Bonanno, Giacomo

Working Paper

Perfect Bayesian equilibrium. Part II: epistemic foundations

Working Paper, No. 11-1

Provided in Cooperation with:
University of California Davis, Department of Economics

Suggested Citation: Bonanno, Giacomo (2011) : Perfect Bayesian equilibrium. Part II: epistemic foundations, Working Paper, No. 11-1, University of California, Department of Economics, Davis, CA

This Version is available at:
https://hdl.handle.net/10419/58372

UCDAVIS
DEPARTMENT OF ECONOMICS
Working Paper Series

Perfect Bayesian equilibrium. Part II: epistemic foundations.

Giacomo Bonanno
University of California Davis

January 19, 2011

Paper # 11-1

In a companion paper we introduced a general notion of perfect Bayesian equilibrium which can be applied to arbitrary extensive-form games. The essential ingredient of the proposed definition is the qualitative notion of AGM-consistency. In this paper we provide an epistemic foundation for AGM-consistency based on the AGM theory of belief revision.

Department of Economics
One Shields Avenue
Davis, CA 95616
(530)752-0741

http://www.econ.ucdavis.edu/working_search.cfm

Perfect Bayesian equilibrium. Part II: epistemic foundations.

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

January 2011

Abstract

In a companion paper we introduced a general notion of perfect Bayesian equilibrium which can be applied to arbitrary extensive-form games. The essential ingredient of the proposed definition is the qualitative notion of AGM-consistency. In this paper we provide an epistemic foundation for AGM-consistency based on the AGM theory of belief revision.

Keywords: belief revision, common prior, plausibility order, perfect Bayesian equilibrium, consistency, sequential equilibrium.

1 Introduction

In an earlier paper ([5]) we introduced a general notion of perfect Bayesian equilibrium, which can be applied to arbitrary extensive-form games and

*Some of the material in this paper was presented at the the ninth conference on Logic and the Foundations of the Theory of Games and Decision (LOFT9), Toulouse, July 2010 and at the Workshop on Epistemic Game Theory, SUNY Stony Brook, July 2010.

is intermediate between subgame-perfect equilibrium and sequential equilibrium. A perfect Bayesian equilibrium is defined in [5] as an assessment[1] (σ, μ) which is (1) sequentially rational, (2) "AGM-consistent" and (3) "Bayes' rule compatible". The essential ingredient is the qualitative notion of AGM-consistency, which requires the assessment (σ, μ) to be rationalizable by a plausibility order on the set of histories, in the sense that (1) the histories that are assigned positive probability by the "system of beliefs" μ are the most plausible in each information set and (2) the choices that are assigned positive probability by the strategy profile σ are those that 'preserve plausibility'. The precise definitions are as follows.[2]

Definition 1 *A* plausibility order *is a total pre-order*[3] $\precsim$ *on the set of histories* H*, that satisfies the following properties:* $\forall h \in D$,[4]

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

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

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 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. 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 $h' \not\precsim h$". By PL1 and PL2, for every decision history h, there is at least one action a at h such that (1) $h \sim ha$ (that is, ha is as plausible as h) and

[1] The definitions of extensive-form game, assessment, etc. are reviewed in Appendix A.

[2] For simplicity in this paper we restrict attention to extensive-form games without chance moves.

[3] That is, a binary relation $\precsim$ on 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 $h \precsim h''$).

[4] D is the set of decision histories. $A(h)$ denotes the set of actions or choices available at $h \in D$ and $I(h)$ denotes the information set that contains h (of the player who moves at h). If $a \in A(h)$, ha denotes the history that results from appending a to h. For details see Appendix A.

(2) if h' belongs to the same information set as h, then $h' \sim h'a$. We call such actions *plausibility preserving.*

Definition 2 *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:* $\forall h \in D, \forall a \in A(h)$,

$$\sigma(a) > 0 \text{ if and only if } h \sim ha, \tag{C1}$$

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

$$\mu(h) > 0 \text{ if and only if } h \precsim h', \forall h' \in I(h). \tag{C2}$$

If $\precsim$ *satisfies properties C1 and C2 with respect to* (σ, μ)*, we say that* $\precsim$ rationalizes (σ, μ).

In this paper we provide an epistemic justification for the notion of AGM-consistency in terms of the theory of belief revision proposed by Alchourrón, Gärdenfors and Makinson [1] (see also [9]), known as the AGM theory. In Section 2 we review the syntactic AGM theory and a set-theoretic semantics based on choice frames. In Section 3 we show that choice frames can be used in extensive-form games to encode the players' initial (qualitative) beliefs and disposition to change those beliefs in response to new information. The main result is that if the (conditional, qualitative) beliefs of each player are compatible with the AGM postulates for belief revision and satisfy three natural properties and the players have a "common prior", then the profile of beliefs can be rationalized by a plausibility ordering, thus providing an epistemic basis for the notion of AGM-consistent assessment.

2 AGM belief revision and choice frames

Let Φ be the set of formulas of a propositional language based on a countable set S of atoms.[5] Given a subset $K \subseteq \Phi$, its deductive closure, denoted by

[5] The set Φ is built recursively from the set S of atomic propositions and the connectives $\neg$ (for "not") and $\vee$ (for "or") as follows: if $p \in S$ then $p \in \Phi$ and if $\phi, \psi \in \Phi$ then $\neg\phi \in \Phi$

$[K]$, is defined as follows: $\psi \in [K]$ if and only if there exist $\phi_1, ..., \phi_n \in K$ (with $n \geq 0$) such that $(\phi_1 \wedge ... \wedge \phi_n) \rightarrow \psi$ is a tautology. A set $K \subseteq \Phi$ is *deductively closed* if $K = [K]$ and it is *consistent* if $[K] \neq \Phi$. Let K be a consistent and deductively closed set representing the agent's initial beliefs and let $\Psi \subseteq \Phi$ be a set of formulas representing possible items of information. A *belief revision function based on* K is a function $B_K : \Psi \rightarrow 2^\Phi$ (where 2^Φ denotes the set of subsets of Φ) that associates with every formula $\phi \in \Psi$ (thought of as new information) a set $B_K(\phi) \subseteq \Phi$ (thought of as the revised beliefs). If $\Psi \neq \Phi$ then B_K is called a *partial* belief revision function, while if $\Psi = \Phi$ then B_K is called a *full* belief revision function.

Let $B_K : \Psi \rightarrow 2^\Phi$ be a (partial) belief revision function and $B^*_K : \Phi \rightarrow 2^\Phi$ a full belief revision function. We say that B^*_K is an *extension* of B_K if, for every $\phi \in \Psi$, $B^*_K(\phi) = B_K(\phi)$.

A *full* belief revision function is called an *AGM function* if it satisfies the following properties, known as the AGM postulates: $\forall \phi, \psi \in \Phi$,

(AGM1) $B_K(\phi) = [B_K(\phi)]$.
(AGM2) $\phi \in B_K(\phi)$.
(AGM3) $B_K(\phi) \subseteq [K \cup \{\phi\}]$.
(AGM4) if $\neg\phi \notin K$, then $[K \cup \{\phi\}] \subseteq B_K(\phi)$.
(AGM5) $B_K(\phi) = \Phi$ if and only if ϕ is a contradiction.
(AGM6) if $\phi \leftrightarrow \psi$ is a tautology then $B_K(\phi) = B_K(\psi)$.
(AGM7) $B_K(\phi \wedge \psi) \subseteq [B_K(\phi) \cup \{\psi\}]$.
(AGM8) if $\neg\psi \notin B_K(\phi)$, then $[B_K(\phi) \cup \{\psi\}] \subseteq B_K(\phi \wedge \psi)$.

