Representing Bayesian Games Without a Common Prior

Dimitrios Antos and Avi Pfeffer
Harvard University
33 Oxford Street
Cambridge, MA 01239

ABSTRACT

Game-theoretic analyses of multi-agent systems typically assume that all agents have full knowledge of everyone's possible moves, information sets and utilities for each outcome. Bayesian games relax this assumption by allowing agents to have different "types," representing different beliefs about the game being played, and to have uncertainty over other agents' types. However, applications of Bayesian games almost universally assume that all agents share a common prior distribution over everyone's type. We argue, in concord with certain economists, that such games fail to accurately represent many situations. However, when the common prior assumption is abandoned, several modeling challenges arise, one of which is the emergence of complex belief hierarchies. In these cases it is necessary to specify which parts of other agents' beliefs are relevant to an agent's decision-making (or need be known by that agent). We address this issue by suggesting a concise way of representing Bayesian games with uncommon priors. Our representation centers around the concept of a *block*, which groups agents' view of (a) the game being played and (b) their posterior beliefs. This allows us to construct the *belief graph*, a graphical structure that allows agents' knowledge of other agents' beliefs to be carefully specified. Furthermore, when agents' views of the world are represented by extensive form games, our block structure places useful semantic constraints on the extensive form trees. We show through examples that our representation can be used to naturally represent games with rich belief structures and interesting predicted behavior.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Economics, Game Theory, Human Factors

Keywords

Bayesian games, common prior assumption, equilibrium, belief hierarchy, formalism, knowledge requirements

1. INTRODUCTION

Game theory has recently gained wide acceptance as an analysis tool for multi-agent systems. The theory allows intelligent agents to reason strategically about their behavior, as well as the behavior of other agents in the environment. It also offers powerful solution concepts, like Nash equilibrium, that define and predict the choices of rational agents, and are therefore useful to policy-makers who wish to steer the behavior of a population to some desired point. Yet game theory seems to sometimes make excessive and unreasonable assumptions about what agents in a system might know. For instance, it usually assumes that the choices available to every player, the set of outcomes of the game, and each player's utility for every such outcome, are all common knowledge.

Many situations in multi-agent systems are not adequately described by such a framework. Other agents' preferences and utility functions cannot generally be known with any certainty, while sometimes even the choices available to the players, or the observations they receive, might not be fully known. To an extent, Bayesian games (BGs) allow uncertainty over these aspects to be formally captured and reasoned about. In a BG, each player has some private information (e.g., her utility function or available moves), all of which are signified by her "type;" moreover, every player is assumed to know her type, but might be uncertain over other players' types. However, even BGs usually assume that the joint distribution over players' types is common knowledge. This assumption is known as the *common prior assumption* (CPA).

In many practical situations, the use of the CPA is unacceptable or misleading. Suppose, for example, that three people are playing poker. Also assume, for the sake of illustration, that players in poker are either *truthful* or *bluffers*. Alice thinks that Bob is a bluffer with probability 0.4. Carol, however, thinks that Bob is a bluffer, say, with probability only 0.1. The strategies optimally chosen by each player should obviously depend on their beliefs about how likely it is for their opponents to be truthful or bluffers. The question then is: should Alice reason about Carol having a *different* belief about the likelihood of Bob's type? The answer is yes, but a BG with a common prior cannot capture this.[1]

The CPA has been justified on both philosophical and practical grounds. One of the most important arguments in favor of the CPA is that relaxing it introduces serious modeling challenges. For example, in a BG with common priors, agents' inferences are "compatible," e.g., what agent i believes agent j of a particular type believes is exactly what agent j of that type in fact believes. In contrast, dropping the CPA may lead to infinitely nested beliefs

Cite as: The title of your paper should be written here, Author1, Author2 and Author3, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

[1] A BG with a common prior could capture this if the priors themselves were drawn from a hyper-prior that is common knowledge, but this is not natural, and the assumption of a common-knowledge hyper-prior is in general not easier to justify than that of a common prior.

of the form "*i* believes that *j* believes that *k* believes that *i* believes…" that are quite hard to formally reason with. Another issue with relaxing the CPA is that doing so seems to lead to one of two undesirable possibilities. We could assume that while players' priors differ, these different priors are common knowledge; this, however, is as difficult to justify as the CPA. Alternatively, we could make their priors private, in which case each player is playing a different, completely subjective game, and there is no interaction between them.

This paper addresses both these challenges. First, to make BGs with uncommon priors easier to model, it presents a novel technique for representing such games. This technique offers several advantages: Defining a game is conceptually easy, as the agents' private information and beliefs are conveniently grouped into *blocks*. Moreover, the beliefs across the various blocks can be used to construct the *belief graph*, a structure that describes which parts of other agents' priors are relevant to the decision-making of a particular agent. The belief graph therefore reveals how much an agent needs to know about others' belief hierarchies to compute a solution (equilibrium) to the game.

Second, the paper provides a middle road between the two undesirable alternatives discussed above. Rather than having to choose between the priors being common knowledge or completely subjective, our formalism allows precise control of which priors a player of a particular type knows (or needs to know). It is possible in our formalism for one type of a player to be more knowledgeable about other players' priors than another type. This is especially useful when some types of players are "boundedly rational," and hence have limited knowledge of the game. Such a representational capability is achieved through the belief graph.

