

A logic for reasoning about ambiguity

Joseph Y. Halpern
Computer Science Dept.
Cornell University
Ithaca, NY

E-mail: halpern@cs.cornell.edu

Willemien Kets
Kellogg School of Management
Northwestern University
Evanston, IL

E-mail: w-kets@kellogg.northwestern.edu

December 20, 2013

Abstract

Standard models of multi-agent modal logic do not capture the fact that information is often *ambiguous*, and may be interpreted in different ways by different agents. We propose a framework that can model this, and consider different semantics that capture different assumptions about the agents' beliefs regarding whether or not there is ambiguity. We examine the expressive power of logics of ambiguity compared to logics that cannot model ambiguity, with respect to the different semantics that we propose.

1 Introduction

In the study of multi-agent modal logics, it is implicitly assumed that all agents interpret all formulas the same way. While they may have different beliefs regarding whether a formula φ is true, they agree on what φ means. Formally, this is captured by the fact that the truth of φ does not depend on the agent.

Of course, in the real world, there is *ambiguity*; different agents may interpret the same utterance in different ways. For example, consider a public announcement p . Each player i may interpret p as corresponding to some event E_i , where E_i may be different from E_j if $i \neq j$. This seems natural: even if people have a common background, they may still disagree on how

to interpret certain phenomena or new information. Someone may interpret a smile as just a sign of friendliness; someone else may interpret it as a “false” smile, concealing contempt; yet another person may interpret it as a sign of sexual interest.

To model this formally, we can use a straightforward approach already used in [Halpern 2009; Grove and Halpern 1993]: formulas are interpreted relative to a player. But once we allow such ambiguity, further subtleties arise. Returning to the announcement p , not only can it be interpreted differently by different players, it may not even occur to the players that others may interpret the announcement in a different way. Thus, for example, i may believe that E_i is common knowledge. The assumption that each player believes that her interpretation is how everyone interprets the announcement is but one assumption we can make about ambiguity. It is also possible that player i may be aware that there is more than one interpretation of p , but believes that player j is aware of only one interpretation. For example, think of a politician making an ambiguous statement which he realizes that different constituencies will interpret differently, but will not realize that there are other possible interpretations. In this paper, we investigate a number of different semantics of ambiguity that correspond to some standard assumptions that people make with regard to ambiguous statements, and investigate their relationship.

Our interest in ambiguity was originally motivated by a seminal result in game theory: Aumann’s [1976] theorem showing that players cannot “agree to disagree.” More precisely, this theorem says that agents with a common prior on a state space cannot have common knowledge that they have different posteriors. This result has been viewed as paradoxical in the economics literature. Trade in a stock market seems to require common knowledge of disagreement (about the value of the stock being traded), yet we clearly observe a great deal of trading. One well known explanation for the disagreement is that we do not in fact have common priors: agents start out with different beliefs. In a companion paper [Halpern and Kets 2013], we provide a different explanation, in terms of ambiguity. It is easy to show that we can agree to disagree when there is ambiguity, even if there is a common prior.

Although our work is motivated by applications in economics, ambiguity has long been a concern in philosophy, linguistics, and natural language processing. For example, there has been a great deal of work on word-sense disambiguation (i.e., trying to decide from context which of the multiple meanings of a word are intended); see Hirst [1988] for a seminal contribution, and Navigli [2009] for a recent survey. However, there does not seem to be much work on incorporating ambiguity into a logic. Apart from the literature on the logic of context and on underspecification (see Van Deemter and Peters [1996]), the only papers that we are aware of that does this are ones by Monz [1999] and Kuijer [2013]. Monz allows for statements that have multiple interpretations, just as we do. But rather than incorporating the ambiguity directly into the logic, he considers updates by ambiguous statements.

Kuijer models the fact that ambiguous statements can have multiple meanings by using a nondeterministic propositional logic, which, roughly speaking allows him to consider all the meanings simultaneously. He then defines a notion of implication such that an ambiguous statement A entails another ambiguous statement B if and only if every possible interpretations of A

entails every possible interpretation of B . This idea of considering all possible interpretations of an ambiguous statement actually has a long tradition in the philosophy literature. For example, Lewis [1982] considers assigning truth values to an ambiguous formula ϕ by considering all possible disambiguations of ϕ . This leads to a semantics where a formula can, for example, have the truth value $\{\text{true}, \text{false}\}$. Lewis views this as a potential justification for *relevance logic* (a logic where a formula can be true, false, both, or neither; cf. [Rescher and Brandom 1979]). Our approach is somewhat different. We assume that each agent uses only one interpretation of a given ambiguous formula ϕ , but an agent may consider it possible that another agent interprets ϕ differently. In our applications, this seems to be the most appropriate way to dealing with ambiguity (especially when it comes to considering the strategic implications of ambiguity).

There are also connections between ambiguity and vagueness. Although the two notions are different—a term is *vague* if it is not clear what its meaning is, and is *ambiguous* if it can have multiple meanings, Halpern [2009] also used agent-dependent interpretations in his model of vagueness, although the issues that arose were quite different from those that concern us here.

Given the widespread interest in ambiguity, in this paper we focus on the logic of ambiguity. We introduce the logic in Section 2. The rest of the paper is devoted to arguing that, in some sense, ambiguity is not necessary. In Section 3, we show that a formula is satisfiable in a structure with ambiguity (i.e., one where different agents interpret formulas differently) if and only if it is satisfiable in a structure without ambiguity. Then in Section 4, we show that, by extending the language so that we can talk explicitly about how agents interpret formulas, we do not need structures with ambiguity. Despite that, we argue in Section 5 that we it is useful to be able to model ambiguity directly, rather than indirectly.

2 Syntax and Semantics

2.1 Syntax

We want a logic where players use a fixed common language, but each player may interpret formulas in the language differently. Although we do not need probability for the points we want to make in this paper, for the applications that we have in mind it is also important for the agents to be able to reason about their probabilistic beliefs. Thus, we take as our base logic a propositional logic for reasoning about probability.

The syntax of the logic is straightforward (and is, indeed, essentially the syntax already used in papers going back to Fagin and Halpern [1994]). There is a finite, nonempty set $N = \{1, \dots, n\}$ of players, and a countable, nonempty set Φ of primitive propositions. Let $\mathcal{L}_n^C(\Phi)$ be the set of formulas that can be constructed starting from Φ , and closing off under conjunction, negation, the modal operators $\{CB_G\}_{G \subseteq N, G \neq \emptyset}$, and the formation of probability formulas. (We omit the Φ if it is irrelevant or clear from context.) Probability formulas are constructed as

follows. If $\varphi_1, \dots, \varphi_k$ are formulas, and $a_1, \dots, a_k, b \in \mathbb{Q}$, then for $i \in N$,

$$a_1 pr_i(\varphi_1) + \dots + a_k pr_i(\varphi_k) \geq b$$

is a probability formula, where $pr_i(\varphi)$ denotes the probability that player i assigns to a formula φ . Note that this syntax allows for nested probability formulas. We use the abbreviation $B_i\varphi$ for $pr_i(\varphi) = 1$, $EB_G^1\varphi$ for $\bigwedge_{i \in G} B_i\varphi$, and $EB_G^{m+1}\varphi$ for $EB_G^m EB_G^1\varphi$ for $m = 1, 2, \dots$. Finally, we take *true* to be the abbreviation for a fixed tautology such as $p \vee \neg p$.