AGM1 requires the revised belief set to be deductively closed. AGM2 requires that the information be believed. AGM3 says that beliefs should be revised minimally, in the sense that no new formula should be added unless it can be deduced from the information received and the initial beliefs. AGM4 says that if the information received is compatible with the initial beliefs, then any formula that can be deduced from the information and the initial beliefs should be part of the revised beliefs. AGM5 requires the revised beliefs to be consistent, unless the information ϕ is a contradiction (that is, $\neg\phi$ is a tautology). AGM6 requires that if ϕ is equivalent to ψ then the result of revising by ϕ be identical to the result of revising by ψ. AGM7 and AGM8

and $(\phi \vee \psi) \in \Phi$. The connectives $\wedge$ (for "and"), $\rightarrow$ (for "if ... then ...") and $\leftrightarrow$ (for "if and only if") are defined as usual: $\phi \wedge \psi = \neg(\neg\phi \vee \neg\psi)$, $\phi \rightarrow \psi = \neg\phi \vee \psi$ and $\phi \leftrightarrow \psi = (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$.

are a generalization of AGM3 and AGM4 that requires $B_K(\phi \wedge \psi)$ to coincide with the expansion of $B_K(\phi)$ by ψ, as long as ψ is compatible with $B_K(\phi)$.

Applying the AGM theory of belief revision to the analysis of extensive-form games is problematic, since the AGM postulates are expressed in a syntactic framework, while extensive-form games are set-theoretic constructs. In [4] the notion of choice frame is used to provide a link between the syntactic AGM approach and set-theoretic structures.

Definition 3 *A* choice frame *is a triple* $\langle \Omega, \mathcal{E}, f \rangle$ *where*

Ω *is a non-empty set of* states*; subsets of* Ω *are called* events*.*

$\mathcal{E} \subseteq 2^\Omega$ *is a collection of events such that* $\varnothing \notin \mathcal{E}$ *and* $\Omega \in \mathcal{E}$*.*

$f : \mathcal{E} \to 2^\Omega$ *is a function that associates with every event* $E \in \mathcal{E}$ *an event* $f(E)$ *satisfying the following properties: (1)* $f(E) \subseteq E$ *and (2)* $f(E) \neq \varnothing$*.*

In rational choice theory a set $E \in \mathcal{E}$ is interpreted as a set of available alternatives and $f(E)$ is interpreted as the subset of E which consists of the chosen alternatives (see, for example, [17] and [18]). In our case, we think of the elements of $\mathcal{E}$ as potential items of information and the interpretation of $f(E)$ is that, if informed that event E has occurred, the agent considers as doxastically possible all and only the states in $f(E)$.[6] The set $f(\Omega)$ is interpreted as the set of states that are *initially* considered doxastically possible (that is, before the receipt of information).

Note that in the rational choice literature (see, for example, [17]) it is common to impose some structure on the collection of events $\mathcal{E}$ (for example, that it be closed under finite unions). On the contrary, we allow $\mathcal{E}$ to be an arbitrary subset of 2^Ω and typically think of $\mathcal{E}$ as containing only a small number of events. This is characteristically the case in extensive-form games, as shown in the next section.

In order to interpret a choice frame $\langle \Omega, \mathcal{E}, f \rangle$ in terms of belief revision we need to add a *valuation* $V : S \to 2^\Omega$ that associates with every atomic formula $p \in S$ the set of states at which p is true. The quadruple $\langle \Omega, \mathcal{E}, f, V \rangle$

[6] In order to avoid ambiguity, we use the expression 'doxastically possible' to distinguish between possibility in terms of information (or "objective" possibility) and possibility in tems of beliefs (or "subjective" possibility or "doxastic" possibility). Thus a state ω may be possible according to the information received ($\omega \in E$) but may be ruled out by the agent's beliefs ($\omega \notin f(E)$); the doxastically possible states - when informed that E - are precisely those in $f(E)$. In a framework where beliefs are represented by a probability measure, a state is doxastically possible if and only if it is assigned positive probability.

is called a *model* (or an *interpretation*) *of* $\langle \Omega, \mathcal{E}, f \rangle$. Given a model $\mathcal{M} = \langle \Omega, \mathcal{E}, f, V \rangle$, truth of an arbitrary formula at a state is defined recursively as follows ($\omega \models_{\mathcal{M}} \phi$ means that formula ϕ is true at state ω in model $\mathcal{M}$): (1) for $p \in S$, $\omega \models_{\mathcal{M}} p$ if and only if $\omega \in V(p)$, (2) $\omega \models_{\mathcal{M}} \neg\phi$ if and only if $\omega \not\models_{\mathcal{M}} \phi$ and (3) $\omega \models_{\mathcal{M}} (\phi \vee \psi)$ if and only if either $\omega \models_{\mathcal{M}} \phi$ or $\omega \models_{\mathcal{M}} \psi$ (or both). The truth set of formula ϕ in model $\mathcal{M}$ is denoted by $\|\phi\|_{\mathcal{M}}$, that is, $\|\phi\|_{\mathcal{M}} = \{\omega \in \Omega : \omega \models_{\mathcal{M}} \phi\}$.

Given a model $\mathcal{M} = \langle \Omega, \mathcal{E}, f, V \rangle$ and formulas ϕ and ψ, we say that

- the agent *initially believes that* ψ if and only if $f(\Omega) \subseteq \|\psi\|_{\mathcal{M}}$,[7]
- the agent *believes that* ψ *upon learning that* ϕ if and only if (1) $\|\phi\|_{\mathcal{M}} \in \mathcal{E}$ and (2) $f(\|\phi\|_{\mathcal{M}}) \subseteq \|\psi\|_{\mathcal{M}}$.[8]

Accordingly, we can associate with every model $\mathcal{M}$ a (partial) belief revision function as follows. Let

$$\begin{aligned} &K_{\mathcal{M}} = \{\phi \in \Phi : f(\Omega) \subseteq \|\phi\|_{\mathcal{M}}\}, \\ &\Psi_{\mathcal{M}} = \{\phi \in \Phi : \|\phi\|_{\mathcal{M}} \in \mathcal{E}\}, \\ &B_{K_{\mathcal{M}}} : \Psi_{\mathcal{M}} \to 2^{\Phi} \text{ given by } B_{K_{\mathcal{M}}}(\phi) = \{\psi \in \Phi : f(\|\phi\|_{\mathcal{M}}) \subseteq \|\psi\|_{\mathcal{M}}\}. \end{aligned} \tag{1}$$

Definition 4 *A choice frame* $\langle \Omega, \mathcal{E}, f \rangle$ *is* AGM-consistent *if, for every model* $\mathcal{M} = \langle \Omega, \mathcal{E}, f, V \rangle$ *based on it, the (partial) belief revision function* $B_{K_{\mathcal{M}}}$ *associated with* $\mathcal{M}$ *(given by (1)) can be extended to a full belief revision function that satisfies the AGM postulates.*

Definition 5 *A choice frame* $\langle \Omega, \mathcal{E}, f \rangle$ *is* rationalizable *if there exists a total pre-order* $\precsim$ *on* Ω *such that, for every* $E \in \mathcal{E}$, $f(E) = \{\omega \in E : \omega \precsim \omega', \forall \omega' \in E\}$.

The interpretation of $\omega \precsim \omega'$ is that state ω is at least as plausible as state ω'. Thus in a rationalizable choice frame $\langle \Omega, \mathcal{E}, f \rangle$, for every $E \in \mathcal{E}$, $f(E)$ is the set of most plausible states in E. The following proposition, which establishes the equivalence of AGM-consistency and rationalizability, is proved in [4].

[7] That is, if ψ is true at every state initially considered doxastically possible.

[8] That is, if the truth set of ϕ is one of the potential items of information and ψ is true at every state considered doxastically possible given the information that ϕ is the case.

Proposition 6 *Let $\langle \Omega, \mathcal{E}, f \rangle$ be a choice frame where Ω is finite. Then $\langle \Omega, \mathcal{E}, f \rangle$ is AGM-consistent if and only if it is rationalizable.*

On the basis of Proposition 6, rationalizable choice frames can be viewed as providing a set-theoretic semantics for AGM belief revision. In the next section we use choice frames to model belief revision in extensive-form games.

3 Choice frames and belief revision in extensive-form games

Choice frames can be used in extensive-form games to represent, for every player, her initial beliefs as well as her disposition to change those beliefs when informed that it is her turn to move. We make use of the history-based definition of extensive-form game, which is reviewed in Appendix A. For simplicity, in this paper we focus on games without chance moves.

