

Ambiguous Language and Consensus*

Joseph Y. Halpern[†]

Willemien Kets[‡]

October 28, 2013

Abstract

Standard economic models cannot capture the fact that information is often *ambiguous*, and is interpreted in multiple ways. Using a framework that distinguishes between the language in which statements are made and the interpretation of statements, we show that ambiguity can have important consequences. We show that players can agree to disagree in the presence of ambiguity, even if there is a common prior, but that allowing for ambiguity is more restrictive than assuming heterogeneous priors. We also demonstrate that, unlike in the case where there is no ambiguity, players may come to have different beliefs starting from a common prior, even if they have received exactly the same information, unless the information is common knowledge. Taken together, these results suggest that ambiguity provides a potential explanation for heterogeneous beliefs. At the same time, it imposes nontrivial restrictions on the situations that can be modeled, so that it is not the case that “anything goes” once we allow for ambiguous statements.

*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. We thank Moshe Vardi for helpful comments. Halpern’s work was supported in part by NSF grants IIS-0534064, IIS-0812045, IIS-0911036, and CCF-1214844, 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.

[†]Computer Science Dept., Cornell University, Ithaca, NY. E-mail: halpern@cs.cornell.edu.

[‡]Kellogg School of Management, Northwestern University, Evanston, IL. E-mail: w-kets@kellogg.northwestern.edu.

1 Introduction

Natural language is often ambiguous; the same message can be interpreted in different ways by different people.¹ Ambiguous language can lead to misunderstandings, and strategic actors may try to exploit ambiguity to their advantage, for example, when writing contracts (Scott and Triantis, 2006),² choosing a campaign platform (Aragonès and Neeman, 2000), or communicating policy intentions (Blinder et al., 2008).

Such ambiguity is hard to model using standard models, which do not separate meaning from message. We therefore develop a framework that distinguishes between the language that players use, and the interpretation of the terms. The language that players use is common, but players may interpret terms differently. A clause in a contract that requires that a widget is of “merchantable quality,” for example, can have a different meaning to different parties. More formally, the clause may be true in one set of states of the world for one party, but in an altogether different set of states according to another.

Allowing for ambiguity has some important consequences, as we demonstrate. It is easy to show that we can “agree to disagree” when there is ambiguity, even if there is a common prior, where players agree to disagree if it is common knowledge that they have different posteriors; in his seminal work, Aumann (1976) has shown that this is not possible without ambiguity. Thus, ambiguity and lack of common priors provide two ways of explaining the fact that we can agree to disagree. Interestingly, as we show, these two explanations are closely related, but not identical. We can convert an explanation in terms of ambiguity to an explanation in terms of lack of common priors.³ Importantly, however, the converse does not hold; there are models in which players have a common interpretation that cannot in general be converted into an equivalent model with ambiguity and a common prior. In other words, using heterogeneous priors may be too permissive if we are interested in modeling a situation where differences in beliefs are due to differences in interpretation.

We go on to show that ambiguity can in fact offer a plausible explanation for differences in beliefs. Aumann (1987) argued that “there is no rational basis for people who have always been fed precisely the same information to [entertain different beliefs].” We show that this no longer holds in the presence of ambiguity. More precisely, it is easy to see that if there is no ambiguity, players with a common prior who have received the

¹We thus use the term ambiguity in a different sense than the decision-theory literature, where ambiguous events are events that the decision-maker cannot assign a precise probability to.

²Also see Bernheim and Whinston (1998), who study the incompleteness of contracts, thus leaving some obligations ambiguous.

³More precisely, we can convert a model with ambiguity and a common prior to an equivalent model—equivalent in the sense that the same formulas are true—where there is no ambiguity but no common prior.

same information have the same posterior. On the other hand, merely receiving the same signals is not sufficient for players to have identical posteriors when there is ambiguity. Intuitively, players may be updating their beliefs on the basis of very different events if there is ambiguity. We show that a necessary and sufficient condition for players to have the same beliefs is that the content of the shared signal must be common knowledge.

Ambiguity thus provides an intuitive explanation of how people may come to have different beliefs about statements such as “the car is blue,” even if they have exactly the same background. Since our results show that not every heterogeneous prior can be explained by ambiguity, ambiguity thus meets the criterion of [Morris \(1995\)](#), who argued that “[n]ot *any* heterogeneous prior beliefs should be acceptable as explanations. We should resort to unmodelled heterogeneities in prior beliefs only when we can imagine an origin for the differences in beliefs.”

A number of authors have considered the role of ambiguity in explaining economic phenomena. [Harris and Raviv \(1993\)](#), for instance, show that speculative trade is possible when traders have a common prior and observe public signals if they use different statistical models to update their beliefs, which they interpret as traders interpreting signals differently. [Blume and Board \(2009\)](#) demonstrate that the strategic use of ambiguous messages can mitigate conflict and thus be welfare enhancing. Neither of these papers models ambiguity syntactically, as we do, which makes it hard to apply the models to related but different situations. [Grant et al. \(2009\)](#) model contracting in the face of ambiguity. They model ambiguity by assuming that unambiguous terms do not fully specify the state of the world.⁴

Rather than focusing on specific applications, we characterize the implications of ambiguity generally, and illustrate how ambiguity can give more insight into various questions. In general, there are two significant differences between our approach to ambiguity and that used in other papers: first, what is ambiguous for us are formulas. While, as a special case, formulas can talk about numeric values, so that we can reproduce the ambiguity considered earlier, in many of our examples, the statements that are ambiguous are non-numeric. In addition, we do not assume that the ambiguity is commonly known. Different agents can have quite different beliefs about what other agents believe, even if there is a public announcement.

⁴Also, [Board and Chung \(2009\)](#) use a simplified model of the one presented here to study ambiguous contracts.

2 Framework

To model that statements can be ambiguous, we explicitly model players' language, i.e., the syntax. After defining the syntax in Section 2.1, we specify the corresponding semantic model in Section 2.2. Section 2.3 discusses how we give meaning to the formulas in the language in the semantic model. This material is largely taken from our companion paper Halpern and Kets (2013), where we study the logic of ambiguity. We have simplified some of the definitions here so as to allow us to better focus on the game-theoretic applications.

2.1 Syntax

We start by defining a formal language that describes all relevant situations, i.e., the syntax. We want a logic where players use a fixed common language, but each player may interpret formulas in the language differently. We also want to allow the players to be able to reason about (probabilistic) beliefs, so as to be able to study the possibility of disagreement about those beliefs.

