



Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Games and Economic Behavior 49 (2004) 49–80

GAMES and
Economic
Behavior

www.elsevier.com/locate/geb

Dynamic interactive epistemology

Oliver Board

Department of Economics, Amherst College, Amherst, MA 01002-5000, USA

Received 20 September 2002

Available online 31 January 2004

Abstract

The epistemic program in game theory uses formal models of interactive reasoning to provide foundations for various game-theoretic solution concepts. Much of this work is based around the (static) Aumann structure model of interactive epistemology, but more recently dynamic models of interactive reasoning have been developed, most notably by Stalnaker [Econ. Philos. 12 (1996) 133–163] and Battigalli and Siniscalchi [J. Econ. Theory 88 (1999) 188–230], and used to analyze rational play in extensive form games. But while the properties of Aumann structures are well understood, without a formal language in which belief and belief revision statements can be expressed, it is unclear exactly what are the properties of these dynamic models. Here we investigate this question by defining such a language. A semantics and syntax are presented, with soundness and completeness theorems linking the two.

© 2003 Elsevier Inc. All rights reserved.

JEL classification: C72; D82

Keywords: Interactive epistemology; Beliefs; Belief revision; Semantic; Syntactic; Language; Canonical structure

1. Introduction

It is well established both theoretically and empirically that strategic reasoning requires agents to form not just conjectures about each other's actions, but also about each other's knowledge and beliefs, which can then be used to infer what actions they might take. In particular, the implications of *common knowledge of rationality*, where all the agents are rational, all know they are all rational, all know that they know, and so on, have been extensively analyzed. More recently, epistemic foundations have been provided for game theoretic solution concepts such as Nash equilibrium (Aumann and Brandenburger, 1995).

E-mail address: ojboard@amherst.edu.

0899-8256/\$ – see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/j.geb.2003.10.006

Comprehensive surveys of work in this area are provided by Dekel and Gul (1997) and Battigalli and Bonanno (1999).

Much of this work is based around the *Aumann structure* model (see Aumann, 1976), in which each agent's knowledge is represented by an information partition over a set of states, or possible worlds. For the purposes of the game theorist, however, Aumann structures have several important limitations. First, they describe a very strong concept of knowledge. An implication of modeling agents' epistemic states with information partitions is that everything they know is true, and that they have complete introspective access to this knowledge, i.e. they know everything they know (positive introspection), and they know everything they do not know (negative introspection). Negative introspection in particular has widely been considered inappropriate. More generally, it has been thought important to analyze agents' beliefs as well as their knowledge. And beliefs, unlike knowledge, can be false. These issues can be dealt with by replacing the information partitions with possibility correspondences (see e.g. Samet, 1990). Beliefs modeled by possibility correspondences at their most general do not satisfy any of the properties described above. By imposing certain restrictions on the correspondences we can recover these properties one by one.

The second problem with using Aumann structures to model rational play in games is that they are essentially static: the epistemic states that they model are fixed, while in dynamic games¹ agents have a chance to change their beliefs as the game progresses. In particular, conjectures about what strategies one's opponents might be playing can be revised as moves are observed. A stark illustration of the importance of such revisions is given by Reny (1993), who shows that once the possibility of belief change is taken into account, the game-theoretic wisdom that common knowledge of rationality implies backward induction in games of perfect information is undermined. As long as the information that an agent learns is consistent with what she already knew or believed, this problem can be handled in the existing framework. The agent's partition (or possibility correspondence) can be refined, in a manner analogous to Bayesian updating of probabilities, to take account of the new information. But, like Bayes rule, this process is not well defined when the information learned is incompatible with the agent's previous beliefs, i.e. she is *surprised*. And modeling the response to such surprises is crucial: to evaluate the rationality of strategies in a dynamic game, we must have a theory about what the players would believe at *every* node in the game, even though some of these nodes will typically be ruled out by the players on the basis of the information they possess at the beginning of the game.

Models of dynamic interactive reasoning have thus been developed. Stalnaker (1996) replaces the information partitions of the Aumann structure with *plausibility orderings* on the set of possible worlds, which encode information not just about each agent's current beliefs, but also about how these beliefs will be revised as new information is learned, even if this new information is a surprise (e.g. it takes the form of an unexpected move made by one's opponent). This seems to be a satisfactory resolution to the problem, and models

¹ That is, games in which there is a flow of information as the game proceeds. These games are commonly represented by the extensive form.

of this kind have been used by Stalnaker and others to analyze rational play in dynamic games.

From a philosophical point of view, however, there is something unsatisfactory about the Aumann structure model and all its extensions, as identified by Aumann (1999) himself: "...the whole idea of 'state of the world,' and of a partition structure that reflects the players' knowledge about the other players' knowledge, is not transparent. What *are* the states? Can they be explicitly described? Where do they come from?" (p. 264). Fagin et al. (1999) elaborate further: "If we think of a state as a complete description of the world, then it must capture all of the agents' knowledge. Since the agents' knowledge is defined in terms of the partitions, the state must include a description of the partitions. This seems to lead to circularity, since the partitions are defined over the states, but the states contain a description of the partitions" (p. 332).

Economists have developed an alternative model of interactive beliefs which seems to avoid this circularity. The hierarchical approach (Mertens and Zamir, 1985; Brandenburger and Dekel, 1993) takes as its starting point a set of *states of nature*, which describe facts of interest about the physical world, such as which strategy profile will be played. Each agent's beliefs about the state of nature is represented by a probability distribution over the set of states of nature; their beliefs about these beliefs are then represented by a probability distribution over these distributions and the set of states of nature; and so on. In this way, we build up an infinite hierarchy of beliefs for each player, called her *type* (after Harsanyi, 1968). In contrast to the Aumann structure approach, where the infinite hierarchy of beliefs is generated implicitly by partitions over obscure states of the world, here it is explicitly constructed from levels of probability distributions over clearly defined states of nature.

The question remains, however, as to whether a state of nature together with a description of each agent's type provides a satisfactory description of a state of the world. For it is not clear that an agent's type gives a complete description of her beliefs. Her type specifies what she believes about all the finite-level beliefs of her opponents, but does it actually describe what she believes about their types, what she believes about what they believe about her type, and so on? It turns out that as long as the types satisfy certain coherency conditions, we can answer this question in the affirmative. These coherency conditions amount to assuming that the agents satisfy positive and negative introspection, and guarantee that the belief hierarchies are closed.

Furthermore, the hierarchical model can be extended to deal with the problem of belief revision. Battigalli and Siniscalchi (1999) have shown how to construct hierarchies of *conditional probability systems*; the level-0 probability systems describe each agent's (probabilistic) beliefs about the physical world as before, but they also encode information about how these beliefs are revised. The level-1 systems represent the agents' beliefs over these level-0 systems, and so on. Again, as long as the appropriate coherency conditions are satisfied, these hierarchies are closed and each agent's type describes all of her beliefs.

Any extra clarity these hierarchical constructions might bring, however, is paid for at a price of greatly-increased complexity. The complexity of these models may well be self-defeating: Aumann (1999) describes them as "cumbersome and far from transparent... In fact, the hierarchy construction is so convoluted that we present it here with some diffidence" (pp. 265, 295). In addition, two more specific problems arise. The first concerns the coherency conditions that are required for closure of the hierarchies. As we have

already discussed, it may not always be appropriate to assume that agents have complete introspective access to their epistemic states; this remains true even if we are dealing with belief rather than knowledge. In the case of conditional probability systems, the coherency assumption becomes even stronger: here it is assumed that agents have complete introspective access to their belief revision schemes as well. Ideally we would like to have a system that is flexible enough to work with or without positive and negative introspection. The second problem arises when we consider the non-probabilistic analogue of these belief hierarchies, where each level in the hierarchy describes simply which members of the previous level the agent considers possible, rather than assigning probabilities to each (the former is not generally derivable from the latter: a world may be considered possible even if it is assigned zero probability). In this case it turns out that, even with the appropriate coherency conditions, the infinite hierarchy does not in general provide a complete description of an agent's uncertainty; that is, it does not tell us which types of her opponents she considers possible (Fagin, 1994; Heifetz and Samet, 1998; Brandenburger and Keisler, 1999).

Thankfully there is a path between this Scylla and Charybdis, between the obscurity of Aumann structures and the complexity of belief hierarchies. *Epistemic logic* is based on a formal language which can express statements about the world and what agents believe about the world and about each other. The language is built up from a set of primitive formulas by means of an inductive rule. The primitive formulas and each step of the inductive process are entirely transparent. Hintikka (1962) showed how *Kripke structures* (Kripke, 1963) can be used to provide a *semantics* for this language, i.e. a set of rules for determining the truth or falsity of every sentence or *formula* in the language. Hence there is no issue about whether or not these structures provide a complete description of the agents' uncertainty: the language itself defines the limits of what we can and cannot say about the agents' beliefs.

There is a very close connection between Kripke structures and Aumann structures: the former are a general version of the latter, where the information partitions are replaced by possibility correspondences (traditionally referred to as *accessibility relations*), plus the addition of an *interpretation* which assigns truth values to the primitive formulas. Kripke structures are general enough to model knowledge or belief, with or without the introspection assumptions. Certain properties of Kripke structures correspond to various axioms and rules governing the behavior of formulas in the language: these axioms and rules, jointly referred to as an *axiom system*, give us a precise characterization of sets of formulas which are true in different types of Kripke structure, and hence an elucidation of the particular concept of knowledge or belief that is being modeled. The axiom system and language form a *syntax* for the logic.

But there is a gap still to be filled. In order to extend the results just described to structures such as Stalnaker's, we must develop a language that is richer than that of epistemic logic. In Section 2 of this paper, we define such a language by adding *revised belief* operators to the standard language. Thus, if $B_i\phi$ is a formula of the language, then so is $B_i^\phi\psi$, to be interpreted "*i* believes that ψ on learning that ϕ ." We then present a semantics for this language consisting of *belief revision structures*, which look much like a generalized version of Stalnaker's structures. A theorem links these structures to an axiom system which describes how these revised belief operators, and the rest of

the language, behave. This axiom system is essentially the most basic axiom system of epistemic logic augmented by additional axioms and rules that correspond to some of the AGM axioms of belief revision (Alchourrón et al., 1985). These axioms are reproduced in Appendix A. Several extensions to the model, including the introduction of introspection and consistency axioms, and common belief operators, are developed in Section 3, and Section 4 comments on some issues which are not treated by our formalism.

Before we start, however, we should comment more carefully on the relevance of this work for game theory. The importance of higher-order beliefs in strategic reasoning is well understood,² and in the dynamic setting it is essential to model how these beliefs change as agents learn new information. Battigalli and Siniscalchi (2002) have used their hierarchical models of belief revision to provide an analysis of forward-induction reasoning in its various guises (including the intuitive criterion of Cho and Kreps, 1987), as well as an epistemic characterization of backward induction. The logical approach adopted here, although less direct in application than the hierarchies of Battigalli and Siniscalchi, forms the basis of an alternative, complementary, framework for analyzing dynamic games, and offers simplicity at the same time as transparency. The simplicity comes from the semantic structures that are used to provide truth conditions for formulas of the formal language: these structures are easily adapted to provide epistemic models of games (see Section 5 for an example). Unlike the constructions of Battigalli and Siniscalchi, which are infinite by definition, these models can be very small. And the axiom system and the language itself, which provide the syntax of the logic, are straightforward to interpret and give us transparency. This syntax lays out the “rules of argument” and allows nothing to be hidden in the formalism. Soundness and completeness theorems establish equivalence between what is true in every structure and what can be proved in the axiom system: the notoriously tricky task of proving that a formula can be derived from a given set of axioms and rules is thus reduced to the mathematical problem of checking that our structures have a particular property. This methodology is adopted by Board (2002a) in a companion paper. Other papers which adopt the logical approach to analyze dynamic games include Clausing (2001) and Feinberg (2001); we discuss the logical components of those papers, along with other related literature, in Section 6.

2. Dynamic interactive epistemology

As discussed in the introduction, an important distinction is made in logic between a syntax and a semantics. A syntax consists of a formal language, defined by a set of formulas, and a proof procedure for generating theorems in that language. The proof procedure, usually expressed in the form of an axiom system, is often rather cumbersome: even basic theorems can be very tricky to prove. A semantics is made out of *structures* that give truth conditions for every formula in the language. A structure is a well-defined mathematical object, and usually very easy to work with, but hard to interpret. The task of the logician is to establish a connection between the syntax and the semantics. This

² For a brief survey of the theoretical literature, see Morris and Shin (2003).

can be done by means of *soundness* and *completeness* theorems, which link the theorems generated by the proof procedure with the truth conditions established by the structures. We start by describing the language we shall work with.

2.1. Language

Our language $\mathcal{L}_n(\Phi)$ is built up from a nonempty set Φ of primitive formulas and an inductive rule. The primitive formulas stand for statements expressing basic facts about the world, such as “agent i plays action a_i .” The inductive rule enables us to build up more complex formulas standing for statements such as “agent i plays action a_i and agent j plays action a_j ,” and “agent j believes that agent i plays action a_i .” Formally, $\mathcal{L}_n(\Phi)$ is defined as the smallest set which satisfies the following conditions:

- (a) if $\phi \in \Phi$, then $\phi \in \mathcal{L}_n(\Phi)$;
- (b) if $\phi, \psi \in \mathcal{L}_n(\Phi)$, then $\neg\phi \in \mathcal{L}_n(\Phi)$ and $(\phi \wedge \psi) \in \mathcal{L}_n(\Phi)$;
- (c) if $\phi, \psi \in \mathcal{L}_n(\Phi)$, then $B_i\phi \in \mathcal{L}_n(\Phi)$ and $B_i^\phi\psi \in \mathcal{L}_n(\Phi)$ for $i = 1, \dots, n$.