While the basic formalism does not commit to a particular underlying representation for the individual games within the agents' beliefs, it becomes particularly interesting when the underlying games are in extensive form. In such situations, it utilizes a mapping between information sets across the various game trees that is consistent with the agents' beliefs. Constraints are placed on the extensive form games that allow the entire structure to be studied as a coherent game. We show that when the underlying games satisfy these constraints, the overall game defines a BG with a natural and compact type space.

The structure of the paper is as follows: In Section 2 we discuss the advantages and disadvantages of the common prior assumption, and argue that certain commonly encountered situations render it highly inappropriate. In Section 3 we discuss issues that arise when the CPA is relaxed, and present an example illustrating how interesting belief structures might arise. Section 4 presents our block formalism, and Section 5 explains how it can be used to construct the belief graph, which can concisely represent these interesting belief structures in various types of games without a common prior. In section 6 we give an extended example to illustrate how our formalism can be used in practice. Next, section 7 shows how it fits in a particularly convenient way with the extensive form. Finally, section 8 compares and contrasts our formalism with prior game representations. The paper concludes with a short discussion and directions for future work.

2. THE COMMON PRIOR

The common prior assumption (CPA) has been widely accepted on the basis of three arguments (for three thoughtful treatises of the CPA, see [7], [9] and [12]):

- *Rationality implies the CPA*. This argument states that rational agents having received identical information (or none at all) cannot publicly disagree on the likelihood of any event or outcome in the world—if one agent disagrees, then she must have made a mistake and therefore she is not rational. The above is flawed in three ways: First, probabilities of events are not directly observable in the world, but are instead statistical properties of complex mechanisms that either we do not fully understand or might be too hard to exhaustively model. For example, the fact that a fair coin lands on either side with equal likelihood is a result of small forces applied to its surface, its weight distribution and small air currents in a "typical" environment where the coin is tossed; but the number 0.5 is not directly observable or measurable on the coin itself. If probabilities are then derived from statistical assessments, rationality objections are irrelevant. Second, if one adopts the subjective view of probability (as Bayesians do), there is nothing implying that all agents must agree on the probabilities of all events—otherwise probabilities would not be truly subjective. Finally, rationality tells us how to update a prior given new observations, but in general says nothing about how to choose a prior in the first place.
- *Learning will cause agents' priors to converge to the empirically observed frequencies.* This argument basically states that, even if agents start with differing priors, they should, given enough data, update those to more or less identical posteriors, which they will then use as their new (almost common) priors. Of course such a claim is valid only if we can safely assume that agents have been given enough time in the environment to explore and assess the likelihood of every event with sufficient accuracy. In many multi-agent systems, however, the environment is too complex to even represent in state-space, let alone fully explore and quantitatively assess. Furthermore, agents often engage in localized interactions, and therefore their experiences may be concentrated in only a narrow subspace of the environment.
- *Abandoning the CPA would cause serious modeling issues.* This is a more serious concern. The fact that disagreement between agents arises solely from different private information, as captured in their types, leads to simpler models, since the beliefs of all agents are consistent with their common prior. Therefore, agents' inferences are "compatible," e.g., what agent *i* believes agent *j* of a particular type believes is exactly what agent *j* of that type in fact believes. On the contrary, dropping the CPA may lead to infinitely nested beliefs of the form "*i* believes that *j* believes that *k* believes that *i* believes…" that are quite hard to formally reason with (yet not impossibly hard, e.g., see [8]). Furthermore, as [1] has shown, a repeated Bayesian game with uncommon priors may never converge to any equilibrium under *any* learning process.

3. RELAXING THE CPA

In many cases, forcing the agents to have common priors is an unnatural imposition on the model. Why would they have a common prior if they have different subjective beliefs about how the world works, or the space of possible strategies? In such cases, we can try to relax the CPA. To do so, we assign the agents different prior distributions over the joint types. If we do that, we will be able to say things like "Alice thinks Bob is a bluffer with probability 0.4, but Carol only thinks he's a bluffer with probability 0.1." The question that immediately arises is, "What does Alice think that Carol thinks about Bob?" One possible answer to this question is to say that while the priors are uncommon, they are common knowledge, so Alice knows Carol's prior and the fact that she thinks Bob is a bluffer with probability 0.1. However, if we are not assuming that the agents have a common prior, this assumption is very hard to justify. If the agents have different subjective beliefs that lead them to have different priors, how can we expect them to know each others' subjective beliefs?

Alternatively, we can say that agents do not know other agents' priors. Instead, they must have beliefs about them. One possibility is to say that agents believe other agents' priors are like their own, so Alice believes Carol thinks Bob is a bluffer with probability 0.4. In this approach, each agent solves a separate BG with common priors. The agents do not reason about other agents having different priors from their own. Each agent comes up with a different, completely subjective solution to the game under this assumption. Agents whose beliefs about each other's priors are entirely independent might derive a completely different set of Bayes-Nash equilibria for the game.

Another possibility is to say that each agent has private beliefs about the priors of other agents, which need not necessarily be the same as their own. For example, Alice might believe that Carol thinks Bob is a bluffer with probability 0.3, while Carol herself believes he is a bluffer with probability 0.1. If we follow this option, we then need to ask what Alice believes Carol believes Alice thinks about Bob. We thus end up with a complex hierarchy of beliefs that is very hard to model. Furthermore, each agent's belief hierarchy has no interaction with the other agents, so each agent's solution of the game is again completely subjective.

