

Generalized Solution Concepts in Games with Possibly Unaware Players

Leandro C. Rêgo
Statistics Department
Federal University of Pernambuco
Recife-PE, Brazil
e-mail: leandro@de.ufpe.br

Joseph Y. Halpern
Computer Science Department
Cornell University
Ithaca-NY, U.S.A.
e-mail: halpern@cs.cornell.edu

Abstract

Most work in game theory assumes that players are perfect reasoners and have common knowledge of all significant aspects of the game. In earlier work [Halpern and Rêgo 2006], we proposed a framework for representing and analyzing games with possibly unaware players, and suggested a generalization of Nash equilibrium appropriate for games with unaware players that we called *generalized Nash equilibrium*. Here, we use this framework to analyze other solution concepts that have been considered in the game-theory literature, with a focus on sequential equilibrium. We also provide some insight into the notion of generalized Nash equilibrium by proving that it is closely related to the notion of rationalizability when we restrict the analysis to games in normal form and no unawareness is involved.

1 INTRODUCTION

Game theory has proved to be a useful tool in the modeling and analysis of many phenomena involving interaction between multiple agents. However, standard models used in game theory implicitly assume that agents are perfect reasoners and have common knowledge of all significant aspects of the game. There are many situations where these assumptions are not reasonable. In large games, agents may not be aware of the other players in the game or all the moves a player can make. Recently, we [Halpern and Rêgo 2006] proposed a way of modeling such games. A key feature of this approach is the use of an *augmented game*, which represents what players are aware of at each node of an extensive form representation of a game. Since the game is no longer assumed to be common knowledge, each augmented game represents the game a player considers possible in some situation and describes how he believes each other player's *awareness level* changes over time, where intuitively the awareness level of a player is the set of histories of the game that the player is aware of.

In games with possibly unaware players, standard solution concepts cannot be applied. For example, in a standard game a strategy profile 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 his strategy even if i knew what strategies the other agents were using. In the presence of unawareness this no longer make sense, since the strategies used by other players may involve moves i is unaware of. We proposed a generalization of Nash equilibrium consisting 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' . Roughly speaking, a generalized strategy profile $\vec{\sigma}$, which includes a strategy $\sigma_{i,\Gamma'}$ for each pair (i, Γ') , is a *generalized Nash equilibrium* if $\sigma_{i,\Gamma'}$ 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 showed that every game with awareness has a generalized Nash equilibrium by associating a game with awareness with a standard game (where agents are aware of all moves) and proving that there is a one-to-one correspondence between the generalized Nash equilibria of the game with awareness and the Nash equilibria of the standard game.

Some Nash equilibria seem unreasonable. For example, consider the game shown in Figure 1. One Nash equilibrium of this game has A playing down_A and B playing across_B . Since




Figure 1: A simple game.

player A chooses down_A , player B never gets the chance to move so whatever he does he will get the same payoff of 1. However, if player B is rational and ever gets to move, he will not choose move across_B since it gives a lower payoff for him than choosing move down_B .

There have been a number of variants of Nash equilibrium proposed in the literature, such as *perfect equilibrium* [Selten 1975], *proper equilibrium* [Myerson 1978], *sequential equilibrium* [Kreps and Wilson 1982], and *rationalizability* [Bernheim 1984; Pearce 1984], to name just a few. Each of these solution concepts involves some notion of best response. Our framework allows for straightforward generalizations of all these solution concepts. As in our treatment of Nash equilibrium, if $\Gamma_1 \neq \Gamma_2$, we treat 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. Each version of player i best responds (in the sense appropriate for that solution concept) given his view of the game. In standard games, it has been shown that, in each game, there is a strategy profile satisfying that solution concept. Showing that an analogous result holds in games with

awareness can be nontrivial. Instead of going through the process of generalizing every solution concept, we focus here on *sequential equilibrium* since (a) it is one of the best-known solution concepts for extensive games, (b) the proof that a generalized sequential equilibrium exists suggests an interesting generalization of sequential equilibrium for standard games, and (c) the techniques used to prove its existence in games with awareness may generalize to other solution concepts.

Sequential equilibrium 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). Intuitively, in a sequential equilibrium, every player must make a best response at every information set (even if it is reached with probability 0). In the game shown in Figure 1, the unique sequential equilibrium has *A* choosing across_A and *B* choosing down_B. We propose a generalization of sequential equilibrium to games with possibly unaware players, and show that every game with awareness has a generalized sequential equilibrium. This turns out to be somewhat more subtle than the corresponding argument for generalized Nash equilibrium. Our proof requires us to define a generalization of sequential equilibrium in standard games. Roughly speaking, this generalization relaxes the implicit assumption in sequential equilibrium that every history in an information set is actually considered possible by the player. We call this notion *conditional sequential equilibrium*.

We also provide some insight into the notion of generalized Nash equilibrium by proving that, in a precise sense, it is closely related to the notion of *rationalizability* when we restrict the analysis to games in *normal form* and no unawareness is involved (although the underlying game is no longer common knowledge among the players). Roughly speaking, a normal form game can be thought as a one-shot extensive game where no player knows the move the others made before they make their own move. Intuitively, in standard games, a strategy is rationalizable for a player if it is a best response to some reasonable beliefs he might have about the strategies being played by other players, and a strategy is part of a Nash equilibrium if it is a best response to the strategies actually played by the other players. Since, in games with awareness, the game is not common knowledge, a local strategy for player *i* in Γ^+ is part of a generalized Nash equilibrium if it is a best response to the strategies played by the opponents of player *i* in the games player *i* believes his opponents consider to be the actual one while moving in Γ^+ . Note that the line between rationalizability and generalized Nash equilibrium is not sharp. In fact, we are essentially able to prove that a strategy is rationalizable in a standard game Γ iff it is part of generalized Nash equilibrium of an appropriate game with awareness whose underlying game is Γ .

The rest of this paper is organized as follows. In Section 2, we give the reader the necessary background to understand this paper by reviewing our model for games with awareness. In Section 3, we review the definition of sequential equilibrium for standard games and define its generalization for games with awareness. In Section 4, we define the concept of conditional sequential equilibrium for standard games, and prove that there is a one-to-one correspondence between the generalized sequential equilibria of a game with awareness and the conditional sequential equilibria of the standard game associated with it. In Section 5, we analyze the connection between rationalizability and generalized Nash equilibrium. We conclude with

some discussion of solution concepts in games with unaware players in Section 6.

2 GAMES WITH AWARENESS

In this section, we introduce some notation and give some intuition regarding games with awareness. We encourage the reader to consult our earlier paper for details.

Games with awareness are modeled using *augmented games*. Given a standard extensive-form game described by a game tree Γ , an augmented game Γ^+ based on Γ 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.

Formally, a (*finite*) *extensive game* is a tuple $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$, 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.
- 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 .
- $P : (H - Z) \rightarrow N \cup \{c\}$ 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 $f_c(\cdot \mid h)$ 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 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.

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. There are also extra moves of nature that intuitively capture players' uncertainty regarding the awareness level of their opponents. Formally, given an extensive game $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$, an *augmented game based on* Γ is a tuple $\Gamma^+ = (N^+, M^+, H^+, P^+, f_c^+, \{\mathcal{I}_i^+ : i \in N^+\}, \{u_i^+ : i \in N^+\}, \{A_i^+ : i \in N^+\})$, where $(N^+, M^+, H^+, P^+, f_c^+, \{\mathcal{I}_i^+ : i \in N^+\}, \{u_i^+ : i \in N^+\})$ is a standard extensive game with perfect recall¹ and $A_i^+ : H_i^+ \rightarrow 2^H$ describes i 's awareness level at each history at which he moves. Γ^+ must satisfy some consistency conditions. These conditions basically ensure that

- a player's awareness level depends only on the information she has as captured by her information sets;
- players do not forget histories that they were aware of; and
- there is common knowledge of (1) what the payoffs are in the underlying game and (2) what the information sets are in the underlying game.

The formal conditions are not needed in this paper, so we omit them here.

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. A game with awareness collects all these different views, and describes, in each view, what view other players have. Formally, a *game with awareness based on* $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$ is 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 that i currently considers possible.

The augmented game Γ^m and the mapping \mathcal{F} must satisfy a number of consistency conditions. The conditions on the modeler's game ensures that the modeler is aware of all the players and moves of the underlying game, and that he understands how nature's moves work in the underlying game Γ . 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.

¹A game with perfect recall is one where, players remember all the actions they have performed and all the information sets they have passed through; see [Osborne and Rubinstein 1994] for the formal definition.

There are also ten constraints on the mapping \mathcal{F} that capture desirable properties of awareness. Rather than describing all ten constraints here, we briefly describe a few of them, to give some idea of the intuition behind these constraints. Suppose that $\mathcal{F}(\Gamma^+, h) = (\Gamma^h, I)$ and $A_i^+(h) = a$,² then the following conditions hold.

- C1. $\{\bar{h} : h \in H^h\} = a$, where \bar{h} is the subsequence of h consisting of all the moves in h that are also in the set of moves M of the underlying game Γ .
- C2. If $h' \in H^h$ and $P^h(h') = j$, then $A_j^h(h') \subseteq a$ and $M_{\bar{h}'} \cap \{m : \bar{h}' \cdot \langle m \rangle \in a\} = M_{h'}^h$.
- 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' .
- 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 $h_1 \in I'$, $h_1 \cdot \langle m \rangle \in H'$.

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. 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.

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 conditions on augmented games, 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.

It may seem that by making \mathcal{F} a function we cannot capture a player's uncertainty about the game being played or uncertainty about opponents' unawareness about histories. However, we

²As in our earlier paper, we use the convention that the components of a (standard or augmented) game Γ^+ are labeled with the same superscript $+$, so that we have M^+ , H^+ , A_i^+ , and so on.

can capture such uncertainty by folding it into nature’s initial move in the game the player consider possible while moving. It should be clear that this gives a general approach to capturing such uncertainties.