2.2 Epistemic probability structures

There are standard approaches for interpreting this language [Fagin and Halpern 1994], but they all assume that there is no ambiguity, that is, that all players interpret the primitive propositions the same way. To allow for different interpretations, we use an approach used earlier [Halpern 2009; Grove and Halpern 1993]: formulas are interpreted relative to a player.

An (*epistemic probability*) *structure* (over Φ) has the form

$$M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, (\pi_j)_{j \in N}),$$

where Ω is the state space, and for each $i \in N$, Π_i is a partition of Ω , \mathcal{P}_i is a function that assigns to each $\omega \in \Omega$ a probability space $\mathcal{P}_i(\omega) = (\Omega_{i,\omega}, \mathcal{F}_{i,\omega}, \mu_{i,\omega})$, and π_i is an interpretation that associates with each state a truth assignment to the primitive propositions in Φ . That is, $\pi_i(\omega)(p) \in \{\mathbf{true}, \mathbf{false}\}$ for all ω and each primitive proposition p . Intuitively, π_i describes player i 's interpretation of the primitive propositions. Standard models use only a single interpretation π ; this is equivalent in our framework to assuming that $\pi_1 = \dots = \pi_n$. We call a structure where $\pi_1 = \dots = \pi_n$ a *common-interpretation structure*; we call a structure where $\pi_i \neq \pi_j$ for some agents i and j a structure *with ambiguity*. Denote by $[[p]]_i$ the set of states where i assigns the value **true** to p . The partitions Π_i are called *information partitions*. While it is more standard in the philosophy and computer science literature to use models where there is a binary relation \mathcal{K}_i on Ω for each agent i that describes i 's accessibility relation on states, we follow the common approach in economics of working with information partitions here, as that makes it particularly easy to define a player's probabilistic beliefs. Assuming information partitions corresponds to the case that \mathcal{K}_i is an equivalence relation (and thus defines a partition). The intuition is that a cell in the partition Π_i is defined by some information that i received, such as signals or observations of the world. Intuitively, agent i receives the same information at each state in a cell of Π_i . Let $\Pi_i(\omega)$ denote the cell of the partition Π_i containing ω . Finally, the probability space $\mathcal{P}_i(\omega) = (\Omega_{i,\omega}, \mathcal{F}_{i,\omega}, \mu_{i,\omega})$ describes the beliefs of player i at state ω , with $\mu_{i,\omega}$ a probability measure defined on the subspace $\Omega_{i,\omega}$ of the state space Ω . The σ -algebra $\mathcal{F}_{i,\omega}$ consists of the subsets of $\Omega_{i,\omega}$ to which $\mu_{i,\omega}$ can assign a probability. (If $\Omega_{i,\omega}$ is finite, we typically take $\mathcal{F}_{i,\omega} = 2^{\Omega_{i,\omega}}$, the set of all subsets of $\Omega_{i,\omega}$.) The interpretation is that $\mu_{i,\omega}(E)$ is the probability that i assigns to event $E \in \mathcal{F}_{i,\omega}$ in state ω .

Throughout this paper, we make the following assumptions regarding the probability assignments \mathcal{P}_i , $i \in N$:

- A1.** For all $\omega \in \Omega$, $\Omega_{i,\omega} = \Pi_i(\omega)$.
- A2.** For all $\omega \in \Omega$, if $\omega' \in \Pi_i(\omega)$, then $\mathcal{P}_i(\omega') = \mathcal{P}_i(\omega)$.
- A3.** For all $j \in N$, $\omega, \omega' \in \Omega$, $\Pi_i(\omega) \cap \Pi_j(\omega') \in \mathcal{F}_{i,\omega}$.

Furthermore, we make the following joint assumption on players' interpretations and information partitions:

- A4.** For all $\omega \in \Omega$, $i \in N$, and primitive proposition $p \in \Phi$, $\Pi_i(\omega) \cap [[p]]_i \in \mathcal{F}_{i,\omega}$.

These are all standard assumptions. A1 says that the set of states to which player i assigns probability at state ω is just the set $\Pi_i(\omega)$ of worlds that i considers possible at state ω . A2 says that the probability space used is the same at all the worlds in a cell of player i 's partition. Intuitively, this says that player i knows his probability space. Informally, A3 says that player i can assign a probability to each of j 's cells, given his information. A4 says that primitive propositions (as interpreted by player i) are measurable according to player i .

2.3 Prior-generated beliefs

One assumption that we do not necessarily make, but want to examine in this framework, is the common-prior assumption. The common-prior assumption is an instance of a more general assumption, that beliefs are generated from a prior, which we now define. The intuition is that players start with a prior probability; they then update the prior in light of their information. Player i 's information is captured by her partition Π_i . Thus, if i 's prior is ν_i , then we would expect $\mu_{i,\omega}$ to be $\nu_i(\cdot \mid \Pi_i(\omega))$.

Definition 2.1 An epistemic probability structure $M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, (\pi_j)_{j \in N})$ has *prior-generated beliefs* (generated by $(\mathcal{F}_1, \nu_1), \dots, (\mathcal{F}_n, \nu_n)$) if, for each player i , there exist probability spaces $(\Omega, \mathcal{F}_i, \nu_i)$ such that

- for all $i, j \in N$ and $\omega \in \Omega$, $\Pi_j(\omega) \in \mathcal{F}_i$;
- for all $i \in N$ and $\omega \in \Omega$, $\mathcal{P}_i(\omega) = (\Pi_i(\omega), \mathcal{F}_i \mid \Pi_i(\omega), \mu_{i,\omega})$, where $\mathcal{F}_i \mid \Pi_i(\omega)$ is the restriction of \mathcal{F}_i to $\Pi_i(\omega)$,¹ and $\mu_{i,\omega}(E) = \nu_i(E \mid \Pi_i(\omega))$ for all $E \in \mathcal{F}_i \mid \Pi_i(\omega)$ if $\nu_i(\Pi_i(\omega)) > 0$. (There are no constraints on $\nu_{i,\omega}$ if $\nu_i(\Pi_i(\omega)) = 0$.)

■

It is easy to check that if M has prior-generated beliefs, then M satisfies A1, A2, and A3. More interestingly for our purposes, the converse also holds for a large class of structures. Say that a structure is *countably partitioned* if for each player i , the information partition Π_i has countably many elements, i.e., Π_i is a finite or countably infinite collection of subsets of Ω .

¹Recall that the restriction of \mathcal{F}_i to $\Pi_i(\omega)$ is the σ -algebra $\{B \cap \Pi_i(\omega) : B \in \mathcal{F}_i\}$.

