# Reasoning About Rationality

Adam Bjorndahl, Joseph Y. Halpern, Rafael Pass

September 22, 2017

**Abstract**

We provide a sound and complete axiomatization for a class of logics appropriate for reasoning about the rationality of players in games, and show that essentially the same axiomatization applies to a very wide class of decision rules. We also consider epistemic representations of games in which players may be uncertain as to what decision rules their opponents are using, and define in this context a new solution concept, $\mathfrak{D}$-rationalizability.

## 1 Introduction

Decision-making in an uncertain environment is a fundamental component of game theory: players must choose what to do without necessarily knowing what their opponents will do. Under certainty, decision-making is straightforward: one simply chooses the course of action that leads to the most preferred outcome. Under uncertainty, however, a player must evaluate many possible outcomes in a manner that somehow takes into account her relative degrees of belief. There are many ways to do this. The *maximin* decision rule, for example, focuses entirely on worst-case scenarios. An *expected utility maximizer*, on the other hand, weights each outcome according to her (subjective) assessment of its probability and chooses the course of action that maximizes the corresponding expected value.

One can argue about which decision rules are reasonable and which are not, and the word "rational" might be invoked to denote this very divide. However, this is not the debate that concerns us here. Rather, assuming that we have fixed a decision rule for a given player, we can ask whether that player is, in fact, making choices in accordance with it, and call her *rational* precisely when she is. In classical game theory, for example, rationality is typically identified with expected utility maximization: a player is rational if and only if she is acting to maximize her expected utility.

Rationality in this sense plays a crucial role in the analysis of games; indeed, many notions of equilibrium require not only that each player is rational, but also that each player believes their opponents are rational, believes their opponents believe *their* opponents are rational, and so on. Epistemic game theory concerns itself with models expressive enough to represent the complexities of higher-order beliefs. Formal logic furnishes a powerful and versatile class of such models; namely, modal logics of belief and the Kripke structures typically used to give semantics to these logics. However, while the notion of rationality has been incorporated into these models both syntactically and semantically, no axiomatization of the resulting logical systems has been provided. This paper fills this gap.

We take as our point of departure axioms for rationality in the sense of expected utility maximization first presented in [Bjorndahl, Halpern, and Pass 2011]. We then extend these axioms to *arbitrary* decision rules; this allows us to reason about other standard rules beyond expected utility maximization, such as *maximin* and *minimax regret* (see [Halpern 2003] for a discussion of all the decision rules mentioned in this paper). We also consider the effect of representing uncertainty by a *set* of probability measures rather than a single one; this allows us to capture well-known decision rules such as *maxmin expected utility* and *minimax expected regret*. Finally, we turn our attention to modeling situations where players might be uncertain about which decision rules their opponents are using. Endogenizing decision rules in this way broadens not only the notion of rationality but also that of *iterative rationality*; this, we argue, provides a better epistemic foundation for a number of solution concepts.

The rest of this paper is organized as follows. In Section 2, we define the core concepts formally: games, modal logics of belief appropriate for reasoning about games, and the incorporation of rationality into these logics. Section 3 presents the main axiomatization together with a proof of soundness and completeness; we also extend these core results to logics in which the players' uncertainty is represented by sets of probabilities, and discuss the role of language. In Section 4, we consider logics in which players may be uncertain about the decision rules used by their opponents, and provide a natural application of this framework in the form of a new solution concept, $\mathfrak{D}$-rationalizability, as well as an axiomatization. Section 5 concludes with a discussion of future work. Proofs are collected in Appendix A.

## 2 Reasoning about games

Given a tuple $(X_i)_{i\in I}$ over some finite index set $I = \{1, \ldots, n\}$, we adopt the usual notational convention of writing

$$X \coloneqq \prod_{i\in I} X_i \quad \text{and} \quad X_{-i} \coloneqq \prod_{j\neq i} X_j.$$

We also write $X_i' \times X_{-i}$ for

$$X_1 \times \cdots \times X_{i-1} \times X_i' \times X_{i+1} \times \cdots \times X_n$$

and similarly $(x'_i, x_{-i})$ for $(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$.

A **(normal-form) game** is a tuple $\Gamma = (I, (\Sigma_i)_{i\in I}, (u_i)_{i\in I})$ where $I = \{1, \dots, n\}$ is the set of *players*, $\Sigma_i$ is the (finite) set of *strategies available to player* $i$, and $u_i : \Sigma \to \mathbb{R}$ is *player* $i$*'s utility function*, where $\Sigma = \prod_i \Sigma_i$ denotes the set of *strategy profiles*.

## 2.1 Syntax

One way of reasoning formally about a game is to build a logical language that is expressive enough to capture the aspects of play that we are interested in analyzing. To this end, given a game $\Gamma$, we begin by defining a propositional modal language of belief and then specializing the primitive propositions to correspond to the strategies available to the players.

Given an arbitrary set $\Phi$ of primitive propositions, let $\mathcal{L}_B(\Phi)$ be the language recursively generated by the grammar

$$\varphi ::= p \,|\, \neg\varphi \,|\, \varphi \wedge \psi \,|\, B_i\varphi,$$

where $p \in \Phi$ and $i \in I$. Here $\neg$ and $\wedge$ are the usual Boolean connectives corresponding to logical negation and conjunction, while $B_i\varphi$ is read "player $i$ believes $\varphi$". We define other connectives such as $\vee$, $\to$, and $\leftrightarrow$ as usual in terms of $\wedge$ and $\neg$, and write $\widehat{B}_i\varphi$ for $\neg B_i \neg\varphi$ ("player $i$ consider it possible that $\varphi$"). We also write

$$\begin{array}{rcl} E^1\varphi & \equiv & B_1\varphi \wedge \cdots \wedge B_n\varphi \text{ and} \\ E^k\varphi & \equiv & E^1(E^{k-1}\varphi) \end{array}$$

for "everyone believes $\varphi$" and its $k$-fold iteration, respectively. Set

$$\Phi_\Gamma \coloneqq \{\mathit{play}_i(\sigma_i) \,:\, i \in I,\, \sigma_i \in \Sigma_i\},$$

where we read $\mathit{play}_i(\sigma_i)$ as "player $i$ is playing strategy $\sigma_i$"; we write

$$\begin{array}{rcl} \mathit{play}(\sigma) & \equiv & \mathit{play}_1(\sigma_1) \wedge \cdots \wedge \mathit{play}_n(\sigma_n) \text{ and} \\ \mathit{play}_{-i}(\sigma_{-i}) & \equiv & \displaystyle\bigwedge_{j \neq i} \mathit{play}_j(\sigma_j) \end{array}$$

for "the players are playing according to the strategy profile $\sigma$" and the analogous statement regarding the players other than $i$. The language $\mathcal{L}_B(\Phi_\Gamma)$ is thus appropriate for reasoning about the beliefs of the players with respect to the strategies they are playing.

## 2.2 Semantics

A language of belief can be interpreted using Kripke-style possible world semantics, where associated to each world $\omega$ and each player $i$ is a probability measure on the set of all worlds, used to interpret player $i$'s beliefs at $\omega$. In the case of a language like $\mathcal{L}_B(\Phi_\Gamma)$, we also must take care to interpret the primitive propositions appropriately.