We contend that these three approaches—common knowledge, believing everyone has the same prior, and completely private beliefs—are just three out of many structures that can arise when relaxing the CPA. We will shortly present a formalism that can represent a wide and diverse array of such structures. Before doing that, however, we present an example that shows how such rich structures can come about.

Suppose you are playing Rock-Paper-Scissors (RPS) with a child. Let us also assume that there are only two types of agents in RPS: *maximizer* and *naïf*. The maximizer agents have a correct model for the game, and utility function that gives them +1 for winning, –1 for losing and 0 for a draw. The naïf type, on the other hand, represents how children are expected to play the game: naives model their opponent as an automaton that always repeats its last move. Therefore, naives best-respond to their opponent's last choice, e.g., they play 'scissors' after their opponent has played 'paper.' For completeness, assume that naives on the first round choose any action with equal probability.

Now let us assume that you are a maximizer (you know your type) and you believe that the child is either a maximizer (with probability p), or a naïf (with probability $1 - p$). This belief is consistent with many possible priors. One possibility is shown in Table 1.

Table 1 Your prior in the RPS game (p_1)

(t_1, t_2)	probability
(maximizer, maximizer)	p^2
(maximizer, naïf)	$p(1-p)$
(naïf, maximizer)	$p(1-p)$
(naïf, naïf)	$(1-p)^2$

Table 2: The child's prior in the RPS game (p_2)

(t_1, t_2)	probability
(maximizer, maximizer)	p^2
(maximizer, naïf)	0
(naïf, maximizer)	$p(1-p)$
(naïf, naïf)	$1-p$

However, the child does not share your prior. This is because, if a child is a naïf, it would be rather unreasonable to expect it to even be *aware* of the concept of a game-theoretic maximizer. Hence, the child's prior could look like Table 2.

Clearly, no common prior can adequately capture this situation. Many instances in which bounded rationality is evoked, or in which some agents are only partially aware of the game's structure, can be viewed as generalizations of the above scheme. The question that now arises is "what prior does the child believe you are using?" One option here would be to assume that the child believes that you use your *actual* prior, the one shown in Table 1. In this case, the situation is a game with uncommon, common knowledge priors. However, closer scrutiny would reveal that this turns out to be rather unnatural: the child, when its type is naïf, is assumed not to even be *aware* of the possibility of you being anything but naïf, but also seems to maintain a non-zero probability that you might think it is *not* a naïf.

A better solution here can be achieved by constructing a richer belief structure. In this case, we shall use three types: *maximizer*, *naïf*, and *naïf-adult*. An agent of the naïf-adult type models the game exactly as a naïf (the usefulness of adding this new type will be apparent below). Let us now look at beliefs: You, as an adult, will still use the prior shown in Table 1. The child, however, will use the prior of Table 3 (missing assignments are assumed to have probability zero):

Table 3: The child's updated prior in the RPS game (p_3)

(t_1, t_2)	probability
(maximizer, maximizer)	$0.9p^2$
(maximizer, naïf)	0
(naïf, maximizer)	$0.9p(1-p)$
(naïf, naïf)	0
(naïf-adult, naïf)	$0.9(1-p)$
(naïf-adult, naïf-adult)	0.1

Furthermore, the child shall believe that you are using the prior below (Table 4):

Table 4: The child's belief about your prior (p_4)

(t_1, t_2)	probability
(maximizer, maximizer)	$0.7p^2$
(naïf, maximizer)	$0.7p(1 - p)$
(maximizer, naïf)	$0.7p(1 - p)$
(naïf, naïf)	$0.7(1 - p)^2$
(naïf-adult, naïf-adult)	0.3

This belief structure now describes a child that, if a maximizer, will be aware of the possibility that you are either a maximizer or a naïf (Table 3, first and third rows). If a naïf though, the child will also assume that you are a naïf (Table 3, fifth row) and that you are as oblivious to the existence of any types other than itself. To capture that last piece, the naïf-adult type is introduced, such that it is wholly similar to the naïf type, except it fails to consider other alternative types (Table 4, last row). (The constants in Table 4 are chosen such that the probabilities sum to one and do not affect the posterior distributions.)

4. THE REPRESENTATION

The main idea behind our formalism is that modeling a game becomes simpler if the agents' types and beliefs are captured in a conceptually appealing and graphical way. We therefore introduce the concept of a "block," and define a game as a *collection* of blocks B. A block $b \in B$ consists of two elements: (1) the model $m(b)$ an agent has about the world, and (2) the beliefs $\beta(b)$ the agent assumes, and believes others to assume, in that block. The agent's model is a complete game in normal or extensive form, with everyone's information sets, available moves and utilities fully specified. The beliefs in block b consist of $n(n-1)$ probability distributions over B, indexed p_{ij}^b for all $i, j \in N$, $i \neq j$, where N is the set of agents ($|N| = n$). The distribution p_{ij}^b captures agent i's beliefs over which block agent j might be using. Also, let us denote by $p_{ij}^b(b')$ the probability assigned to block b' by the distribution p_{ij}^b.

It is straightforward to map this construct onto a Bayesian game. For each agent i, her typeset T_i is equivalent to the set of blocks B. When agent i is of a particular type, say $b \in B$, then agent i's private information (utility, observations, etc.) are fully captured by the game $m(b)$. Moreover, i's posterior distribution over the beliefs of all other agents given her type, $p(\mathbf{T}_{-i} = (t_j)_{j \neq i} \mid T_i = b)$, is given by the product $\prod_{j \neq i} p_{ij}^b(t_j)$ of the distributions in $\beta(b)$.