Proposition 2.2 *If a structure M has prior-generated beliefs, then M satisfies A1, A2, and A3. Moreover, every countably partitioned structure that satisfies A1, A2, and A3 is one with prior-generated beliefs, with the priors ν_i satisfying $\nu_i(\Pi_i(\omega)) > 0$ for each player $i \in N$ and state $\omega \in \Omega$.*

Proof. The first part is immediate. To prove the second claim, suppose that M is a structure satisfying A1–A3. Let \mathcal{F}_i be the unique algebra generated by $\cup_{\omega \in \Omega} \mathcal{F}_{i,\omega}$. To define ν_i , if there are $N_i < \infty$ cells in the partition Π_i , define $\nu_i(\omega) = \frac{1}{N_i} \mu_{i,\omega}(\omega)$. Otherwise, if the collection Π_i is countably infinite, order the elements of Π_i as p_i^1, p_i^2, \dots . Choose some state $\omega_k \in p_i^k$ for each k , with associated probability space $\mathcal{P}_i(\omega_k) = (\Omega_{i,\omega_k}, \mathcal{F}_{i,\omega_k}, \mu_{i,\omega_k})$. By A2, each choice of ω_k in p_i^k gives the same probability measure μ_{i,ω_k} . Define $\nu_i = \sum_k \frac{1}{2^k} \mu_{i,\omega_k}$. It is easy to see that ν_i is a probability measure on Ω , and that M is generated by $(\mathcal{F}_1, \nu_1), \dots, (\mathcal{F}_n, \nu_n)$. ■

Note that the requirement that that M is countably partitioned is necessary to ensure that we can have $\nu_i(\Pi_i(\omega)) > 0$ for each player i and state ω .

In light of Proposition 2.2, when it is convenient, we will talk of a structure satisfying A1–A3 as being *generated by* $(\mathcal{F}_1, \nu_1), \dots, (\mathcal{F}_n, \nu_n)$.

The *common-prior assumption* discussed in the introduction is essentially just the special case of prior-generated beliefs where all the priors are identical.

2.4 Capturing ambiguity

We use epistemic probability structures to give meaning to formulas. Since primitive propositions are interpreted relative to players, we must allow the interpretation of arbitrary formulas to depend on the player as well. Exactly how we do this depends on what further assumptions we make about what players know about each other’s interpretations. There are many assumptions that could be made. We focus on two of them here, ones that we believe arise in applications of interest, and then reconsider them under the assumption that there may be some ambiguity about the partitions.

Believing there is no ambiguity The first approach is appropriate for situations where players may interpret statements differently, but it does not occur to them that there is another way of interpreting the statement. Thus, in this model, if there is a public announcement, all players will think that their interpretation of the announcement is common knowledge. We write $(M, \omega, i) \models^{out} \varphi$ to denote that φ is true at state ω according to player i (that is, according to i ’s interpretation of the primitive propositions in φ). The superscript *out* denotes *outermost scope*, since the formulas are interpreted relative to the “outermost” player, namely the player i on the left-hand side of \models^{out} . We define \models^{out} , as usual, by induction.

If p is a primitive proposition,

$$(M, \omega, i) \models^{out} p \text{ iff } \pi_i(\omega)(p) = \mathbf{true}.$$

This just says that player i interprets a primitive proposition p according to his interpretation function π_i . This clause is common to all our approaches for dealing with ambiguity.

For conjunction and negation, as is standard,

$$(M, \omega, i) \models^{out} \neg\varphi \text{ iff } (M, \omega, i) \not\models^{out} \varphi,$$

$$(M, \omega, i) \models^{out} \varphi \wedge \psi \text{ iff } (M, \omega, i) \models^{out} \varphi \text{ and } (M, \omega, i) \models^{out} \psi.$$

Now consider a probability formula of the form $a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b$. The key feature that distinguishes this semantics is how i interprets j 's beliefs. This is where we capture the intuition that it does not occur to i that there is another way of interpreting the formulas other than the way she does. Let

$$[[\varphi]]_i^{out} = \{\omega : (M, \omega, i) \models^{out} \varphi\}.$$

Thus, $[[\varphi]]_i^{out}$ is the event consisting of the set of states where φ is true, according to i . Note that A1 and A3 guarantee that the restriction of $\Omega_{j,\omega}$ to $\Pi_i(\omega)$ belongs to $\mathcal{F}_{i,\omega}$. Assume inductively that $[[\varphi_1]]_i^{out} \cap \Omega_{j,\omega}, \dots, [[\varphi_k]]_i^{out} \cap \Omega_{j,\omega} \in \mathcal{F}_{j,\omega}$. The base case of this induction, where φ is a primitive proposition, is immediate from A3 and A4, and the induction assumption clearly extends to negations and conjunctions. We now define

$$(M, \omega, i) \models^{out} a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b \text{ iff}$$

$$a_1 \mu_{j,\omega}([[\varphi_1]]_i^{out} \cap \Omega_{j,\omega}) + \dots + a_k \mu_{j,\omega}([[\varphi_k]]_i^{out} \cap \Omega_{j,\omega}) \geq b.$$

Note that it easily follows from A2 that $(M, \omega, i) \models^{out} a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b$ if and only if $(M, \omega', i) \models^{out} a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b$ for all $\omega' \in \Pi_j(\omega)$. Thus, $[[a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b]]_i$ is a union of cells of Π_j , and hence $[[a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b]]_i \cap \Omega_{j,\omega} \in \mathcal{F}_{j,\omega}$.

With this semantics, according to player i , player j assigns φ probability b if and only if the set of worlds where φ holds according to i has probability b according to j . Intuitively, although i “understands” j 's probability space, player i is not aware that j may interpret φ differently from the way she (i) does. That i understands j 's probability space is plausible if we assume that there is a common prior and that i knows j 's partition (this knowledge is embodied in the assumption that i intersects $[[\varphi_k]]_i^{out}$ with $\Omega_{j,\omega}$ when assessing what probability j assigns to φ_k).²

Given our interpretation of probability formulas, the interpretation of $B_j\varphi$ and $EB^k\varphi$ follows. For example,

$$(M, \omega, i) \models^{out} B_j\varphi \text{ iff } \mu_{j,\omega}([[\varphi]]_i^{out}) = 1.$$

²Note that at state ω , player i will not in general know that it is state ω . In particular, even if we assume that i knows which element of j 's partition contains ω , i will not in general know which of j 's cells describes j 's current information. But we assume that i does know that if the state is ω , then j 's information is described by $\Omega_{j,\omega}$. Thus, as usual, “ $(M, i, \omega) \models^{out} \varphi$ ” should perhaps be understood as “according to i , φ is true if the actual world is ω ”. This interpretational issue arises even without ambiguity in the picture.

For readers more used to belief defined in terms of a possibility relation, note that if the probability measure $\mu_{j,\omega}$ is *discrete* (i.e., all sets are $\mu_{j,\omega}$ -measurable, and $\mu_{j,\omega}(E) = \sum_{\omega' \in E} \mu_{j,\omega}(\omega')$ for all subsets $E \subset \Pi_j(\omega)$), we can define $\mathcal{B}_j = \{(\omega, \omega') : \mu_{j,\omega}(\omega') > 0\}$; that is, $(\omega, \omega') \in \mathcal{B}_j$ if, in state ω , agent j gives state ω' positive probability. In that case, $(M, \omega, i) \models^{out} B_j \varphi$ iff $(M, \omega', i) \models^{out} \varphi$ for all ω' such that $(\omega, \omega') \in \mathcal{B}_j$. That is, $(M, \omega, i) \models^{out} B_j \varphi$ iff φ is true according to i in all the worlds to which j assigns positive probability at ω .

