assigned under Step 1.[25]

Step 3. After completing Steps 1 and 2, the only histories $h \in E_i$ for which $f_{E_i}(h)$ has not been defined yet are those that satisfy the following properties: (1) there is no prefix h' of h such that $h' \in E_i \cap D_\mu^+$ and (2) there is no $h' \in E_i \cap D_\mu^+$ such that h is a prefix of h'. Let $\hat{E}_i \subseteq E_i$ be the set of such histories (it could be that $\hat{E}_i = \varnothing$). A maximal path in $\hat{E}_i$ is a sequence $\langle h, ha_1, ha_1a_2, ..., ha_1a_2...a_p\rangle$ in $\hat{E}_i$ such that $ha_1a_2...a_p$ is a terminal history and there is no $h' \in \hat{E}_i$ which is a proper prefix of h. Fix an arbitrary maximal path $\langle h, ha_1, ..., ha_1a_2...a_p\rangle$ in $\hat{E}_i$ and define $f_{E_i}(h) = 1$ and, for every $i = 0, ..., p-1$, $f_{E_i}(ha_1...a_{i+1}) = f_{E_i}(ha_1...a_i) \times \sigma(a_{i+1})$ (defining ha_0 to be h; note that, by $P1$ of Definition 2 $\sigma(a_i) > 0$ for all $i = 1, ..., p$).

By construction, the function f_{E_i} satisfies the following property:

$$\begin{aligned} &\forall h \in E_i, \forall a \in A(h), \text{ if } \sigma(a) > 0 \text{ then } (ha \in E_i \text{ and}) \\ &f_{E_i}(ha) = f_{E_i}(h) \times \sigma(a) \text{ and thus } f_{E_i}(ha) \leq f_{E_i}(h). \end{aligned} \tag{8}$$

Note that Steps 1-3 always assign positive values in $(0, 1]$; thus if $h \in E_i$ then $f_{E_i}(h) \in (0, 1]$.

Next we show that there exist weights $\lambda_1, ..., \lambda_m \in (0, 1)$ such that $\lambda_1 + ... + \lambda_m < 1$ and, $\forall i, j \in \{1, ..., m\}$, if $h \in E_i$, $h' \in E_j$ and $i < j$ then $\lambda_j \times f_{E_j}(h') \leq \lambda_i \times f_{E_i}(h)$. Fix an arbitrary $\lambda_1 \in (0, 1)$ and let $a = \min_{h \in E_1} \{\lambda_1 \times f_{E_1}(h)\}$ (clearly, $a \in (0, 1)$). Let $b = \max_{h \in E_2} \{f_{E_2}(h)\}$ and

[24] Kreps and Wilson themselves [10, p. 876] express dissatisfaction with their definition of sequential equilibrium: "We shall proceed here to develop the properties of sequential equilibrium as defined above; however, we do so with some doubts of our own concerning what 'ought' to be the definition of a consistent assessment that, with sequential rationality, will give the 'proper' definition of a sequential equilibrium." In a similar vein, Osborne and Rubinstein [11, p. 225] write "we do not find the consistency requirement to be natural, since it is stated in terms of limits; it appears to be a rather opaque technical assumption". In these quotations "consistency" corresponds to what we called "KW-consistency".

[25] In the sense that If $h, h' \in E \cap D_\mu^+$ and $h' = ha_1...a_m$ then by Step 1 (and $B2$ of Definition 3) $f_{E_i}(h') = f_{E_i}(h) \times \sigma(a_1) \times ... \times \sigma(a_m)$ and by Step 2 - if applicable - for every $j = 1, ..., m$, $f_{E_i}(ha_1...a_j) = f_{E_i}(h) \times \sigma(a_1) \times ... \times \sigma(a_j)$.