We identify a standard extensive game Γ 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 . This is the *canonical representation of Γ* as a game with awareness.

In [Halpern and Rêgo 2006], we discussed generalizations of games with awareness to include situations where players may be aware of their own unawareness and, more generally, games where players may not have common knowledge of the underlying game is; for example, players may disagree about what the payoffs or the information sets are. With these models, we can capture a situation where, for example, player i may think that another player j cannot make a certain a certain move, when in fact j can make such a move. For ease of exposition, we do not discuss these generalizations further here. However, it is not hard to show that the results of this paper can be extended to them in a straightforward way.

Feinberg [2004, 2005] also studied games with awareness. Feinberg [2005] gives a definition of extended Nash equilibrium in normal-form games. Feinberg [2004] 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. Li [2006a] has also provided a model of unawareness in extensive games, based on her earlier work on modeling unawareness [Li 2006b; Li 2006c]. See [Halpern and Rêgo 2006] for some further discussion of the relation between these approaches and ours.

3 GENERALIZED SEQUENTIAL EQUILIBRIUM

To explain generalized sequential equilibrium, we first review the notion of sequential equilibrium for standard games.

3.1 SEQUENTIAL EQUILIBRIUM FOR STANDARD GAMES

Sequential equilibrium is defined with respect to an *assessment*, a pair $(\vec{\sigma}, \mu)$ where $\vec{\sigma}$ is a strategy profile consisting of *behavioral strategies* and μ is a *belief system*, i.e., a function that determines for every information set I a probability μ_I over the histories in I . Intuitively, if I is an information set for player i , μ_I is i ’s subjective assessment of the relative likelihood of the histories in I . Roughly speaking, an assessment is a sequential equilibrium if for all players i , at every i -information set, (a) i chooses a best response given the beliefs he has about the histories in that information set and the strategies of other players, and (b) i ’s beliefs are consistent with the strategy profile being played, in the sense that they are calculated by conditioning the probability distribution induced by the strategy profile over the histories on the information set.

Note that μ_I is defined even if I is reached with probability 0. Defining consistency at an information set that is reached with probability 0 is somewhat subtle. In that case, intuitively, once information set I is reached player i moving at I must believe the game has been played according to an alternative strategy profile. In a sequential equilibrium, that alternative strategy profile consists of a small perturbation of the original assessment where every move is chosen with positive probability.

Given a strategy profile $\vec{\sigma}$, let $\text{Pr}_{\vec{\sigma}}$ be the probability distribution induced by $\vec{\sigma}$ over the possible histories of the game. Intuitively, $\text{Pr}_{\vec{\sigma}}(h)$ is the product of the probability of each of the moves in h . For simplicity we assume $f_c > 0$, so that if $\vec{\sigma}$ is such that every player chooses all of his moves with positive probability, then for every history h , $\text{Pr}_{\vec{\sigma}}(h) > 0$.³ For any history h of the game, define $\text{Pr}_{\vec{\sigma}}(\cdot | h)$ to be the conditional probability distribution induced by $\vec{\sigma}$ over the possible histories of the game given that the current history is h . Intuitively, $\text{Pr}_{\vec{\sigma}}(h' | h)$ is 0 if h is not a prefix of h' , is 1 if $h = h'$, and is the product of the probability of each of the moves in the path from h to h' if h is a prefix of h' . Formally, an assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium if it satisfies the following properties:

- *Sequential rationality.* For every information set I and player i and every behavioral strategy σ for player i ,

$$\text{EU}_i((\vec{\sigma}, \mu) | I) \geq \text{EU}_i((\vec{\sigma}_{-i}, \sigma), \mu | I),$$

where $\text{EU}_i((\vec{\sigma}, \mu) | I) = \sum_{h \in I} \sum_{z \in Z} \mu_I(h) \text{Pr}_{\vec{\sigma}}(z | h) u_i(z)$.

- *Consistency between belief system and strategy profile.* If $\vec{\sigma}$ consists of *completely mixed* (behavior) strategies, that is, ones that assign positive probability to every action at every information set, then for every information set I and history h in I ,

$$\mu_I(h) = \frac{\text{Pr}_{\vec{\sigma}}(h)}{\sum_{h' \in I} \text{Pr}_{\vec{\sigma}}(h')}.$$

Otherwise, there exists a sequence $(\vec{\sigma}^n, \mu^n)$, $n = 1, 2, 3, \dots$, of assessments such that $\vec{\sigma}^n$ consists of completely mixed strategies, $(\vec{\sigma}^n, \mu^n)$ is consistent in the above sense, and $\lim_{n \rightarrow \infty} (\vec{\sigma}^n, \mu^n) = (\vec{\sigma}, \mu)$.

Sequential equilibrium is not a reasonable solution concept for games with awareness for the same reason that Nash equilibrium is not a reasonable solution concept for games with awareness; it requires that a player be aware of the set of possible strategies available to other players and to him. In order to define a generalized notion of sequential equilibrium for games with awareness, we first need to define a generalized notion of assessment for games with awareness. We first need a generalized notion of strategy, which we defined in our earlier paper.

³See [Myerson 1991] for a definition of sequential equilibrium in the case nature chooses some of its move with probability 0.

Intuitively, a strategy describes what i will do in every possible situation that can arise. This no longer makes sense in games with awareness, since a player no longer understands in advance all the possible situations that can arise. For example, player i cannot plan in advance for what will happen if he becomes aware of something he is initially unaware of. We solved this problem in our earlier paper as follows. 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. Rather than considering a single strategy in a game $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ with awareness, we considered a collection $\{\sigma_{i,\Gamma'} : \Gamma' \in \mathcal{G}_i\}$ of *local strategies*. 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 $\mathcal{F}(\Gamma^+, h) = (\Gamma', I)$. Let $(\Gamma^h, I)^* = \{(\Gamma', h) : \mathcal{F}(\Gamma', h) = (\Gamma^h, I)\}$; we call $(\Gamma^h, I)^*$ a *generalized information set*.

Definition 3.1: Given a game $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ with awareness, 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 (Γ_1, h_1) and (Γ_2, h_2) are in the same generalized information set. A *generalized strategy profile* of $\Gamma^* = (\mathcal{G}, \Gamma^m, \mathcal{F})$ is a set of local strategies $\vec{\sigma} = \{\sigma_{i,\Gamma'} : i \in N, \Gamma' \in \mathcal{G}_i\}$.

The belief system, the second component of the assessment, is a function from information sets I to probability distribution over the histories in I . Intuitively, it captures how likely each of the histories in I is for the player moving at I . For standard games this distribution can be arbitrary, since the player considers every history in the information set possible. This is no longer true in games with awareness. It is possible that a player is playing game Γ_1 but believes he is playing a different game Γ_2 . Furthermore, in an augmented game, there may be some histories in an i -information set that include moves of which i is not aware; player i cannot consider these histories possible. To deal with these problems, we define μ to be a *generalized belief system* if it is a function from generalized information sets to a probability distribution over the set of histories in the generalized information set that the player considers possible.

Definition 3.2: A *generalized belief system* μ is a function that associates each generalized information set $(\Gamma', I)^*$ with a probability distribution $\mu_{\Gamma', I}$ over the set $\{(\Gamma', h) : h \in I\}$. A *generalized assessment* is a pair $(\vec{\sigma}, \mu)$, where $\vec{\sigma}$ is a generalized strategy profile and μ is a generalized belief system.

We can now define what it means for a generalized assessment $(\vec{\sigma}^*, \mu^*)$ to be a *generalized sequential equilibrium* of a game with awareness. The definition is essentially identical to that of $(\vec{\sigma}, \mu)$ being a sequential equilibrium; the use of EU_i in the definition of sequential rationality is replaced by $\text{EU}_{i,\Gamma'}$, where $\text{EU}_{i,\Gamma'}((\vec{\sigma}^*, \mu^*) \mid I)$ is the conditional expected payoff for i

in the game Γ' , given that strategy profile $\vec{\sigma}^*$ is used, information set I has been reached, and player i 's beliefs about the histories in I are described by $\mu_{\Gamma', I}^*$. We leave the straightforward modifications of the definition to the reader. It is easy to see that $(\vec{\sigma}, \mu)$ is a sequential equilibrium of a standard game Γ iff $(\vec{\sigma}, \mu)$ is a (generalized) sequential equilibrium of the canonical representation of Γ as a game with awareness. Thus, our definition of generalized sequential equilibrium generalizes the standard definition.

To better understand the concept of generalized sequential equilibrium concept, consider the game shown in Figure 2. Suppose that both players 1 and 2 are aware of all runs of the game, but player 1 (falsely) believes that player 2 is aware only of the runs not involving L and believes that player 2 is aware of these runs as well. Also suppose that player 2 is aware of all of this; that is, player 2's view of the game is the same as the modeler's view of the game Γ^m shown in Figure 2. While moving at node 1.1, player 1 considers the true game to be identical to the modeler's game except that from player 1's point of view, while moving at 2.1, player 2 believes the true game is $\Gamma^{2.2}$, shown in Figure 3.




Figure 2: The modeler's game Γ^m .




Figure 3: Player 2's view of the game from point of view of player 1.

This game has a unique generalized sequential equilibrium where player 2 chooses r and player 1 chooses A in $\Gamma^{2.2}$. Believing that player 2 will move r , player 1 best responds by choosing L at node 1.1. Since player 2 knows all this at node 2.1 in Γ^m , she chooses l at this node. Thus, if players follow their equilibrium strategies, the payoff vector is $(-10, -1)$. In this situation, player 2 would be better off if she could let player 1 know that she is aware of

move L , since then player 1 would play A and both players would receive 1. On the other hand, if we slightly modify the game by making $u_2(\langle L, l \rangle) = 3$, then player 2 would benefit from the fact that 1 believes that she is unaware of move L .

