AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties.

Giacomo Bonanno*

Department of Economics,

University of California,

Davis, CA 95616-8578, USA

gfbonanno@ucdavis.edu

Abstract

We provide a general notion of perfect Bayesian equilibrium which can be applied to arbitrary extensive-form games and is intermediate between subgame-perfect equilibrium and sequential equilibrium. The essential ingredient of the proposed definition is the qualitative notion of AGM-consistency, which has an epistemic justification based on the theory of belief revision introduced by Alchourrón, Gärdenfors and Makinson. AGM-consistency is a generalization of the notion of consistency introduced by Kreps and Wilson as part of the definition of sequential equilibrium.

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

1 Introduction

Attempts to refine the notion of Nash equilibrium in extensive-form (or dynamic) games must deal with the issue of *belief revision*: how should a player revise her beliefs when informed that she has to make a choice at an information set to which she initially assigned zero probability? Kreps and Wilson [14] suggested the notion of assessment as a way of expressing the players' beliefs during an arbitrary play of the game. An *assessment* is a pair (σ, μ) where σ is a strategy profile and μ a "system of beliefs", defined as a collection of probability distributions, one for every information set, over the nodes in that information set. The system of beliefs μ specifies a player's beliefs about past moves, while the strategy profile σ provides the initial beliefs as well as beliefs about future moves conditional on every node. Given the separate roles played by σ and μ in specifying players' beliefs, it is necessary to impose some requirement of compatibility between the two. Kreps and

Wilson proposed the notion of "consistency", which we will call KW-consistency. An assessment (σ, μ) is *KW-consistent* 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, $\lim_{m \to \infty}(\sigma^m, \mu^m) = (\sigma, \mu)$.[1] A number of authors have tried to shed light on this topological notion by relating it to more intuitive concepts,[2] such as (convex) structural consistency ([15]), generally reasonable extended assessment ([10]), stochastic independence ([4, 13]).[3]. In applications, checking the KW-consistency requirement has proved to be rather complex and simpler notions of equilibrium have been sought. A drastic simplification is the notion of weak sequential equilibrium ([16], p. 170), which is defined as a sequentially rational assessment (σ, μ) where the beliefs expressed by μ are obtained using Bayes' rule at all the information sets that are reached by σ with positive probability (while no restrictions are imposed on the beliefs at information sets that are not reached by σ). However, this notion is too weak in the sense that (σ, μ) can be a weak sequential equilibrium without σ being a subgame-perfect equilibrium.[4] Thus attempts have been made to find an intermediate notion between subgame-perfect equilibrium and sequential equilibrium, incorporating a simpler requirement than KW-consistency. Such an intermediate notion was proposed by Fudenberg and Tirole [10] and called perfect Bayesian equilibrium. Unfortunately this new notion was defined only for a small subset of extensive-form games (namely the class of multi-stage games with observed actions) and extending it to arbitrary games proved to be problematic.[5]

In this paper we propose a new intermediate notion between subgame-perfect equilibrium and sequential equilibrium, which - in order to avoid introducing a new expression - we will also call perfect Bayesian equilibrium. The advantages of this equilibrium concept are that (1) it is a general notion that can be applied to arbitrary extensive-form games and (2) its main ingredient is a purely qualitative condition - we call it "AGM-consistency" - which is simple, easy to verify and a generalization of KW-consistency. An assessment (σ, μ) is *AGM-consistent* if there is a total pre-order - which we call a *plausibility order* - on the set of histories such that: (1) for every information set I the histories that are assigned positive probability by μ are precisely those that are most plausible in I and (2) at every information set I the choices that are assigned positive

---

[1] Kreps and Wilson then proceeded to define a sequential equilibrium as an assessment which is KW-consistent and sequentially rational. All the definitions will be spelled out in the next section. Sequential rationality requires that, at every information set, the choice(s) prescribed by the relevant component of σ be optimal, given the corresponding player's beliefs at that information set and the hypothesis that future behavior of the other players is as specified by σ.

[2] Kreps and Wilson themselves ([14], p. 872) remark that the "definition of consistency is not completely intuitive on its own" and proceed to offer a more intuitive interpretation in terms of the notion of *structural consistency* and a refinement of it, which they call *lexicographic consistency*. However, Keps and Ramey [15] later showed that the attempted justification of consistency was incorrect. In particular, they produce an example where every sequential equilibrium violates structural consistency. Kreps and Ramey derive the following conclusion from their analysis ([15], p. 1333): "We believe we strengthen the doubts expressed by Kreps and Wilson concerning consistency of beliefs; consistency itself does not encompass all the properties it was orginally thought to, and no single alternate notion of "consistent beliefs" seems ideal."

[3] Perea *et al* [19] offer an algebraic characterization of KW-consistency.

[4] A strategy profile σ is a subgame-perfect equilibrium if, for every subgame, the restriction of σ to the subgame yields a Nash equilibrium of the subgame. The notion of subgame-perfect equilibrium was introduced by Selten [21].

[5] Battigalli [4] introduced a notion of perfect Bayesian equilibrium that can be applied to general extensive-form games (although he only considered games without chance moves). However, his proposed definition is in terms of a more complex object, namely a "tree-extended assessment" (ν, σ, μ) where ν is a conditional probability system on the set of terminal nodes.

In a different direction, some authors have sought a coarsening, rather than a refinement, of Nash equilibrium in extensive-form games. For example, [22] provides a definition of correlated equilibrium which can be applied to arbitrary extensive-form games.

probability by σ are precisely those that "preserve plausibility", in the sense that if h is a history in I and a is a choice at h then ha is as plausible as h. An attractive feature of the notion of AGM-consistency is that it has an independent justification based on the so called AGM theory of belief revision introduced by Alchourrón, Gärdenfors and Makinson [1]. The epistemic foundations of AGM-consistency are developed in [9].

We define an assessment (σ, μ) to be a *perfect Bayesian equilibrium* if it is AGM-consistent, sequentially rational and the probabilities specified by μ are compatible with Bayes' rule "throughout the game". We show that if (σ, μ) is a perfect Bayesian equilibrium then σ is a subgame-perfect equilibrium and that the set of sequential equilibria is a subset of the set of perfect Bayesian equilibria. We also show that the proposed notion of perfect Bayesian equilibrium yields a strict refinement of subgame-perfect equilibrium and, in general, is weaker than sequential equilibrium.

2 Extensive forms, assessments and KW-consistency

We shall use the history-based definition of extensive-form game (see, for example, [18]). If A is a set, we denote by A^* the set of finite sequences in A. If $h = \langle a_1, ..., a_k \rangle \in A^*$ and $1 \leq j \leq k$, the sequence $h' = \langle a_1, ..., a_j \rangle$ is called a *prefix* of h. If $h = \langle a_1, ..., a_k \rangle \in A^*$ and $a \in A$, we denote the sequence $\langle a_1, ..., a_k, a \rangle \in A^*$ by ha.

A *finite extensive form* is a tuple $\langle A, H, N, \iota, \{\approx_i\}_{i \in N} \rangle$ whose elements are:

- A finite set of actions A.

- A finite set of histories $H \subseteq A^*$ which is closed under prefixes (that is, if $h \in H$ and $h' \in A^*$ is a prefix of h, then $h' \in H$). The null history $\langle \rangle$, denoted by $\emptyset$, is an element of H and is a prefix of every history. A history $h \in H$ such that, for every $a \in A$, $ha \notin H$, is called a *terminal history*. The set of terminal histories is denoted by Z. $D = H \backslash Z$ denotes the set of non-terminal or *decision* histories. For every history $h \in H$, we denote by $A(h)$ the set of actions available at h, that is, $A(h) = \{a \in A : ha \in H\}$. Thus $A(h) \neq \varnothing$ if and only if $h \in D$. We assume that $A = \bigcup_{h \in D} A(h)$ (that is, we restrict attention to actions that are available at some decision history).

- A finite set $N = \{1, ..., n\}$ of players. In some cases there is also an additional, fictitious, player called *chance*.

- A function $\iota : D \rightarrow N \cup \{chance\}$ that assigns a player to each decision history. Thus $\iota(h)$ is the player who moves at history h. A game is said to be *without chance moves* if $\iota(h) \in N$ for every $h \in D$. For every $i \in N \cup \{chance\}$, let $D_i = \iota^{-1}(i)$ be the histories assigned to player i. Thus $\{D_{chance}, D_1, ..., D_n\}$ is a partition of D. If history h is assigned to chance, then a probability distribution over $A(h)$ is given that assigns positive probability to every $a \in A(h)$.

- For every player $i \in N$, $\approx_i$ is an equivalence relation on D_i. The interpretation of $h \approx_i h'$ is that, when choosing an action at history $h \in D_i$, player i does not know whether she is moving at h or at h'. The equivalence class of $h \in D_i$ is denoted by $I_i(h)$ and is called an *information set of player i*; thus $I_i(h) = \{h' \in D_i : h \approx_i h'\}$. The following restriction applies: if $h' \in I_i(h)$ then $A(h') = A(h)$, that is, the set of actions available to a player is the same at any two histories that belong to the same information set of that player.

- The following property, known as *perfect recall*, is assumed: for every player $i \in N$, if $h_1, h_2 \in D_i$, $a \in A(h_1)$ and $h_1 a$ is a prefix of h_2 then for every $h' \in I_i(h_2)$ there exists an $h \in I_i(h_1)$ such that ha is a prefix of h'. Intuitively, perfect recall requires a player to remember what she knew in the past and what actions she took previously.[6]

Given an extensive form, one obtains an *extensive game* by adding, for every player $i \in N$, a *utility (or payoff) function* $U_i : Z \to \mathbb{R}$ (where $\mathbb{R}$ denotes the set of real numbers; recall that Z is the set of terminal histories).
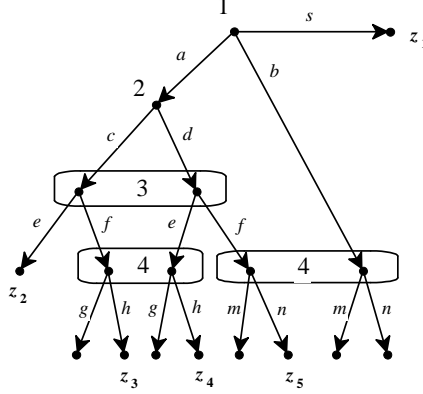
An extensive form without chance moves
Figure 1

Figure 1 shows an extensive form without chance moves where
$A = \{a, b, s, c, d, e, f, g, h, m, n\}$, $H = D \cup Z$ with (to simplify the notation we write a instead of $\langle \emptyset, a \rangle$, ac instead of $\langle \emptyset, a, c \rangle$, etc.) $D = \{\emptyset, a, b, ac, ad, acf, ade, adf\}$, $Z = \{s, ace, acfg, acfh, adeg,$ $adeh, adfm, adfn, bm, bn\}$, $A(\emptyset) = \{a, b, s\}$, $A(a) = \{c, d\}$, $A(ac) = A(ad) = \{e, f\}$, $A(acf) = A(ade) = \{g, h\}$, $A(adf) = A(b) = \{m, n\}$, $N = \{1, 2, 3, 4\}$, $\iota(\emptyset) = 1$, $\iota(a) = 2$, $\iota(ac) = \iota(ad) = 3$, $\iota(acf) = \iota(ade) = \iota(adf) = \iota(b) = 4$, $\approx_1 = \{(\emptyset, \emptyset)\}$, $\approx_2 = \{(a, a)\}$, $\approx_3 = \{(ac, ac), (ac, ad), (ad, ac), (ad, ad)\}$ and $\approx_4 = \{(acf, acf), (acf, ade), (ade, acf), (ade, ade),$ $(adf, adf), (adf, b), (b, adf), (b, b)\}$. The information sets containing more than one history (for example, $I_4(b) = \{adf, b\}$) are shown as rounded rectangles. The root of the tree represents the null history $\emptyset$.