There is a finite, nonempty set $N = \{1, \dots, n\}$ of players, and a countable, nonempty set Φ of primitive propositions, such as “the economic outlook is improving, but growth remains subdued.” Let \mathcal{L}_n^C be the set of formulas that can be constructed starting from Φ , and closing off under the following operations:

- conjunction (i.e., if φ and ψ are formulas, then so is $\varphi \wedge \psi$ (read “ φ and ψ ”));
- negation (i.e., if φ is a formula, then so is $\neg\varphi$ (read “not φ ”));
- the modal operator CB (i.e., if φ is a formula, then so is $CB\varphi$ (read “ φ is common belief”));
- the formation of probability formulas, defined below.

Probability formulas describe players' beliefs, and 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. Note that we allow for nested probability formulas. The intended reading of $pr_i(\varphi) = x$ is that player i assigns probability x to a formula φ .

It will be convenient to use some abbreviations. We take $\varphi \vee \psi$ (read “ φ or ψ ”) to be the abbreviation for $\neg(\neg\varphi \wedge \neg\psi)$, and $\varphi \Rightarrow \psi$ (“ φ implies ψ ”) to be the abbreviation

for $\neg\varphi \vee \psi$. We use the abbreviation $B_i\varphi$ for the formula $pr_i(\varphi) = 1$ that i believes φ (with probability 1), and we use the abbreviation $EB\varphi$ for the formula $\bigwedge_{i \in N} B_i\varphi$ that all players believe φ , i.e., the formula φ is *mutual belief*. We write $EB^m\varphi$ for the formula $EBEB^{m-1}\varphi$ that φ is *mth-order mutual belief*, where $m = 2, 3, \dots$, and where we write $EB^1\varphi$ for $EB\varphi$.

2.2 Epistemic probability structures

The intended reading of formulas like $\varphi \vee \psi$ and $CB\varphi$ we gave above is supposed to correspond to intuitions that we have regarding words like “or” and “common belief.” These intuitions are captured by providing a semantic model for the formulas in the language, i.e., a method for deciding whether a given formula is true or false.

To model that statements can be ambiguous, we want to allow for the possibility that players interpret statements differently. We build on an approach used earlier (Halpern, 2009; Grove and Halpern, 1993), where formulas are interpreted relative to a player. This means that players can disagree on the meaning of a statement.

More specifically, the semantic model we adopt is an epistemic probability structure. An (*epistemic probability*) *structure* (over a set of primitive propositions Φ) has the form

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

where Ω is the state space and \mathcal{F} is a σ -algebra on Ω , and for each player $i \in N$, Π_i is a partition of Ω , μ_i is i 's prior on Ω (defined on the σ -algebra \mathcal{F}),⁵ 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 , where $\pi_i(\omega)(p) = \mathbf{true}$ means that the primitive proposition p is true in state ω according to i , and $\pi_i(\omega)(p) = \mathbf{false}$ meaning that p is false in state ω according to i . 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*. Otherwise, we say that it is a structure with *ambiguity*.

The information partitions describe the information that player i has in each state: every cell in Π_i is defined by some information that i received, such as signals or observations of the world. Intuitively, player i receives the same information at each state in a

⁵In the companion paper Halpern and Kets (2013), we consider a more general model. In the more general model, we do not assume that agents have a prior on Ω . Rather, we associate with each state $\omega \in \Omega$ each agent's beliefs on ω . The special case where i 's beliefs at ω are obtained from the prior by conditioning on $\Pi_i(\omega)$, i 's information at ω , is what we consider in this paper. There would be no technical difficulty in considering the more general model here, but it would make the presentation a bit more complicated, and is an issue orthogonal to the points we want to make here.

cell of Π_i . As is standard, player i 's posterior beliefs are derived from his prior μ_i and his information partition Π_i . To define these beliefs, we need some more notation and some additional assumptions. Let $\Pi_i(\omega)$ denote the cell of the partition Π_i containing ω , and denote by $[[p]]_i$ the set of states where i assigns the value **true** to p . We assume that for all $i \in N$ and $\omega \in \Omega$:

A1. $\Pi_i(\omega) \in \mathcal{F}$ and $\mu_i(\Pi_i(\omega)) > 0$.

A2. For all primitive proposition $p \in \Phi$, we have $\Pi_i(\omega) \cap [[p]]_i \in \mathcal{F}$.

These are both standard assumptions. A1 says that each element of the information partition for i is measurable, and has positive probability ex ante. Assumption A2 says that primitive propositions (as interpreted by a given player) are measurable.

Player i 's (*posterior*) *belief* in state ω that a given event $E \in \mathcal{F}$ is the case is then simply the conditional probability $\mu_i(E \mid \Pi_i(\omega))$ that E holds given his information, that is, given that the state belongs to $\Pi_i(\omega)$. By A1, players' posterior beliefs are well-defined, and by A2, a player can assign a probability to every primitive proposition (as interpreted by some player).

2.3 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. We write $(M, \omega, i) \models \varphi$ to denote that the formula φ is true at state ω according to player i (that is, according to i 's interpretation). We define \models , as usual, by induction. We start with the primitive propositions $p \in \Phi$. Suppose that p is the statement "the car is blue." Then $(M, \omega, i) \models p$ is true precisely when i would say that the car is blue if he knew the state ω of the world. That is, player i might not be able to see the car, and thus is uncertain as to whether or not it is blue. But if player i saw the car, he would call it blue (even if another agent would call it purple).

Formally, if p is a primitive proposition,

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

This just says that player i interprets a primitive proposition p according to his interpretation function π_i .