3.2 EXISTENCE OF GENERALIZED EQUILIBRIA

We now want to show that every game with awareness Γ^* has at least one generalized sequential equilibrium. To prove that a game with awareness Γ^* has a generalized Nash equilibrium, we constructed a standard game Γ^ν with perfect recall and showed that there exists a one-to-one correspondence between the set of generalized Nash equilibrium of Γ^* and the set of Nash equilibrium of Γ^ν . Intuitively, Γ^ν is constructed by essentially “gluing together” all the games $\Gamma' \in \mathcal{G}$, except that only histories in Γ' that can actually be played according to the players’ awareness level are considered.

More formally, 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.) For each $\Gamma' \in \mathcal{G}$, let $\lfloor H^{\Gamma'} \rfloor = \{h \in H^{\Gamma'} : \text{for every prefix } h_1 \cdot \langle m \rangle \text{ of } h, \text{ if } P'(h_1) = i \in N \text{ and } \mathcal{F}(\Gamma', h_1) = (\Gamma'', I), \text{ then for all } h_2 \in I, h_2 \cdot \langle m \rangle \in H^{\Gamma''}\}$. The histories in $\lfloor H^{\Gamma'} \rfloor$ are the ones that can actually be played according to the players’ awareness levels.

Let Γ^ν be the standard game such that

- $N^\nu = \{(i, \Gamma') : \Gamma' \in \mathcal{G}_i\}$;
- $M^\nu = \mathcal{G} \cup_{\Gamma' \in \mathcal{G}} \lfloor M^{\Gamma'} \rfloor$, where $\lfloor M^{\Gamma'} \rfloor$ is the set of moves that occur in $\lfloor H^{\Gamma'} \rfloor$;
- $H^\nu = \langle \rangle \cup \{\langle \Gamma' \rangle \cdot h : \Gamma' \in \mathcal{G}, h \in \lfloor H^{\Gamma'} \rfloor\}$;
- $P^\nu(\langle \rangle) = c$, 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 $P^h(h') = c$;
- $\mathcal{I}_{i, \Gamma'}^\nu$ is a partition of $H_{i, \Gamma'}^\nu$ where two histories $\langle \Gamma^1 \rangle \cdot h^1$ and $\langle \Gamma^1 \rangle \cdot h^1$ are in the same information set $\langle \Gamma', I \rangle^*$ iff (Γ^1, h^1) and (Γ^2, h^2) are in the same generalized information set $(\Gamma', I)^*$;
- $u_{i, \Gamma'}^\nu(\langle \Gamma^h \rangle \cdot z) = \begin{cases} u_i^h(z) & \text{if } \Gamma^h = \Gamma', \\ 0 & \text{if } \Gamma^h \neq \Gamma'. \end{cases}$

Unfortunately, while it is the case that there is a 1-1 correspondence between the Nash equilibria of Γ^p and the generalized Nash equilibria of Γ^* , this correspondence breaks down for sequential equilibria. To see why consider the modified version of prisoner's dilemma Γ^p , described in Figure 4.




Figure 4: Γ^p : modified prisoner's dilemma.

Besides being able to cooperate (C_A) or defect (D_A), player A who moves first has also the option of escaping (E_A). If player A escapes, then the game is over; if player A cooperates or defects, then player B may also cooperate (C_B) or defect (D_B). Suppose further that in the modeler's game,

- both A and B are aware of all histories of Γ^p ;
- with probability p , A believes that B is unaware of the extra move E_A , and with probability $1 - p$, A believes B is aware of all histories;
- if A believes B is unaware of E_A , then A believes that B believes that it is common knowledge that the game being played contains all histories but E_A ;
- if A believes B is aware of E_A , then A believes that B believes that there is common knowledge that the game being played is Γ^p ;
- B believes that it is common knowledge that the game being played is Γ^p .

We need four augmented games to model this situation:

- Γ^m is the game from the modeler's point of view;
- Γ^A is the game from A 's point of view when she is called to move in the modeler's game;

- $\Gamma^{B.1}$ is the game from B 's point of view when he is called to move in game Γ^A after nature chooses he is unaware of E_A and is also the game from A 's point of view when she is called to move in $\Gamma^{B.1}$; and
- $\Gamma^{B.2}$ is the game from B 's point of view when he is called to move in game Γ^A after nature chooses aware $_B$; $\Gamma^{B.2}$ is also the game from B 's point of view when he is called to move at Γ^m and the game from A 's point of view when she is called to move in $\Gamma^{B.2}$.

Although Γ^m and $\Gamma^{B.2}$ have the same game tree as Γ^p , they are different augmented games, since the \mathcal{F} function is defined differently at histories in these games. For example, $\mathcal{F}(\Gamma^m, \langle \rangle) = (\Gamma^A, \{\text{unaware}_B, \text{aware}_B\}) \neq (\Gamma^{B.2}, \langle \rangle) = \mathcal{F}(\Gamma^{B.2}, \langle \rangle)$. For this reason, we use different labels for the nodes of these games. Let $A.3$ and $B.2$ (resp., $A.2$ and $B.2$) be the labels of the nodes in game Γ^m (resp., $\Gamma^{B.2}$) corresponding to A and B in Γ^p , respectively. Γ^A and $\Gamma^{B.1}$ are shown in Figures 5 and 6, respectively.⁴

- In the modeler's game Γ^m , A believes she is playing game Γ^A , and B believes he is playing game $\Gamma^{B.2}$.
- In game Γ^A , nature chooses move unaware_B with probability p and aware_B with probability $1 - p$. Then A moves and believes she is playing Γ^A . At node $B.1$, B believes he is playing $\Gamma^{B.1}$, and at node $B.2$, B believes he is playing $\Gamma^{B.2}$.
- In game $\Gamma^{B.1}$, A and B both believe that the game is $\Gamma^{B.1}$.
- In game $\Gamma^{B.2}$, A and B both believe that the game is $\Gamma^{B.2}$.




Figure 5: Γ^A .

The game Γ^ν is the result of pasting together Γ^A , $\Gamma^{B.1}$, and $\Gamma^{B.2}$. There are 5 players: (A, Γ^A) , $(A, \Gamma^{B.1})$, $(A, \Gamma^{B.2})$, $(B, \Gamma^{B.1})$, and $(B, \Gamma^{B.2})$. $(A, \Gamma^{B.1})$ and $(B, \Gamma^{B.1})$ are playing

⁴We abuse notation and use the same label for nodes in different augmented games that are in the same generalized information set. For example, $A.3$ is a label at both Γ^m and Γ^A .




Figure 6: $\Gamma^{B.1}$.

standard prisoner's dilemma and therefore both should defect with probability 1; $(B, \Gamma^{B.1})$ must believe he is in the history where $(A, \Gamma^{B.1})$ defected with probability 1. (A, Γ^A) and $(A, \Gamma^{B.2})$ choose the extra move E_A with probability 1, since it gives A a payoff of 5. The subtlety arises in the beliefs of $(B, \Gamma^{B.2})$ in the generalized information set $(\Gamma^{B.2}, \{C_A, D_A\})^*$, since this generalized information set is reached with probability zero. Note that $(\Gamma^{B.2}, \{C_A, D_A\})^* = \{\langle \Gamma^m, C_A \rangle, \langle \Gamma^m, D_A \rangle, \langle \Gamma^A, \text{aware}_B, C_A \rangle, \langle \Gamma^A, \text{aware}_B, D_A \rangle, \langle \Gamma^{B.2}, C_A \rangle, \langle \Gamma^{B.2}, D_A \rangle\}$. By the definition of sequential equilibrium, player $(B, \Gamma^{B.2})$ will have to consider a sequence of strategies where all these histories are assigned positive probability. Although in general this is not a problem, note that $(B, \Gamma^{B.2})$ is meant to represent the type of player B that considers only histories in game $\Gamma^{B.2}$ possible. Thus, intuitively, he should assign positive probability only to the histories $\{\langle \Gamma^{B.2}, C_A \rangle, \langle \Gamma^{B.2}, D_A \rangle\}$.

To see how this leads to a problem, first note that there is a sequential equilibrium of Γ^ν where $(B, \Gamma^{B.2})$ believes with probability 1 that the true history is $\langle \Gamma^A, \text{aware}_B, C_A \rangle$, $(A, \Gamma^{B.2})$ chooses E_A with probability 1, and $(B, \Gamma^{B.2})$ chooses C_B with probability 1. It is rational for $(B, \Gamma^{B.2})$ to choose C_B because $(B, \Gamma^{B.2})$ assigns probability 1 to the first move of nature in Γ^ν be Γ^A . Since his utility is 0 for every run in Γ^ν whose first move is Γ^A , his expected utility is 0 no matter what move he makes at the generalized information set, given his beliefs.

There is no reasonable definition of generalized sequential equilibrium corresponding to this sequential equilibrium of Γ^ν . Player B while moving at node $B.2$ would never cooperate, since this is a strictly dominated strategy for him in the game that he considers to be the actual game, namely $\Gamma^{B.2}$.

The problem is that there is nothing in the definition of sequential equilibrium that guarantees that the belief system of a sequential equilibrium in Γ^ν assigns probability zero to histories that players are unaware of in the game Γ^* with awareness. We want to define a modified notion of sequential equilibrium for standard games that guarantees that the belief system in Γ^ν associates each information set with a probability distribution over a pre-specified subset of the histories in the information, which consists only of the histories in the information set player i actually considers possible. In this example, the pre-specified subset would be $\{\langle \Gamma^{B.2}, C_A \rangle, \langle \Gamma^{B.2}, D_A \rangle\}$.

4 CONDITIONAL SEQUENTIAL EQUILIBRIUM

In the standard definition of sequential equilibrium for extensive games, it is implicitly assumed that every player considers all histories in his information set possible. This is evident from the fact that if a strategy profile that is part of a sequential equilibrium assigns positive probability to every move, then by the consistency requirement the belief system also assigns positive probability to every history of every information set of the game. Therefore, this notion of equilibrium is not strong enough to capture situations where a player is certain that some histories in his information set will not occur. The notion of *conditional sequential equilibrium*, which we now define, is able to deal with such situations. It generalizes sequential equilibrium: in a game where every player considers every history in his information set possible, the set of conditional sequential equilibria and the set of sequential equilibria coincide.