Notice how in our formalism the modeling is performed in terms of the posterior distributions $p_i(T_j \mid T_i = b) \equiv p_{ij}^b$, not the priors $p_i(T)$. Given these posteriors, *any* prior that is consistent with them will be essentially expressing the same game.

In each block b, the set of pure strategies for player i contains all her pure strategies in the model $m(b)$. For the game as a whole, a pure strategy for i is then a choice of pure strategy for every block $b \in B$. Moreover, if the models $m(b)$ are represented in extensive (tree) form, a pure strategy for i for the whole game is a mapping from all information sets of all trees $m(b)$ to an action available to her in every such information set. Similarly, *mixed* strategies are probability distributions over pure strategies, and *behavioral* strategies can be defined as mappings from information sets to probability distributions over available actions. Finally, a strategy profile σ denotes, for every agent i and every type $b \in B$, a choice of mixed (or behavioral) strategy $\sigma_{i,b}$.

The main solution concept for a Bayesian game is a Bayes-Nash equilibrium. A strategy profile σ is a Bayes-Nash equilibrium if, for all agents i and for all types b, the strategy $\sigma_{i,b}$ maximizes i's expected utility against strategies $\sigma_{j,b'}$, where each is weighted according to the posterior distribution $p_{ij}^b(b')$.

Observe here that the notion of a Bayes-Nash equilibrium does not change when the common prior assumption is dropped. For a strategy profile to be in equilibrium, all that is required is that each agent, given her beliefs (about herself and others' beliefs), considers everyone to be best-responding to everybody else's expected behavior. (For a deep discussion of equilibria in non-common prior games see also [10]. For approaches towards computing them see [11].) There are, however, subtle changes in the *interpretation* of equilibrium with and without a common prior, as well as with and without common knowledge of priors.

If the CPA is adopted, all agents agree on the game being played and therefore the equilibrium represents an optimal solution to it. Replacing the common prior with commonly known, differing priors maintains the agents' belief that this equilibrium is an optimal solution, but each of them thinks that only *her* utility is maximized in expectation under the equilibrium. Others' utilities are not necessarily maximized; they only *think*, using their erroneous priors, that their utilities are maximized. Hence equilibria are in a sense subjective solutions. On the other hand, if priors are also private, then it need not necessarily hold that agents even *agree* on what the equilibria of the game are. If the prior of agent i is very different from the prior agent j assumes for i, then clearly the equilibria of the game, as computed by the two agents, might be completely unrelated.

5. THE BELIEF GRAPH

One useful property of our formalism is that it allows for belief dependencies to be uncovered easily. In particular, it can help a modeler answer the question "Which of the beliefs of other agents are relevant to agent i's decision-making?" This is performed by constructing the game's *belief graph*. Alternatively, this question can be stated as "Which of the other agents' beliefs need i know to compute her optimal decision?"

The belief graph is constructed as follows: Its nodes are the set of blocks B. Then, we add an edge (b, b') and we label it "$i.j$" if $p_{ij}^b(b') > 0$. In other words, the edge (b, b') denotes that agent i in block b assumes that j might be using block b' as his model of the world. The destination block b' may be the same as the source b (self-edge). Next, we define a path $\pi = (b_1, \ldots, b_m)$ such that, for every node b_k, where $k \in [1, m-1]$, there is an edge (b_k, b_{k+1}) and, for each consecutive edge pair $\{(b_k, b_{k+1}), (b_{k+1}, b_{k+2})\}$, where $k \in [1, m-2]$, the label of the first edge is "$i.j$" and the label of the second is "$j.k$" for some agents i, j and k. (A path may very well contain self-edges.) We say that a block-agent pair (b', j)

is *reachable* from pair (b, i) if there is a path from b to b' in which the first agent is i and the last agent is j. The set of reachable blocks from (b, i) is denoted by $R(b, i)$.

The belief graph captures which distributions an agent needs to take into account in its decision-making. Only those posterior distributions $p_l^{b'''}$, where $(b''',l) \in R(b,i)$, are relevant to agent i's decision-making, when that agent is in block b. This is because, in block b, which is deemed by agent i to be the "true world," some other agent j will be modeled *as if* he was using one of the blocks for which an edge (b, b') exists with label "$i.j$," hence j's beliefs in b' need to be considered by i in b. Furthermore, agent j in b' might be modeling k (who could be the same as i), *as if* she were using some block b'', which j needs to consider in order to predict their behavior and hence best-respond to it. Therefore i in b, who is best-responding to j, must also consider k in b''. By induction, if $(b''',l) \in R(b,i)$, the game tree $m(b''')$ and the posterior $p_l^{b'''}$ are potentially relevant to i's decision-making problem. On the other hand, if $(b''',l) \notin R(b,i)$, there is no path of reasoning by which agent i in b needs to take agent l in b''' into account, so $p_l^{b'''}$ is irrelevant to i in b.