For economy of notation, we take Φ and n to be fixed henceforth and omit them from the notation. We also omit parentheses whenever there is no risk of confusion, and use the following standard abbreviations: $\phi \vee \psi$ for $\neg(\neg\phi \wedge \neg\psi)$; $\phi \Rightarrow \psi$ for $\neg\phi \vee \psi$; and $\phi \Leftrightarrow \psi$ for $(\phi \Rightarrow \psi) \wedge (\psi \Rightarrow \phi)$. As discussed in the introduction, \mathcal{L} is the language of epistemic logic augmented by adding modal operators B_i^ϕ that tell us what the agents believe after receiving the information that ϕ . Notice that the language cannot express iterated belief revisions; that is, there are no formulas expressing statements such as “agent i believes that χ on learning that ϕ and then learning that ψ .” We comment on this restriction in Section 4.2.

We now present an axiom system and semantics for \mathcal{L} .

2.2. Axiom system

An axiom system AX consists of a set of axioms and inference rules. An axiom is simply a formula or set of formulas, and an inference rule allows us to infer one formula from a set of other formulas. A *proof* in AX is a finite sequence of formulas, each of which is either an (instance of) an axiom or follows from some of the preceding formulas by applying an inference rule. A *proof of ϕ* is a proof whose last formula is ϕ . We say that ϕ is provable in AX (or ϕ is a *theorem* of AX), and write $AX \vdash \phi$, if there is a proof of ϕ in AX .

We shall consider the axiom system BRS for \mathcal{L} , consisting of the following axioms and inference rules:

Taut	$true$
Dist	$(B_i^\phi\psi \wedge B_i^\phi(\psi \Rightarrow \chi)) \Rightarrow B_i^\phi\chi$
Triv	$B_i\phi \Leftrightarrow B_i^{true}\phi$
Succ	$B_i^\phi\phi$
IE(a)	$B_i^\phi\psi \Rightarrow (B_i^{\phi \wedge \psi}\chi \Leftrightarrow B_i^\phi\chi)$

IE(b)	$\neg B_i^\phi \neg \psi \Rightarrow (B_i^{\phi \wedge \psi} \chi \Leftrightarrow B_i^\phi (\psi \Rightarrow \chi))$
MP	from ϕ and $\phi \Rightarrow \psi$ infer ψ
RE	from ψ infer $B_i^\phi \psi$
LE	from $\phi \Leftrightarrow \psi$ infer $B_i^\phi \chi \Leftrightarrow B_i^\psi \chi$

Note that: any formulas in \mathcal{L} may be substituted for $\phi, \psi, \chi; i \in \{1, \dots, n\}$; *true* stands for any propositional tautology, and *false* stands for \neg *true*. This system is close to the system K of epistemic logic with extra axioms and rules, corresponding roughly to the AGM axioms of belief revision, to describe the behavior of the revised belief operators B_i^ϕ .

Taut, Dist (the *distribution axiom*), MP (*modus ponens*), and RE (the *rule of epistemization*) are familiar from epistemic logic, and need no further comment (but note that Dist and RE apply only to the *revised* belief operators). Jointly, these correspond to AGM axiom (K^*1).³ Triv, the *triviality axiom*, says that if the information received by an agent is trivial (i.e. if it is a propositional tautology), then she does not revise her beliefs;⁴ a corresponding condition is implied by the AGM axioms. This also ensures that ordinary beliefs satisfies the same properties as revised beliefs. Succ, the analogue of (K^*2), is the *success axiom*, which guarantees that the information received is indeed believed in the revised belief state. IE(a) and IE(b) are the axioms of *informational economy*, motivated by the *criterion of informational economy*: our beliefs are not in general gratuitous, and so when we change them in response to new evidence, the change should be no greater than is necessary to incorporate that new evidence. More specifically, IE(a) says that if an agent learns something she already believed, she doesn't revise her beliefs at all; and IE(b) says that if she learns something consistent with her original beliefs, then her revised beliefs are formed simply by adding the new information to her existing stock of beliefs and closing under *modus ponens*. IE(b) corresponds directly to (K^*7) and (K^*8), and, in the presence of Triv, also to (K^*3) and (K^*4); IE(a) is implied by (K^*7) and (K^*8) in the presence of (K^*5). Finally, LE, the *rule of logical equivalence*, corresponds to (K^*6), and says that logically equivalent formulas should lead to identical belief revisions: it only the *content* of the information and not the way it is expressed that determines how beliefs are revised.⁵

2.3. Semantics

The semantics for \mathcal{L} is provided by a *belief revision structure*. This is based on a combination of Grove's (1988) *spheres* model and the Kripke structure framework. The Kripkean accessibility relations are replaced by *plausibility orderings* at every world for

³ Under the numbering system of Gärdenfors (1988), reproduced in Appendix A.

⁴ A referee has pointed out that the Triv axiom could be replaced by the definition $B_i(\cdot) \equiv B_i^{true}(\cdot)$ (for some instance of *true*), simplifying the language as well as axiom system, since now only the revised belief operator would be a primitive of the system. With game-theoretic applications in mind, however, there may be some advantage in retaining a distinction between prior and revised belief operators, to be interpreted as pre-game and in-game beliefs respectively. In games with absent-mindedness, for example, pre-game beliefs and uninformed in-game beliefs may differ: in this case the triviality axiom must be rejected. See Board (2003) for a logical analysis of such games.

⁵ This rules out what psychologists call *framing effects*.

each agent, with the most plausible worlds for a given agent at a particular world taking the role of the accessible worlds for that agent at that world. But a plausibility ordering for an agent tells us not only her current epistemic state, it also encodes information about her belief revision policy. In turn, the structure generates the other agents' beliefs about this belief revision policy, and so on, thus providing truth conditions for each formula of \mathcal{L} . Formally, a *belief revision structure* M over Φ for n agents is a ordered triple $\langle W, \pi, \preceq \rangle$, where W is a set of possible worlds; $\pi : W \times \Phi \rightarrow \{\mathbf{true}, \mathbf{false}\}$ is an interpretation; and \preceq is a vector of binary relations over W , giving the plausibility ordering of each agent at each world. We use \preceq_i^w to denote the plausibility ordering of agent i at world w . Intuitively, $x \preceq_i^w y$ means "from the point of view of agent i at world w , world x is at least as plausible as world y ."

Belief revision structures are used to give truth conditions to formulas. Formally, truth of formulas is characterized by the \models relation: $(M, w) \models \phi$ means that ϕ is true at world w in structure M . We use $[\phi]_M$ to denote the set of worlds in which ϕ is true (the *truth set* of ϕ), i.e. $[\phi]_M = \{w \mid (M, w) \models \phi\}$. If ϕ is true at every world of a given structure, we say that ϕ is *valid in* M , and write $M \models \phi$. Finally, for a given class of structures \mathcal{C} , we say that ϕ is *valid with respect to* \mathcal{C} , and write $\mathcal{C} \models \phi$, if $M \models \phi$ for all $M \in \mathcal{C}$.

Before giving the formal definition of \models , we impose several restrictions on the form of belief revision structures. Define $W_i^w = \{x \mid x \preceq_i^w y \text{ for some } y\}$, the set of worlds which are *conceivable* to agent i at world w , though not necessarily accessible. Then, we assume that:

- R1 for all i, w : \preceq_i^w is complete and transitive on W_i^w ;
 R2 for all i, w : \preceq_i^w is well-founded.

R1 ensures that each plausibility ordering divides all the worlds into ordered equivalence classes; the inconceivable worlds, i.e. those not in W_i^w , are a class unto themselves and are to be considered least plausible. If \preceq_i^w is *well-founded* (R2), then there are no infinitely descending sequences of the form $\dots w_n \prec_i^w w_{n-1} \prec_i^w \dots \prec_i^w w_0$ (where $x \prec_i^w y$ if and only if $x \preceq_i^w y$ and not $y \preceq_i^w x$). This guarantees that for every set $X \subseteq W$, if $X \cap W_i^w \neq \emptyset$, then $\min_i^w \{X \cap W_i^w\} \neq \emptyset$, where \min_i^w is defined in the usual way (i.e. $\min_i^w(X) = \{x \in X \mid \text{for all } y \in X, x \preceq_i^w y\}$); intuitively, it says that if there are any conceivable worlds in a certain set, then there is a most plausible world in that set.⁶ Well-foundedness is satisfied automatically in the case where W is finite. We call a belief revision structure *ordered* if it satisfies R1, and *focused* if it satisfies R2. Let \mathcal{M} denote the class of all belief revision structures that are ordered and focused.

We are now in a position to define \models . The definition proceeds by induction on the form of ϕ :

⁶ Well-foundedness of \preceq_i^w is stronger than the limit assumption proposed by Lewis (1973): if $[\phi]_M \cap W_i^w \neq \emptyset$, then $\min_i^w \{[\phi]_M \cap W_i^w\} \neq \emptyset$. Well-foundedness requires that *every* set which has a nonempty intersection with W_i^w has a least element, while the limit assumption applies this condition only to sets which represent formulas. We impose the stronger condition in order to preserve the clean cut between extra-linguistic reality (as represented by the *frame* (W and \preceq)) and the semantics (which maps the language into the frame).

$$\begin{aligned}
(M, w) \models \phi \quad (\text{for } \phi \in \Phi) \quad &\text{iff } \pi(w)(\phi) = \mathbf{true}; \\
(M, w) \models \phi \wedge \psi \quad &\text{iff } (M, w) \models \phi \quad \text{and} \quad (M, w) \models \psi; \\
(M, w) \models \neg\phi \quad &\text{iff } \text{not } (M, w) \models \phi; \\
(M, w) \models B_i\phi \quad &\text{iff } (M, x) \models \phi \quad \text{for all } x \in \min_i^w(W_i^w); \\
(M, w) \models B_i^\phi\psi \quad &\text{iff } (M, x) \models \psi \quad \text{for all } x \in \min_i^w\{[\phi]_M \cap W_i^w\}.
\end{aligned}$$

The first three rules are straightforward. The fourth rule gives truth conditions for formulas of the form $B_i\phi$ in much the same way as the Kripke semantics, with the most plausible worlds playing the role of the accessible worlds: agent i believes that ϕ if and only if ϕ is true in all the most plausible worlds. The fifth rule operates similarly: the worlds accessible to the agent when she learns that ϕ are precisely the most plausible worlds that are consistent with ϕ ; thus agent i believes that ψ on learning that ϕ if and only if ψ is true in all the most plausible worlds in which ϕ is true.⁷ The five rules provide truth conditions for every formula in \mathcal{L} .

2.4. Soundness and completeness

Before stating the main theorem of the paper, we need to introduce some more terminology.

An axiom system AX is said to be *sound* for a language \mathcal{L} with respect to a class \mathcal{C} of structures if every formula in \mathcal{L} that is provable in AX is valid with respect to \mathcal{C} . The system AX is said to be *complete* for \mathcal{L} with respect to \mathcal{C} if every formula in \mathcal{L} that is valid with respect to \mathcal{C} is provable in AX . We can think of AX as characterizing the class \mathcal{C} if it provides a sound and complete axiomatization of that class, i.e. for all $\phi \in \mathcal{L}$, we have $AX \vdash \phi$ if and only if $\mathcal{C} \models \phi$. Soundness and completeness provide a tight connection between the syntactic notion of provability, which is hard to use but easy to understand, and the semantic notion of validity, which is easy to use but hard to understand.

It turns out that a precise connection can be made between the axiom system BRS and belief revision structures. The following theorem tells us that every theorem ϕ that is provable in BRS is also true in every world of every belief revision structure:

Theorem 1. *BRS is a sound and complete axiomatization w.r.t. \mathcal{M} for formulas in \mathcal{L} .*

The proof of this and all other theorems is given in Appendix B.

To get a more concrete feel for the formalism, the reader may at this point wish to skip forward to Section 5 where there is an example of the logic applied to game theory.

⁷ The purpose of the well-foundedness condition should now be clear: if it does not hold, $B_i^\phi\psi$ could be (vacuously) true even though ψ was not true in any sufficiently plausible ϕ -world, because there might be no *most* plausible ϕ -world. Thus we would clearly have the wrong truth conditions for sentences of this form, and in fact for sentences of the form $B_i\phi$.

2.5. The canonical structure

Before moving on to discuss various extensions of the logic presented in this section, we show how to construct a particularly important belief revision structure M^c , called the *canonical structure for BRS*. To understand what the canonical structure is, we need some more definitions.

For a given axiom system AX , we say that a formula ϕ is *AX-consistent* if $\neg\phi$ is not provable in AX . A finite set of formulas $\{\phi_1, \dots, \phi_k\}$ is *AX-consistent* exactly if $\phi_1 \wedge \dots \wedge \phi_k$ is *AX-consistent*, and an infinite set of formulas is *AX-consistent* exactly if all its finite subsets are *AX-consistent*. Finally, a set of formulas $S \subseteq \mathcal{L}$, is a *maximal AX-consistent set* if (a) it is *AX-consistent*, and (b) for all ϕ in \mathcal{L} but not in S , the set $S \cup \{\phi\}$ is not *AX-consistent*.

The canonical structure has a world w_S corresponding to every maximal *BRS-consistent* set S . This structure is analogous to the *universal* type space construction of Battigalli and Siniscalchi (1999), who extend the work of Mertens and Zamir (1985) and Brandenburger and Dekel (1993) to the dynamic setting. In both cases every *allowable* profile of epistemic types is represented: in the canonical structure by sets of formulas describing each agent's beliefs and how these beliefs will be revised; and in the universal type space an infinite hierarchy of conditional probability systems for each agent. Both approaches rule out certain beliefs: according to the former, an epistemic type is allowable only if the formulas describing it are logically consistent according to the axiom system; in the hierarchical construction of Battigalli and Siniscalchi, the representation of beliefs by conditional probability systems allows only beliefs that satisfy an appropriate set of probability axioms, and additional *coherency* conditions are imposed on the hierarchies to ensure that the various levels of each hierarchy agree with each other.