Notation 1 *If h and h' are decision histories not assigned to chance, we write $h' \in I(h)$ as a short-hand for $h' \in I_{\iota(h)}(h)$. Thus $h' \in I(h)$ means that h and h' belong to the same information set (of the player who moves at h). If h is a history assigned to chance, we use the convention that $I(h) = \{h\}$.*

Remark 2 *In order to simplify the notation in the proofs, we shall assume that no action is available at more than one information set: $\forall h, h' \in H, \forall a \in A,$ if $a \in A(h) \cap A(h')$ then $h' \in I(h)$.*

---

[6] For an investigation of the conceptual content of the property of perfect recall see [6].

4

Given an extensive form, a *pure strategy* of player $i \in N$ is a function that associates with every information set of player i an action at that information set, that is, a function $s_i : D_i \to A$ such that (1) $s_i(h) \in A(h)$ and (2) if $h' \in I_i(h)$ then $s_i(h') = s_i(h)$. For example, one of the pure strategies of Player 4 in the extensive form illustrated in Figure 1 is $s_4(acf) = s_4(ade) = g$ and $s_4(adf) = s_4(b) = m$. A *behavior strategy* of player i is a collection of probability distributions, one for each information set, over the actions available at that information set; that is, a function $\sigma_i : D_i \to \Delta(A)$ (where $\Delta(A)$ denotes the set of probability distributions over A) such that (1) $\sigma_i(h)$ is a probability distribution over $A(h)$ and (2) if $h' \in I_i(h)$ then $\sigma_i(h') = \sigma_i(h)$. If the game does not have chance moves, we define a behavior strategy *profile* as an n-tuple $\sigma = (\sigma_1, ..., \sigma_n)$ where, for every $i \in N$, σ_i is a behavior strategy of player i. If the game has chance moves then we use the convention that a behavior strategy profile is an $(n+1)$-tuple $\sigma = (\sigma_1, ..., \sigma_n, \sigma_{chance})$ where, if h is a history assigned to chance and $a \in A(h)$ then $\sigma_{chance}(h)(a)$ is the probability associated with a. Given our assumption that no action is available at more than one information set, without risking ambiguity we shall denote by $\sigma(a)$ the probability assigned to action a by the relevant component of the strategy profile σ.[7] Note that a pure strategy is a special case of a behavior strategy where each probability distribution is degenerate. A behavior strategy is *completely mixed* at history $h \in D$ if, for every $a \in A(h)$, $\sigma(a) > 0$.

For example, in the extensive form of Figure 1 a possible behavior strategy for Player 1 is $\begin{pmatrix} a & b & s \\ 0 & 0 & 1 \end{pmatrix}$, which we will more simply denote by s (and which coincides with a pure strategy of Player 1) and a possible behavior strategy of Player 2 is $\begin{pmatrix} c & d \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}$, which we will more simply denote by $\left(\frac{1}{3}c, \frac{2}{3}d\right)$ (and which is a completely mixed strategy).

A *system of beliefs*, is a collection of probability distributions, one for every information set, over the elements of that information set, that is, a function $\mu : D \to \Delta(H)$ such that if $h \in D_i$ ($i \in N$) then $\mu(h)$ is a probability distribution over $I_i(h)$ and if $h' \in I_i(h)$ then $\mu(h) = \mu(h')$. If the game has chance moves, then we use the convention that $\mu(h) = 1$ for every history h assigned to chance. Without risking ambiguity we shall denote by $\mu(h)$ the probability assigned to history h by the system of beliefs μ.[8]

An *assessment* is a pair (σ, μ) where σ is a behavior strategy profile and μ is a system of beliefs. An assessment represents the beliefs of the players: the strategy profile σ yields the initial beliefs as well as conditional beliefs about the future, while the system of beliefs μ gives conditional beliefs about the past. For example, consider the following assessment for the extensive form of Figure 1: $\sigma = \left(s, \left(\frac{1}{3}c, \frac{2}{3}d\right), f, (h, n)\right)$ and $\mu = \left(\left(\frac{1}{2}ac, \frac{1}{2}ad\right), acf, \left(\frac{1}{4}adf, \frac{3}{4}b\right)\right)$ and consider Player 3. By σ, the initial beliefs of Player 3 (that is, before the game begins) is that Player 1 will play s and thus the outcome of the game will be z_1; in particular, Player 3 believes that she will not be asked to move. If she were asked to move, then, by μ, she would assign equal probability to the event that Player 2 played c and the event that Player 2 played d (while knowing that Player 1 played a) and, by σ, she would believe that her choice of e would lead to either outcome z_2 or outcome z_4 (with equal probability) and that her choice of f would lead to either outcome z_3 or outcome z_5 (with equal probability).

---

[7] If $h \in D_i$ and σ_i is the *ith* component of σ, then $\sigma_i(h)$ is a probability distribution over $A(h)$ and if $a \in A(h)$ then $\sigma_i(h)(a)$ is the probability assigned to action a by $\sigma_i(h)$. Thus we denote $\sigma_i(h)(a)$ more simply by $\sigma(a)$.

[8] A more precise notation would be $\mu(h)(h)$: if $h \in D_i$ then $\mu(h)$ is a probability distribution over $I_i(h)$ and, for every $h' \in I(h)$, $\mu(h) = \mu(h')$ so that $\mu(h)(h) = \mu(h')(h)$. With slight abuse of notation we denote this common probability by $\mu(h)$.

Given that σ and μ play separate roles in the representation of the players' initial beliefs and disposition to change those beliefs, it is necessary to impose some requirement of compatibility between the two. Kreps and Wilson [14] proposed the following notion of compatibility, which they called consistency; we will call it KW-consistency. An assessment (σ, μ) is KW-consistent 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,[9] $\lim_{m \to \infty}(\sigma^m, \mu^m) = (\sigma, \mu)$.

In this paper we provide an independently motivated (see [9]) qualitative notion of compatibility, which we call *AGM-consistency*, and use it to define a notion of equilibrium which is intermediate between subgame-perfect equilibrium and sequential equilibrium.

3   AGM-consistency

Given a set H, a total pre-order on H is a binary relation $\precsim \ \subseteq H \times H$ 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 $h \precsim h''$).

Definition 3 *Given an extensive form, a* plausibility order *is a total pre-order on the set of histories H that satisfies the following properties:* $\forall h \in D$,

    PL1. $h \precsim ha, \quad \forall a \in A(h)$.

    PL2.   *(i)* $\exists a \in A(h)$ *such that* $ha \precsim h$,
            (ii) $\forall a \in A(h)$, *if* $ha \precsim h$ *then* $h'a \precsim h'$, $\forall h' \in I(h)$.

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

If $h \precsim h'$ we say that history h is *at least as plausible* as history h'. 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 some action a such that adding a to h yields a history which is at least as plausible as h and, 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" (see Remark 5 below).

Remark 4 *In [9] the notion of plausibility order is derived from the primitive concept of a player's (qualitative) epistemic state, which consists of the player's* initial beliefs *(or* prior beliefs, *that is, her beliefs before the game is played) and* disposition to revise those beliefs *when informed that it is her turn to move. In an extensive-form game, a player might find herself having to move at an information set that - according to her prior beliefs - should not have been reached; in such a case the player will have to revise her initial beliefs by formulating a hypothesis about the past moves of the other players and a prediction about future moves by herself and the other players. In [9] the epistemic state of player i is expressed as a triple $\langle H, \mathcal{E}_i, f_i \rangle$, where H is the set of histories, $\mathcal{E}_i$ is the set of possible items of information for player i (consisting of the information sets of player i augmented with all possible continuation histories) and $f_i : \mathcal{E}_i \to H$ (which satisfies the property that, for every $E \in \mathcal{E}_i$, $\varnothing \neq f_i(E) \subseteq E$) is a function that gives the player's conditional beliefs.*

---

[9]That is, for every $h \in D$, $\mu^m(h) = \dfrac{\prod_{a \in h} \sigma^m(a)}{\sum_{x \in I(h)} \prod_{a \in x} \sigma^m(a)}$, where $a \in h$ means that action a occurs 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 $h \in D$.

The interpretation of $f_i(E)$ is that, if informed that event E has occurred, player i considers as doxastically (that is, according to her beliefs) possible all and only the histories in $f_i(E)$. Four natural properties of the individual epistemic states are given and it is shown that those properties, together with the hypothesis of a common prior, are necessary and sufficient for the existence of a plausibility order $\precsim\ \subseteq H \times H$ (Definition 3 above) that rationalizes the epistemic state of every player, in the sense that, for every $i \in N$ and for every $E \in \mathcal{E}_i$, $f_i(E) = \{h \in E : h \precsim h', \forall h' \in E\}$. That is, when informed that event E has occurred, player i considers as doxastically possible precisely the histories that are most plausible within the set E. For more details the reader is referred to [9].

We write $h \sim h'$ (with the interpretation that h is *as plausible* as h') as a short-hand for "$h \precsim h'$ and $h' \precsim h$" and we write $h \prec h'$ (with the interpretation that h is *more plausible* than h') as a short-hand for "$h \precsim h'$ and $h' \not\precsim h'$".

Remark 5 *It follows from Property PL1 of Definition 3 that, for every $h, h' \in H$, if h' is a prefix of h then $h' \precsim h$.[10] Furthermore, by Properties PL1 and PL2, for every decision history h, there is at least one action a at h such that $h \sim ha$, that is, ha is as plausible as h and, furthermore, if h' belongs to the same information set as h, then $h' \sim h'a$. We call such actions* plausibility preserving.

Definition 6 *Fix an extensive-form. An assessment (σ, μ) is* AGM-consistent *if there exists a plausibility order $\precsim$ on H such that:*

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

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

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

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

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

Remark 7 *Given the justification of plausibility orders in terms of belief revision (see Remark 4), the definition of AGM-consistent assessment captures a subjective interpretation of assessments. This is along the lines of a recent literature (see, for example, [3] and the references given there in Footnote 7) where mixed strategies are treated not as conscious randomization but as conjectures. In particular, we subscribe to the view expressed in [5] (p.2) that "strategies as plans cannot be anything but beliefs of players about their own behavior". Definition 6 says that - when informed that an information set of hers has been reached - a player forms her beliefs about what has happened so far by focusing on the most plausible histories within the information set (Property P2) and forms beliefs about her own and the other players' future choices by focusing on the continuation histories obtained by following the plausibility-preserving actions (Property P1). For more details the reader is referred to [9].*

---

[10]If h' is a prefix of h then $h = h'a_1...a_m$ for some (possibly none) $a_1, ..., a_m \in A$, so that, by Property PL1, $h' \precsim h'a_1 \precsim h'a_1a_2 \precsim ... \precsim h'a_1...a_m = h$ and thus, by transitivity of $\precsim$, $h' \precsim h$.