In addition, if there is an edge from b to b' labeled "$i.j$," then agent i in block b believes that agent j's model of the world is $m(b')$ and his beliefs are $\beta(b')$. Thus agent i in b knows the model and beliefs of j in b'. Likewise, if there is an edge from b' to b'' labeled "$j.k$," agent j in b' knows agent k's model and beliefs in b''. It follows that agent i in b knows agent k's model and beliefs in b''. Recursively, agent i in block b knows the beliefs of all agents l in blocks b''' such that (b''',l) is in $R(b, i)$. Thus the belief graph precisely captures what agents must know about the beliefs of other agents. A subtle point must be made, however. If (b', j) is not reachable from (b, i), the graph does not preclude the possibility that agent i in b knows the beliefs of agent j in b'. It merely says that i does not *need* to know them, so it is not necessary to assume that i knows them.

The different relaxations of the CPA discussed in Section 3 form special cases of the belief graph. If in the graph every block-agent pair is reachable from every other block-agent pair, then the agents' beliefs are common knowledge. The case in which priors are completely private is captured by a belief graph containing n disconnected subgraphs, one for each agent.

In **Figure 1** we present the belief graph for the Rock-Paper-Scissors example of section 3. **Figure 1**(a) shows the graph for the case of common knowledge priors. The game, $m(b_1)$ is a standard description of rock-paper-scissors, while $m(b_2)$ is an alternative game in which the adult always repeats her last move. In b_1, you (the adult) model the child as using b_2, while in b_2 the child models you as using b_1.

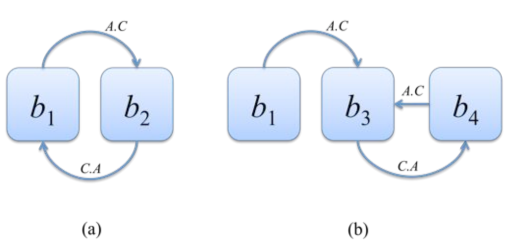

Figure 1: The belief graph of the RPS game

Alternatively, in **Figure 1**(b) the case with the richer belief structure is considered, whereby you model the child as using b_3, and the child models you as using b_4. Notice how in this case the child need not be aware of your beliefs in b_1 to compute his optimal strategy. On the other hand, you in b_1 do consider the child's beliefs in b_3, and therefore also your own beliefs in b_4.

6. AN EXTENDED EXAMPLE

We now present an extended example that further illustrates these concepts, and shows the interesting equilibrium analysis that our framework supports. Suppose Alice works for company X and all her retirement savings are tied to X's stock, therefore she would benefit greatly from an increase in their value. Bob is a billionaire considering to buy a large number of X's stocks, but lacks the expertise to make this decision. Hence, he relies on a predicting agency C that informs him whether the stock price will go up (in which case it is optimal for him to buy) or down (in which case he should optimally rest). Clearly, Alice would significantly prefer Bob buying the stock. Let us also assume that Bob believes that C, the predicting agency, has prior probability 0.2 to suggest 'buy' and 0.8 to suggest 'rest.'

Suppose now that, in Alice's model of the world, she can "threaten" the predicting agency in some frowned-upon way, which Alice thinks is entirely effective, i.e., a threatened C will do as Alice says, i.e., suggest 'buy' if Alices dictates it to say 'buy,' and suggest 'rest' if Alice dictates it to say 'rest.' (Assume that either action has zero cost for Alice.) In Bob's model, Alice can choose either action, but the agency cannot be bullied by her threats (in this case, the scenario is equivalent to Bob being unaware of Alice's presence); this means that Bob considers, in his model, C's suggestion to be trustworthy and predictive of the stock's future price. Alice, however, is uncertain of whether Bob is aware of her manipulative power—in fact, she believes that the probability of him being aware of her actions' effects is 0.3.

We shall create two blocks to represent this situation ($B = \{K, L\}$). In blocks K and L the models, which are represented by extensive form game trees, look like in Figure 2. At first Alice decides whether to dictate 'buy' or 'rest.' Then Nature, labeled C, representing the agency, makes a suggestion to Bob. Finally, Bob has to decide whether to buy the stock or not. He has two information sets, one representing the history "C suggested 'buy'," and one capturing the alternative "C suggested 'rest'." The only difference between the two trees $m(K)$ and $m(L)$ is the fact that the probability of C making a 'buy' suggestion to Bob in $m(L)$ is 0.2 indepenently of Alice's action, whereas in $m(K)$ it is either one or zero, depending on Alice's threat.

Let us now define the beliefs β. In block K we need to have a distribution $p_{A,B}^{K}$ = <0.3 : K, 0.7 : L> to capture the fact that Alice of type K thinks that Bob is likely to use K (and hence be aware of her threats) with probability 0.3, or use L (and therefore be oblivious to her actions) with probability 0.7. Bob in K will have what we call a *trivial* belief, i.e., $p_{B,A}^{K}$ = <1 : K>, meaning that a Bob who is aware of Alice's threats also considers her to be aware of them, which makes sense. In block L both agents have *trivial* beliefs, such that Bob of type L is unaware of Alice.

The belief graph for the game is presented in Figure 3. Notice here how Bob of type L has no edges going back to K, hence Bob of that type need not even consider (or be aware of what happens in) block K. This makes sense, as Bob in L is modeled as an agent who is truly oblivious of Alice's threats.