For negation and conjunction, as is standard,

$$\begin{aligned} (M, \omega, i) \models \neg\varphi &\text{ iff } (M, \omega, i) \not\models \varphi, \\ (M, \omega, i) \models \varphi \wedge \psi &\text{ iff } (M, \omega, i) \models \varphi \text{ and } (M, \omega, i) \models \psi. \end{aligned}$$

This immediately fixes the interpretation of disjunction, given that $\varphi \vee \psi$ is just $\neg(\neg\varphi \wedge \neg\psi)$:

$$(M, \omega, i) \models \varphi \vee \psi \text{ iff } (M, \omega, i) \models \varphi \text{ or } (M, \omega, i) \models \psi.$$

A critical question is how to interpret probability formulas such as $pr_j(p) \geq b$. As discussed above, a natural reading of $(M, \omega, i) \models p$ is that “if i had all the relevant information about the state of the world, i would say that p is true.” What relevant information would i need to determine whether or not $pr_j(p) \geq b$ is true? If i knew the state of the world, then he would know whether j would say that p was true. Hence, player i uses player j 's interpretation in determining whether $pr_j(p) \geq b$ is true. Thus,

$$(M, \omega, i) \models a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b \text{ iff} \\ a_1 \mu_j([\varphi_1]_j \mid \Pi_j(\omega)) + \dots + a_k \mu_j([\varphi_k]_j \mid \Pi_j(\omega)) \geq b,$$

where $[[\varphi]]_j$ is the set of states ω of the world such that $(M, \omega, j) \models \varphi$.⁶ Hence, according to player i , player j assigns φ probability at least b if and only if the set of worlds where φ holds according to j has probability at least b according to j . Thus, player i “understands” j 's probability space, in the sense that i uses j 's partition element $\Pi_j(\omega)$ and j 's probability measure μ_j in assessing the probability that j assigns to each event.

Given our interpretation of probability formulas, the interpretation of $B_j\varphi$ and $EB^k\varphi$ follows immediately:

$$(M, \omega, i) \models B_j\varphi \text{ iff } \mu_j([\varphi]_j \mid \Pi_j(\omega)) = 1,$$

and

$$(M, \omega, i) \models EB\varphi \text{ iff } \mu_j([\varphi]_j \mid \Pi_j(\omega)) = 1 \text{ for all } j \in N.$$

It is important to note that $(M, \omega, i) \models \varphi$ does not imply $(M, \omega, i) \models B_i\varphi$: while $(M, \omega, i) \models \varphi$ 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

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

⁶To see that all relevant sets are measurable, note that by A2, we have $\Pi_j(\omega) \in \mathcal{F}$. Furthermore, we can use induction to show that $[[\varphi]]_j \cap \Pi_j(\omega) \in \mathcal{F}$ for any formula φ . To see this, assume inductively that $[[\varphi_1]]_j \cap \Pi_j(\omega), \dots, [[\varphi_k]]_j \cap \Pi_j(\omega) \in \mathcal{F}$. The base case of this induction, where φ is a primitive proposition, is immediate from A1 and A2, and the induction assumption clearly extends to negations and conjunctions. To see that the claim holds for probability formulas, note that $(M, \omega, i) \models a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b$ if and only if $(M, \omega', i) \models a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b$ 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 $[[a_1 pr_j(\varphi_1) + \dots + a_k pr_j(\varphi_k) \geq b]]_i \cap \Pi_j(\omega) \in \mathcal{F}$ by A1. It is then immediate from the definitions below that if φ is a formula of the form $CB\psi$, then $[[\varphi]]_i \in \mathcal{F}$.

If all players interpret a formula ψ in the same way in a given structure M (i.e., for any $i, j \in N$, $(M, \omega, i) \models \psi$ if and only if $(M, \omega, j) \models \psi$), we sometimes write $(M, \omega) \models \psi$ for $(M, \omega, \ell) \models \psi$ (where, of course, player ℓ can be chosen arbitrarily). If φ is a probability formula or a formula of the form $CB\varphi'$, then it is easy to see that all players interpret φ the same way: $(M, \omega, i) \models \varphi$ if and only if $(M, \omega, j) \models \varphi$.⁷

One assumption that we do not necessarily make, but want to examine in this framework, is the common-prior assumption. An epistemic probability structure M satisfies the *common prior assumption (CPA)* if $\mu_1 = \dots = \mu_n$.⁸

The following example illustrates how our framework can be applied.

Example 2.1. When negotiators announce the outcome of peace treaties, they may realize that hotheads on either side may not be happy with provisions of the treaty, even if the negotiators from both sides believe that the treaty is in the best interests of all parties. Thus, the negotiators may describe the treaty using deliberately ambiguous terms, in the hope that the hotheads that the on each side will hear what they want to hear.

Formally, suppose that there are three players: one representing each of the two parties, and one representing the negotiators. Suppose that the negotiators make a (public) statement regarding water rights. If the statement is made in a sufficiently obscure way, there may be ambiguity about exactly what was said. Consider the following three primitive propositions:

- $said(p_0)$ – the negotiators said that the water rights are split;
- $said(p_1)$ – the negotiators said that player 1 gets the water rights;
- $said(p_2)$ – the negotiators said that player 2 gets the water rights.

(We can have three further primitive propositions, p_0 , p_1 , and p_2 , with the obvious meanings, but they are not relevant to this discussion.) Taking the negotiator to be

⁷In the companion paper [Halpern and Kets \(2013\)](#), we also consider an alternative semantics that is intended to capture the intuition that, although may players may interpret statements differently, it does not occur to them that there is another way of interpreting the statement; we call this *outermost semantics*. We do not consider outermost scope here, because the semantics presented here is well-suited to model many phenomena of economic interest, as illustrated by our examples.

⁸Since we separate message from meaning, an alternative (syntactic) version of the CPA would require that all players assign the same prior probability to (their interpretation of) all propositional formulas. Formally, let Φ^* be the set of formulas that is obtained from Φ by closing off under negation and conjunction. Then the syntactic version of the CPA requires that there exist priors ν_1, \dots, ν_n such that $\nu_i([\varphi]_i) = \nu_j([\varphi]_j)$ for all formulas $\varphi \in \Phi^*$, and i 's posterior at ω is $\nu_i(\cdot | \Pi_i(\omega))$. We do not investigate this assumption further, as it is unclear why players would assign the same probability to formulas if their interpretations can differ.

player 0, after hearing the negotiator’s statement player i takes $said(p_i)$ to be true, and the other two statements to be false. For simplicity, we assume that it does not occur to either player 1 or 2 that there is any other way to interpret the negotiator’s statement than the way they did. The negotiator is more sophisticated, and correctly understands how the other two players will interpret things.

This can be formalized in a structure where there are three states, ω_0 , ω_1 , and ω_2 . In state ω_0 , i interprets $said(p_i)$ as true, and the other two propositions as false. In state ω_j for $j = 1, 2$, all the players interpret $said(p_j)$ as true. All players have the trivial partition. But they have quite different beliefs. In all states, player i assigns probability 1 to ω_i . The true world is ω_0 (so the negotiator has an accurate picture of the situation, and the other two players do not).