We restrict our attention to finite models. A **finite $\Gamma$-structure** is a tuple $M = (\Omega, (Pr_i)_{i \in I}, s)$ satisfying the following conditions:

(C1) $\Omega$ is a nonempty, finite set;

(C2) $Pr_i$ associates with each $\omega \in \Omega$ a probability measure $Pr_i(\omega)$ on $\Omega$;

(C3) $Pr_i(\omega)(\{\omega' \in \Omega \,:\, Pr_i(\omega') = Pr_i(\omega)\}) = 1$;

(C4) $s : \Omega \to \Sigma$ satisfies $Pr_i(\omega)(\{\omega' \in \Omega \,:\, s_i(\omega') = s_i(\omega)\}) = 1$.

Conditions (C1) and (C2) set the stage to interpret player $i$'s beliefs at $\omega$ by the measure $Pr_i(\omega)$. Condition (C3) then ensures that at each world $\omega$, each player is sure of (i.e. assigns probability 1 to) her own beliefs. Finally, condition (C4) establishes that the *strategy function* $s$ assigns to each world $\omega$ a strategy profile $s(\omega)$ in game $\Gamma$—intuitively, the strategy that each player is playing at $\omega$—and moreover, that each player is sure of her own strategy.

A finite $\Gamma$-structure $M$ induces an **interpretation** $[\![\cdot]\!]_M : \mathcal{L}_B(\Phi_\Gamma) \to 2^\Omega$ defined recursively as follows:

$$\begin{aligned} [\![play_i(\sigma_i)]\!]_M &:= \{\omega \in \Omega \,:\, s_i(\omega) = \sigma_i\} \\ [\![\varphi \wedge \psi]\!]_M &:= [\![\varphi]\!]_M \cap [\![\psi]\!]_M \\ [\![\neg\varphi]\!]_M &:= \Omega \setminus [\![\varphi]\!]_M \\ [\![B_i\varphi]\!]_M &:= \{\omega \in \Omega \,:\, Pr_i(\omega)([\![\varphi]\!]_M) = 1\}. \end{aligned}$$

Thus, the primitive propositions are interpreted in the obvious way using the strategy function ($s_i$ denotes the $i$th component function of $s$), the Boolean connectives are interpreted classically, and the formula $B_i\varphi$ holds at all and only those worlds $\omega$ at which $Pr_i(\omega)$ assigns probability 1 to $\varphi$. As is standard, we often write $(M, \omega) \models \varphi$ or just $\omega \models \varphi$ to indicate that $\omega \in [\![\varphi]\!]_M$. Similarly, we write $M \models \varphi$ if $[\![\varphi]\!]_M = \Omega$, and when $(M, \omega) \not\models \varphi$, we say that $M$ *refutes* $\varphi$ *(at* $\omega$*)*, or just that $\omega$ refutes $\varphi$.

## 2.3 Rationality

Informally, a player is *rational* if the strategy she is playing is a best response to her beliefs about the outcome of the game, given her preferences. But there is no single conception of what constitutes a "best response"; a wide variety of principles of decision-making have been proposed and studied.

One influential notion, particularly in game theory, is that of expected utility maximization. Given a game $\Gamma$ and a probability measure $\mu$ on $\Sigma_{-i}$ (thought of as representing player $i$'s beliefs about the strategies her opponents will play), the **expected utility** of a strategy $\sigma_i \in \Sigma_i$ is just the expected value of the function $u_i(\sigma_i, \cdot\,) : \Sigma_{-i} \to \mathbb{R}$ with respect to $\mu$:

$$EU_i(\sigma_i; \mu) \coloneqq \sum_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) \cdot \mu(\sigma_{-i}).$$

A best response for an expected utility maximizer is a strategy that maximizes this value. Abstractly, we might identify the mandate "maximize expected utility" for player $i$ (in the game $\Gamma$) with a function $\mathfrak{d}_i^{eu}$ that takes as input a belief $\mu$ on $\Sigma_{-i}$ and returns as output the set of strategies $\sigma_i \in \Sigma_i$ that maximize player $i$'s expected utility given $\mu$:

$$\mathfrak{d}_i^{eu}(\mu) \coloneqq \{\sigma_i \in \Sigma_i \;:\; (\forall \sigma_i' \in \Sigma_i)(EU_i(\sigma_i; \mu) \geq EU_i(\sigma_i'; \mu))\}.$$

Let $\Delta(\Sigma_{-i})$ denote the set of all probability measures on $\Sigma_{-i}$. Then, generalizing the above, we define a **decision rule for player $i$ (in $\Gamma$)** to be a function $\mathfrak{d}_i : \Delta(\Sigma_{-i}) \to 2^{\Sigma_i} \setminus \{\emptyset\}$. Intuitively, $\sigma_i \in \mathfrak{d}_i(\mu)$ just in case $\sigma_i$ is a best response for player $i$ to the belief $\mu$ according to the decision rule $\mathfrak{d}_i$. For another example, the "maximin" mandate for player $i$, which says to maximize the worst-case outcome among those considered possible, corresponds to the decision rule defined as follows: let

$$WC_i(\sigma_i; \mu) \coloneqq \min\{u_i(\sigma_i, \sigma_{-i}) \;:\; \mu(\sigma_{-i}) > 0\},$$

and set

$$\mathfrak{d}_i^{m}(\mu) \coloneqq \{\sigma_i \in \Sigma_i \;:\; (\forall \sigma_i' \in \Sigma_i)(WC_i(\sigma_i; \mu) \geq WC_i(\sigma_i'; \mu))\}.$$

Decision rules can be interpreted in $\Gamma$-structures; roughly speaking, for each world $\omega$, we can define the set of $\mathfrak{d}_i$-best responses for player $i$ at $\omega$, and thus determine whether or not player $i$ is being $\mathfrak{d}_i$-rational (i.e., acting in accordance with the decision rule $\mathfrak{d}_i$) at $\omega$. Formally, given a $\Gamma$-structure $M$, for each player $i$ and each world $\omega$, the probability measure $Pr_i(\omega)$ induces a probability measure $\mu_{i,\omega}$ defined on $\Sigma_{-i}$ as follows:

$$\mu_{i,\omega}(\sigma_{-i}) \coloneqq Pr_i(\omega)(\llbracket play_{-i}(\sigma_{-i}) \rrbracket_M).$$

In fact, $\mu_{i,\omega}$ is the *pushforward* of $Pr_i(\omega)$ by $s_{-i}$. Since $Pr_i(\omega)$ is interpreted as representing player $i$'s beliefs at $\omega$, it makes sense to apply her decision rule to $\mu_{i,\omega}$. This leads us to define the set of **$\mathfrak{d}_i$-best responses for player $i$ at $\omega$** to be $\mathfrak{d}_i(\mu_{i,\omega})$. We say that player $i$ is **$\mathfrak{d}_i$-rational at $\omega$** just in case $s_i(\omega) \in \mathfrak{d}_i(\mu_{i,\omega})$.

Since we wish to reason formally about games using the logic developed above, and rationality is an important concept for such analyses, we expand the language to include primitive propositions denoting rationality of the players. Given a profile of decision rules $\mathfrak{d} = (\mathfrak{d}_i)_{i \in I}$, let

$$\Phi_\Gamma^{\mathfrak{d}} \coloneqq \Phi_\Gamma \cup \{RAT_i^{\mathfrak{d}_i} \;:\; i \in I\},$$

where $RAT_i^{\mathfrak{d}_i}$ is read "player $i$ is $\mathfrak{d}_i$-rational". Henceforth, except where noted otherwise, we work with an arbitrary but fixed profile of decision rules $\mathfrak{d}$. For notational convenience, we omit the $\mathfrak{d}_i$ when possible, referring to "best responses" and "rationality" instead of "$\mathfrak{d}_i$-best responses" and "$\mathfrak{d}_i$-rationality", and writing $RAT_i$ instead of $RAT_i^{\mathfrak{d}_i}$. We also make use of the syntactic abbreviation

$$RAT \quad \equiv \quad RAT_1 \wedge \cdots \wedge RAT_n$$

for "everyone is rational".

