A characterization of sequential equilibrium in terms of AGM belief revision

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

Extended Abstract

1 Introduction

In economics and applied game theory, the most widely used solution concept for extensive-form games is sequential equilibrium, introduced by Kreps and Wilson [8]. Given an extensive-form game, an *assessment* is a pair (σ, μ), where σ is a behavior strategy profile (which gives, for every information set, a probability distribution over the choices at that information set) and μ is a system of beliefs (which gives, for every information set, a probability distribution over the nodes that constitute that information set). A *sequential equilibrium* is an assessment which satisfies two properties: sequential rationality and consistency. Sequential rationality requires that, at each information set, the strategy of the player who moves there be optimal given the player's beliefs (as captured by the relevant part of μ) and the strategies of the other players. While sequential rationality has a clear conceptual content, the notion of consistency is purely technical. An assessment (σ, μ) is consistent if there is an infinite sequence $\langle \sigma^1, ..., \sigma^m, ... \rangle$ of completely mixed strategy profiles (that is, every choice is assigned positive probability) such that, letting μ^m be the unique system of beliefs derived from σ^m by using Bayes' rule, $\lim_{m \to \infty}(\sigma^m, \mu^m) = (\sigma, \mu)$. Kreps and Wilson proposed the notion of consistent assessment as an attempt to capture the concept of "minimal" belief revision. A number of authors have tried to shed light on the technical notion of consistent assessment by relating it to more intuitive concepts, such as "structural consistency" ([9]), "generally reasonable extended assessment" ([6]), "stochastic independence" ([2], [7]).[1]

In this paper we provide a purely qualitative characterization of consistent assessments in terms of the AGM theory of belief revision [1], through the notion of AGM-consistent choice frame. For the purpose of this extended abstract we introduce several simplifying assumptions. The general case and the proofs of all the results can be found in the full paper [3].

[1] Perea et al [11] offer an algebraic characterization of consistent assessments

2 Choice frames and AGM-consistent beliefs

We start with a brief review of the notion of rationalizable choice frame and its relationship to the AGM theory of belief revision.

Definition 1 *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. In our case, we think of the elements of $\mathcal{E}$ as possible items of information and the interpretation of $f(E)$ is that, if informed that event E has occurred, the agent considers as possible all and only the states in $f(E)$. The set $f(\Omega)$ is interpreted as the states that are initially considered possible.

Note that in the literature (see, for example [12]) 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 typically the case in extensive-form games, as shown in the following 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 $p \in S$ (in a given propositional language) the set of states at which p is true. The quadruple $\langle \Omega, \mathcal{E}, f, V \rangle$ 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 $\mathcal{M} = \langle \Omega, \mathcal{E}, f, V \rangle$ we say that

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

Accordingly, we can associate with every model $\mathcal{M}$ a (partial) belief revision function as follows (to simplify the notation we will drop the subscript $\mathcal{M}$). Let

$$\begin{array}{l} K = \{\phi \in \Phi : f(\Omega) \subseteq \|\phi\|\}\,, \\ \Psi = \{\phi \in \Phi : \|\phi\| \in \mathcal{E}\}\,, \\ B_K : \Psi \to 2^{\Phi} \text{ given by } B_K(\phi) = \{\psi \in \Phi : f(\|\phi\|) \subseteq \|\psi\|\}\,. \end{array} \quad (1)$$

Thus K is the initial belief set and, for every item of information $\phi \in \Psi$, $B_K(\phi)$ is the revised belief set.

Definition 2 *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 *associated with* $\mathcal{M}$ *(see (1)) can be extended to a full-domain belief revision function that satisfies the AGM postulates.*[2]

Recall that a binary relation $\precsim$ on Ω is a *total pre-order* if it is complete ($\forall \omega, \omega' \in \Omega$, either $\omega \precsim \omega'$ or $\omega' \precsim \omega$) and transitive ($\forall \omega, \omega', \omega'' \in \Omega$, if $\omega \precsim \omega'$ and $\omega' \precsim \omega''$ then $\omega \precsim \omega''$). Given a total pre-order $\precsim$ of Ω and an event $E \subseteq \Omega$, let $Min_{\precsim}\ E = \{\omega \in E : \omega \precsim \omega', \forall \omega' \in E\}$.