The notion of AGM-consistency imposes natural restrictions on assessments. Consider, for example, the extensive form of Figure 2 and any assessment (σ, μ) where $\sigma = (c, d, f)$ (highlighted by double edges) and μ assigns positive probability to history be. Any such assessment is *not* AGM-consistent. In fact, if there were a plausibility order $\precsim$ that satisfied Definition 6, then, by *P1*, $b \sim bd$ (since $\sigma(d) = 1 > 0$) and $b \prec be$ (since $\sigma(e) = 0$)[11] and, by *P2*, $be \precsim bd$ (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.
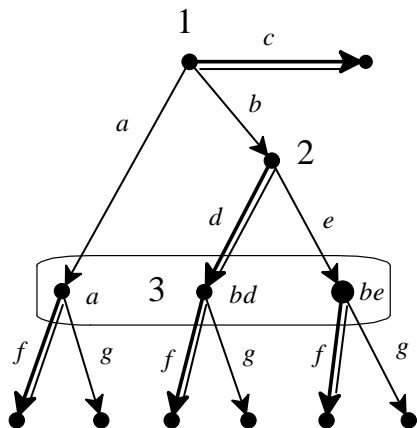
Any assessment (σ, μ) where $\sigma = (c, d, f)$ and μ assigns
positive probability to history be is not AGM-consistent.

Figure 2

On the other hand, consider the partial extensive form of Figure 3 (in order to simplify the figure, the choices of Player 3 have been omitted) and an arbitrary assessment of the form (σ, μ) where $\sigma = (a, (g, r))$ (highlighted by double edges) and μ assigns positive probability to the black nodes and zero probability to the gray nodes (thus $\mu(c) = \mu(e) = \mu(ch) = 0$, while every other history is assigned positive probability). Any such assessment is AGM-consistent. In fact it is rationalized by the following plausibility order (h is on the same row as h' if and only if $h \sim h'$ and h is above h' if and only if $h \prec h'$):

$$
\begin{pmatrix}
\emptyset, a \\
b, bg, d, dr \\
bh, ds \\
c, cg, e, er \\
es \\
ch
\end{pmatrix}
$$

Note that there may be several plausibility orders that rationalize a given assessment (for instance, an alternative plausibility order to the one given above is obtained by switching the rows (bh, ds) and (c, cg, e, er)).

---

[11] By definition of plausibility order, $b \precsim be$ and, by *P1*, it is not the case that $b \sim be$ because e is not assigned positive probability by σ. Thus $b \prec be$.
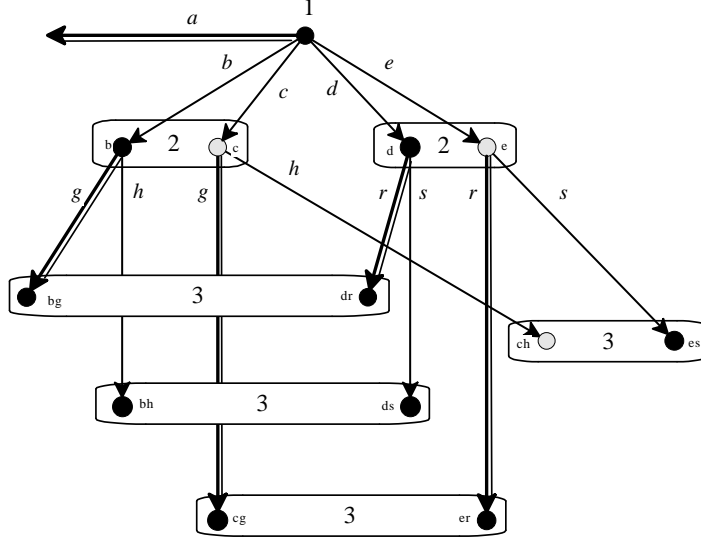
Any assessment (σ, μ) where $\sigma = (a, (g, r))$ and μ assigns zero
probability to, and only to, histories c, e and ch is AGM-consistent.

Figure 3

For every extensive form, the set of plausibility orders is non-empty and, for every plausibility order, there is at least one assessment which is rationalized by that plausibility order.[12]

We now show that the notion of AGM-consistency generalizes the notion of KW-consistency, in the sense that *the set of KW-consistent assessments is a subset of the set of AGM-consistent assessments.*

Given a plausibility order $\precsim$ on the set of histories H, a function $F : H \to \mathbb{N}$ (where $\mathbb{N}$ denotes the set of non-negative integers) is an *integer-valued representation* of $\precsim$ if $F(\emptyset) = 0$ (recall that $\emptyset$ denotes the null history) and, $\forall h, h' \in H$, $F(h) \le F(h')$ if and only if $h \precsim h'$. Since H is finite, the set of integer-valued representations of $\precsim$ is non-empty.[13]

Definition 8 *A plausibility order $\precsim$ on the set of histories H is* choice-measurable *if it has at least one integer-valued representation F that satisfies the following property: $\forall h \in D$, $\forall h' \in I(h)$, $\forall a \in A(h)$,*

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

Note that not every integer-valued representation of a choice measurable plausibility order need satisfy Property CM. For example, consider the plausibility order and the two integer-valued

---

[12] If $h \in D_i$ let $\sigma_i(h)$ be an arbitrary probability distribution whose support coincides with the set of plausibility preserving actions at h and if I is an information set then let μ assign positive probability to all and only the most plausible histories in I.

[13] A natural integer-valued representation is the following. Define $H_0 = \{h \in H : h \precsim x, \ \forall x \in H\}$, $H_1 = \{h \in H \setminus H_0 : h \precsim x, \ \forall x \in H \setminus H_0\}$ and, in general, for every integer $k \ge 1$, $H_k = \{h \in H \setminus H_0 \cup ... \cup H_{k-1} : h \precsim x, \ \forall x \in H \setminus H_0 \cup ... \cup H_{k-1}\}$. Since H is finite, there is an $m \in \mathbb{N}$ such that $\{H_0, ..., H_m\}$ is a partition of H and, for every $j, k \in \mathbb{N}$, with $j < k \le m$, and for every $h, h' \in H$, if $h \in H_j$ and $h' \in H_k$ then $h \prec h'$. Define $F : H \to \mathbb{N}$ as follows: $F(h) = k$ if and only if $h \in H_k$. The function F so defined is an integer-valued representation of $\precsim$.

representations F_1 and F_2 shown in Figure 4a. F_1 is the representation described in Footnote 13 and does not satisfy Property *CM*, since $c \in I(b)$, $f \in A(b) = A(c)$ and $F_1(bf) - F_1(b) = 2 - 1 = 1$ while $F_1(cf) - F_1(c) = 5 - 3 = 2$. On the other hand, F_2 does satisfy *CM*.

$\precsim$	F_1	F_2
$\emptyset$, a	0	0
b, be	1	1
bf	2	3
c, ce	3	4
d	4	5
cf	5	6

(a) A plausibility order and two integer-valued representations of it.

(b) The assessment $\sigma = (a, e)$, $\mu(b) = 1$, $\mu(c) = 0$ is rationalized by the choice-measurable plausibility order shown on the left.
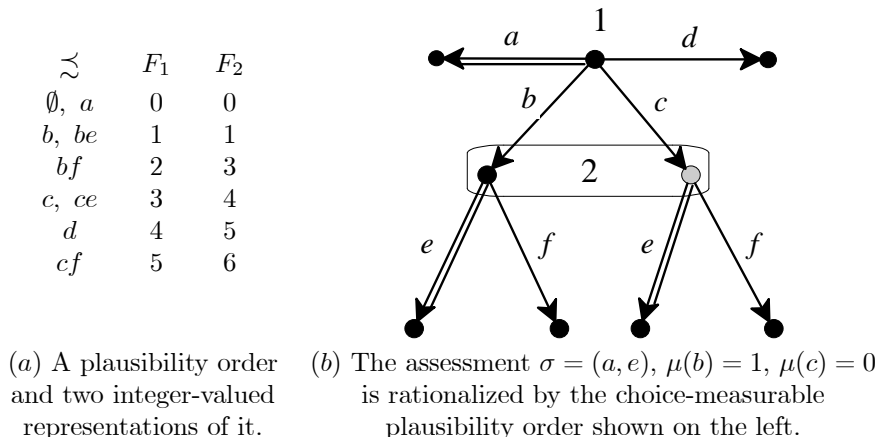
Figure 4

The following lemma (which follows easily from Lemma A.1 in Kreps and Wilson [14], p. 887) shows that AMG-consistency is a generalization of KW-consistency. The proof is given in the Appendix.

Lemma 9 *Fix an extensive form. If (σ, μ) is a KW-consistent assessment then there is a choice-measurable plausibility order that rationalizes it (and thus (σ, μ) is AGM consistent).*

Lemma 9 makes it easy to prove that an assessment is not KW-consistent. For example, consider again the extensive form of Figure 3 and an arbitrary assessment (σ, μ) where $\sigma = (a, (g, r))$ (highlighted by double edges) and μ assigns positive probability to the black nodes and zero probability to the gray nodes (thus $\mu(c) = \mu(e) = \mu(ch) = 0$, while every other history is assigned positive probability). We saw above that any such assessment is AGM-consistent. On the other hand, no such assessment is KW-consistent, because there is no choice-measurable plausibility order that rationalizes it. To see this, let $\precsim$ be a plausibility order that rationalizes the assessment under consideration and let F be an arbitrary integer-valued representation of $\precsim$. By *P2* of Definition 6, $es \prec ch$ and thus

$$F(es) < F(ch). \tag{1}$$

By *P1* of Definition 6, $c \sim cg$ and $e \sim er$ and by *P2* $cg \sim er$. Thus, by transitivity of $\precsim$, $c \sim e$ so that $F(c) = F(e)$. Hence, by (1),

$$F(es) - F(e) < F(ch) - F(c). \tag{2}$$

Similarly, by *P1*, $b \sim bg$ and $d \sim dr$ and, by *P2*, $bg \sim dr$. Thus, by transitivity of $\precsim$, $b \sim d$ so that

$$F(b) = F(d). \tag{3}$$

10

By $P2$, $bh \sim ds$ and thus $F(bh) = F(ds)$. Hence, by (3),

$$F(bh) - F(b) = F(ds) - F(d). \tag{4}$$

Now, if $F(ch) - F(c) = F(bh) - F(b)$ (as required by Property CM of Definition 8) then, by (2) and (4), $F(es) - F(e) < F(ds) - F(d)$, violating Property CM.

4 Perfect Bayesian equilibrium

We now define a notion of perfect Bayesian equilibrium which is very general, in the sense that it can be applied to every finite extensive-form game. The essential ingredient is the qualitative notion of AGM-consistency to which we add a Bayesian updating requirement for probabilities and sequential rationality.

Given a total pre-order $\precsim$ on H and a subset $I \subseteq H$ we denote by $Min_{\precsim} I$ the set of most plausible histories in I, that is,

$$Min_{\precsim} I = \{h \in I : h \precsim h', \forall h' \in I\}.$$

Definition 10 *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 $\precsim$-equivalence class E that contains some decision history $h \in D_i$ ($i \in N$) with $\mu(h) > 0$, there exists a probability measure $\nu_E : H \to [0,1]$ such that:*

B1. $Supp(\nu_E) = \{h \in E : \mu(h) > 0\}$.