Given a $\Gamma$-structure $M$, the interpretation $[\![\cdot]\!]_M$ is extended to $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ in the obvious way, by setting

$$[\![RAT_i]\!]_M \quad \coloneqq \quad \{\omega \in \Omega \, : \, s_i(\omega) \in \mathfrak{d}_i(\mu_{i,\omega})\}.$$

When $\mathfrak{d} = \mathfrak{d}^{eu}$, so each player's decision rule is given by expected utility maximization, we regain the traditional notion of rationality in game theory. Rationality so defined can be used to characterize several well-known solution concepts in terms of $\Gamma$-structures. For example, as shown by Tan and Werlang [1988] and Brandenburger and Dekel [1987], given a game $\Gamma$, a strategy $\sigma_i$ is *rationalizable* if and only if there exists a $\Gamma$-structure $M$ and a state $\omega$ therein such that $\omega \models play_i(\sigma_i)$ and for every $k \in \mathbb{N}$, $\omega \models E^k(RAT)$ (i.e., it is common belief that everyone is rational).

It is worth noting that decision rules are general enough to represent processes that fall outside the traditional purview of "rationality". For example, suppose that player $i$ is a computer system and $\Sigma_i$ is a collection of actions it can execute. Suppose also that the system maintains a database consisting of estimates of the values of certain variables, which can be represented as beliefs about the strategies of "opponents". In this context, a decision rule can be thought of as a (nondeterministic) process that the computer system might use to choose which action to execute on the basis of the information in its database. In particular, $RAT_i^{\mathfrak{d}_i}$ asserts that the system has executed an action consistent with the process $\mathfrak{d}_i$.[1]

**Example 2.1:** In order to help solidify the framework just introduced, we present a simple example. Consider the standard Bach-or-Stravinsky game $\Gamma_{\mathrm{BoS}}$ [Osborne and Rubinstein 1994], in which each of two players must choose which of two concerts to attend this evening: one featuring the music of Bach, and one of Stravinsky. Player 1 prefers to attend the Bach concert, while player two prefers the Stravinsky; moreover, each much prefers to attend the same concert as the other. We can represent these preferences with the utility functions summarized in Table 1.

[1] In this sense, we can think of a decision rule as a *knowledge-based program* [Fagin, Halpern, Moses, and Vardi 1995]. Knowledge-based programs have previously been used to characterize rationality and solution concepts [Halpern and Moses 2007].

| | Bach | Stravinsky |
|---|---|---|
| Bach | 3,2 | 0,0 |
| Stravinsky | 0,0 | 2,3 |

Table 1: Bach or Stravinsky?

We now describe a $\Gamma_{\mathrm{BoS}}$-structure in which we can reason about the beliefs and rationality of the players. Of course, there are many such $\Gamma_{\mathrm{BoS}}$-structures, each representing a different configuration of facts and beliefs; the one we consider here is chosen simply to provide a concrete illustration of the connection between the logical formalism and the game.

Let $\Omega = \{\omega_0, \omega_1, \omega_2\}$ and let $Pr_1(\omega_0)$ be the uniform distribution on $\{\omega_1, \omega_2\}$, so player 1 considers $\omega_1$ and $\omega_2$ equally likely. Consider a strategy function $s$ defined such that $s_2(\omega_1) = \text{Bach}$ and $s_2(\omega_2) = \text{Stravinsky}$. It is easy to see that in this case, the induced measure $\mu_{1,\omega_0}$ assigns probability .5 to Bach and .5 to Stravinsky. It follows easily that player 1's expected utility on choosing Bach is 1.5, while his expected utility on choosing Stravinsky is 1. Thus,

$$\omega_0 \models RAT_1^{\mathfrak{d}_1^{eu}} \leftrightarrow \mathit{play}_1(\text{Bach}).$$

On the other hand, the worst-case outcome for player 1 on choosing either Bach or Stravinsky yields a utility of 0, so we have

$$\omega_0 \models RAT_1^{\mathfrak{d}_1^{m}},$$

regardless of which strategy player 1 actually chooses at $\omega_0$. ∎

# 3 Axiomatization

Although reasoning formally about beliefs and rationality has long been recognized as important to game theory (see, e.g., [Aumann 1999]), to the best of our knowledge, rationality has not been axiomatized in a logic of belief. In this section we provide an axiomatization and prove that it is sound and complete.

While most of the axioms are easy to state and understand, two require some preliminary definitions. Fix a game $\Gamma$ and a profile of decision rules $\mathfrak{d}$. Given a subset $S \subseteq \Sigma_{-i}$, let

$$\delta_{i,S} \equiv \bigwedge_{\sigma_{-i} \in S} \widehat{B}_i\, \mathit{play}_{-i}(\sigma_{-i}) \;\wedge \bigwedge_{\sigma_{-i} \notin S} \neg \widehat{B}_i\, \mathit{play}_{-i}(\sigma_{-i}).$$

Intuitively, the formula $\delta_{i,S}$ says that player $i$ considers possible *all* and *only* the strategy profiles for his opponents that are elements of $S$. It is easy to see that for each player $i$ and any world $\omega$, exactly one of the formulas $\delta_{i,S}$ holds.

Given a measure $\mu$ on a finite set $X$, let

$$supp(\mu) \coloneqq \{x \in X \,:\, \mu(x) > 0\},$$

the **support** of $\mu$. For each player $i$ and each $\sigma_i \in \Sigma_i$, let $\mathcal{S}_i^+(\sigma_i)$ denote the collection of all $S \subseteq \Sigma_{-i}$ such that there exists a probability measure $\mu$ with $supp(\mu) = S$ and with respect to which $\sigma_i$ is a best response; that is, $\sigma_i \in \mathfrak{d}_i(\mu)$. Similarly, define $\mathcal{S}_i^-(\sigma_i)$ to be the collection of all $S \subseteq \Sigma_{-i}$ such that there exists a probability measure $\mu$ with $supp(\mu) = S$ and $\sigma_i \notin \mathfrak{d}_i(\mu)$.

Consider the following axiom schemes:

**G1.** $\bigvee_{\sigma_i \in \Sigma_i} play_i(\sigma_i)$

**G2.** $\neg(play_i(\sigma_i) \wedge play_i(\sigma_i'))$, for $\sigma_i \neq \sigma_i'$

**G3.** $play_i(\sigma_i) \leftrightarrow B_i\, play_i(\sigma_i)$

**G4.** $RAT_i \leftrightarrow B_i(RAT_i)$

**G5.** $(play_i(\sigma_i) \wedge RAT_i) \rightarrow \bigvee_{S \in \mathcal{S}_i^+(\sigma_i)} \delta_{i,S}$

**G6.** $(play_i(\sigma_i) \wedge \neg RAT_i) \rightarrow \bigvee_{S \in \mathcal{S}_i^-(\sigma_i)} \delta_{i,S}$.

**G1**–**G4** are straightforward. **G1** and **G2** say that, at each state, each player plays exactly one strategy. **G3** and **G4** say that each player is certain of her strategy and of whether or not she is rational. The interesting axioms are **G5** and **G6**. Intuitively, **G5** says that if $RAT_i$ holds and player $i$ is playing $\sigma_i$, then player $i$ must consider possible a collection of strategy profiles[2] on which she could place a probability that would justify her playing $\sigma_i$. **G6** is interpreted analogously. Notice that player $i$'s actual (i.e. quantitative) beliefs are not fully specified. This may be somewhat surprising, given that the semantic interpretation of $RAT_i$ is defined in terms of the quantitative probabilities that constitute player $i$'s beliefs; we return to this point in Section 3.2.