There are, however, important differences between the two approaches. While the universal type space of Battigalli and Siniscalchi describes the probabilistic beliefs of each agent, the canonical structure presented here tells us what the agents consider possible. It is clear that probabilistic beliefs cannot be recovered from the canonical structure. Nor can possibility correspondences be recovered from the universal type space, unless possibility is identified with strictly positive probability. In addition, the conditional probability systems used by Battigalli and Siniscalchi specify beliefs conditional on *observable* events only; in our terminology, this means that information can take the form of propositional formulas only (i.e. primitive formulas and their conjunctions and negations), which describe the physical world. Our language places no restrictions on the kind of information that may be received; in particular, the possibility that one agent may learn another's beliefs is not ruled out. As far as we are aware, the canonical structure presented here is the first example of a belief-complete construction that is general in this sense.

For the construction of M^c , we introduce some new notation: let $S/B_i^\phi = \{\psi \mid B_i^\phi \psi \in S\}$, i.e. S/B_i^ϕ is the set of formulas believed by i when she learns that ϕ at world w_S . Let $M^c = \langle W, \pi, \preceq \rangle$, where

$$W = \{w_S : S \text{ is a maximal BRS-consistent set}\},$$

$$\pi(w_S)(\phi) = \begin{cases} \text{true} & \text{if } \phi \in S \\ \text{false} & \text{if } \phi \notin S \end{cases} \quad \text{for } \phi \in \Phi,$$

$w_T \preceq_i^{w_S} w_U$ if there is some $\phi \in T \cap U$ such that $S/B_i^\phi \subseteq T$.

Note that for each maximal *BRS*-consistent set of formulas S we have precisely one world w_S . To show that w_S really does correspond to S , we must prove the following proposition:

Proposition 1. $(M^c, w_S) \models \phi$ if and only if $\phi \in S$.

Proposition 1 says that S contains exactly those formulas which are true at w_S . The proof is given in Appendix B. Another soundness and completeness theorem emerges as a corollary of this proposition.

Corollary 1. *BRS* is a sound and complete axiomatization w.r.t. M^c for formulas in \mathcal{L} .

For (soundness) if ϕ is provable in *BRS*, it must be contained in every maximal *BRS*-consistent set (see proof of Theorem 1), and by Proposition 1, it is therefore valid with respect to M^c . And (completeness) if ϕ is valid with respect to M^c , Proposition 1 tells us that it must be contained in every maximal *BRS*-consistent set. It follows that ϕ is provable in *BRS*: if not, $\neg\phi$ would be *BRS*-consistent and thus contained in *some* maximal *BRS*-consistent set (see again proof of Theorem 1); but $\{\phi, \neg\phi\}$ is not *BRS*-consistent, and so ϕ and $\neg\phi$ cannot be contained in the same maximal *BRS*-consistent set.

Although it might seem that the soundness part of Corollary 1 follows from Theorem 1, and that the completeness part of Theorem 1 follows from Corollary 1, this is not the case because $M^c \notin \mathcal{M}$. The reason is that some of the $\preceq_i^{w_S}$ relations are not well-founded (though they are complete and transitive on $W_i^{w_S}$ in each case). Construct a sequence of maximal *BRS*-consistent sets containing the formulas in Table 1 and a maximal *BRS*-consistent set S such that $S/B_i^{\neg\phi} = T_1, S/B_i^{\neg B_i\phi} = T_2, S/B_i^{\neg B_i B_i\phi} = T_3, \dots, S/B_i = T_\infty$. Then it follows from Proposition 1 and the definition of \models that $\dots w_{T_3} \prec_i^{w_S} w_{T_2} \prec_i^{w_S} w_{T_1}$, i.e. we have an infinitely descending sequence and $\preceq_i^{w_S}$ is not well-founded. Nonetheless Corollary 1 tells us that the tight connection between valid formulas and formulas that are provable in *BRS* still holds.

The canonical structure is useful for certain game-theoretic applications. Both forward and backward induction are based on the premise that players try to interpret their opponents' strategy choices as rational whenever possible. But what it is rational for a player to do depends on her beliefs. Using a structure that rules out certain beliefs to analyze a game restricts the set of available explanations for a particular action. So if we are interested in which strategies are compatible with rationality and which are not we must

Table 1

T_1	T_2	T_3	T_4	...	T_∞
$\neg\phi$	ϕ	ϕ	ϕ		ϕ
$\neg B_i\phi$	$\neg B_i\phi$	$B_i\phi$	$B_i\phi$		$B_i\phi$
$\neg B_i B_i\phi$	$\neg B_i B_i\phi$	$\neg B_i B_i\phi$	$B_i B_i\phi$		$B_i B_i\phi$
$\neg B_i B_i B_i\phi$	$\neg B_i B_i B_i\phi$	$\neg B_i B_i B_i\phi$	$\neg B_i B_i B_i\phi$		$B_i B_i B_i\phi$
\vdots	\vdots	\vdots	\vdots	\vdots	

work with a structure that includes all possible beliefs. The canonical structure does just this. See Battigalli and Siniscalchi (2002) and Board (2002a) for further elaboration of this point.

We finish this section with a brief comment on the impossibility results of Fagin (1994), Heifetz and Samet (1998), Brandenburger and Keisler (1999) and others, which show that if epistemic types are represented by possibility sets (as they are here) rather than probability distributions, then a structure containing all epistemic types cannot exist. These results would seem to contradict our claim that the canonical structure does contain a representation of every epistemic type. Brandenburger (2002) explains how the two can be reconciled: "...completeness is impossible if literally all possibility sets are wanted. But if we make topological assumptions that serve to rule out certain kinds of possibility sets, then a (restrictedly) complete structure may exist" (p. 4). Working with a formal language has precisely this effect. It is formulas of this language and not arbitrary sets of worlds that are the content of beliefs (and of information), and in any given model there may be sets of worlds that do not represent any formula of the language. Similarly, in the context of a hierarchical framework, Mariotti et al. (2002) evade the impossibility result by assuming that the underlying space of uncertainty (the set of states of nature) is compact and Hausdorff and the possibility sets are compact. Like Battigalli and Siniscalchi, they consider beliefs conditional on non-epistemic events only.

3. Extensions

3.1. Introspection

Consider the following additional axioms:

$$\begin{aligned} \text{TPI} \quad & B_i^\phi \psi \Rightarrow B_i^X B_i^\phi \psi, \\ \text{TNI} \quad & \neg B_i^\phi \psi \Rightarrow B_i^X \neg B_i^\phi \psi. \end{aligned}$$

TPI and TNI are the axioms of *total positive introspection* and *total negative introspection*, and state that agents have complete introspective access to their own minds, including not only their current beliefs but also how these beliefs would be or would have been revised. Let *BRSI* be the axiom system formed by the addition of TPI and TNI to *BRS*. To illustrate the strength of these axioms, we consider three implications. The first is introspection of current beliefs: $B_i^\phi \psi \Rightarrow B_i^\phi B_i^\phi \psi$ and $\neg B_i^\phi \psi \Rightarrow B_i^\phi \neg B_i^\phi \psi$. The knowledge analogues of these principles emerge as properties of the Aumann structure model discussed in the introduction and widely used in economic theory, but their universal applicability has been questioned by Geanakoplos (1992), among others. Second, TPI and TNI imply that agents have correct beliefs about their future beliefs, whatever information they receive: $B_i^\phi \psi \Rightarrow B_i B_i^\phi \psi$ and $\neg B_i^\phi \psi \Rightarrow B_i \neg B_i^\phi \psi$. Finally, it is implied that agents can recall their prior beliefs: $B_i \psi \Rightarrow B_i^\phi B_i \psi$ and $\neg B_i \psi \Rightarrow B_i^\phi \neg B_i \psi$. This assumption is inappropriate in certain games and decision problems, such as the absent-minded driver

paradox of Piccione and Rubinstein (1997). Bonanno (2003) provides a careful analysis of this and other memory axioms in the context of extensive form games.

Imposing an additional restriction on the form of the \preceq_i^w relations provides a semantic characterization of TPI and TNI:

R3 for all i, w, x, y, z : if $x \in W_i^w$, then $y \preceq_i^x z$ if and only if $y \preceq_i^w z$.

Intuitively, R3 says an agent has the same plausibility ordering in every world that is conceivable to her. If a belief revision structure satisfies R3 we call it *absolute*, and let \mathcal{A} be the set of belief revision structures which satisfy R1–R3. Then the following result formalizes the link between TPI and TNI, and absoluteness.

Theorem 2. *BRSI is a sound and complete axiomatization w.r.t. \mathcal{A} for formulas in \mathcal{L} .*

3.2. Consistency

The observant reader will have noticed that there is no axiom in *BRS* corresponding to the AGM *consistency axiom* (K^*5). In the AGM system, this axiom ensures that agents' beliefs are logically consistent whenever possible, i.e. whenever the information learned is logically consistent (if the information is not consistent, then (K^*2) forces inconsistent beliefs on the agent). But any attempt to axiomatize this in our logic will lead to circularity, since the notion of logical consistency presupposes a particular axiom system. The AGM system and Friedman and Halpern's (1994) more expressive logic of belief change avoid this problem by working with two languages, one for describing facts about the world which can be learned, and another for talking about beliefs. Their consistency axioms ((K^*5) and PS respectively) apply to the second language, and make reference to the logical consistency only of formulas of the first. In addition to the analytical convenience of working with one language rather than two, an advantage of our approach is that no restrictions are imposed on the form that information may take. This issue is discussed in more detail in Section 6.

An independent reason to be suspicious of the AGM consistency axiom is that it affords no purely semantic representation. To guarantee the validity of this axiom, we would need to restrict our attention to belief revision structures which contain *enough* worlds: for each logically consistent formula, we need at least one world in which that formula is true. Hence in Grove's (1988) semantics for the AGM system, the set of worlds is identified with the set of maximal consistent sets of formulas of the object language, and Friedman and Halpern make the assumption that their structures are *saturated*, i.e. there is at least one minimal world for every consistent formula of the object language. But logicians tend to think of the *frame* (in this case the worlds and the plausibility orderings) as a representation of the extra-linguistic reality, which is mapped onto a formal language by an interpretation and semantic rules. A reality that can be described only in syntactic terms seems artificial.⁸

In the place of (K^*5), we consider the *weak consistency axiom*:

⁸ It may, however, be reasonable to impose linguistic constraints on belief revision structures for the sake of particular applications. For example, if we wish to model rational play in a game, we may wish to assume that

WCon $\phi \Rightarrow \neg B_i^\phi \text{ false.}$

WCon says that as long as the information an agent receives is true, her revised beliefs are consistent,⁹ and is represented by the following assumption, which says that the actual world is always conceivable:

R4 for all $i, w: w \in W_i^w$.

Call a belief revision structure satisfying R4 *inclusive*. Let \mathcal{I} the class of all inclusive belief revision structures which also satisfy R1 and R2, and let *BRSC* be the axiom system consisting of *BRS* and WCon. Then:

Theorem 3. *BRSC is a sound and complete axiomatization with respect to \mathcal{I} for formulas in the language \mathcal{L} .*

3.3. A simplification

If we are willing to accept the introspection axioms, TPI and TNI, and consistency axiom, WCon, discussed above, the belief revision structures which provide the semantics for \mathcal{L} can be greatly simplified. Let *BSIC* be the resulting axiom system (i.e. *BSIC* = *BRS* + TPI + TNI + WCon). Theorem 4 says that R1–R4 give a semantic characterization of *BSIC*:

Theorem 4. *BSIC is a sound and complete axiomatization with respect to $\mathcal{A} \cap \mathcal{I}$ for formulas in the language \mathcal{L} .*

It turns out that if R1–R4 are satisfied, the plausibility orderings of each agent can be replaced by a *single* binary relation, \triangleleft_i , defined as follows:

$$w \triangleleft_i x \quad \text{if and only if} \quad w \triangleleft_i^x x.$$

The intuition is as follows: recall that R3 says that each agent has the same plausibility ordering at every world conceivable to her, and R4 says that the actual world is always conceivable. If both conditions are satisfied, the plausibility orderings divide the worlds into distinct subsets, which can then be described by a single relation. Although some information is lost by this transformation, since many different families of plausibility orderings for a given agent map onto the same \triangleleft_i relation, all of the *semantically relevant* information is preserved. Truth conditions can be given in terms of the \triangleleft_i relations

there is at least one world corresponding to every strategy profile, or even that there is a world for every consistent set of beliefs the players might hold (as in the canonical structure of Section 2.5). These restrictions represent contingent facts about particular situations, not matters of logic alone.

⁹ For the purposes of modeling rational play in extensive games, replacing the AGM consistency axiom WCon is without loss of generality: the information structure of extensive form games is such that the information received is always true, and so WCon guarantees that agents maintain consistency of beliefs. For more on this point, see Board (2002b).

which for every formula match the standard truth conditions. First observe that the set of conceivable worlds can be defined as follows:

$$W_i^w = \{x \mid x \triangleleft_i w \text{ or } w \triangleleft_i x\}.$$

To show that this definition is correct, we must prove

Proposition 2. $\{x \mid x \triangleleft_i w \text{ or } w \triangleleft_i x\} = \{x \mid x \preceq_i^w y \text{ for some } y\}$.

Next, we show how the \triangleleft_i relation can be used to define truth of formulas. The truth conditions for primitive formulas, conjunctions and negations are the same as before; truth of formulas of the forms $B_i\phi$ and $B_i^\phi\psi$ are defined as follows:

$$\begin{aligned} (M, w) \models B_i\phi & \text{ iff } (M, x) \models \phi \text{ for all } x \in \min_i(W_i^w), \\ (M, w) \models B_i^\phi\psi & \text{ iff } (M, x) \models \psi \text{ for all } x \in \min_i\{[\phi]_M \cap W_i^w\}, \end{aligned}$$

where $\min_i(X) = \{x \in X \mid \text{for all } y \in X, x \triangleleft_i y\}$. The equivalence of these truth conditions and the standard conditions follows immediately from Proposition 3:

Proposition 3. For all $X \subseteq W_i^w$, $\min_i(X \cap W_i^w) = \min_i^w(X \cap W_i^w)$.