B2. If $h, h' \in Supp(\nu_E)$ 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 $h' \in Supp(\nu_E)$, then, $\forall h \in I(h')$

$$\mu(h) = \nu_E\left(h \mid I(h')\right) = \frac{\nu_E(h)}{\nu_E(I(h'))}.$$

Property $B1$ requires that $\nu_E(h) > 0$ if and only if $h \in E$ and $\mu(h) > 0$ (thus $\nu_E(h) = 0$ if and only if either $h \in H \backslash E$ or $\mu(h) = 0$). Property $B2$ requires ν_E to be consistent with the strategy profile σ in the sense that if $h, h' \in E$, $\nu_E(h) > 0$, $\nu_E(h') > 0$ and $h' = ha_1...a_m$ then the probability of h' (according to ν_E) is equal to the probability of h multiplied by the probabilities (according to σ) of the actions that lead from h to h'.[14] Property $B3$ requires the system of beliefs μ to satisfy Bayes' rule in the sense that if history h belongs to information set I then $\mu(h)$ (the probability assigned to h by μ) is the probability of h conditional on I using the probability measure ν_E, where

---

[14]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 6, $\sigma(a_j) > 0$.

E is the equivalence class of the most plausible elements of I. In fact, Property $B3$ is equivalent to the following:[15]

$$\text{For every information set } I \text{ such that } Min_{\precsim}I \subseteq E \qquad (B3')$$
$$\text{and, for every } h \in I, \ \mu(h) = \nu_E\left(h \mid I\right) = \frac{\nu_E(h)}{\nu_E(I)}.$$

Definition 11 *An assessment (σ, μ) is* Bayesian AGM-consistent *if it is rationalized by a plausibility order $\precsim$ on the set of histories H and it is Bayesian relative to $\precsim$.*

Consider, for example, the extensive form of Figure 3 and the following assessment:

$$\sigma = (a, (g, r)) \text{ and } \mu = \left(b, (\tfrac{1}{2}bg, \tfrac{1}{2}dr), (\tfrac{1}{2}bh, \tfrac{1}{2}ds), (\tfrac{1}{2}cg, \tfrac{1}{2}er), d, es\right).$$

Then, as we saw above, (σ, μ) is rationalized by

$$\precsim \ = \ \begin{pmatrix} \emptyset, a \\ b, bg, d, dr \\ bh, ds \\ c, cg, e, er \\ es \\ ch \end{pmatrix}.$$

Furthermore, (σ, μ) is Bayesian relative to $\precsim$: for every (relevant) $\precsim$-equivalence class E let ν_E to be the uniform measure on $\{h \in E : \mu(h) > 0\}$. Take, for instance, $E = \{b, bg, d, dr\}$. Then ν_E assigns probability $\frac{1}{4}$ to every element of E and zero to every history in $H \backslash E$. Property $B2$ of Definition 10 is satisfied, because $b, bg \in E$, b is a prefix of bg and $\nu_E(bg) = \frac{1}{4} = \nu_E(b) \times \sigma(g) = \frac{1}{4} \times 1 = \frac{1}{4}$ (and similarly for d and dr). Property $B3$ is also satisfied. For instance, $\mu(bg) = \frac{1}{2} = \frac{\nu_E(bg)}{\nu_E(\{bg,dr\})} = \frac{\frac{1}{4}}{\frac{1}{4}+\frac{1}{4}}$ and $\mu(b) = 1 = \frac{\nu_E(b)}{\nu_E(\{b,c\})} = \frac{\frac{1}{4}}{\frac{1}{4}+0}$. For the equivalence class $F = \{c, cg, e, er\}$ we take $\nu_F = \begin{pmatrix} c & cg & e & er \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix}$ since $\mu(c) = \mu(e) = 0$. Property $B2$ is not relevant, since, although c is a prefix of cg, $\nu_F(c) = 0$. Property $B3$ is only relevant for histories cg and er (and indeed $\mu(cg) = \frac{1}{2} = \frac{\nu_F(cg)}{\nu_F(\{cg,er\})} = \frac{\frac{1}{2}}{\frac{1}{2}+\frac{1}{2}}$), while ν_F is not the relevant measure for $I(c) = \{b, c\}$, since $Min_{\precsim}\{b, c\} = \{b\} \nsubseteq F$ (the relevant measure for $I(c)$ is the measure ν_E described above).

As another example, consider the extensive form of Figure 5 and the following assessment:

---

[15] Proof that $B3 \Rightarrow B3'$. Let I be such that $Min_{\precsim}I \subseteq E$. Fix an arbitrary $h \in I$. If $h \in Min_{\precsim}I$ then, by Property $P2$ of Definition 6, $\mu(h) > 0$ and thus, by $B1$, $h \in Supp(\nu_E)$ so that, by $B3$, $\mu(h) = \nu_E\left(h \mid I\right) = \frac{\nu_E(h)}{\nu_E(I)}$. If $h \notin Min_{\precsim}I$ then, by Property $P2$ of Definition 6, $\mu(h) = 0$ and thus, by $B1$, $\nu_E(h) = 0$, so that $\mu(h) = \nu_E\left(h \mid I\right) = \frac{\nu_E(h)}{\nu_E(I)}$ (note that $\nu_E(I) > 0$ since $\varnothing \neq Min_{\precsim}I \subseteq \{h \in E : \mu(h) > 0\}$).
Proof that $B3' \Rightarrow B3$. Let $h' \in Supp(\nu_E)$. By $B1$, $\mu(h') > 0$ so that, by Property $P2$ of Definition 6, $h' \in Min_{\precsim}I(h')$ and thus $Min_{\precsim}I(h') \subseteq E$ (since $h' \in E$ by hypothesis). Fix an arbitrary $h \in I(h')$. By $B3'$, $\mu(h) = \nu_E\left(h \mid I(h')\right) = \frac{\nu_E(h)}{\nu_E(I(h'))}$.
Note that the $\precsim$-equivalence class that contains $Min_{\precsim}I$ may be a proper superset of $Min_{\precsim}I$ and is not necessarily a subset of I. For example, if $H = \{h_1, ..., h_4\}$ and $\precsim$ is given by $h_1 \prec h_2 \sim h_3 \prec h_4 \sim h_5$ and $I = \{h_2, h_4\}$ then $Min_{\precsim}I = \{h_2\}$ and the equivalence class that contains h_2 is $\{h_2, h_3\}$.

$$\sigma = \left(c, d, (\tfrac{1}{3}f, \tfrac{2}{3}g)\right), \ \mu = \left((\tfrac{1}{4}a, \tfrac{1}{4}bf, \tfrac{2}{4}bg), (\tfrac{1}{4}ad, \tfrac{3}{4}b)\right).$$

Then (σ, μ) is rationalized by

$$\precsim \ = \left(\begin{array}{c} \emptyset, c \\ a, b, ad, bf, bg, adf, adg, bfd, bgd \\ ae, aef, aeg, bfe, bge \end{array} \right)$$

and it is Bayesian relative to $\precsim$. First of all, note that, letting $I_2 = \{a, bf, bg\}$ be the information set of Player 2 and $I_3 = \{ad, ae, b\}$ the information set of Player 3, we have that $Min_{\prec}I_2 = I_2$ and $Min_{\prec}I_3 = \{ad, b\}$ and the equivalence class that contains these two sets is $E = \{a, b, ad, bf, bg, adf, adg, bfd, bgd\}$. Then all we need to specify is a probability measure over E. Let

$$\nu_E = \left(\begin{array}{ccccccccc} a & b & ad & bf & bg & adf & adg & bfd & bgd \\ \frac{1}{8} & \frac{3}{8} & \frac{1}{8} & \frac{1}{8} & \frac{2}{8} & 0 & 0 & 0 & 0 \end{array} \right).$$

Then all the properties of Definition 10 are satisfied. For instance, concerning Property $B2$, b is a prefix of bf and $\nu_E(bf) = \frac{1}{8} = \nu_E(b) \times \sigma(f) = \frac{3}{8} \times \frac{1}{3}$, and, concerning Property $B3$, $\mu(ad) = \frac{1}{4} = \frac{\nu_E(ad)}{\nu_E(I_3)} = \frac{\frac{1}{8}}{\frac{1}{8} + 0 + \frac{3}{8}}$.

The assessment $\sigma = \left(c, d, (\tfrac{1}{3}f, \tfrac{2}{3}g)\right), \ \mu = \left((\tfrac{1}{4}a, \tfrac{1}{4}bf, \tfrac{2}{4}bg), (\tfrac{1}{4}ad, \tfrac{3}{4}b)\right)$
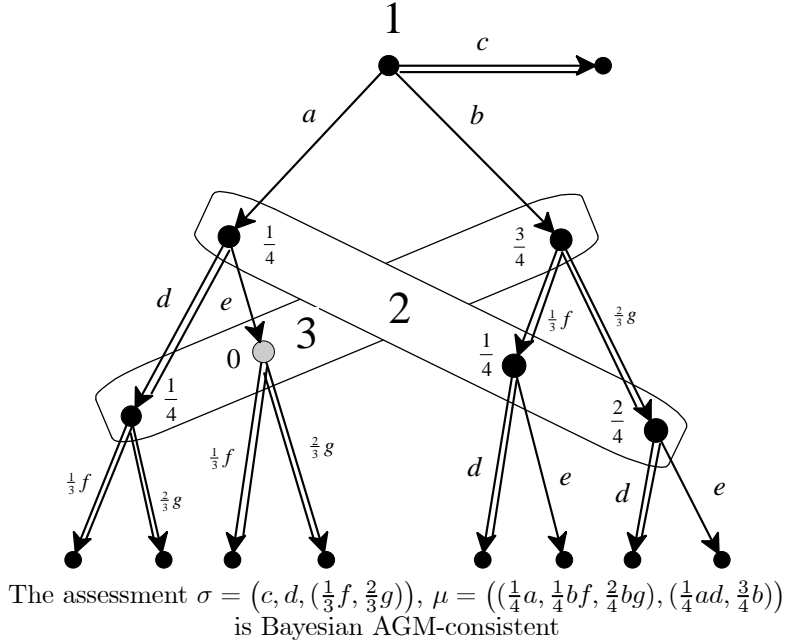is Bayesian AGM-consistent
Figure 5

The last ingredient of the definition of perfect Bayesian equilibrium is the standard requirement of sequential rationality. An assessment (σ, μ) is sequentially rational if, for every player $i \in N$ and every information set I_i of player i, player i's expected payoff, given her beliefs at I_i (as specified

by μ) and the strategy profile σ, cannot be increased by unilaterally changing her choice at I_i and possibly at information sets of hers that follow I_i.[16] In order to define sequential rationality precisely we need to introduce more notation. Recall that Z denotes the set of terminal histories and, for every player $i \in N$, $U_i : Z \to \mathbb{R}$ denotes player i's von Neumann-Morgenstern utility function. Given a decision history h, let $Z(h)$ denote the terminal histories which have h as a prefix. Let $\mathbb{P}_{h,\sigma}$ denote 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_i be an information set of player i and let $u_i(I_i|\sigma,\mu) = \sum_{h \in I_i} \mu(h) \sum_{z \in Z(h)} \mathbb{P}_{h,\sigma}(z)U_i(z)$ be player $i's$ expected payoff at I_i if σ is played, given her beliefs at I_i (as specified by μ). We say that player i's strategy σ_i is *sequentially rational at* I_i if $u_i(I_i|(\sigma_i,\sigma_{-i}),\mu) \geq u_i(I_i|(\tau_i,\sigma_{-i}),\mu)$ for every strategy τ_i of player i (where σ_{-i} denotes the strategy profile of the players other than i, that is, $\sigma_{-i} = (\sigma_1, ..., \sigma_{i-1}, \sigma_{i+1}, ..., \sigma_n)$). An assessment (σ,μ) is *sequentially rational* if, for every player $i \in N$ and for every information set I_i of player i, σ_i is sequentially rational at I_i.