choose a $\lambda_2 \in (0,1)$ such that (1) $\lambda_1 + \lambda_2 < 1$ and (2) $\lambda_2 \times b \leq a$. Then, for every $h \in E_1$ and $h' \in E_2$, $\lambda_2 \times f_{E_2}(h') \leq \lambda_1 \times f_{E_1}(h)$. Repeat this procedure for choosing a weight λ_i for every $i \in \{3,...,m\}$. Define $\alpha_i = \frac{\lambda_i}{\lambda_1+...+\lambda_m}$ and $\bar{\nu} : H \to (0,1]$ by $\bar{\nu}(h) : \sum_{i=1}^{m} [\alpha_i \times f_{E_i}(h)]$. We want to show that, $\forall h \in H$, $\forall a \in A(h)$

$$\begin{array}{c} \text{if } ha \in D \text{ then } (A)\ \bar{\nu}(ha) \leq \bar{\nu}(h) \text{ and} \\ (B) \text{ if } \sigma(a) > 0 \text{ then } \bar{\nu}(ha) = \bar{\nu}(h) \times \sigma(a). \end{array} \tag{9}$$

Fix an arbitrary $ha \in D$. Suppose first that $h \prec ha$. Then, by Property $P1$ of Definition 2, $\sigma(a) = 0$, so that (B) of (9) is trivially satisfied. Let E_i be the equivalence class to which h belongs and E_j the equivalence class to which ha belongs, so that $i < j$. Then $\bar{\nu}(h) = \alpha_i \times f_{E_i}(h)$ and $\bar{\nu}(ha) = \alpha_j \times f_{E_j}(ha)$ and thus, since $\lambda_j \times f_{E_j}(ha) \leq \lambda_i \times f_{E_i}(h)$, dividing both sides by $(\lambda_1+...+\lambda_m)$ we get that $\bar{\nu}(ha) \leq \bar{\nu}(h)$. Suppose now that $h \sim ha$ (so that, by Property $P1$ of Definition 2, $\sigma(a) > 0$). Let E_i be the equivalence class to which both h and ha belong; then, $\bar{\nu}(h) = \alpha_i \times f_{E_i}(h)$ and $\bar{v}(ha) = \alpha_i \times f_{E_i}(ha)$ and thus (9) follows from (8).

Finally, define $\nu : D \to (0,1]$ by $\nu(h) = \frac{\bar{\nu}(h)}{\sum_{h' \in D} \bar{\nu}(h')}$. Then it follows from (9) that if $ha \in D$ then (A) $\nu(ha) \leq \nu(h)$ and (B) if $\sigma(a) > 0$ then $\nu(ha) = \nu(h) \times \sigma(a)$. It only remains to prove that ν is a common prior of $\mathcal{N} = \{\nu_E\}_{E \in \mathcal{E}^+}$, that is, that, $\forall i \in \{1,...,m\}$ and $\forall h \in E_i \cap D_\mu^+$, $v(h \mid E_i \cap D_\mu^+) \overset{def}{=} \frac{\nu(h)}{\sum_{h' \in E_i \cap D_\mu^+} \nu(h')} = \nu_{E_i}(h)$, where $\nu_{E_i}(\cdot)$ is the relevant element of $\mathcal{N}$. Now, $\frac{\nu(h)}{\sum_{h' \in E_i \cap D_\mu^+} \nu(h')} = \frac{\nu(h) \times \sum_{h'' \in D} \bar{\nu}(h'')}{\sum_{h' \in E_i \cap D_\mu^+} \nu(h') \times \sum_{h'' \in D} \bar{\nu}(h'')} = \frac{\bar{\nu}(h)}{\sum_{h' \in E_i \cap D_\mu^+} \left[\nu(h') \times \sum_{h'' \in D} \bar{\nu}(h'')\right]} = \frac{\bar{\nu}(h)}{\sum_{h' \in E_i \cap D_\mu^+} \bar{\nu}(h')} = \frac{\alpha_i \times f_{E_i}(h)}{\sum_{h' \in E_i \cap D_\mu^+} \left[\alpha_i \times f_{E_i}(h')\right]} = \frac{\alpha_i \times \nu_{E_i}(h)}{\sum_{h' \in E_i \cap D_\mu^+} \left[\alpha_i \times \nu_{E_i}(h')\right]} = \frac{\alpha_i \times \nu_{E_i}(h)}{\alpha_i \times \sum_{h' \in E_i \cap D_\mu^+} \nu_{E_i}(h')} = \nu_{E_i(h)}$ since, by construction, for all $h \in E_i \cap D_\mu^+$, $f_{E_i}(h) = \nu_{E_i}(h)$ and $Supp(\nu_{E_i}) = E_i \cap D_\mu^+$. ■