Definition 3 *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* $E \in \mathcal{E}$, $f(E) = Min_{\precsim}\ E$.

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 is proved in [5]:

Proposition 4 *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 4, rationalizable choice frames can be viewed as providing a semantics for one-stage partial belief revision functions that obey the AGM postulates. In the next section we use choice frames to analyze extensive-form games.

3 Choice frames in extensive-form games

We adopt the history-based definition of extensive-form game (see, for example, [10]). For the purpose of this extended abstract we restrict attention to games without chance moves. 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*[3] is a tuple $\left\langle A, H, N, P, \{\approx_i\}_{i \in N} \right\rangle$ where:

- A is a finite set of actions.
- $H \subseteq A^*$ is a finite set of histories 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 empty history $\langle\ \rangle$ is denoted by $\emptyset$ and is an element of H. 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

[2] Because of space limitations we shall not review the AGM postulates which define the class of full AGM revision functions. See [3] for details.

[3] 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 and Z the set of terminal histories).

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 $h \in D$.

- $N = \{1, ...n\}$ is a set of players.
- $P : D \to N$ is a function that assigns a player to each non-terminal 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 set of histories assigned to player i. Thus $\{D_1, ..., D_n\}$ is a partition of D.
- For each 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 satisfied: 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.

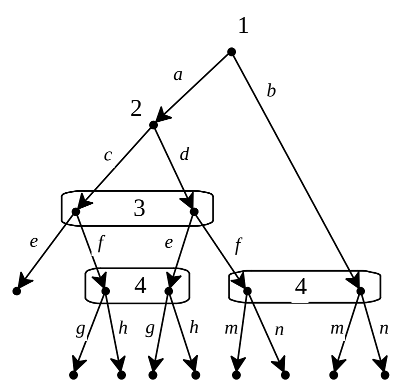

Figure 1

Figure 1 shows an extensive form where $A = \{a, b, c, d, e, f, g, h, m, n\}$, $H = D \cup Z$ with (to simplify the notation we write a instead of $\langle a \rangle$, ac instead of $\langle a, c \rangle$, etc.) $D = \{\emptyset, a, b, ac, ad, acf, ade, adf\}$, $Z = \{ace, acfg, acfh, adeg, adeh, adfm, adfn, bm, bn\}$, $A(\emptyset) = \{a, b\}$, $A(a) = \{c, d\}$, $A(ac) = A(ad) = \{e, f\}$, $A(acf) = A(ade) = \{g, h\}$, $A(adf) = A(b) = \{m, n\}$, $N = \{1, 2, 3, 4\}$, $P(\emptyset) = 1$, $P(a) = 2$, $P(ac) = P(ad) = 3$, $P(acf) = P(ade) = P(adf) = P(b) = 4$, $\approx_1 = \{(\emptyset, \emptyset)\}$, $\approx_2 = \{(a, a)\}$, $\approx_3 = \{(ac, ac), (ac, ad), (ad, ac), (ad, ad)\}$ 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 are shown as rounded rectangles. Thus, for example, $I_4(b) = \{adf, b\}$. The root of the tree represents the empty history $\emptyset$.

For the purpose of this extended abstract we restrict attention to the class of extensive forms where no player moves more than once along any history. That is, for every history h, if h_1 and h_2 are prefixes of h with $P(h_1) = P(h_2)$ then $h_1 = h_2$ (recall that $P(h)$ is the player who moves at h). The extensive form represented in Figure 1 satisfies this property.

Choice frames can be used to represent, for every player, her initial beliefs and her disposition to change those beliefs when it is her turn to move. Given an extensive form, we can associate with every $i \in N$ 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. Recall that, if $h \in D_i$, player i's information set that contains h is denoted by $I_i(h)$; that is, $I_i(h) = \{h' \in H : h' \approx_i h\}$. We shall denote by $\vec{I_i}(h)$ the set $I_i(h)$ together with the continuation histories: for $h \in D_i$,

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

Thus

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

For example, in the extensive form of Figure 1, $\mathcal{E}_4 = \{H, E_1, E_2\}$, where $E_1 = \{acf, ade, acfg, acfh, adeg, adeh\}$ and $E_2 = \{adf, b, adfm, adfn, bm, bn\}$.