Let $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ be the axiom system that results from adding **G1**–**G6** to the standard $\mathsf{KD45}$ axioms and rules of inference of belief logic (see, e.g., [Fagin, Halpern, Moses, and Vardi 1995]). These axioms completely characterize the logical properties of rationality as expressible in the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$.

**Theorem 3.1:** $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ *is a sound and complete axiomatization of the language* $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ *with respect to the class of all finite* $\Gamma$*-structures.*

[2] Note that the sets $\mathcal{S}_i^+(\sigma_i)$ and $\mathcal{S}_i^-(\sigma_i)$ could be empty; by convention, we define the empty disjunction to be $\bot$ (or *false*). If, for example, $\mathcal{S}_i^+(\sigma_i) = \emptyset$, this means that $\sigma_i$ is not a best response to any beliefs, in which case the corresponding axiom $(play_i(\sigma_i) \wedge RAT_i) \rightarrow \bot$ is intuitively correct.

It bears emphasizing that $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ is parametrized by two variables: the underlying game $\Gamma$ that determines the players, their strategies, and their preferences, and the profile of decision rules $\mathfrak{d}$ that determines the meaning of "rationality" for each player. Thus, what we are axiomatizing here is not a single logic but a class of logics. For each fixed $\Gamma$ and $\mathfrak{d}$, the corresponding axiom system $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ is trivially decidable (i.e., we can effectively determine whether a formula is an instance of an axiom scheme) because $\mathsf{KD45}$ is decidable and there are only finitely-many instances of the axiom schemes **G1**–**G6** (the implicit quantification in these schemes ranges in some cases over players and in others over strategies, all of which are finite sets).

## 3.1 Belief as lower probability

In the above we take for granted that each player's uncertainty is represented by a probability measure. While this is a very standard assumption, it is by no means the only framework that has been considered; see [Halpern 2003] for an overview of different ways of modeling uncertainty. Here we show that, with very minor modifications, the axiomatization given above also works in the more general context where beliefs are represented using *sets* of probability measures.

Given a set $\mathcal{P}$ of probability measures, the **lower probability** of an event $E$, denoted $\mathcal{P}_*(E)$, is defined to be the infimum of the probabilities assigned to $E$ by members of $\mathcal{P}$:

$$\mathcal{P}_*(E) \coloneqq \inf\{\mu(E) \,:\, \mu \in \mathcal{P}\}.$$

Fix a game $\Gamma$. A **finite lower $\Gamma$-structure** is a tuple $M = (\Omega, (\mathcal{PR}_i)_{i \in I}, s)$ satisfying the following conditions:

(L1) $\Omega$ is a nonempty, finite set;

(L2) $\mathcal{PR}_i$ associates to each $\omega \in \Omega$ a set $\mathcal{PR}_i(\omega)$ of probability measures on $\Omega$;

(L3) $\mathcal{PR}_i(\omega)_*(\{\omega' \in \Omega \,:\, \mathcal{PR}_i(\omega') = \mathcal{PR}_i(\omega)\}) = 1$;

(L4) $s : \Omega \to \Sigma$ satisfies $\mathcal{PR}_i(\omega)_*(\{\omega' \in \Omega \,:\, s_i(\omega') = s_i(\omega)\}) = 1$.

These conditions are simply the analogues of conditions (C1) through (C4) where uncertainty is represented by sets of probability measures and certainty is identified with lower probability 1. Accordingly, we define the interpretation $[\![\cdot]\!]_M$ as before, except for the clause corresponding to the belief modalities, which is replaced by the following:

$$[\![B_i\varphi]\!]_M \quad \coloneqq \quad \{\omega \in \Omega \,:\, \mathcal{PR}_i(\omega)_*([\![\varphi]\!]_M) = 1\}.$$

Finally, a **decision rule** for player $i$ in this context is a function $\mathfrak{d}_i : 2^{\Delta(\Sigma_{-i})} \to 2^{\Sigma_i} \setminus \{\emptyset\}$, since player $i$ must make her choice based on the uncertainty given by a set of probability measures. For example, the "maximin expected utility" decision rule for player $i$ would be given by the following:

$$\mathfrak{d}_i^{meu}(\mathcal{P}) := \big\{\sigma_i \in \Sigma_i \,:\, (\forall \sigma_i' \in \Sigma_i)\big(\min_{\mu \in \mathcal{P}}\{EU_i(\sigma_i;\mu)\} \geq \min_{\mu \in \mathcal{P}}\{EU_i(\sigma_i';\mu)\}\big)\big\}.$$

Other rules, such as minimax expected regret [Hayashi 2008], can also easily be defined in this setting.

As before, such decision rules makes sense in a $\Gamma$-structure $M$: for each player $i$ and each world $\omega$, the probability measures in the set $\mathcal{PR}_i(\omega)$ can be pushed forward by $s_{-i}$ to probability measures on $\Sigma_{-i}$. Let $\mathcal{P}_{i,\omega}$ denote the set of all such pushforwards:

$$\mathcal{P}_{i,\omega} := \{\mu_{i,\omega} \,:\, \mu \in \mathcal{PR}_i(\omega)\}.$$

Then we can define **$\mathfrak{d}_i$-rationality** for player $i$ at $\omega$ in the obvious way, namely, by the requirement that $s_i(\omega) \in \mathfrak{d}_i(\mathcal{P}_{i,\omega})$.

The axiomatization of Section 3 can be generalized as well. First observe that the dual belief modality, $\widehat{B}_i \equiv \neg B_i \neg$, is interpreted as positive *upper* probability, where the **upper probability** of an event $E$ with respect to a set $\mathcal{P}$ of probability measures, denoted $\mathcal{P}^*(E)$, is given by

$$\mathcal{P}^*(E) := \sup\{\mu(E) \,:\, \mu \in \mathcal{P}\}.$$

Accordingly, given a set $\mathcal{P}$ of probability measures on a finite space $X$, we define

$$\mathit{supp}(\mathcal{P}) := \{x \in X \,:\, (\exists \mu \in \mathcal{P})(\mu(x) > 0)\}.$$

For each $\sigma_i \in \Sigma_i$, let $\mathcal{S}_i^+(\sigma_i)$ denote the collection of all $S \subseteq \Sigma_{-i}$ such that there exists a set of probability measures $\mathcal{P}$ with $\mathit{supp}(\mathcal{P}) = S$ and $\sigma_i \in \mathfrak{d}_i(\mathcal{P})$. Similarly, define $\mathcal{S}_i^-(\sigma_i)$ to be the collection of all $S \subseteq \Sigma_{-i}$ such that there exists a set $\mathcal{P}$ of probability measures with $\mathit{supp}(\mathcal{P}) = S$ and $\sigma_i \notin \mathfrak{d}_i(\mathcal{P})$. It is not hard to see that all of these definitions generalize what was presented in Sections 2 and 3; indeed, by considering the special case where all sets of probability measures are singletons, we recover that framework exactly. Moreover, the axiom system $\mathsf{GL}_\Gamma^\mathfrak{d}$, interpreted in this more general setting using the definitions above, remains sound and complete.