Definition 12 *An assessment (σ,μ) is a* perfect Bayesian equilibrium *if it is Bayesian AGM-consistent and sequentially rational.*

The following propositions are proved in the Appendix.

Proposition 13 *If (σ,μ) is a perfect Bayesian equilibrium then σ is a subgame-perfect equilibrium.*

Note that not every subgame-perfect equilibrium is part of an assessment which is a perfect Bayesian equilibrium. That is, the notion of perfect Bayesian equilibrium is a *strict refinement of subgame-perfect equilibrium*. This can be shown with the aid of the extensive form of Figure 2. Turn it into a game by adding the following payoffs:

	af	ag	bdf	bdg	bef	beg	c
U_1	0	2	0	0	0	0	1
U_2	0	2	1	0	0	0	0
U_3	0	2	0	1	1	0	0

Let $\sigma = (c,d,f)$. Then σ is a Nash equilibrium and thus - since there are no proper subgames - also a subgame-perfect equilibrium. Add to σ a system of beliefs μ that makes (σ,μ) a sequentially rational assessment. Then it must be that $\mu(be) > 0$ (since, for Player 3, g is strictly better than f at both history a and history bd). As showed above, such an assessment is not AGM-consistent.

Proposition 14 *If (σ,μ) is a sequential equilibrium then it is a perfect Bayesian equilibrium.*

Note that not every perfect Bayesian equilibrium is a sequential equilibrium (this can be shown with the aid of a completion of the extensive form of Figure 3). Thus *sequential equilibrium is a strict refinement of perfect Bayesian equilibrium.*

Kreps and Wilson [14] proved that every finite extensive-form game has at least one sequential equilibrium. From this existence result and Proposition 14 one thus obtains the following.

---

[16] 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 12 below) is modified by using the strictly local definition of sequential rationality, then an extra condition needs to be added to Definition 10, namely the "pre-consistency" condition on μ identified in [12] and [20] as being necessary and sufficient for the equivalence of the two notions. For simplicity we have chosen the weakly local definition.

Corollary 15 *Every finite extensive-form game has at least one perfect Bayesian equilibrium.*

5 Conclusion

The central elements of the proposed definition of perfect Bayesian equilibrium are the qualitative notions of plausibility order and AGM-consistency of assessments. As shown in [9], these notions can be derived from the primitive concept of a player's epistemic state, which encodes the player's initial (or prior) 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 (see Remarks 4 and 7) guarantees that the belief revision policy of every player satisfies the axioms for belief revision introduced by Alchourrón, Gärdenfors and Makinson [1].[17]

The notion of plausibility order and the associated notion of AGM-consistent assessment are qualitative notions. It is well-known that there is a correspondence between the set of total pre-orders on a set S and the set of conditional probability systems (CPS) on $2^S \times 2^S \backslash \varnothing$ (see, for example, [2]).[18] Thus one could alternatively develop the analysis of this paper in terms of a CPS on $2^H \times 2^H \backslash \varnothing$ (where H is the set of histories in the given extensive-form game). Indeed several papers in the literature study solution concepts for extensive-form games in terms of conditional probability systems.[19] However, conditional probability systems are much more complex objects than plausibility orders. In a CPS the set of conditioning events is very "large" as compared to the typically very few items of information that a player might receive in a game. Thus a CPS introduces much more structure than is needed. We feel that the approach proposed in this paper offers a considerable gain in simplicity by avoiding unnecessary structure.[20]

It is important to note that *the notion of AGM-consistency allows for correlation in a player's beliefs about the choices of her opponents* (and possibly her own choices), while KW-consistency has several independence properties built into it. Some of these independence properties can be expressed in purely qualitative terms (that is, as properties of the underlying plausibility order)

---

[17] The AGM theory of belief revision is not directly applicable to extensive-form games for two reasons: (1) the AGM postulates are formulated within a syntactic framework, while extensive-form games are set-theoretic constructs, and (2) the AGM belief revision functions have as domain the set of all conceivable items of information (the set of all formulas) while in an extensive game there are typically very few items of information that a player might receive (represented by her information sets). The link between the syntactic and the set-theoretic frameworks is provided by a *valuation* which gives, for every atomic formula p, the set of histories at which p is true. A valuation gives rise to a *partial* syntactic belief revision function. It is shown in [7, 8] that such a function can be extended to a full-domain function that satisfies the AGM postulates if and only if the epistemic state of the player can be rationalized by a total pre-order.

[18] Given a CPS $P : 2^S \times 2^S \backslash \varnothing \to [0,1]$, one can extract from it a total pre-order $\precsim$ on S by letting $s \precsim s'$ if and only if $P(\{s\}|\{s,s'\}) > 0$. Conversely, given a total pre-order $\precsim$ on S on can construct at least one (typically several) CPS such that, for every $s, s' \in S$, $P(\{s\}|\{s,s'\}) > 0$ if and only if $s \precsim s'$.

[19] For example, [4] uses either CPSs on the set of strategy profiles or (in Section 5) CPSs on the set of terminal nodes and [5] uses more complex CPSs defined over the product space of terminal histories and profiles of players' types.

[20] It could be objected that the collection of probability measures postulated in Definition 10 is close to a lexicographic probability system, which - as is well known (see, for example, [11]) - is equivalent to a CPS. However, Definition 10 requires specifying a measure only for the $\precsim$-equivalence classes that contain some history $h \in D_i$ ($i \in N$) such that $\mu(h) > 0$. Furthermore, the plausibility order already takes care of the basic structure of histories and assessments and the probabilities can be added with ease, while starting with a lexicographic probability system on the set of histories as a primitive notion would require more complex notation and less transparent definitions. Furthermore, the qualitative notion of AGM-consistency has an appealing independent justification in terms of the AGM theory of belief revision.

and used to refine the notion of AGM-consistency and thus of perfect Bayesian equilibrium.[21]

A   Appendix

Although Lemma 9 follows easily from a result in Kreps and Wilson [14] we provide a proof for completeness.

Recall our assumption that $A = \bigcup_{h \in D} A(h)$. Thus, for every $a \in A$ there is an $h \in D$ such that $a \in A(h)$. Recall also the assumption that no action is available at more than one information set, that is, if $h, h' \in H$ are such that $A(h) \cap A(h') \neq \varnothing$ then $h' \in I(h)$.

For every $h \in H$ and $a \in A$ we write $a \in h$ if there exists an $h' \in H$ such that $h'a$ is a prefix of h (thus $a \in A(h')$), that is, if a is an action that occurs in history h.

Definition 16 *An A-weighting is a function $\lambda : A \to \mathbb{N}$ such that, for every $h \in D$, there is at least one $a \in A(h)$ with $\lambda(a) = 0$. Furthermore, if history h is assigned to chance, then $\lambda(a) = 0$ for every $a \in A(h)$.*

The following result is proved in Kreps and Wilson ([14], Lemma A.1, p. 887; we have re-written the result in terms of the notation used in this paper and slightly reworded it).

Lemma 17 *Fix an extensive form. If (σ, μ) is a KW-consistent assessment then there exists an A-weighting $\lambda : A \to \mathbb{N}$ such that, $\forall h \in D, \forall a \in A(h)$, (i) $\lambda(a) = 0$ if and only if $\sigma(a) > 0$, and (ii) $\mu(h) > 0$ if and only if $\Lambda(h) \leq \Lambda(h')$ for all $h' \in I(h)$, where $\Lambda(h) = \sum_{a \in h} \lambda(a)$.*

Proof of Lemma 9. Let (σ, μ) be a KW-consistent assessment. By Lemma 17 there exists an A-weighting $\lambda : A \to \mathbb{N}$ such that, $\forall h \in D, \forall a \in A(h)$, (i) $\lambda(a) = 0$ if and only if $\sigma(a) > 0$, and (ii) $\mu(h) > 0$ if and only if $\Lambda(h) \leq \Lambda(h')$ for all $h' \in I(h)$. Define the following total pre-order on the set of histories H: $h \precsim h'$ if and only if $\Lambda(h) \leq \Lambda(h')$. Then $\precsim$ is a plausibility order. In fact, if $h \in D$ and $a \in A(h)$ then $\Lambda(ha) = \Lambda(h) + \lambda(a) \geq \Lambda(h)$ (since $\lambda(a) \geq 0$), so that $h \precsim ha$; thus Property $PL1$ of Definition 3 is satisfied. By Definition 16, for every $h \in D$ there is an $a \in A(h)$ such that $\lambda(a) = 0$ and thus $\Lambda(ha) = \Lambda(h) + \lambda(a) = \Lambda(h)$, so that $ha \precsim h$; furthermore, if $a \in A(h)$ is such that $ha \precsim h$ then $\Lambda(ha) = \Lambda(h) + \lambda(a) = \Lambda(h)$, so that $\lambda(a) = 0$ and thus, for every $h' \in I(h)$, $\Lambda(h'a) = \Lambda(h') + \lambda(a) = \Lambda(h')$ and therefore $h'a \precsim h'$. Hence Property $PL2$ of Definition 3 is satisfied. Finally, by Definition 16, if h is assigned to chance then, for every $a \in A(h)$, $\lambda(a) = 0$ and thus $\Lambda(ha) = \Lambda(h) + \lambda(a) = \Lambda(h)$ so that $ha \precsim h$. Hence Property $PL3$ of Definition 3 is also satisfied. It is also clear that $\precsim$ is choice measurable (the function Λ provides an integer-valued representation of $\precsim$ that satisfies Property CM of Definition 8) and that $\precsim$ rationalizes (σ, μ) (Definition 6).[22] ∎

Before we prove Proposition 13 we recall the definition of subgame and of subgame-perfect equilibrium. Let $G = \left\langle A, H, N, \iota, \{\approx_i\}_{i \in N} \right\rangle$ be an extensive form and let $h_0 \in D$. We say that $G' = \left\langle A', H', N', \iota', \{\approx'_i\}_{i \in N'} \right\rangle$ is a *subgame of G with root h_0* if: (1) $I(h_0) = \{h_0\}$ (that is, the information set that contains h_0 consists of h_0 only; recall the convention - see Notation 1 - that is

---

[21] An in-depth analysis of appropriate qualitative notions of independence is left for future work. It is also possible to obtain a simple characterization of the backward-induction solution(s) in perfect information games in terms of AGM-consistency.