Finally, the function f_i provides conditional beliefs about past and future moves. For example, in the extensive form of Figure 1 one possibility for Player 4 is: $f_4(H) = \{a, ac, ace\}$, $f_4(E_1) = \{acf, acfh\}$ and $f_4(E_2) = \{b, bm\}$, where E_1 and E_2 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). If informed that she is at her information set on the left then she 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. If informed that she is at her information set on the right then she would believe that Player 1 played b and she herself plans to choose m.

If we assume that the choice frame of player i is AGM consistent, then, by Proposition 4, there exists a total pre-order $\precsim_i$ on H that rationalizes f_i (that is, for every $E \in \mathcal{E}_i$, $f_i(E) = Min_{\precsim_i} E$).

What are natural properties to impose on these total pre-orders, that is, on the associated beliefs? We introduce four properties and show that they characterize the notion of consistent assessment.

The first property expresses the notion of *agreement of beliefs*, in the sense that the players share the same initial beliefs and the same disposition to change those beliefs in response to the same information:[4]

[4] This property can be viewed as an expression of the notion of a "common prior" (see, for example, [4]), which is pervasive in game theory.

$$\exists \precsim \ \subseteq H \times H: \ \ \forall i \in N, \quad \precsim_i \ = \ \precsim . \tag{$P1$}$$

Note that $P1$ is consistent with the players holding different beliefs during any particular play of the game, since they will typically receive different information.

The remaining properties will be stated in terms of the common pre-order $\precsim$ given by $P1$. Recall that the interpretation of $h \precsim h'$ is that history h is at least as plausible as history h'. 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 $h \prec h'$ (with the interpretation that h is more plausible than h') as a short-hand for "$h \precsim h'$ and $h' \not\precsim h$".

The second property says that adding an action to a history h cannot yield a more plausible history than h itself:

$$\forall h \in D, \quad \forall a \in A(h), \quad h \precsim ha. \tag{$P2$}$$

The third property 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; furthermore, any such action a performs the same role with any other history that belongs to the same information set:

$$\begin{array}{c} \forall i \in N, \forall h \in D_i, \ (1) \ \exists a \in A(h): \ ha \precsim h \text{ and} \\ (2) \ \forall a \in A(h), \ \forall h' \in I_i(h), \text{ if } ha \precsim h \text{ then } h'a \precsim h'. \end{array} \tag{$P3$}$$

Remark 5 *It follows from Properties $P2$ and $P3$ that, for every decision history h, there is at least one action a at h such that, for every h' in the same information set as h, $h'a$ is as plausible as h'. We call such actions* plausibility preserving.

A function $F : H \to \mathbb{N}$ (where $\mathbb{N}$ denotes the set of natural numbers) is an *integer-valued representation* of $\precsim$ if $F(\emptyset) = 0$ and, $\forall h, h' \in H$, $h \precsim h'$ if and only if $F(h) \leq F(h')$. Let $\mathcal{R}$ be the set of integer-valued representations of $\precsim$. Since H is finite, $\mathcal{R} \neq \varnothing$. We call an integer-valued representation F of $\precsim$ *action-consistent* if,

$$\begin{array}{l} \forall i \in N, \ \forall h, h' \in D_i, \forall a \in A(h), \text{ if } h' \in I_i(h) \text{ then} \\ F(ha) - F(h) = F(h'a) - F(h'). \end{array} \tag{4}$$

For example, consider the extensive form represented in Figure 2 and the following total pre-order: $\emptyset \sim a \prec b \sim be \prec bf \prec c \sim ce \prec d \prec cf$. The first column in Table 1 reproduces this total pre-order with the convention that if x and y are on the same line, then $x \sim y$ and if x is above y then $x \prec y$; the second and third columns give two integer-valued representations of $\precsim$, F_1 and F_2. F_1 is not action-consistent, since $c \in I(b)$ 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 is action-consistent.