Of course, we could easily modify this structure to capture, for example, a situation where player 1 continues to be certain that $said(p_1)$ is true, but is uncertain about 2’s beliefs, and ascribing positive probability both to states where 2 is certain that $said(p_1)$ is true, and states where 2 believes that $said(p_1)$ is true with only small probability. (This requires player 2 to have at least two information sets.) \triangleleft

Examples of strategic use of ambiguity abound; see Section 1 for references.

3 Agreeing to disagree

Aumann (1976) shows that players cannot “agree to disagree” if they have a common prior. As we show now, this is no longer true if players can have different interpretations. Importantly, though, there are some restrictions on what players can agree to disagree about. In that sense, allowing language to be ambiguous is less permissive than assuming heterogeneous priors.

We start by showing that players can agree to disagree, i.e., it can be common belief that players have different posteriors, even if they have a common prior.⁹

Example 3.1. [Agreeing to Disagree] Consider a structure M with a single state ω , such that $\pi_1(\omega)(p) = \mathbf{true}$ and $\pi_2(\omega)(p) = \mathbf{false}$. Clearly M satisfies the CPA. The fact

⁹Note that in the model of Example 3.1, there is maximal ambiguity: the players disagree with probability 1. We also have complete disagreement. In fact, the less disagreement there is in the interpretation of events, the closer the players come to not being able to agree to disagree. Suppose that M satisfies the CPA, where ν is the common prior, and that $\varphi \in \Phi$. Suppose that the set of states where the players disagree on the interpretation of φ has ν -measure at most $\epsilon \geq 0$. Then one can show that there cannot exist players i and j , numbers b and b' with $b' > b + \epsilon$, and a state ω such that all states that are reachable from ω and $(M, \omega) \models CB((pr_i(\varphi) < b) \wedge (pr_j(\varphi) > b'))$. (In fact, the result holds more generally: it holds for any propositional formula, i.e., any formula that can be composed from the primitive propositions using negation and conjunction.)

that there is only a single state in M means that, although the players interpret p differently, there is perfect understanding of how p is interpreted by each player. Specifically, we have that $(M, \omega) \models CB((pr_i(p) = 1) \wedge (pr_2(p) = 0))$. Thus, according to each player, there is common belief that they have different beliefs at state ω ; that is, they agree to disagree. \triangleleft

Example 3.1 shows that there can be agreement to disagree about an ambiguous statement. One might argue that for economic applications, we are more interested in situations where the disagreement is about formulas that are unambiguous: for example, a trader who offers to sell a risky asset makes an unambiguous statement that he believes the asset is worth less than the price. Rather, there will often be ambiguity about the underlying economic conditions. Indeed, as Brunnermeier (2001) writes, “[e]ven if all traders hear the same news in the form of a public announcement, they still might interpret it differently. [...] Typically one has to make use of other information to figure out the impact of this news on the asset’s value. Thus, traders with different background information might draw different conclusions from the same public announcement.”

The next example shows that we can even have agreement to disagree about unambiguous statements, provided that players condition their beliefs on ambiguous information.

Example 3.2. There are three states, ω_1, ω_2 , and ω_3 . Players believe that they will be informed (privately) about a proposition if and only if p is true. Player 1 thinks that the proposition p is true in states ω_1 and ω_2 (and false in ω_3), and player 2 thinks p is true in ω_2 and ω_3 . There is also an unambiguous proposition q : each player thinks q is true in ω_1 and ω_2 , and false in ω_3 . Thus, player 1’s partition consists of $\{\omega_1, \omega_2\}$ (the states where 1 believes that p is true and $\{\omega_3\}$, while player 2’s partition is $\{\omega_2, \omega_3\}$ and $\{\omega_1\}$. Suppose that the true state is ω_2 , so both players 1 and 2 are informed about p . Then player 1 assigns probability 1 to q , and player 2 assigns probability $\frac{1}{2}$ to q , and this is common belief. \triangleleft

Examples 3.1 and 3.2 demonstrate that there can be agreement to disagree when language is ambiguous, even if players have a common prior and receive the same signals. Of course, as is well-known, players can also agree to disagree if they have heterogeneous priors. This raises the question what the relation is between the two assumptions.

We first note that it is easy to construct a structure M' with heterogeneous priors and common interpretations that is equivalent to the structure M in Example 3.1, in the sense that the same formulas are valid in both (see below for the precise definition).

Example 3.3. To construct the structure M' , let $\Omega' = \{\omega_1, \omega_2\}$, and take $\Pi'_1(\omega') = \Pi'_2(\omega') = \Omega'$ for every state $\omega' \in \Omega'$. Assume that players have the common interpretation

π , which interprets p to be true in ω_1 and false in ω_2 . Player 1 assigns probability 1 to ω_1 , and player 2 assigns probability 1 to ω_2 (i.e., $\mu'_1(\{\omega_1\} \mid \Pi'_1(\omega')) = 1$ and $\mu'_2(\{\omega_2\} \mid \Pi'_2(\omega')) = 1$ for each $\omega' \in \Omega'$). Clearly, we have $(M', \omega', i) \models CB((pr_1(p) = 1) \wedge (pr_2(p) = 0))$ for every state ω' and player $i \in N$, just like in the structure M in Example 3.1. \triangleleft

A key feature of the construction of the equivalent structure M' is that we enlarge the state space: rather than having one state ω , as in M , the structure M' has two states, ω_1 and ω_2 , with ω_i representing the situation from player i 's perspective. As shown in the companion paper Halpern and Kets (2013), this can be done more generally. To state the result, we need some more definitions. A formula $\varphi \in \mathcal{L}_n^C$ is *valid* in a structure M if for every state ω and each player i , we have $(M, \omega, i) \models \varphi$. Two structures are *equivalent* if for every formula φ , φ is valid in M if and only if φ is valid in M' . We then have the following result:

Proposition 3.4. (Halpern and Kets, 2013) For any structure M with ambiguity and a common prior, there is an equivalent common-interpretation structure M' with heterogeneous priors.

Roughly speaking, anything that can be modeled with ambiguous language can be modeled with heterogeneous priors. The converse does not hold, as the next example illustrates: there is a common-interpretation structure with heterogeneous priors that cannot be converted into an equivalent structure with ambiguity that satisfies the CPA.

Example 3.5. We construct a structure M with heterogeneous priors for which there is no equivalent ambiguous structure that satisfies the CPA. The structure M has three players, one primitive proposition p , and two states, ω_1 and ω_2 . In ω_1 , p is true according to all players; in ω_2 , the proposition is false according to all players. Player 1 knows the state: his information partition is $\Pi_1 = \{\{\omega_1\}, \{\omega_2\}\}$. The other players have no information on the state, that is, $\Pi_i = \{\{\omega_1, \omega_2\}\}$ for $i = 2, 3$. Player 2 assigns probability $\frac{2}{3}$ to ω_1 , and player 3 assigns probability $\frac{3}{4}$ to ω_1 . Hence, M is a common-interpretation structure with heterogeneous priors. We claim that there is no equivalent structure M' that satisfies the CPA.