**Theorem 3.2:** *$\mathsf{GL}_\Gamma^\mathfrak{d}$ is a sound and complete axiomatization of the language $\mathcal{L}_B(\Phi_\Gamma^\mathfrak{d})$ with respect to the class of all finite lower $\Gamma$-structures.*

## 3.2 The role of language

In this section we focus on the profile of decision rules $\mathfrak{d} = \mathfrak{d}^{eu}$, with respect to which each player is rational precisely if they are playing a strategy that maximizes their expected utility. As noted, it is somewhat surprising that **G5** and **G6** are sufficient to capture this notion of rationality. Whether or not a player is maximizing their expected utility depends on their *quantitative* beliefs; however, while **G5** and **G6** specify the possible supports for player $i$'s beliefs, they say nothing about the actual weights placed on the individual outcomes. Nor could they—the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ cannot express anything beyond such qualitative properties of the measures $Pr_i(\omega)$.

Expressivity, however, is a double-edged sword: when working with a less expressive language, though we are more limited in the possible axioms we have available, there are also fewer validities to worry about proving. This, in essence, is why $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ can be a complete axiomatization: the properties of rationality it fails to encode are precisely those properties that are not expressible in the language at all.

A richer language—in particular, one with a finer-grained representation of belief—may not be axiomatizable at all, or at least not using the techniques in this paper. Consider, for example, a language with belief modalities $B_i^\alpha$ for each $\alpha \in [0,1]$, where $B_i^\alpha \varphi$ is interpreted as saying that player $i$ assigns probability $\alpha$ to $\varphi$. In this case, $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ (replacing $B_i$ by $B_i^1$) is sound but certainly not complete. It can easily happen, for example, that in the game $\Gamma$ it is only rational for player $i$ to play $\sigma_i$ if she assigns probability $\frac{1}{2}$ to $\sigma_{-i}$; however, the corresponding validity

$$(play_i(\sigma_i) \wedge RAT_i) \rightarrow B_i^{\frac{1}{2}}\, play_{-i}(\sigma_{-i})$$

is clearly not a theorem of $\mathsf{GL}_\Gamma^{\mathfrak{d}}$. Moreover, extending $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ to this richer language runs into difficulties. The axiom schemes **G5** and **G6** essentially work by insisting that the players' beliefs be compatible with rationality or its negation, respectively. In the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$, this amounts to specifying the possible supports for the players' beliefs, which can be written using finite formulas since each $\Sigma_{-i}$ is finite and therefore has only finitely-many subsets. By contrast, in the language with belief modalities $B_i^\alpha$ for every $\alpha \in [0,1]$, the "formula" that says that player $i$'s beliefs are compatible with rationality may be infinitely long.

A still richer language, however, can circumvent these difficulties entirely. Fix a game $\Gamma$ and consider the language of *linear likelihood inequalities* defined by the grammar

$$\varphi ::= p \,|\, \neg\varphi \,|\, \varphi \wedge \psi \,|\, a_1\ell_i(\varphi_1) + \cdots + a_k\ell_i(\varphi_k) \geq b,$$

where $p \in \Phi_\Gamma$, $i \in I$, $k \in \mathbb{N}$, and $a_1, \ldots, a_k, b \in \mathbb{R}$. The *likelihood terms* $\ell_i(\varphi)$ are meant to be read as "the probability of $\varphi$ according to player $i$", and a likelihood formula $a_1\ell_i(\varphi_1)+\cdots+a_k\ell_i(\varphi_k) \geq b$ should be thought of as asserting the corresponding inequality. More precisely, we interpret such formulas in a $\Gamma$-structure $M$ as follows:

$$[\![a_1\ell_i(\varphi_1) + \cdots + a_k\ell_i(\varphi_k) \geq b]\!]_M \quad := \quad \{\omega \in \Omega \,:\, \sum_{j=1}^{k} a_j Pr_i(\omega)([\![\varphi_j]\!]_M) \geq b\}.$$

For example, the formula $\ell_i(\varphi) \geq 1$ says that player $i$ assigns probability at least (and therefore exactly) 1 to $\varphi$, while the formula $(\ell_i(\varphi) \geq \frac{1}{2}) \wedge (\ell_i(\neg\varphi) \geq \frac{1}{2})$ says that player $i$ assigns probability $\frac{1}{2}$ to $\varphi$. See [Halpern 2003] for a thorough discussion of this and related logics; a sound and complete axiomatization is given in [Fagin, Halpern, and Megiddo 1990].

In this language, rationality in the sense of expected utility maximization can be *defined*, thus obviating the need for a separate axiomatization. Indeed, if we let $\sigma_i \succeq \sigma'_i$ be an abbreviation for the formula

$$\sum_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i})\ell_i(play_{-i}(\sigma_{-i})) \; - \sum_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma'_i, \sigma_{-i})\ell_i(play_{-i}(\sigma_{-i})) \geq 0,$$

which says that the expected value according to player $i$ of playing $\sigma_i$ is no less than the expected value of playing $\sigma'_i$, then it is easy to see that

$$RAT_i^{\mathfrak{d}_i^{eu}} \leftrightarrow \bigvee_{\sigma_i \in \Sigma_i} \Big( play_i(\sigma_i) \; \wedge \bigwedge_{\sigma'_i \in \Sigma_i} \sigma_i \succeq \sigma'_i \Big)$$

is valid. Since rationality is expressible in this language, it is axiomatized as well.

# 4 Endogenizing decision rules

In general, we may wish to reason about players who are uncertain about which decision rules their opponents are using. For example, player $i$ might believe that if player $j$ is maximizing her expected utility, then she will play $\sigma_j$, but if she plays $\sigma'_j$, then she might instead be minimizing the worst-case outcome. One way to try to model such uncertainty is to expand the logic so that there is a *set* $\mathfrak{D}_i$ of decision rules associated with each player $i$. Consider the collection of primitive propositions

$$\Phi_\Gamma^{\mathfrak{D}} \coloneqq \Phi_\Gamma \cup \{RAT_i^{\mathfrak{d}_i} \; : \; i \in I, \, \mathfrak{d}_i \in \mathfrak{D}_i\}.$$

Interpret $RAT_i^{\mathfrak{d}_i}$ as before. Then $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{D}})$ is a language for reasoning about the strategies and beliefs of the players $i \in I$ as well as their adherence to the various decision rules $\mathfrak{d}_i \in \mathfrak{D}_i$.

In Appendix A.3, we show how to modify **G5** and **G6** to obtain a sound and complete axiomatization of $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{D}})$ with respect to the class of all finite $\Gamma$-structures. But there is something unsatisfying about using this language to model players' uncertainty about decision rules: the propositions $RAT_i^{\mathfrak{d}_i}$ say that player $i$ is playing a $\mathfrak{d}_i$-best-response, but *not* that player $i$ is actually *using* the rule $\mathfrak{d}_i$ to decide her strategy. To see the difference, consider a player $i$ who is trying to maximize her expected utility (i.e. using $\mathfrak{d}_i^{eu}$), and happens to also play a maximin strategy; contrast this with a scenario in which she is actively seeking to maximize the worst-case outcome (i.e. using $\mathfrak{d}_i^m$), and in so doing happens to play a strategy that maximizes her expected utility. Although player $i$ is following different decision rules in these two cases, the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{D}})$ cannot express this difference; the formula $RAT_i^{\mathfrak{d}_i^{eu}} \wedge RAT_i^{\mathfrak{d}_i^m}$ holds either way. For instance, in Example 2.1, it is not hard to check that

$$\omega_0 \models play_1(\text{Bach}) \to \big(RAT_1^{\mathfrak{d}_1^{eu}} \wedge RAT_1^{\mathfrak{d}_1^m}\big).$$

In sum, the propositions $RAT_i^{\mathfrak{d}_i}$ do not say anything about *how* player $i$ is making her decision, but simply record whether or not the decision she does make is compatible with the rule $\mathfrak{d}_i$. What we want is a different kind of proposition, say $rule_i(\mathfrak{d}_i)$, that says that player $i$ really is *using* the rule $\mathfrak{d}_i$ in deciding her strategy.