Given an extensive form, we can associate with every player $i \in N$ a choice frame $\langle \Omega, \mathcal{E}_i, f_i \rangle$ as follows: $\Omega = H$ (the set of histories), $E \in \mathcal{E}_i$ if and only if either $E = H$ or E consists of an information set of player i together with all the continuation histories, as explained below. If h is a decision history of player i ($h \in D_i$), player i's information set that contains h is denoted by $I_i(h)$. We shall denote by $\overrightarrow{I_i}(h)$ the set $I_i(h)$ together with the continuation histories:

$$\overrightarrow{I_i}(h) = \{x \in H : \exists h' \in I_i(h) \text{ such that } h' \text{ is a prefix of } x\}. \tag{2}$$

Thus

$$\mathcal{E}_i = \{H\} \cup \{\overrightarrow{I_i}(h) : h \in D_i\}. \tag{3}$$

We call $\overrightarrow{I_i}(h)$ the *augmented* information set of player i at decision history $h \in D_i$.[9] For example, in the extensive form of Figure 1, $\mathcal{E}_4 = \{H, E, F\}$,

[9]Because of the property of perfect recall (see Appendix A), for every player $i \in N$ and for every $h, h' \in D_i$, either $\overrightarrow{I_i}(h) \cap \overrightarrow{I_i}(h') = \varnothing$ or $\overrightarrow{I_i}(h) \subseteq \overrightarrow{I_i}(h')$ or $\overrightarrow{I_i}(h') \subseteq \overrightarrow{I_i}(h)$. That is, any two different augmented information sets of the same player are either disjoint or one is a subset of the other. Thus if $E, F \in \mathcal{E}_i$ are such that $E \cap F \neq \varnothing$, then either $E \subseteq F$ or $F \subseteq E$. Furthermore, if $h, h' \in D_i$ and h is a prefix of h', then $\overrightarrow{I_i}(h') \subseteq \overrightarrow{I_i}(h)$. Hence, during any play of the game, player i never receives contradictory information; in fact if information F follows information E then $F \subseteq E$, that is, F is a refinement of E.

where $E = \{acf, ade, acfg, acfh, adeg, adeh\}$ and $F = \{adf, b, adfm, adfn, bm, bn\}$. Finally, the function f_i provides initial beliefs as well as revised beliefs about past and future moves.

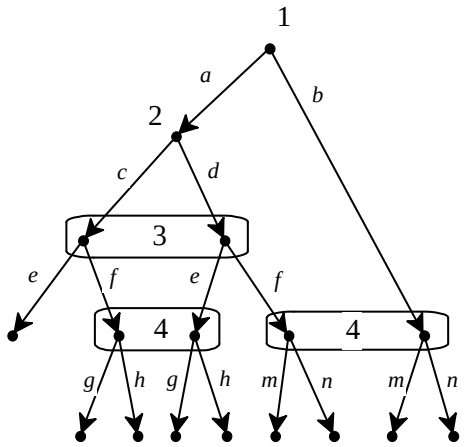

Figure 1

For example, in the extensive form of Figure 1 possible beliefs for Player 4 are as follows: $f_4(H) = \{a, ac, ace\}$, $f_4(E) = \{acf, acfh\}$ and $f_4(F) = \{b, bm\}$, where E and F are as given above. The interpretation of this is that Player 4 initially believes that Player 1 will play a, Player 2 will follow with c and Player 3 with e (so that Player 4 does not expect to be asked to make any choices; all this is encoded in $f_4(H)$). If informed that she is at her information set on the left, Player 4 would continue to believe that Player 1 played a and Player 2 followed with c, but she would now believe that Player 3 chose f and she herself plans to choose h (this is encoded in $f_4(E)$). On the other hand, if informed that she is at her information set on the right, Player 4 would believe that Player 1 played b and she herself plans to choose m (this is encoded in $f_4(F)$).

We shall make the following natural assumptions about the beliefs of the players. Let $\langle H, \mathcal{E}_i, f_i \rangle$ be the choice frame of player i representing the player's initial beliefs and disposition to change those beliefs. We assume that, for every $E \in \mathcal{E}_i$ and for every $h, h' \in H$,[10]

[10]Recall that D_i is the set of decision histories of player i and if $h \in D_i$ then $I_i(h)$ denotes the information set of player i that contains h.

If $h \in f_i(E)$ and $h' \in E$ is a prefix of h then $h' \in f_i(E)$. (A1)

If $h \in D_i \cap f_i(E)$ then $\exists a \in A(h)$ such that $ha \in f_i(E)$. (A2)

If $h \in D_i,\ h, ha \in f_i(E)$ and $h' \in I_i(h) \cap f_i(E)$ then $h'a \in f_i(E)$. (A3)

Assumption A1 says that the player's beliefs are closed under prefixes: if, when informed that event E has occurred, the player considers history h doxastically possible, and history h' is a prefix of h (that is, h' is a necessary condition for h to be reached) then she also considers history h' doxastically possible, as long as h' is compatible with the information received (that is, as long as $h' \in E$).

Assumption A2 says that if h is a decision history of player i, which she considers doxastically possible, then she also considers ha doxastically possible for some action a available at h. The interpretation of this is that the player has a belief, that is a plan, about how she would play at h.[11]

Finally, Assumption A3 states that the player's beliefs about her own choices are consistent in the sense that if she considers histories h and ha doxastically possible (where h is a decision history of hers and a an action available at h) and h' belongs to the same information set (of hers) as h, then she cannot consider h' doxastically possible without also considering $h'a$ doxastically possible.

Remark 7 *If we assume that player i's beliefs encoded in the choice frame $\langle H, \mathcal{E}_i, f_i \rangle$ are compatible with the AGM postulates for belief revision (see Definition 4), then, by Proposition 6, there exists a total pre-order $\precsim_i$ on H that rationalizes f_i, in the sense that, for every $E \in \mathcal{E}_i$, $f_i(E) = \{h \in E : h \precsim_i h', \forall h' \in E\}$).*[12]

[11] The view that "strategies as plans cannot be anything but beliefs of players about their own behavior" is also adopted in [3].

[12] It should be noted that the AGM theory deals with 'one-stage' belief revision, while in extensive-form games a player might receive information sequentially (when one of her information sets is preceded by another). Thus, in general, in extensive-form games one needs to consider what has been called in the literature 'iterated' belief revision. As noted in Footnote 9, because of the property of perfect recall, if a player receives two sequential pieces of information, E and F, then the latter is a refinement of the former (that is, $F \subseteq E$). In all the theories of iterated belief revision that have been proposed (see, for instance [6, 7, 13, 14]) it is postulated that when information E precedes information F

If $h \precsim_i h'$ player i judges history h to be at least as plausible as h'. An item of information $E \in \mathcal{E}_i$ lists all the histories that are still possible and $f_i(E)$ gives the histories that player i considers most plausible, given the information. Since, by definition of choice frame, $H \in \mathcal{E}_i$, the set $f_i(H)$ gives player i's *initial beliefs*, that is, her beliefs before the game is played, while for $E \in \mathcal{E}_i \backslash \{H\}$, $f_i(E)$ gives player i's *revised beliefs* if informed that E has occurred.

Remark 8 *An AGM-consistent choice frame $\langle H, \mathcal{E}_i, f_i \rangle$ of player i contains both (conditional) beliefs about the past and (conditional) beliefs about her own future choices. Given a decision history h of player i and the corresponding information set $I_i(h)$, player i's beliefs about past moves are given by the set $\{x \in I_i(h) : x \precsim_i y,\ \forall y \in I_i(h)\}$, that is, the most plausible histories in $I_i(h)$.*[13] *Furthermore, by Assumptions A2 and A3, for every $h \in D_i$ there exists at least one plausibility preserving action; the plausibility preserving actions at h represent the beliefs − and thus plans − of player i about her own choice at h.*[14]

Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i. The following property is known as *Arrow's Axiom* (see [18], p. 25): $\forall E, F \in \mathcal{E}_i$

$$\text{if } E \subseteq F \text{ and } f_i(F) \cap E \neq \varnothing \text{ then } f_i(E) = f_i(F) \cap E. \qquad (AA)$$

In order to simplify the proofs of the following propositions, we shall restrict attention to extensive forms that satisfy the following condition: $\forall i \in N, \forall h \in D, \forall a \in A(h)$,

$$\text{if } h \in D_i \text{ then } ha \notin D_i. \qquad (C)$$

and the latter is a refinement of the former, then the revised beliefs after the sequence $\langle E, F \rangle$ are the same as in the (possibly hypothetical) case where information F is received without it being preceded by E. Our analysis implicitly makes use of this assumption about iterated belief revision.

