Extensive Games with Possibly Unaware Players*

Joseph Y. Halpern

Computer Science Department

Cornell University, U.S.A.

e-mail: halpern@cs.cornell.edu

Leandro Chaves Rêgo †

Statistics Department

Federal University of Pernambuco, Brazil

e-mail: leandro@de.ufpe.br

*A preliminary version of this work was presented at AAMAS06 conference in Hakodate, Japan, in May of 2006. This work was supported in part by NSF under grants CTC-0208535, ITR-0325453, and IIS-0534064, by ONR under grants N00014-00-1-03-41 and N00014-01-10-511, and by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grant N00014-01-1-0795. The second author was also supported in part by a scholarship from the Brazilian Government through the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq).

†Most of this work was done while the author was at the School of Electrical and Computer Engineering at Cornell University, U.S.A.

Abstract

Standard game theory assumes that the structure of the game is common knowledge among players. We relax this assumption by considering extensive games where agents may be unaware of the complete structure of the game. In particular, they may not be aware of moves that they and other agents can make. We show how such games can be represented; the key idea is to describe the game from the point of view of every agent at every node of the game tree. We provide a generalization of Nash equilibrium and show that every game with awareness has a generalized Nash equilibrium. Finally, we extend these results to games with *awareness of unawareness*, where a player i may be aware that a player j can make moves that i is not aware of, and to *subjective games*, where payers may have no common knowledge regarding the actual game and their beliefs are incompatible with a common prior.

keywords: Economic Theory, Foundations of Game Theory, Awareness, Solution Concepts.

1 Introduction

Standard game theory models implicitly assume that all significant aspects of the game (payoffs, moves available, etc.) are common knowledge among the players. While such common knowledge may seem unreasonable, there are well-known techniques going back to Harsanyi [12] for transforming a game where some aspects are not common knowledge to one where they are common knowledge. All these techniques assume that players are at least *aware* of all possible moves in the game. However, this is not always a reasonable assumption. For example, sleazy companies assume that consumers are not aware that they can lodge complaints if there are problems; in a war setting, having technology that an enemy is unaware of (and thus being able to make moves that the enemy is unaware of) can be critical; in financial markets, some investors may not be aware of certain investment strategies (complicated hedging strategies, for example, or tax-avoidance strategies).

In a standard game, a set of strategies is a Nash equilibrium if each agent's strategy is a best response to the other agents' strategies, so each agent i would continue playing its strategy even if i knew what strategies the other agents were using. To understand the relevance of adding the possibility of unawareness to the analysis of games, consider the game shown in Figure 1. One Nash equilibrium of this game has A playing across$_A$ and B playing down$_B$. However, suppose that A is not aware that B can play down$_B$. In that case, if A is rational, A will play down$_A$. Therefore, Nash equilibrium does not seem to be the appropriate solution concept here. Although A would play across$_A$ if A knew that B were going to play down$_B$, A cannot even contemplate this possibility, let alone know it.

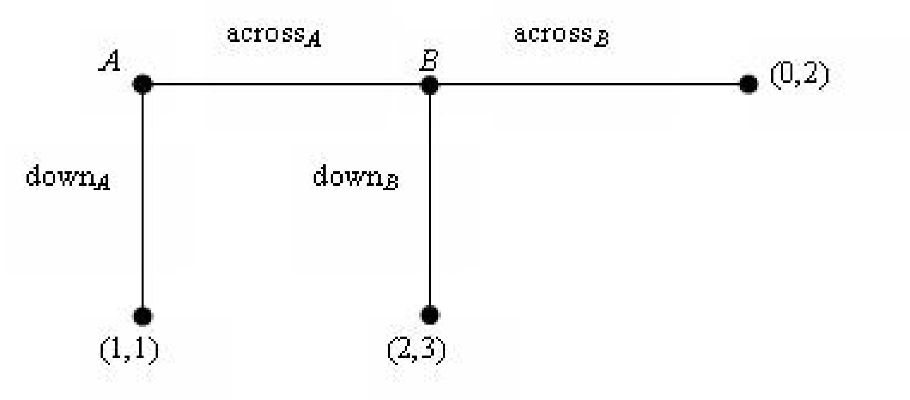

Figure 1: A simple game.

Our goal is to find appropriate solution concepts for extensive games with possibly unaware players, and more generally, to find ways of representing multiagent systems where some agents may not be aware of features of the system. To do this, we must first find an appropriate representation for such games. The first step in doing so is to explicitly represent what players are aware of at each node. We do this by using what we call an *augmented game*. An augmented game describes how awareness changes over time. For example, perhaps A playing across$_A$ will result in B becoming aware of the possibility of playing down$_B$. In financial settings, one effect of players using certain investment strategies is that other players become aware of the possibility of using that strategy. Strategic thinking in such games must involve taking this possibility into account.

We cannot in general represent what is going on using only one augmented game. The standard representation of a game implicitly assumes that (it is common knowledge that) the modeler and the players all understand the game the same way. This is no longer true once we allow for the possibility of unawareness, since a player's description of the game can now involve only those aspects of the game that he is aware of. Thus, the full description of the game with awareness is given by a set of augmented games, one for the modeler and one for each game that at least one of the agents thinks might be the true game in some situation.

Continuing with the game in Figure 1, the augmented game from the point of view of the type of B that is unaware of the possibility of playing down_B would just include A's moves down_A and across_A and the move across_B. In that augmented game, player A is also unaware of the move down_B. By way of contrast, the augmented game from the point of view of the type of B that is aware of down_B would include the move down_B, but may also allow for the possibility that A is not aware that B is aware of this move.

The standard notion of Nash equilibrium consists of a collection of strategies, one for each player. Our generalization consists of a collection of strategies, one for each pair (i, Γ'), where Γ' is a game that agent i considers to be the true game in some situation. Intuitively, the strategy for a player i at Γ' is the strategy i would play in situations where i believes that the true game is Γ'. To understand why we may need to consider different strategies consider, for example, the game of Figure 1. B would play differently depending on whether or not he was aware of down_B. Roughly speaking, a set of strategies, one for each pair (i, Γ'), is a *generalized Nash equilibrium* if the strategy for (i, Γ') is a best response for player i if the true game is Γ', given the strategies being used by the other players in Γ'.

We argue that this notion of equilibrium correctly captures our intuitions. We then show that every game with awareness has a generalized Nash equilibrium by associating with a game with awareness a standard game (where agents are aware of all moves) such that there is a one-to-one correspondence between generalized Nash equilibria of the game with awareness and Nash equilibria of the standard game.

For ease of exposition, for most of the paper we focus on games where agents are not aware of their lack of awareness. That is, we do not consider games where one player might be aware that there are moves that another player (or even she herself) might be able to make, although she is not aware of what they are. Such awareness of unawareness can be quite relevant in practice. For example, in the war setting described above, even if one side cannot conceive of a new technology available to the enemy, they might believe that there is some move available to the enemy without understanding what that particular move is. This, in turn, may encourage peace overtures. To take another example, an agent might delay making a decision because she considers it possible that she might learn about more possible moves, even if she is not aware of what these moves are.

If we interpret "lack of awareness" as "unable to compute" (cf. [2]), then awareness of unawareness becomes even more significant. Consider a chess game. Although all players understand in principle all the moves that can be made, they are certainly not aware of all consequences of all moves. A more accurate representation of chess would model this computational unawareness explicitly. We provide such a representation.

Roughly speaking, we capture the fact that player i is aware that, at a node h in the game tree, there is a move that j can make she (i) is not aware by having i's subjective representation of the game include a "virtual" move for j at node h. Since i might have only an incomplete understanding of what can happen after this move,

i simply describes what she believes will be the game after the virtual move, to the extent that she can. For example, if she has no idea what will happen after the virtual move, then she can describe her beliefs regarding the payoffs of the game. Thus, our representation can be viewed as a generalization of how chess programs analyze chess games. They explore the game tree up to a certain point, and then evaluate the board position at that point. We can think of the payoffs following a virtual move by j in i's subjective representation of a chess game as describing the evaluation of the board from i's point of view. This seems like a much more reasonable representation of the game than the standard complete game tree!

Our framework is flexible enough to deal with games where there is lack of common knowledge about what is the game being played. In particular, we can deal with lack of common knowledge regarding the utilities, who moves next, the structure of other players' information sets, and the probability of nature's moves (even in cases where there is no common prior compatible with the players' beliefs regarding nature).

Recently, Feinberg [3, 4] also studied games with awareness. Feinberg [4] gives a definition of extended Nash equilibrium in normal-form games. Although his definition stems from much the same intuitions as ours (although some details are different—see Section 6), it is expressed syntactically. Each player is characterized by a complete description of what moves and players he is aware of, what moves and players he is aware that each other player is aware of, and so on through all levels of iterated awareness. Feinberg [3] deals with extensive-form games and defines solution concepts only indirectly, via a syntactic epistemic characterization. His approach lacks a more direct semantic framework, which our model provides. He also does not deal with awareness of unawareness.

The rest of this paper is organized as follows. In Section 2, we describe how we represent different awareness levels in a game. In Section 3, we use our representation to define a generalized notion of Nash equilibrium, and we prove its existence in games with awareness. In Section 4, we describe how we can extend our approach to deal with awareness of unawareness. In Section 5, we describe how to extend our framework to deal with games where there is lack of common knowledge, even if awareness is not an issue. We compare our work to others in the literature, particularly Feinberg's, in Section 6, and conclude in Section 7.

2 Modeling awareness

The first step in dealing with awareness is modeling it. To this end, we consider *augmented games*. We start with a standard game, described by a game tree Γ (as in Figure 1). An augmented game Γ^+ *based on* Γ essentially augments Γ by describing each agent's *awareness level* at each node, where player i's awareness level at a node h is essentially the set of runs (complete histories) in Γ that i is aware of at node h. A player's awareness level may change over time, as the player becomes aware of more moves.

Our formal definition of augmented game is based on the definition of extensive game given by Osborne and Rubinstein [22]. We start by briefly reviewing Osborne and Rubinstein's definition.

A *(finite) extensive game* is a tuple $(N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$, where

- N is a finite set consisting of the players of the game.
- M is a finite set whose elements are the moves (or actions) available to players (and nature) during the game.[1]

[1]Osborne and Rubinstein did not make M explicit in their definition of an extensive game; we find it convenient to make

- H is a finite set of finite sequences of moves (elements of M) that is closed under prefixes, so that if $h \in H$ and h' is a prefix of H, then $h' \in H$. Intuitively, each member of H is a *history*. We can identify the nodes in a game tree with the histories in H. Each node n is characterized by the sequence of moves needed to reach n. A *run* in H is a terminal history, one that is not a strict prefix of any other history in H. Let Z denote the set of runs of H. Let $M_h = \{m \in M : h \cdot \langle m \rangle \in H\}$ (where we use $\cdot$ to denote concatenation of sequences); M_h is the set of moves that can be made after history h.

