

Extensive Games with Possibly Unaware Players*

Joseph Y. Halpern
Computer Science
Cornell University, USA
halpern@cs.cornell.edu

Leandro C. Rêgo
School of Electrical and Computer Engineering
Cornell University, USA
lcr26@cornell.edu

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.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems

General Terms

Economics, Theory

Keywords

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 [7] 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

*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).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

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. [1]), 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 does not understand what can happen after this move, i does not develop the game tree after this move; i simply describes what she believes will be the payoffs after the move. Thus, our representation is quite similar in spirit to 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!

Recently, Feinberg [2, 3] also studied games with awareness. Feinberg [3] gives a definition of extended Nash equilibrium in normal-form games. Although his definition stems from much the same intuitions as ours, 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 [2] 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.

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 [8]. 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 available to players (and nature) during the game.¹
- H is a finite set of finite sequences of moves (elements of M) that is closed under prefixes, so that if $h \in H$

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

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 .

- $P : (H - Z) \rightarrow N \cup \{c\}$ 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 $f_c(\cdot | h)$ on M_h . Intuitively, $f_c(\cdot | h)$ 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 for 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.
- $u_i : Z \rightarrow \mathbb{R}$ is a payoff function for player i , assigning a real number (i 's payoff) to each run of the game.

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 finite extensive game with 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 $P^+(h) = P(\bar{h})$, where $\bar{h}$ is the subsequence of h consisting of all the moves in h that are also in M , and $M_h^+ \subseteq M_{\bar{h}}$. 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 $P(\bar{h}) = c$ and $M_h^+ \subseteq M_{\bar{h}}$, 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 $\bar{h}$ and $\bar{h}'$ are in the same information set in Γ .
- A9. If h and h' are histories in both Γ^+ and Γ (i.e., $h = \bar{h}$ and $h' = \bar{h}'$), and $\bar{h}$ and $\bar{h}'$ 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. $\{\bar{z} : z \in Z^+\} \subseteq Z$; 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 $\bar{z} \in Z$, then $u_i^+(z) = u_i(\bar{z})$. Thus, a player's utility just depends on the moves made in the underlying game. (By A11, we have $\bar{z} \in Z$. We have included the clause "if $\bar{z} \in Z$ " 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 $P^+(h) = P(\bar{h})$ unless $P^+(h) = c$ and $M_h^+ \cap M = \emptyset$), (2) what the payoffs are in the underlying game (since $u^+(\bar{z}) = u(z)$), 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 the full paper for details.

To fully understand A8–A10, we must first discuss our view of information sets. As pointed out by Halpern [5], 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. We interpret an i -information set to be the set of all histories where player i receives the same sequence of signals. A8–A10 reflect this interpretation under our assumption that (modulo awareness) information sets are common knowledge.

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 $\langle \text{across}_A, \text{down}_B \rangle$ 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 .

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.




Figure 2: The augmented game Γ^A .




Figure 3: The augmented game Γ^B .

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

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 $\{\bar{z} : z \in Z^m\} = Z$.

M3. If $P^m(h) \in N$, then $M_h^m = M_{\bar{h}}$. If $P^m(h) = c$, then either $M_h^m \cap M = \emptyset$ or $M_h^m = M_{\bar{h}}$ and $f_c^m(\cdot | h) = f_c(\cdot | \bar{h})$.

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_c^m) 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. $\{\bar{h} : h \in H^h\} = a$.
- C2. If $h' \in H^h$ and $P^h(h') = j$, then $A_j^h(h') \subseteq a$ and $M_{\bar{h}'} \cap \{m : \bar{h}' \cdot \langle m \rangle \in a\} = M_{\bar{h}'}^h$.
- C3. If h' and h are in the same information set in Γ^+ and $\bar{h}' \in a$, then there exists $h'' \in I$ such that $\bar{h}'' = \bar{h}'$.
- 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 $\bar{h}$ and $\bar{h}'$ 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 $h_1 \in I'$, $h_1 \cdot \langle m \rangle \in H^I$.
- 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 for player j at h' is just the set of moves that player i is aware of that are available for j at $\bar{h}'$ 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 , $\mathcal{F}(\Gamma^+, h)$ is not uniquely determined by Γ^+ and h . Moreover, 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.

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 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.

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 [4] notion *Networks of Influence Diagrams*; we discuss the connections in detail in the full paper.

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 $\mathcal{G}_i = \{\Gamma' \in \mathcal{G} : \text{for some } \Gamma^+ \in \mathcal{G} \text{ and } h \text{ in } \Gamma^+, P^+(h) = i \text{ and } \mathcal{F}(\Gamma^+, h) = (\Gamma', \cdot)\}$. 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 Γ^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.² These considerations suggest that our notion of Nash equilibrium is appropriate.