It is important to note that $(M, \omega, i) \models \varphi$ does not imply $(M, \omega, i) \models B_i \varphi$: while $(M, \omega, i) \models^{out} \varphi$ means “ φ is true at ω according to i ’s interpretation,” this does not mean that i believes φ at state ω . The reason is that i can be uncertain as to which state is the actual state. For i to believe φ at ω , φ would have to be true (according to i ’s interpretation) at all states to which i assigns positive probability.

Finally, we define

$$(M, \omega, i) \models^{out} CB_G \varphi \text{ iff } (M, \omega, i) \models^{out} EB_G^k \varphi \text{ for } k = 1, 2, \dots$$

for any nonempty subset $G \subseteq N$ of players.

Awareness of possible ambiguity We now consider the second way of interpreting formulas. This is appropriate for players who realize that other players may interpret formulas differently. We write $(M, \omega, i) \models^{in} \varphi$ to denote that φ is true at state ω according to player i using this interpretation, which is called *innermost scope*. The definition of \models^{in} is identical to that of \models^{out} except for the interpretation of probability formulas. In this case, we have

$$(M, \omega, i) \models^{in} a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b \text{ iff} \\ a_1 \mu_{j,\omega}([\varphi_1]_j^{in} \cap \Omega_{j,\omega}) + \dots + a_k \mu_{j,\omega}([\varphi_k]_j^{in} \cap \Omega_{j,\omega}) \geq b,$$

where $[\varphi]_j^{in}$ is the set of states ω such that $(M, \omega, j) \models^{in} \varphi$. Hence, according to player i , player j assigns φ probability b if and only if the set of worlds where φ holds according to j has probability b according to j . Intuitively, now i realizes that j may interpret φ differently from the way that she (i) does, and thus assumes that j uses his (j ’s) interpretation to evaluate the probability of φ . Again, in the case that $\mu_{j,\omega}$ is discrete, this means that $(M, \omega, i) \models^{in} B_j \varphi$ iff $(M, \omega', j) \models^{in} \varphi$ for all ω' such that $(\omega, \omega') \in \mathcal{B}_j$.

Note for future reference that if φ is a probability formula or a formula of the form $CB_G \varphi'$, then it is easy to see that $(M, \omega, i) \models^{in} \varphi$ if and only if $(M, \omega, j) \models^{in} \varphi$; we sometimes write $(M, \omega) \models^{in} \varphi$ in this case. Clearly, \models^{out} and \models^{in} agree in the common-interpretation case, and we can write \models . There is a sense in which innermost scope is able to capture the intuitions behind outermost scope. Specifically, we can capture the intuition that player i is convinced that all players interpret everything just as he (i) does by assuming that in all worlds ω' that player i considers possible, $\pi_i(\omega') = \pi_j(\omega')$ for all players j .

Ambiguity about information partitions Up to now, we have assumed that players “understand” each other’s probability spaces. This may not be so reasonable in the presence of

ambiguity and prior-generated beliefs. We want to model the following type of situation. Players receive information, or signals, about the true state of the world, in the form of strings (formulas). Each player understands what signals he and other players receive in different states of the world, but players may interpret signals differently. For instance, player i may understand that j sees a red car if ω is the true state of the world, but i may or may not be aware that j has a different interpretation of “red” than i does. In the latter case, i does not have a full understanding of j ’s information structure.

We would like to think of a player’s information as being characterized by a formula (intuitively, the formula that describes the signals received). Even if the formulas that describe each information set are commonly known, in the presence of ambiguity, they might be interpreted differently.

To make this precise, let Φ^* be the set of formulas that is obtained from Φ by closing off under negation and conjunction. That is, Φ^* consists of all propositional formulas that can be formed from the primitive propositions in Φ . Since the formulas in Φ^* are not composed of probability formulas, and thus do not involve any reasoning about interpretations, we can extend the function $\pi_i(\cdot)$ to Φ^* in a straightforward way, and write $[[\varphi]]_i$ for the set of the states of the world where the formula $\varphi \in \Phi^*$ is true according to i .

The key new assumption that we make to model players’ imperfect understanding of the other players’ probability spaces is that i ’s partition cell at ω is described by a formula $\varphi_{i,\omega} \in \Phi^*$. Roughly speaking, this means that $\Pi_i(\omega)$ should consist of all states where the formula $\varphi_{i,\omega}$ is true. More precisely, we take $\Pi_i(\omega)$ to consist of all states where $\varphi_{i,\omega}$ is true according to i . If player j understands that i may be using a different interpretation than he does (i.e., the appropriate semantics are the innermost-scope semantics), then j correctly infers that the set of states that i thinks are possible in ω is $\Pi_i(\omega) = [[\varphi_{i,\omega}]]_i$. But if j does not understand that i may interpret formulas in a different way (i.e., under outermost scope), then he thinks that the set of states that i thinks are possible in ω is given by $[[\varphi_{i,\omega}]]_j$. Of course, $[[\varphi_{i,\omega}]]_j$ does not in general coincide with $\Pi_i(\omega)$. Indeed, $[[\varphi_{i,\omega}]]_j$ may even be empty. If this happens, j might well wonder if i is interpreting things the same way that he (j) is. In any case, we require that j understand that these formulas form a partition and that ω belongs to $[[\varphi_{i,\omega}]]_j$. Thus, we consider structures that satisfy A1–A5, and possibly A6 (when we use outermost scope semantics).

A5. For each $i \in N$ and $\omega \in \Omega$, there is a formula $\varphi_{i,\omega} \in \Phi^*$ such that $\Pi_i(\omega) = [[\varphi_{i,\omega}]]_i$.

A6. For each $i, j \in N$, the collection $\{[[\varphi_{i,\omega}]]_j : \omega \in \Omega\}$ is a partition of Ω and for all $\omega \in \Omega$, $\omega \in [[\varphi_{i,\omega}]]_j$.

Assumption A6 ensure that the signals for player i define an information partition according to every player j when we consider the outermost scope semantics. With innermost scope, this already follows from A5 and the definition of $\Pi_i(\omega)$.

We can now define analogues of outermost scope and innermost scope in the presence of ambiguous information. Thus, we define two more truth relations, $\models^{out,ai}$ and $\models^{in,ai}$. (The “ai” here stands for “ambiguity of information”.) The only difference between $\models^{out,ai}$ and

\models^{out} is in the semantics of probability formulas. In giving the semantics in a structure M , we assume that M has prior-generated beliefs, generated by $(\mathcal{F}_1, \nu_1), \dots, (\mathcal{F}_n, \nu_n)$. As we observed in Proposition 2.2, this assumption is without loss of generality as long as the structure is countably partitioned. However, the choice of prior beliefs *is* relevant, as we shall see, so we have to be explicit about them. When i evaluates j 's probability at a state ω , instead of using $\mu_{j,\omega}$, player i uses $\nu_j(\cdot \mid [[\varphi_{j,\omega}]]_i)$. When $i = j$, these two approaches agree, but in general they do not. Thus, assuming that M satisfies A5 and A6 (which are the appropriate assumptions for the outermost-scope semantics), we have

$$(M, \omega, i) \models^{out, ai} a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b \text{ iff} \\ a_1 \nu_j([[\varphi_1]]_i^{out, ai} \mid [[\varphi_{j,\omega}]]_i^{out, ai}) + \dots \\ + a_k \nu_j([[\varphi_k]]_i^{out, ai} \mid [[\varphi_{j,\omega}]]_i^{out, ai}) \geq b,$$

where $[[\psi]]_i^{out, ai} = \{\omega' : (M, \omega, i) \models^{out, ai} \psi\}$.