In order to prove Proposition 11 we will exploit the characterization of sequential equilibrium given in [12]. First some notation and terminology. Let A be the set of actions (as in [6], we assume that no action is available at more than one information set, that is, if $h' \notin I(h)$ then $A(h') \neq A(h)$). If h is a history, we denote by A_h the set of actions that occur in history h (thus while h is a *sequence* of actions, A_h is the *set* of actions in that sequence; note that, for every history h, $A_h \neq \varnothing$ if and only if $h \neq \emptyset$). Given an assessment (σ, μ) we denote by $A^0 = \{a \in A : \sigma(a) = 0\}$ the set of actions that are assigned zero probability by the strategy profile σ. Recall that D_μ^+ denotes the set of decision histories to which μ assigns positive probability ($D_\mu^+ = \{h \in D : \mu(h) > 0\}$) and that $h' \in I(h)$ mans that h and h' belong to the same information set. A *pseudo behavior strategy profile* (PBSP) is a generalization of the notion of behavior strategy profile that allows the sum of the "probabilities" over the actions at an information set to be larger than 1, that is, a PBSP is a function $\bar{\sigma} : A \to [0,1]$. A PBSP $\bar{\sigma}$ is a *completely mixed extension* of a behavior strategy profile σ if, $\forall a \in A$, (1) $\bar{\sigma}(a) > 0$ and (2) if $\sigma(a) > 0$ then $\bar{\sigma}(a) = \sigma(a)$. Given a PBSP $\bar{\sigma}$, for every history h let $\mathbb{P}_{\bar{\sigma}}(h) = \begin{cases} 1 & \text{if } h = \emptyset \\ \bar{\sigma}(a_1) \times ... \times \bar{\sigma}(a_m) & \text{if } h = a_1...a_m \end{cases}$.

The following proposition is proved in [12] (see also [13, p. 74] and [17]).

Proposition 13 *([12, Theorem 3.1, p. 241]) Fix an extensive game and let (σ, μ) be an assessment. Then the following are equivalent:*

(A) $(A.1)$ *There exists a function $\varepsilon : A^0 \to (0,1)$ such that, $\forall h, h' \in D$ with $h' \in I(h)$,*

$(A.1a)$ *If $h, h' \in D_\mu^+$ then $\prod_{a\in A^0\cap A_h} \varepsilon(a) = \prod_{a\in A^0\cap A_{h'}} \varepsilon(a)$, and*

$(A.1b)$ *if $h \in D_\mu^+$ and $h' \notin D_\mu^+$ then $\prod_{a\in A^0\cap A_h} \varepsilon(a) > \prod_{a\in A^0\cap A_{h'}} \varepsilon(a)$, and*

$(A.2)$ *there is a PBSP $\bar{\sigma}$ which is a completely mixed extension of σ and is such that, $\forall h, h' \in D_\mu^+$ with $h' \in I(h)$, $\frac{\mathbb{P}_{\bar{\sigma}}(h)}{\mathbb{P}_{\bar{\sigma}}(h')} = \frac{\mu(h)}{\mu(h')}$.*

(B) *(σ, μ) is KW-consistent.*

Proof of Proposition 11. $(I) \Rightarrow (II)$. Let (σ, μ) be a perfect Bayesian equilibrium which is rationalized by a choice measurable plausibility order $\precsim$ and is uniformly Bayesian relative to it. We need to show that (σ, μ) is a sequential equilibrium. Since (σ, μ) is a perfect Bayesian equilibrium, it is sequentially rational and thus we only need to show that (σ, μ) is KW-consistent. We shall use choice measurability (Definition 7) to obtain the function ε of Proposition 13 (a similar argument can be found in [18]) and the full-support common prior ν of Definition 9 to obtain the PBSP $\bar{\sigma}$.