The structures resulting from this simplification bear a very close resemblance to the belief revision models developed by Stalnaker (1996). Our \triangleleft_i relations work in exactly the same way as the reverse of his Q_i relations: if we define $w Q_i x$ if and only $x \triangleleft_i w$, then the truth conditions for formulas expressing beliefs and revised beliefs are identical. Furthermore, it can be shown our \triangleleft_i relations satisfy the same properties he requires of his Q_i relations (i.e. they are reflexive and transitive, and if two worlds are related (in either direction) to a third world, then those two worlds are related (in some direction) to each other). Thus it follows from Theorem 4 that the axiom system *BRSIC* provides a precise syntactic characterization of Stalnaker's purely semantic logic.

3.4. Common belief

One of the strengths of the logic we are developing here, and of epistemic logic in general, is that they give us a complete description of agents' beliefs about agents' beliefs. As we discussed in the introduction, this is particularly useful for game theory since such beliefs are often considered necessary for sophisticated strategic reasoning. In particular, we can give an account of the notion of *common belief* frequently used by economists.

But common belief cannot be expressed in the language \mathcal{L} defined above, since infinite conjunctions of formulas in \mathcal{L} are not themselves formulas of \mathcal{L} . To remedy this problem, we augment the language with the modal operators E ("everyone believes that..."), and C ("it is common belief that...").¹⁰ Formally, \mathcal{L}^C is defined by adding the following condition to the definition of \mathcal{L} in Section 2.1:

¹⁰ Operators representing mutual and common *revised* beliefs could be added in the same way, with a corresponding set of axioms. But if our interest is in providing epistemic characterization results for solution concepts in game theory, this is probably not necessary: these results typically consider restrictions on the *prior*

(d) if $\phi \in \mathcal{L}^C$, then $E\phi \in \mathcal{L}^C$ and $C\phi \in \mathcal{L}^C$.

It is straightforward to extend our axiom system to incorporate the E operator:

$$\text{E} \quad E\phi \Leftrightarrow \bigwedge_{i \in \{1, \dots, n\}} B_i \phi.$$

E says simply that everyone believes that ϕ if and only if every agent believes that ϕ . Unfortunately, the axiomatic characterization of common belief is trickier. The problem is that although common belief is an infinite concept, our axioms must be finite in length. It turns out that the following axiom-rule pair (familiar from epistemic logic) will serve our purpose:

$$\begin{aligned} \text{FP} \quad & C\phi \Rightarrow E(\phi \wedge C\phi), \\ \text{IR} \quad & \text{from } \phi \Rightarrow E(\psi \wedge \phi) \text{ infer } \phi \Rightarrow C\psi. \end{aligned}$$

FP and IR which are known as the *fixed-point axiom* and the *induction rule*, are harder to interpret. We shall merely remark that jointly they imply that common belief has all the properties of (individual) belief. For example, if B_i satisfies TPI, so too does C . Let BRS^C (respectively, $BRSI^C$, $BRSC^C$, and $BRSIC^C$) denote the axiom system formed by adding E, FP, and RI to BRS (respectively, $BRSI$, $BRSC$, and $BRSIC$).

The definition of truth for the augmented language, \mathcal{L}^C is extended exactly as we would expect. $E\phi$ is true just if everyone believes that ϕ :

$$(M, w) \models E\phi \quad \text{iff} \quad (M, w) \models B_i \phi \quad \text{for all } i \in \{1, \dots, n\};$$

and $C\phi$ is true if everyone believes that ϕ , everyone believes that everyone believes that ϕ , and so on. So, letting $E^0\phi$ be an abbreviation for ϕ , and $E^{k+1}\phi$ be an abbreviation for $E E^k \phi$, we have:

$$(M, w) \models C\phi \quad \text{iff} \quad (M, w) \models E^k \phi \quad \text{for } k = 1, 2, \dots$$

The following theorem confirms the equivalence of the syntactic and semantic characterization of common belief:

Theorem 5. (a) BRS^C is a sound and complete axiomatization w.r.t. \mathcal{M} for formulas in \mathcal{L}^C .
 (b) $BRSI^C$ is a sound and complete axiomatization w.r.t. \mathcal{A} for formulas in \mathcal{L}^C .
 (c) $BRSC^C$ is a sound and complete axiomatization w.r.t. \mathcal{I} for formulas in \mathcal{L}^C .
 (d) $BRSIC^C$ is a sound and complete axiomatization w.r.t. $\mathcal{A} \cap \mathcal{I}$ for formulas in \mathcal{L}^C .

beliefs of the players (i.e. their beliefs before the game is played), and the way they are disposed to revise these beliefs.

4. Comments

4.1. Knowledge

While our language allows us to make statements about agents' beliefs (and how these beliefs are revised), economists often make assumptions about agents' *knowledge*. Knowledge could be modeled by adding another set of modal operators to our language: K_i ("i knows that...").

As we discussed in the introduction, in the economics literature knowledge is most commonly analyzed using Aumann's (1976) information partition model. The properties of this model are well understood. An Aumann structure can be provided with an interpretation and used to provide truth conditions for a language containing knowledge operators. The properties of the knowledge operators can then be precisely described by a set of axioms which are sound and complete with respect to the class of all (enriched) Aumann structures. In addition to the appropriate analogues of Taut, Dist, MP and RE, this axiom system contains:

T	$K_i\phi \Rightarrow \phi,$
PI	$K_i\phi \Rightarrow K_i K_i\phi,$
NI	$\neg K_i\phi \Rightarrow K_i\neg K_i\phi.$

T, the *truth* axiom, is uncontroversial: it says simply that what is known must be true; PI (*positive introspection*) and NI (*negative introspection*), on the other hand, are even more controversial in the context of knowledge than in the context of belief. They say respectively that an agent knows what she knows and knows what she does not know. The problem is the following: as long as we accept the truth axiom, the concept of knowledge imposes an external condition on the agent's cognitive state. Thus even if the agent has complete introspective access to what she believes and does not believe, the introspection axioms do to carry over to knowledge through logic alone.

So it seems that we must reject PI and NI, and take T as a starting point in the analysis of knowledge. But philosophers have long argued that true belief, while necessary, is not a sufficient condition for knowledge. For a true belief to be classified as knowledge, it is required in addition that it be somehow justified in an appropriate manner. Reflecting this requirement, Stalnaker (1996) appeals to an analysis of knowledge called the *defeasibility analysis*. The idea behind this account is that "if a person has knowledge, then that person's justification must be sufficiently strong that it is not capable of being defeated by evidence that he does not possess" (Pappas and Swain, 1978).

Stalnaker uses his (semantic) model of belief revision to formalize this idea, by defining knowledge as follows: an agent knows that ϕ if and only if ϕ is true, she believes that ϕ , and she continues to believe that ϕ if any true information is received. Truth conditions for formulas of the form $K_i\phi$ can be provided as follows:¹¹

$$(M, w) \models K_i\phi \quad \text{iff} \quad (M, w) \models \phi \quad \text{and} \quad \text{if } (M, w) \models \psi, \text{ then } (M, w) \models B_i^\psi \phi.$$

¹¹ If we assume that R4 holds, this can be replaced by the much simpler $(M, w) \models K_i\phi$ iff $(M, x) \models \phi$ for all $x \preceq_i^w w$, which is precisely the condition given by Stalnaker.

But we have been unable to find a finite axiomatization: the direct translation of Stalnaker's definition involves a formula of infinite length. Let ψ_1, ψ_2, \dots be an enumeration of all the formulas of \mathcal{L} . Then the appropriate axiom would be:

$$\text{Know} \quad K_i \phi \Leftrightarrow (\phi \wedge (\psi_1 \Rightarrow B_i^{\psi_1} \phi) \wedge (\psi_2 \Rightarrow B_i^{\psi_2} \phi) \wedge \dots).$$

(Note that we do not need to include the formula $B_i \phi$ on the right-hand side, since, in the presence of Taut and Triv, it is implied by $\psi \Rightarrow B_i^{\psi} \phi$ when *true* is substituted for ψ .)

An alternative approach would be to treat Stalnaker's definition as providing a necessary but not sufficient condition for knowledge:

$$\text{Know}' \quad K_i \phi \Rightarrow (\phi \wedge (\psi \Rightarrow B_i^{\psi} \phi)).$$

It is an open question whether the axiom system consisting of *BRS* and *Know'* is sound and complete given the proposed semantics.

4.2. Iterated belief revision

In an extensive form game, of course, beliefs may need to be revised more than once: new information may be received at each round of the game. But the language \mathcal{L} is not rich enough to express iterated belief revisions. Although it would be a simple task to augment the language to allow such iterations, it is not obvious what the axioms governing these iterations should be. Nor is it clear how the semantics should be extended: the rules for revision give us a new set of most plausible worlds after a formula ϕ is learned (that is, the lowest ranking members of $[\phi]$), but we are not told the relative plausibilities of all the other worlds. So we need some way of preserving the relative plausibility data given by the \preceq_i^w relations, while taking into account the new information ϕ . More precisely, we need to construct a new ordering that represents the revised epistemic state, so that we can re-apply the revision rule as more information is learned. Further discussion of these issues is beyond the scope of this paper. The interested reader is referred to Spohn (1988).

These issues can be avoided if we restrict our attention to extensive form games of perfect recall, which are the focus of much of modern game theory. Such games can be analyzed without loss of generality by considering only single revisions if we are willing to accept a plausible assumption. In such games, the sequence of information that the players receive as the game progresses and they observe moves made by their opponents is of a very particular kind: each new piece of information logically implies every previous piece, since the set of strategies consistent with any given information set for a player is a subset of the set of strategies consistent with any predecessor information sets.¹² If ψ logically implies ϕ , it seems reasonable to assume that learning ϕ and then ψ will generate the same beliefs as if one learns ψ at first: in both cases exactly the same information is learned. This assumption is adopted by Board (2002a).

¹² Board (2002b) suggests that this information structure is unduly restrictive, even under the assumption of perfect recall, and shows that it rules out certain interesting situations.




Fig. 1. Centipede game.

5. An example: the failure of backward induction

In this section we present a simple example to show how our logic can be applied to game theory. Consider the three-round centipede game in Fig. 1.

We shall show that common belief in rationality at the beginning of the game does not imply that the backward induction outcome will result.¹³ Our language consists of only three primitive formulas, l_1 , l_2 , and l_3 , to be interpreted as “player 1 chooses L_1 ,” “player 2 chooses L_2 ,”¹⁴ and “player 1 chooses L_3 ,” respectively. Statements about which strategy profile is chosen and what the outcome is can be built up from these three formulas using conjunctions and negations: for example, $l_1 \wedge \neg l_2 \wedge \neg l_3$ represents the proposition “player 1 chooses $L_1 T_3$ and player 2 chooses T_2 ,” and $\neg l_1$ represents the proposition “the backward induction outcome is realized.” Player i is said to be rational (Rat_i) if she believes her chosen strategy is certain to yield a higher payoff than any other, at every node at which she is on move.¹⁵ Note that we are not introducing any new formulas into the language: whether or not a player is rational is determined completely by her strategy choice and her first-order beliefs, i.e. her beliefs about which strategy profile will be chosen. Thus Rat_i is simply an abbreviation for a longer formula, as follows:

$$Rat_1 \equiv (l_1 \wedge \neg l_3 \wedge B_1 l_2) \vee (\neg l_1 \wedge \neg l_3 \wedge B_1 \neg l_2),$$

$$Rat_2 \equiv (l_2 \wedge B_2^1 l_3) \vee (\neg l_2 \wedge B_2^1 \neg l_3).$$

These equivalences state that player 1 is rational if and only if she chooses $L_1 T_3$ and she believes (at node 1) that player 2 will choose L_2 , or if she chooses $T_1 T_3$ and she believes that player 2 will choose T_2 ; and player 2 is rational if and only if he chooses L_2 and he believes (at node 2) that player 1 will choose L_3 , or if he chooses T_2 and he believes that player 1 will choose T_3 . Finally, we shall use CBR as shorthand for $Rat_1 \wedge Rat_2 \wedge C(Rat_1 \wedge Rat_2)$ (“everyone is rational and there is common belief in rationality”).

¹³ Of course, this is not a new result (see, e.g., Reny, 1993 and Ben-Porath, 1997). The aim here is simply to show the logic at work.

¹⁴ Or “player 2 chose/will choose/would have chosen/would choose L_2 .” There is no notion of time in our logic, so formulas should be interpreted as past, present or future, indicative or subjunctive depending on the viewpoint. The primitive formulas play two roles: they describe the way the game is actually played, and they provide a set of counterfactuals for evaluating the payoffs if the action taken at any node deviates from the specified action.

¹⁵ In Board (2002a) we use the more standard notion of expected utility maximization. Adopting the current definition allows us to avoid the use of probabilities while strengthening the result.

A natural belief revision structure, M , for this game consists of one world for each strategy profile, labeled LLL, LLT, \dots, TTT , with π defined in the obvious way (i.e. $\pi(LTT)(l_3) = \mathbf{false}$, for example). Now suppose that LLT is the actual world, and that the plausibility ordering satisfies:

$$\begin{aligned} LLT &<_1^{LLT} w \quad \text{for all } w \neq LLT, \\ TLT &<_2^{LLT} LLL <_2^{LLT} w \quad \text{for all } w \neq TLT, LLL, \\ TTT &<_1^{TLT} w \quad \text{for all } w \neq TTT, \\ TLT &<_2^{TLT} LLL <_2^{TLT} w \quad \text{for all } w \neq TLT, LLL, \\ TTT &<_1^{TTT} w \quad \text{for all } w \neq TTT, \\ TTT &<_2^{TTT} LTT <_2^{TTT} w \quad \text{for all } w \neq TTT, LTT. \end{aligned}$$

The remainder of the ordering can be specified in any way that satisfies R1–R4.