That is, at $\omega \in \Omega$, player j receives the information (a string) $\varphi_{j,\omega}$, which he interprets as $[[\varphi_{j,\omega}]]_j$. Player i understands that j receives the information $\varphi_{j,\omega}$ in state ω , but interprets this as $[[\varphi_{j,\omega}]]_i$. This models a situation such as the following. In state ω , player j sees a red car, and thinks possible all states of the world where he sees a car that is red (according to j). Player i knows that at world ω player j will see a red car (although she may not know that the actual world is ω , and thus does not know what color of car player j actually sees). However, i has a somewhat different interpretation of “red car” (or, more precisely, of j seeing a red car) than j ; i 's interpretation corresponds to the event $[[\varphi_{j,\omega}]]_i$. Since i understands that j 's beliefs are determined by conditioning her prior ν_j on her information, i can compute what she believes j 's beliefs are.

We can define $\models^{in, ai}$ in an analogous way. Thus, the semantics for formulas that do not involve probability formulas are as given by \models^{in} , while the semantics of probability formulas is defined as follows (where M is assumed to satisfy A5, which is the appropriate assumption for the innermost-scope semantics):

$$(M, \omega, i) \models^{in, ai} a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b \text{ iff} \\ a_1 \nu_j([[\varphi_1]]_j^{in, ai} \mid [[\varphi_{j,\omega}]]_j^{in, ai}) + \dots \\ + a_k \nu_j([[\varphi_k]]_j^{in, ai} \mid [[\varphi_{j,\omega}]]_j^{in, ai}) \geq b.$$

Note that although we have written $[[\varphi_{j,\omega}]]_i^{in, ai}$, since $\varphi_{j,\omega}$ is a propositional formula, $[[\varphi_{j,\omega}]]_i^{in, ai} = [[\varphi_{j,\omega}]]_i^{out, ai} = [[\varphi_{j,\omega}]]_i^{out} = [[\varphi_{j,\omega}]]_i^{in}$. It is important that $\varphi_{j,\omega}$ is a propositional formula here; otherwise, we would have circularities in the definition, and would somehow need to define $[[\varphi_{j,\omega}]]_i^{in, ai}$.

Again, here it may be instructive to consider the definition of $B_j \varphi$ in the case that $\mu_{j,\omega}$ is discrete for all ω . In this case, \mathcal{B}_j becomes the set $\{(\omega, \omega') : \nu_j(\omega' \mid [[\varphi_{j,\omega}]]_j^{in, ai}) > 0\}$. That is, state ω' is considered possible by player j in state ω if agent j gives ω' positive probability after conditioning his prior ν_j on (his interpretation of) the information $\varphi_{j,\omega}$ he receives in state ω . With this definition of \mathcal{B}_j , we have, as expected, $(M, \omega, i) \models^{in, ai} B_j \varphi$ iff $(M, \omega', i) \models^{in, ai} \varphi$ for all ω' such that $(\omega, \omega') \in \mathcal{B}_j$.

The differences in the different semantics arise only when we consider probability formulas. If we go back to our example with the red car, we now have a situation where player j sees a red car in state ω , and thinks possible all states where he sees a red car. Player i knows that in state ω , player j sees a car that he (j) interprets to be red, and that this determines his posterior. Since i understands j 's notion of seeing a red car, she has a correct perception of j 's posterior in each state of the world. Thus, the semantics for $\models^{in,ai}$ are identical to those for \models^{in} (restricted to the class of structures with prior-generated beliefs that satisfy A5), though the information partitions are not predefined, but rather generated by the signals.

Note that, given an epistemic structure M satisfying A1–A4, there are many choices for ν_i that allow M to be viewed as being generated by prior beliefs. All that is required of ν_j is that for all $\omega \in \Omega$ and $E \in \mathcal{F}_{j,\omega}$ such that $E \subseteq [[\varphi_{j,\omega}]_j^{out,ai}]$, it holds that $\nu_j(E \cap [[\varphi_{j,\omega}]_j^{out,ai}]) / \nu_j([[\varphi_{j,\omega}]_j^{out,ai}]) = \mu_{j,\omega}(E)$. However, because $[[\varphi_{j,\omega}]_i^{out,ai}]$ may not be a subset of $[[\varphi_{j,\omega}]_j^{out,ai}] = \Pi_j(\omega)$, we can have two prior probabilities ν_j and ν'_j that generate the same posterior beliefs for j , and still have $\nu_j([[\varphi_k]_i^{out,ai}] \mid [[\varphi_{j,\omega}]_i^{out,ai}]) \neq \nu'_j([[\varphi_k]_i^{out,ai}] \mid [[\varphi_{j,\omega}]_i^{out,ai}])$ for some formulas φ_k . Thus, we must be explicit about our choice of priors here.

3 Common interpretations suffice

In this section, we show in there is a sense in which we do not need structures with ambiguity. Specifically, we show that the same formulas are valid in common-interpretation structures as in structures that do not have a common interpretation, no matter what semantics we use, even if we have ambiguity about information partitions.

To make this precise, we need some notation. Fix a nonempty, countable set Ψ of primitive propositions, and let $\mathcal{M}(\Psi)$ be the class of all structures that satisfy A1–A4 and that are defined over some nonempty subset Φ of Ψ such that $\Psi \setminus \Phi$ is countably infinite.³ Given a subset Φ of Ψ , a formula $\varphi \in \mathcal{L}_n^C(\Phi)$, and a structure $M \in \mathcal{M}(\Psi)$ over Φ , we say that φ is *valid in M according to outermost scope*, and write $M \models^{out} \varphi$, if $(M, \omega, i) \models^{out} \varphi$ for all $\omega \in \Omega$ and $i \in N$. Given $\varphi \in \Psi$, say that φ is *valid according to outermost scope* in a class $\mathcal{N} \subseteq \mathcal{M}(\Psi)$ of structures, and write $\mathcal{N} \models^{out} \varphi$, if $M \models^{out} \varphi$ for all $M \in \mathcal{N}$ defined over a set $\Phi \subset \Psi$ of primitive propositions that includes all the primitive propositions that appear in φ .

We get analogous definitions by replacing \models^{out} by \models^{in} , $\models^{out,ai}$ and $\models^{in,ai}$ throughout (in the latter two cases, we have to restrict \mathcal{N} to structures that satisfy A5 and A6 or just A5, respectively, in addition to A1–A4). Finally, given a class of structures \mathcal{N} , let \mathcal{N}_c be the subclass of \mathcal{N} in which players have a common interpretation. Thus, $\mathcal{M}_c(\Psi)$ denotes the structures in $\mathcal{M}(\Psi)$ with a common interpretation. Let $\mathcal{M}^{ai}(\Psi)$ denote all structures in $\mathcal{M}(\Psi)$

³Most of our results hold if we just consider the set of structures defined over some fixed set Φ of primitive propositions. However, for one of our results, we need to be able to add fresh primitive propositions to the language. Thus, we allow the set Φ of primitive propositions to vary over the structures we consider, but require $\Psi \setminus \Phi$ to be countably infinite so that there are always “fresh” primitive propositions that we can add to the language.

with prior-generated beliefs that satisfy A5 and A6 (where we assume that the prior ν that describes the initial beliefs is given explicitly).⁴

Theorem 3.1 *For all formulas $\varphi \in \mathcal{L}_n^C(\Psi)$, the following are equivalent:*

- (a) $\mathcal{M}_c(\Psi) \models \varphi$;
- (b) $\mathcal{M}(\Psi) \models^{out} \varphi$;
- (c) $\mathcal{M}(\Psi) \models^{in} \varphi$;
- (d) $\mathcal{M}_c^{ai}(\Psi) \models \varphi$;
- (e) $\mathcal{M}^{ai}(\Psi) \models^{out,ai} \varphi$;
- (f) $\mathcal{M}^{ai}(\Psi) \models^{in,ai} \varphi$.