To see this, suppose that M' is an equivalent structure that satisfies the CPA, with a common prior ν and a state space Ω' . Because M and M' are equivalent, we must have $M' \models pr_2(p) = \frac{2}{3}$ and $M' \models pr_3(p) = \frac{3}{4}$, and therefore

$$\nu(\{\omega' \in \Omega' : (M', \omega', 2) \models p\}) = \frac{2}{3}, \quad (3.1)$$

$$\nu(\{\omega' \in \Omega' : (M', \omega', 3) \models p\}) = \frac{3}{4}. \quad (3.2)$$

Note that $M \models B_2(p \Leftrightarrow B_1p) \wedge B_3(p \Leftrightarrow B_1p)$. Thus, since M and M' are equivalent, we must have that the same formula is valid in M' , i.e., that

$$M' \models B_2(p \Leftrightarrow B_1p) \wedge B_3(p \Leftrightarrow B_1p). \quad (3.3)$$

But the interpretation of a formula of the form $B_i\psi$ does not depend on the player, so if we define $E = \{\omega' \in \Omega' : (M', \omega', 1) \models B_1p\}$, then (3.1)–(3.3) imply that we must have $\nu(E) = 2/3$ and $\nu(E) = 3/4$, a contradiction. \triangleleft

This means that allowing for ambiguity (under the CPA) is less permissive than assuming heterogeneous priors. The reason that allowing for ambiguity puts more restrictions on models than heterogeneous priors is that players interpret statements involving others’ beliefs in the same way: all players agree on the set of states where a formula of the form $B_i\phi$ is true.¹⁰ In the companion paper (Halpern and Kets, 2013), we discuss an alternative logic that has the feature that players may not fully understand others’ beliefs. The logic we discuss here has the advantage that it is closest to the standard model: agents are sophisticated in their understanding of the situation.

One might ask how natural the CPA is when players may interpret information in different ways. In the next section, we study the general conditions under which players with a common prior can come to have different beliefs when they receive ambiguous information.

4 Understanding differences in beliefs

Since our framework separates meaning from message, it is worth asking what happens if players receive the same statement, but interpret it differently (as is the case in Example 3.2).¹¹ Aumann (1987) has argued that “people with different information may legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so.” Here we show that this is no longer true when information is ambiguous, even if players have a common prior and fully understand the ambiguity that they face, unless certain strong assumptions on players’ beliefs about the information that others receive are satisfied.

To formalize the argument of Aumann, we assume that information partitions are generated by signals, which are truthful but may be ambiguous. That is, players receive information, or signals, about the true state of the world, in the form of strings (formulas). Each player understands what signals she and other players receive in different states of the world, but players may interpret signals differently.

¹⁰This also means, for example, that if the information that players receive comes in the form of announcements of beliefs (as opposed to ambiguous information, as in Example 3.2 and Section 4 below), as in Geanakoplos and Polemarchakis (1982), they cannot “agree to disagree” forever, just like in the standard framework.

¹¹Al-Najjar (2009) and Acemoglu et al. (2008), among others, have studied the evolution of beliefs when players may interpret information differently. They consider an environment with i.i.d. numeric signals drawn from a fixed but unknown distribution.

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^* do not involve probability formulas, 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 is that in each state of the world ω , each player i receives some signal $\sigma_{i,\omega}$ that determines the states of the world he thinks possible; that is, $\Pi_i(\omega) = [[rec_i(\sigma_{i,\omega})]]_i$, where $rec_i(\sigma_{i,\omega}) \in \Phi^*$ is “ i received $\sigma_{i,\omega}$.” Different players may receive different signals in state ω ; moreover, the formula $\sigma_{i,\omega}$ may be interpreted differently by each player. We assume that player j understands that i may be using a different interpretation than he does, so that j correctly infers that the set of states that i thinks are possible in ω is $\Pi_i(\omega) = [[rec_i(\sigma_{i,\omega})]]_i$. Since we are interested in understanding how players can come to have different beliefs, we restrict attention to structures that satisfy the common-prior assumption in this section. For simplicity, we assume that the set Ω of states of nature is (at most) countable, and take \mathcal{F} to be the power set.

For example, suppose that two players i, j are observing a car’s color. In this case, “receiving a signal” means observing the car’s color. Even if they agree on what blue means (i.e., $[[blue]]_i = [[blue]]_j$, where $blue$ means that the car is blue), they might disagree on whether i has observed a blue car; that is, we might have $[[rec_i(blue)]]_i \neq [[rec_j(blue)]]_j$.¹² Note that in a state where $rec_i(blue)$ is true according to player i , $B_i(blue)$ is also true. However, note that while the interpretation of $B_i(blue)$ is player-independent, like that of all formulas of the form $B_i\psi$, the interpretation of $rec_i(blue)$ may depend on the player.

In any given state, the signals that determine the states that players think are possible may be the same or may differ across players. We are particularly interested in the former case. Formally, we say that a propositional formula $\sigma_\omega \in \Phi^*$ is a *common signal* at ω if $\sigma_{i,\omega} = \sigma_\omega$ for all $i \in N$. As noted earlier, players may interpret the event that a given player receives a signal differently, that is, we may have $[[rec_i(\sigma_\omega)]]_i \neq [[rec_j(\sigma_\omega)]]_j$ even if $[[\sigma_\omega]]_i = [[\sigma_\omega]]_j$. In addition, it may be that some player thinks possible states where the other players have received a signal other than σ_ω , so that he does not know that the signal is in fact common to all players. Similarly, a player may think possible states where another player thinks possible states where other players have received a signal other than σ_ω , and so on.