[13] Where $\precsim_i$ is the total pre-order that rationalizes the player's beliefs. Hence the i^{th} component μ_i of a "system of beliefs" μ would be such that, for every $h \in D_i$, $\mu_i(h) > 0$ if and only if $h \precsim_i y,\ \forall y \in I_i(h)$.

[14] Hence the i^{th} component σ_i of a strategy profile σ would be such that, for every $a \in A(h)$ (with $h \in D_i$), $\sigma_i(a) > 0$ if and only if $h \sim_i ha$. Thus implicit in an AGM-consistent choice frame $\langle H, \mathcal{E}_i, f_i \rangle$ of player i are the i^{th} components of an assessment (σ, μ).

Condition C rules out situations where two consecutive actions are taken by the same player. Thus if a player takes several actions in a sequence then between any two of them there is an action taken by another - possibly fictitious - player.[15]

The following proposition, which is proved in Appendix B, characterizes AGM-consistency of beliefs under Assumptions A1-A3.

Proposition 9 *Fix an extensive form that satisfies Condition C. Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame representing player i's initial beliefs and disposition to change those beliefs. Then the following are equivalent:*

(a) There is a total pre-order $\precsim_i$ that rationalizes $\langle H, \mathcal{E}_i, f_i \rangle$ (that is, $\forall E \in \mathcal{E}_i$, $f_i(E) = \{h \in E : h \precsim_i h', \forall h' \in E\}$) and satisfies Property PL1 of Definition 1 (that is, $\forall h \in D$, $\forall a \in A(h)$, $h \precsim_i ha$) as well as the following property:

PL2i. $\forall h \in D_i$, *(1)* $\exists a \in A(h)$ *such that* $ha \precsim_i h$ *and,*
(2) $\forall a \in A(h)$, *if* $ha \precsim_i h$ *then* $h'a \precsim_i h'$, $\forall h' \in I_i(h)$.

(b) $\langle H, \mathcal{E}_i, f_i \rangle$ satisfies Arrow's Axiom and Assumptions A1-A3.

Let $\{\langle H, \mathcal{E}_i, f_i \rangle\}_{i \in N}$ be a profile of choice frames representing the initial beliefs and disposition to revise those beliefs of all the players. Let $\mathcal{P}_i$ be the set of total pre-orders that rationalize $\langle H, \mathcal{E}_i, f_i \rangle$ and satisfy Properties PL1 and PL2i. Proposition 9 above gives necessary and sufficient conditions for $\mathcal{P}_i \neq \varnothing$.

Definition 10 *We say that the profile $\{\langle H, \mathcal{E}_i, f_i \rangle\}_{i \in N}$* admits a common prior *if $\bigcap_{i \in N} \mathcal{P}_i \neq \varnothing$, that is, if there exists a total pre-order $\precsim$ on H that rationalizes the beliefs of all the players*[16] *and satisfies Properties PL1 and PL2i for every $i \in N$. We call any element of $\bigcap_{i \in N} \mathcal{P}_i$ a* common prior.[17]

[15] If an extensive form does not satisfy Condition C then one can transform it into one that does, by adding a fictitious player between two consecutive actions of the same player and assigning to the fictitious player only one action. Such a trasformation would be "inessential" in the set that, for example, it would not affect the set of sequential equilibria.

[16] That is, $\forall i \in N$, $\forall E \in \mathcal{E}_i$, $f_i(E) = \{h \in E : h \precsim h', \ \forall h' \in E\}$.

[17] There may be several total pre-orders that play the role of a common prior, but they all yield the same conditional beliefs, given the possible items of information encoded in the extensive form.

If the players have a common prior then they share the same initial beliefs and the same disposition to change those beliefs in response to the same information. However, the existence of a common prior is consistent with the players holding different beliefs during any particular play of the game, since they will typically receive different information. A common prior can also be viewed as encoding the initial beliefs and belief revision policy of an external observer (the external-observer point of view is pursued in [10]).

The following proposition, which is proved in Appendix B, provides an epistemic justification for the notion of AGM-consistent assessment (Definition 2), used in the definition of perfect Bayesian equilibrium. The epistemic justification is based on compatibility of individual beliefs with the AGM theory of belief revision and the existence of a common prior.

Proposition 11 *Fix an extensive form that satisfies Condition C. Let* $\{\langle H, \mathcal{E}_i, f_i\rangle\}_{i\in N}$ *be a profile of AGM-consistent choice frames representing the initial beliefs and disposition to revise those beliefs of all the players. If* $\{\langle H, \mathcal{E}_i, f_i\rangle\}_{i\in N}$ *admits a common prior then every common prior is a plausibility order (see Definition 1).*

Let $\{\langle H, \mathcal{E}_i, f_i\rangle\}_{i\in N}$ be a profile of rationalizable choice frames that admits a common prior $\precsim$. By Proposition 11, $\precsim$ is a plausibility order. Corresponding to $\precsim$ there will be many AGM-consistent assessments (σ, μ), all of which share the same support (for σ the support is given by the plausibility-preserving actions and for μ the support is given by the most plausible histories in each information set). The definition of perfect Bayesian equilibrium put forward in [5] specifies a way in which the probabilities can be chosen on these supports so as to make μ compatible with σ and Bayes' rule.

4 Conclusion

As shown in [5], the qualitative notion of AGM-consistency of assessments is a generalization of the notion of consistency proposed by Kreps and Wilson [11] as part of the definition of sequential equilibrium. The conceptual content of the notion of Kreps-Wilson consistency is not clear and several attempts have been made to clarify it by relating it to more intuitive notions, such as 'structural consistency' ([12]), 'generally reasonable extended assessment' ([8]), 'stochastic independence' ([2, 10]).[18] In this paper we have provided

[18] Perea *et al* [16] offer an algebraic characterization of consistent assessments.

an independent justification for AGM consistency, based on the AGM theory of belief revision, thus providing an epistemic foundation for the notion of perfect Bayesian equilibrium proposed in [5]. A third paper will be devoted to studying various qualitative notions of independence and the relationship between perfect Bayesian equilibrium and sequential equilibrium.

A Appendix: Extensive forms and assessments

In this appendix we review the history-based definition of extensive-form game (see, for example, [15]). 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.[19] 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* without chance moves is a tuple $\left\langle A, H, N, P, \{\approx_i\}_{i \in N} \right\rangle$ 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. Let $D = H \setminus Z$ denote 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.
- A function $P : D \to N$ that assigns a player to each decision history; thus $P(h)$ is the player who moves at history h. For every $i \in N$, let $D_i = P^{-1}(i)$ be the histories assigned to player i. Thus $\{D_1, ..., D_n\}$ is a partition of D.

[19] In particular, every history is a prefix of itself.

- 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.

Given an extensive form, one obtains an *extensive-form 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).

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 $s_i(h') = s_i(h)$. 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 $\sigma_i(h') = \sigma_i(h)$. Note that a pure strategy is a special case of a behavior strategy where each probability distribution is degenerate. A *behavior-strategy profile* is an n-tuple $\sigma = (\sigma_1, ..., \sigma_n)$ where, for every $i \in N$, σ_i is a behavior strategy of player i. 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.

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 (1) if $h \in D_i$ then $\mu(h)$ is a probability distribution over $I_i(h)$ and (2) if $h \in D_i$ and $h' \in I_i(h)$ then $\mu(h) = \mu(h')$.

Without risking ambiguity we shall denote by $\mu(h)$ the probability assigned to history h by the system of beliefs μ.[20]

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

B Appendix: Proofs

The proof of Proposition 9 requires several preliminary results. The idea of the proof is to construct a binary relation on the set of histories H that satisfies Properties PL1 and PL2i and extend it to a total pre-order which is then shown to rationalize the given choice frame. The extension is obtained by invoking Proposition 13 below, which is known as Szpilrajn's theorem (for a proof see [18], p. 14). First we give the definition of extension. Given a binary relation R on H (thus $R \subseteq H \times H$) we shall interchangeably use the notation hRh' and $(h, h') \in R$.