First observe that Rat_1 and Rat_2 hold at LLT, TLT , and TTT . At TLT , for instance, player 1 considers only world TTT to be possible: she believes that player 2 will choose T_2 , so T_1T_3 is indeed the rational choice. Player 2, on the other hand, believes only world TLT to be possible; but at node 2, when he is on move, he believes only world LLL to be possible: here he believes that player 1 will choose L_3 , and therefore L_2 is rational. It follows that CBR holds at LLT .¹⁶ But the backward induction outcome is not chosen. Formally, we have shown that

$$(M, LLT) \models \neg(CBR \Rightarrow \neg l_1).$$

We can now appeal to Theorem 5(d) to show that

$$BRSIC^C \not\vdash CBR \Rightarrow \neg l_1.$$

We have proved that we cannot prove that common belief in rationality implies backward induction. The intuition behind our counterexample is as follows: even though player 2 believes at the beginning of the game that player 1 is rational and therefore will choose T_3 if node 3 is reached, he no longer believes this at node 2: if player 1 were rational, according to the beliefs player 2 believes she has, she should not have played L_1 , and node 2 should not have been reached. So he plays L_2 . Player 1 correctly believes that player 2 will reason this way, so she plays L_1 after all, tricking player 2 into thinking that she is irrational.

A more detailed account of how our logic of belief revision can be applied to games, along with an analysis of the positive implications of common belief in rationality, can be found in Board (2002a).

6. Related literature

Much of the related literature has been discussed in the main text of this paper. Here we provide a summary and mention some important omissions.

¹⁶ This can be shown quickly using the notion of *reachability* (see Fagin et al., 1995, p. 23): only worlds LLT, TLT and TTT are reachable from world LLT .

The semantics of our logic bear a close resemblance to those of conditional logic (Lewis, 1973; Burgess, 1981). The belief revision structures considered above are essentially a multi-agent version of Burgess' *models*. R1 corresponds to his transitivity and connectivity requirements. The counterpart of R2 is Lewis' *limit assumption* (L), and R3 and R4 correspond to *local absoluteness* (A-) and *total reflexivity* (T), respectively. But their models are used to provide truth conditions for conditional formulas, while here we are interested in belief revision; the axioms of conditional logic are very different from the axioms of belief revision.

Stalnaker (1996) also uses semantic structures very similar to those employed here; the results of Section 3.3 show that his models are essentially belief revision structures which satisfy R1–R4. But he provides no formal language and no syntax. Thus the results of this paper are complementary to his: our axiom system *BRSIC* can be thought of as characterizing his models. Friedman and Halpern's (1994) logic of belief change does provide a syntax as well as a semantics for characterizing the belief revision process, with soundness and completeness theorems linking the two. The key difference between their work and our own is that they use two distinct formal languages, one for describing facts about the physical world which can be learned, and another for talking about beliefs. Thus in their system agents cannot learn about each other's beliefs. Friedman and Halpern suggest that this restriction is necessary to avoid a triviality result established by Gärdenfors (1986), but the results of this paper show that this is not the case. Triviality can be avoided as long as the *Ramsey test* ($B_i^\phi B_i^\psi \chi \Leftrightarrow B_i^{\phi \wedge \psi} \chi$) is not adopted as an axiom.

It could be argued, however, that there is no advantage to this extra flexibility, since epistemic events are not observable and therefore cannot play the role of inputs in the belief revision process. Clearly this argument does not apply to beliefs about oneself, which can be acquired by reflection over the course of time. But whether or not we can have direct perceptual access to others' mental states is a philosophically controversial issue. McDowell (1982), for example, claims that we can. This is not a form of mind reading. Our perceptions of each others' minds can be epistemically direct and yet causally mediated, as indeed are our perceptions of physical objects.

If we cannot learn directly about each others' beliefs, as Levi (1988) seems to argue, then we must explain where our beliefs about each others' mental states come from. Many of these beliefs are acquired through conversation. If the learning procedure is not direct, then presumably Alice acquires the belief that Bob believes that ϕ not by updating on "Bob believes that ϕ " but by updating on "Bob said ' ϕ '" and combining it with the prior belief "Bob said ' ϕ ' \Rightarrow Bob believes that ϕ ." But certainly we are not usually conscious of making such inferences; and if the way we learn about the beliefs of others when we converse with them seems non-inferential, then unless there is some positive reason for doubting the veracity of this seeming, it is rational to accept this as the real situation.

Even if we reject the view that others' mental states are directly observable, however, we could accept that we sometimes behave as if they are, and treat what people say as a direct expression of their beliefs. Indeed, Lewis (1975) takes such behavior to be part of the definition of language use. The *cheap talk* literature (e.g. Crawford and Sobel (1982)) attempts to explain language use as an equilibrium outcome, but elsewhere it has proved fruitful to adopt an assumption of sincerity (e.g. Geanakoplos and Polemarchakis (1982)). Less abstract applications where sincerity may be a reasonable assumption include

consultancy advice, where agents pay to receive specialist opinions, and expert testimony in the courtroom, which yields information about the expert's beliefs, and not hard facts about the physical world.

Working with a language that is rich enough to describe a wide range of revised beliefs may be useful for another, theoretical, reason. The results of Battigalli (1996) and Kohlberg and Reny (1997) show that the consistency requirement of sequential equilibrium can be characterized by independence conditions imposed on the conditional beliefs of the players of a game. These results require that the conditional probability measures be defined for every nonempty set of strategy profiles; yet many of these subsets do not correspond to events that could be observed on any play path of the game. Allowing "virtual learning" is essential for their approach. It seems probable that extending the class of virtually learnable events to include epistemic events would make no difference in this case. But in psychological games (see Geanakoplos et al., 1989; Dufwenberg, 2002) the payoffs themselves depend on players' beliefs. In this setting, it is quite plausible that interesting restrictions on the way players revise their beliefs in response to *all* types of information would generate restrictions on actual beliefs conditional on observable events that would not obtain if one considered only non-epistemic information.

Clausing (2001) formulates a logic of belief revision which is general in the sense we have been discussing: no restrictions are imposed on what can be learned.¹⁷ In addition, Clausing's language allows him to express conditional statements of the form "if ϕ were true, then it would be the case that ψ ." He states (but does not prove) soundness and completeness theorems for a weak and a strong version of his logic. There are close parallels between this work and the current paper; in particular, Clausing's axioms of belief revision (B1–B7) and the belief revision component of the axiom system *BRS* developed here are both based on the AGM axioms. Where the logics differ most is in the semantics. In the place of the plausibility orderings employed here, Clausing uses *state selection functions* to provide truth conditions for formulas involving belief revision operators.¹⁸ For each world w and every formula ϕ , an agent's state selection function picks out a set of worlds, $f_i(w, \phi) \subseteq 2^W$. $f_i(w, \phi)$ is interpreted as the set of worlds the agent considers possible at world w if she learns that ϕ , and plays precisely the same role as the most plausible worlds at which ϕ is true in our semantic framework. But plausibility orderings, unlike state selection functions, make no reference to the language. This allows us to establish results about games by purely mathematical reasoning, which can then be translated (by soundness and completeness) into theorems of logic, as in the example of Section 5.

Alternative models of interactive belief revision have been developed by Battigalli and Siniscalchi (1999) and Brandenburger and Keisler (2002), who show how the hierarchical approach of Mertens and Zamir (1985) can be extended to the dynamic setting. The differences between these models and the current work have been discussed in the introduction and in Section 2.5.

¹⁷ We are grateful to an anonymous referee for pointing out this reference.

¹⁸ The results of Halpern (1998) suggest that it may be possible to replace the selection functions and with plausibility orderings.

Before concluding this section we should mention the work of Feinberg (2002). Feinberg also develops a logic for reasoning about dynamic games, and uses it to establish various game-theoretic results. The language that Feinberg works with is richer than that considered here in two respects: it can express probabilistic beliefs, and it can express statements about utilities (though completeness of the logic is proved only for formulas that do not contain these operators). On the other hand, there is no belief revision component to the logic. Rather each player is represented by a different agent for each node at which she is on move, and any links between the beliefs of these agents are expressed in a meta language. Although the logic is dynamic in the sense that it describes player's beliefs throughout the game, it does not attempt to model the development of these beliefs as new information is learned.

7. Conclusion

The aim of this paper has been to develop a dynamic model of interactive reasoning which combines analytical simplicity with clarity of interpretation. Belief revision structures are similar to the models used very successfully by Stalnaker to analyze rational play in extensive form games (Stalnaker, 1996), and to shed light on the forward and backward induction procedures (Stalnaker, 1998). These structures provide truth conditions for a formal language. Soundness and completeness theorems establish tight connections between the formulas that are true in various classes of belief revision structure, and those that are provable in certain axiom systems, thereby giving us a precise understanding of what the structures mean.

Acknowledgments

I thank Paolo Battigalli, Timothy Williamson, an associate editor and two anonymous referees, whose detailed comments led to substantial improvements in the content and exposition of this paper. Helpful comments from Michael Bacharach, John Foster, Meg Gleason, Graham Hubbs, Mamoru Kaneko, Bob Stalnaker and Johan van Benthem, are also gratefully acknowledged.

Appendix A. The AGM axioms for belief revision

This appendix gives a short description of the AGM belief revision theory. A more complete account is provided by Gärdenfors (1988).

The AGM theory (so-called after Alchourrón et al., 1985) provides a set of axioms which, they argue, any reasonable belief revision system should satisfy. The axioms do not determine a unique belief revision function. An agent's epistemic state at any one point in time is represented by a set of formulas of propositional calculus, the agent's *belief set* K . This is interpreted as the set of all formulas the agent believes. Belief statements are thus

expressed in a meta-language (“inclusion in K ”). It is assumed that belief sets are closed under logical consequence.

The AGM axioms impose restrictions on the form of K_ϕ^* , the agent’s revised belief set after information ϕ is learned. Although we normally assume that belief sets are logically consistent, it is convenient to define the *absurd* belief set K_{false} in which everything is believed (i.e. K_{false} contains every formula of propositional calculus). Finally, K_ϕ^+ is used to denote the *expansion of K by ϕ* , formed by adding ϕ to K and closing under logical consequence. We can now state the axioms:

- (K^*1) K_ϕ^* is a belief set;
- (K^*2) $\phi \in K_\phi^*$;
- (K^*3) $K_\phi^* \subseteq K_\phi^+$;
- (K^*4) if $\neg\phi \notin K$, then $K_\phi^+ \subseteq K_\phi^*$;
- (K^*5) $K_\phi^* = K_{false}$ if and only if ϕ is logically inconsistent;
- (K^*6) if $\phi \Leftrightarrow \psi$, then $K_\phi^* = K_\psi^*$;
- (K^*7) $K_{\phi \wedge \psi}^* \subseteq (K_\phi^*)_\psi^+$;
- (K^*8) if $\neg\psi \notin K_\phi^*$, then $(K_\phi^*)_\psi^+ \subseteq K_{\phi \wedge \psi}^*$.

Appendix B. Proofs

Proof of Theorem 1. Soundness. The proof of soundness is straightforward, and proceeds by induction on the length of a proof of ϕ . Every element of a proof is either an axiom or follows from previous elements by the application of a rule, so we must show that every axiom is valid with respect to \mathcal{M} and that each rule is truth preserving. We consider the cases of Triv, IE(a) and LE, and leave the rest as an exercise.

Triv: we must show that $\mathcal{M} \models B_i \phi \Leftrightarrow B_i^{true} \phi$. Propositional reasoning and the definition of \models imply that $(M, w) \models true$, for all M and w . Thus $[true]_M = W$, and $W_i^w = [true]_M \cap W_i^w$. From the definition of \models again, it follows that $(M, w) \models B_i \phi$ iff $(M, w) \models B_i^{true} \phi$.

IE(a): we must show that $\mathcal{M} \models B_i^\phi \psi \Rightarrow (B_i^{\phi \wedge \psi} \chi \Leftrightarrow B_i^\phi \chi)$. Suppose $(M, w) \models B_i^\phi \psi$; then $(M, x) \models \psi$ for all $x \in \min_i^w \{[\phi]_M \cap W_i^w\}$, and $\min_i^w \{[\phi]_M \cap W_i^w\} = \min_i^w \{[\phi]_M \cap [\psi]_M \cap W_i^w\}$. From the definition of \models , we have $[\phi]_M \cap [\psi]_M = [\phi \wedge \psi]_M$; therefore $\min_i^w \{[\phi]_M \cap [\psi]_M \cap W_i^w\} = \min_i^w \{[\phi \wedge \psi]_M \cap W_i^w\}$. It follows immediately that $(M, w) \models B_i^{\phi \wedge \psi} \chi$ iff $(M, w) \models B_i^\phi \chi$, as required.

LE: we must show that if $M \models \phi \Leftrightarrow \psi$, then $M \models B_i^\phi \chi \Leftrightarrow B_i^\psi \chi$. Suppose that $M \models \phi \Leftrightarrow \psi$. Then by the definition of \models , $[\phi]_M = [\psi]_M$, and so $\min_i^w \{[\phi]_M \cap W_i^w\} = \min_i^w \{[\psi]_M \cap W_i^w\}$. It follows immediately that $(M, w) \models B_i^\phi \chi$ iff $(M, w) \models B_i^\psi \chi$, as required.