Given a standard extensive game Γ , define a *possibility system* \mathcal{K} on Γ to be a function that determines for every information set I a nonempty subset of I consisting of the histories in I that the player moving at I considers possible. We assume that \mathcal{K} is common knowledge among players of the game, so that every player understands what histories are considered possible by everyone else in the game. If I is an i -information set, intuitively i should be indifferent among all runs that go through histories in $I - \mathcal{K}(I)$, since i believes that those runs will not occur and every other player knows that. Thus, for a given Γ , a possibility system \mathcal{K} must satisfy the following requirement: if z and z' are two runs going through histories in $I - \mathcal{K}(I)$ and I is an i -information set, then $u_i(z) = u_i(z')$.

Given a pair (Γ, \mathcal{K}) , a \mathcal{K} -*assessment* is a pair $(\vec{\sigma}, \mu)$, where $\vec{\sigma}$ is a strategy profile of Γ , and μ is a *restricted belief system*, i.e., a function that determines for every information set I of Γ a probability μ_I over the histories in $\mathcal{K}(I)$. Intuitively, if I is an information set for player i , μ_I is i 's subjective assessment of the relative likelihood of the histories player i considers possible while moving at I , namely $\mathcal{K}(I)$. As in the definition of sequential equilibrium, a \mathcal{K} -assessment $(\vec{\sigma}, \mu)$ is a *conditional sequential equilibrium with respect to \mathcal{K}* if (a) at every information set where a player moves he chooses a best response given the beliefs he has about the histories that he considers possible in that information set and the strategies of other players, and (b) his restricted beliefs must be consistent with the strategy profile being played and the possibility system, in the sense that they are calculated by conditioning the probability distribution induced by the strategy profile over the histories considered possible on the information set. Formally, the definition of $(\vec{\sigma}, \mu)$ being a conditional sequential equilibrium is identical to that of sequential equilibrium, except that the summation in the definition of $EU_i(\vec{\sigma}, \mu \mid I)$ and $\mu_I(h)$ is taken over histories in $\mathcal{K}(I)$ rather than histories in I . It is immediate that if $\mathcal{K}(I) = I$ for every information set I of the game, then the set of conditional sequential equilibria with respect to \mathcal{K} coincides with the set of sequential equilibria. The next theorem shows that set of conditional sequential equilibria for a large class of extensive games that includes Γ^ν is nonempty.

Theorem 4.1: *Let Γ be an extensive game with perfect recall and countably many players such that (a) each player has only finitely many pure strategies and (b) each player's payoff depends*

only on the strategy of finitely many other players. Let \mathcal{K} be an arbitrary possibility system. Then there exists at least one \mathcal{K} -assessment that is a conditional sequential equilibrium of Γ with respect to \mathcal{K} .

We now prove that every game of awareness has a generalized sequential equilibrium by defining a possibility system \mathcal{K} on Γ^ν and showing that there is a one-to-one correspondence between the set of conditional sequential equilibria of Γ^ν with respect to \mathcal{K} and the set of generalized sequential equilibria of Γ^* .

Theorem 4.2: *For all probability measures ν on \mathcal{G} , if ν gives positive probability to all games in \mathcal{G} , and $\mathcal{K}(\langle \Gamma', I \rangle^*) = \{ \langle \Gamma', h \rangle : h \in I \}$ for every information set $\langle \Gamma', I \rangle^*$ of Γ^ν , then $(\vec{\sigma}', \mu')$ is a generalized sequential equilibrium of Γ^* iff $(\vec{\sigma}, \mu)$ is a conditional sequential equilibrium of Γ^ν with respect to \mathcal{K} , where $\sigma_{i, \Gamma'}(\langle \Gamma^h \rangle \cdot h') = \sigma'_{i, \Gamma'}(\Gamma^h, h')$ and $\mu'_{\Gamma', I} = \mu_{\langle \Gamma', I \rangle^*}$.*

Since Γ^ν satisfies all the conditions of Theorem 4.1, it easily follows from Theorems 4.1 and 4.2 that every game with awareness has at least one generalized sequential equilibrium.

Although it is not true that every conditional sequential equilibrium is also a sequential equilibrium of an arbitrary game, the next theorem shows there is a close connection between these notions of equilibrium. If $(\vec{\sigma}, \mu)$ is a conditional sequential equilibrium with respect to some possibility system \mathcal{K} , then there exists a belief system μ' such that $(\vec{\sigma}, \mu')$ is a sequential equilibrium.

Theorem 4.3: *For every extensive game Γ with countably many players where each player has finitely many pure strategies and for every possibility system \mathcal{K} , if $(\vec{\sigma}, \mu)$ is a conditional sequential equilibrium of Γ with respect to \mathcal{K} , then there exists a belief system μ' such that $(\vec{\sigma}, \mu')$ is a sequential equilibrium of Γ .*

5 RATIONALIZABILITY AND GENERALIZED NASH EQUILIBRIUM

In this section, we analyze the relationship between the notions of rationalizability and generalized Nash equilibrium, providing some more intuition about the latter.

The usual justification for Nash equilibrium is that a player's strategy must be a best response to the strategies selected by other players in the equilibrium, because he can deduce what those strategies are. However, in most strategic situations, it is not the case that a player can deduce the strategies used by other players. Since every player tries to maximize his expected payoff and this is common knowledge, the best that a player can hope to do is to deduce a set of reasonable strategies for the other players. Here, we take a "reasonable strategy" to be a best response to some reasonable beliefs a player might hold about the strategy profile being played. This is the intuition that the *rationalizability* solution concept tries to capture. Even though a notion of rationalizability for extensive-form games was proposed by Pearce [1984], rationalizability is more widely applied in normal-form games. In this section, we

explore the relationship between rationalizability and generalized Nash equilibrium in games with awareness where in the underlying game each player moves only once, and these moves are made simultaneously (or, equivalently, a player does not know the moves made by other players before making his own move). We show that, given an underlying game Γ satisfying this requirement, a pure strategy profile contains only rationalizable strategies iff it is the strategy profile used by the players in the modeler's game in some (pure) generalized Nash equilibrium of a game Γ^* with awareness. If we think of rationalizability as characterizing “best response to your beliefs” and Nash equilibrium characterizing “best response to what is actually played”, then this result shows that in the framework of games with awareness, since the game is not common knowledge, the line between these two notions is somewhat blurred.

We start by reviewing the notion of rationalizability for standard normal-form games. Let \mathcal{C}_i be the set of available pure strategies for player i ; $\mathcal{C} = \times_{i \in N} \mathcal{C}_i$ is thus the set of pure strategy profiles. Let $\Delta(M)$ denote the set of all probability distributions on M . Suppose that each player i is rational and is commonly known to choose a strategy from a subset \mathcal{D}_i of \mathcal{C}_i . Let $\mathcal{D}_{-i} = \times_{j \neq i} \mathcal{D}_j$ and

$$B(\mathcal{D}_{-i}) = \{ \operatorname{argmax}_{s_i \in \mathcal{C}_i} \operatorname{EU}_i((s_i, \pi(\mathcal{D}_{-i}))) : \text{for some } \pi \in \Delta(\mathcal{D}_{-i}) \};$$

that is, $B(\mathcal{D}_{-i})$ consists of the strategies in \mathcal{C}_i that are best responses to some belief that player i could have about the strategies other players are using.

The set $\mathcal{S} = \times_{i \in N} \mathcal{S}_i$ of *correlated rationalizable strategies* is characterized by the following two properties: (a) for all $i \in N$, $\mathcal{S}_i \subseteq B(\mathcal{S}_{-i})$ and (b) \mathcal{S} is the largest set satisfying condition (a), in the sense that, for every set of strategy profiles \mathcal{D} satisfying (a), we have that $\mathcal{D} \subseteq \mathcal{S}$. It is not hard to show that for every player i , $\mathcal{S}_i = B(\mathcal{S}_{-i})$. A strategy $s_i \in \mathcal{S}_i$ is called a *correlated rationalizable strategy for player i* .^{5 6}

It turns out that we can construct \mathcal{S} by the following iterative procedure. Let $\mathcal{C}_i^0 = \mathcal{C}_i$ for all $i \in N$. Define $\mathcal{C}_i^j = B(\mathcal{C}_{-i}^{j-1})$ for $j \geq 1$. Since there are finitely many strategies it is easy to see that there exists a finite k such that $\mathcal{C}_i^j = \mathcal{C}_i^k$ for all $j \geq k$. It can be shown that $\mathcal{S}_i = \lim_{j \rightarrow \infty} \mathcal{C}_i^j = \mathcal{C}_i^k$. It is also easy to see that if $\vec{\sigma}$ is a (behavioral) Nash equilibrium, then every pure strategy that is played with positive probability according to $\vec{\sigma}$ is rationalizable (where the probability with which a pure strategy is played according to $\vec{\sigma}$ is the product of the probability of each of its moves according to $\vec{\sigma}$).

We now explore the relationship between rationalizability in an underlying game Γ in normal-form and generalized Nash equilibrium in a special class of games with awareness based on Γ . Given a standard game Γ and a pure strategy profile \vec{s} consisting of rationalizable

⁵From now on, we use $\vec{s} = (s_1, \dots, s_n)$ to denote pure strategy profiles, and will continue to use $\vec{\sigma}$ for possibly nonpure strategy profiles.