We first consider [Aumann](#)’s argument in the context of common-interpretation struc-

¹²In addition, one player may think that he observed a blue car while the other might not, even if they both agree on what “blue” is. That is, taking $blue$ to represent that the car is blue, we might have $[[rec_i(blue)]]_i \neq [[rec_j(blue)]]_j$, although $[[blue]]_i = [[blue]]_j$.

tures. One natural formalization of [Aumann](#)'s condition that players are always "fed the same information" is that players believe that all players have received the common signal. Formally, a common signal σ_ω at state ω is a *shared signal* at ω if

$$(M, \omega) \models EB(\bigwedge_{i \in N} \text{rec}_i(\sigma_\omega)).$$

When there is no ambiguity, it is sufficient for players to receive a shared signal for [Aumann](#)'s claim to hold:

Proposition 4.1. Suppose M is a common-interpretation structure with a common prior ν , that is, $\mu_i = \nu$ for $i \in N$, and that σ_ω is a shared signal at ω . Then players' posteriors are identical at ω : for all $i, j \in N$ and $E \in \mathcal{F}$,

$$\mu_i(E \mid \Pi_i(\omega)) = \mu_j(E \mid \Pi_j(\omega)).$$

In particular, for any formula φ ,

$$\mu_i([\varphi])_i \mid \Pi_i(\omega) = \mu_j([\varphi])_j \mid \Pi_j(\omega).$$

Proof. Suppose that $(M, \omega) \models EB(\bigwedge_{i, j \in N} \text{rec}_i(\sigma_\omega))$, and that $\Pi_i(\omega) = [[\text{rec}_i(\sigma_\omega)]]_i$ for all $i \in N$. We first show that $\nu(\Pi_i(\omega)) = \nu(\Pi_i(\omega) \cap \Pi_j(\omega))$ for all players $i, j \in N$. Let $i \in N$. Then, by assumption,

$$\nu(\{\omega' : (M, \omega') \models \bigwedge_j \text{rec}_j(\sigma_\omega)\} \mid \Pi_i(\omega)) = 1,$$

and it follows that

$$\begin{aligned} \nu(\{\omega' : (M, \omega') \models \text{rec}_j(\sigma_\omega)\} \mid \Pi_i(\omega)) &= \nu(\Pi_j(\omega) \mid \Pi_i(\omega)) \\ &= 1. \end{aligned}$$

By the definition of conditional probability, we thus have $\nu(\Pi_i(\omega)) = \nu(\Pi_i(\omega) \cap \Pi_j(\omega))$.

Since $\nu(\Pi_i(\omega) \cap (\Omega \setminus \Pi_j(\omega))) = 0$, it easily follows that, for all events $E \in \mathcal{F}$, $\nu(E \cap \Pi_i(\omega)) = \nu(E \cap \Pi_i(\omega) \cap \Pi_j(\omega))$. Since we also have $\nu(E \cap \Pi_j(\omega)) = \nu(E \cap \Pi_i(\omega) \cap \Pi_j(\omega))$, it follows that $\nu(E \cap \Pi_i(\omega)) = \nu(E \cap \Pi_j(\omega))$. Moreover, taking $E = \Omega$, we have that $\nu(\Pi_i(\omega)) = \nu(\Pi_j(\omega))$. Thus,

$$\nu(E \mid \Pi_i(\omega)) = \frac{\nu(E \cap \Pi_i(\omega))}{\nu(\Pi_i(\omega))} = \frac{\nu(E \cap \Pi_j(\omega))}{\nu(\Pi_j(\omega))} = \nu(E \mid \Pi_j(\omega)).$$

The second claim in the statement of the proposition now follows immediately from the fact that M is a common-interpretation structure. \square

One might think that posteriors coincide if players are merely fed the same information, even if they are uncertain whether others have received the same information, but this is not the case, even in a common-interpretation structure, as the following example demonstrates:

Example 4.2. There are two players, 1 and 2, and two states, labeled ω and ω' . The common prior gives each state equal probability, and players have the same interpretation. The propositional formula σ_ω is a common signal in ω ; and players' information is given by $\Pi_1(\omega) = [[rec_1(\sigma_\omega)]]_1 = \{\omega\}$ and $\Pi_2(\omega) = [[rec_2(\sigma_\omega)]]_2 = \{\omega, \omega'\}$. In state ω , both players receive signal σ_ω , so that $(M, \omega) \models rec_1(\sigma_\omega) \wedge rec_2(\sigma_\omega)$. However, players 1 and 2 assign different probabilities to the event $E = \{\omega\}$. \triangleleft

As we show now, there are conditions under which the claim of [Aumann \(1987\)](#) holds even if there is ambiguity. However, as the next example illustrates, receiving shared signals is no longer sufficient for players' posteriors to coincide:

Example 4.3. As in the previous example, there are two players, 1 and 2, and two states, labeled ω and ω' . The common prior gives each state equal probability. Each player believes that the other player receives the signal σ if and only if she herself does:

$$\begin{aligned} [[rec_1(\sigma)]]_1 &= [[rec_2(\sigma)]]_1 = \{\omega\}; \\ [[rec_1(\sigma)]]_2 &= [[rec_2(\sigma)]]_2 = \{\omega, \omega'\}. \end{aligned}$$

Let $\Pi_i(\omega) = [[rec_i(\sigma)]]_i$, for $i = 1, 2$. Note that this is not a common-interpretation structure. Nevertheless, σ is a shared signal at ω , that is,

$$(M, \omega) \models EB(rec_1(\sigma) \wedge rec_2(\sigma)).$$

But the posteriors differ: player 1 assigns probability 1 to ω , and player 2 assigns probability $\frac{1}{2}$ to ω . \triangleleft

The problem with [Example 4.3](#) is that, although the signal is shared, the players don't interpret receiving the signal the same way. It is not necessarily the case that player 1 received σ from player 1's point of view if and only if player 2 received σ from player 2's point of view. One way of strengthening the condition that all players believe that each player has received the common signal is to require that all players believe that each player has received the common signal, all players believe that all players believe that, and so on, that is, it is common belief that each player has received the common signal. That is, a common signal σ_ω at ω is a *public signal* at ω if

$$(M, \omega) \models CB(\bigwedge_{i \in N} rec_i(\sigma_\omega)).$$

This condition fails in [Example 4.3](#): while player 2 believes at ω that both have received the common signal σ_ω , she does not believe that player 1 believes that: she assigns probability $\frac{1}{2}$ to the state ω' , in which player 1 believes that neither player has received σ_ω . The next result shows that if there is a public signal at ω , then players' posteriors coincide on events, but may differ over formulas.