- $P : (H - Z) \to N \cup \{c\}$ is a function that assigns to each nonterminal history h a member of $N \cup \{c\}$. (We can think of c as representing nature.) If $P(h) = i$, then player i moves after history h; if $P(h) = c$, then nature moves after h. Let $H_i = \{h : P(h) = i\}$ be the set of all histories after which player i moves.

- f_c is a function that associates with every history for which $P(h) = c$ a probability measure $f_c(\cdot \mid h)$ on M_h. Intuitively, $f_c(\cdot \mid h)$ describes the probability of nature's moves once history h is reached.

- $\mathcal{I}_i$ is a partition of H_i with the property that if h and h' are in the same cell of the partition then $M_h = M_{h'}$, i.e., the same set of moves is available at every history in a cell of the partition. Intuitively, if h and h' are in the same cell of $\mathcal{I}_i$, then h and h' are indistinguishable from i's point of view; i considers history h' possible if the actual history is h, and vice versa. A cell $I \in \mathcal{I}_i$ is called an (i-)*information set*.

- $u_i : Z \to \mathrm{R}$ is a payoff function for player i, assigning a real number (i's payoff) to each run of the game.

---

it explicit here.

In the game of Figure 1,

- $N = \{A, B\}$, $H = \{\langle\,\rangle, \langle \text{down}_A\rangle, \langle \text{across}_A, \text{down}_B\rangle, \langle \text{across}_A, \text{across}_B\rangle\}$,
- $P(\langle\,\rangle) = A$, $P(\langle \text{across}_A\rangle) = B$,
- $I_A = \{\langle\,\rangle\}$, $I_B = \{\langle \text{across}_A\rangle\}$,
- $u_A(\langle \text{down}_A\rangle) = u_B(\langle \text{down}_A\rangle) = 1$,
- $u_A(\langle \text{across}_A, \text{across}_B\rangle) = 0$, and
- $u_B(\langle \text{across}_A, \text{across}_B\rangle) = 2$.

In this paper, as in most work in game theory, we further assume that players have *perfect recall*: they remember all the actions that they have performed and the information sets they passed through. Formally, we require that

- if h and h' are in the same i-information set and h_1 is a prefix of h such that $P(h_1) = i$, then there is a prefix h_1' of h' such that h_1 and h_1' are in the same information set; moreover, if $h_1 \cdot \langle m\rangle$ is a prefix of h (so that m was the action performed when h_1 was reached in h) then $h_1' \cdot \langle m\rangle$ is a prefix of h'.

An *augmented game* is defined much like an extensive game; the only essential difference is that at each nonterminal history we not only determine the player moving but also her awareness level. Since the awareness level is a set of runs in a game Γ, we say that $\Gamma^+ = (N^+, M^+, H^+, P^+, f_c^+, \{\mathcal{I}_i^+ : i \in N^+\}, \{u_i^+ : i \in N^+\}, \{A_i^+ : i \in N^+\})$ is an *augmented game based on the (standard) extensive game* $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$ if the following conditions are satisfied:

A1. $(N^+, M^+, H^+, P^+, f_c^+, \{\mathcal{I}_i^+ : i \in N^+\}, \{u_i^+ : i \in N^+\})$ is a (standard) finite extensive game where players have perfect recall.

A2. $A_i^+ : H_i^+ \rightarrow 2^H$ describes i's awareness level at each nonterminal history after which player i moves. For each $h \in H_i^+$, $A_i^+(h)$ consists of a set of histories in H and all their prefixes. Intuitively, $A_i^+(h)$ describes the set of histories of Γ that i is aware of at history $h \in H_i^+$. (Having $A_i^+(h)$ consist of histories rather than just runs makes it easier to deal with awareness of unawareness.)

A3. $N^+ \subseteq N$.

A4. If $P^+(h) \in N^+$, then $P^+(h) = P(\overline{h})$, where $\overline{h}$ is the subsequence of h consisting of all the moves in h that are also in M, and $M_h^+ \subseteq M_{\overline{h}}$. Intuitively, all the moves available to i at h must also be available to i in the underlying game Γ.

A5. If $P^+(h) = c$, then either $P(\overline{h}) = c$ and $M_h^+ \subseteq M_{\overline{h}}$, or $M_h^+ \cap M = \emptyset$. The moves in M_h^+ in the case where $M_h^+ \cap M = \emptyset$ intuitively capture uncertainty regarding a player's awareness level.

A6. If h and h' are in the same information set in $\mathcal{I}_i^+$, then $A_i^+(h) = A_i^+(h')$. Intuitively, i's awareness level depends only on the information that i has.

A7. If h is a prefix of h' and $P^+(h) = P^+(h')$, then $A_i^+(h) \subseteq A_i^+(h')$. This is a perfect recall requirement; players do not forget histories that they were aware of.

A8. If h and h' are in the same information set in Γ^+, then $\overline{h}$ and $\overline{h}'$ are in the same information set in Γ.

A9. If h and h' are histories in both Γ^+ and Γ, and $\overline{h}$ and $\overline{h}'$ are in the same information set in Γ, then h and h' are in the same information set in Γ^+.

A10. For all $i \in N^+$ and $h \in H_i^+$, if $h', h'' \in A_i(h)$, h' and h'' are in the same information set in Γ, then $h' \cdot \langle m \rangle \in A_i(h)$ iff $h'' \cdot \langle m \rangle \in A_i(h)$.

A11. $\{\overline{z} : z \in Z^+\} \subseteq Z$; moreover, for all $i \in N^+$, $h \in H_i^+$, if z is a terminal history in $A_i^+(h)$ (i.e., if $z \in A_i^+(h)$ and z is not a strict prefix of another element of $A_i^+(h)$), then $z \subseteq Z$. Thus, the runs in Z^+ correspond to runs in Z, and players understand this fact.

A12. For all $i \in N^+$ and runs z in Z^+, if $\overline{z} \in Z$, then $u_i^+(z) = u_i(\overline{z})$. Thus, a player's utility just depends on the moves made in the underlying game. (By A11, we have $\overline{z} \in Z$. We have included the clause "if $\overline{z} \in Z$" so that A12 is applicable when we consider awareness of unawareness, where we drop A11.)

Conditions A1–A12 are intended to capture our intuitions regarding information sets, awareness, and common knowledge. To allow us to focus on issues directly related to awareness, we have implicitly assumed that there is common knowledge of (1) who moves at histories in the underlying game (this is captured by the fact that $P^+(h) = P(\overline{h})$ unless $P^+(h) = c$ and $M_h^+ \cap M = \emptyset$ —either the same player or nature moves at both h and $\overline{h}$ unless nature makes an "awareness" move at h), (2) what the payoffs are in the underlying game (since $u^+(\overline{z}) = u(z)$), and (3) what the information sets are in the underlying game (see A8–A10). Our approach is flexible enough to allow us to drop these assumptions; see Section 5.

To understand A8–A10, we must first discuss our view of information sets. As pointed out by Halpern [7], special attention must be given to the interpretation of information sets in game trees. This issue requires even more care in games with awareness. The standard intuition for information sets is that a player considers all the histories in his information set possible. But this intuition does not apply in augmented games. In an augmented game, there may be some histories in an i-information set that i

is not aware of; player i cannot consider these histories possible. For example, consider finitely repeated prisoners dilemma where Alice and Bob each move twice before their moves are revealed. Even if Bob is not aware of defection, his information set after Alice's first move in the modeler's game will still contain the history where Alice defects.

We interpret an i-information set to be the set of all histories where player i has the same *local state*. Intuitively, this local states encodes all the information that i has about the moves he can make, what moves have been made, the other players in the game, his strategy, and so on. We assume that player i's local state is characterized by the sequence of signals that that i has received in the course of the game. Therefore, h and h' are in the same i-information set in Γ iff i received the same sequence of signals in both histories.

In standard extensive games, the sequence of signals a player receives after every history h is assumed to be common knowledge. (This assumption is implicit in the assumption that the game, is common knowledge, and hence so are the information sets.) As we said, we continue to assume this in games with awareness (although we show how the assumption can be dropped in Section 5). That is why we require in A8 that if h and h' are in the same i-information in an augmented game, then $\overline{h}$ and $\overline{h}'$ must be in the same i-information set in the underlying game. The converse of A8 does not necessarily hold. It could well be the case that $\overline{h}$ and $\overline{h}'$ are in the same i-information set, but since i receives different signals from nature, h and h' are not in the same information set. On the other hand, if all the moves in h and h' are already in the underlying game, then if h and h' are in the same information set of Γ, they should be in the same information set of Γ^+. This is the content of A9. Since, the signals received by a player determine the moves he has available, if player i is aware of two histories in the same information set in Γ, he must be aware of the same set of moves available at both of these histories. A10 captures that intuition.

For the remainder of the paper, we use the following notation: for a (standard or augmented) game Γ^s, we denote the components of Γ^s with the same superscript s, so that we have M^s, H^s, and so on. Thus, from here on we do not explicitly describe the components of a game.

An augmented game describes either the modeler's view of the game or the subjective view of the game of one of the players, and includes both moves of the underlying game and moves of nature that change awareness. For example, consider again the game shown in Figure 1 and suppose that

- players A and B are aware of all histories of the game;
- player A is uncertain as to whether player B is aware of run $\langle \text{across}_A, \text{down}_B \rangle$ and believes that he is unaware of it with probability p; and
- the type of player B that is aware of the run $\langle$across$_A$, down$_B\rangle$ is aware that player A is aware of all histories, and he knows A is uncertain about his awareness level and knows the probability p.

Because A and B are actually aware of all histories of the underlying game, from the point of view of the modeler, the augmented game is essentially identical to the game described in Figure 1, with the awareness level of both players A and B consisting of all histories of the underlying game. However, when A moves at the node labeled A in the modeler's game, she believes that the actual augmented game is Γ^A, as described in Figure 2. In Γ^A, nature's initial move captures A's uncertainty about B's awareness level. At the information set labeled $A.1$, A is aware of all the runs of the underlying game. Moreover, at this information set, A believes that the true game is Γ^A.

At the node labeled $B.1$, B is aware of all the runs of the underlying game and believes that the true game is the modeler's game; but at the node labeled $B.2$, B is not aware that he can play down$_B$, and so believes that the true game is the augmented game Γ^B described in Figure 3. At the nodes labeled $A.3$

and $B.3$ in the game Γ^B, neither A nor B is aware of the move down$_B$. Moreover, both players think the true game is Γ^B.