Decision rules interpreted in this sense are particularly relevant in a dynamic setting. When an opponent does something unexpected and seemingly irrational, there is the question of how to update your beliefs. One option is to abandon the belief that your opponent is rational, but this is unsatisfying both conceptually and methodologically. An alternative response is to update your beliefs about your opponent's beliefs: what they did actually was rational with respect to their beliefs, you had just misjudged what those beliefs were (see, e.g., [Battigalli and Siniscalchi 2002]). But in some cases, this too is unsatisfying: for example, "continuing" at the second-last stage of the centipede game (see Example 4.3) is only rational for a player who believes his opponent to be irrational. When decision rules are present in the model as objects of belief, however, a third option becomes available: abandon the belief that your opponent is $\mathfrak{d}_i$-rational, but not that they are behaving rationally with respect to some *other* decision rule. Though an analysis of decision rules in extensive-form games is beyond the scope of this paper, the groundwork for such a study can be laid by formalizing them in a static context.

Fix a game $\Gamma$ and a profile $\mathfrak{D} = (\mathfrak{D}_i)_{i \in I}$ of sets of decision rules for each player $i \in I$. Expand the set $\Phi_\Gamma$ of primitive propositions that we considered earlier by taking

$$\Phi_{\Gamma,\mathfrak{D}} := \Phi_\Gamma \cup \{rule_i(\mathfrak{d}_i) \,:\, i \in I,\, \mathfrak{d}_i \in \mathfrak{D}_i\}.$$

In order to interpret the primitive propositions $rule_i(\mathfrak{d}_i)$, we must extend the semantic model so that it associates with each world $\omega$ the decision rule that each player $i$ is using at that world; furthermore, we must constrain the strategies used at each world so that they are compatible with the corresponding decision rules. Formally, a $(\Gamma, \mathfrak{D})$**-structure** is a tuple $M = (\Omega, (Pr_i)_{i \in I}, s, r)$ satisfying (C1) through (C4) as well as the following additional conditions:

(C5) $r : \Omega \to \mathfrak{D}$ satisfies $Pr_i(\omega)(\{\omega' \in \Omega \,:\, r_i(\omega') = r_i(\omega)\}) = 1$;

(C6) $s_i(\omega) \in r_i(\omega)(\mu_{i,\omega})$.

Condition (C5) says that the *decision function* $r$ assigns to each world $\omega$ a profile of decision rules $r(\omega)$—intuitively, $r_i(\omega) \in \mathfrak{D}_i$ is the rule that player $i$ is using at

$\omega$—and moreover, each player is sure of her own decision rule. Condition (C6) requires that, at each world $\omega$, the strategy $s_i(\omega)$ is an $r_i(\omega)$-best response for player $i$; in other words, player $i$ really is following the decision rule $r_i(\omega)$ at $\omega$.

The language $\mathcal{L}_B(\Phi_{\Gamma,\mathfrak{D}})$ can be interpreted in a $(\Gamma, \mathfrak{D})$-structure $M$ as before, with the additional clause

$$[\![rule_i(\mathfrak{d}_i)]\!]_M \quad := \quad \{\omega \in \Omega \, : \, r_i(\omega) = \mathfrak{d}_i\}.$$

The resulting logic can be axiomatized using essentially the same technique as in Section 3 (see Section 4.2). But perhaps more interesting than axiomatizing this logic is the prospect of applying it to the analysis of games.

## 4.1 $\mathfrak{D}$-rationalizability

It is quite natural in certain strategic contexts for players to reason not only about their opponents' strategies and beliefs, but also the decision-making process that they might be using. A decision rule like minimax regret, for instance, can lead to very different behaviour in games like the centipede game or the traveler's dilemma [Halpern and Pass 2012]; as such, it is reasonable in such contexts to wonder, for example, whether an opponent is motivated to maximize utility or to avoid regret.

Recall that strategies that are consistent with common belief of rationality are called *rationalizable*. Common belief of rationality in games—the requirement that every player is rational, believes their opponents are rational, believes their opponents believe their opponents are rational, and so on—is often conceived of as a kind of "minimal" condition for equilibrium. But games like the traveler's dilemma, where the rationalizable strategies are far from optimal and quite distinct from the typical strategies employed by human players [Capra, Goeree, Gomez, and Holt 1999], belie this intuition of minimality. However, by decoupling the meaning of rationality from expected utility maximization, the notion of "rationalizability" can be expanded to other decision rules, thereby providing what is arguably a better epistemic foundation for equilibrium theory.

More precisely, generalizing the traditional epistemic characterization of rationalizability, we define a strategy $\sigma_i$ to be **$\mathfrak{D}$-rationalizable (in $\Gamma$)** just in case there exists a $(\Gamma, \mathfrak{D})$-structure in which $\sigma_i$ is played at some state. Of course, the standard notion arises as the special case where each $\mathfrak{D}_i = \{\mathfrak{d}_i^{eu}\}$. It is easy to see (via a straightforward iterated deletion argument[3]) that when the strategy sets are finite, $\mathfrak{D}$-rationalizable strategies must exist. Moreover, if for each player $i$ we have $\mathfrak{D}_i \subseteq \mathfrak{D}'_i$, then clearly every $(\Gamma, \mathfrak{D})$-structure is also a $(\Gamma, \mathfrak{D}')$-structure; this immediately establishes the following:

[3] Set $\Sigma_i^{(0)} = \Sigma_i$, and inductively define $\Sigma_i^{(k+1)} = \bigcup_{\mathfrak{d}_i \in \mathfrak{D}_i} \{\sigma_i : (\exists \mu \in \Delta(\Sigma_{-i}^{(k)}))(\sigma_i \in \mathfrak{d}_i(\mu))\}$. Then $\Sigma^{(0)}, \Sigma^{(1)}, \ldots$ is a nested decreasing sequence that cannot include the empty set, so it must stabilize if $\Sigma$ is finite. From such a stable $\Sigma^{(K)}$, it is easy to construct a $(\Gamma, \mathfrak{D})$-structure in which each $\sigma_i \in \Sigma_i^{(K)}$ is played at some state.