Proposition 4.4. If M has a common prior ν , and σ is a public signal at ω , then players' posteriors over events are identical at ω : for all $i, j \in N$ and $E \in \mathcal{F}$,

$$\mu_i(E \mid \Pi_i(\omega)) = \mu_j(E \mid \Pi_j(\omega)).$$

However, players' posteriors on formulas may differ; that is, for some formula ψ , we could have that

$$\mu_i([\psi]_i \mid \Pi_i(\omega)) \neq \mu_j([\psi]_j \mid \Pi_j(\omega)).$$

Proof. To prove the first claim, we first show that for all $i, j \in N$, we have that $\nu(\Pi_j(\omega) \mid \Pi_i(\omega)) = 1$. To see this, note that $(M, \omega) \models CB(\bigwedge_{\ell \in N} rec_\ell(\sigma))$ implies that we have $(M, \omega) \models B_i(rec_j(\sigma))$ for all $i, j \in N$. For any $i, j \in N$ and ω' in the support of $\nu(\cdot \mid \Pi_i(\omega))$, we thus have $(M, \omega') \models rec_j(\sigma)$. In other words, the support of $\nu(\cdot \mid \Pi_i(\omega))$ is contained in $\Pi_j(\omega)$, so that $\nu(\Pi_j(\omega) \mid \Pi_i(\omega)) = 1$. The rest of the proof is now analogous to the proof of Proposition 4.1, and therefore omitted.

As for the second claim, it is straightforward to construct an example with public signals and a common prior, but different posteriors over formulas. For example, consider a structure M such that the common prior ν has full support, and in which players have no information, i.e., $\Pi_i(\omega) = \Omega$ for all $i \in N$ and $\omega \in \Omega$. (Thus, the public signal is trivial: “something happened.”) Suppose there is a formula ψ such that $[\psi]_i \neq [\psi]_j$ for some $i, j \in N$. Clearly, i and j have different posterior beliefs about ψ . \square

The next result shows the converse: if players' information comes from a common signal, and their posteriors coincide, then the signal is public.

Proposition 4.5. If M has a common prior ν , players receive a common signal σ at ω , and players' posteriors over events are identical at ω , then σ is a public signal at ω .

Proof. As posteriors coincide, we have that $\nu(\Pi_j(\omega) \mid \Pi_i(\omega)) = 1$ for each $i, j \in N$, so $\nu(\bigcap_j \Pi_j(\omega) \mid \Pi_i(\omega)) = 1$. Consequently, for all $i \in N$, $\omega' \in \Pi_i(\omega)$, we have that $(M, \omega') \models EB(\bigwedge_j rec_j(\sigma))$. For $k > 1$, suppose, inductively, that for all $i \in N$, $\omega' \in \Pi_i(\omega)$, we have that $(M, \omega') \models EB^{k-1}(\bigwedge_j rec_j(\sigma))$. Using again that $\nu(\Pi_j(\omega) \mid \Pi_i(\omega)) = 1$ for each $i, j \in N$, it follows that $(M, \omega') \models EB^k(\bigwedge_j rec_j(\sigma))$ for all $i \in N$ and $\omega' \in \Pi_i(\omega)$. Consequently, $(M, \omega) \models CB(\bigwedge_j rec_j(\sigma))$. \square

Together, Propositions 4.4 and 4.5 demonstrate that when players have a common prior and their information comes from common signals that are potentially ambiguous, players' posterior beliefs coincide if and only if signals are public.

The assumption that players receive a public signal is, of course, very strong: players receive a common signal, and it is commonly believed that they receive that signal. While

the conditions for players' posteriors to coincide seems weaker when there is no ambiguity, it is in fact equally strong: the next result shows that for common-interpretation structures, a signal is shared if and only if it is public.

Proposition 4.6. If M is a common-interpretation structure with a common prior ν , and σ is a common signal at ω , then the following are equivalent:

- $(M, \omega) \models EB(\wedge_j rec_j(\sigma))$; and
- $(M, \omega) \models CB(\wedge_j rec_j(\sigma))$.

Proof. Clearly, if $(M, \omega) \models CB(\wedge_j rec_j(\sigma))$, then $(M, \omega) \models EB(\wedge_j rec_j(\sigma))$. So suppose that $(M, \omega) \models EB(\wedge_j rec_j(\sigma))$. Hence, there is some $\omega' \in \Omega$ such that $(M, \omega') \models \wedge_j rec_j(\sigma)$.

Let $i \in N$, and suppose $(M, \omega') \models \wedge_j rec_j(\sigma)$. The first step is to show that $(M, \omega') \models B_i(\wedge_j rec_j(\sigma))$. As $(M, \omega') \models \wedge_j rec_j(\sigma)$, we have that $(M, \omega') \models rec_i(\sigma)$, so $\omega' \in \Pi_i(\omega)$. Since $(M, \omega) \models B_i(\wedge_j rec_j(\sigma))$, we have that

$$\nu(\{\omega'' : (M, \omega'') \models \wedge_j rec_j(\sigma)\} \mid \Pi_i(\omega')) = \nu(\{\omega'' : (M, \omega'') \models \wedge_j rec_j(\sigma)\} \mid \Pi_i(\omega)) = 1.$$

It follows that $(M, \omega') \models B_i(\wedge_j rec_j(\sigma))$. Hence, for each $i \in N$,

$$\{\omega' : (M, \omega') \models \wedge_j rec_j(\sigma)\} \subseteq \{\omega' : (M, \omega') \models B_i(\wedge_j rec_j(\sigma))\},$$

and it follows that

$$(M, \omega) \models EB^2(\wedge_j rec_j(\sigma)).$$

For $k > 0$, suppose that for each $i \in N$ and $\ell \leq k - 1$,

$$\{\omega' : (M, \omega') \models EB^\ell(\wedge_j rec_j(\sigma))\} \subseteq \{\omega' : (M, \omega') \models B_i(EB^\ell(\wedge_j rec_j(\sigma)))\},$$

and it follows that

$$(M, \omega) \models EB^{\ell+1}(\wedge_j rec_j(\sigma)).$$

Let $i \in N$ and suppose $(M, \omega') \models EB^k(\wedge_j rec_j(\sigma))$. We want to show that $(M, \omega') \models B_i(EB^k(\wedge_j rec_j(\sigma)))$. Since $(M, \omega') \models EB^k(\wedge_j rec_j(\sigma))$, we have that $(M, \omega') \models B_i(EB^{k-1}(\wedge_j rec_j(\sigma)))$, so that

$$\nu(\{\omega'' : (M, \omega'') \models EB^{k-1}(\wedge_j rec_j(\sigma))\} \mid \Pi_i(\omega')) = 1.$$

Hence, by the induction hypothesis,

$$\nu(\{\omega'' : (M, \omega'') \models EB^k(\wedge_j rec_j(\sigma))\} \mid \Pi_i(\omega')) = 1,$$

and $(M, \omega') \models B_i(EB^k(\wedge_j rec_j(\sigma)))$. It follows that for all $i \in N$,

$$\{\omega' : (M, \omega') \models EB^k(\wedge_j rec_j(\sigma))\} \subseteq \{\omega' : (M, \omega') \models B_i(EB^k(\wedge_j rec_j(\sigma)))\},$$

so that

$$(M, \omega) \models EB^{k+1}(\bigwedge_j \text{rec}_j(\sigma)).$$

Consequently, $(M, \omega) \models CB(\bigwedge_j \text{rec}_j(\sigma))$, that is, σ is a public signal at ω . □

As Example 4.3 shows, Proposition 4.6 does not hold when there is ambiguity; in that example, the signal σ is shared, but it is not public.