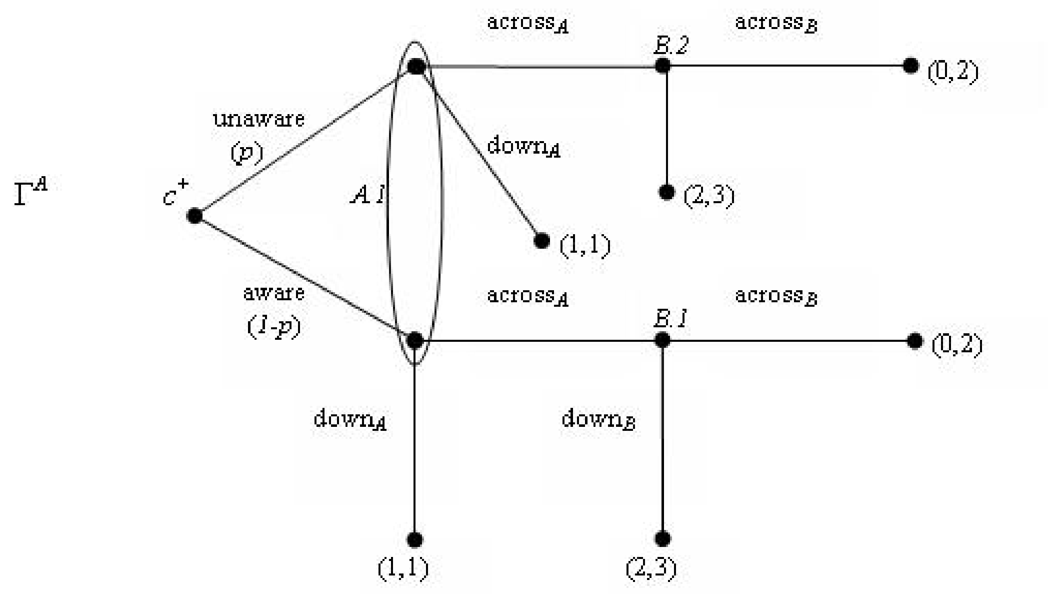

Figure 2: The augmented game Γ^A.

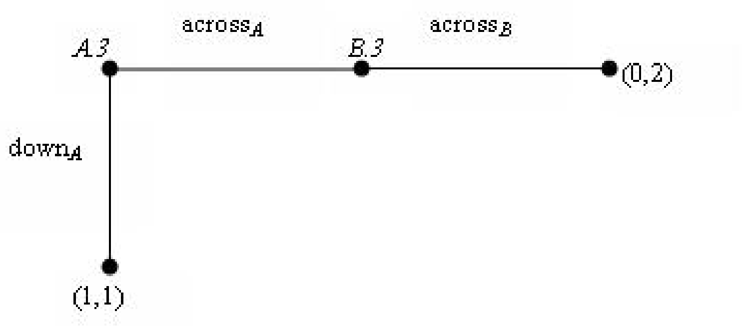

Figure 3: The augmented game Γ^B.

As this example should make clear, to model a game with possibly unaware players, we need to consider not just one augmented game, but a collection of them. Moreover, we need to describe, at each history in an augmented game, which augmented game the player playing at that history believes is the actual augmented game being played.

To capture these intuitions, we define a *game with awareness based on* $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$ to be a tuple $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$, where

- $\mathcal{G}$ is a countable set of augmented games based on Γ, of which one is Γ^m;
- $\mathcal{F}$ maps an augmented game $\Gamma^+ \in \mathcal{G}$ and a history h in Γ^+ such that $P^+(h) = i$ to a pair (Γ^h, I), where $\Gamma^h \in \mathcal{G}$ and I is an i-information set in game Γ^h.

Intuitively, Γ^m is the game from the point of view of an omniscient modeler. If player i moves at h in game $\Gamma^+ \in \mathcal{G}$ and $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$, then Γ^h is the game that i believes to be the true game when the history is h, and I consists of the set of histories in Γ^h he currently considers possible. For example, in the examples described in Figures 2 and 3, taking Γ^m to be the augmented game in Figure 1, we have $\mathcal{F}(\Gamma^m, \langle\,\rangle) = (\Gamma^A, I)$, where I is the information set labeled $A.1$ in Figure 2, and $\mathcal{F}(\Gamma^A, \langle \text{unaware,across}_A \rangle) = (\Gamma^B, \{\langle \text{across}_A \rangle\})$.

It may seem that by making $\mathcal{F}$ a function we cannot capture a player's uncertainty about the game being played. However, we can capture such uncertainty by folding it into nature's move. For example, we capture A's uncertainty about whether B is aware of being able to move down$_B$ in the augmented game Γ^A illustrated in Figure 2 by having nature decide this at the first step. It should be clear that this gives a general approach to capturing such uncertainty.

The augmented game Γ^m and the mapping $\mathcal{F}$ must satisfy a number of consistency conditions. The first set of conditions applies to Γ^m. Since the modeler is presumed to be omniscient, the conditions say that the modeler is aware of all the players and moves of the underlying game.

M1. $N^m = N$.

M2. $M \subseteq M^m$ and $\{\overline{z} : z \in Z^m\} = Z$.

M3. If $P^m(h) \in N$, then $M^m_h = M_{\overline{h}}$. If $P^m(h) = c$, then either $M^m_h \cap M = \emptyset$ or $M^m_h = M_{\overline{h}}$ and $f^m_c(\cdot \mid h) = f_c(\cdot \mid \overline{h})$.

M1, M2 and M3 enforce the intuition that the modeler understands the underlying game. He knows all the players and possible moves, and understands how nature's moves work in the underlying game Γ. It may seem somewhat surprising that there is no analogue of the second part of M3 (i.e., the constraint of f^m_c) for all augmented games. While it makes sense to have such an analogue if nature's moves are in some sense objective, it seems like an unreasonable requirement that all player's should agree on these probabilities in general. This is especially so in the case that a player suddenly becomes aware of some moves of nature that he was not aware of before. It does not seem reasonable to assume that this awareness should come along with an understanding of the probabilities of these moves. Of course, we could require such an analogue of M3. Since the set of games that have such a requirement is a subset of the games we consider, all our results apply without change if such a requirement is imposed.

Although the modeler understands the underlying game Γ, Γ^m is not uniquely determined by Γ. There may be many modeler's games based on Γ, where the players have different awareness levels and the awareness changes in different ways.

The game Γ^m can be thought of as a description of "reality"; it describes the effect of moves in the underlying game and how players' awareness levels change. The other games in $\mathcal{G}$ describe a player's subjective view of the situation. The constraints on the mapping $\mathcal{F}$ that we now describe capture desirable properties of awareness.

Consider the following constraints, where $\Gamma^+ \in \mathcal{G}$, $h \in H^+$, $P^+(h) = i$, $A^+_i(h) = a$, and $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$.

C1. $\{\overline{h} : h \in H^h\} = a$.

C2. If $h' \in H^h$ and $P^h(h') = j$, then $A_j^h(h') \subseteq a$ and $M_{\overline{h}'} \cap \{m : \overline{h}' \cdot \langle m \rangle \in a\} = M_{h'}^h$.

C3. If h' and h are in the same information set in Γ^+ and $\overline{h}' \in a$, then there exists $h'' \in I$ such that $\overline{h}'' = \overline{h}'$.

C4. If $h' \in I$, then $A_i^h(h') = a$ and $\mathcal{F}(\Gamma^h, h') = (\Gamma^h, I)$.

C5. If $h' \in H^+$, $P^+(h') = i$, $A_i^+(h') = a$, then if h and h' are in the same information set of Γ^+, then $\mathcal{F}(\Gamma^+, h') = (\Gamma^h, I)$, while if h is a prefix or a suffix of h', then $\mathcal{F}(\Gamma^+, h') = (\Gamma^h, I')$ for some i-information set I'.

C6. If $h' \in I$, then $\overline{h}$ and $\overline{h}'$ are in the same information set in Γ;

C7. If $\Gamma^h = \Gamma^+$, then $h' \in I$ iff h and h' are in the same i-information set in Γ^+.

C8. For all histories $h' \in I$, there exists a prefix h_1' of h' such that $P^h(h_1') = i$ and $\mathcal{F}(\Gamma^h, h_1') = (\Gamma', I')$ iff there exists a prefix h_1 of h such that $P^+(h_1) = i$ and $\mathcal{F}(\Gamma^+, h_1) = (\Gamma', I')$. Moreover, $h_1' \cdot \langle m \rangle$ is a prefix of h' iff $h_1 \cdot \langle m \rangle$ is a prefix of h.

C9. There exists a history $h' \in I$ such that for every prefix $h'' \cdot \langle m \rangle$ of h', if $P^h(h'') = j \in N^h$ and $\mathcal{F}(\Gamma^h, h'') = (\Gamma', I')$, then for all $h_1 \in I'$, $h_1 \cdot \langle m \rangle \in H'$.

C10. If h' and h'' are histories in both Γ^+ and Γ^h, then h' and h'' are in the same i-information set in Γ^+ iff h' and h'' are in the same i-information set in Γ^h.

Suppose that $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$. Player i moving at history h in Γ^+ thinks the actual game is Γ^h. Moreover, i thinks he is in the information set of I of Γ^h. C1 guarantees that the set of histories of the underlying game player i is aware of is exactly the set of histories of the underlying game that appear in Γ^h. C2 states that no player in Γ^h can be aware of histories not in a. The second part of C2 implies that the set of moves available to player j at h' is just the set of moves that player i is aware of that are available to j at $\overline{h}'$ in the underlying game. C3 guarantees that for all histories h' indistinguishable from h that player i is aware of, there exists some history $h'' \in I$ differing from h' at most in some moves of nature that change awareness levels. C4 says that at all histories in I player i indeed thinks the game is Γ^h and that the information set is I. C5 says that player i's subjective view of the game changes only if i becomes aware of more moves and is the same at histories in H^+ that i cannot distinguish. C6 captures the assumption that at all histories i considers possible, he must have gotten the same signals as he does in the actual history.

C7 says that if while moving at history h player i thinks that Γ^+ is the actual game, then he considers possible all and only histories in the information set containing h. C8 is a consequence of the perfect recall assumption. C8 says that if, at history h, i considers h' possible, then for every prefix h'_1 of h' there is a corresponding prefix of h where i considers himself to be playing the same game, and similarly, for every prefix of h there is a prefix of h' where i considers himself to be playing the same game. Moreover, i makes the same move at these prefixes.

The intuition behind condition C9 is that player i knows that player j only make moves that j is aware of. Therefore, player i must consider at least one history h' where he believes that every player j made a move that j was aware of. It follows from A11, C1, C2, and C9 that there is a run going through I where every player j makes a move that player i believes that j is aware of.