**Proposition 4.1:** *For each player $i$, let $\mathfrak{D}_i \subseteq \mathfrak{D}'_i$. Then if $\sigma_i$ is $\mathfrak{D}$-rationalizable, it is also $\mathfrak{D}'$-rationalizable.*

We illustrate these concepts with two examples. It will be useful first to formally define the minimax regret decision rule in our setting. Given a game $\Gamma$ and probability measure $\mu$ on $\Sigma_{-i}$, let

$$MR_i(\sigma_i;\mu) := \max\{\max_{\sigma'_i\in\Sigma_i} u_i(\sigma'_i,\sigma_{-i}) - u_i(\sigma_i,\sigma_{-i}) \,:\, \mu(\sigma_{-i})>0\},$$

corresponding to the maximum "regret" player $i$ might feel having played $\sigma_i$, where "regret" is interpreted as the difference between the best-case payoff and the actual payoff (with respect to the strategy profiles $\sigma_{-i}$ that player $i$ considers possible). The minimax regret decision rule $\mathfrak{d}^r_i$ seeks to minimize this value:

$$\mathfrak{d}^r_i(\mu) := \{\sigma_i \in \Sigma_i \,:\, (\forall \sigma'_i \in \Sigma_i)(MR_i(\sigma_i;\mu) \le MR_i(\sigma'_i;\mu))\}.$$

**Example 4.2:** Consider the traveller's dilemma: each of two players must name an amount in $\Sigma_1 = \Sigma_2 = \{2,3,\dots,100\}$, which is the reimbursement they are requesting for luggage that was damaged by their airline. The airline will reimburse them both by the minimum amount requested, with one catch: the person who asks for less receives a \$2 bonus, while the person who asks for more receives a \$2 penalty (if they ask for the same amount, no bonuses or penalties are applied). Thus, payoffs are defined as follows:

$$u_i(\sigma) = \begin{cases} \sigma_i & \text{if } \sigma_i = \sigma_{-i} \\ \sigma_i + 2 & \text{if } \sigma_i < \sigma_{-i} \\ \sigma_{-i} - 2 & \text{if } \sigma_i > \sigma_{-i}. \end{cases}$$

Clearly the best payoff is achieved by undercutting one's fellow traveler by 1 if possible, and otherwise (if the other traveler plays 2), playing 2. It is easy to see that playing 100 is never a $\mathfrak{d}^{eu}_i$-best response; an iterative deletion argument then shows that the only rationalizable strategy is to play 2. By contrast, when each $\mathfrak{D}_i = \{\mathfrak{d}^r_i\}$, playing 100 is $\mathfrak{D}$-rationalizable. To prove this, by definition, it suffices to exhibit a $(\Gamma,\mathfrak{D})$-structure in which 100 is played at some state. Consider the structure presented in Figure 1:

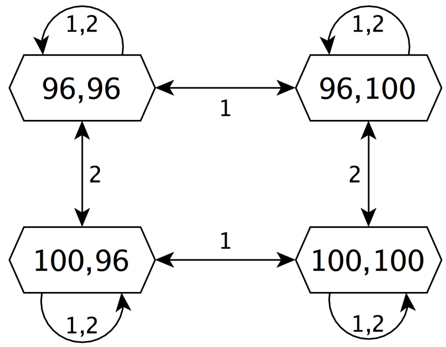


Figure 1: A $(\Gamma, \mathfrak{D})$-structure for the traveler's dilemma

Each of the four states of this structure is labeled with the strategy profile being played at that state, while the edges labeled $i$ represent which states player $i$ considers possible (i.e., assigns positive probability to) from which other states (numerical probabilities are irrelevant for this analysis and so are suppressed). We must show that each player is playing according to minimax regret. Take player 1's perspective (the argument for player 2 is analogous); observe first that she considers 96 and 100 to be the only possible plays her opponent might make. Given this, player 1's maximum regret on playing $\sigma_1 > 96$ must be at least 3, since $u_1(\sigma_1, 96) = 94$, whereas $u_1(95, 96) = 97$. Similarly, player 1's maximum regret on playing $\sigma_1 \leq 96$ must be at least 3, since $u_1(\sigma_1, 100) = \sigma_1 + 2 \leq 98$, whereas $u_1(99, 100) = 101$. Moreover, it is straightforward to check that player 1's maximum regret on playing either 96 or 100 is exactly 3. It follows that each of 96 and 100 constitutes a $\mathfrak{d}_1^r$-best response. ∎

**Example 4.3:** Consider the normal-form version of the centipede game [Rosenthal 1982] depicted in Figure 2: each player must choose whether to quit at some stage or play to the end. Let $\Sigma_1 = \{Q_1, Q_3, Q_*\}$ and $\Sigma_2 = \{Q_2, Q_4, Q_*\}$, where $Q_k$ stands for quitting at stage $k$ and $Q_*$ stands for playing to the end. Payoffs are determined by the the minimal stage that some player quit at, as shown in Figure 2. For instance: $u(Q_1, Q_2) = u(Q_1, Q_4) = (1, 0)$, since in either case player 1 quits at the first stage (making player 2's choice irrelevant); on the other hand, $u(Q_*, Q_4) = (2, 8)$, since player 1 never quits and player 2 quits at the fourth stage.

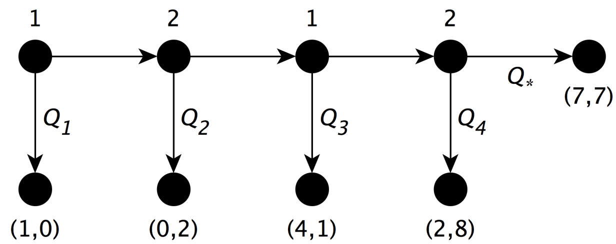


Figure 2: The centipede game

It is well known that all pure strategies in this game are rationalizable; however, the only strategy that is rationalizable for player 1 under *conservative* beliefs—namely, beliefs that ascribe positive probability to the actual state—is $Q_1$, quitting immediately [Halpern and Pass 2013]. By contrast, we now show that when each $\mathfrak{D}_i = \{\mathfrak{d}_i^{eu}, \mathfrak{d}_i^r\}$, all strategies are $\mathfrak{D}$-rationalizable even under conservative beliefs.

It is easy to see that $Q_1$ and $Q_2$ are $\mathfrak{D}$-rationalizable with conservative beliefs: indeed, the structure with exactly one state $\omega$ where $s(\omega) = (Q_1, Q_2)$ is a $(\Gamma, \mathfrak{D})$-structure because each player $i$ must be sure of the actual state, and is easily seen to be playing a $\mathfrak{d}_i^{eu}$-best response to this belief. This observation is an instance of Proposition 4.1 applied to the fact that $Q_1$ and $Q_2$ are rationalizable (with conservative beliefs) in the traditional sense of rationalizability.

To show that the remaining strategies are $\mathfrak{D}$-rationalizable under conservative beliefs, it suffices to construct a $(\Gamma, \mathfrak{D})$-structure at which each of these strategies is played and all beliefs are conservative. Consider the structure presented in Figure 3:

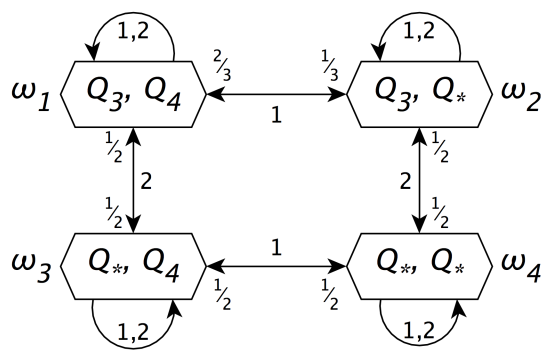


Figure 3: A $(\Gamma, \mathfrak{D})$-structure for the centipede game

As in Figure 1, each of the four states $\omega_1$, $\omega_2$, $\omega_3$, and $\omega_4$ of this structure is labeled with the strategy profile being played at that state, while the edges labeled $i$ represent which states player $i$ considers possible (i.e., assigns positive probability to) from which other states. In addition, the fractions adjacent to the arrowheads specify the numeric probability of each state; for example, the fractions $\frac{2}{3}$ and $\frac{1}{3}$ indicate that $Pr_1(\omega_1)(\{\omega_1\}) = Pr_1(\omega_2)(\{\omega_1\}) = \frac{2}{3}$ and $Pr_1(\omega_1)(\{\omega_2\}) = Pr_1(\omega_2)(\{\omega_2\}) = \frac{1}{3}$, respectively.