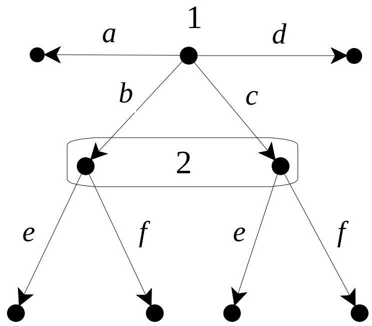

Figure 2

$\precsim$	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

Table 1

The fourth (and last) property says that among the integer-valued representations of $\precsim$ there is at least one which is action-consistent:

$$\text{There exists an } F \in \mathcal{R} \text{ which is action-consistent.} \qquad (P4)$$

Note that Properties $P2$-$P4$ are independent of each other.

If the beliefs of player i are rationalized by a total pre-order $\precsim$ on H, then the following holds: if the play of the game reaches history $h \in D_i$ then player i receives information $\vec{I}_i(h)$ and revises her previous beliefs to $f_i(\vec{I}_i(h)) = Min_{\precsim}\ \vec{I}_i(h)$, that is, the histories that are most plausible given her information constitute her revised beliefs.

Before we proceed to our main result, we recall in more detail the notion of sequential equilibrium ([8]). Given an extensive form, a *pure strategy* of player $i \in N$ is a function that associates with every information set of player i a choice 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 $s_4(acf) = s_4(ade) = g$, $s_4(adf) = s_4(b) = m$. 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 $\sigma_i(h') = \sigma_i(h)$. We denote by $\sigma_i(h)(a)$ the probability assigned to $a \in A(h)$ by $\sigma_i(h)$. Note that a pure strategy is a special case of a behavior strategy where each probability distribution is degenerate. A behavior strategy σ_i of player i is *completely mixed* if, for every $h \in D_i$ and for every $a \in A(h)$, $\sigma_i(h)(a) > 0$. 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.

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 $\mu(h) = \mu(h')$. Note that a completely mixed behavior strategy profile yields, using Bayes' rule, a unique system of beliefs.

An *assessment* is a pair (σ, μ) where σ is a behavior strategy profile and μ is a system of beliefs. An assessment (σ, μ) is *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 using Bayes' rule, $\lim_{m \to \infty}(\sigma^m, \mu^m) = (\sigma, \mu)$.

Kreps and Wilson [8] proposed the notion of consistent assessment as an attempt to capture the concept of "minimal" belief revision. The following proposition provides a characterization of consistent assessments in terms of the AGM theory of belief revision, through the notion of AGM-consistency of choice frames. Sequential rationality is discussed in Section 5.

Given an extensive form, we say that a profile $\{\langle \Omega, \mathcal{E}_i, f_i \rangle\}_{i \in N}$ of AGM-consistent choice frames (where $\Omega = H$ and $\mathcal{E}_i$ is given by (3)) satisfies properties $P1$-$P4$ if the collection of total pre-orders $\{\precsim_i\}_{i \in N}$ that rationalize $\{\langle \Omega, \mathcal{E}_i, f_i \rangle\}_{i \in N}$ (whose existence is guaranteed by Proposition 4) satisfies properties $P1$-$P4$ (that is, there exists a common total pre-order $\precsim$ on H that rationalizes those choice frames and satisfies properties $P2 - P4$).