To summarize, in an environment where players have a common prior and receive information in the form of common signals, the conditions for players to have identical beliefs are very strong: posterior beliefs coincide if and only if signals are public.

5 Concluding remarks

While it has been widely recognized that language or, more generally, signals can be ambiguous, existing game-theoretic models cannot capture this in a natural way. In the standard model, information structures can be extremely general, making it possible to model almost any situation. This generality is also a weakness, however: it is not a priori clear what restrictions to impose on beliefs to capture a phenomenon such as ambiguity in meaning. We circumvent this problem by modeling ambiguity directly, using a formal logic. This automatically delivers the restrictions on beliefs that we need when modeling settings with ambiguous signals.

In formulating our logic, we have stayed as close as possible to the standard framework, by assuming that players fully understand the ambiguity that they face: they understand that others may interpret a statement differently than they do, even if they are uncertain as to the precise interpretation of the other players. This minimal departure from the standard framework already allows for new insights, such as the finding that the common prior assumption can be expected to hold only under fairly restrictive circumstances when signals can be ambiguous. However, one might also expect interesting interactions between ambiguity and unawareness:¹³ a player may not be aware that others may interpret a statement differently.

In politics, for example, dog whistles play an important role. A *dog whistle* is a political message (typically made in a public speech) that will seem innocuous to most listeners, while delivering a message to a specific subset of the electorate. One example, due to Safire (2008), is of a speech given by George W. Bush during the 2004 presidential campaign, where he criticized the U.S. Supreme Court’s 1857 Dred-Scott decision

¹³See, e.g., Feinberg (2005), Halpern and Rêgo (2013a), and Heifetz et al. (2006) for a discussion of unawareness in a strategic context, and Halpern and Rêgo (2013b) for a logic of awareness that could be combined with our logic of ambiguity.

upholding slavery. Most observers might would interpret Bush’s comments as innocuous, but, according to [Safire](#), “sharp-eared observers” (in particular, anti-abortionists) interpreted the remark to be a reminder that Supreme Court decisions can be reversed, and a signal that, if re-elected, Bush might nominate to the Supreme Court a justice who would overturn *Roe v. Wade*. Thus, there is a public announcement p that is interpreted by many people as innocuous (i.e., true in all worlds); moreover, these people are unaware that others can interpret it in a different way. But for a small subset of the population, the true meaning is that Bush will nominate a judge that will overturn *Roe v. Wade*. The people in this latter group realize that that others in their group will interpret p as they do, but that everyone else will interpret it as an innocuous statement.

In such a case, players do not completely understand the ambiguity that they face, unlike in the framework studied here. It would be interesting to study how unawareness interacts with ambiguity, and how game-theoretic predictions change when we vary the level of sophistication of players, from the fully sophisticated ones studied here, to players who understand that there is ambiguity, but do not understand that others understand that, and so on, to the fully unsophisticated, who are unaware of any ambiguity. We leave this for future research.

References

- Acemoglu, D., V. Chernozhukov, and M. Woldz (2008). Fragility of asymptotic agreement under Bayesian learning. Working Paper 08-09, MIT Department of Economics.
- Al-Najjar, N. (2009). Decision makers as statisticians: Diversity, ambiguity, and learning. *Econometrica* 77, 1371–1401.
- Aragonès, E. and Z. Neeman (2000). Strategic ambiguity in electoral competition. *Journal of Theoretical Politics* 12, 183–204.
- Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics* 4(6), 1236–1239.
- Aumann, R. J. (1987). Correlated equilibria as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
- Bernheim, B. D. and M. D. Whinston (1998). Incomplete contracts and strategic ambiguity. *American Economic Review* 88, 902–932.
- Blinder, A. S., M. Ehrmann, M. Fratzscher, J. de Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature* 46, 910–945.

- Blume, A. and O. Board (2009). Intentional vagueness. Working paper, University of Pittsburgh.
- Board, O. and K.-S. Chung (2009). Object-based unawareness: Theory and applications. Working paper, University of Pittsburgh and University of Minnesota.
- Brunnermeier, M. K. (2001). *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding*. Oxford: Oxford University Press.
- Feinberg, Y. (2005). Games with unawareness. Working paper, Stanford University.
- Geanakoplos, J. D. and H. M. Polemarchakis (1982). We can't disagree forever. *Journal of Economic Theory* 28, 192–200.
- Grant, S., J. Kline, and J. Quiggin (2009). A matter of interpretation: Bargaining over ambiguous contracts. working paper, Bond University.
- Grove, A. J. and J. Y. Halpern (1993). Naming and identity in epistemic logics, Part I: the propositional case. *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). A logic for reasoning about ambiguity. Working paper, Cornell University and Northwestern University.
- Halpern, J. Y. and L. C. Rêgo (2013a). Extensive games with possibly unaware players. *Mathematical Social Sciences*. To appear.
- Halpern, J. Y. and L. C. Rêgo (2013b). Reasoning about knowledge of unawareness revisited. *Mathematical Social Sciences* 66(2), 73–84.
- Harris, M. and A. Raviv (1993). Differences of opinion make a horse race. *Review of Financial Studies* 6, 473–506.
- Heifetz, A., M. Meier, and B. Schipper (2006). Interactive unawareness. *Journal of Economic Theory* 130, 78–94.
- Morris, S. (1995). The common prior assumption in economic theory. *Economics and Philosophy* 11, 227–253.
- Safire, W. (2008). *Safire's Political Dictionary (revised edition)*. New York: Oxford University Press.
- Scott, R. E. and G. G. Triantis (2006). Anticipating litigation in contract design. *Yale Law Journal* 115, 814–879.