By hypothesis $\precsim$ is choice measurable. Fix a cardinal integer-valued representation F of $\precsim$ and normalize it so that $F(\emptyset) = 0$.[26] For every action $a \in A$, define $\varepsilon(a) = e^{[F(h) - F(ha)]}$ for some h such that $a \in A(h)$. By definition of choice measurability, if $h' \in I(h)$ then $F(h) - F(ha) = F(h') - F(h'a)$ and thus the function $\varepsilon(a)$ is well defined. Furthermore, if $a \in A^0$ (that is, $\sigma(a) = 0$) it follows from $P1$ of Definition 2 that $h \prec ha$ and thus $F(h) - F(ha) < 0$ so that $0 < \varepsilon(a) < 1$, while if $a \notin A^0$ (that is, $\sigma(a) > 0$) then, by $P1$ of Definition 2, $h \sim ha$ and thus $F(h) - F(ha) = 0$ so that $\varepsilon(a) = 1$; hence,

$$\text{if } A^0 \cap A_h \neq \varnothing, \text{ then } \prod_{a\in A_h} \varepsilon(a) = \prod_{a\in A^0\cap A_h} \varepsilon(a). \tag{10}$$

We want to show that the restriction of $\varepsilon(\cdot)$ to A^0 satisfies $(A.1)$ of Proposition 13. Fix an arbitrary history $h = a_1 a_2 ... a_m$. Since $[F(\emptyset) - F(a_1)] + [F(a_1) - F(a_1 a_2)] + ... + [F(a_1 ... a_{m-1}) - F(h)] = -F(h)$ (recall that $F(\emptyset) = 0$), it follows that,

$$\forall h \in H \backslash \{\emptyset\}, \quad \prod_{a\in A_h} \varepsilon(a) = e^{-F(h)}. \tag{11}$$

Fix arbitrary $h, h' \in D\backslash\{\emptyset\}$ with $h' \in I(h)$ and $h \in D_\mu^+$ (that is, $\mu(h) > 0$).[27] Suppose first that $h' \in D_\mu^+$. Then, by $P2$ of Definition 2, $h \sim h'$ and thus $F(h) = F(h')$, so that, by (11), $\prod_{a\in A_h} \varepsilon(a) = \prod_{a\in A_{h'}} \varepsilon(a)$. Thus, by (10), $(A.1a)$ of Proposition 13 is satisfied. Suppose now that $h' \notin D_\mu^+$. Then, by $P2$ of Definition 2, $h \prec h'$ and thus $F(h) < F(h')$ so that $e^{-F(h')} < e^{-F(h)}$ and thus, by (11), $\prod_{a\in A_h} \varepsilon(a) > \prod_{a\in A_{h'}} \varepsilon(a)$. Thus, by (10), $(A.1b)$ of Proposition 13 is also satisfied.

[26] Recal that $\emptyset$ denotes the null history, that is, the root of the tree. If $\hat{F}$ is an integer-valued representation of $\precsim$ that satisfies property CM, then F defined by $F(h) = \hat{F}(h) - \hat{F}(\emptyset)$ is also an integer-valued representation of $\precsim$ that satisfies property CM; clearly, $F(\emptyset) = 0$.

[27] If $h = \emptyset$ then $h' = h$ and there is nothing to prove because $A_h = A^0 \cap A_h = \varnothing$.