Since we assume that players have (modulo awareness) common knowledge about information sets, if Γ^+ is the game from the point of view of player j (or the modeler) and there are histories h' and h'' in both Γ^+ and Γ^h, then player j (or the modeler) knows that player i gets the same signals in both h' and

h'' iff he knows that player i knows that he gets the same signals in those histories. C10 captures that intuition.

Just as Γ^m is not uniquely determined by Γ, $\mathcal{F}(\Gamma^+, h)$ depends on more than just the awareness level of the player who moves at h. That is, even if $A_i(h) = A_i(h')$, we may have $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$ and $\mathcal{F}(\Gamma^+, h') = (\Gamma^{h'}, I')$ with $\Gamma^h \neq \Gamma^{h'}$. We do not require that the awareness level determines the game a player considers possible. This extra flexibility allows us to model a situation where, for example, players 2 and 3, who have the same awareness level and agree on the awareness level of player 1, have different beliefs about the game player 1 considers possible.[2]

A standard extensive game Γ can be identified with the game $(\{\Gamma^m\}, \Gamma^m, \mathcal{F})$, where (abusing notation slightly) $\Gamma^m = (\Gamma, \{A_i : i \in N\})$ and, for all histories h in an i-information set I in Γ, $A_i(h) = H$ and $\mathcal{F}(\Gamma^m, h) = (\Gamma^m, I)$. Thus, all players are aware of all the runs in Γ, and agree with each other and the modeler that the game is Γ. We call this the *canonical representation of* Γ as a game with awareness.

One technical issue: We have assumed that the set $\mathcal{G}$ of games in a game Γ^* with awareness is countable. For our purposes, this is without loss of generality. We are ultimately interested in what happens in the game Γ^m, since this is the game actually being played. However, to analyze that, we need to consider what happens in other games in $\mathcal{G}$. For example, if h is a history in Γ^m where i moves, we need to understand what happens in the game Γ^h such that $\mathcal{F}(\Gamma^m, h) = (\Gamma^h, \cdot)$, since Γ^h is the game that i thinks is being played at history h in Γ^m. It is not hard to see that the set of games we need

[2]If the beliefs of players 2 and 3 regarding 1 are compatible with a common prior, then we can view players 2 and 3 as considering different information sets in the same game possible. However, if their beliefs are not compatible with a common prior, for example, if player 2 believes that player 1 believes that, in history h, Γ_1 is the actual game with probability 1, and player 3 believes that, in history h, player 1 believes that Γ_2 is the actual game with probability 1, where $\Gamma_1 \neq \Gamma_2$, then we cannot view players 2 and 3 as considering the same game possible.

to consider is the least set $\mathcal{G}'$ such that $\Gamma^m \in \mathcal{G}'$ and, for every $\Gamma' \in \mathcal{G}$ and history h in Γ' such that $\mathcal{F}(\Gamma', h) = (\Gamma'', \cdot)$, $\Gamma'' \in \mathcal{G}'$. $\mathcal{G}'$ is guaranteed to be countable, even if $\mathcal{G}$ is not.

3 Local strategies and generalized Nash equilibrium

3.1 Local Strategies

In this section, we generalize the notion of Nash equilibrium to games with awareness. To do that, we must first define what a strategy is in a game with awareness. Recall that in a standard game, a strategy for player i is a function from i-information sets to a move or to a distribution over moves, depending on whether we are considering *pure* (i.e., deterministic) strategies or *behavioral* (i.e., randomized) strategies. The intuition is that player i's actions depend on what i knows; the strategy can be viewed as a *universal plan*, describing what i will do in every possible situation that can arise. This makes sense only because i is presumed to know the game tree, and thus to know in advance all the situations that can arise.

In games with awareness, this intuition no longer makes sense. For example, player i cannot plan in advance for what will happen if he becomes aware of something he is initially unaware of. We must allow i's strategy to change if he becomes aware of more moves. Let $\mathcal{G}_i = \{\Gamma' \in \mathcal{G} :$ for some $\Gamma^+ \in \mathcal{G}$ and h in Γ^+, $P^+(h) = i$ and $\mathcal{F}(\Gamma^+, h) = (\Gamma', \cdot)\}$. Intuitively, $\mathcal{G}_i$ consists of the games that i views as the real game in some history. Thus, rather than considering a single strategy in a game $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ with awareness, we consider a collection $\{\sigma_{i,\Gamma'} : \Gamma' \in \mathcal{G}_i\}$ of what we call *local strategies*, one for each augmented game in $\mathcal{G}_i$. Intuitively, a local strategy $\sigma_{i,\Gamma'}$ for game Γ' is the strategy that i would use if i were called upon to play and i thought that the true game was Γ'. Thus, the domain of $\sigma_{i,\Gamma'}$ consists of pairs (Γ^+, h) such that $\Gamma^+ \in \mathcal{G}$, h is a history in Γ^+, $P^+(h) = i$, and $\mathcal{F}(\Gamma^+, h) = (\Gamma', I)$.

Define an equivalence relation $\sim_i$ on pairs (Γ', h) such that $\Gamma' \in \mathcal{G}$ and h is a history in Γ' where i moves such that $(\Gamma_1, h_1) \sim_i (\Gamma_2, h_2)$ if $\mathcal{F}(\Gamma_1, h_1) = \mathcal{F}(\Gamma_2, h_2)$. We can think of $\sim_i$ as defining a *generalized information partition* in Γ^*. It is easy to check that a $\sim_i$ equivalence class consists of a union of i-information sets in individual games in $\mathcal{G}$. Moreover, if some element of a $\sim_i$ equivalence class is in the domain of $\sigma_{i,\Gamma'}$, then so is the whole equivalence class. At all pairs (Γ', h') in a $\sim_i$ equivalence class, if $\mathcal{F}(\Gamma', h') = (\Gamma^{h'}, I)$, player i thinks he is actually playing in the information set I of $\Gamma^{h'}$. Thus, we require that $\sigma_{i,\Gamma'}(\Gamma_1, h_1) = \sigma_{i,\Gamma'}(\Gamma_2, h_2)$ if (Γ_1, h_1) and (Γ_2, h_2) are both in the domain of $\sigma_{i,\Gamma'}$ and $(\Gamma_1, h_1) \sim_i (\Gamma_2, h_2)$.

The following definition summarizes this discussion.

Definition 3.1 *Given a game with awareness $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$, a* local strategy *$\sigma_{i,\Gamma'}$ for agent i is a function mapping pairs (Γ^+, h) such that h is a history where i moves in Γ^+ and $\mathcal{F}(\Gamma^+, h) = (\Gamma', I)$ to a probability distribution over $M'_{h'}$, the moves available at a history $h' \in I$, such that $\sigma_{i,\Gamma'}(\Gamma_1, h_1) = \sigma_{i,\Gamma'}(\Gamma_2, h_2)$ if $(\Gamma_1, h_1) \sim_i (\Gamma_2, h_2)$.*

Note that there may be no relationship between the strategies $\sigma_{i,\Gamma'}$ for different games Γ'. Intuitively, this is because discovering about the possibility of a different move may cause agent i to totally alter his strategy. We could impose some consistency requirements, but we have not found any that we believe should hold in all games. We believe that all our results would continue to hold in the presence of reasonable additional requirements, although we have not explored the space of such requirements.

3.2 Generalized Nash Equilibrium

We want to define a notion of generalized Nash equilibrium so as to capture the intuition that for every player i, if i believes he is playing game Γ', then his local strategy $\sigma_{i,\Gamma'}$ is a best response to the local strategies of other players in Γ'.

Define a *generalized strategy profile* of $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ to be a set of local strategies $\vec{\sigma} = \{\sigma_{i,\Gamma'} : i \in N, \Gamma' \in \mathcal{G}_i\}$. Let $EU_{i,\Gamma'}(\vec{\sigma})$ be the expected payoff for i in the game Γ' given that strategy profile $\vec{\sigma}$ is used. Note that the only strategies in $\vec{\sigma}$ that are needed to compute $EU_{i,\Gamma'}(\vec{\sigma})$ are the strategies actually used in Γ'; indeed, all that is needed is the restriction of these strategies to information sets that arise in Γ'.

A *generalized Nash equilibrium* of $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ is a generalized strategy profile $\vec{\sigma}$ such that for all $\Gamma' \in \mathcal{G}_i$, the local strategy $\sigma_{i,\Gamma'}$ is a best response to $\vec{\sigma}_{-(i,\Gamma')}$, where $\vec{\sigma}_{-(i,\Gamma')}$ is the set of all local strategies in $\vec{\sigma}$ except $\sigma_{i,\Gamma'}$.

Definition 3.2 *A generalized strategy profile $\vec{\sigma}^*$ is a* generalized Nash equilibrium *of a game $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ with awareness if, for every player i, game $\Gamma' \in \mathcal{G}_i$, and local strategy σ for i in Γ',*

$$EU_{i,\Gamma'}(\vec{\sigma}^*) \geq EU_{i,\Gamma'}((\vec{\sigma}^*_{-(i,\Gamma')}, \sigma)).$$

The standard definition of Nash equilibrium would say that $\vec{\sigma}$ is a Nash equilibrium if σ_i is a best response to $\vec{\sigma}_{-i}$. This definition implicitly assumes that player i can choose a whole strategy. This is inappropriate in our setting. An agent cannot anticipate that he will become aware of more moves. Essentially, if $\Gamma_1 \neq \Gamma_2$, we are treating player i who considers the true game to be Γ_1 to be a different agent from the version of player i who considers Γ_2 to be the true game. To understand why this is appropriate, suppose that player i considers Γ_1 to be the true game, and then learns about more moves, and so considers Γ_2 to be the true game. At that point, it is too late for player i to change the strategy he was playing when he thought the game was Γ_1. He should just try to play optimally for what he now considers the true game. Moreover, while player i thinks that the game Γ_1 is the true game, he never considers it possible that he will ever be playing a different game, so that he cannot "prepare himself" for a change in his subjective view of the game.[3] These considerations suggest that our notion of Nash equilibrium is appropriate.

It is easy to see that $\vec{\sigma}$ is a Nash equilibrium of a standard game iff $\vec{\sigma}$ is a (generalized) Nash equilibrium of the canonical representation of Γ as a game with awareness. Thus, our definition of generalized Nash equilibrium generalizes the standard definition.