Completeness. We start with some definitions. For a given axiom system AX , we say that a formula ϕ is *AX-consistent* if $\neg\phi$ is not provable in AX . A finite set of formulas $\{\phi_1, \dots, \phi_k\}$ is *AX-consistent* exactly if $\phi_1 \wedge \dots \wedge \phi_k$ is *AX-consistent*, and an infinite set of formulas is *AX-consistent* exactly if all its finite subsets are *AX-consistent*. Finally, given

two sets of formulas S, T with $S \subseteq T \subseteq \mathcal{L}$, we say that S is a *maximal AX-consistent subset of T* if (a) it is AX-consistent, and (b) for all ϕ in T but not in S , the set $S \cup \{\phi\}$ is not AX-consistent.

Now, to prove completeness, we must show that every formula in \mathcal{L} that is valid with respect to \mathcal{M} is provable in *BRS*. It is sufficient to prove that

(*) Every *BRS*-consistent formula in \mathcal{L} is satisfiable with respect to \mathcal{M} .

For assume that we can prove (*), and that ϕ is a valid formula in \mathcal{L} . If ϕ is not provable in *BRS*, then neither is $\neg\neg\phi$, so, by definition, $\neg\phi$ is *BRS*-consistent. It follows from (*) that $\neg\phi$ is satisfiable with respect to \mathcal{M} , contradicting the validity of ϕ with respect to \mathcal{M} .

Before proceeding, we need another round of definitions. Let $Sub(\phi)$ be the set of all subformulas of ϕ ; formally, $\psi \in Sub(\phi)$ if either (a) $\psi = \phi$, or (b) ϕ is of the form $\neg\phi', \phi' \wedge \phi'', B_i\phi'$, or $B_i^{\phi'}\phi''$, and $\psi \in Sub(\phi')$ or $\psi \in Sub(\phi'')$; and let $Sub^+(\phi)$ consist of all the formulas in $Sub(\phi)$ and their negations and conjunctions, i.e. $Sub^+(\phi)$ is the smallest set such that (a) if $\psi \in Sub(\phi)$ then $\psi \in Sub^+(\phi)$; and (b) if $\psi, \chi \in Sub^+(\phi)$, then $\neg\psi, \psi \wedge \chi \in Sub^+(\phi)$. Let $Sub^{++}(\phi)$ consist of all formulas of $Sub^+(\phi)$ together with all formulas of the form $B_i\psi$ and $B_i^{\chi}\psi$, where $\psi, \chi \in Sub^+(\phi)$; and let $Sub_{neg}^{++}(\phi)$ consist of all the formulas in $Sub^{++}(\phi)$ and their negations. Finally, let $Con(\phi)$ be the set of maximal *BRS*-consistent subsets of $Sub_{neg}^{++}(\phi)$. It is easy to show¹⁹ that every *BRS*-consistent subset of $Sub_{neg}^{++}(\phi)$ can be extended to an element of $Con(\phi)$ by addition of formulas; and if S is a member of $Con(\phi)$, it must satisfy the following properties:

- for every $\psi \in Sub^{++}(\phi)$, exactly one of ψ and $\neg\psi$ is in S ;
- if $\psi \wedge \chi \in S$ then $\psi \in S$ and $\chi \in S$;
- if $\psi \vee \chi \in S$ then $\psi \in S$ or $\chi \in S$;
- if $\psi \in S$ and $\psi \Rightarrow \chi \in S$, then $\chi \in S$;
- if $\psi \Leftrightarrow \chi$ then $\psi \in S$ if and only if $\chi \in S$;
- if $\psi \in Sub_{neg}^{++}(\phi)$ and $BRS \vdash \psi$, then $\psi \in S$.

To prove (*), we construct a special structure $M_\phi \in \mathcal{M}$ for each ϕ . M_ϕ has a world w_S corresponding to every $S \in Con(\phi)$; we show that for all $\psi \in Sub(\phi)$, we have

(**) $(M_\phi, w_S) \models \psi$ if and only if $\psi \in S$,

i.e. a formula in $Sub(\phi)$ is true at world w_S exactly if it is one of the formulas in S . This is sufficient to prove (*), since if ϕ is *BRS*-consistent, it is contained in some set $S \in Con(\phi)$; it then follows from (**) that $(M_\phi, w_S) \models \phi$, and so ϕ is satisfiable with respect to \mathcal{M} as required.

¹⁹ See, e.g., Fagin et al. (1995, pp. 52, 53).

For construction of M_ϕ we introduce some new notation: let $S/B_i^\phi = \{\psi \mid B_i^\phi \psi \in S\}$, i.e. S/B_i^ϕ is the set of formulas believed by i when she learns that ϕ . We now define M_ϕ . Let $M_\phi = \langle W, \pi, \preceq \rangle$, where

$$W = \{w_S \mid S \in \text{Con}(\phi)\},$$

$$\pi(w_S)(\psi) = \begin{cases} \text{true} & \text{if } \psi \in S \\ \text{false} & \text{if } \psi \notin S \end{cases} \quad \text{for all } \psi \in \Phi,$$

$$w_T \preceq_i^{w_S} w_U \quad \text{if there is some } \psi \in \text{Sub}^+(\phi) \cap T \cap U \text{ such that } S/B_i^\psi \subseteq T.$$

We prove (**) by induction on the structure of formulas: supposing that it holds for all subformulas of $\psi \in \text{Sub}(\phi)$, we show it holds for ψ . The cases where ψ is a primitive formula, a conjunction or a negation are straightforward. Suppose ψ is of the form $B_i^\chi \zeta$. We prove the “if” direction first, and assume that $\psi \in S$. This implies that $\zeta \in S/B_i^\chi$. Consider the set $\min_i^{w_S} \{[\chi]_{M_\phi} \cap W_i^{w_S}\}$. If this set is empty, then it follows immediately from the definition of \models that $(M_\phi, w_S) \models B_i^\chi \zeta$.

Suppose then that there is some $w_T \in \min_i^{w_S} \{[\chi]_{M_\phi} \cap W_i^{w_S}\}$, i.e. $w_T \preceq_i^{w_S} w_U$ for all $w_U \in \{[\chi]_{M_\phi} \cap W_i^{w_S}\}$. Then there is some $\xi \in \text{Sub}^+(\phi) \cap T$ such that $S/B_i^\xi \subseteq T$. We must show that $\zeta \in T$. Since $S/B_i^\xi \subseteq T$, S/B_i^ξ must be a BRS-consistent set. It follows that S/B_i^χ is a BRS-consistent set too. Suppose not: then there is some finite set of formulas $F = \{\phi_1, \phi_2, \dots, \phi_k\} \subseteq S/B_i^\chi$ such that $\text{BRS} \vdash \neg(\phi_1 \wedge \phi_2 \wedge \dots \wedge \phi_k)$. Letting η denote $(\phi_1 \wedge \phi_2 \wedge \dots \wedge \phi_k)$, we have:

1. $\text{BRS} \vdash \neg\eta$ assumption
2. $\text{BRS} \vdash \eta \Rightarrow \xi$ 1, Taut, MP
3. $\text{BRS} \vdash B_i^\chi \eta \Rightarrow B_i^\chi \xi$ 2, RE, Dist, Taut, MP
4. $\text{BRS} \vdash B_i^\chi \xi \Rightarrow (B_i^{\chi \wedge \xi} \eta \Leftrightarrow B_i^\chi \eta)$ IE(a)
5. $\text{BRS} \vdash B_i^\chi \eta \Rightarrow B_i^{\chi \wedge \xi} \eta$ 3, 4, Taut, MP
6. $\text{BRS} \vdash \neg B_i^\xi \neg\chi \Rightarrow (B_i^{\xi \wedge \chi} \eta \Leftrightarrow B_i^\xi (\chi \Rightarrow \eta))$ IE(b)
7. $\text{BRS} \vdash \neg B_i^\xi \neg\chi \Rightarrow (B_i^{\chi \wedge \xi} \eta \Leftrightarrow B_i^\xi (\chi \Rightarrow \eta))$ 6, LE, Taut, MP
8. $\text{BRS} \vdash (\chi \Rightarrow \eta) \Rightarrow \neg\chi$ 1, Taut, MP
9. $\text{BRS} \vdash B_i^\xi (\chi \Rightarrow \eta) \Rightarrow B_i^\xi \neg\chi$ 8, RE, Dist, Taut, MP
10. $\text{BRS} \vdash (B_i^\chi \eta \wedge \neg B_i^\xi \neg\chi) \Rightarrow B_i^\xi \eta$ 5, 7, 9, Taut, MP
11. $\text{BRS} \vdash (B_i^\chi \phi_1 \wedge \dots \wedge B_i^\chi \phi_k \wedge \neg B_i^\xi \neg\chi)$
 $\Rightarrow (B_i^\xi \phi_1 \wedge \dots \wedge B_i^\xi \phi_k)$ 10, Dist, Taut, MP

By the hypothesis of induction, we know that $\chi \in T$, since $w_T \in [\chi]_{M_\phi}$. It follows that $\neg B_i^\xi \neg\chi \in S$ since $B_i^\xi \neg\chi \in \text{Sub}^{++}(\phi)$. Since $B_i^\chi \phi_1, \dots, B_i^\chi \phi_k \in S$, we have $\neg B_i^\xi \phi_1, \dots, \neg B_i^\xi \phi_k \notin S$, or else S would be inconsistent according to line 11 above. Since $B_i^\xi \phi_1, \dots, B_i^\xi \phi_k \in \text{Sub}^{++}(\phi)$, it follows that $B_i^\xi \phi_1, \dots, B_i^\xi \phi_k \in S$. Thus $F \subseteq S/B_i^\xi$, i.e. S/B_i^ξ is not a BRS-consistent set, contradicting our original assumption.

So S/B_i^χ is a BRS-consistent set, and it therefore has a maximal BRS-consistent extension, U . And since $B_i^\chi \chi \in S$ (single instance of Succ), we have $\chi \in (S/B_i^\chi) \subseteq U$.

Thus $w_U \preceq_i^{w_S} w_T$, by construction; and since $w_T \in \min_i^{w_S} \{[\chi]_{M_\phi} \cap W_i^{w_S}\}$, $w_T \preceq_i^{w_S} w_U$. So there is some $\rho \in \text{Sub}^+(\phi) \cap T \cap U$ such that $S/B_i^\rho \subseteq T$.

We started off by assuming that $B_i^\chi \zeta \in S$; we also know that $\neg B_i^\chi \neg \rho \in S$, since $\rho \in U$ and $S/B_i^\chi \subseteq U$; and that $\neg B_i^\rho \neg \chi \in S$, since $\chi \in T$ and $S/B_i^\rho \subseteq T$. Furthermore,

12. $BRS \vdash \neg B_i^\chi \neg \rho \Rightarrow (B_i^{\chi \wedge \rho} \zeta \Leftrightarrow (B_i^\chi \zeta \vee B_i^\rho (\rho \Rightarrow \zeta)))$ IE(b)
13. $BRS \vdash (\neg B_i^\chi \neg \rho \wedge B_i^\chi \zeta) \Rightarrow B_i^{\chi \wedge \rho} \zeta$ 12, Taut, MP
14. $BRS \vdash (\neg B_i^\chi \neg \rho \wedge B_i^\chi \zeta) \Rightarrow B_i^{\rho \wedge \chi} \zeta$ 13, LE, Taut, MP
15. $BRS \vdash \neg B_i^\rho \neg \chi \Rightarrow (B_i^{\rho \wedge \chi} \zeta \Leftrightarrow (B_i^\rho \zeta \vee B_i^\rho (\chi \Rightarrow \zeta)))$ IE(b)
16. $BRS \vdash (\neg B_i^\chi \neg \rho \wedge B_i^\chi \zeta \wedge \neg B_i^\rho \neg \chi) \Rightarrow (B_i^\rho \zeta \vee B_i^\rho (\chi \Rightarrow \zeta))$ 14, 15, Taut, MP

Since $\chi, \zeta, \rho \in \text{Sub}^+(\phi)$, $B_i^\rho \zeta$ and $B_i^\rho (\chi \Rightarrow \zeta)$ are both in $\text{Sub}^{++}(\phi)$. So either $B_i^\rho \zeta \in S$ or $B_i^\rho (\chi \Rightarrow \zeta) \in S$: if $\neg B_i^\rho \zeta \in S$ and $\neg B_i^\rho (\chi \Rightarrow \zeta) \in S$, S would be inconsistent according to line 16 above. But $S/B_i^\rho \subseteq T$, so either $\zeta \in T$, or $\chi \Rightarrow \zeta \in T$, in which case $\zeta \in T$ again because $\chi \in T$.

Thus for every $w_T \in \min_i^{w_S} \{[\chi]_{M_\phi} \cap W_i^{w_S}\}$, $\zeta \in T$, and so $(M_\phi, w_T) \models \zeta$ by the hypothesis of induction. It follows from the definition of \models that $(M_\phi, w_S) \models B_i^\chi \zeta$.

For the “only if” direction, assume $(M_\phi, w_S) \models B_i^\chi \zeta$. It follows that the set $(S/B_i^\chi) \cup \{\neg \zeta\}$ is not *BRS*-consistent. If it were, it would have a maximal *BRS*-consistent extension T . Now, $B_i^\chi \chi \in S$ (single instance of Succ), and so $\chi \in (S/B_i^\chi) \subseteq T$. Thus by construction, $w_T \preceq_i^{w_S} w_U$ for all U such that $\chi \in U$, and (given the hypothesis of induction), $w_T \in \min_i^{w_S} \{[\chi]_{M_\phi} \cap W_i^{w_S}\}$. The hypothesis of induction also tells us that $(M_\phi, w_T) \models \neg \zeta$, since $\neg \zeta \in T$, and it follows immediately from the definition of \models that $(M_\phi, w_S) \models \neg B_i^\chi \zeta$, contradicting our original assumption. So $(S/B_i^\chi) \cup \{\neg \zeta\}$ is not *BRS*-consistent. It follows from Taut, Dist, MP and RE that $B_i^\chi \zeta \in S$, as desired (see Fagin et al., 1995, p. 55).