[22] By Definition 16, $\sigma(a) > 0$ if and only if $\lambda(a) = 0$ and, by definition of $\Lambda(\cdot)$, $ha \sim h$ if and only if $\lambda(a) = 0$. Furthermore, $\forall h \in D$, $\mu(h) > 0$ if and only if $\Lambda(h) \leq \Lambda(h')$ (and thus $h \precsim h'$) for all $h' \in I(h)$.

h_0 is assigned to chance then $I(h_0) = \{h_0\}$), (2) $H' = \{h \in H : h_0$ is a prefix of $h\}$ (that is, H' is the subset of H consisting of histories that have h_0 as a prefix), (3) $\forall h' \in H'$, $\forall i \in N$, $\forall h \in H$, if $h \approx_i h'$ then $h \in H'$ (that is, the information sets of G that contain an element of H' are entirely included in H'), (4) $A' = \{a \in A : a \in A(h')$ for some $h' \in H'\}$ (that is, A' is the subset of A consisting of those actions that are available at histories in H'), (5) $N' = \{i \in N : \iota^{-1}(i) \subseteq H'\}$ (that is, N' is the subset of players that are assigned to histories in H'), (6) ι' is the restriction of ι to H' and (7) $\forall i \in N'$, $\approx_i'$ is the restriction of $\approx_i$ to H', that is, $\approx_i' = \approx_i \cap (H' \times H')$

Let G be a game and σ a profile of behavior strategies. Then σ is a *subgame-perfect equilibrium* of G if, for every subgame G' of G, the restriction of σ to G' is a Nash equilibrium of G'.

Next we introduce some notation. Fix a game G and a profile of behavior strategies σ. Let $h, h' \in H$ with h a prefix of h', that is, $h' = ha_1...a_m$ for some actions $a_1, ..., a_m \in A$. Define $\mathbb{Q}_{h,\sigma}(h') = \sigma(a_1) \times ... \times \sigma(a_m)$ (that is, $\mathbb{Q}_{h,\sigma}(h')$ is the probability of reaching h' from h if play is according to σ) and if I is an information set let $\mathbb{Q}_{h,\sigma}(I) = \sum_{x \in I} \mathbb{Q}_{h,\sigma}(x)$ be the probability of reaching I from h, if play is according to σ. An assessment (σ, μ) is *weakly consistent* if for every information set I, whenever the probability of reaching I *from the null history* (that is, from the root of the tree) is positive (that is, $\mathbb{Q}_{\emptyset,\sigma}(I) > 0$) then, for every $h \in I$, $\mu(h)$ is obtained from $\mathbb{Q}_{\emptyset,\sigma}(\cdot)$ by using Bayes' rule:

$$\text{if } I \text{ is an information set and } \mathbb{Q}_{\emptyset,\sigma}(I) > 0 \text{ then, } \forall h \in I, \quad \mu(h) = \frac{\mathbb{Q}_{\emptyset,\sigma}(h)}{\mathbb{Q}_{\emptyset,\sigma}(I)}. \tag{WC}$$

The following result is well-known (see, for example, [17] p. 329).

Lemma 18 *Fix an extensive-form game G and an assessment (σ, μ). If (σ, μ) is weakly consistent and sequentially rational, then σ is a Nash equilibrium.*

Proof of Proposition 13 . Fix an extensive-form game G and let (σ, μ) be a perfect Bayesian equilibrium, that is, (σ, μ) is Bayesian AGM-consistent and sequentially rational (Definition 12). Let $\precsim$ be a plausibility order on the set of histories H that rationalizes (σ, μ) (Definition 6) and relative to which $\precsim$ is Bayesian (Definition 10). Let G' be an arbitrary subgame of G. We need to show that the restriction of σ to G' is a Nash equilibrium of G'. Since (σ, μ) is sequentially rational in G, the restriction of (σ, μ) to G' is sequentially rational in G'. Hence, by Lemma 18, it is sufficient to show that the restriction of (σ, μ) to G' is weakly consistent in G', that is, if h_0 is the root of G', then for every information set $I \subseteq H'$,

$$\text{if } \mathbb{Q}_{h_0,\sigma}(I) > 0 \text{ then, } \forall h \in I, \quad \mu(h) = \frac{\mathbb{Q}_{h_0,\sigma}(h)}{\mathbb{Q}_{h_0,\sigma}(I)}. \tag{WC'}$$

Fix an arbitrary information set $I \subseteq H'$ and suppose that $\mathbb{Q}_{h_0,\sigma}(I) > 0$, that is, I is reached with positive probability if σ is played *starting from h_0*. Let E be the $\precsim$-equivalence class that contains h_0 and let $E^+ = \{h \in E : \mu(h) > 0\}$. Since $I(h_0) = \{h_0\}$, $\mu(h_0) = 1$.[23] Thus $h_0 \in E^+$. Next we show that

$$Min_{\precsim}I \subseteq E^+. \tag{5}$$

Since $\mathbb{Q}_{h_0,\sigma}(I) > 0$, there is an $h_1 \in I$ such that $h_1 = h_0a_1...a_m$ for some actions $a_1, ..., a_m$ such that $\sigma(a_j) > 0$ for all $j = 1, ..., m$. Thus, by Property $P1$ of Definition 6, every action a_j is plausibility

---

[23] Recall the convention that if h is assigned to chance, then $I(h) = \{h\}$ and $\mu(h) = 1$.

preserving (see Remark 5) and hence, by transitivity of $\precsim$, $h_1 \sim h_0$, so that $h_1 \in E$. Since (by definition of subgame) every history in H' has h_0 as a prefix, by Remark 5 $h_0 \precsim x$, for every $x \in H'$. From this and the fact that $h_1 \sim h_0$, it follows that $h_1 \in Min_{\precsim}I$ (since $I \subseteq H'$), so that, by Property $P2$ of Definition 6, $\mu(h_1) > 0$. Thus $h_1 \in E^+$. Now fix an arbitrary $h \in Min_{\precsim}I$. Then $h \sim h_1$ and thus (since $h_1 \in E$), $h \in E$. By Property $P2$ of Definition 6, $\mu(h) > 0$. Hence $h \in E^+$. Thus (5) holds. By $B1$ of Definition 10, there is a probability measure ν_E on H whose support coincides with E^+. Next we show that

$$\forall h \in I, \quad \nu_E(h) = \nu_E(h_0) \times \mathbb{Q}_{h_0,\sigma}(h). \tag{6}$$

Fix an arbitrary $h \in I$. Then $h = h_0 a_1 ... a_m$ for some actions $a_1, ..., a_m$, so that $\mathbb{Q}_{h_0,\sigma}(h) = \sigma(a_1) \times ... \times \sigma(a_m)$. If $h \in Min_{\precsim}I$, then, by (5), $h \in E^+$. Thus, since $h_0 \in E^+$, by $B2$ of Definition 10, $\nu_E(h) = \nu_E(h_0) \times (\sigma(a_1) \times ... \times \sigma(a_m)) = \nu_E(h_0) \times \mathbb{Q}_{h_0,\sigma}(h)$. If $h \notin Min_{\precsim}I$ then $h_0 \prec h$.[24] Hence $h = h_0 c_1 ... c_r$ for some actions $c_1, ..., c_r$ at least one of which - say c_j - is not plausibility preserving.[25] Hence, by Property $P1$ of Definition 6, $\sigma(c_j) = 0$ and thus $\mathbb{Q}_{h_0,\sigma}(h) = \sigma(c_1) \times ... \times \sigma(c_r) = 0$. Hence $\nu_E(h) = \mathbb{Q}_{h_0,\sigma}(h) = 0 = \nu_E(h_0) \times \mathbb{Q}_{h_0,\sigma}(h)$. Thus (6) holds.

To complete the proof of (WC'), fix an arbitrary $h \in I$. Then, by (6),

$$\frac{\nu_E(h)}{\nu_E(I)} = \frac{\nu_E(h_0) \times \mathbb{Q}_{h_0,\sigma}(h)}{\sum\limits_{x \in I} \nu_E(h_0) \times \mathbb{Q}_{h_0,\sigma}(x)} = \frac{\nu_E(h_0) \times \mathbb{Q}_{h_0,\sigma}(h)}{\nu_E(h_0) \sum\limits_{x \in I} \mathbb{Q}_{h_0,\sigma}(x)} = \frac{\mathbb{Q}_{h_0,\sigma}(h)}{\mathbb{Q}_{h_0,\sigma}(I)}. \tag{7}$$

If $h \in Min_{\precsim}I$, then, by $B3$ of Definition 10, $\mu(h) = \frac{\nu_E(h)}{\nu_E(I)}$ and thus, by (7), $\mu(h) = \frac{\mathbb{Q}_{h_0,\sigma}(h)}{\mathbb{Q}_{h_0,\sigma}(I)}$. If $h \notin Min_{\precsim}I$, then, by Property $P2$ of Definition 6, $\mu(h) = 0$ and, as shown above $\mathbb{Q}_{h_0,\sigma}(h) = 0$ so that also in this case $\mu(h) = \frac{\mathbb{Q}_{h_0,\sigma}(h)}{\mathbb{Q}_{h_0,\sigma}(I)}$ (recall that, by hypothesis, $\mathbb{Q}_{h_0,\sigma}(I) > 0$). Thus (WC') holds and, therefore, the restriction of (σ, μ) to the subgame G' with root h_0 is weakly consistent in G' so that, by Lemma 18, the restriction of σ to G' is a Nash equilibrium of G'. $\blacksquare$

In order to prove Proposition 14 we need some preliminary definitions and lemmas.

Definition 19 *Fix an extensive-form game. Let (σ, μ) be an assessment and $\precsim$ a plausibility order on the set of histories H that rationalizes (σ, μ). Let $E \subseteq H$ be an $\precsim$-equivalence class and let $E^+ = \{h \in E : \mu(h) > 0\}$. A path in E^+ is a sequence $\langle h_1, ..., h_m \rangle$ in E^+ such that, for every $i = 1, ..., m - 1$, either h_i is a prefix of h_{i+1} or h_{i+1} is a prefix of h_i or $h_{i+1} \in I(h_i)$. We say that $h, h' \in E^+$ are* linked *if there exists a path $\langle h_1, ..., h_m \rangle$ in E^+ with $h_1 = h$ and $h_m = h'$. A subset $F \subseteq E^+$ is* connected *if every pair in F is linked. A subset $F \subseteq E^+$ is* maximally connected *if it is connected and, for every connected $G \subseteq E^+$, if $F \subseteq G$ then $F = G$.*

For example, consider the (partial) extensive form of Figure 6 and a plausibility order $\precsim$ for which the following is an equivalence class: $E = \{a, b, c, d, em, en, bf, cg\}$. Suppose that $E^+ = E$. Then histories a and d are linked by the path $\langle a, b, bf, cg, c, d \rangle$ ($b \in I(a)$, b is a prefix of bf, $cg \in I(bf)$, c is a prefix of cg and $d \in I(c)$), while histories a and em are not linked; furthermore, the maximally connected subsets of E are $F_1 = \{a, b, c, d, bf, cg\}$ and $F_2 = \{em, en\}$.[26]

---

[24] If $h \notin Min_{\precsim}I$, then $h_1 \prec h$ (where $h_1 \in Min_{\precsim}I$: see above). Thus, since $h_0 \sim h_1$, $h_0 \prec h$.

[25] If all those actions were plausibility preserving than we would have $h_0 \sim h$.