Consider the game with awareness shown in Figures 1 (taking this to be Γ^m), 2, and 3. We have $\mathcal{G}_A = \{\Gamma^A, \Gamma^B\}$ and $\mathcal{G}_B = \{\Gamma^m, \Gamma^B\}$. Taking $dom(\sigma_{i,\Gamma'})$ to denote the domain of the strategy $\sigma_{i,\Gamma'}$, we have

$$dom(\sigma_{A,\Gamma^A}) = \{(\Gamma^m, \langle\,\rangle), (\Gamma^A, \langle\text{unaware}\rangle), (\Gamma^A, \langle\text{aware}\rangle)\},$$

$$dom(\sigma_{B,\Gamma^m}) = \{(\Gamma^m, \langle\text{across}_A\rangle), (\Gamma^A, \langle\text{aware, across}_A\rangle)\},$$

$$dom(\sigma_{A,\Gamma^B}) = \{(\Gamma^B, \langle\,\rangle)\}, \text{ and}$$

$$dom(\sigma_{B,\Gamma^B}) = \{(\Gamma^A, \langle\text{unaware,across}_A\rangle), (\Gamma^B, \langle\text{across}_A\rangle)\}.$$

Each of these domains consists of a single generalized information set. If $p < 1/2$, then there exists a generalized Nash equilibrium where $\sigma_{A,\Gamma^A} = \text{across}_A$, $\sigma_{A,\Gamma^B} = \text{down}_A$, $\sigma_{B,\Gamma^m} = \text{down}_B$, $\sigma_{B,\Gamma^B} = \text{across}_B$. Thus, in the modeler's game, A plays across$_A$, B plays down$_B$, and the resulting payoff vector is $(2, 3)$. On the other hand, if $p > 1/2$, then there exists a generalized Nash equilibrium where $\sigma_{A,\Gamma^A} = \text{down}_A$, $\sigma_{A,\Gamma^B} = \text{down}_A$, $\sigma_{B,\Gamma^m} = \text{down}_B$, $\sigma_{B,\Gamma^B} = \text{across}_B$. Thus, in the modeler's game, A plays down$_A$, and the payoff vector is $(1, 1)$. Intuitively, even though both A and B are aware of all the moves in the modeler's game, A considers it sufficiently likely that B is not aware of down$_B$, so A plays down$_A$. There exists another generalized Nash equilibrium where $\sigma_{A,\Gamma^A} = \text{down}_A$, $\sigma_{A,\Gamma^B} = \text{down}_A$, $\sigma_{B,\Gamma^m} = \text{across}_B$, and $\sigma_{B,\Gamma^B} = \text{across}_B$ that holds for any value of p. Intuitively, A believes B will

[3] In games with awareness of unawareness, an agent may consider it possible that he will become aware of more information. But this too is incorporated in his view of the game, so he can still do no better than playing optimally in his current view of the game.

play across$_B$ no matter what he (B) is aware of, and therefore plays down$_A$; given that A plays down$_A$, B cannot improve by playing down$_B$ even is he is aware of that move.[4]

We now show that every game with awareness has at least one generalized Nash equilibrium. We proceed as follows. Given a game $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ with awareness, let ν be a probability on $\mathcal{G}$ that assigns each game in $\mathcal{G}$ positive probability. (Here is where we use the fact that $\mathcal{G}$ is countable.) We construct a standard extensive game Γ^ν by essentially "gluing together" all the games $\Gamma' \in \mathcal{G}$, except that we restrict to the histories in Γ' that can actually be played according to the players' awareness level. Formally, for each $\Gamma' \in \mathcal{G}$, we restrict to the histories $\lfloor H' \rfloor = \{h \in H'$: for every prefix $h_1 \cdot \langle m \rangle$ of h, if $P'(h_1) = i \in N$ and $\mathcal{F}(\Gamma', h_1) = (\Gamma'', I)$, then for all $h_2 \in I$, $h_2 \cdot \langle m \rangle \in H''\}$. As we shall see, all the components of Γ^ν are independent of ν except for nature's initial move (as encoded by f_c^ν). In Γ^ν, the set of players is $\{(i, \Gamma') : \Gamma' \in \mathcal{G}_i\}$. The game tree of Γ^ν can be viewed as the union of the pruned game trees of $\Gamma' \in \mathcal{G}$. The histories of Γ^ν have the form $\langle \Gamma' \rangle \cdot h$, where $\Gamma' \in \mathcal{G}$ and $h \in \lfloor H^h \rfloor$. The move that a player or nature makes at a history $\langle \Gamma' \rangle \cdot h$ of Γ^ν is the same as the move made at h when viewed as a history of Γ'. The only move in Γ^ν not determined by Γ^* is nature's initial move (at the history $\langle \, \rangle$), where nature chooses the game $\Gamma' \in \mathcal{G}$ with probability $\nu(\Gamma')$.

Formally, let Γ^ν be a standard game such that

- $N^\nu = \{(i, \Gamma') : \Gamma' \in \mathcal{G}_i\}$;
- $M^\nu = \mathcal{G} \cup_{\Gamma' \in \mathcal{G}} \lfloor M' \rfloor$, where $\lfloor M' \rfloor$ is the set of moves that occur in $\lfloor H' \rfloor$;
- $H^\nu = \langle \, \rangle \cup \{\langle \Gamma' \rangle \cdot h : \Gamma' \in \mathcal{G}, h \in \lfloor H' \rfloor\}$;

[4] We did not discuss this latter equilibrium in the preliminary version of this paper.

- $P^\nu(\langle\,\rangle) = c$, and

$$P^\nu(\langle\Gamma^h\rangle \cdot h') = \begin{cases} (i, \Gamma^{h'}) & \text{if } P^h(h') = i \in N \text{ and} \\ & \mathcal{F}(\Gamma^h, h') = (\Gamma^{h'}, \cdot), \\ c & \text{if } P^h(h') = c; \end{cases}$$

- $f_c^\nu(\Gamma'|\langle\,\rangle) = \nu(\Gamma')$ and $f_c^\nu(\cdot|\langle\Gamma^h\rangle \cdot h') = f_c^h(\cdot|h')$ if $P^h(h') = c$;

- $\mathcal{I}^\nu_{i,\Gamma'}$ is just the $\sim_i$ relation restricted to histories $(\Gamma'', h) \in H^\nu$ where i moves and $\mathcal{F}(\Gamma'', h)$ has the form $(\Gamma', \cdot)$;

- $u^\nu_{i,\Gamma'}(\langle\Gamma^h\rangle \cdot z) = \begin{cases} u_i^h(z) & \text{if } \Gamma^h = \Gamma', \\ 0 & \text{if } \Gamma^h \neq \Gamma'. \end{cases}$

Theorem 3.1 *For all probability measures ν on $\mathcal{G}$*

(a) Γ^ν is a standard extensive game with perfect recall; and

(b) if ν gives positive probability to all games in $\mathcal{G}$, then $\vec{\sigma}$ is a Nash equilibrium of Γ^ν iff $\vec{\sigma}'$ is a generalized Nash equilibrium of Γ^*, where $\sigma_{i,\Gamma'}(\langle\Gamma^h\rangle \cdot h') = \sigma'_{i,\Gamma'}(\Gamma^h, h')$.

Although a Nash equilibrium does not necessarily exist in games with infinitely many players, Γ^ν has three special properties: (a) each player has only finitely many information sets, and (b) for each player (i, Γ'), there exists a finite subset $N(i, \Gamma')$ of N^ν such that (i, Γ)'s payoff in Γ^ν depends only on the strategies of the players in $N(i, \Gamma')$, and (c) Γ^ν is a game with perfect recall. This turns out to be enough to show that Γ^ν has at least one Nash equilibrium. Thus, we get the following corollary to Theorem 3.1.

Corollary 3.1 *Every game with awareness has a generalized Nash equilibrium.*

4 Modeling Awareness of Unawareness

In this section, we describe how to extend our representation of games with awareness to deal with awareness of unawareness. In an augmented game that represents player i's subjective view of the game, we want to model the fact that i may be aware of the fact that j can make moves at a history h that i is not aware of. We do this by allowing j to make a "virtual move" at history h. Histories that contain virtual moves are called *virtual histories*. These virtual histories do not necessarily correspond to a history in the underlying game Γ (i.e., i may falsely believe that j can make a move at h that he is unaware of), and even if a virtual history does correspond to a history in Γ, the subgame that follows that virtual history may bear no relationship to the actual subgame that follows the corresponding history in the underlying game Γ. Intuitively, the virtual histories describe agent i's (possibly incomplete and possibly incorrect) view of what would happen in the game if some move she is unaware of is made by agent j. Player j may have several virtual moves available at history h, and may make virtual moves at a number of histories in the augmented game.[5] Note that agent i's subjective game may include virtual moves for i himself; i may believe that he will become aware of more moves (and may take active steps to try and learn about these moves).

To handle awareness of unawareness, we consider a generalization of the notion of augmented game. We continue to refer to the generalized notion as an augmented game, using "augmented game without awareness of unawareness" to refer to the special case we have focused on up to now. Formally, $\Gamma^+ = (N^+, M^+, H^+, P^+, f_c^+, \{\mathcal{I}_i^+ : i \in N^+\}, \{u_i^+ : i \in N^+\}, \{A_i^+ : i \in N^+\})$ is an *augmented game*

[5]In the preliminary version of the paper, we assumed that all virtual moves were terminal moves. This is appropriate if i has no idea at all of what will happen in the game after a virtual move is made. The greater generality we allow here is useful to model situations where player i has some partial understanding of the game. For example, i may know that he can move left after j's virtual move, no matter what that virtual move is.

based on the (standard) finite extensive game $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$ if it satisfies conditions A1–A3, A6–A10 and A12 of augmented games, and variants of A4, A5, and A8. [6] Before stating these variants we need to define formally the set of virtual histories of Γ^+. The set of *virtual histories* V^+ *of* Γ^+ is defined by induction on the length of histories as follows:

1. if $m \in H^+$, $m \in M^+ - M$, and either $P^+(\emptyset) \in N^+$ or $P^+(\emptyset) = c = P(\emptyset)$, then $m \in V^+$;
2. if $h \cdot \langle m \rangle \in H^+$ and $h \in V^+$, then $h \cdot \langle m \rangle \in V^+$;
3. if $h \cdot \langle m \rangle \in H^+$, $m \in M^+ - M$, $h \notin V^+$, and either $P^+(h) \in N^+$ or $P^+(h) = c = P(\overline{h})$, then $m \in V^+$, where if $h \notin V^+$, then $\overline{h}$ is the subsequence of h consisting of all moves in h that are also in M, and if $h \in V^+$, then $\overline{h} = h$.

We can now state the variants of A4, A5, and A8.

A4$'$. If $P^+(h) \in N^+$ and $h \notin V^+$, then $P^+(h) = P(\overline{h})$ and $M_h^+ \subseteq M_{\overline{h}} \cup (M^+ - M)$.

A5$'$. If $P^+(h) = c$ and $h \notin V^+$, then either $P(\overline{h}) = c$ and $M_h^+ \subseteq M_{\overline{h}} \cup (M^+ - M)$, or $P(\overline{h}) \neq c$ and $M_h^+ \cap M = \emptyset$.