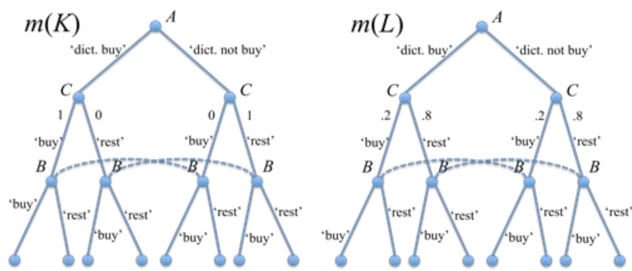

Figure 2: The trees for the Alice-Bob example

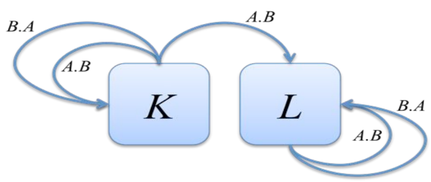

Figure 3: The belief graph for the Alice-Bob example

Identifying the game's Bayes-Nash equilibrium could be done as follows: First of all, Alice has two pure strategies in either block, 'dictate buy' and 'dictate rest.' Bob in turn has four strategies, 'do what C suggests,' 'do the opposite of what C suggests,' 'always buy' and 'always rest.' In tree $m(L)$ it is clearly optimal for Bob to 'do what C suggests,' as the agency's suggestion is predictive of the stock's performance. And of course, Alice cannot affect the situation, so any randomization between 'dictate buy' and 'dictate rest' constitutes an equilibrium strategy. In block K, however, things are different. First, the strategy 'do the opposite of what C suggests' is weakly dominated by 'never buy' for Bob; once removed, then 'dictate buy' becomes weakly dominant for Alice, in which case the dominant strategy for Bob becomes 'never buy.'

These would be the equilibria in each block, if that block was commonly held to be true by both agents—but how about the game as a whole, given the agents' complex beliefs? Clearly, the Bayes-Nash equilibrium would have Bob of type L playing 'do what C says,' as in that case Bob is unaware that Alice might even be of a different type. For Bob of type K, similarly, he should 'never buy,' as his beliefs in K are also trivial. As for Alice, if she is of type L she could do whatever, but if she is of type K, she needs to best-respond to a mixture consisting of Bob of type K (with 0.3 probability) and Bob of type L (with 0.7 probability). Hence, her optimal strategy overall would be to 'dictate buy.'

Notice that the equilibria are interdependent but subjective. We have not taken a position which of these blocks, K or L, accurately represents the real world (if any does). The notion of "truth" is not employed in the analysis. If K happens to be true, then Alice's threats will indeed be effective and Bob of type L will be most likely misled into an ill-adivsed purchase. If L happens to be true, there will be no effect and Bob's decision will be profitable to him. If it was important to us to identify what would happen in the real world, we could introduce a modeler agent whose beliefs correspond to what we believe the truth actually is.

7. SEMANTICS IN EXTENSIVE FORM

Thus far, we have allowed each block to be any game, with no connection stipulated between them. However, in order for the belief structures to make sense, there has to be some connection between the strategies and information in different games. In this section, we make this precise for extensive form games by describing how the beliefs of agents induce a *mapping* of information sets, and we give an interpretation to that mapping. First, let us define $F : N \times \mathcal{I} \to \Delta(\mathcal{I})$, where N is the set of agents, $\mathcal{I}$ is the set of information sets (across all trees $m(b)$), and $\Delta(\cdot)$ denotes the infinite set of probability distributions over the elements of a set. We shall constrain F such that $F(i, I)$, where i is an agent and I is an information set of agent $j \neq i$ in, say, tree b, is a probability distribution over "corresponding" information sets I' of j in trees b', such that I' of tree b' is assigned probability mass equal to $p^b_{ij}(b')$. The notion of "corresponding" information sets relies on the assumption that, for each information set I in each tree there exists another information set in every other tree that is indistinguishable from I with respect to the set of histories leading up to it, as well as the set of available moves to the agent at I; this is the *corresponding* information set of I.[2] In essence F describes the mapping among information sets *implied* by the beliefs $\beta(b)$.

This mapping is useful because it allows for (a) interpreting the beliefs in a conceptually simple and appealing way, and (b) verifying equilibria easily. In a game where the CPA holds, the various trees representing agents' views of the worlds can be brought together to form a single large tree. This can be done by introducing Nature nodes in the beginning of the tree. Nature would choose types according to the common prior. After Nature has assigned types to all players, we could append replicas of the game tree (as many as there are possible assignment vectors), then connect the information sets of all agents across all trees where they ought not to be able to distinguish between them (due to not being aware of others' actual types).

However, when the CPA is dropped, compiling a single tree is in general not possible. This is essentially because agents disagreeing on priors will disagree on the probabilities of Nature's moves when she assigns their types. Hence, we need a *collection* of trees. However, these trees should not be independent, and this is what our function F rectifies. Essentially, when $F(i, I)$ maps information set I to different information sets, agent i expects agent j, the one choosing at I, to behave as a mixture of agents optimally choosing at the information sets defined by the distribution $F(i, I)$.

Let us revisit our example to make this clearer. Let us denote as $K.A$ the information set of agent A in tree K, as $K.B1$ the leftmost information set of agent B in K (the one corresponding to the history "C suggests buy") and as $K.B2$ the rightmost information set of B in K. Then:

$F(B, K.A) = \langle 1 : K.A \rangle$

$F(A, K.B1) = \langle 0.3 : K.B1, 0.7 : L.B1 \rangle$

$F(A, K.B2) = \langle 0.3 : K.B2, 0.7 : L.B2 \rangle$

$F(B, L.A) = \langle 1 : L.A \rangle$

---

[2] In essence this restricts the various game trees to differ only in the utility values at their leaf nodes and the probabilities of Nature's moves.