Next we prove $(A.2)$. Denote by $\bar{A}$ the set of "non-terminal actions", that is, $\bar{A} = \{a \in A : ha \in D \text{ for some } h \text{ with } a \in A(h)\}$. Let ν be a uniform full-support common prior (Definition 9). Define $\bar{\sigma} : \bar{A} \to (0,1]$ as follows: $\bar{\sigma}(a) = \frac{\nu(ha)}{\nu(h)}$ for some h such that $a \in A(h)$ and $ha \in D$. By Property $UB2$ of Definition 9, if $h' \in I(h)$ then $\frac{\nu(h'a)}{\nu(h')} = \frac{\nu(ha)}{\nu(h)}$ and thus $\bar{\sigma}$ is well defined; furthermore, since $\nu(h) > 0$ for all $h \in D$, $\bar{\sigma}(a) > 0$. By (A) of Property $UB1$ of Definition 9, $\bar{\sigma}(a) \leq 1$. Finally, by (B) of Property $UB1$ of Definition 9, if $\sigma(a) > 0$ then $\bar{\sigma}(a) = \sigma(a)$. Thus $\bar{\sigma}$ is a PBSP which is a completely mixed extension of σ. We need to show that $\forall h, h' \in D_\mu^+$ with $h' \in I(h)$, $\frac{\mathbb{P}_{\bar{\sigma}}(h)}{\mathbb{P}_{\bar{\sigma}}(h')} = \frac{\mu(h)}{\mu(h')}$. If $h = \emptyset$ it is trivially true because $h' = h$ and $\mathbb{P}_{\bar{\sigma}}(\emptyset) = \mu(\emptyset) = 1$. Fix arbitrary $h, h' \in D_\mu^+ \setminus \{\emptyset\}$ with $h' \in I(h)$. Let $h = a_1 a_2 ... a_p$ $(p \geq 1)$ and $h' = b_1 b_2 ... b_r$ $(r \geq 1)$. By definition of $\bar{\sigma}$, $\mathbb{P}_{\bar{\sigma}}(h) = \bar{\sigma}(a_1) \times \bar{\sigma}(a_2) \times ... \times \ \bar{\sigma}(a_p) = \frac{\nu(a_1)}{\nu(\emptyset)} \times \frac{\nu(a_1 a_2)}{\nu(a_1)} \times ... \times \frac{\nu(h)}{\nu(a_1 a_2 ... a_{p-1})} = \frac{\nu(h)}{\nu(\emptyset)}$. Similarly, $\mathbb{P}_{\bar{\sigma}}(h') = \frac{\nu(h')}{\nu(\emptyset)}$. Thus $\frac{\mathbb{P}_{\bar{\sigma}}(h)}{\mathbb{P}_{\bar{\sigma}}(h')} = \frac{\nu(h)}{\nu(h')}$. Dividing numerator and denominator of the right-hand-side by $\sum\limits_{h'' \in E_i \cap D_\mu^+} \nu(h'')$ and using the fact that (since ν is a common prior) $\frac{\nu(h)}{\sum\limits_{h'' \in E_i \cap D_\mu^+} \nu(h'')} = \nu_{E_i}(h)$ and $\frac{\nu(h')}{\sum\limits_{h'' \in E_i \cap D_\mu^+} \nu(h'')} = \nu_{E_i}(h')$, where E_i is the equivalence class to which both h and h' belong, we get that $\frac{\mathbb{P}_{\bar{\sigma}}(h)}{\mathbb{P}_{\bar{\sigma}}(h')} = \frac{\nu_{E_i}(h)}{\nu_{E_i}(h')}$; now, dividing numerator and denominator of the right-hand-side by $\sum\limits_{h'' \in I(h)} \nu_{E_i}(h'')$ and using the fact that, by $B3$ of Definition 3, $\frac{\nu_{E_i}(h)}{\sum\limits_{h'' \in I(h)} \nu_{E_i}(h'')} = \mu(h)$ and $\frac{\nu_{E_i}(h')}{\sum\limits_{h'' \in I(h)} \nu_{E_i}(h'')} = \mu(h')$, we obtain $\frac{\mathbb{P}_{\bar{\sigma}}(h)}{\mathbb{P}_{\bar{\sigma}}(h')} = \frac{\mu(h)}{\mu(h')}$, so that $(A.2)$ of Proposition 13 also holds. Hence, by Proposition 13, (σ, μ) is KW-consistent.

$(II) \Rightarrow (I)$. Let (σ, μ) be a sequential equilibrium. That (σ, μ) is rationalized by a choice measurable plausibility order $\precsim$ and is Bayesian relative to it was proved in [6]. Thus we only need to show that it is uniformly Bayesian. By Proposition 13, there exists a completely mixed PBSP $\bar{\sigma}$ that extends σ and is such that,

$$\forall h, h' \in D_\mu^+ \text{ with } h' \in I(h), \quad \frac{\mathbb{P}_{\bar{\sigma}}(h)}{\mathbb{P}_{\bar{\sigma}}(h')} = \frac{\mu(h)}{\mu(h')}. \tag{12}$$