Proof. Since the set of structures with a common interpretation is a subset of the set of structures, it is immediate that (c) and (b) both imply (a). Similarly, (e) and (f) both imply (d). The fact that (a) implies (b) is also immediate. For suppose that $\mathcal{M}_c(\Psi) \models \varphi$ and that $M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, (\pi_j)_{j \in N}) \in \mathcal{M}(\Psi)$ is a structure over a set $\Phi \subset \Psi$ of primitive propositions that contains the primitive propositions that appear in φ . We must show that $M \models^{out} \varphi$. Thus, we must show that $(M, \omega, i) \models^{out} \varphi$ for all $\omega \in \Omega$ and $i \in N$. Fix $\omega \in \Omega$ and $i \in N$, and let $M'_i = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, (\pi'_j)_{j \in N})$, where $\pi'_j = \pi_i$ for all j . Thus, M'_i is a common-interpretation structure over Φ , where the interpretation coincides with i 's interpretation in M . Clearly M'_i satisfies A1–A4, so $M'_i \in \mathcal{M}_c(\Psi)$. It is easy to check that $(M, \omega, i) \models^{out} \psi$ if and only if $(M'_i, \omega, i) \models \psi$ for all states $\omega \in \Omega$ and all formulas $\psi \in \mathcal{L}_n^C(\Phi)$. Since $M'_i \models \varphi$, we must have that $(M, \omega, i) \models^{out} \varphi$, as desired.

To see that (a) implies (c), given a structure $M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, (\pi_j)_{j \in N}) \in \mathcal{M}(\Psi)$ over some set $\Phi \subset \Psi$ of primitive propositions and a player $j \in N$, let Ω_j be a disjoint copy of Ω ; that is, for every state $\omega \in \Omega$, there is a corresponding state $\omega_j \in \Omega_j$. Let $\Omega' = \Omega_1 \cup \dots \cup \Omega_n$. Given $E \subseteq \Omega$, let the corresponding subset $E_j \subseteq \Omega_j$ be the set $\{\omega_j : \omega \in E\}$, and let E' be the subset of Ω' corresponding to E , that is, $E' = \{\omega_j : \omega \in E, j \in N\}$.

Define $M' = (\Omega', (\Pi'_j)_{j \in N}, (\mathcal{P}'_j)_{j \in N}, (\pi'_j)_{j \in N})$, where $\Omega' = \Omega_1 \cup \dots \cup \Omega_n$ and, for all $\omega \in \Omega$ and $i, j \in N$, we have

- $\Pi'_i(\omega_j) = (\Pi_i(\omega))'$;
- $\pi_i(\omega_j)(p) = \pi_j(\omega)(p)$ for a primitive proposition $p \in \Phi$;
- $\mathcal{P}'_i(\omega_j) = (\Omega'_{i,\omega_j}, \mathcal{F}'_{i,\omega_j}, \mu'_{i,\omega_j})$, where $\Omega'_{i,\omega_j} = \Omega'_{i,\omega}$, $\mathcal{F}'_{i,\omega_j} = \{E_\ell : E \in \mathcal{F}_{i,\omega}, \ell \in N\}$, $\mu'_{i,\omega_j}(E_i) = \mu_{i,\omega}(E)$, $\mu'_{i,\omega_j}(E_\ell) = 0$ if $\ell \neq i$.

⁴For ease of exposition, we assume A6 even when dealing with innermost scope.

Thus, $\pi_1 = \dots = \pi_n$, so that M' is a common-interpretation structure; on a state ω_j , these interpretations are all determined by π_j . Also note that the support of the probability measure μ'_{i,ω_j} is contained in Ω_i , so for different players i , the probability measures μ'_{i,ω_j} have disjoint supports. Now an easy induction on the structure of formulas shows that $(M', \omega_j) \models \psi$ if and only if $(M, \omega, j) \models^{in} \psi$ for any formula $\psi \in \mathcal{L}_n^C(\Phi)$. It easily follows that if $M' \models \varphi$, then $M \models^{in} \varphi$ for all $\varphi \in \mathcal{L}_n^C(\Phi)$.

The argument that (d) implies (e) is essentially identical to the argument that (a) implies (b); similarly, the argument that (d) implies (f) is essentially the same as the argument that (a) implies (c). Since $\mathcal{M}_c^{ai}(\Psi) \subseteq \mathcal{M}_c(\Psi)$, (a) implies (d). To show that (d) implies (a), suppose that $\mathcal{M}_c^{ai}(\Psi) \models \varphi$ for some formula $\varphi \in \mathcal{L}_n^C(\Psi)$. Given a structure $M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, \pi) \in \mathcal{M}_c(\Psi)$ over a set $\Phi \subset \Psi$ of primitive propositions that includes the primitive propositions that appear in φ , we want to show that $(M, \omega, i) \models \varphi$ for each state $\omega \in \Omega$ and player i . Fix ω . Recall that $R_N(\omega)$ consists of the set of states N -reachable from ω . Let $M' = (R_N(\omega), (\Pi'_j)_{j \in N}, (\mathcal{P}'_j)_{j \in N}, \pi')$, with Π'_j and \mathcal{P}'_j the restriction of Π_j and \mathcal{P}_j , respectively, to the states in $R_N(\omega)$, be a structure over a set Φ' of primitive propositions, where Φ' contains Φ and new primitive propositions that we call $p_{i,\omega}$ for each player i and state $\omega \in R_N(\omega)$.⁵ Note that there are only countably many information sets in $R_N(\omega)$, so Φ' is countable. Define π' so that it agrees with π (restricted to $R_N(\omega)$) on the propositions in Φ , and so that $[[p_{i,\omega}]]_i = \Pi_i(\omega)$. Thus, M' satisfies A5 and A6. It is easy to check that, for all $\omega' \in R_N(\omega)$ and all formulas $\psi \in \mathcal{L}_n^C(\Phi)$, we have that $(M, \omega', i) \models \psi$ iff $(M', \omega', i) \models \psi$. Since $M' \models \varphi$, it follows that $(M, \omega, i) \models \varphi$, as desired. ■

From Theorem 3.1 it follows that for formulas in $\mathcal{L}_n^C(\Psi)$, we can get the same axiomatization with respect to structures in $\mathcal{M}(\Psi)$ for both the \models^{out} and \models^{in} semantics; moreover, this axiomatization is the same as that for the common-interpretation case. An axiomatization for this case is already given in [Fagin and Halpern 1994]; there is also a complete characterization of the complexity of determining whether a formula is valid.