Proposition 6 *Fix an extensive form. Then*

(a) If the players' initial beliefs and disposition to revise those beliefs are represented by a profile of AGM-consistent choice frames that satisfies properties $P1$-$P4$ *then there exists a consistent assessment* (σ, μ) *such that (letting* $\precsim$ *be a total pre-order on* H *that rationalizes those choice frames), for all* $i \in N$, $h \in D_i$ *and* $a \in A(h)$, *(1)* $\sigma_i(h)(a) > 0$ *if and only if* $h \sim ha$ *and (2)* $\mu(h) > 0$ *if and only if* $h \in Min_{\precsim} \vec{I}_i(h)$;[5]

(b) if (σ, μ) *is a consistent assessment then there exists a profile of AGM-consistent choice frames that satisfies properties* $P1$-$P4$ *such that (letting* $\precsim$ *be a total pre-order on* H *that rationalizes those choice frames), for every* $i \in N$, $h \in D_i$ *and* $a \in A(h)$, *(1)* $h \sim ha$ *if and only if* $\sigma_i(h)(a) > 0$ *and (2)* $h \in Min_{\precsim} \vec{I}_i(h)$ *if and only if* $\mu(h) > 0$.

4 General extensive games and iterated belief revision

In an arbitrary extensive form there may be players who move more than once along some histories. If i is such a player, then the set $\mathcal{E}_i$ defined above (see (3)) will contain two sets E and F such that, along some history, player i receives first information E and then, at a later moment, information F. Because of the property of perfect recall, in such a situation it will the case that $F \subseteq E$ (for every player $i \in N$ and for every $h, h' \in D_i$, if h is a prefix of h' then

[5] Recall - see (2) - that $\vec{I}_i(h)$ is the information set that contains h together with the continuation histories and that if $\langle \Omega, \mathcal{E}_i, f_i \rangle$ is the choice frame of player i then $\Omega = H$ and $\mathcal{E}_i = \{H\} \cup \{\vec{I}_i(h) : h \in D_i\}$ (see (3)).

$\vec{I}_i(h') \subseteq \vec{I}_i(h)$). We call this property *information refinement*. Because of the possibility of sequential informational inputs, we are outside the scope of one-shot belief revision and it is no longer sufficient to appeal to AGM consistency in order to guarantee the existence of a total pre-order that rationalizes the beliefs of a player. In the full version of the paper [3] it is argued that, within the context of information refinement, rationalizability of the choice frame of each player captures a basic principle of belief revision that is common to the may theories of iterated belief revision that have been proposed in the literature. The full paper contains the extension of Proposition 6 to arbitrary extensive-form games.

5 Sequential rationality, pure sequential equilibria, backward induction

A sequential equilibrium is an assessment (σ, μ) which is consistent and sequentially rational. Sequential rationality requires that - at each information set - the strategy of each player be optimal starting from there according to the player's beliefs over the nodes in that information set (as captured by the relevant part of μ) and the strategies of the other players. Conceptually, little is gained by expressing sequential rationality in terms of the total pre-order underlying the consistent assessment (σ, μ). However, there is one case where sequential rationality can be expressed very simply and that is the case where the restriction of the total pre-order $\precsim$ to the set Z of terminal histories is antisymmetric:

$$\text{if } z, z' \in Z, \quad \text{and} \quad z \sim z' \quad \text{then} \quad z = z'. \tag{$P5$}$$

Lemma 7 *Let $\precsim$ be a total pre-order on H that satisfies Properties $P2$, $P3$ and $P5$. Then, for every history $h \in H$, there is a unique terminal history z such that $h \sim z$. Call this terminal history $z(h)$ (if $h \in Z$ then $z(h) = h$). Furthermore, for every decision history $h \in D$, (a) there is a unique action $a \in A(h)$ such that $h \sim ha$ and (b) for all $h' \in \vec{I}_i(h)$, if $h \sim h'$ then h is a prefix of h'.*

Under the hypotheses of Lemma 7, sequential rationality can be expressed as follows (recall - see Footnote 3 - that, for every player $i \in N$, $U_i : Z \to \mathbb{R}$ is i's payoff function):

$$\forall i \in N, \forall h \in D_i, \forall a \in A(h), \quad U_i(z(h)) \geq U_i(z(ha)). \tag{$P6$}$$

Call a sequential equilibrium (σ, μ) *pure* if the strategy σ_i of each player $i \in N$ is a pure strategy and μ consists of degenerate probability distributions (that is, if $h \in D_i$ and $h' \in I_i(h)$ then either $\mu(h)(h') = 0$ or $\mu(h)(h') = 1$).