Definition 12 *Let R be a binary relation on H and $\precsim$ a total pre-order on H. We say that $\precsim$* extends *R if (1) if $(h, h') \in R$ then $(h, h') \in \precsim$ and (2) if $(h, h') \in R$ and $(h', h) \notin R$ then $(h', h) \notin \precsim$.*

Proposition 13 *(Szpilrajn's theorem) Let R be a binary relation on H which is reflexive and transitive. Then there exists a total pre-order $\precsim$ on H which extends R.*

The following proposition is more general than Proposition 9 in that it applies to arbitrary extensive forms (that is, Condition C is *not* assumed), but it is weaker since it only refers to Property PL1 and Assumption A1. The proof illustrates the strategy used in proving Proposition 9.

Proposition 14 *Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i. The following are equivalent:*

(a) There is a total pre-order $\precsim$ on H that satisfies property PL1 of Definition 1 ($\forall h \in D$, $\forall a \in A(h)$, $h \precsim ha$) and rationalizes $\langle H, \mathcal{E}_i, f_i \rangle$ ($\forall E \in \mathcal{E}_i$, $f_i(E) = \{h \in E : h \precsim h', \forall h' \in E\}$),

(b) $\langle H, \mathcal{E}_i, f_i \rangle$ satisfies Arrow's Axiom and Assumption A1.

[20] 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 $\mu(h)$.

Proof. $(a) \Rightarrow (b)$. Let $\precsim$ be a total pre-order on H that satisfies property PL1 and is such that

$$\forall E \in \mathcal{E}_i, \quad f_i(E) = \{h \in E : h \precsim h', \forall h' \in E\}. \tag{4}$$

First we show that Arrow's Axiom (AA) holds. Let $F, G \in \mathcal{E}_i$ be such that $F \subseteq G$ and $f_i(G) \cap F \neq \varnothing$. We need to show that $f_i(F) = f_i(G) \cap F$. Fix an arbitrary $h \in f_i(G) \cap F$. By (4), $h \precsim h', \forall h' \in G$ and thus, since $F \subseteq G$, $h \precsim h', \forall h' \in F$. Hence, by (4) and the fact that $h \in F$, $h \in f_i(F)$. Conversely, fix an arbitrary $h \in f_i(F)$. Then, by (4),

$$h \precsim h', \forall h' \in F. \tag{5}$$

By hypothesis, $f_i(G) \cap F \neq \varnothing$. Fix an arbitrary $h_0 \in f_i(G) \cap F$. Since $h_0 \in f_i(G)$, by (4), $h_0 \precsim h', \forall h' \in G$. Since $h_0 \in F$, by (5) $h \precsim h_0$. Thus, by transitivity of $\precsim$, $h \precsim h', \forall h' \in G$, so that, by (4), $h \in f_i(G)$ (note that $h \in G$ since $h \in f_i(F) \subseteq F$ and $F \subseteq G$). Hence $h \in f_i(G) \cap F$.

Next we prove that Assumption A1 is satisfied. Fix arbitrary $E \in \mathcal{E}_i$ and $h \in f_i(E)$. Let $h' \in E$ be a prefix of h. We need to show that $h' \in f_i(E)$. By (4) (since $h \in f_i(E)$), $h \precsim y, \forall y \in E$. By Property PL1 and transitivity of $\precsim$, $h' \precsim h$.[21] Thus, by transitivity of $\precsim$, $h' \precsim y, \forall y \in E$, so that, by (4), $h' \in f_i(E)$.

$(b) \Rightarrow (a)$. Let $\langle H, \mathcal{E}_i, f_i \rangle$ satisfy Arrow's Axiom and Assumption A1. Define the following binary relation S on H:

$$(h, h') \in S \text{ if and only if } \begin{cases} \text{either} & \text{(a)} & h \text{ is a prefix of } h' \\ \text{or} & \text{(b)} & \exists h_1 \in H, \exists E \in \mathcal{E}_i : \\ & & h \text{ is a prefix of } h_1, \\ & & h_1 \in f_i(E) \text{ and } h' \in E. \end{cases} \tag{6}$$

First we show that S is reflexive and transitive. Reflexivity follows from (a) of (6) and the fact that, by definition of prefix, every history is a prefix of itself. To prove transitivity, fix arbitrary $h, h', h'' \in H$ and suppose that hSh' and $h'Sh''$. We need to show that hSh''. If h is a prefix of h' and h' is a prefix of h'', then h is a prefix of h'' and thus hSh''. If h is a prefix of h' while

[21] Since h' is a prefix of h, there exist $a_1, ..., a_m \in A$ $(m \geq 0)$ such that $h = h'a_1...a_m$. By PL1 $h' \precsim h'a_1 \precsim h'a_1a_2 \precsim ... \precsim h'a_1...a_m = h$. Thus, by transitivity of $\precsim$, $h' \precsim h$.

h' is not a prefix of h'', then $\exists h_1 \in H$, $\exists E \in \mathcal{E}_i$ such that h' is a prefix of h_1, $h_1 \in f_i(E)$ and $h'' \in E$. Then (since h is a prefix of h' and h' is a prefix of h_1) h is a prefix of h_1 and thus hSh'' by (b) of (6). If h is not a prefix of h' while h' is a prefix of h'', then $\exists h_1 \in H$, $\exists E \in \mathcal{E}_i$ such that h is a prefix of h_1, $h_1 \in f_i(E)$ and $h' \in E$. Then, since $h' \in E$ and h' is a prefix of h'', $h'' \in E$ (this follows from the definition of $\mathcal{E}_i$: see (2) and (3)). Thus hSh'' by (b) of (6). We are left with the case where h is not a prefix of h' and h' is not a prefix of h''. Then $\exists x_1, y_1 \in H$, $\exists E, F \in \mathcal{E}_i$ such that (i) h is a prefix of x_1, (ii) $x_1 \in f_i(E)$, (iii) $h' \in E$, (iv) h' is a prefix of y_1, (v) $y_1 \in f_i(F)$ and (vi) $h'' \in F$. By (iii) and (iv) $y_1 \in E$. Hence, by (v) (since $f_i(F) \subseteq F$), $E \cap F \neq \varnothing$ so that either $F \subseteq E$ or $E \subseteq F$ (see Footnote 9). Consider first the case where $F \subseteq E$. Then, since $h'' \in F$, we have that $h'' \in E$. By (b) of (6), it follows from this, (i) and (ii) that hSh''. Now consider the case where $E \subseteq F$. Since $y_1 \in f_i(F)$ and $y_1 \in E$, $f_i(F) \cap E \neq \varnothing$. Thus, by Arrow's Axiom, $f_i(E) = f_i(F) \cap E$. Hence, since $x_1 \in f_i(E)$, $x_1 \in f_i(F)$. Thus, since since h is a prefix of x_1, $x_1 \in f_i(F)$ and $h'' \in F$, by (b) of (6) hSh''.

Since S is reflexive and transitive, by Proposition 13, there exists a total pre-order $\precsim$ on H which extends S (see Definition 12). Fix an arbitrary such total pre-order $\precsim$. We want to show that $\precsim$ satisfies Property PL1 and rationalizes $\langle H, \mathcal{E}_i, f_i \rangle$. Since, for every $h \in D$ and $a \in A(h)$, h is a prefix of ha, $(h, ha) \in S$ and thus, since S is a subset of $\precsim$, $h \precsim ha$ so that $\precsim$ satisfies Property PL1. Now fix an arbitrary $E \in \mathcal{E}_i$. We need to show that $f_i(E) = \{h \in E : h \precsim h', \forall h' \in E\}$. Fix arbitrary $h \in f_i(E)$ and $h' \in E$. Then (since h is a prefix of itself) by (b) of (6) hSh' and thus, since S is a subset of $\precsim$, $h \precsim h'$. Hence $f_i(E) \subseteq \{h \in E : h \precsim h', \forall h' \in E\}$. For the converse, let $h \in E$ be such that $h \precsim h'$ for all $h' \in E$; we need to show that $h \in f_i(E)$. Fix an arbitrary $h_0 \in f_i(E)$ (recall that, by definition of choice frame, $f_i(E) \neq \varnothing$). If h is a prefix of h_0 then, by Assumption A1, $h \in f_i(E)$. Suppose that h is not a prefix of h_0. By definition of S, $(h_0, h) \in S$. If $(h, h_0) \notin S$, then, since $\precsim$ is an extension of S (see Definition 12), $(h, h_0) \notin \precsim$, contradicting our hypothesis that $h \precsim h', \forall h' \in E$. Thus it must be that $(h, h_0) \in S$. Then (since h is not a prefix of h_0) there exist $h_1 \in H$ and $F \in \mathcal{E}_i$ such that (i) h is a prefix of h_1, (ii) $h_1 \in f_i(F)$ and (iii) $h_0 \in F$. Then (since $h_0 \in F$ and $h_0 \in f_i(E) \subseteq E$), $E \cap F \neq \varnothing$ and thus (see Footnote 9) either $E \subseteq F$ or $F \subseteq E$. Suppose first that $E \subseteq F$. Since $h \in E$ and h is a prefix of h_1, $h_1 \in E$. Thus, since $h_1 \in f_i(F)$, $f_i(F) \cap E \neq \varnothing$ and, by Arrow's Axiom, $f_i(E) = f_i(F) \cap E$. Hence $h_1 \in f_i(E)$ and thus, by