A8$'$. If h and h' are in the same information set in Γ^+ and $h, h' \notin V^+$, then $\overline{h}$ and $\overline{h}'$ are in the same information set in Γ.

A *game with awareness of unawareness based on* Γ is defined as a tuple $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ just as before. The modeler's extended game Γ^m satisfies the same conditions M1-M3 as before, and the mapping $\mathcal{F}$ satisfies C3–C5 and C7–C10 and the following variants of C1, C2, and C6:

C1$'$. $\{\overline{h} : h \in H^h, h \notin V^h\} = a$.

[6] We could also relax A3 to allow some "virtual players". We do not do that here for ease of exposition.

C2$'$. If $h' \in H^h$ and $P^h(h') = j$, then (a) $A_j^h(h') \subseteq a$, (b) if $h' \notin V^h$, then $(M_{\overline{h}'} \cap \{m : \overline{h}' \cdot \langle m \rangle \in a\}) \cup (M_{h'}^h - M_{\overline{h}'}) = M_{h'}^h$, and (c) if $\mathcal{F}(\Gamma^h, h') = (\Gamma', I')$, then for all $h'' \in I'$, we have $M'_{h''} \subseteq M_{h'}^h$.

C6$'$. If $h' \in I$ and $h, h' \notin V^h$, then $\overline{h}$ and $\overline{h}'$ are in the same information set in Γ.

C1$'$ and C6$'$ have been weakened so that these restrictions only apply to non-virtual histories of Γ^h. Part (a) of C2$'$ is the same as the first part of C2; part (b) implies that the set of moves available to player j at a non-virtual history h' is the set of moves that player i is aware of that are available to j at $\overline{h}'$ in the underlying game together with some virtual moves. It is not hard to check that in games without awareness of unawareness, part (c) follows from A4, C1, and C2, so it does not need to be explicitly stated in C2. However, now that A4 has been weakened to A4$'$, we must mention it explicitly.

Note that Γ^m is an augmented game with no awareness of unawareness; there are no virtual moves, since the modeler is indeed aware of all possible moves (and knows it). We can now define local strategies, generalized strategy profiles, and generalized Nash equilibrium just as we did for games with awareness. The same technique as that used to show Corollary 3.1 can be used to prove the following.

Theorem 4.1 *Every game with awareness of unawareness has a generalized Nash equilibrium.*

5 Modeling Lack of Common Knowledge

Game theorists have long searched for good approaches to modeling games where there is no common knowledge among players regarding the game being played. Our approach is flexible enough to handle such lack of common knowledge. In this section, we discuss the changes needed to handle lack of common knowledge. We remark that what we do here makes perfect sense even in games where there is full awareness.

We can modify our model to accommodate four different aspects of lack of common knowledge.

- *Lack of common knowledge regarding who moves.* We assumed that every player understands who moves in each history he is aware of. Although we still need to require that every player knows when it is his turn to move, we can handle the case where a player has false beliefs about who moves after a history that is not in one of his information sets. For example, we are interested in modeling the case where player i may be confused after some history h as to whether player j or player k moves, but in both cases i still believes that the same moves are available. That is, player i knows what could happen next, but he does not know who is going to do it. (Later we model uncertainty not only regarding who moves but also regarding what the move is.)

  To explain the necessary modifications, we need one more definition. Let $\mathcal{G}_{m,i}$ be the smallest subset of $\mathcal{G}$ such that if either $\Gamma^+ = \Gamma^m$ or $\Gamma^+ \in \mathcal{G}_{m,i}$, $h \in H^+$, $P^+(h) = i$, and $\mathcal{F}(\Gamma^+, h) = (\Gamma', \cdot)$, then $\Gamma' \in \mathcal{G}_{m,i}$. Intuitively, $\mathcal{G}_{m,i}$ consists of all games player i considers possible, or considers possible that he considers possible, and so on, at some history of the modeler's game.

  We can model lack of common knowledge about who moves by replacing A4 by

  A4$''$. If $P^+(h) = i \in N^+$, then $M_h^+ \subseteq M_{\overline{h}}$.

  Thus, we no longer require that the player who moves at history h is necessarily the one who moves at $\overline{h}$. However, we do make this requirement for the modeler's game, since the modeler is assumed to understand the underlying game. Thus, we must add a requirement M4 for the modeler's game that is identical to A4 except that we replace Γ^+ by Γ^m.

  Player i must also understand that he moves at a history h iff he moves at $\overline{h}$ for games in $\mathcal{G}_{m,i}$.

  C11. If $\Gamma^+ \in \mathcal{G}$, $h \in H^+$, $P^+(h) = i$, $A_i^+(h) = a$, $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$, and $h' \in H^h$, then if

$\Gamma^h \in \mathcal{G}_{m,i}$ and $P^h(h') = i$, then $P(\overline{h}') = i$. Conversely, if $P(\overline{h}') = i$, then there exists a prefix or suffix h'' of h' such that $\overline{h}'' = \overline{h}'$ and $P^h(h'') = i$.

We also need to make modifications to A5. Since we want to allow a player to have false beliefs about when nature moves, we replace A5 with

A5$''$. If $P^+(h) = c$, then either $M_h^+ \subseteq M_{\overline{h}}$, or $M_h^+ \cap M = \emptyset$.

As before, the moves in M_h^+ in the case where $M_h^+ \cap M = \emptyset$ intuitively capture uncertainty regarding a player's awareness level. But now it may be the case that a player i falsely believes that nature moves after history $\overline{h}$ in the underlying game. Just as with A4, we must add a condition M5 to the modeler's game that is identical to A5, except that Γ^+ is replaced by Γ^m.

- *Lack of common knowledge about the information sets.* We assumed that every player understand the signals every other player receives in every history he is aware of. We can weaken this assumption by allowing a player to have false beliefs about the signals received by other players, or equivalently, by allowing a player to have false beliefs about the information sets of other players.

  We can model lack of common knowledge about the information sets by removing conditions A8–A10. Again, because we assume that the modeler understands the information sets, we would add analogues of A8–A10 to the conditions on the modeler's game (replacing Γ^+ by Γ^m, of course). Similarly, we would require analogues of A8 and A9 to hold in the "C-list" of conditions for games $\Gamma^h \in \mathcal{G}_{m,i}$, and we weaken C6 so that it also holds only for $\Gamma^h \in \mathcal{G}_{m,i}$. We must also add an analogue of A10 to the "C-list" for games $\Gamma^h \in \mathcal{G}_{m,i}$ for histories h' and h'' in an i-information set.

- *Lack of common knowledge about payoffs.* We assumed that payoffs depended only on moves of the underlying game and that they were common knowledge among players. By dropping condition A12, we remove both of these assumptions. If we want to require that payoffs depend only on the underlying game, but still want to allow players to have false beliefs about the utilities, we would add an analogue of A12 in the modeler's game and use the following weakening of A12:

A12$'$. If $\Gamma^+ \in \mathcal{G}$, $z, z' \in Z^+$, and $\overline{z} = \overline{z}'$, then for all $i \in N^+$, $u_i^+(z) = u_i^+(z')$.

Although the player j whose view of the game is Γ^+ may have false beliefs about the payoffs, player j knows that the payoffs depend only on the moves made in the underlying game. A12$'$ captures that intuition.

- *Lack of common knowledge of the underlying game.* We assumed players have common knowledge about the structure of the underlying game. Our framework can model a situation where each player has a completely different conception of what game is actually being played, which may have very little relationship to the actual underlying game (although we still assume that the modeler's game corresponds to the actual game). The key idea is to drop the assumption that all augmented games are based on the same game Γ.

  To formalize this intuition, we modify A2 so that the A_i^+ function does not necessarily map histories of an augmented game to histories of the same game Γ. Rather, $A^+(h)$ is the set of histories of some game $\Gamma(h)$ that, intuitively, i considers to be the true underlying game. Thus, if h and h' are two histories in Γ^+, then $A_i^+(h)$ and $A_i^+(h')$ may be histories in two completely different games. Since $\Gamma(h)$ is viewed as i's subjective view of the true underlying game, we assume that he understands it perfectly. Thus, we retain A1, A6 and A8–A12 and replace conditions A3–A5 by M1–M3 (where the set of players is the set of players in $\Gamma(h)$ and the projection function maps a history h to a history $\overline{h}$ in $\Gamma(h)$). With regard to A7, note that, even if a player intuitively has perfect recall, he may realize in the future that he does not consider possible a history he considered possible in the past.

In the definition of games with awareness, we allow $\mathcal{G}$ to contain augmented games based on standard games different from the game on which Γ^m, the modeler's game, is based. We continue to require conditions C1, C3–C5, C7–C9, and C11, but we weaken C2. In C2 we required that a player i cannot consider possible a game Γ^h where one of the players j moving in Γ^h is aware of more runs than i is. In this setting, we allow i to consider possible a game Γ^h where one of the players j moving in Γ^h believes (falsely, from i's point of view) that some runs are possible that i does not consider possible. However, we require that the set of moves that i believes that j believes are available to him while moving at history h' in Γ^h is a subset of the moves i believes are available to j while moving at h'. We thus replace C2 by the following condition C2$''$, which is the analogue of parts (b) and (c) of C2$'$.

C2$''$. If $h' \in H^h$, $P^h(h') = j$, and $\mathcal{F}(\Gamma^h, h') = (\Gamma', I')$, then for all $h'' \in I'$, $M'_{h''} \subseteq M^h_{h'}$, and $M_{\overline{h}'} \cap \{m : \overline{h}' \cdot \langle m \rangle \in a\} = M^h_{h'}$.

Since we allow players to have false beliefs about information sets, we drop conditions C6 and C10. However, since we have dropped A7 and weakened C2, we now need the following condition, which requires that if a player considers possible a set of histories of the underlying game, then he cannot believe that in the future he will consider possible a different set of histories.[7]

C12. If $h' \in I$, $h'' \in H^h$, $P^h(h'') = i$, and h'' is a suffix of h', then $A^h_i(h') = A^h_i(h'')$.

[7] Note that this does not rule out a situation where a player i realizes at history h' that his view of the game will change at a future history h'' when he receives some additional information. If this is the case, then this should already be described in the set of histories that i considers possible at h'.

It is easy to see what C12 follows from A7, C2, and C4, which is why we did not list it explicitly earlier.

This approach of allowing the augmented games in $\mathcal{G}$ to be based on different underlying games actually subsumes our earlier approach and allows us to capture lack of common knowledge about who moves, what the information sets are, and what the payoffs are. For example, note that despite the fact that we have replaced A3-5 by M1-3, we can also model games with awareness using this approach by taking the game $\Gamma(h)$ to be the game consisting only of the runs of Γ that are in $A^+(h)$. (Of course, if we do that, we need to reinstate A7 and replace C2$''$ with C2.) To capture lack of common knowledge about who moves, we take $\Gamma(h)$ to be identical to Γ except that different agents may move at a given information set. Similarly, we can model lack of common knowledge about what the information sets and what the payoffs are by restricting $\Gamma(h)$ appropriately.