Proposition 8 *Fix an extensive-form game without chance moves. Then,*

(a) If the players' beliefs and belief revision policies are represented by a profile of rationalizable choice frames that satisfies properties $P1$-$P6$ then the

assessment (σ, μ) *given by (letting* $\precsim$ *be a total pre-order on* H *that rationalizes those choice frames), for all* $i \in N$, $h \in D_i$ *and* $a \in A(h)$, *(1)* $\sigma_i(h)(a) > 0$ *if and only if* $h \sim ha$ *and (2)* $\mu(h) > 0$ *if and only if* $h \in Min_{\precsim} \vec{I}_i(h)$, *is a pure sequential equilibrium.*

(b) if (σ, μ) *is a pure sequential equilibrium then there exists a profile of rationalizable choice frames that satisfies properties* $P1$*-*$P6$ *such that (letting* $\precsim$ *be a total pre-order on* H *that rationalizes those choice frames), for every* $i \in N$, $h \in D_i$ *and* $a \in A(h)$, *(1)* $h \sim ha$ *if and only if* $\sigma_i(h)(a) = 1$ *and (2)* $h \in Min_{\precsim} \vec{I}_i(h)$ *if and only if* $\mu(h) = 1$.

Note that, since we ruled out chance moves, Proposition 8 does not require the payoff function U_i of player i to satisfy the von Neumann-Morgenstern axioms of expected utility; indeed, it could be an ordinal payoff function (that is, a numerical representation of a total pre-order over Z expressing player i's preferences over the elements of Z).

An extensive form has *perfect information* if every information set is a singleton. The solution concept that is most commonly used for perfect-information games is that of *backward induction* (for a review of the backward-induction algorithm see the full paper [3]). In perfect-information games Property $P4$ is trivially satisfied, since $h' \in I_i(h)$ implies that $h' = h$. We now show that, for every perfect-information game, there is a one-to-one correspondence between the set of backward-induction solutions and the set of total pre-orders on H that satisfy properties $P2$, $P3$ and the following property, which is a generalization of $P6$. First of all some notation. Fix a total pre-order $\precsim$ on H. For every decision history $h \in D$, let $A_0(h) = \{a \in A(h) : h \sim ha\}$ and let $Z(h) = \{z \in Z : z \sim ha$ for some $a \in A_0(h)\}$. If $\precsim$ satisfies Properties $P2$ and $P3$ then $A_0(h) \neq \varnothing$ and $Z(h) \neq \varnothing$. We can now introduce the generalization of $P6$:

$$\forall i \in N,\ \forall h \in D_i,\ \forall z \in Z(h),\ \forall a \in A(h),\ \forall z' \in Z(ha),\quad U_i(z) \geq U_i(z'). \quad (P7)$$

Property $P7$ says that if h is a decision history of player i then the utility of any terminal history reached from h by following only plausibility preserving actions is not less than the utility of a terminal node reached by taking an arbitrary action a at h and then continuing from ha by following only plausibility preserving actions.

Proposition 9 *Fix an arbitrary finite perfect-information game. There is a one-to-one correspondence between the set of backward-induction solutions and the set of total pre-orders on* H *that satisfy properties* $P2$, $P3$ *and* $P7$.

Proposition 9 thus provides a characterization of backward induction in terms of beliefs and belief revision policies that are represented by profiles of rationalizable choice frames that satisfy Properties $P1$-$P3$ and $P7$.

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] Bonanno, Giacomo, A characterization of sequential equilibrium in terms of AGM belief revision, Working Paper, University of California, Davis, 2009 (http://www.econ.ucdavis.edu/faculty/bonanno/PDF/ibr_se.pdf).

[4] Bonanno, Giacomo and Klaus Nehring, How to make sense of the common prior assumption under incomplete information, *International Journal of Game Theory*, 1999, 28: 409-434.

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

[6] Fudenberg, Drew and Jean Tirole, Perfect Bayesian Equilibrium and Sequential Equilibrium, *Journal of Economic Theory*, 1991, 53: 236-260.

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

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

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

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

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

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