[26] On the other hand, if $E^+ = \{a, c, d, bf, cg, em\}$ then the maximally connected subsets of E^+ are $\{a\}$ (note that $b \notin E^+$), $\{c, d, bf, cg\}$ and $\{em\}$.
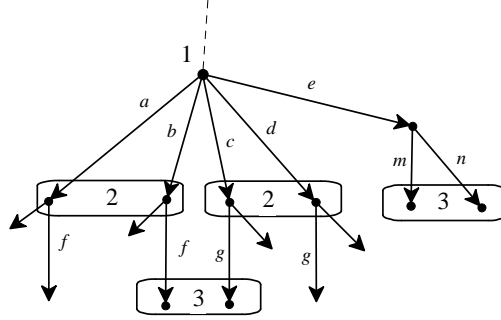
Illustration of Definition 19

Figure 6

Lemma 20 *Fix an extensive-form game. Let (σ, μ) be an assessment and $\precsim$ a plausibility order on the set of histories H that rationalizes (σ, μ). Let E be an $\precsim$-equivalence class and let $E^+ = \{h \in E : \mu(h) > 0\}$. If F and G are two maximally connected subsets of E^+ then either $F = G$ or $F \cap G = \varnothing$.*

Proof. Suppose that F and G are two maximally connected subsets of E^+ with $F \neq G$ and $F \cap G \neq \varnothing$. Fix an arbitrary $h \in F \cap G$ and arbitrary $h_1 \in F$ and $h_2 \in G$. Then there is a path from h_1 to h and there is a path from h to h_2. Joining these two paths we get that h_1 is linked to h_2. Thus every two histories in $F \cup G$ are linked and therefore $F \cup G$ is connected. Since $F \neq G$, $F \cup G$ is a proper superset of either F or G, contradicting the hypothesis that F and G are maximally connected. ∎

Corollary 21 *Fix an extensive-form game. Let (σ, μ) be an assessment and $\precsim$ a plausibility order on the set of histories H that rationalizes (σ, μ). Let E be an $\precsim$-equivalence class and let $E^+ = \{h \in E : \mu(h) > 0\}$. Then E^+ can be partitioned into a collection $\{E_1, ..., E_r\}$ of maximally connected subsets. Furthermore, the collection $\{E_1, ..., E_r, E^0\}$, where $E^0 = \{h \in E : \mu(h) = 0\}$, is a partition of E.*

Corollary 21 is an immediate consequence of Lemma 20 and finiteness of H.

Lemma 22 *Fix an extensive-form game. Let (σ, μ) be a KW-consistent assessment and let $\langle \sigma^1, ..., \sigma^m, ...\rangle$ be a sequence of completely mixed strategy profiles such that, letting (for every $h \in D$ and for every $m \geq 1$) $\mu^m(h) = \frac{\mathbb{Q}_{\varnothing, \sigma^m}(h)}{\sum\limits_{x \in I(h)} \mathbb{Q}_{\varnothing, \sigma^m}(x)}$, $\lim_{m \to \infty}(\sigma^m, \mu^m) = (\sigma, \mu)$. Let I be an information set and $I^+ = \{h \in I : \mu(h) > 0\}$. Then,*

$$\forall h \in I^+, \quad \mu(h) = \lim_{m \to \infty} \frac{\mathbb{Q}_{\varnothing, \sigma^m}(h)}{\sum\limits_{x \in I^+} \mathbb{Q}_{\varnothing, \sigma^m}(x)}. \tag{8}$$

Proof. To simplify the notation we write $\nu^m(h)$ instead of $\mathbb{Q}_{\varnothing, \sigma^m}(h)$. By hypothesis, for every $h \in I$, $\mu(h) = \lim_{m \to \infty} \frac{\nu^m(h)}{\sum\limits_{x \in I} \nu^m(x)}$. Thus we need to show that, for every $h \in I^+$, $\lim_{m \to \infty} \frac{\nu^m(h)}{\sum\limits_{x \in I} \nu^m(x)} = $

19

$\lim_{m\to\infty} \dfrac{\nu^m(h)}{\sum\limits_{x\in I^+}\nu^m(x)}$. Fix an arbitrary $h \in I^+$. Since $I^+ \subseteq I$, $\sum\limits_{x\in I^+}\nu^m(x) \leq \sum\limits_{x\in I}\nu^m(x)$. Thus, for every $y \in I^+$, $\dfrac{\nu^m(y)}{\sum\limits_{x\in I}\nu^m(x)} \leq \dfrac{\nu^m(y)}{\sum\limits_{x\in I^+}\nu^m(x)}$. Hence

$$\lim_{m\to\infty} \frac{\nu^m(y)}{\sum\limits_{x\in I^+}\nu^m(x)} \geq \lim_{m\to\infty} \frac{\nu^m(y)}{\sum\limits_{x\in I}\nu^m(x)} = \mu(y), \quad \forall y \in I^+. \tag{9}$$

In particular, $\lim_{m\to\infty} \dfrac{\nu^m(h)}{\sum\limits_{x\in I^+}\nu^m(x)} \geq \mu(h)$. Suppose that $\lim_{m\to\infty} \dfrac{\nu^m(h)}{\sum\limits_{x\in I^+}\nu^m(x)} > \mu(h)$. Then, summing over I^+ and using (9) we get that

$$\sum_{y\in I^+} \lim_{m\to\infty} \frac{\nu^m(y)}{\sum\limits_{x\in I^+}\nu^m(x)} > \sum_{y\in I^+} \mu(y). \tag{10}$$

By definition of I^+, the right-hand side of (10) is equal to 1 and the left-hand side is equal to $\lim_{m\to\infty} \sum_{y\in I^+} \dfrac{\nu^m(y)}{\sum\limits_{x\in I^+}\nu^m(x)} = \lim_{m\to\infty} 1 = 1$, yielding a contradiction. Thus $\lim_{m\to\infty} \dfrac{\nu^m(h)}{\sum\limits_{x\in I^+}\nu^m(x)} = \mu(h)$. ∎

Proof of Proposition 14. Fix an extensive-form game and let (σ, μ) be a sequential equilibrium, that is, (σ, μ) is sequentially rational and KW-consistent. By Lemma 9, KW-consistency implies AGM-consistency. Fix an arbitrary plausibility order $\precsim$ that rationalizes (σ, μ) (e.g. the one described in the proof of Lemma 9). We need to show that (σ, μ) is Bayesian relative to $\precsim$ (see Definition 10). Fix an arbitrary $\precsim$-equivalence class E that contains some decision history $h \in D_i$ ($i \in N$) such that $\mu(h) > 0$. Let $E^+ = \{h \in E : \mu(h) > 0\}$. Thus $E^+ \neq \varnothing$. We need to show that there is a probability measure ν_E on H that satisfies the three properties of Definition 10. By definition of KW-consistency, there is a sequence $\langle \sigma^1, ..., \sigma^m, ... \rangle$ of completely mixed behavior-strategy profiles such that (1) $\lim_{m\to\infty} \sigma^m = \sigma$ and (2) for every $h \in H$, $\lim_{m\to\infty} \mu^m(h) = \mu(h)$ where $\mu^m(h) = \dfrac{\mathbb{Q}_{\varnothing,\sigma^m}(h)}{\sum\limits_{x\in I(h)}\mathbb{Q}_{\varnothing,\sigma^m}(x)}$. As before, to simplify the notation, we denote $\mathbb{Q}_{\varnothing,\sigma^m}(h)$ by $\nu^m(h)$.

Let $\{E_1, ..., E_r\}$ be a partition of E^+ into maximally connected subsets (see Corollary 21). Fix an arbitrary E_i ($i = 1, ..., r$) and define the following measure $\nu_{E_i} : H \to [0, 1]$:[27]

$$\nu_{E_i}(h) = \begin{cases} 0 & \text{if } h \notin E_i \\ \lim_{m\to\infty} \dfrac{\nu^m(h)}{\sum\limits_{y\in E_i}\nu^m(y)} & \text{if } h \in E_i \end{cases} \tag{11}$$

---

[27] Note that, since σ^m is completely mixed, $0 < \nu^m(h) < 1$ for every $h \in H$ and thus, for every non-empty $F \subseteq H$ and every $h \in F$, $0 < \dfrac{\nu^m(h)}{\sum\limits_{y\in F}\nu^m(y)} \leq 1$. Hence, by the Bolzano–Weierstrass theorem, the sequence $\left\langle \dfrac{\nu^m(h)}{\sum\limits_{y\in E_i}\nu^m(y)} \right\rangle_{m=1,2,...}$ has a convergent subsequence. Thus, if necessary, we can take a convergent subsequence to define the $\lim_{m\to\infty} \dfrac{\nu^m(h)}{\sum\limits_{y\in E_i}\nu^m(y)}$.

First we show that

$$Supp(\nu_{E_i}) = E_i \tag{12}$$

By construction $Supp(\nu_{E_i}) \subseteq E_i$. Fix an arbitrary $h \in E_i$. We need to show that $\nu_{E_i}(h) > 0$. Fix an arbitrary $h' \in Supp(\nu_{E_i})$. Then (see Definition 19) there is a path $\langle h_1, ..., h_m \rangle$ in E_i with $h_1 = h'$ and $h_m = h$ and for every $j \in \{1, ..., m-1\}$ either h_j is a prefix of h_{j+1} or h_{j+1} is a prefix of h_j or $h_{j+1} \in I(h_j)$. Since $h_1 \in Supp(v_{E_i})$, $v_{E_i}(h_1) > 0$. Next we show that, for every $j \in \{1, ..., m-1\}$ if $v_{E_i}(h_j) > 0$ then $v_{E_i}(h_{j+1}) > 0$. Fix $j \in \{1, ..., m-1\}$ and suppose that $v_{E_i}(h_j) > 0$. We need to consider three cases.

Case 1: h_j is a prefix of h_{j+1}. Then $h_{j+1} = h_j a_1 ... a_s$ for some actions $a_1, ..., a_s$. Since $h_j \sim h_{j+1}$ (they belong to the same $\precsim$-equivalence class E) every action a_k ($k = 1, ..., s$) is plausibility preserving and thus, by Property $P1$ of Definition 6, $\sigma(a_k) > 0$. Now, by definition of $\nu^m(\cdot)$, $\nu^m(h_{j+1}) = \nu^m(h_j) \times \sigma^m(a_1) \times ... \times \sigma^m(a_s)$. Thus $\nu_{E_i}(h_{j+1}) = \lim_{m\to\infty} \frac{\nu^m(h_{j+1})}{\sum_{y \in E_i} \nu^m(y)} =$

$\lim_{m\to\infty} \frac{\nu^m(h_j)}{\sum_{y \in E_i} \nu^m(y)} (\sigma^m(a_1) \times ... \times \sigma^m(a_s)) = \lim_{m\to\infty} \frac{\nu^m(h_j)}{\sum_{y \in E_i} \nu^m(y)} \times \lim_{m\to\infty} (\sigma^m(a_1) \times ... \times \sigma^m(a_s)) =$

$\nu_{E_i}(h_j) \times \sigma(a_1) \times ... \times \sigma(a_s)$. Thus $\nu_{E_i}(h_{j+1}) > 0$ (since $\nu_{E_i}(h_j) > 0$ by hypothesis and $\sigma(a_k) > 0$ for every $k = 1, ..., s$).