It is easy to see that $\bar{\sigma}$ is a Nash equilibrium of a standard game iff $\bar{\sigma}$ 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 $\text{dom}(\sigma_{i,\Gamma'})$ to denote the domain of the strategy $\sigma_{i,\Gamma'}$, we have

$$\begin{aligned} \text{dom}(\sigma_{A,\Gamma^A}) &= \{(\Gamma^m, \langle \rangle), (\Gamma^A, \langle \text{unaware} \rangle), (\Gamma^A, \langle \text{aware} \rangle)\}, \\ \text{dom}(\sigma_{B,\Gamma^m}) &= \{(\Gamma^m, \langle \text{across}_A \rangle), (\Gamma^A, \langle \text{aware, across}_A \rangle)\}, \\ \text{dom}(\sigma_{A,\Gamma^B}) &= \{(\Gamma^B, \langle \rangle)\}, \text{ and} \\ \text{dom}(\sigma_{B,\Gamma^B}) &= \{(\Gamma^A, \langle \text{unaware, across}_A \rangle), (\Gamma^B, \langle \text{across}_A \rangle)\}. \end{aligned}$$

Each of these domains consists of a single generalized information set. If $p < 1/2$, then the unique generalized Nash equilibrium is $\sigma_{A,\Gamma^A} = \text{across}_A$, $\sigma_{A,\Gamma^B} = \text{down}_A$, $\sigma_{B,\Gamma^m} = \text{down}_B$, $\sigma_{B,\Gamma^B} = \text{across}_B$. 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 the unique generalized Nash equilibrium is $\sigma_{A,\Gamma^A} = \text{down}_A$, $\sigma_{A,\Gamma^B} = \text{down}_A$, $\sigma_{B,\Gamma^m} = \text{down}_B$, $\sigma_{B,\Gamma^B} = \text{across}_B$. 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 .

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. 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 $\langle \Gamma^h \rangle \cdot h$, where $\Gamma^h \in \mathcal{G}$ and h is a “pruned” history in H^h . The move that a player or nature makes at a history $\langle \Gamma^h \rangle \cdot h$ of Γ^ν is the same as the move made at h when viewed as a history of Γ^h . 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 $\nu(\Gamma')$. We leave the details of the construction of Γ^ν to the full paper.

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 $\bar{\sigma}$ is a Nash equilibrium of Γ^ν iff $\bar{\sigma}'$ is a generalized Nash equilibrium of Γ^* , where $\sigma_{i,\Gamma'}(\langle \Gamma^h \rangle \cdot h') = \sigma'_{i,\Gamma'}(\Gamma^h, h')$.

²In 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.

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 . A virtual move is always a terminal move; associated with it there is a payoff vector. Intuitively, these payoffs are agent i ’s subjective feeling about what would happen 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. These virtual moves do not necessarily correspond to a move 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 j ’s virtual move does correspond to a move in Γ , the payoffs after the virtual move may bear no relationship to the actual payoffs after that move in Γ . 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 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 the following variants of A4 and A5:

- A4'. If $P^+(h) \in N^+$, then $P^+(h) = P(\bar{h})$, and $M_h^+ \subseteq M_{\bar{h}} \cup (M^+ - M)$. Moreover, every history $h \cdot \langle m \rangle \in H^+$ with $m \in M_h^+ - M_{\bar{h}}$ is terminal. Intuitively, the moves in $M_h^+ - M_{\bar{h}}$ are virtual moves.
- A5'. If $P^+(h) = c$ and $M_h^+ \cap M \neq \emptyset$, then $M_h^+ \subseteq M_{\bar{h}} \cup (M^+ - M)$. Moreover, if $m \in M_h^+ - M$ and $M_h^+ \cap M \neq \emptyset$, then $h \cdot \langle m \rangle$ is terminal. Intuitively, if $M_h^+ \cap M \neq \emptyset$, then all the moves in $M_h^+ - M$ are virtual moves.

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 C1 and C3–C10 and the following variant of C2:

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

Part (a) of C2' is the same as the first part of C2. Part (b) implies that the set of moves available for player j at h' is the set of moves that player i is aware of that are available for j at $\bar{h}'$ 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. CONCLUSION

We have generalized the representation of games to take into account agents who may not be aware of all the moves or agents, but may be aware of their lack of awareness. As we show in the full paper, 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, by giving awareness a more computational interpretation, we can provide a more realistic model of a game like chess.

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:

- 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. There are refinements of Nash equilibrium that are arguably more appropriate for an extensive-form game [8]. 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.
- We have analyzed situations where agents may be unaware of some moves in the underlying game and may be aware of their unawareness. 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 extensions 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 [6] showed that, in analyzing co-ordinated attack, no finite level of knowledge suffices; common knowledge is needed for coordination. Stopping at any finite level has major implications. On the other hand, Weinstein and Yildiz [9] provide a condition under which the effect of players' k th order beliefs is exponentially decreasing in k . While we strongly suspect that there are games in which higher-order unawareness will be quite relevant, there may be conditions under which higher-order awareness becomes less important, and a simpler representation may suffice.

6. REFERENCES

- [1] R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.
- [2] Y. Feinberg. Subjective reasoning—games with unawareness. Technical Report Research Paper Series #1875, Stanford Graduate School of Business, 2004.
- [3] Y. Feinberg. Games with incomplete awareness. Technical Report Research Paper Series #1894, Stanford Graduate School of Business, 2005.
- [4] Y. Gal and A. Pfeffer. Reasoning about agents' beliefs and decision-making processes in games. Unpublished manuscript, 2005.
- [5] J. Y. Halpern. On ambiguities in the interpretation of game trees. *Games and Economic Behavior*, 20:66–96, 1997.
- [6] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
- [7] J. Harsanyi. Games with incomplete information played by 'Bayesian' players, parts I–III. *Management Science*, 14:159–182, 320–334, 486–502, 1968.
- [8] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Mass., 1994.
- [9] J. Weinstein and M. Yildiz. Impact of higher-order uncertainty. Unpublished manuscript., 2003.