⁶In the literature, it is often assumed that each player chooses his strategy independently of the others and that this is common knowledge. If we make this assumption, we get a somewhat stronger solution concept (at least, if $|N| \geq 3$), which we call *uncorrelated rationalizability*. Essentially the same results as we prove here for correlated rationalizability hold for uncorrelated rationalizability; we omit further details here.

strategies, we define a game with awareness $\Gamma^*(\vec{s}) = (\mathcal{G}, \Gamma^m, \mathcal{F})$ such that (a) there exists a generalized Nash equilibrium of $\Gamma^*(\vec{s})$, where s_i is the strategy followed by player i in Γ^m , and (b) every local strategy in every pure generalized Nash equilibrium of $\Gamma^*(\vec{s})$ is rationalizable in Γ . To understand the intuition behind the construction, note that if s_i is a pure correlated rationalizable strategy of player i in Γ , then s_i must be a best response to some probability distribution π^{s_i} over the set \mathcal{S}_{-i} of pure correlated rationalizable strategies of i 's opponents. The idea will be to include a game Γ^{s_i} in \mathcal{G} that captures the beliefs that make s_i a best response. Let $\vec{s}^1, \dots, \vec{s}^m$ be the strategy profiles in \mathcal{S}_{-i} that get positive probability according to π^{s_i} . (There are only finitely many, since \mathcal{S}_{-i} consists of only pure strategies.) Let Γ^{s_i} be the game where nature initially makes one of m moves, say c_1, \dots, c_m (one corresponding to each strategy that gets positive probability according to π^{s_i}), where the probability of move c_j is $\pi^{s_i}(\vec{s}^j)$. After nature's choice a copy of Γ is played. All the histories in Γ^{s_i} in which player i is about to move are in the same information set of player i ; that is, player i does not know nature's move. However, all the other players know nature's move. Finally, all players are aware of all runs of Γ at every history in Γ^{s_i} . Note that if h is a history where player i thinks game Γ^{s_i} is the actual game, and believes that other players will play \vec{s}_{-i}^j after nature's move c_j , then player i believes that s_i is a best response at h .

Given a pure strategy profile \vec{s} of the game Γ , let $\Gamma^*(\vec{s}) = (\mathcal{G}, \Gamma^m, \mathcal{F})$ be the following game with awareness:

- $\Gamma^m = (\Gamma, \{A_i : i \in N\})$, where for every player i and every history $h \in H_i^m$, $A_i(h) = H$ (the set of all histories in Γ);
- $\mathcal{G} = \{\Gamma^m\} \cup \{\Gamma^{s'_i} : s'_i \in \mathcal{S}_i, i \in N\}$;
- for an augmented game in $\Gamma^+ \in (\mathcal{G} - \{\Gamma^m\})$ and a history h of Γ^+ of the form $\langle s'_{-i} \rangle \cdot h'$,
 - if $P^+(h) = i$, then $\mathcal{F}(\Gamma^+, h) = (\Gamma^+, I)$ where I is the information set containing h ;
 - if $P^+(h) = j \in N - \{i\}$ and s'_j is the strategy of player j specified by s'_{-i} , then $\mathcal{F}(\Gamma^+, h) = (\Gamma^{s'_j}, I)$, where I is the unique j -information set in game $\Gamma^{s'_j}$;
- for $h \in H_i^m$, $\mathcal{F}(\Gamma^m, h) = (\Gamma^{s_i}, I)$, where I is the unique i -information set in game Γ^{s_i} .

The intuition is that if \vec{s} is a strategy profile such that, for all $i \in N$, s_i is a rationalizable strategy for player i in Γ , then at the (unique) i -information set of Γ^m , i considers the actual game to be Γ^{s_i} . For this particular game with awareness, there exists a generalized Nash equilibrium where the strategy for each player i in the modeler's game is s_i . Conversely, only rationalizable strategies are used in any pure generalized Nash equilibrium of $\Gamma^*(\vec{s})$. There is only one small problem with this intuition: strategies in Γ and local strategies for augmented games in $\Gamma^*(\vec{s})$ are defined over different objects. The former are defined over information sets of the underlying game Γ and the latter are defined over generalized information sets of $\Gamma^*(s)$.

Fortunately, this problem is easy to deal with: we can identify a local strategy in $\Gamma^*(\vec{s})$ with a strategy in Γ in the obvious way. By definition of $\Gamma^*(\vec{s})$, for every player i and augmented game $\Gamma' \in \mathcal{G}_i$, the domain of the local strategy $\sigma_{i,\Gamma'}$ consists of a unique generalized information set. We denote this information set by $I_{i,\Gamma'}$. For each local strategy $\sigma_{i,\Gamma'}$ of $\Gamma^*(\vec{s})$, we associate the strategy $\underline{\sigma}_{i,\Gamma'}$ in the underlying game Γ such that $\sigma_{i,\Gamma'}(I_{i,\Gamma'}) = \underline{\sigma}_{i,\Gamma'}(I)$, where I is the unique i -information set in Γ .

The following theorem summarizes the relationship between correlated rationalizable strategies in Γ and generalized Nash equilibrium of games with awareness.

Theorem 5.1: *If Γ is a standard normal-form game and \vec{s} is a (pure) strategy profile such that for all $i \in N$, s_i is a correlated rationalizable strategy of player i in Γ , then*

- (i) *there is a (pure) generalized Nash equilibrium \vec{s}^* of $\Gamma^*(\vec{s})$ such that for every player i , $\underline{s}_{i,\Gamma}^* s_i = s_i$;*
- (ii) *for every (pure) generalized Nash equilibrium \vec{s}^* of $\Gamma^*(\vec{s})$, for every local strategy $s_{i,\Gamma'}^*$ for every player i in \vec{s}^* , the strategy $\underline{s}_{i,\Gamma'}^*$ is correlated rationalizable for player i in Γ .*

Note that Theorem 5.1 does not imply that for a fixed game with awareness, (pure) generalized Nash equilibrium and generalized rationalizability coincide. These notions are incomparable for standard extensive games (cf. [Battigalli 1997; Pearce 1984]), so the corresponding generalized notions are incomparable when applied to the canonical representation of a standard game as a game with awareness. If we restrict the underlying game to be in normal form, it can be shown that, just as in standard games, every strategy in a pure generalized Nash equilibrium is (generalized) rationalizable. Since rationalizability is usually defined for pure strategies in the literature [Myerson 1991; Osborne and Rubinstein 1994], we focused on that case here. But it is not hard to show that an analogue of Theorem 5.1 holds for behavioral rationalizable strategies as well.

6 DISCUSSION

Our focus in this paper has been on refinements of generalized Nash equilibrium in games with awareness. It is worth reconsidering here the conceptual basis for these solution concepts in the presence of awareness.⁷ As we noted, equilibrium refinements are used in standard games to eliminate some “undesirable” or “unreasonable” equilibria. Arguably, unreasonable equilibria pose an even deeper problem with unaware players. For example, one standard interpretation of a Nash equilibrium is that a player chooses his strategy at the beginning of the game, and then does not change it because he has no motivation for doing so (since his payoff is no higher when he changes strategies). But this interpretation is suspect in extensive-form games when a player makes a move that takes the game off the equilibrium path. It may seem

⁷We thank Aviad Heifetz and an anonymous referee for raising some of these issues.

unreasonable for a player to then play the move called for by his strategy (even if the strategy is part of a Nash equilibrium). A threat to blow up the world if I catch you cheating in a game may be part of a Nash equilibrium, and does not cause problems if in fact no one cheats, but it hardly seems credible if someone does cheat.

One way to justify the existence of incredible threats off the equilibrium path in a Nash equilibrium is to view the player as choosing a computer program that will play the game for him, and then leaving. Since the program is not changed once it is set in motion, threats about moves that will be made at information sets off the equilibrium path become more credible. However, in a game with awareness, a player cannot write a program to play the whole game at the beginning of the game because, when his level of awareness changes, he realizes that there are moves available to him that he was not aware of at the beginning of the game. He thus must write a new program that takes this into account. But this means we cannot sidestep the problem of incredible threats by appealing to the use of a pre-programmed computer to play a strategy. Once we allow a player to change his program, threats that were made credible because the program could not be rewritten become incredible again. Thus, the consideration of equilibrium refinements such as sequential equilibrium, which block incredible threats, becomes even more pertinent with awareness.

Moving up a level, we might ask more generally for the appropriate interpretation of Nash equilibrium in games with awareness. In standard games with a unique Nash equilibrium, we could perhaps argue that rational players will play their component of the equilibrium, since they can compute it and realize that it is the only stable strategy. In games with several Nash equilibria, perhaps one can be singled out as most salient, or some can be eliminated by using refinements of Nash equilibria.

To some extent, these considerations apply in games with awareness as well. If there is a unique generalized Nash equilibrium, although a player cannot necessarily compute the whole equilibrium (for example, if it involves moves that he is not aware of), he can compute that part of the equilibrium that is within the scope of his awareness. Thus, this argument for playing a Nash equilibrium lifts from standard games to games with awareness. However, other arguments do not lift so well. For example, in standard games, one argument for Nash equilibrium is that, over time, players will learn to play a Nash equilibrium, for example, by playing a best response to their current beliefs. This argument will not work in the presence of awareness, since playing the game repeatedly can make players aware of moves or of other players awareness, and thus effectively change the game altogether.

Another way of interpreting Nash equilibrium in standard games is in terms of evolutionary game theory. This approach works with awareness as well. Suppose that we have populations consisting of each awareness type of each player, and that at each time step we draw without replacement one individual of each of these populations and let them play the game once. If the sample individuals are playing an equilibrium strategy, they do not have incentive to deviate unilaterally given their beliefs that the other players will continue to follow the equilibrium strategies.

Other issues arise when considering sequential equilibrium in games with awareness. For

example, in a standard game, when a player reaches a history that is not on the equilibrium path, he must believe that his opponent made a mistake. However, in games with awareness, a player may become aware of her own unawareness and, as a result, switch strategies. In the definition of sequential equilibrium in standard games, play off the equilibrium path is dealt with by viewing it as the limit of “small mistakes” (i.e., small deviations from the equilibrium strategy). Given that there are alternative ways of dealing with mistakes in games with awareness, perhaps other approaches for dealing with off-equilibrium play might be more appropriate. While other ways of dealing with mistakes may well prove interesting, we would argue that our generalization of sequential equilibrium can be motivated the same way as in standard games. Roughly speaking, for us, how a player’s awareness level changes over time is not part of the equilibrium concept, but is given as part of the description of the game.