So we have proved the inductive step for the case where ψ is of the form $B_i^\chi \zeta$. The case where ϕ is of the form $B_i \chi$ follows quickly. Pick any instance of *true* $\in \text{Sub}^+(\phi)$. Since $BRS \vdash B_i \chi \Leftrightarrow B_i^{\text{true}} \chi$ (instance of Triv), $B_i \chi \in S$ if and only if $B_i^{\text{true}} \chi \in S$. Furthermore, since *true* is a propositional tautology, it follows from the definition of \models that $(M_\phi, w_T) \models \text{true}$ for all w_T , i.e. $[\text{true}] = W$. Again from the definition of \models , we have $(M_\phi, w_T) \models B_i \chi$ if and only if $(M_\phi, w_T) \models B_i^{\text{true}} \chi$.

We have shown that (**) holds for all formulas $\psi \in \text{Sub}(\phi)$. To complete the proof of completeness, we need to show that $M_\phi \in \mathcal{M}$, i.e. that M_ϕ really is a belief revision structure.

It is clear that W is a well-defined set of possible worlds, and π is an interpretation. We need to show that for all S , $\preceq_i^{w_S}$ is complete and transitive on $W_i^{w_S}$, and is well-founded. For completeness, assume that $w_T, w_U \in W_i^{w_S}$. We need to show that either $w_T \preceq_i^{w_S} w_U$ or $w_U \preceq_i^{w_S} w_T$. From the definitions of W_i^w and $\preceq_i^{w_S}$, there is some $\psi \in \text{Sub}^+(\phi) \cap T$ such that $S/B_i^\psi \subseteq T$, and some $\chi \in \text{Sub}^+(\phi) \cap U$ such that $S/B_i^\chi \subseteq U$. Since S is a maximal *BRS*-consistent set, either (a) $B_i^{\psi \vee \chi} \neg \psi \in S$ or (b) $\neg B_i^{\psi \vee \chi} \neg \psi \in S$. We consider each case in turn. First (a) $B_i^{\psi \vee \chi} \neg \psi \in S$. We know that $B_i^{\psi \vee \chi} (\psi \vee \chi) \in S$ (instance of Succ) and

17. $BRS \vdash (B_i^{\psi \vee \chi} \neg \psi \wedge B_i^{\psi \vee \chi} (\psi \vee \chi)) \Rightarrow B_i^{\psi \vee \chi} \chi$ Dist
18. $BRS \vdash B_i^{\psi \vee \chi} \chi \Rightarrow (B_i^{(\psi \vee \chi) \wedge \chi} \zeta \Leftrightarrow B_i^{\psi \vee \chi} \zeta)$ IE(a)
19. $BRS \vdash ((\psi \vee \chi) \wedge \chi) \Leftrightarrow \chi$ Taut
20. $BRS \vdash B_i^{(\psi \vee \chi) \wedge \chi} \zeta \Leftrightarrow B_i^{\chi} \zeta$ 19, LE

Line 17 implies that $B_i^{\psi \vee \chi} \chi \in S$. So it follows from line 18 that $B_i^{(\psi \vee \chi) \wedge \chi} \zeta \in S$ if and only if $B_i^{\psi \vee \chi} \zeta \in S$, i.e. $S/B_i^{(\psi \vee \chi) \wedge \chi} = S/B_i^{\psi \vee \chi}$. Furthermore, line 20 implies that $S/B_i^{(\psi \vee \chi) \wedge \chi} = S/B_i^{\chi}$. So $S/B_i^{\chi} = S/B_i^{\psi \vee \chi} \subseteq U$. But $\psi \vee \chi \in T \cap U$, so $w_U \preceq_i^{w_S} w_T$.

Next (b) $\neg B_i^{\psi \vee \chi} \neg \psi \in S$. Since

21. $BRS \vdash \neg B_i^{\psi \vee \chi} \neg \psi \Rightarrow (B_i^{(\psi \vee \chi) \wedge \psi} \zeta \Leftrightarrow B_i^{\psi \vee \chi} (\chi \Rightarrow \zeta))$ IE(b)
22. $BRS \vdash B_i^{\psi \vee \chi} \zeta \Rightarrow B_i^{\psi \vee \chi} (\chi \Rightarrow \zeta)$ RE, Dist, Taut, MP

if $B_i^{\psi \vee \chi} \zeta \in S$ then $B_i^{\psi \vee \chi} (\chi \Rightarrow \zeta) \in S$ (from line 22) and $B_i^{(\psi \vee \chi) \wedge \psi} \zeta \in S$ (from line 21), so $S/B_i^{\psi \vee \chi} \subseteq S/B_i^{(\psi \vee \chi) \wedge \psi} = S/B_i^{\psi} \subseteq T$. But $\psi \vee \chi \in T \cap U$, so $w_T \preceq_i^{w_S} w_U$.

To show transitivity, suppose that for some $w_T, w_U, w_V \in W_i^{w_S}$, $w_T \preceq_i^{w_S} w_U$ and $w_U \preceq_i^{w_S} w_V$. Then there is some $\psi \in T \cap U$ such that $S/B_i^{\psi} \subseteq T$, and there is some $\chi \in U \cap V$ such that $S/B_i^{\chi} \subseteq U$. Since S is a maximal *BRS*-consistent set, either (a) $B_i^{\psi \vee \chi} \neg \psi \in S$ or (b) $\neg B_i^{\psi \vee \chi} \neg \psi \in S$. We consider each case in turn.

(a) $B_i^{\psi \vee \chi} \neg \psi \in S$. We have just shown in part (a) above that this implies $S/B_i^{\chi} = S/B_i^{\psi \vee \chi} \subseteq U$. But $\psi \in U$, contradicting the assumption that $B_i^{\psi \vee \chi} \neg \psi \in S$.

(b) $\neg B_i^{\psi \vee \chi} \neg \psi \in S$. We have just shown in part (b) above that this implies $S/B_i^{\psi \vee \chi} \subseteq S/B_i^{\psi} \subseteq T$. But $\psi \vee \chi \in T \cap U \cap V$, so $w_T \preceq_i^{w_S} w_V$ as required. Note that nothing in this proof makes use of the fact that $w_V \in W_i^w$: it follows that $\preceq_i^{w_S}$ is in fact transitive on the entire domain W .

Well-foundedness of $\preceq_i^{w_S}$ follows immediately from the finiteness of W . To show that W is finite, we must show that $Con(\phi)$ is finite, since $\|W\| = \|Con(\phi)\|$. It is clear that there is a finite number maximal *BRS*-consistent subsets of $Sub(\phi)$, since $Sub(\phi)$ is itself a finite set. And each maximal *BRS*-consistent subset of $Sub(\phi)$ has a unique extension to a maximal *BRS*-consistent subset of $Sub^+(\phi)$: suppose S is a maximal *BRS*-consistent subset of $Sub(\phi)$, and $S^+ \supseteq S$ is a maximal *BRS*-consistent subset of $Sub^+(\phi)$; then $\psi \in S^+$ only if either (a) $\psi \in S$, or (b) ψ is of the form $\neg \chi$ and $\chi \notin S^+$, or (c) ψ is of the form $\chi \wedge \zeta$ and $\chi, \zeta \in S^+$. And if there is no maximal *BRS*-consistent subset S of $Sub(\phi)$ such that $S \subseteq S^+$, S^+ cannot be a maximal *BRS*-consistent subset of $Sub^+(\phi)$. Thus there is a one-to-one correspondence between the maximal *BRS*-subsets of $Sub(\phi)$ and the maximal *BRS*-consistent subsets of $Sub^+(\phi)$. Propositional reasoning tells us that there are at most $2^{\|Sub(\phi)\|}$ logically distinct formulas in $Sub^+(\phi)$, corresponding to the combinations of truth-value assignments to the formulas in $Sub(\phi)$. And if ψ and χ are logically equivalent (i.e. $BRS \vdash \psi \Leftrightarrow \chi$), and S_{neg}^{++} is a maximal *BRS*-consistent subset of $Sub_{neg}^{++}(\phi)$, then $B_i \psi \in S_{neg}^{++}$ if and only if $B_i \chi \in S_{neg}^{++}$ (Triv, RE and Dist); $B_i^{\zeta} \psi \in S_{neg}^{++}$ if and only if $B_i^{\zeta} \chi \in S_{neg}^{++}$ (RE and Dist); and $B_i^{\psi} \zeta \in S_{neg}^{++}$ if and only if $B_i^{\chi} \zeta \in S_{neg}^{++}$. So each maximal

BRS-consistent subset of $Sub^+(\phi)$ has a finite number of extensions to maximal BRS-consistent subsets of $Sub_{neg}^{++}(\phi)$. Thus $Con(\phi)$ is a finite set as required. \square

Proof of Proposition 1. The proof of Proposition 1 follows the same steps as the proof of (**) in Theorem 1, and is not repeated here. \square

Proof of Theorem 2. Soundness. We must check that TPI and TNI are valid with respect to \mathcal{A} , i.e. $\mathcal{A} \models B_i^\phi \psi \Rightarrow B_i^\chi B_i^\phi \psi$ and $\mathcal{A} \models \neg B_i^\phi \psi \Rightarrow B_i^\chi \neg B_i^\phi \psi$.

(TPI) Suppose that $M \in \mathcal{A}$, and $(M, w) \models B_i^\phi \psi$. Then for every $x \in \min_i^w \{[\phi]_M \cap W_i^w\}$ we have $(M, x) \models \psi$. Now for every $y \in \min_i^w \{[\chi]_M \cap W_i^w\}$, $z \preceq_i^y u$ if and only if $z \preceq_i^w u$, so $\min_i^w \{[\phi]_M \cap W_i^w\} = \min_i^y \{[\phi]_M \cap W_i^y\}$. It follows that $(M, y) \models B_i^\phi \psi$, and therefore $(M, w) \models B_i^\chi B_i^\phi \psi$.

(TNI) Suppose that $M \in \mathcal{A}$, and $(M, w) \models \neg B_i^\phi \psi$. Then there is some $x \in \min_i^w \{[\phi]_M \cap W_i^w\}$ such that $(M, x) \not\models \psi$. Now for every $y \in \min_i^w \{[\chi]_M \cap W_i^w\}$, $z \preceq_i^y u$ if and only if $z \preceq_i^w u$, so $\min_i^w \{[\phi]_M \cap W_i^w\} = \min_i^y \{[\phi]_M \cap W_i^y\}$. It follows that $(M, y) \models \neg B_i^\phi \psi$, and therefore $(M, w) \models B_i^\chi \neg B_i^\phi \psi$.

Completeness. To prove completeness, we must show that every BRSI-consistent formula in \mathcal{L} is satisfiable with respect to \mathcal{A} . We proceed in the same way as in the proof of Theorem 1, and construct a structure $M_\phi \in \mathcal{A}$ for every formula $\phi \in \mathcal{L}$. The construction of M_ϕ is exactly the same as before, except that the set $Sub^{++}(\phi)$ is enlarged: $\psi \in Sub^{++}(\phi)$ is the smallest set of formulas such that (a) if $\psi, \chi \in Sub^+(\phi)$ then $\psi, B_i \psi, B_i^\chi \psi \in Sub^{++}(\phi)$; (b) if $\xi \in Sub^+(\phi)$ and $B_i^\chi \psi \in Sub^{++}(\phi)$, then $B_i^\xi B_i^\chi \psi \in Sub^{++}(\phi)$ and $B_i^\xi \neg B_i^\chi \psi \in Sub^{++}(\phi)$. The proof that $(M_\phi, w_S) \models \psi$ if and only if $\psi \in S$ is unaffected, as is the proof that $\preceq_i^{w_S}$ is complete and transitive on $W_i^{w_S}$ for all S . Finiteness of W still holds as well, since for every formula of the form $B_i^\xi B_i^\chi \dots B_i^\chi \psi \in S$ if and only if $B_i^\xi \dots B_i^\chi \psi \in S$, given TPI and TNI. It remains to show that M is absolute.

Suppose that $w_T \in W_i^{w_S}$. Then there is some $\psi \in Sub^+(\phi) \cap T$ such that $S/B_i^\psi \subseteq T$. We must show that $w_U \preceq_i^{w_S} w_V$ if and only if $w_U \preceq_i^{w_T} w_V$. If $w_U \preceq_i^{w_S} w_V$, then there is some $\chi \in Sub^+(\phi) \cap U \cap V$ such that $S/B_i^\chi \subseteq U$. Suppose that $B_i^\chi \zeta \notin S$. Then $\neg B_i^\chi \zeta \in S$, and since $BRSI \vdash \neg B_i^\chi \zeta \Rightarrow B_i^\psi \neg B_i^\chi \zeta$ (instance of TNI), we also have $B_i^\psi \neg B_i^\chi \zeta \in S$. But $S/B_i^\psi \subseteq T$, so $\neg B_i^\chi \zeta \in T$, and $B_i^\chi \zeta \notin T$. Thus $T/B_i^\chi \subseteq S/B_i^\chi \subseteq U$, and $w_U \preceq_i^{w_T} w_V$, as required. If $w_U \not\preceq_i^{w_T} w_V$, then there is some $\chi \in Sub^+(\phi) \cap U \cap V$ such that $T/B_i^\chi \subseteq U$. Suppose that $B_i^\chi \zeta \in S$. Since $BRSI \vdash B_i^\chi \zeta \Rightarrow B_i^\psi B_i^\chi \zeta$ (instance of TPI), we also have $B_i^\psi B_i^\chi \zeta \in S$. But $S/B_i^\psi \subseteq T$, so $B_i^\chi \zeta \in T$. Thus $S/B_i^\chi \subseteq T/B_i^\chi \subseteq U$, and $w_U \preceq_i^{w_S} w_V$, completing the proof. \square

Proof of Theorem 3. Soundness. We must check that WCon is valid with respect to \mathcal{I} , i.e. $\mathcal{I} \models \phi \Rightarrow \neg B_i^\phi$ false. Suppose that $M \in \mathcal{I}$, and $(M, w) \models \phi$, i.e. $w \in [\phi]_M$. By the inclusion assumption, $w \in W_i^w$, so $[\phi]_M \cap W_i^w \neq \emptyset$, and by well-foundedness of \preceq_i^w , $\min_i^w \{[\phi]_M \cap W_i^w\} \neq \emptyset$. So there is some world $x \in \min_i^w \{[\phi]_M \cap W_i^w\}$. Since it is not

the case that $(M, x) \models \text{false}$, it is not the case that $(M, w) \models B_i^\phi \text{false}$, and it follows from the definition of \models that $(M, w) \models \neg B_i^\phi \text{false}$.