Case 2: h_{j+1} is a prefix of h_j. Then $h_j = h_{j+1} a_1 ... a_s$ for some actions $a_1, ..., a_s$. Since $h_j \sim h_{j+1}$ every action a_k is plausibility preserving and thus $\sigma(a_k) > 0$. By definition of $\nu^m(\cdot)$, $\nu^m(h_j) = \nu^m(h_{j+1}) \times \sigma^m(a_1) \times ... \times \sigma^m(a_s)$. Thus
$\nu_{E_i}(h_j) = \lim_{m\to\infty} \frac{\nu^m(h_j)}{\sum_{y \in E_i} \nu^m(y)} = \lim_{m\to\infty} \frac{\nu^m(h_{j+1})}{\sum_{y \in E_i} \nu^m(y)} (\sigma^m(a_1) \times ... \times \sigma^m(a_s)) =$

$\lim_{m\to\infty} \frac{\nu^m(h_{j+1})}{\sum_{y \in E_i} \nu^m(y)} \times \lim_{m\to\infty} (\sigma^m(a_1) \times ... \times \sigma^m(a_s)) = \nu_{E_i}(h_{j+1}) \times \sigma(a_1) \times ... \times \sigma(a_s)$

and therefore $\nu_{E_i}(h_{j+1}) > 0$.

Case 3: $h_{j+1} \in I(h_j)$. First of all, note that

$$\lim_{m\to\infty} \frac{\sum_{x \in I(h_j)} \nu^m(x)}{\sum_{y \in E_i} \nu^m(y)} > 0. \tag{13}$$

In fact, since $h_j \in I(h_j)$, $\frac{\sum_{x \in I(h_j)} \nu^m(x)}{\sum_{y \in E_i} \nu^m(y)} \geq \frac{\nu^m(h_j)}{\sum_{y \in E_i} \nu^m(y)}$ and $\lim_{m\to\infty} \frac{\nu^m(h_j)}{\sum_{y \in E_i} \nu^m(y)} = \nu_{E_i}(h_j)$ and, by

hypothesis, $\nu_{E_i}(h_j) > 0$. Now,

$$\nu_{E_i}(h_{j+1}) = \lim_{m\to\infty} \frac{\nu^m(h_{j+1})}{\sum_{y \in E_i} \nu^m(y)} = \lim_{m\to\infty} \left(\frac{\nu^m(h_{j+1})}{\sum_{x \in I(h_j)} \nu^m(x)} \times \frac{\sum_{x \in I(h_j)} \nu^m(x)}{\sum_{y \in E_i} \nu^m(y)} \right) =$$

$\lim_{m\to\infty} \frac{\nu^m(h_{j+1})}{\sum_{x \in I(h_j)} \nu^m(x)} \times \lim_{m\to\infty} \frac{\sum_{x \in I(h_j)} \nu^m(x)}{\sum_{y \in E_i} \nu^m(y)}$. Since $\lim_{m\to\infty} \frac{\nu^m(h_{j+1})}{\sum_{x \in I(h_j)} \nu^m(x)} = \mu(h_{j+1})$ (recall that,

by hypothesis, $h_{j+1} \in I(h_j)$ and thus $I(h_j) = I(h_{j+1})$) and $\mu(h_{j+1}) > 0$ (because $h_{j+1} \in E^+$) it follows, by (13), that $\nu_{E_i}(h_{j+1}) > 0$. This completes the proof of (12). Note that

21

$$\text{if } h, h' \in E_i \text{ and } h' = ha_1...a_s, \text{ then } \nu_{E_i}(h') = \nu_{E_i}(h) \times \sigma(a_1) \times ... \times \sigma(a_s) \tag{14}$$

(the proof is a repetition of the argument used above in Case 1). Next we show that

$$\text{if } h \in Min_{\precsim}I(h) \cap E_i \text{ then } Min_{\precsim}I(h) \subseteq E_i. \tag{15}$$

Let $h \in Min_{\precsim}I(h) \cap E_i$ and fix an arbitrary $h' \in Min_{\precsim}I(h)$. Then $h' \sim h$ and thus $h' \in E$. By Property $P2$ of Definition 6, $\mu(h') > 0$, so that $h' \in E^+$. Furthermore $\langle h, h' \rangle$ is a path in E^+ (see Definition 19) and thus, since $h \in E_i$, it follows that $h' \in E_i$ (because E_i is a maximally connected subset of E). Next we show that

$$\begin{aligned} &\text{if } I \text{ is an information set such that } Min_{\precsim}I \subseteq E_i \text{ then,} \\ &(a) \quad \sum_{x \in I} \nu_{E_i}(x) = \sum_{x \in Min_{\precsim}I} \nu_{E_i}(x), \text{ and} \\ &(b) \quad \forall h \in I, \quad \mu(h) = \nu_{E_i}(h \mid I) = \frac{\nu_{E_i}(h)}{\nu_{E_i}(I)}. \end{aligned} \tag{16}$$

Let I be such that $Min_{\precsim}I \subseteq E_i$. If $h \notin Min_{\precsim}I$ then, by Property $P2$ of Definition 6, $\mu(h) = 0$. Hence $h \notin E^+$ (and thus $h \notin E_i$) so that $\nu_{E_i}(h) = 0$. Thus part (a) of (16) holds. Hence to prove part (b) it is sufficient to show that $\forall h \in Min_{\precsim}I, \quad \mu(h) = \frac{\nu_{E_i}(h)}{\nu_{E_i}(Min_{\precsim}I)}$.[28] Furthermore, by Lemma 9, $\forall h \in Min_{\precsim}I, \mu(h) = \lim_{m \to \infty} \frac{\nu^m(h)}{\sum_{x \in Min_{\precsim}I} \nu^m(x)}$.[29] Thus it will be sufficient to prove that,

$\forall h \in Min_{\precsim}I, \frac{\nu_{E_i}(h)}{\nu_{E_i}(Min_{\precsim}I)} = \lim_{m \to \infty} \frac{\nu^m(h)}{\sum_{x \in Min_{\precsim}I} \nu^m(x)}$. We have that

$$\frac{\nu_{E_i}(h)}{\nu_{E_i}(Min_{\precsim}I)} = \frac{\nu_{E_i}(h)}{\sum_{x \in Min_{\precsim}I} \nu_{E_i}(x)} = \sum_{x \in Min_{\precsim}I} \frac{\lim_{m \to \infty} \frac{\nu^m(h)}{\sum_{y \in E_i} \nu^m(y)}}{\lim_{m \to \infty} \frac{\nu^m(x)}{\sum_{y \in E_i} \nu^m(y)}} = \frac{\lim_{m \to \infty} \frac{\nu^m(h)}{\sum_{y \in E_i} \nu^m(y)}}{\lim_{m \to \infty} \sum_{x \in Min_{\precsim}I} \frac{\nu^m(x)}{\sum_{y \in E_i} \nu^m(y)}}$$

$$= \lim_{m \to \infty} \frac{\frac{\nu^m(h)}{\sum_{y \in E_i} \nu^m(y)}}{\sum_{x \in Min_{\precsim}I} \frac{\nu^m(x)}{\sum_{y \in E_i} \nu^m(y)}} = \lim_{m \to \infty} \frac{\frac{1}{\sum_{y \in E_i} \nu^m(y)} \nu^m(h)}{\frac{1}{\sum_{y \in E_i} \nu^m(y)} \sum_{x \in Min_{\precsim}I} \nu^m(x)} = \lim_{m \to \infty} \frac{\nu^m(h)}{\sum_{x \in Min_{\precsim}I} \nu^m(x)}.$$

If $E_i = E^+$ then the proof of Proposition 14 is complete: by (12), (14) and (16) the measure ν_{E_i} satisfies the three properties of Definition 10. If E_i is a proper subset of E^+, then let $\alpha_1, ..., \alpha_r$ be arbitrary numbers such that $0 < \alpha_i < 1$, for all $i = 1, ..., r$, and $\sum_{i=1}^{r} \alpha_i = 1$ and let $\nu_E : H \to [0,1]$ be defined by $\nu_E(h) = \sum_{i=1}^{r} \alpha_i v_{E_i}(h)$. Then, by (12), (14) and (16) ν_E satisfies properties of Definition 10 [30] and thus (σ, μ) is Bayesian relative to $\precsim$. ∎

---

[28] Note that $\nu_{E_i}(Min_{\precsim}I) > 0$. Proof: fix an arbitrary $h \in Min_{\precsim}I$ (since I is finite $Min_{\precsim}I \neq \varnothing$). Then $h \in E_i$ and thus, by (12), $\nu_{E_i}(h) > 0$. Hence, since $\nu_{E_i}(Min_{\precsim}I) \geq \nu_{E_i}(h)$, $\nu_{E_i}(Min_{\precsim}I) > 0$.

[29] By Property $P2$ of Definition 6, $I^+ \equiv \{h \in I : \mu(h) > 0\} = Min_{\precsim}I$.

[30] If I is an information set such that $Min_{\precsim}I \subseteq E_i$ (for some $i = 1, ..., r$) and $h \in I$ then $\nu_E(h|I) = \frac{\alpha_i \nu_{E_i}(h)}{\alpha_i \nu_{E_i}(I)} = \nu_{E_i}(h|I)$.

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] Arló-Costa, Horacio and Rohit Parikh, Conditional probability and defeasible inference, *Journal of Philosophical Logic*, 2005, 34: 97-119.

[3] Aumann, Robert and Adam Brandenburger, Epistemic conditions for Nash equilibrium, *Econometrica*, 1995, 63: 1161-1180.

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

[5] Battigalli, Pierpaolo, Alfredo Di Tillio and Dov Samet, Strategies and interactive beliefs in dynamic games, Technical Report IGIER WP 375, Bocconi University, January 2011.

[6] Bonanno, Giacomo, Memory and perfect recall in extensive games, *Games and Economic Behavior*, 2004, 47: 237-256.

[7] Bonanno, Giacomo, Rational choice and AGM belief revision, *Artificial Intelligence*, 2009, 173: 1194-1203.

[8] Bonanno, Giacomo, Perfect Bayesian equilibrium. Part II: epistemic foundation and refinements, Working Paper, University of California Davis, 2011.

[9] 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.

[10] Fudenberg, Drew and Jean Tirole, Perfect Bayesian equilibrium and sequential equilibrium, *Journal of Economic Theory*, 1991, 53: 236-260.

[11] Hammond, Peter, Elementary non-Archimedean representations of probability for decision theory and games, in: P. Humphreys (ed.), *Patrick Suppes: scientific philosopher. Vol. I: probability and probabilistic causality*, Kluwer Academic Publishers, 1994, chapter 2, 25–59.

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

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

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

[15] Kreps, David and Garey Ramey, Structural consistency, consistency, and sequential rationality, *Econometrica*, 1987, 55: 1331-1348.

[16] Myerson, Roger, *Game Theory*, Harvard University Press, 1991.

[17] Osborne, Martin, *An introduction to game theory*, Oxford University Press, 2004.

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

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

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

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

[22] von Stengel, Bernhard and Françoise Forges, Extensive form correlated equilibrium: definition and computational complexity. *Mathematics of Operations Research*, 2008, 33: 1002-1022.