$F(A, L.B1) = < 1 : L.B1 >$

$F(A, L.B2) = < 1 : L.B2 >$

where $< \{c_i : X_i\} >$ represents a probability distribution where each X_i has probability mass c_i.

This is to be interpreted as follows. In the game represented by tree $m(K)$, which is the game that agent A of type K believes, agent B is expected to behave *as if* the game was represented by the mapping of F. That is, although A believes that $m(K)$ is the true game, she also believes with probability 0.3 that B agrees, and with probability 0.7 that B thinks the game is actually $m(L)$.

When trees are linked this way, the typespace of the underlying BG is kept simple. This is because, at any point throughout the game tree $m(b)$, each player i believes another player j to be acting *as if* the actual world is given by a block drawn using the same distribution p_{ij}^b. Hence, it is straightforward to define b as a type for i, and denote its posterior $p(T_j \mid T_i = b) = p_{ij}^b$. Therefore, the typespace $T \equiv \times_i T_i$ is equivalent to B^n. If, on the other hand, each information set of j in $m(b)$ was mapped using a different distribution, then defining typesets ceases to be as straightforward and the resulting typespace might be quite large.

There is an important semantic distinction here. It is different to say "Alice believes that Bob thinks $m(K)$ is true with probability 0.3 and $m(L)$ is true with probability 0.7," and to say "Alice believes with probability 0.3 that Bob thinks $m(K)$ is true and with probability 0.7 that Bob thinks $m(L)$ is true." The former statement represents the case where Alice believes Bob to be himself uncertain of the true state of the world, hence his optimal strategy will take this uncertainty into consideration. This is *not* allowed by our formalism, in which agents are assumed to know their type (and hence their block). The second statement is what our structure captures. Here the uncertainty is not in Bob's mind, but in Alice's. Bob either uses $m(K)$ and behaves optimally with respect to that, or uses $m(L)$. It is Alice that doesn't know which block he uses and has uncertainty over this information.

Although the above distinction may seem like a technicality, it does profoundly affect the definition of equilibrium. If the former interpretation were true, Bob might reason that it is too risky to follow C's suggestion, as with 30% probability this is dictated by Alice. Hence, Alice of type K would have no reason to prefer either of her strategies over the other. As we have seen, following the second interpretation, we deduce that it is a dominant strategy for Alice of type K to 'dictate buy,' and for Bob of type L to follow his agency's suggestion—two vastly opposing results.

8. RELATED WORK

Other representations for games have been proposed as a solution to representational and conceptual aspects of game-theoretic analysis. In this section we review the literature and discuss how our formalism differs from existing approaches.

8.1 Relation to games of awareness

In [4] the authors extend game trees to represent situations where the available actions in each information set are not common knowledge to all players. In particular, player i might not know the available actions of player j, and he might be aware of this fact or not. In their formalism, a game consists of a collection of trees, each representing different levels of awareness, plus a function that maps each information set I of agent i in a tree T to an information set I' of i in another tree T', with the interpretation that i in T believes that the game actually being played is the one represented by T' (and hence the awareness of himself and all other players is given by the structure of T').

Our formalism is similar to these games in two ways: First, we also define our game to be a collection of blocks, and we map information sets in one tree $m(b)$ to information sets in other trees $m(b')$, $m(b'')$, etc. Second, our Bayes-Nash equilibrium coincides with the equilibrium defined for games of awareness, for all games that can be represented in both formalisms. However, there are a number of fundamental differences between games of awareness and our work. First, games of awareness deal with ordinary extensive form games rather than Bayesian games. Second, despite the lack of agreement and common knowledge about the structure of the game being played, a game of awareness can be converted to a single (possibly large) extensive form game. In contrast, as we have seen, a BG with uncommon priors cannot easily be represented as a single extensive form game. Third, while games of awareness allow for situations in which the agents do not have common knowledge about the game being played, in such cases [4] appears to view the game as being entirely subjective, rather than considering a rich belief structure as we do. Fourth, in some ways games of awareness are more expressive than our formalism. This is because we do not explicitly model awareness of available actions, as we assume that between every two trees there is a mapping of information sets that are indistinguishable with respect to available actions; hence, action sets are common knowledge. However, our lack of expressiveness in this regard is also a feature. We constrain the mapping function to express the beliefs $\beta(b)$ associated with each tree $m(b)$; therefore, if b is the block used by agent i, we do not allow two information sets of agent j in $m(b)$ to be mapped to information sets of other trees drawn with differing distributions; both need to be drawn according to the same distribution p_{ij}^b. This restriction allows us to define the typeset of the underlying Bayesian game in a straightforward, concise and natural way, something that we suspect is much harder in games of awareness. This is because, if the information sets of player j in b are mapped to information sets of other trees with differing distributions, we can no longer use the set of blocks B as the typeset T_i. Instead, we need to construct a typeset that is rich enough (and hence, large enough) to take into account all the compound probabilistic mappings implied by the underlying game of awareness.