Assumption A1 (since h is a prefix of h_1 and $h \in E$), $h \in f_i(E)$. Suppose now that $F \subseteq E$. Then, since $h_0 \in f_i(E) \cap F$, $f_i(E) \cap F \neq \varnothing$ and thus, by Arrow's Axiom, $f_i(F) = f_i(E) \cap F$. Thus, since $h_1 \in f_i(F)$, $h_1 \in f_i(E)$ and therefore, by Assumption A1 (since h is a prefix of h_1), $h \in f_i(E)$. ■

The proof of Proposition 9 follows the same strategy, starting from a relation that satisfies also Property PL2i. In order to do this we need several preliminary lemmas. Note that Condition C is used only in the proof of Lemma 20 and is not needed for any other result (except, of course, for Proposition 9, whose proof makes use of Lemma 20).

Lemma 15 *Fix an arbitrary choice frame $\langle H, \mathcal{E}_i, f_i \rangle$ of player i. Let $h \in D_i$ be a decision history of player i and let $F \in \mathcal{E}_i$ be such that $h \in F$. Then $F \supseteq E$, where $E = \overrightarrow{I_i}(h) \in \mathcal{E}_i$.*

Proof. Since $h \in F \in \mathcal{E}_i$, there exists an $x \in D_i$ such that x is a prefix of h and $F = \overrightarrow{I_i}(x)$ (see (2) and (3)). If $x = h$ then $F = E$. If $x \neq h$, then, by perfect recall, every $h' \in I_i(h)$ has a prefix in $I_i(x)$ and thus $E = \overrightarrow{I_i}(h) \subseteq \overrightarrow{I_i}(x) = F$. ■

Fix an abitrary choice frame $\langle H, \mathcal{E}_i, f_i \rangle$ of player i. Define the following binary relations on H:

$$(x, y) \in R_1 \text{ if and only if } \begin{cases} x \in D_i,\ y \in I_i(x) \text{ and} \\ x \in f_i(E) \text{ where } E = \overrightarrow{I_i}(x). \end{cases} \tag{7}$$

$$(x, y) \in R_2 \text{ if and only if } \begin{cases} y \in D_i,\ x = ya \text{ for some } a \in A(y) \\ \text{and } \exists h \in I_i(y) \text{ such that: } h, ha \in f_i(E) \\ \text{where } E = \overrightarrow{I_i}(y) \end{cases} \tag{8}$$

$$(x, y) \in R_3 \text{ if and only if } x \in f_i(H) \text{ and } y \text{ is a prefix of } x. \tag{9}$$

$$(x, y) \in R_4 \text{ if and only if } \begin{cases} y \in D_i,\ y \text{ is a prefix of } x \text{ and} \\ x \in f_i(E) \text{ where } E = \overrightarrow{I_i}(y). \end{cases} \tag{10}$$

$$(x, y) \in R_5 \text{ if and only if } x \text{ is a prefix of } y. \tag{11}$$

$$R = \bigcup_{j=1}^{5} R_j. \tag{12}$$

$$R^* \text{ transitive closure of } R. \tag{13}$$

Remark 16 *The relations R_1 and R_5 are transitive. Furthermore, R_5 is reflexive (since every history is a prefix of itself).*

Remark 17 *Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i. If $h \in D_i$ and $a \in A(h)$ are such that $(ha, h) \in R_2$ then $(h'a, h') \in R_2$, for every $h' \in I_i(h)$.*[22]

Lemma 18 *Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i that satisfies Arrow's Axiom and Assumption A1. Let $h \in D_i$ and $a \in A(h)$ be such that $(ha, h) \notin R_2$. Let $\langle x_1, ..., x_m \rangle$ $(m \geq 2)$ be a sequence in H such that $x_1 = ha$ and, $\forall j = 1, ..., m-1$, $(x_j, x_{j+1}) \in R$ (where R is give by (12)). Then, $\forall j = 1, ..., m$, $\exists h_j \in I_i(h)$ such that $h_j a$ is a prefix of x_j.*

Proof. This is clearly true for $j = 1$ (take $h_1 = h$). We now show that if the statement is true for $j \geq 1$ then it is true for $j + 1$. Let $h_j \in H$ be such that

$$h_j \in I_i(h) \text{ and } h_j a \text{ is a prefix of } x_j. \tag{14}$$

By hypothesis, $(x_j, x_{j+1}) \in R$. We need to consider all the possible cases.

Case 1: $(x_j, x_{j+1}) \in R_5$. Then x_j is a prefix of x_{j+1} and thus, since $h_j a$ is a prefix of x_j, $h_j a$ is a prefix of x_{j+1}.

Case 2: $(x_j, x_{j+1}) \in R_4$. Then $x_{j+1} \in D_i$, x_{j+1} is a prefix of x_j and

$$x_j \in f_i(F) \text{ where } F = \overrightarrow{I_i}(x_{j+1}). \tag{15}$$

Since both $h_j a$ and x_{j+1} are prefixes of x_j, either $h_j a$ is a prefix of x_{j+1} (with $x_{j+1} = h_j a$ as a special case), and thus the claim is true (take $h_{j+1} = h_j$), or x_{j+1} is a prefix of $h_j a$ and $x_{j+1} \neq h_j a$. Consider the latter case; then x_{j+1} is a prefix of h_j. Let $E = \overrightarrow{I_i}(h) = \overrightarrow{I_i}(h_j)$. Then, by perfect recall (since $x_{j+1} \in D_i$), $E \subseteq F$. Thus, since, by (14) and (15), $x_j \in f_i(F) \cap E$, by Arrow's Axiom $f_i(E) = f_i(F) \cap E$ so that $x_j \in f_i(E)$. Hence, by Assumption

[22] Proof: since $(ha, h) \in R_2$, $\exists h_0 \in I_i(h)$ such that $h_0, h_0 a \in f_i(E)$ where where $E = \overrightarrow{I_i}(h_0)$. Hence, by definition of R_2, $(h'a, h') \in R_2$, for every $h' \in I_i(h_0) = I_i(h)$.

A1, since $h_j a$ is a prefix of x_j, $h_j a \in f_i(E)$ and thus also $h_j \in f_i(E)$; but this implies, by definition of R_2 (see (8)), that $(ha, h) \in R_2$, contrary to our hypothesis. Thus if $(x_j, x_{j+1}) \in R_4$ then $h_j a$ is a prefix of x_{j+1}.

Case 3: we show that it cannot be that $(x_j, x_{j+1}) \in R_3$. In fact, $(x_j, x_{j+1}) \in R_3$ requires that $x_j \in f_i(H)$ so that, by Arrow's Axiom, $f_i(E) = f_i(H) \cap E$ (where $E = \overrightarrow{I_i}(h)$; note that $x_j \in E$). Hence $x_j \in f_i(E)$ and, by Assumption A1 (since $h_j a$ is a prefix of x_j), $h_j a \in f_i(E)$ an thus also $h_j \in f_i(E)$; but this implies, by definition of R_2, that $(ha, h) \in R_2$, contrary to our hypothesis.