Completeness. The proof for completeness is the same as for Theorem 1, except that it must also be shown $M_\phi \in \mathcal{I}$, i.e. that every plausibility ordering \preceq_i^{wS} in M_ϕ satisfies the inclusion assumption. Let ϕ_1, \dots, ϕ_n be an enumeration of the formulas in $\text{Sub}(\phi)$, and for some $S \in \text{Con}(\phi)$ let $\phi'_1 = \phi_1$ if $\phi_1 \in S$, and $\phi'_1 = \neg\phi_1$ if $\phi_1 \notin S$. Propositional reasoning implies that $\phi'_1 \wedge \dots \wedge \phi'_n \in S$, and since $\text{BRSC} \vdash (\phi'_1 \wedge \dots \wedge \phi'_n) \Rightarrow \neg B_i^{\phi'_1 \wedge \dots \wedge \phi'_n} \text{false}$ (instance of WCon), we have $\neg B_i^{\phi'_1 \wedge \dots \wedge \phi'_n} \text{false} \in S$ for some $\text{false} \in \text{Sub}^+(\phi)$. It follows that $S/B_i^{\phi'_1 \wedge \dots \wedge \phi'_n}$ is a BRSC-consistent set. Furthermore, since $\text{BRS} \vdash B_i^{\phi'_1 \wedge \dots \wedge \phi'_n} \phi'_1 \wedge \dots \wedge \phi'_n$ (instance of Succ), $\phi'_1 \wedge \dots \wedge \phi'_n \in S/B_i^{\phi'_1 \wedge \dots \wedge \phi'_n}$. Now suppose that $B_i^{\phi'_1 \wedge \dots \wedge \phi'_n} \psi \in S$. Then $\psi \in S/B_i^{\phi'_1 \wedge \dots \wedge \phi'_n}$ and $\psi \in \text{Sub}^+(\phi)$. Since $\text{Sub}^+(\phi)$ consists only of formulas in $\text{Sub}(\phi)$ and their conjunctions and negations, it follows from propositional reasoning that either $\text{BRSC} \vdash (\phi'_1 \wedge \dots \wedge \phi'_n) \Rightarrow \psi$ or $\text{BRSC} \vdash (\phi'_1 \wedge \dots \wedge \phi'_n) \Rightarrow \neg\psi$. But $S/B_i^{\phi'_1 \wedge \dots \wedge \phi'_n}$ is a BRSC-consistent set, so we must have $\text{BRSC} \vdash (\phi'_1 \wedge \dots \wedge \phi'_n) \Rightarrow \psi$, and therefore $\psi \in S$. We have shown that $S/B_i^{\phi'_1 \wedge \dots \wedge \phi'_n} \subseteq S$. The definition of \preceq_i^{wS} implies that $w_S \preceq_i^{wS} w_S$, and so $w_S \in W_i^{wS}$ as required. \square

Proof of Theorem 4. Soundness follows immediately from Theorem 2 and Theorem 3. To prove completeness, we follow the construction of M_ϕ described in the proof of Theorem 2, and the same steps imply that $M_\phi \in \mathcal{A}$. The completeness part of the proof of Theorem 3 can then be followed to show that $M_\phi \in \mathcal{I}$. \square

Proof of Proposition 2. First suppose that $x \in \{x \mid x \triangleleft_i w \text{ or } w \triangleleft_i x\}$. Then either (a) $x \triangleleft_i w$, in which case it follows immediately from the definition of \triangleleft_i that $x \preceq_i^w w$, as required; or (b) $w \triangleleft_i x$; from the definition of \triangleleft_i we have $w \preceq_i^x x$; $x \preceq_i^x y$ for some y from R4, and so $x \preceq_i^w y$ from R3, as required.

Now suppose that $x \in \{x \mid x \preceq_i^w y \text{ for some } y\}$. We know from R4 that $w \preceq_i^w z$ for some z . It follows from R1 that either $x \preceq_i^w w$, in which case $x \triangleleft_i w$ by definition; or $w \preceq_i^w x$, in which case R3 gives us $w \preceq_i^x x$ as so $w \triangleleft_i x$. In both cases $x \in \{x \mid x \triangleleft_i w \text{ or } w \triangleleft_i x\}$ as required. \square

Proof of Proposition 3. First suppose $x \in \min_i(X \cap W_i^w)$. Then $x \triangleleft_i y$ for all $y \in X \cap W_i^w$, and from the definition of \triangleleft_i , $x \preceq_i^y y$ for all $y \in X \cap W_i^w$. It follows from R3 that $x \preceq_i^w y$ for all $y \in X \cap W_i^w$ and so $x \in \min_i^w(X \cap W_i^w)$.

Now suppose that $x \in \min_i^w(X \cap W_i^w)$. Then $x \preceq_i^w y$ for all $y \in X \cap W_i^w$, and from R3 we have $x \preceq_i^y y$ for all $y \in X \cap W_i^w$. The definition of \triangleleft_i gives us $x \triangleleft_i y$ for all $y \in X \cap W_i^w$, and so $x \in \min_i(X \cap W_i^w)$. \square

Proof of Theorem 5. The proof of Theorem 5 follows the same steps as the proof of Theorem 3.3.1 in Fagin et al. (1995), and is omitted. \square

References

- Alchourrón, C.E., Gärdenfors, P., Makinson, D., 1985. On the logic of theory change: partial meet functions for contraction and revision. *J. Symbolic Logic* 50, 510–530.
- Aumann, R.J., 1976. Agreeing to disagree. *Ann. Statist.* 4, 1236–1239.
- Aumann, R.J., 1999. Interactive epistemology I: knowledge. *Int. J. Game Theory* 28, 263–300.
- Aumann, R.J., Brandenburger, A., 1995. Epistemic conditions for Nash equilibrium. *Econometrica* 63, 1161–1180.
- Battigalli, P., 1996. Strategic independence and perfect Bayesian equilibria. *J. Econ. Theory* 70, 201–234.
- Battigalli, P., Bonanno, G., 1999. Recent results on belief, knowledge and the epistemic foundations of game theory. *Res. Econ.* 53, 149–225.
- Battigalli, P., Siniscalchi, M., 1999. Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *J. Econ. Theory* 88, 188–230.
- Battigalli, P., Siniscalchi, M., 2002. Strong belief and forward induction reasoning. *J. Econ. Theory* 106, 356–391.
- Ben-Porath, E., 1997. Rationality, Nash equilibrium and backwards induction in perfect-information games. *Rev. Econ. Stud.* 64, 23–46.
- Board, O.J., 2002a. Algorithmic characterization of rationalizability in extensive form games. Working paper. Department of Economics, Oxford.
- Board, O.J., 2002b. The deception of the Greeks: generalizing the information structure of extensive form games. *Greek Econ. Rev.* 22, 1–16.
- Board, O.J., 2003. The not-so-absent-minded driver. *Res. Econ.* 57, 189–200.
- Bonanno, G., 2003. Memory and perfect recall in extensive games. *Games Econ. Behav.* In press.
- Brandenburger, A., 2002. On the existence of a ‘complete’ possibility structure. Working paper. Harvard Business School.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* 59, 189–198.
- Brandenburger, A., Keisler, H.J., 1999. An impossibility theorem on beliefs in games. Working paper. Harvard Business School.
- Brandenburger, A., Keisler, H.J., 2002. Epistemic conditions for iterated admissibility. Working paper. Harvard Business School.
- Burgess, J.P., 1981. Quick completeness proofs for some logics of conditionals. *Notre Dame J. Formal Logic* 22, 76–84.
- Cho, I.K., Kreps, D.M., 1987. Signalling games and stable equilibria. *Quart. J. Econ.* 102, 179–221.
- Clausing, T., 2001. Belief revision in games of perfect information. Working paper. Univ. of Magdeburg.
- Crawford, V., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50, 1431–1451.
- Dekel, E., Gul, F., 1997. Rationality and knowledge in game theory. In: Kreps, D.M., Wallis, K.W. (Eds.), *Advances in Economics and Econometrics: Theory and Applications: Seventh World Congress, Volume I*. Cambridge Univ. Press, Cambridge, UK, pp. 87–172.
- Dufwenberg, M., 2002. Marital investments, time consistency and emotions. *J. Econ. Behav. Organ.* 48, 57–69.
- Fagin, R., 1994. A quantitative analysis of modal logic. *J. Symbolic Logic* 59, 209–252.
- Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y., 1995. *Reasoning About Knowledge*. MIT Press, Cambridge, MA.
- Fagin, R., Geanakoplos, J., Halpern, J.Y., Vardi, M.Y., 1999. The hierarchical approach to modeling knowledge and common knowledge. *Int. J. Game Theory* 28, 331–365.
- Feinberg, Y., 2001. Subjective formulation and analysis of games and solutions—Part 2. Working paper. Stanford University.
- Feinberg, Y., 2002. Subjective reasoning in dynamic games. Working paper. Stanford University.
- Friedman, N., Halpern, J.Y., 1994. Conditional logics of belief change. In: *Proceedings of the Twelfth National Conference on Artificial Intelligence*. AAAI Press, Menlo Park, CA, pp. 915–921.
- Gärdenfors, P., 1986. Belief revisions and the Ramsey test for conditionals. *Philos. Rev.* XCV, 81–93.
- Gärdenfors, P., 1988. *Knowledge in Flux*. MIT Press, Cambridge, MA.
- Geanakoplos, J.D., Pearce, D., Stacchetti, E., 1989. Psychological games and sequential rationality. *Games Econ. Behav.* 1, 60–79.
- Geanakoplos, J.D., 1992. Common knowledge. *J. Econ. Perspect.* 6, 53–82.
- Geanakoplos, J.D., Polemarchakis, H.M., 1982. We can’t disagree forever. *J. Econ. Theory* 28, 192–200.

- Grove, A., 1988. Two modellings for theory change. *J. Philos. Logic* 17, 157–170.
- Halpern, J.Y., 1998. Set-theoretic completeness for epistemic and conditional logic. *Ann. Math. Artificial Intelligence* 26, 1–27.
- Harsanyi, J., 1968. Games with incomplete information played by ‘Bayesian’ players, parts I–III. *Manage. Sci.* 14, 159–182, 320–334, 486–502.
- Heifetz, A., Samet, D., 1998. Knowledge spaces with arbitrarily high rank. *Games Econ. Behav.* 22, 260–273.
- Hintikka, J., 1962. *Knowledge and Belief*. Cornell Univ. Press, Ithaca, NY.
- Kohlberg, E., Reny, P.J., 1997. Independence on relative probability spaces and consistent assessments in game trees. *J. Econ. Theory* 75, 280–313.
- Kripke, S., 1963. A semantical analysis of modal logic I: normal modal propositional calculi. *Z. Math. Logik Grund. Math.* 9, 67–96.
- Levi, I., 1988. Iterations of conditionals and the Ramsey test. *Synthese* 76, 49–81.
- Lewis, D., 1973. *Counterfactuals*. Basil Blackwell, Oxford.
- Lewis, D., 1975. Languages and Language. In: Gunderson, K. (Ed.), *Minnesota Studies in the Philosophy of Science*, vol. VII. Univ. of Minnesota Press, pp. 3–35.
- Mariotti, T., Meier, M., Piccioni, M., 2002. Hierarchies of beliefs for compact possibility models. Working paper. London School of Economics.
- Mertens, J.-F., Zamir, S., 1985. Formulation of Bayesian analysis for games of incomplete information. *Int. J. Game Theory* 14, 1–29.
- McDowell, J., 1982. Criteria, defeasibility, and knowledge. *Proc. Br. Acad.* 68, 455–479.
- Morris, S., Shin, H.S., 2003. Global games: theory and applications. In: Dewatripont, M., Hansen, L.P., Turnovsky, S.J. (Eds.), *Advances in Economics and Econometrics: Theory and Applications: Eighth World Congress*, vol. 1. Cambridge Univ. Press, Cambridge, UK, pp. 56–114.
- Pappas, G., Swain, M. (Eds.), 1978. *Essays on Knowledge and Justification*. Cornell Univ. Press, Ithaca, NY.
- Piccione, M., Rubinstein, A., 1997. On the interpretation of decision problems with imperfect recall. *Games Econ. Behav.* 20, 3–24.
- Reny, P.J., 1993. Common belief and the theory of games with perfect information. *J. Econ. Theory* 59, 257–274.
- Samet, D., 1990. Ignoring ignorance and agreeing to disagree. *J. Econ. Theory* 52, 190–207.
- Spohn, W., 1988. Ordinal conditional functions: a dynamic theory of epistemic states. In: Harper, W.L., Skyrms, B. (Eds.), *Causation in Decision, Belief Change, and Statistics*, vol. II. Kluwer, Dordrecht, pp. 105–134.
- Stalnaker, R., 1996. Knowledge, belief and counterfactual reasoning in games. *Econ. Philos.* 12, 133–163.
- Stalnaker, R., 1998. Belief revision in games: forward and backward induction. *Math. Soc. Sci.* 36, 31–56.