However, the equivalence in Theorem 3.1 does not extend to subclasses of \mathcal{M} , \mathcal{M}_c , and \mathcal{M}^{ai} . As shown in our companion paper [Halpern and Kets 2013], the equivalence result does not hold if we consider the innermost scope semantics and restrict attention to the subclasses of \mathcal{M} and \mathcal{M}_c that satisfy the common-prior assumption. We defer a further discussion of the modeling implications of this result to Section 5.

4 A more general language

Although, when considering innermost scope, we allowed for agents that were sophisticated enough to realize that different agents might interpret things in different ways, our syntax did

⁵This is the one argument that needs the assumption that the set of primitive propositions can be different in different structures in $\mathcal{M}(\Psi)$, and the fact that every $\Psi \setminus \Phi$ is countable. We have assumed for simplicity that the propositions $p_{i,\omega}$ are all in $\Psi \setminus \Phi$, and that they can be chosen in such a way so that $\Psi \setminus (\Phi \cup \{p_{i,\omega} : i \in \{1, \dots, n\}, \omega \in \Omega\})$ is countable.

not reflect that sophistication. Specifically, the language does not allow the modeler (or the agents) to reason about how other agents interpret formulas. Here we consider a language that is rich enough to allow this. Specifically, we have primitive propositions of the form (p, i) , that can be interpreted as “ i ’s interpretation of p .” With this extended language, we do not need to have a different interpretation function π_i for each i ; it suffices in a precise sense to use a single (common) interpretation function. We now make this precise, and show that this approach is general enough to capture both outermost and innermost scope.

More precisely, we consider the same syntax as in Section 2.1, with the requirement that the set Φ of primitive propositions have the form $\Phi' \times N$, for some set Φ' ; that is, primitive propositions have the form (p, i) for some $p \in \Phi'$ and some agent $i \in N$. We interpret these formulas using a standard epistemic probability structure $M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, \pi)$, with a common interpretation π , as in [Fagin and Halpern 1994]. Thus, truth is no longer agent-dependent, so we have only (M, ω) on the left-hand side of \models , not (M, ω, i) . In particular, if (p, i) is a primitive proposition,

$$(M, \omega) \models (p, i) \text{ iff } \pi(\omega)((p, i)) = \mathbf{true}.$$

As expected, we have

$$(M, \omega) \models a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b \text{ iff} \\ a_1 \mu_{j, \omega}([\varphi_1] \cap \Omega_{j, \omega}) + \dots + a_k \mu_{j, \omega}([\varphi_k] \cap \Omega_{j, \omega}) \geq b.$$

We no longer need to write $[[\varphi_j]]_i^{ou}$ or $[[\varphi_j]]_i^{in}$, since all agents interpret all formulas the same way.

We now show how we can capture innermost and outermost scope using this semantics. Specifically, suppose that we start with an epistemic probability structure $M = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, (\pi_j)_{j \in N})$ over some set Φ of primitive propositions. Consider the corresponding common-interpretation structure $M_c = (\Omega, (\Pi_j)_{j \in N}, (\mathcal{P}_j)_{j \in N}, \pi)$ over $\Phi \times N$, where $\pi(\omega)(p, i) = \pi_i(\omega)(p)$. Thus, M and M_c are identical except in the primitive propositions that they interpret, and how they interpret them. In M_c , the primitive proposition $(p, i) \in \Phi \times N$ is interpreted the same way that i interprets p in M .

We can now define, for each formula ϕ , two formulas ϕ_i^{in} and ϕ_i^{out} with the property that $(M, \omega, i) \models^{in} \phi$ iff $(M_c, \omega) \models \phi_i^{in}$ and $(M, \omega, i) \models^{out} \phi$ iff $(M_c, \omega) \models \phi_i^{out}$. We start with ϕ_i^{in} , defining it by induction on structure:

- $p_i^{in} = (p, i)$
- $(\psi \wedge \psi')_i^{in} = \psi_i^{in} \wedge (\psi')_i^{in}$
- $(a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b)_i^{in} = a_1 pr_j((\varphi_1)_j^{in}) + \dots + a_k pr_j((\varphi_k)_j^{in}) \geq b$
- $(CB_G \psi)_i^{in} = CB_G(\bigwedge_{j \in G} B_j \psi_j^{in})$.

Note that ϕ_i^{in} is independent of i if ϕ is a probability formula or of the form $CB_G\psi$. This is to be expected, since, as we have seen, with innermost scope, the semantics of such formulas is independent of i . The definition of $(CB_G\psi)_i^{in}$ is perhaps the only somewhat surprising clause here; as we discuss after the proof of Theorem 4.1 below, the more natural definition, $(CB_G\psi)_i^{in} = CB_G(\psi_i^{in})$, does not work.

For outermost scope, the first two clauses of the translation are identical to those above; the latter two change as required for outermost scope. Thus, we get

- $p_i^{out} = (p, i)$
- $(\psi \wedge \psi')_i^{out} = \psi_i^{out} \wedge (\psi')_i^{out}$
- $(a_1pr_j(\varphi_1) + \dots + a_kpr_j(\varphi_k) \geq b)_i^{out} = a_1pr_j((\varphi_1)_i^{out}) + \dots + a_kpr_j((\varphi_k)_i^{out}) \geq b$
- $(CB_G\psi)_i^{out} = CB_G(\psi_i^{out})$.

Interestingly, here the natural definition of $(CB_G\psi)_i^{out}$ does work.

Theorem 4.1 *If M is a probabilistic epistemic structure over Φ and M_c is the corresponding common-interpretation structure over $\Phi \times N$, then*

- (a) $(M, \omega, i) \models^{in} \phi$ iff $(M_c, \omega) \models \phi_i^{in}$;
- (b) $(M, \omega, i) \models^{out} \phi$ iff $(M_c, \omega) \models \phi_i^{out}$.

Proof. We prove the result by induction on the structure of ϕ . The argument for outermost scope is completely straightforward, and left to the reader. The argument for innermost scope is also straightforward, except for the case that ϕ has the form $CB_G\psi$. We now consider this case carefully.

By definition,

$$\begin{aligned} & (M_c, \omega) \models (CB_G\psi)_i^{in} \\ \text{iff } & (M_c, \omega) \models CB_G(\bigwedge_{j \in G} B_j \psi_j^{in}) \\ \text{iff } & (M_c, \omega) \models (EB_G)^k(\bigwedge_{j \in G} B_j \psi_j^{in}) \text{ for } k = 1, 2, 3, \dots \end{aligned}$$

Note that, by definition, $(EB_G\psi)_i^{in} = \bigwedge_{j \in G} B_j \psi_j^{in}$. Thus, by the induction hypothesis, it follows that

$$(M_c, \omega) \models \bigwedge_{j \in G} B_j \psi_j^{in} \text{ iff } (M, \omega, i) \models^{in} EB_G.$$

Now by a straightforward induction on k , we can show that

$$(M_c, \omega) \models EB^k(\bigwedge_{j \in G} B_j \psi_j^{in}) \text{ iff } (M, \omega, i) \models^{in} EB_G^{k+1}\psi.$$

That is,

$$(M_c, \omega) \models CB(\bigwedge_{j \in G} B_j \psi_j^{in}) \text{ iff } (M, \omega, i) \models^{in} EB_G^k\psi \text{ for } k = 2, 3, 4, \dots \quad (1)$$