More generally, we have focused here on generalizing solution concepts that have proved useful in standard games, where there is no lack of awareness. The discussion above suggests that introducing awareness allows us to consider other solution concepts. For example, Ozbay [2006] proposes an approach where a player’s beliefs about the probability of revealed moves of nature, that the player was initially unaware of, are formed as part of the equilibrium definition. We hope to explore the issue of which solution concepts are most appropriate in games with awareness in future work.

7 CONCLUSIONS

In this paper, we further developed the framework of games with awareness by analyzing how to generalize sequential equilibrium to such games. Other solution concepts can be generalized in a similar way. Although we have not checked all the details for all solution concepts, we believe that techniques like those used in our earlier paper to prove existence of generalized Nash equilibrium and ones similar to those used in this paper for generalized sequential equilibrium will be useful for proving the existence of other generalized solution concepts. For example, consider the notion of (*trembling hand*) *perfect equilibrium* [Selten 1975]. In normal-form games. A strategy profile $\vec{\sigma}$ is a perfect equilibrium if there exists some sequence of strategies $(\vec{\sigma}^k)_{k=0}^{\infty}$, each assigning positive probability to every available move, that converges pointwise to $\vec{\sigma}$ such that for each player i , the strategy σ_i is a best response to $\vec{\sigma}_{-i}^k$ for all k . The definition of *generalized perfect equilibrium* in games with awareness is the same as in standard games, except that we use generalized strategies rather than strategies, and require that for every local strategy $\sigma_{i,\Gamma'}$ of every player i , $\sigma_{i,\Gamma'}$ is a best response to $\vec{\sigma}_{-(i,\Gamma')}^k$ in game Γ' for all k . To prove that every game with awareness in normal form has a generalized perfect equilibrium, we prove an analogue of Theorem 3.1(b) in [Halpern and Rêgo 2006], giving a correspondence between the set of generalized perfect equilibria of Γ^* and the set of perfect equilibria of Γ^ν . The existence of a generalized perfect equilibrium follows from the existence of a perfect equilibrium in Γ^ν ; the existence of a perfect equilibrium in Γ^ν follows from Lemma 7.2.

As a byproduct of our proof of existence of a generalized sequential equilibrium in each game with awareness, we proposed a concept of conditional sequential equilibrium for standard

games. This solution is more appropriate than standard sequential equilibrium if there are histories in an information set that, even though indistinguishable from a player's point of view, are not considered possible. We showed how to construct a standard game given a game with awareness such that there is a one-to-one correspondence between the set of generalized sequential equilibria of the game with awareness and the set of conditional sequential equilibria of the standard game. Roughly speaking, this result shows that a game with awareness is equivalent to a standard game where there are multiple versions of a player and it is possible that a player does not consider the actual history possible if it involves moves he is unaware of.

In our earlier work, we showed that our definitions could be extended in a straightforward way to games with awareness of unawareness; that is, 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. We captured 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 of was 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. We showed that a generalized Nash equilibrium exists in games with awareness of unawareness. It is straightforward to define generalized sequential equilibrium for those games and to prove its existence using the techniques of this paper; we leave details to the reader.

We also provided a further insight into the notion of generalized Nash equilibrium by analyzing its relationship with the notion of rationalizability for standard games. We showed that, in a sense, generalized Nash equilibrium can be viewed as a generalization of rationalizability. In particular, this shows that, unlike in the standard case where Nash equilibrium is characterized by every player best responding to the actual strategies played by their opponents, in games with awareness (or, more generally, in games with lack of common knowledge) a generalized Nash equilibrium is characterized by every player best responding to the strategy they believe their opponents are playing.

We have focused here on generalizing solution concepts that have proved useful in games where there is no lack of awareness. It may well be the case that, in games with (lack of) awareness, other solution concepts may also be appropriate. We hope to explore the issue of which solution concepts are most appropriate in games with awareness in future work.

Acknowledgments We thank Yossi Feinberg and Aviad Heifetz for useful discussions on awareness. 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. Some of this work was done while the first author was at the School of Electrical and Computer Engineering at Cornell University, U.S.A., supported in part by a scholarship from the Brazilian Government through the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq).

References

- Battigalli, P. (1997). On rationalizability in extensive games. *Journal of Economic Theory* 74, 40–61.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica* 52(4), 1007–1028.
- Feinberg, Y. (2004). Subjective reasoning—games with unawareness. Technical Report Research Paper Series #1875, Stanford Graduate School of Business.
- Feinberg, Y. (2005). Games with incomplete awareness. Technical Report Research Paper Series #1894, Stanford Graduate School of Business.
- Halpern, J. Y. and L. C. Rêgo (2006). Extensive games with possibly unaware players. In *Proc. Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 744–751. Full version available at arxiv.org/abs/0704.2014.
- Kreps, D. M. and R. B. Wilson (1982). Sequential equilibria. *Econometrica* 50, 863–894.
- Li, J. (2006a). Dynamic games with unawareness. Unpublished manuscript.
- Li, J. (2006b). Information structures with unawareness. Unpublished manuscript.
- Li, J. (2006c). Modeling unawareness without impossible states. Unpublished manuscript.
- Myerson, R. (1978). Refinements of the Nash equilibrium concept. *International Journal of Game Theory* 7, 73–80.
- Myerson, R. B. (1991). *Game Theory*. Cambridge, Mass.: Harvard University Press.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge, Mass.: MIT Press.
- Ozbay, E. (2006). Unawareness and strategic announcements in games with uncertainty. Unpublished manuscript.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4), 1029–1050.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 25–55.

7.1 PROOF OF THEOREMS 4.1, 4.2, AND 4.3

Theorem 4.1: *Let Γ be an extensive game with perfect recall and countably many players such that (a) each player has only finitely many pure strategies and (b) each player’s payoff depends only on the strategy of finitely many other players. Let \mathcal{K} be an arbitrary possibility system. Then there exists at least one \mathcal{K} -assessment that is a conditional sequential equilibrium of Γ with respect to \mathcal{K} .*

Proof: We use the same ideas that are used to prove existence of standard sequential equilibrium, following closely the presentation in [Myerson 1991]. The proof goes as follows. Given

Γ , let Γ_M be the *multiagent representation* of Γ in normal form. We prove (Lemma 7.1) that for every *perfect equilibrium* σ of Γ_M and every possibility system \mathcal{K} , there exists a restricted belief system μ such that (σ, μ) is a conditional sequential equilibrium of Γ with respect to \mathcal{K} . Then we show (Lemma 7.2) that for Γ satisfying the hypothesis of the theorem, Γ_M has at least one perfect equilibrium.

We now review the relevant definitions. A *normal-form game* is a tuple $(N, \times_{i \in N} \mathcal{C}_i, \{u_i : i \in N\})$, where N is the set of players of the game, \mathcal{C}_i is the collection of pure strategies available for player i in the game, and u_i is a payoff function that determines for each strategy profile in $\times_{i \in N} \mathcal{C}_i$ the payoff for player i .

Given a standard extensive-form game $\Gamma = (N, M, H, P, f_c, \{\mathcal{I}_i : i \in N\}, \{u_i : i \in N\})$, let $S^* = \cup_{i \in N} \mathcal{I}_i$. Intuitively, we associate with each i -information set $I \in \mathcal{I}_i$ a *temporary player* that has $M(I)$ as its set possible strategies; S^* is just the set of all temporary players in Γ . For each temporary player I we associate a payoff function $v_I : \times_{I \in S^*} M(I) \rightarrow \mathbb{R}$ such that if each temporary player I chooses action a_I , and σ is the pure strategy profile for Γ such that for every $i \in N$ and $I \in \mathcal{I}_i$, $\sigma_i(I) = a_I$, then $v_I(\times_{I \in S^*} a_I) = u_i(\sigma)$. The *multiagent representation for Γ in normal form* is the tuple $(S^*, \times_{I \in S^*} M(I), \{v_I : I \in S^*\})$.

Given any countable set B , let $\Delta(B)$ be the set of all probability measures over B , and let $\Delta^0(B)$ be the set of all probability measures over B whose support is all of B . Given a game in normal form $\Gamma = (N, \times_{i \in N} \mathcal{C}_i, \{u_i : i \in N\})$, a mixed strategy profile $\sigma \in \times_{i \in N} \Delta(\mathcal{C}_i)$ is a *perfect equilibrium* of Γ iff there exists a sequence $(\hat{\sigma}^k)_{k=1}^\infty$ such that (a) $\hat{\sigma}^k \in \times_{i \in N} \Delta^0(\mathcal{C}_i)$ for $k \geq 1$, (b) $\hat{\sigma}^k$ converges pointwise to σ , and (c) $\sigma_i \in \operatorname{argmax}_{\tau_i \in \Delta(\mathcal{C}_i)} \operatorname{EU}_i(\hat{\sigma}_{-i}^k, \tau_i)$ for all $i \in N$.

The following lemmas are analogues of Theorems 5.1 and 5.2 in [Myerson 1991].