Define $\bar{\nu} : D \to (0,1]$ recursively as follows: $\nu(\emptyset) = 1$ and, if $a \in A(h)$ and $ha \in D$, $\bar{\nu}(ha) = \bar{\nu}(h) \times \bar{\sigma}(a)$. Since, $\forall a \in A$, $\bar{\sigma}(a) \in (0,1]$ and $\bar{\sigma}(a) = \sigma(a)$ whenever $\sigma(a) > 0$, it follows that

$$\begin{array}{c} \text{if } a \in A(h) \text{ and } ha \in D, \text{ then } (A)\ \bar{\nu}(ha) \leq \bar{\nu}(h) \text{ and} \\ (B) \text{ if } \sigma(a) > 0 \text{ then } \bar{\nu}(ha) = \bar{\nu}(h) \times \sigma(a). \end{array} \tag{13}$$

Define the probability density function $\nu : D \to (0,1]$ by $\nu(h) = \frac{\bar{\nu}(h)}{\sum\limits_{h' \in D} \bar{\nu}(h')}$. Then, by (13), ν satisfies Property $UB1$ of Definition 9. Furthermore, if $a \in A(h)$, h and h' belong to the same information set and $ha, h'a \in D$, then $\frac{\bar{\nu}(ha)}{\bar{\nu}(h'a)} = \frac{\bar{\nu}(h) \times \bar{\sigma}(a)}{\bar{\nu}(h') \times \bar{\sigma}(a)} = \frac{\bar{\nu}(h)}{\bar{\nu}(h')}$ and thus, dividing numerator and denominator by $\sum\limits_{h'' \in D} \bar{\nu}(h'')$, we get that ν satisfies Property $UB2$ of Definition 9.

Furthermore, as shown above,

$$\forall h \in D\backslash\{\emptyset\},\ \mathbb{P}_{\bar{\sigma}}(h) \overset{def}{=} \prod_{a\in A_h} \bar{\sigma}(a) = \frac{\bar{\nu}(h)}{\bar{\nu}(\emptyset)} = \bar{\nu}(h) \ \ (\text{since } \nu(\emptyset)=1), \tag{14}$$

it follows from (12) and (14) that

$$\forall h, h' \in D_\mu^+ \text{ with } h' \in I(h),\ \ \frac{\bar{\nu}(h)}{\bar{\nu}(h')} = \frac{\mu(h)}{\mu(h')} \text{ and thus } \frac{\nu(h)}{\nu(h')} = \frac{\mu(h)}{\mu(h')}. \tag{15}$$

Fix an arbitrary equivalence class E of the plausibility order that rationalizes (σ, μ) such that $E \cap D_\mu^+ \neq \varnothing$ and define $\nu_E : H \to [0,1]$ as follows:

$$\nu_E(h) = \begin{cases} \frac{\nu(h)}{\sum_{h'\in E\cap D_\mu^+} \nu(h')} & \text{if } h \in E \cap D_\mu^+ \\ 0 & \text{if } h \notin E \cap D_\mu^+. \end{cases} \tag{16}$$

By construction ν_E satisfies Property $B1$ of Definition 3 and, by $UB1$ of Definition 9 (proved above), ν_E satisfies also $B2$ of Definition 3 (recall that if $h, h' \in E$ with $h' = ha_1...a_m$ then, by $P1$ of Definition 2, $\sigma(a_i) > 0$ for all $i = 1, ..., m$). It only remains to prove that Property $B3$ of Definition 3 is satisfied, namely that if $h \in E\cap D_\mu^+$ then, for every $h' \in I(h)$, $\mu(h') = \frac{\nu_E(h')}{\sum_{h''\in I(h)} \nu_E(h'')}$. Number the elements of $E\cap D_\mu^+$ from 1 to m in such a way that $h_1 = h$ and the first p elements belong to $I(h_1)$ and the remaining elements (if any) do not belong to $I(h_1)$, that is, $E \cap D_\mu^+ = \{h_1, ..., h_p, h_{p+1}, ..., h_m\}$ with $h_1 = h$, $I(h_1) \cap E \cap D_\mu^+ = \{h_1, ..., h_p\}$ and, for $i > p$, $h_i \notin I(h_1)$. We shall prove that

$$\frac{\nu_E(h_1)}{\sum_{h''\in I(h_1)} \nu_E(h'')} = \mu(h_1). \tag{17}$$

The proof for $1 < j \leq p$ is similar. By (15), for every $j = 1, ..., m$, $\frac{\nu(h_j)}{\nu(h_1)} = \frac{\mu(h_j)}{\mu(h_1)}$. Thus

$$\frac{\sum_{j=1}^p \nu(h_j)}{\nu(h_1)} = \frac{\sum_{j=1}^p \mu(h_j)}{\mu(h_1)}. \tag{18}$$