Case 4: $(x_j, x_{j+1}) \in R_2$. Then, by definition of R_2, either (i) $x_{j+1} = h_j$ (if $x_j = h_j a$) or (ii) $x_{j+1} = h_j a b_1 ... b_{m-1}$ (if $x_j = h_j a b_1 ... b_m$ for some $b_1, ..., b_m \in A$, $m \geq 1$). In case (i), by definition of R_2, $\exists h_0 \in I_i(h_j)$ such that $h_0, h_0 a \in f_i(E)$ (where $E = \overrightarrow{I_i}(h_j)$). Since $I_i(h_j) = I_i(h)$, it would follow that $(ha, h) \in R_2$, contradicting our hypothesis. In case (ii) $h_j a$ is a prefix of x_{j+1}.

Case 5: $(x_j, x_{j+1}) \in R_1$. Then $x_j \in D_i$ and $x_{j+1} \in I_i(x_j)$. By perfect recall, since $h_j a$ is a prefix of x_j, $\exists h' \in I_i(h_j) = I_i(h)$ such that $h'a$ is a prefix of x_{j+1}. ■

Corollary 19 *Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i that satisfies Arrow's Axiom and Assumption A1. Let $h \in D_i$ and $a \in A(h)$. Then $(ha, h) \in R^*$ if and only if $(ha, h) \in R_2$.*

Proof. If $(ha, h) \in R_2$ then, since $R_2 \subseteq R \subseteq R^*$, $(ha, h) \in R^*$. To prove the converse, suppose that $(ha, h) \in R^*$. Then there exists a sequence $\langle x_1, ..., x_m \rangle$ ($m \geq 2$) in H such that $x_1 = ha$, $x_m = h$ and, $\forall j = 1, ..., m-1$, $(x_j, x_{j+1}) \in R$. If $(ha, h) \notin R_2$ then, by Lemma 18, $\forall j = 1, ..., m$, $\exists h_j \in I_i(h)$ such that $h_j a$ is prefix of x_j. In particular, $\exists h_m \in I_i(h)$ such that $h_m a$ is prefix of $x_m = h$, but this violates perfect recall. ■

Lemma 20 *Fix an extensive form that satisfies Condition C. Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i that satisfies Arrow's Axiom and Assumptions A1 and A3. Let $F \in \mathcal{E}_i$ and $x, y \in H$ be such that $x \in f_i(F)$ and $y \in F \backslash f_i(F)$. Then $(x, y) \in R^*$ and $(y, x) \notin R^*$ (where R^* is given by (13)).*

Proof. First we show that $(x, y) \in R^*$. If $F = H$ then $x \in f_i(H)$. Let $\emptyset$ denote the empty history (recall $\emptyset$ is a prefix of every history). Then $(x, \emptyset) \in R_3$ and $(\emptyset, y) \in R_5$. Thus $(x, y) \in R^*$. Consider now the case where

$F \neq H$. Then (since $x, y \in F$) there exist $x_0, y_0 \in D_i$ such that $F = \overrightarrow{I_i}(x_0)$, $y_0 \in I_i(x_0)$, x_0 is a prefix of x and y_0 is a prefix of y. Since $x \in f_i(F)$,

$$(x, x_0) \in R_4 \tag{16}$$

and, by Assumption A1, $x_0 \in f_i(F)$. Thus

$$(x_0, y_0) \in R_1. \tag{17}$$

Hence, since $(y_0, y) \in R_5$, it follows from (16) and (17) that $(x, y) \in R^*$.

Next we show that $(y, x) \notin R^*$. We will show that if $\langle x_1, ..., x_m \rangle$ $(m \geq 2)$ is a sequence in H with $x_1 \in F \backslash f_i(F)$ for some $F \in \mathcal{E}_i$, and, for all $j = 1, ..., m-1$, $(x_j, x_{j+1}) \in R$ (where R is defined in (12)) then $x_m \notin f_i(F)$. For this purpose it will be sufficient to prove the following: $\forall F \in \mathcal{E}_i$, $\forall h_1, h_2 \in H$

$$\text{if } h_1 \in F \backslash f_i(F) \text{ and } (h_1, h_2) \in R \text{ then } h_2 \in F \backslash f_i(F). \tag{18}$$

Let $F \in \mathcal{E}_i$, $h_1 \in F \backslash f_i(F)$ and $(h_1, h_2) \in R$. We need to consider all the possible cases.

Suppose that $(h_1, h_2) \in R_1$. Then $h_1 \in D_i$, $h_2 \in I_i(h_1)$ and $h_1 \in f_i(E)$ where $E = \overrightarrow{I_i}(h_1)$. Since $h_1 \in F$, by Lemma 15 $F \supseteq E$. Thus, since $h_2 \in E$, $h_2 \in F$. Suppose that $h_2 \in f_i(F)$. Then $h_2 \in f_i(F) \cap E$ and thus, by Arrow's Axiom, $f_i(E) = f_i(F) \cap E$, so that $h_1 \in f_i(F)$, contradicting our hypothesis. Hence $h_2 \in F \backslash f_i(F)$.

Suppose that $(h_1, h_2) \in R_2$. Then $h_2 \in D_i$ and $h_1 = h_2 a$ for some $a \in A(h_2)$ and

$$\exists h \in I_i(h_2) \text{ such that } h, ha \in f_i(E) \text{ where } E = \overrightarrow{I_i}(h_2). \tag{19}$$

Let $x \in H$ be the prefix of $h_2 a$ such that $F = \overrightarrow{I_i}(x)$. By Condition C (since $h_2 \in D_i$), $h_2 a \notin D_i$ and thus x is a prefix of h_2, so that $F \supseteq E$.[23] Thus $h_2 \in F$. If $h_2 \in f_i(F)$ then $f_i(F) \cap E \neq \varnothing$ and thus, by Arrow's Axiom, $f_i(E) = f_i(F) \cap E$, so that $h_2 \in f_i(E)$. It follows from this, (19) and Assumption A3 that $h_2 a \in f_i(E)$ and thus $h_2 a \in f_i(F)$, contradicting the hypothesis that $h_2 a = h_1 \in F \backslash f_i(F)$. Hence $h_2 \in F \backslash f_i(F)$.

[23] Without Condition C it is possible that $h_2 a \in D_i$ and that $F = \overrightarrow{I_i}(h_2 a)$, in which case $h_2 \notin F$.

Next we show that $(h_1, h_2) \notin R_3$. If $(h_1, h_3) \in R_3$ then $h_1 \in f_i(H)$ and thus (since $h_1 \in F$) $f_i(H) \cap F \neq \varnothing$ and by Arrow's Axiom $f_i(F) = f_i(H) \cap F$ so that $h_1 \in f_i(F)$, contradicting the hypothesis that $h_1 \in F \backslash f_i(F)$.

Suppose that $(h_1, h_2) \in R_4$. Then $h_2 \in D_i$, h_2 is a prefix of h_1 and $h_1 \in f_i(E)$ where $E = \overrightarrow{I_i}(h_2)$. Since $h_1 \in E \cap F$, $E \cap F \neq \varnothing$ and thus (see Footnote 9) either $E \subseteq F$ or $F \subseteq E$. It cannot be that $F \subseteq E$ because in this case (since $h_1 \in f_i(E) \cap F$) by Arrow's Axiom $f_i(F) = f_i(E) \cap F$ and thus $h_1 \in f_i(F)$, contradicting our hypothesis. Hence it must be $E \subseteq F$ so that, since $h_2 \in E$, $h_2 \in F$. Suppose that $h_2 \in f_i(F)$. Then $h_2 \in f_i(F) \cap E$ and thus, by Arrow's Axiom, $f_i(E) = f_i(F) \cap E$; hence, since $h_1 \in f_i(E)$, $h_1 \in f_i(F)$, contradicting our hypothesis. Hence $h_2 \in F \backslash f_i(F)$.