It immediately follows from (1) that if $(M, \omega, i) \models^{in} CB_G\psi$, then $(M_c, \omega) \models CB(\bigwedge_{j \in G} B_j \psi_j^{in})$. The converse also follows from (1), once we show that $(M, \omega, i) \models^{in} EB_G^2\psi$ implies $(M, \omega, i) \models^{in} EB_G\psi$. But this too follows easily since

$$\begin{array}{l}
(M, \omega, i) \models^{in} EB_G^2\psi \\
\text{implies } (M, \omega, i) \models^{in} \bigwedge_{j \in G} B_j(\bigwedge_{j \in G} B_j \psi) \\
\text{implies } (M, \omega, i) \models^{in} \bigwedge_{j \in G} B_j(B_j \psi) \\
\text{iff } (M, \omega, i) \models^{in} \bigwedge_{j \in G} B_j \psi \\
\text{iff } (M, \omega, i) \models^{in} EB_G\psi.
\end{array}$$

This completes the argument. ■

To see why we need the more complicated definition of $(CB_G\psi)_i^{in}$, it is perhaps best to consider an example. By definition, $(CB_{\{1,2\}}p)_1^{in} = CB_{\{1,2\}}(B_1(p, 1) \wedge B_2(p, 2))$. By way of contrast, $CB_{\{1,2\}}(p_1^{in}) = CB_{\{1,2\}}(p, 1)$, which (using arguments similar in spirit to those used above) can be shown to be equivalent to $CB_{\{1,2\}}(B_1(p, 1) \wedge B_2(p, 1))$. The key point here is whether we have $B_2(p, 1)$ or $B_2(p, 2)$. We want the latter, which is what we get from the more complicated translation that we use; it is easy to show that the former does not give the desired result. These issues do not arise with outermost scope.

Theorem 4.1 shows that, from a modeler's point of view, there is no loss in working with common-interpretations structures. Any structure that uses ambiguous propositions can be converted to one that uses unambiguous propositions of the form (p, i) . In a sense, this can be viewed as a strengthening of Theorem 3.1. Theorem 3.1 says that any formula that is satisfiable using innermost or outermost semantics in the presence of ambiguity is also satisfiable in a common-interpretation structure. However, that common-interpretation structure might be quite different from the original structure. Theorem 4.1 shows that if a formula ϕ is true according to agent i at a state ω in a structure M , then a variant of ϕ (namely, ϕ_i^{in} or ϕ_i^{out}) is true at state ω in essentially the same structure.

Moreover, once we add propositions of the form (p, i) to the language, we have a great deal of additional expressive power. For example, we can say directly that agent i believes that all agents interpret p the same way that he does by writing $B_i(\bigwedge_j((p, i) \Leftrightarrow (p, j)))$. We can also make more complicated statements, such as “agent i believes that agents j and k interpret p the same way, although they interpret p differently from him: $B_i((p, j) \Leftrightarrow (p, k)) \wedge \neg B_i((p, i) \Leftrightarrow (p, j))$. Clearly, far more subtle relationships among agents' interpretations of primitive propositions can be expressed in this language.

5 Discussion

We have defined a logic for reasoning about ambiguity, and then showed that, in two senses, we really do not need structures with ambiguity: (1) the same axioms hold whether or not we have ambiguity, and (2) we can use a richer language to talk about the ambiguity, while giving

an unambiguous interpretation to all formulas. So why do we bother using structures with ambiguity? Perhaps the main reason is that it allows us to describe the situation from the agent’s point of view. For example, if we are dealing with outermost scope, an agent does not realize that there are other interpretations possible other than the one he is using. Thus, the simpler language more directly captures agents’ assertions. Similarly, a structure with ambiguity may more accurately describe a situation than a structure with a common interpretation. We thus believe that structures with ambiguity will prove to be a useful addition to a modeler’s toolkit. In any case, whatever modeling framework and language is used, it is clear that we need to take ambiguity into account, and reason explicitly about it.

There are two extensions of our framework that we have not considered. First, we model ambiguity by allowing a formula to be interpreted differently by different agents, we assume that each individual agent disambiguates each formula. That is, no agent says “I’m not sure how to disambiguate ϕ . It could correspond to the U of worlds, or it could correspond to U' ; I’m not sure which is right.” As we mentioned earlier, this view is closer to that of Lewis [1982] and Kuijer [2013]. It would involve a nontrivial change to our framework to capture this. Second, we have allowed only ambiguity about the meaning of primitive propositions (which then extends to ambiguity about the meaning of arbitrary formulas). But we have not considered ambiguity about the meaning of belief; for example, i might interpret belief in ϕ terms of having a proof of ϕ in some axiom system, while j might use a possible-worlds interpretation (as we do in this paper). Capturing this seems interesting, but quite difficult. Indeed, even without ambiguity, it is not nontrivial to design a logic that captures various resource-bounded notions of belief. (See [Fagin, Halpern, Moses, and Vardi 1995][Chapters 9–10] for more on this topic.)

Acknowledgments: We thank Moshe Vardi and the anonymous reviewers of this paper for helpful comments. Halpern’s work was supported in part by NSF grants IIS-0534064, IIS-0812045, IIS-0911036, and CCF-1214844, A preliminary version of this work appeared as “Ambiguous language and differences in beliefs” in the *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference*, 2012, pp. 329–338. by AFOSR grants FA9550-08-1-0438, FA9550-12-1-0040, and FA9550-09-1-0266, and by ARO grant W911NF-09-1-0281. The work of Kets was supported in part by AFOSR grant FA9550-08-1-0389.

References

- Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics* 4(6), 1236–1239.
- Fagin, R. and J. Y. Halpern (1994). Reasoning about knowledge and probability. *Journal of the ACM* 41(2), 340–367.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995). *Reasoning About Knowledge*. Cambridge, Mass.: MIT Press. A slightly revised paperback version was published in

2003.

- Grove, A. J. and J. Y. Halpern (1993). Naming and identity in epistemic logics, Part I: the propositional case. *Journal of Logic and Computation* 3(4), 345–378.
- Halpern, J. Y. (2009). Intransitivity and vagueness. *Review of Symbolic Logic* 1(4), 530–547.
- Halpern, J. Y. and W. Kets (2013). Ambiguous language and consensus.
- Hirst, G. (1988). Semantic interpretation and ambiguity. *Artificial Intelligence* 34(2), 131–177.
- Kuijjer, L. (2013). Sequent systems for nondeterministic propositional logic without reflexivity. In *Fourth International Workshop on Logic, Rationality, and Interaction (LORI-IV)*, LNCS, Volume 8196, pp. 190–203. Springer-Verlag.
- Lewis, D. (1982). Logic for equivocators. *Noûs* 16(3), 431–441.
- Monz, C. (1999). Modeling ambiguity in a multi-agent system. In P. Dekker (Ed.), *Proc. 12th Amsterdam Colloquium*, pp. 43–48.
- Navigli, R. (2009). Word sense disambiguation: a survey. *ACM Computing Surveys* 41(2), 1–69.
- Rescher, N. and R. Brandom (1979). *The Logic of Inconsistency*. Totowa, N.J.: Rowman and Littlefield.
- van Deemter, K. and S. Peters (Eds.) (1996). *Semantic Ambiguity and Underspecification*. Cambridge University Press.