Despite all the changes to the conditions, the definitions of local strategies and generalized Nash equilibrium, and the theorems and their proofs remain unchanged. Thus, our techniques can deal with highly subjective games as well as awareness.

6 Related Work

There have been a number of models for unawareness in the literature (see, for example, [2; 14; 20; 21; 1]). Halpern [8] and Halpern and Rêgo [10] showed that in a precise sense all those models are special cases of Fagin and Halpern's [2] approach where they modeled awareness syntactically by introducing a new modal operator for it. Halpern and Rêgo [11] extended Fagin and Halpern's logic of awareness to deal with knowledge of unawareness. All of these papers focused on logic, and did not analyze the impact of unawareness in a strategic setting.

Feinberg's [3, 4] work is most similar work to ours. We discussed the high-level difference between our work and that of Feinberg in the introduction. Here we focus on some of the more detailed differences:

- Feinberg does not model games semantically. He encodes all the information in the $\mathcal{F}$ function syntactically, by describing each player's awareness level and iterated nested awareness levels (e.g., what player 1 is aware that player 2 is aware that player 3 is aware of).

- In dealing with extensive games, Feinberg [3] assumes that the runs that a player is aware of completely determine what game he believes he is playing. It cannot be the case that there are two distinct "identities" of a player that have the same awareness level. As we discussed in Section 2, this assumption limits the applicability of the model.

- Feinberg assumes that if player i is aware of player j, then i must be aware of some move of player j. We do not require such a condition since the analogous condition is not typically assumed in standard extensive games. For example, in a standard extensive game, a player may get a payoff even though there is no node where he can move. But it is trivial to add this requirement (as it would be trivial to drop in Feinberg's framework), and making it has no impact on the results.

- Feinberg [4] defines payoffs for player i by using what he calls "default actions" for players that i is unaware of. He says that this default action will be context dependent. We do not have such default actions in our setting; the payoff of a player in our framework is independent of the payoff of the players he is unaware of. The assumption of a default action seems somewhat problematic to us; it is not clear what the default move should be in general. Moreover, if two different players are unaware of player j, it is not clear why (or whether) they should assume the same default action.

- In dealing with extensive games, Feinberg [3] defines moves of nature by conditioning on the set of moves of nature the player is aware of. In our framework, this would amount to the following requirement:

  C13 If $\Gamma^+ \in \mathcal{G}$, $h \in H^+$, $P^+(h) = i$, $A_i^+(h) = a$, $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$, $h' \in H^h$, $P^h(h') = c$, and $M_{h'}^h \cap M_{\overline{h}'} \neq \emptyset$, then $f_c^h(m \mid h') = \frac{f_c(m|\overline{h}')}{f_c(M_{h'}^h|\overline{h}')}$ for every $m \in M_{h'}^h$ and $f_c^h(m \mid h') = 0$ if $m \notin M_{h'}^h$.

  As Feinberg did, for that condition to be well defined we require that $f_c(m \mid \overline{h}) \neq 0$ for all $m \in M_{\overline{h}}$ and histories h. As we discussed in Section 2, while we believe such a requirement makes sense if nature's move is interpreted objectively, it does not make sense in general so we do not assume this in every augmented game.

Sadzik [26] considers a logic of awareness, knowledge, and probability based on that of Heifetz, Meier, and Schipper [14], and uses it to give a definition of Bayesian equilibrium in normal-form games with awareness. Heifetz, Meier and Schipper [13] also consider a generalized state-space model with interactive unawareness and probabilistic beliefs and give a definition of Bayesian equilibrium in normal-form games, without assuming Feinberg's restriction. Li [17] has also provided a model of unawareness in extensive games, based on her earlier work on modeling unawareness [18; 19]. Although her representation of a game with unawareness is quite similar to ours, her notion of generalized Nash equilibrium is different from ours. Just as we do, she requires that every player i make a best response with respect to his beliefs regarding other player's strategies in the game Γ^i that i considers possible. However, unlike us, she requires that these beliefs satisfy a consistency requirement that implies, for example, that if a player i is aware of the same set of moves for him at both information set I_1 in game Γ_1 and information set I_2 in Γ_2, and these information sets correspond to the same information set in the underlying game Γ, then the local strategies σ_{i,Γ_1} and σ_{i,Γ_2} must agree at these information sets; that is, $\sigma_{i,\Gamma_1}(I_1) = \sigma_{i,\Gamma_2}(I_2)$.

Ozbay [23] proposes a model for games with uncertainty where players may have different awareness levels regarding a move of nature. He assumes that one of the players is fully aware, and can tell the other player about these moves before the second player moves. Although our model can easily capture this setting, what is interesting about Ozbay's approach is that the second player's beliefs about the probability of these revealed moves of are formed as part of the equilibrium definition. Filiz [5] uses Ozbay's model in the context of incomplete contracts in the presence of unforseen contingencies. In this setting, the insurer is assumed to be fully aware of the contingencies, and to decide strategically which contingencies to include in a contract, while the insuree may not be aware of all possible contingencies.

Finally, we remark that our notion of a game with awareness as consisting of the modeler's game together with description of which game each agent thinks is the actual game at each history has much in common with the intuition behind Gal and Pfeffer's [6] notion of a *Network of Influence Diagrams (NID)*. Formally, NIDs are a graphical language for representing uncertainty over decision-making models. A node in a NID (called a *block* by Gal and Pfeffer) represents an agent's subjective belief about the underlying game and what the strategies used by agents depend on. Each node (game) in a NID is associated with a *multiagent influence diagram* [15] (MAID), which is a compact representation of a game. A NID has directed edges between nodes labeled by pairs of the form (i, H), where i is an agent and (in our language) H is a set of histories. Intuitively, if there an edges from a node (game) Γ to a node Γ' in a NID labeled by a pair (i, H), then H is a set a set of histories in Γ, there is an agent j that moves at all the histories in H, and in game Γ, i believes that j believes that Γ' is the true game when moving at a history $h \in H$.

Although Gal and Pfeffer do not try to handle notions of awareness with NIDs, it seems possible to extend them to handle awareness. To do this appropriately, consistency requirements similar to C1–C10 will need to be imposed.

7 Conclusion

We have generalized the representation of games to take into account agents who may not be aware of all the moves or all the other agents, but may be aware of their lack of awareness. Moreover, our representation is also flexible enough to deal with subjective games when there is lack of common knowledge about the game, even if awareness is not an issue. We have also shown how to define strategies and Nash equilibrium in such settings. These generalizations greatly increase the applicability of game-theoretic notions in multiagent systems. In large games involving many agents, agents will almost certainly not be aware of all agents and may well not be aware of all the moves that agents can make. Moreover, as we suggested in the introduction, even in well-understood games like chess, by giving awareness a more computational interpretation, we can provide a more realistic model of the game from the agents' perspective. We remark that although we focus on generalizing extensive-form games, our framework is able to deal with normal-form games as well, since we can view normal-form games as a special case of extensive-form games.

There is clearly much more to be done to understand the role of awareness (and lack of awareness) in multiagent systems. We list some of the many issues here:

- We have assumed perfect recall here. But in long games, it seems more reasonable to assume that agents do not have perfect recall. In a long chess game, typical players certainly do not remember all the moves that have been played and the order in which they were played. It is well known that even in single-agent games, considering agents with imperfect recall leads to a number of subtleties (c.f. [7; 24]). We suspect that yet more subtleties will arise when combining imperfect recall with lack of awareness.

- In a Nash equilibrium of an extensive-form game, it may be the case that the move made at an information set is not necessarily a best response if that information set is not reached. For example, in the game described in Figure 1, even if both players have common knowledge of the game, the profile where A moves down and B moves across is a Nash equilibrium. Nevertheless moving down is not a best response for B if B is actually called upon to play. The only reason that this is a Nash equilibrium is that B does not in fact play. *Sequential equilibrium* [16] is a solution concept that is arguably more appropriate for an extensive-form game; it refines Nash equilibrium (in the sense that every sequential equilibrium is a Nash equilibrium) and does not allow solutions such as (down$_A$, across$_B$). Our representation of games with awareness (of unawareness) allows for relatively straightforward generalizations of such refinements of Nash equilibrium. However, there are subtleties involved in showing that generalized versions of these refinements always exist. For example, we no longer have a one-to-one correspondence between the generalized sequential equilibria of the game Γ^* and the sequential equilibria of the corresponding standard game Γ^ν. Nevertheless, we believe that we should be able to use a more refined construction to show that a generalized sequential equilibrium exists in every game with awareness.

- We have analyzed situations where agents may be unaware of some moves in the underlying game, may be aware of their unawareness, and may have completely false beliefs about the underlying game. Of course, there are other cases of interest where additional properties may hold. For example, consider a large geographically-dispersed game where agents interact only with nearby neighbors. In such a game, an agent may be unaware of exactly who is playing the game (although she may realize that there are other agents besides her neighbors, and even realize that the moves made by distant agents may have an indirect effect on her). To model such a situation, we may want to have virtual moves after which the game does not end, and to allow agents to be aware of subsequences of histories in the underlying game. We suspect that a straightforward extension of the ideas in this paper can deal with such situations, but we have not worked out the details.

- There has been a great deal of work on computing Nash equilibria. As we have shown, a generalized Nash equilibrium of a game with awareness is a Nash equilibrium of a standard game. However, this standard game can be rather large. Are there efficient computational techniques for computing generalized Nash equilibrium in interesting special cases?

- If there is little shared knowledge regarding the underlying game, the set $\mathcal{G}$ of augmented games can be quite large, or even infinite. Is it important to consider all the iterated levels of unawareness encoded in $\mathcal{G}$? Halpern and Moses [9] showed that, in analyzing coordinated attack, no finite level of knowledge suffices; common knowledge is needed for coordination. Stopping at any finite level has major implications. Rubinstein [25] considered a variant of the coordinated attack problem with probabilities, and again showed that no finite level suffices (and significant qualitative differences arise if only a finite part of hierarchy of knowledge is considered). On the other hand, Weinstein and Yildiz [28] provide a condition under which the effect of players' kth order beliefs is exponentially decreasing in k. While we strongly suspect that there are games in which higher-order unawareness will be quite relevant, just as with the Weinstein-Yildiz result, there may be conditions under which higher-order awareness becomes less important, and a simpler representation may suffice. Moreover, it may be possible to use NIDs to provide a more compact representation of games of awareness in many cases of interest (just as Bayesian networks provide a compact representation of probability distributions in many cases of interest), leading to more efficient techniques for computing generalized Nash equilibrium.