Suppose that $(h_1, h_2) \in R_5$. Then h_1 is a prefix of h_2 and thus, since $h_1 \in F$, $h_2 \in F$. If $h_2 \in f_i(F)$ then, by Assumption A1, $h_1 \in f_i(F)$, contradicting our hypothesis. Hence $h_2 \in F \backslash f_i(F)$. ■

Proof of Proposition 9. $(a) \Rightarrow (b)$ Let $\precsim_i$ be a total pre-order that rationalizes $\langle H, \mathcal{E}_i, f_i \rangle$ and satisfies Property PL1 of Definition 1 and Property PL2i. By Proposition 14, $\langle H, \mathcal{E}_i, f_i \rangle$ satisfies Arrow's Axiom and Assumption A1. We need to show that Assumptions A2 and A3 are also satisfied. Let $h \in D_i$ and $F \in \mathcal{E}_i$ and suppose that $h \in f_i(F)$. We want to show that $ha \in f_i(F)$ for some $a \in A(h)$. Let $E = \overrightarrow{I_i}(h)$. By Lemma 15, $F \supseteq E$. Thus, by Arrow's Axiom (since $h \in f_i(F) \cap E$), $f_i(E) = f_i(F) \cap E$. Hence $h \in f_i(E)$ and it will be enough to show that $ha \in f_i(E)$ for some $a \in A(h)$. Since $h \in f_i(E)$ and, by hypothesis, $f_i(E) = \{x \in E : x \precsim_i y, \forall y \in E\}$,

$$h \precsim_i y, \forall y \in E. \tag{20}$$

By (1) of Property PL2i there exists an $a \in A(h)$ such that $ha \precsim_i h$. Thus, by (20) and transitivity of $\precsim_i$, $ha \precsim_i y, \forall y \in E$ and thus $ha \in f_i(E)$. Thus Assumption A2 holds. To prove that Assumption A3 is satisfied, let $h \in D_i$, $a \in A(h)$ and $F \in \mathcal{E}_i$ be such that $h, ha \in f_i(F)$. Fix an arbitrary $h' \in I_i(h) \cap f_i(F)$. We need to show that $h'a \in f_i(F)$. Letting $E = \overrightarrow{I_i}(h)$, by the same argument used above we have that $f_i(E) = f_i(F) \cap E$, so that $h, ha \in f_i(E)$ and $h' \in I_i(h) \cap f_i(E)$ and it is thus sufficient to show that $h'a \in f_i(E)$. Since $ha \in f_i(E)$ and, by hypothesis, $f_i(E) = \{x \in E : x \precsim_i y, \forall y \in E\}$, $ha \precsim_i h$. Thus, by (2) of Property PL2i, $h'a \precsim_i h'$. Since $h' \in f_i(E)$, $h' \precsim_i y, \forall y \in E$. Thus, by transitivity of $\precsim_i$, $h'a \precsim_i y, \forall y \in E$ and therefore $h'a \in f_i(E)$.

$(b) \Rightarrow (a)$ Let $\langle H, \mathcal{E}_i, f_i \rangle$ be a choice frame of player i that satisfies Arrow's Axiom and Assumptions A1-A3. Let R^* be the relation defined in (13). Then R^* is transitive as well as reflexive (because R_5 is reflexive - see Remark 16 - and $R_5 \subseteq R^*$). Let $\succsim_i$ be a total pre-order that extends R^* (see Definition 12 and Proposition 13). Since $R_5 \subseteq R^* \subseteq \ \succsim_i$, $\succsim_i$ satisfies Property PL1. Next we show that $\succsim_i$ satisfies Property PL2i. Fix an arbitrary $h \in D_i$ and let $E = \overrightarrow{I_i}(h) \in \mathcal{E}_i$. By definition of choice frame, $f_i(E) \neq \varnothing$. Fix an arbitrary $x_0 \in f_i(E)$ and let $h_0 \in I_i(h)$ be the prefix of x_0 in $I_i(h)$. Then, by Assumption A1, $h_0 \in f_i(E)$. Thus, by Assumption A2, there exists an $a \in A(h_0) = A(h)$ such that $h_0 a \in f_i(E)$. Hence, by (8), $(ha, h) \in R_2$ and therefore (since R_2 is a subset of $\succsim_i$) $ha \succsim_i h$. Thus we have proved part (1) of Property PL2i. To prove part (2) of Property PL2i, fix an arbitrary $h \in D_i$ and an arbitrary $a \in A(h)$ and suppose that $ha \succsim_i h$. We have to show that $h'a \succsim_i h'$ for all $h' \in I_i(h)$. Since $(h, ha) \in R_5 \subseteq R^*$ if $(ha, h) \notin R^*$ then, by definition of extension (see Definition 12) $ha \not\precsim_i h$, contradicting our supposition. Thus $(ha, h) \in R^*$. Hence, by Corollary 19, $(ha, h) \in R_2$ and thus (see Remark 17) $(h'a, h') \in R_2$, for all $h' \in I_i(h)$. Since $R_2 \subseteq R^* \subseteq \ \succsim_i$, $h'a \succsim_i h'$ for all $h' \in I_i(h)$.

It remains to show that $\succsim_i$ rationalizes $\langle H, \mathcal{E}_i, f_i \rangle$. Fix an arbitrary $E \in \mathcal{E}_i$, $h \in f_i(E)$ and $h' \in E$. Then $(h, h') \in R^*$.[24] Thus $f_i(E) \subseteq \{h \in E : hR^*h', \forall h' \in E\}$ so that, since R^* is a subset of $\succsim$, $f_i(E) \subseteq \{h \in E : h \succsim h', \forall h' \in E\}$. Conversely, let $h \in E$ be such that $h \succsim h', \forall h' \in E$. We need to show that $h \in f_i(E)$. Fix an arbitrary $h_0 \in f_i(E)$. Suppose that $h \notin f_i(E)$. Then, by Lemma 20, $(h_0, h) \in R^*$ and $(h, h_0) \notin R^*$. Thus, since $\succsim_i$ is an extension of R^* (see Definition 12), $(h, h_0) \notin \ \succsim_i$, contradicting our hypothesis that $h \succsim h', \forall h' \in E$.

Proof of Proposition 11. Let $\succsim \ \in \bigcap_{i \in N} \mathcal{P}_i$. By Proposition 9, for every $i \in N$, every element of $\mathcal{P}_i$ satisfies Property PL1. Thus $\succsim$ satisfies PL1. Now fix an arbitrary decision history h and let i be the player to whom it belongs. By Property PL2i of Proposition 9, $\exists a \in A(h)$ such that $ha \succsim h$ and, $\forall a \in A(h)$, if $ha \succsim h$ then $h'a \succsim h'$, $\forall h' \in I_i(h)$. Thus $\succsim$ satisfies also Property PL2. Hence $\succsim$ is a plausibility order. ■

[24] The argument is the same as in the first part of the proof of Lemma 20: if $E = H$ then $(h, \emptyset) \in R_3$ and $(\emptyset, h') \in R_5$; if $E \neq H$ then, $(h, x_0) \in R_4$, $(x_0, y_0) \in R_1$ and $(y_0, h') \in R_5$, where $x_0, y_0 \in D_i$ are such that $E = \overrightarrow{I_i}(x_0)$, $y_0 \in I_i(x_0)$, x_0 is a prefix of h and y_0 is a prefix of h'.

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

[3] Battigalli, Pierpaolo, Alfredo Di Tillio and Dov Samet, Strategies and interactive beliefs in dynamic games, Working Paper, Bocconi University, November 2010.

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

[5] Bonanno, Giacomo, AGM-consistency and Perfect Bayesian equilibrium. Part I: definition and properties, Working Paper 10-17, University of California Davis, 2010.

[6] Boutilier, Craig, Iterated revision and minimal change of conditional beliefs, *Journal of Philosophical Logic*, 1996, 25: 263-305.

[7] Darwiche, Adnan and Judea Pearl, On the logic of iterated belief revision, *Artificial Intelligence,* 1997, 89: 1-29.

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

[9] Gärdenfors, Peter, *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press, 1988.

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

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

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

[13] Jin, Yi and Michael Thielscher, Iterated belief revision, revised, *Artificial Intelligence*, 2007, 171: 1-18.

[14] Nayak, Abhaya, Maurice Pagnucco and Pavlos Peppas, Dynamic belief revision operators, *Artificial Intelligence*, 2003, 146: 193-228.

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

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

[17] Rott, Hans, *Change, choice and inference*, Clarendon Press, 2001.

[18] Suzumura, Kotaro, *Rational choice, collective decisions and social welfare*, Cambridge University Press, 1983.