8.2 Relation to NIDs

Networks of Influence Diagrams (NIDs, [3]) are a formalism much in the same spirit as our own. Like us, the authors use a collection of blocks to represent a game. Each block is a Multi-Agent Influence Diagram [6], with the addition of the so-called "mod nodes." For each decision node (rectangle) d in the MAID and each agent i, there is a mod node $M[d,i]$ indicating which block (or distribution of blocks) i believes the owner of d to be using when making her decision. NIDs come with an algorithm that transforms the NID into a MAID, which is then solved for an equilibrium. A powerful notion in solving NIDs is the distinction between the "optimal strategy," which is the best strategy for all players in a block, given the beliefs represented by the mod nodes, and the "actually-played strategy," the strategy the player will actually draw her actions from. These two might be different, because the player under discussion may be "irrationally" considering another block for her own decision-making purposes.

NIDs are different from our formalism in several ways: First, NIDs are reducible to MAIDs, which are further reducible to extensive form trees; hence, as we explained in the previous section, they cannot represent Bayesian games with non-common priors. Second, because each decision in a block has its own mod node, like games of awareness, deriving the corresponding typeset for the underlying Bayesian game is extremely unintuitive and it might be unnaturally large. Third, our formalism does not allow players to be unaware of (or uncertain about) their own type (block), while NIDs allow for these kinds of irrationalities. Finally, NIDs have a notion of "truth," namely, one of the blocks (called the top-level block) is how the real world is believed to be by an outside modeler agent. Our formalism has no explicit notion of a "true world", but a modeler agent can be added if desired.

8.3 Relation to other subjective games

Purely subjective Nash equilibria have been considered elsewhere. However, the subjectivity of equilibrium has received various interpretations. For example, in [5], equilibria are subjective because they arise through a dynamic process of learning and adaptation, and not through a full consideration of the game parameters. Also, in [2], an equilibrium is "self-confirming" in the subjective sense when the agents' actions in equilibrium (in a repeated setting) are consistent with (i.e., do not contradict) their subjective beliefs and observations. In our model the notion of subjectivity stems solely from the difference in the agents' prior distributions and is not related to any learning process. Furthermore, our framework provides a bridge from purely subjective, completely independent equilibria, through subjective but interdependent equilibria, to objective equilibria.

9. CONCLUSION

In this paper we have argued in favor of abandoning the common prior assumption in constructing Bayesian games when the situation being modeled makes it appropriate. We aimed to expose the representational complexities arising from abandoning the common prior, as well as the nuanced differences in how the equilibria of Bayesian games are to be interpreted in every case. We have also presented a formalism that makes representing games without a common prior more natural and conceptually appealing. The formalism is simple to use and helps expose features of the game that are otherwise difficult to consider, such the dependencies between agents' beliefs.

One issue has been, however, downplayed in our exposition. In particular, we have not addressed the issue of learning. Do agents update their beliefs over others' types as the game unfolds and how do they do it? This is a difficult issue that comes with additional challenges in our domain. For example, if i believes that agent j is using either block b or b', and he knows that j's optimal action is X if she is using b, or Y if she is using b', what should i believe after j has chosen X? It is not obvious that he should be certain that j is of type b, as this might create an easy way for j to manipulate i's beliefs in her favor and mask her true type. Moreover, if j actually plays Z, which is her optimal action in a block b'' deemed to have zero probability by i, what should i now believe? Furthermore, if j's optimal action in b is not unique, but there are two equilibria with optimal actions X and Y, respectively, what is the "rational" likelihood of j choosing X, given she is in fact using b? This last question cannot be answered without an equilibrium selection theory, but has profound implications for learning (specifically, Bayesian updating) in this model. Again, these learning difficulties are not so surprising considering the theoretical results in [1].

10. REFERENCES

[1] Eddie Dekel, Drew Fudenberg, and David K. Levine. Learning how to play Bayesian games. Discussion Papers 1322, Northwestern University, Center for Mathematical Studies in Economics and Management Science, December 2000.

[2] Drew Fudenberg and D. K. Levine. Self-confirming equilibrium. *Econometrica*, (61)523–546, 1993.

[3] Ya'akov Gal and Avi Pfeffer. Networks of influence diagrams: Reasoning about agents' beliefs and decision-making processes. *Journal of Artificial Intelligence*, 33:109-147, 2008.

[4] Joseph Y. Halpern and Leandro Chaves Rêgo. Extensive games with possibly unaware players. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 744–751, 2006

[5] Ehud Kalai and Ehud Lehrer. Subjective games and equilibria. Working Papers 875, California Institute of Technology, Division of the Humanities and Social Sciences, August 1993.

[6] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45(1):118–221, 2003.

[7] Stephen Morris. The common prior assumption in economic theory. *Economics and Philosophy*, 11:227–253, 1995.

[8] Luke S. Zettlemoyer, Brian Milch, and Lesle Pack Kaebling. Multi-agent filtering with infinitely nested beliefs. In Daphne Koller, Dale Schuurmans, Yoshua Bengio, and Léon Bottou, editors, *NIPS*, pages 1905–1912. MIT Press, 2008.

[9] Josheph Sakovics. Games of incomplete information without common knowledge priors. *ESE Discussion Papers 77*. Edinburgh School of Economics, University of Edimnburg, 2004.

[10] Roger B. Myerson. Harsanyi's Games with Incomplete Information, In *Management Science*, 2004 (special issue).

[11] Prashant Doshi and Piotr Gmytrasiewicz, On the Difficulty of Achieving Equilibria in Interactive POMDPs, *in AAAI*, 2006.

[12] Giacomo Bonanno, Klaus Nehring. *How to make sense of the common prior assumption under incomplete information*. In Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge, 3:133–146, Evanston, IL, 1998.