We hope to explore some of these issues in forthcoming work.

References

[1] E. Dekel, B. Lipman, and A. Rustichini. Standard state-space models preclude unawareness. *Econometrica*, 66:159–173, 1998.

[2] R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.

[3] Y. Feinberg. Subjective reasoning—games with unawareness. Technical Report Resarch Paper Series #1875, Stanford Graduate School of Business, 2004.

[4] Y. Feinberg. Games with incomplete awareness. Technical Report Resarch Paper Series #1894, Stanford Graduate School of Business, 2005.

[5] E. Filiz. Incorporating unawareness into ontract theory. Unpublished manuscript, 2006.

[6] Y. Gal and A. Pfeffer. Reasoning about agents' beliefs and decision-making processes in games. Unpublished manuscript, 2005.

[7] J. Y. Halpern. On ambiguities in the interpretation of game trees. *Games and Economic Behavior*, 20:66–96, 1997.

[8] J. Y. Halpern. Alternative semantics for unawareness. *Games and Economic Behavior*, 37:321–339, 2001.

[9] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990. A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.

[10] J. Y. Halpern and L. C. Rêgo. Interactive unawareness revisited. In *Theoretical Aspects of Rationality and Knowledge: Proc. Tenth Conference (TARK 2005)*, pages 78–91, 2005.

[11] J. Y. Halpern and L. C. Rêgo. Reasoning about knowledge of unawareness, 2006. Available at arxiv.org/cs.LO/0603020.

[12] J. Harsanyi. Games with incomplete information played by 'Bayesian' players, parts I–III. *Management Science*, 14:159–182, 320–334, 486–502, 1968.

[13] A. Heifetz, M. Meier, and B. Schipper. Unawareness, belief and games. Unpublished manuscript., 2005.

[14] A. Heifetz, M. Meier, and B. Schipper. Interactive unawareness. *Journal of Economic Theory*, 130:78–94, 2006.

[15] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. In *Proc. Seventeenth International Joint Conference on Artificial Intelligence (IJCAI '01)*, pages 1027–1036, 2001.

[16] D. M. Kreps and R. B. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.

[17] J. Li. Dynamic games with unawareness. Unpublished manuscript, 2006.

[18] J. Li. Information structures with unawareness. Unpublished manuscript, 2006.

[19] J. Li. Modeling unawareness without impossible states. Unpublished manuscript, 2006.

[20] S. Modica and A. Rustichini. Awareness and partitional information structures. *Theory and Deci-*

sion, 37:107–124, 1994.

[21] S. Modica and A. Rustichini. Unawareness and partitional information structures. *Games and Economic Behavior*, 27(2):265–298, 1999.

[22] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Mass., 1994.

[23] E. Ozbay. Unawareness and strategic announcements in games with uncertainty. Unpublished manuscript, 2006.

[24] M. Piccione and A. Rubinstein. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior*, 20(1):3–24, 1997.

[25] A. Rubinstein. The electronic mail game: strategic behavior under "almost common knowledge". *American Economic Review*, 79:385–391, 1989.

[26] T. Sadzik. Knowledge, awareness and probabilistic beliefs. Unpublished manuscript., 2005.

[27] H. Salonen. Nash equilibria in normal form games. Unpublished manuscript; available at http://aws.tt.utu.fi/salo022.html, 2005.

[28] J. Weinstein and M. Yildiz. Impact of higher-order uncertainty. Unpublished manuscript., 2003.

A Proofs

Theorem 3.1: *For all probability measures ν on $\mathcal{G}$*

(a) Γ^ν *is a standard extensive game with perfect recall;*

(b) if ν gives positive probability to all games in $\mathcal{G}$, then $\vec{\sigma}$ is a Nash equilibrium of Γ^ν iff $\vec{\sigma}'$ is a generalized Nash equilibrium of Γ^*, where $\sigma_{i,\Gamma'}(\langle\Gamma^h\rangle \cdot h') = \sigma'_{i,\Gamma'}(\Gamma^h, h')$.

Proof: For part (a), suppose that $\langle\Gamma'\rangle \cdot h_1'$) and $\langle\Gamma''\rangle \cdot h_1''$) are in the same (i,Γ^+)-information set of Γ^ν and that h_2' is a prefix of h_1' such that $P^\nu(\langle\Gamma'\rangle h_2')) = (i,\Gamma^+)$. By definition of Γ^ν, it must be the case that there exist i-information sets I_1 and I_2 in Γ^+ such that $\mathcal{F}(\Gamma',h_1') = \mathcal{F}(\Gamma'',h_1'') = (\Gamma^+,I_1)$ and $\mathcal{F}(\Gamma',h_2') = (\Gamma^+,I_2)$. If h_1 is a history in I_1, C8 implies that there exists a prefix h_2 of h_1 such that $P^+(h_2) = i$, $\mathcal{F}(\Gamma^+,h_2) = (\Gamma^+,I_2)$ and if $h_2' \cdot \langle m\rangle$ is a prefix of h_1', then $h_2 \cdot \langle m\rangle$ is a prefix of h_1. Applying C8 again, it follows that there exists a prefix h_2'' of h_1'' such that $P^{\Gamma''}(h_2'') = i$ and $\mathcal{F}(\Gamma'',h_2'') = (\Gamma^+,I_2)$ and if $h_2 \cdot \langle m\rangle$ is a prefix of h_1, then $h_2'' \cdot \langle m\rangle$ is a prefix of h_1''. Therefore, by definition of Γ^ν, (Γ'',h_2'') and (Γ',h_2') are in the same information set.

Suppose further that $h_2' \cdot \langle m\rangle$ is a prefix of h_1'. Thus, $h_2 \cdot \langle m\rangle$ is a prefix of h_1, which implies that $h_2'' \cdot \langle m\rangle$ is a prefix of h_1''. This proves part (a).

For part (b), let $Pr^\nu_{\vec{\sigma}}$ be the probability distribution over the runs in Γ^ν induced by the strategy profile $\vec{\sigma}$ and f_c^ν. $Pr^\nu_{\vec{\sigma}}(z)$ is the product of the probability of each of the moves in z. (It is easy to define this formally by induction on the length of z; we omit details here.) Similarly, let $Pr^h_{\vec{\sigma}'}$ be the probability distribution over the runs in $\Gamma^h \in \mathcal{G}$ induced by the generalized strategy profile $\vec{\sigma}'$ and f_c^h. Note that if $Pr^h_{\vec{\sigma}'}(z) > 0$, then $z \in \lfloor H^h \rfloor$. Thus, $\langle\Gamma^h\rangle \cdot z \in H^\nu$.

For all strategy profiles σ and generalized strategy profiles σ', if $\sigma'_{i,\Gamma'}(\Gamma^h,h') = \sigma_{i,\Gamma'}(\langle\Gamma^h\rangle \cdot h')$, then it is easy to see that for all $z \in Z^h$ such that $Pr^h_{\vec{\sigma}'}(z) > 0$, we have that $Pr^\nu_{\vec{\sigma}}(\langle\Gamma^h\rangle \cdot z) = \nu(\Gamma^h)Pr^h_{\vec{\sigma}'}(z)$. And since ν is a probability measure such that $\nu(\Gamma^h) > 0$ for all $\Gamma^h \in \mathcal{G}$, we have that $Pr^\nu_{\vec{\sigma}}(\langle\Gamma^h\rangle \cdot z) > 0$ iff $Pr^h_{\vec{\sigma}'}(z) > 0$. Suppose that $\vec{\sigma}$ is a Nash equilibrium of Γ^ν. Suppose, by way of contradiction, that $\vec{\sigma}'$ such that $\sigma'_{i,\Gamma'}(\Gamma^h,h') = \sigma_{i,\Gamma'}(\langle\Gamma^h\rangle \cdot h')$ is not a generalized Nash equilibrium of Γ^*. Thus, there exists a player i, a game $\Gamma^+ \in \mathcal{G}_i$, and a local strategy s' for player i in Γ^+ such that

$$\sum_{z \in Z^+} Pr^+_{\vec{\sigma}'}(z) u_i^+(z) < \sum_{z \in Z^+} Pr^+_{(\vec{\sigma}'_{-(i,\Gamma^+)}, s')}(z) u_i^+(z). \tag{1}$$

Define s to be a strategy for player (i, Γ^+) in Γ^ν such that $s(\langle \Gamma^h \rangle \cdot h') = s'(\Gamma^h, h')$. Multiplying (1) by $\nu(\Gamma^+)$ and using the observation in the previous paragraph, it follows that

$$\sum_{z \in \lfloor Z^+ \rfloor} Pr^\nu_{\vec{\sigma}}(\langle \Gamma^+ \rangle \cdot z) u_i^+(z) < \sum_{z \in \lfloor Z^+ \rfloor} Pr^\nu_{(\vec{\sigma}_{-(i,\Gamma^+)}, s)}(\langle \Gamma^+ \rangle \cdot z) u_i^+(z). \tag{2}$$

By definition of $u^\nu_{i,\Gamma'}$, (2) holds iff

$$\sum_{z^\nu \in Z^\nu} Pr^\nu_{\vec{\sigma}}(z^\nu) u^\nu_{i,\Gamma^+}(z^\nu) < \sum_{z^\nu \in Z^\nu} Pr^\nu_{(\vec{\sigma}_{-(i,\Gamma^+)}, s)}(z^\nu) u^\nu_{i,\Gamma^+}(z^\nu). \tag{3}$$

Therefore, $\vec{\sigma}$ is not a Nash equilibrium of Γ^ν, a contradiction. The proof of the converse is similar; we leave details to the reader. ∎

Corollary 3.1: *Every game with awareness has a generalized Nash equilibrium.*

Proof: For games with perfect recall, there is a natural isomorphism between mixed strategies and behavioral strategies, so a Nash equilibrium in behavior strategies can be viewed as a Nash equilibrium in mixed strategies [22]. Moreover, mixed-strategy Nash equilibria of an extensive-form game are the same as the mixed-strategy Nash equilibria of its normal-form representation. Salonen [27] showed that there exists a Nash equilibrium in mixed strategies in a normal form games with an arbitrary set N of players if, for each player i, the set S_i of pure strategies of player i is a compact metric space, and the utility functions $u_i : S \rightarrow I\!R$ are continuous for all $i \in N$, where $I\!R$ is the set of real numbers and $S = \Pi_{i \in N} S_i$, the set of pure strategies, is endowed with the product topology. Since in Γ^ν, every player has a finite number of pure strategies, S_i is clearly a compact metric space. Moreover, since each player's utility depends only on the strategies of a finite number of other players, it is easy to see that $u_i : S \to I\!R$ is continuous for each player $i \in N$. It follows that there exists a Nash equilibrium of Γ^ν. Thus, the corollary is immediate from Theorem 3.1. ▮