By definition of μ, $\sum_{j=1}^p \mu(h_j) = 1$ (since, for any $h' \in I(h_1)$ that does not belong to $E \cap D_\mu^+$, $\mu(h') = 0$: recall that, by Property $P2$ of Definition 2, if $h' \in I(h_1)$ is such that $\mu(h') > 0$ then $h' \sim h$, that is, $h' \in E$). Hence $\frac{\nu(h_1)}{\sum_{j=1}^p \nu(h_j)} = \mu(h_1)$. By (16), dividing numerator and denominator of left-hand-side by $\sum_{i=1}^m \nu(h_i)$ we obtain

$$\frac{\nu_E(h_1)}{\sum_{j=1}^p \nu_E(h_j)} = \mu(h_1) \tag{19}$$

Since, by (16), for any $h' \in I(h_1)$ that does not belong to $E \cap D_\mu^+$, $\nu_E(h') = 0$, $\sum_{h''\in I(h_1)} \nu_E(h'') = \sum_{j=1}^p \nu_E(h_j)$. Thus (19) yields the desired (17). Since, by construction, ν is a full-support common prior to the collection of probability density functions ν_E given in (16), which have been shown to satisfy the properties of Definition 3, the proof that (σ, μ) is uniformly Bayesian is complete. ■

References

[1] Alchourrón, Carlos, Peter Gärdenfors and David Makinson, On the logic of theory change: partial meet contraction and revision functions, *The Journal of Symbolic Logic*, 1985, 50: 510-530.

[2] Aumann, Robert, Sergiu Hart and Motty Perry, The absent-minded driver, *Games and Economic Behavior*, 1997, 20: 102-116.

[3] Battigalli, Pierpaolo, Strategic independence and perfect Bayesian equilibria, *Journal of Economic Theory*, 1996, 70: 201-234.

[4] Blume, L. and W. R. Zame, The algebraic geometry of perfect and sequential equilibrium, *Econometrica*, 1994, 62: 783-794.

[5] Bonanno, Giacomo, AGM belief revision in dynamic games, in: K. R. Apt (ed.), *Proceedings of the 13th conference on theoretical aspects of rationality and knowledge (TARK XIII)*, ACM, New York, 2011, 37-45.

[6] Bonanno, Giacomo, AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties, *International Journal of Game Theory*, 2013, 42: 567-592.

[7] Grove, Adam and Joseph Halpern, On the expected value of games with absentmindedness, *Games and Economic Behavior*, 1997, 20: 51-65.

[8] Hendon, Ebbe, Jørgen Jacobsen and Brigitte Sloth, The one-shot-deviation principle for sequential rationality, *Games and Economic Behavior*, 1996, 12: 274-282.

[9] Kohlberg, Elon and Philip Reny, Independence on relative probability spaces and consistent assessments in game trees, *Journal of Economic Theory*, 1997, 75: 280-313.

[10] Kreps, David and Robert Wilson, Sequential equilibrium, *Econometrica,* 1982, 50: 863-894.

[11] Osborne, Martin and Ariel Rubinstein, *A course in game theory,* MIT Press, 1994.

[12] Perea, Andrés, Mathijs Jansen and Hans Peters, Characterization of consistent assessments in extensive-form games, *Games and Economic Behavior*, 1997, 21: 238-252.

[13] Perea, Andrés, *Rationality in extensive form games*, Kluwer Academic Publishers, 2001.

[14] Perea, Andrés, A note on the one-deviation property in extensive form games, *Games and Economic Behavior*, 2002, 40: 322-338.

[15] Piccione, Michele and Ariel Rubinstein, On the interpretation of decision problems with imperfect recall, *Games and Economic Behavior*, 1997, 20: 3-24.

[16] Selten, Reinhard, Re-examination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory,* 1975, 4: 25-55.

[17] Streufert, Peter, Characterizing consistency with monomials, working paper, University of Western Ontario, 2007

[18] Streufert, Peter, Additive plausibility characterizes the supports of consistent assessments, Research Report # 2012-3, University of Western Ontario, 2012.