Lemma 7.1: *If Γ_M is a multiagent representation of Γ in normal form, then for every perfect equilibrium σ of Γ_M and every possibility system \mathcal{K} , there exists a restricted belief system μ such that (σ, μ) is a conditional sequential equilibrium of Γ with respect to \mathcal{K} .*

Proof: The proof is almost identical to that of Theorem 5.1 in [Myerson 1991]. We focus on the necessary changes, leaving the task of verifying that the rest of the proof goes without change to the reader. Let $(\hat{\sigma}^k)_{k=1}^\infty \in \times_{I \in S^*} \Delta(M(I))$ be a sequence of behavioral strategy profiles satisfying conditions (a), (b), and (c) of the definition of perfect equilibrium. For each k , define a belief system μ^k such that, for each information set I , $\mu^k(I)$ is the probability over histories in $\mathcal{K}(I)$ defined as

$$\mu_I^k(h) = \frac{\operatorname{Pr}_{\hat{\sigma}^k}(h)}{\sum_{h' \in \mathcal{K}(I)} \operatorname{Pr}_{\hat{\sigma}^k}(h')}.$$

If \mathcal{I} is the set all information sets in Γ_M , for each k , $\mu^k : \mathcal{I} \rightarrow [0, 1]$. Thus, $\mu^k \in [0, 1]^{\mathcal{I}}$; and, by Tychonoff's Theorem, $[0, 1]^{\mathcal{I}}$ is compact. Thus, there must be a convergent subsequence of μ^1, μ^2, \dots . Suppose that this subsequence converges to μ . It is easy to see that μ is consistent with σ and \mathcal{K} .

Let $Z(I)$ denote the set of runs that do not contain any prefix in I . Let I be an arbitrary i -information set of Γ^ν . When agent $I \in S^*$ uses the randomized strategy $\rho_I \in \Delta(M(I))$ against the strategies specified by $\hat{\sigma}^k$ for all other agents, his expected payoff is

$$\text{EU}_I(\hat{\sigma}_{-I}^k, \rho_I) = \sum_{h \in I} \Pr_{(\hat{\sigma}_{-I}^k, \rho_I)}(h) \text{EU}_I(\hat{\sigma}_{-I}^k, \rho_I | h) + \sum_{z \in Z(I)} \Pr_{(\hat{\sigma}_{-I}^k, \rho_I)}(z) u_i(z).$$

Note that for $h \in I$ or $h \in Z(I)$, $\Pr_{(\hat{\sigma}_{-I}^k, \rho_I)}(h) = \Pr_{\hat{\sigma}^k}(h)$, since this probability is independent of the strategy used by player I . Also note that for all $h \in I - \mathcal{K}(I)$, $\text{EU}_I(\hat{\sigma}_{-I}^k, \rho_I | h)$ is independent of ρ_I . Thus,

$$\begin{aligned} & \text{EU}_I(\hat{\sigma}_{-I}^k, \rho_I) \\ &= \sum_{h \in \mathcal{K}(I)} \Pr_{\hat{\sigma}^k}(h) \text{EU}_I(\hat{\sigma}_{-I}^k, \rho_I | h) + \sum_{z \in Z(I)} \Pr_{\hat{\sigma}^k}(z) u_i(z) + C' \\ &= \left(\sum_{h \in \mathcal{K}(I)} \mu_I^k(h) \text{EU}_I(\hat{\sigma}_{-I}^k, \rho_I | h) \right) \left(\sum_{h \in \mathcal{K}(I)} \Pr_{\hat{\sigma}^k}(h) \right) + C'', \end{aligned}$$

where C' and C'' are two constants independent of ρ_I .

The rest of the proof proceeds just as the proof of Theorem 5.1 in [Myerson 1991]; we omit details here. ■

Lemma 7.2: *If Γ is an extensive-form game with perfect recall such that (a) there are at most countably many players, (b) each player has only finitely many pure strategies, and (c) the payoff of each player depends only on the strategy of finitely many other players, then Γ_M has at least one perfect equilibrium.*

Proof: The proof is almost identical to that of Theorem 5.2 in [Myerson 1991]. Again, we focus on the necessary changes, leaving it to the reader to verify that the rest of the proof goes without change. We need to modify some of the arguments since Γ_M is not a finite game; since it may contain countably many players. First, by the same argument used to prove that Γ^ν has at least one Nash equilibrium in our earlier work [Halpern and Rêgo 2006], we have that for any Γ satisfying the hypothesis of the lemma, Γ_M has at least one Nash equilibrium.

Let \mathcal{C}_i be the set of pure strategies available for player i in Γ_M . By Tychonoff's Theorem $\times_{i \in N} [0, 1]^{\mathcal{C}_i}$ is compact. Since $\times_{i \in N} \Delta(\mathcal{C}_i)$ is a closed subset of $\times_{i \in N} [0, 1]^{\mathcal{C}_i}$, it is also compact. All the remaining steps of the proof of Theorem 5.2 in [Myerson 1991] apply here without change; we omit the details. ■

The proof of Theorem 4.1 follows immediately from Lemmas 7.1 and 7.2. ■

Theorem 4.2: *For all probability measures ν on \mathcal{G} , if ν gives positive probability to all games in \mathcal{G} , and $\mathcal{K}(\langle \Gamma', I \rangle^*) = \{ \langle \Gamma', h \rangle : h \in I \}$ for every information set $\langle \Gamma', I \rangle^*$ of Γ^ν , then $(\vec{\sigma}', \mu')$*

is a generalized sequential equilibrium of Γ^* iff $(\vec{\sigma}, \mu)$ is a conditional sequential equilibrium of Γ^ν with respect to \mathcal{K} , where $\sigma_{i, \Gamma'}(\langle \Gamma^h \rangle \cdot h') = \sigma'_{i, \Gamma'}(\Gamma^h, h')$ and $\mu'_{\Gamma', I} = \mu_{\langle \Gamma', I \rangle^*}$.

Proof: Let $\Pr_{\vec{\sigma}}^\nu$ be the probability distribution over the histories in Γ^ν induced by the strategy profile $\vec{\sigma}$ and f_c^ν . For a history h of the game, define $\Pr_{\vec{\sigma}}^\nu(\cdot | h)$ to be the conditional probability distribution induced by $\vec{\sigma}$ and f_c^ν over the possible histories of the game given that the current history is h . Similarly, let $\Pr_{\vec{\sigma}'}^h$ be the probability distribution over the histories in $\Gamma^h \in \mathcal{G}$ induced by the generalized strategy profile $\vec{\sigma}'$ and f_c^h . Note that if $\Pr_{\vec{\sigma}'}^h(h') > 0$, then $h' \in \lfloor H^h \rfloor$. Thus, $\langle \Gamma^h \rangle \cdot h' \in H^\nu$.

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 $h' \in H^h$ such that $\Pr_{\vec{\sigma}'}^h(h') > 0$, we have that $\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^h \rangle \cdot h') = \nu(\Gamma^h) \Pr_{\vec{\sigma}'}^h(h')$. And since ν is a probability measure such that $\nu(\Gamma^h) > 0$ for all $\Gamma^h \in \mathcal{G}$, we have that $\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^h \rangle \cdot h') > 0$ iff $\Pr_{\vec{\sigma}'}^h(h') > 0$. It is also easy to see that for all $h' \neq \langle \rangle$ and all $h'' \in H^h$ such that $\Pr_{\vec{\sigma}'}^h(h'' | h') > 0$, $\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^h \rangle \cdot h'' | h') = \Pr_{\vec{\sigma}'}^h(h'' | h')$.

Suppose that $(\vec{\sigma}, \mu)$ is a conditional sequential equilibrium of Γ^ν with respect to \mathcal{K} . We first prove that $(\vec{\sigma}', \mu')$ satisfies generalized sequential rationality. Suppose, by way of contradiction, that it does not. Thus, there exists a player i , a generalized i -information set $(\Gamma^+, I)^*$, and a local strategy s' for player i in Γ^+ such that

$$\sum_{h \in I} \sum_{z \in Z^+} \mu'_{\Gamma^+, I}(h) \overset{+}{\Pr}_{\vec{\sigma}'}(z | h) u_i^+(z) < \sum_{h \in I} \sum_{z \in Z^+} \mu'_{\Gamma^+, I}(h) \overset{+}{\Pr}_{(\vec{\sigma}'_{-(i, \Gamma')}, s')} (z | h) u_i^+(z).$$

Define s to be a strategy for player (i, Γ^+) in Γ^ν such that $s(\langle \Gamma^h \rangle \cdot h') = s'(\Gamma^h, h')$. Using the observation in the previous paragraph and the fact that $\mu'_{\Gamma^+, I} = \mu_{\langle \Gamma^+, I \rangle^*}$ and $\mathcal{K}(\langle \Gamma^+, I \rangle^*) = \{ \langle \Gamma^+, h \rangle : h \in I \}$, it follows that

$$\begin{aligned} & \sum_{\langle \Gamma^+, h \rangle \in \mathcal{K}(\langle \Gamma^+, I \rangle^*)} \sum_{z \in \lfloor Z^+ \rfloor} \mu_{\langle \Gamma^+, I \rangle^*}(h) \overset{\nu}{\Pr}_{\vec{\sigma}}(\langle \Gamma^+ \rangle \cdot z | h) u_i^+(z) \\ < & \sum_{\langle \Gamma^+, h \rangle \in \mathcal{K}(\langle \Gamma^+, I \rangle^*)} \sum_{z \in \lfloor Z^+ \rfloor} \mu_{\langle \Gamma^+, I \rangle^*}(h) \overset{\nu}{\Pr}_{(\vec{\sigma}_{-(i, \Gamma')}, s)}(\langle \Gamma^+ \rangle \cdot z | h) u_i^+(z). \end{aligned} \tag{1}$$

By definition of u_{i, Γ^+}^ν , $u_i^+(z) = u_{i, \Gamma^+}^\nu(\langle \Gamma^+ \rangle, z)$. Replacing $u_i^+(z)$ by $u_{i, \Gamma^+}^\nu(\langle \Gamma^+ \rangle, z)$ in (1), it follows that $(\vec{\sigma}, \mu)$ does not satisfy sequential rationality in Γ^ν , a contradiction. So, $(\vec{\sigma}', \mu')$ satisfies generalized sequential rationality. It remains to show that μ' is consistent with $\vec{\sigma}'$.

Suppose that, for every generalized information set $(\Gamma^+, I)^*$, $\sum_{h \in I} \Pr_{\vec{\sigma}'}^+(h) > 0$. By definition of \mathcal{K} and the fact that for all $h' \in H^\nu$, $\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^+ \rangle \cdot h') > 0$ iff $\Pr_{\vec{\sigma}'}^+(h') > 0$, we have that for every information set $\langle \Gamma^+, I \rangle^*$ of Γ^ν ,

$$\sum_{\langle \Gamma^+, h \rangle \in \mathcal{K}(\langle \Gamma^+, I \rangle^*)} \overset{\nu}{\Pr}_{\vec{\sigma}}(\langle \Gamma^+ \rangle \cdot h) > 0.$$

Thus, by consistency of μ , $\vec{\sigma}$, and \mathcal{K} , it follows that for every information set $\langle \Gamma^+, I \rangle^*$ of Γ^ν and every $h \in \mathcal{K}(\langle \Gamma^+, I \rangle^*)$, we have

$$\mu_{\langle \Gamma^+, I \rangle^*}(h) = \frac{\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^+ \rangle \cdot h)}{\sum_{h' \in \mathcal{K}(\langle \Gamma^+, I \rangle^*)} \Pr_{\vec{\sigma}}^\nu(\langle \Gamma^+ \rangle \cdot h')}.$$

Since $\mu'_{\Gamma^+, I} = \mu_{\langle \Gamma^+, I \rangle^*}$, $\mathcal{K}(\langle \Gamma', I \rangle^*) = \{\langle \Gamma', h \rangle : h \in I\}$, and for all $h' \in H^h$ such that $\Pr_{\vec{\sigma}'}^h(h') > 0$, we have that $\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^h \rangle \cdot h') = \nu(\Gamma^h) \Pr_{\vec{\sigma}'}^h(h')$, it is easy to see that for every generalized information set $(\Gamma^+, I)^*$ and every $h \in I$,

$$\mu'_{\Gamma^+, I}(h) = \frac{\Pr_{\vec{\sigma}'}^+(h)}{\sum_{h' \in I} \Pr_{\vec{\sigma}'}^+(h')}.$$

Thus, μ' is consistent with $\vec{\sigma}'$.

Finally, suppose that there exists a generalized information set $(\Gamma^+, I)^*$ such that $\sum_{h \in I} \Pr_{\vec{\sigma}'}^+(h) = 0$. By definition of \mathcal{K} and the fact that for all $h' \in H^\nu$, $\Pr_{\vec{\sigma}}^\nu(\langle \Gamma^+ \rangle \cdot h') > 0$ iff $\Pr_{\vec{\sigma}'}^+(h') > 0$, we have that $\sum_{\langle \Gamma^+, h \rangle \in \mathcal{K}(\langle \Gamma^+, I \rangle^*)} \Pr_{\vec{\sigma}}^\nu(\langle \Gamma^+ \rangle \cdot h) = 0$. Thus, by the consistency of μ , $\vec{\sigma}$, and \mathcal{K} , there exists a sequence of \mathcal{K} -assessments $(\vec{\sigma}^n, \mu^n)$ such that $\vec{\sigma}^n$ consists of completely mixed strategies, μ^n is consistent with $\vec{\sigma}^n$ and \mathcal{K} , and $(\vec{\sigma}^n, \mu^n)$ converges pointwise to $(\vec{\sigma}, \mu)$.

Define a sequence of \mathcal{K} -assessments $(\vec{\tau}^n, \nu^n)$ such that $\nu^n_{\Gamma', I} = \mu^n_{\langle \Gamma', I \rangle^*}$ and $\sigma^n_{j, \Gamma'}(\langle \Gamma^h \rangle \cdot h') = \tau^n_{j, \Gamma'}(\Gamma^h, h')$ for all n . Since $\vec{\sigma}^n$ is completely mixed, so is $\vec{\tau}^n$; it also follows from the earlier argument that ν^n is consistent with $\vec{\tau}^n$ for all n . Since $(\vec{\sigma}^n, \mu^n)$ converges pointwise to $(\vec{\sigma}, \mu)$, it is easy to see that $(\vec{\tau}^n, \nu^n)$ converges pointwise to $(\vec{\sigma}', \mu')$. Thus, μ' is consistent with $\vec{\sigma}'$, and $(\vec{\sigma}', \mu')$ is a generalized sequential equilibrium of Γ^* , as desired. The proof of the converse is similar; we leave details to the reader. ■

Theorem 4.3: *For every extensive game Γ with countably many players where each player has finitely many pure strategies and for every possibility system \mathcal{K} , if $(\vec{\sigma}, \mu)$ is a conditional sequential equilibrium of Γ with respect to \mathcal{K} , then there exists a belief system μ' such that $(\vec{\sigma}, \mu')$ is a sequential equilibrium of Γ .*

Proof: Since (σ, μ) is a conditional sequential equilibrium of Γ with respect to \mathcal{K} , by the consistency of μ , σ , and \mathcal{K} , there exists a sequence of \mathcal{K} -assessments $(\hat{\sigma}^k, \hat{\mu}^k)$ such that $\hat{\sigma}^k$ is completely mixed, $\hat{\mu}^k$ is consistent with $\hat{\sigma}^k$ and \mathcal{K} , and $(\hat{\sigma}^k, \hat{\mu}^k)$ converges pointwise to (σ, μ) . Let $\hat{\nu}^k$ be the belief system consistent with $\hat{\sigma}^k$. Using the same techniques as in the proof of Lemma 7.1, we can construct a subsequence of $(\hat{\sigma}^k, \hat{\nu}^k)$ that converges pointwise to (σ, μ') . Thus, μ' is consistent with σ . It remains to show that (σ, μ') satisfies sequential rationality.

Since, by definition of \mathcal{K} , for every i -information set I of Γ , player i has the same utility for every run extending a history in $I - \mathcal{K}(I)$, it is not hard to show that

$$\text{EU}_i((\sigma, \mu') \mid I) = C + \mu'(K(I))\text{EU}_i((\sigma, \mu) \mid I),$$

where C and $\mu'(K(I))$ are independent of $\sigma_i(I)$. Since, by sequential rationality, $\sigma_i(I)$ is a best response given μ , it is also a best response given μ' . It follows that (σ, μ') is a sequential equilibrium of Γ , as desired. ■

7.2 PROOF OF THEOREM 5.1

Theorem 5.1: *If Γ is a standard normal-form game and \vec{s} is a (pure) strategy profile such that for all $i \in N$, s_i is a correlated rationalizable strategy of player i in Γ , then*

- (i) *there is a (pure) generalized Nash equilibrium \vec{s}^* of $\Gamma^*(\vec{s})$ such that for every player i , $\underline{s}_{i,\Gamma^{s_i}}^* = s_i$;*
- (ii) *for every (pure) generalized Nash equilibrium \vec{s}^* of $\Gamma^*(\vec{s})$, for every local strategy $s_{i,\Gamma'}^*$ for every player i in \vec{s}^* , the strategy $\underline{s}_{i,\Gamma'}^*$ is correlated rationalizable for player i in Γ .*

Proof: Let $\Gamma^*(\vec{s}) = (\mathcal{G}, \Gamma^m, \mathcal{F})$ be as defined in Section 5. For part (i), consider the generalized strategy profile \vec{s}^* where, for every player i and every $\Gamma^{s'_i} \in \mathcal{G}_i$, i makes the same move according to both s_i^* and s'_i . Note that, by definition of $\Gamma^*(\vec{s})$, for all $h \in H_i^m$ we have that $\mathcal{F}(\Gamma^m, h) = (\Gamma^{s_i}, \cdot)$. Thus, by definition of \vec{s}^* , for every player i , $\underline{s}_{i,\Gamma^{s_i}}^* = s_i$. It is easy to check using the definition of $\Gamma^*(\vec{s})$, that \vec{s}^* is a generalized Nash equilibrium, and that, for all $i \in N$, s_i is a rationalizable strategy for player i ; we leave details to the reader.

For part (ii), let $\mathcal{D}_i = \{\underline{s}_{i,\Gamma'}^* : \Gamma' \in \mathcal{G}_i\}$, i.e., \mathcal{D}_i consists of the strategies in the underlying game Γ corresponding to some local strategy of player i in $\Gamma^*(\vec{s})$. We claim that $\mathcal{D}_i \subseteq B(\mathcal{D}_{-i})$. To see this, let $s_{i,\Gamma'}^*$ be any local strategy for player i in \vec{s}^* . Since \vec{s}^* is a generalized Nash equilibrium, $s_{i,\Gamma'}^*$ is a best response to the local strategies used by other players in Γ' . Note that, by definition of \mathcal{D}_i , for every other player $j \neq i$, there is a strategy $\underline{s}_{j,\Gamma'} \in \mathcal{D}_j$ corresponding to the local strategy $s_{j,\Gamma'}$ player j follows in game Γ' . Since, by definition of $\Gamma^*(\vec{s})$, in game Γ' nature makes an initial choice and then a copy of Γ is played, and all players but i know the move made by nature, this initial move by nature can be seen as a distribution over the local strategies used by the other players in the different copies of Γ contained in Γ' . Thus, it is easy to see that the strategy $\underline{s}_{i,\Gamma'}^*$ corresponding to $s_{i,\Gamma'}^*$ is in $B(\mathcal{D}_{-i})$. Finally, since $s_{i,\Gamma'}^*$ is an arbitrary local strategy of player i in \vec{s}^* , it follows that $\mathcal{D}_i \subseteq B(\mathcal{D}_{-i})$. By the definition of correlated rationalizable strategies, it follows that $\mathcal{D}_i \subseteq \mathcal{S}_i$. Thus, for player i in $\Gamma(\vec{s})$ and every local strategy $s_{i,\Gamma'}^*$ for i in \vec{s}^* , $\underline{s}_{i,\Gamma'}^*$ is correlated rationalizable for player i in Γ , as desired. ■