We must show that at each state, each player $i$ is playing according to either $\mathfrak{d}_i^{eu}$ or $\mathfrak{d}_i^r$. First we show that player 1 is maximizing expected utility in states $\omega_1$ and $\omega_2$. In these states player 1 quits at stage 3, which yields an expected utility of 4 (since player 1 is sure that player 2 will not quit beforehand). Playing $Q_1$ has an expected utility of 1, so $Q_1$ is strictly dominated by $Q_3$. Finally, playing $Q_*$ results in a $\frac{2}{3}$ chance of a utility of 2, and a $\frac{1}{3}$ chance of a utility of 7, for an expected utility of $\frac{11}{3}$, so $Q_*$ is dominated by $Q_3$.

Next we show that player 1 is minimizing her maximum regret in states $\omega_3$ and $\omega_4$. In these states, player 1 plays $Q_*$ and believes that player 2 will play either $Q_4$ or $Q_*$. In the first case, player 1's payoff is 2, but it could have been as high as 4 had she played $Q_3$; in the second case, her payoff is 7, but it could have been as high as 8 had she played $Q_4$. Thus her maximum regret is 2. How does this compare to her maximum regret on choosing an alternative strategy? If she plays $Q_3$, her maximum regret is 3, which arises when player 2 plays $Q_*$: in this case, her payoff is 4 but it could have been 7 had she played $Q_*$ instead. Even worse, her maximum regret on playing $Q_1$ is 7 (arising as above when player 2 plays $Q_*$). This shows that $Q_*$ is indeed a $\mathfrak{d}_1^r$-best response in states $\omega_3$ and $\omega_4$.

Similar arguments show that $Q_4$ is a $\mathfrak{d}_2^{eu}$-best response in states $\omega_1$ and $\omega_3$, and $Q_*$ is a $\mathfrak{d}_2^r$-best response in states $\omega_2$ and $\omega_4$. Thus, we can set

$$\begin{array}{rcl} r(\omega_1) & = & (\mathfrak{d}_1^{eu}, \mathfrak{d}_2^{eu}) \\ r(\omega_2) & = & (\mathfrak{d}_1^{eu}, \mathfrak{d}_2^{r}) \\ r(\omega_3) & = & (\mathfrak{d}_1^{r}, \mathfrak{d}_2^{eu}) \\ r(\omega_4) & = & (\mathfrak{d}_1^{r}, \mathfrak{d}_2^{r}) \end{array}$$

to make Figure 3 into a $(\Gamma, \mathfrak{D})$-structure, which proves that each of the strategies played therein is $\mathfrak{D}$-rationalizability under conservative beliefs. ∎

## 4.2 Axiomatization

Consider the following axiom schemes:

**P1.** $\bigvee_{\sigma_i \in \Sigma_i} play_i(\sigma_i)$

**P2.** $\neg(play_i(\sigma_i) \wedge play_i(\sigma_i'))$, for $\sigma_i \neq \sigma_i'$

**P3.** $play_i(\sigma_i) \leftrightarrow B_i\, play_i(\sigma_i)$

**P4.** $\bigvee_{\mathfrak{d}_i \in \mathfrak{D}_i} rule_i(\mathfrak{d}_i)$

**P5.** $\neg(rule_i(\mathfrak{d}_i) \wedge rule_i(\mathfrak{d}'_i))$, for $\mathfrak{d}_i \neq \mathfrak{d}'_i$

**P6.** $rule_i(\mathfrak{d}_i) \leftrightarrow B_i\, rule_i(\mathfrak{d}_i)$

**P7.** $(play_i(\sigma_i) \wedge rule_i(\mathfrak{d}_i)) \rightarrow \bigvee_{S \in \mathcal{S}_i^+(\sigma_i)} \delta_{i,S}$.

Here, as before, $\mathcal{S}_i^+(\sigma_i)$ denotes the collection of all $S \subseteq \Sigma_{-i}$ such that there exists a probability measure $\mu$ with $supp(\mu) = S$ and such that $\sigma_i \in \mathfrak{d}_i(\mu)$. Note that there is no need for a symmetric axiom involving $\mathcal{S}_i^-(\sigma_i)$ for this logic, because the formula $\neg rule_i(\mathfrak{d}_i)$, unlike $\neg RAT_i$, does not say that $\sigma_i$ is *incompatible* with player $i$'s beliefs and the decision rule $\mathfrak{d}_i$; it simply says that player $i$ did not *use* the rule $\mathfrak{d}_i$ to help choose her strategy $\sigma_i$ (though she may, coincidentally, have beliefs with respect to which $\sigma_i$ is a $\mathfrak{d}_i$-best response).

Let $\mathsf{GL}_{\Gamma,\mathfrak{D}}$ be the axiom system that results from adding **P1**–**P7** to the $\mathsf{KD45}$ axioms and rules of inference of belief logic. Then we have the following result, the proof of which proceeds analogously to that of Theorem A.1 given in Appendix A.1.

**Theorem 4.4:** $\mathsf{GL}_{\Gamma,\mathfrak{D}}$ *is a sound and complete axiomatization of the language* $\mathcal{L}_B(\Phi_{\Gamma,\mathfrak{D}})$ *with respect to the class of all finite* $(\Gamma, \mathfrak{D})$*-structures.*

# 5 Discussion

Almost all solution concepts in game theory are grounded in the idea of rationality and best responding. Thus, one natural application of a logic of rationality is to the analysis of solution concepts. But doing so raises a number of research issues.

One subtlety involves the use of mixed strategies. The language $\mathcal{L}_B(\Phi_\Gamma)$ has formulas that represent pure strategy choices, but not mixed strategies. In the context of Nash equilibrium, this difference turns out to be (at least formally) innocuous: one can view a mixed strategy $\ell_i \in \Delta(\Sigma_i)$ either as a conscious randomization on the part of player $i$, or as the common conjecture of the players $j \neq i$ about what pure strategy $i$ will choose—either way, the set of mixed strategy Nash equilibria stays the same. However, this insensitivity is, in part, dependent on that fact that rationality in the sense of expected utility maximization "plays well" with mixing: $\ell_i \in \Delta(\Sigma_i)$ maximizes player $i$'s expected utility (with respect to some fixed beliefs) if and only if every pure strategy $\sigma_i$ in the support of $\ell_i$ maximizes expected utility. But this correspondence breaks down when "expected utility maximization" is replaced with the generalized notion of rationality presented in Section 2.3: in the context of an arbitrary decision rule $\mathfrak{d}_i : \Delta(\Sigma_{-i}) \to 2^{\Sigma_i} \setminus \emptyset$, there is no principled way to extend the notion of "best response" from $\Sigma_i$ to $\Delta(\Sigma_i)$. This suggests that further research into the interaction between pure and mixed strategies under general decision rules may be fruitful.

A second issue involves reconsidering what happens to various solution concepts when we replace maximizing expected utility by another decision rule. Consider, for example, Nash equilibrium. In principle, it makes sense to consider "$\mathfrak{d}$-Nash equilibria", defined by replacing $\mathfrak{d}^{eu}$ with an arbitrary profile of decision rules $\mathfrak{d}$ in the definition of Nash equilibrium. It is certainly too much to hope that Nash's famous existence theorem applies in full force to this wider concept; however, properties of $\mathfrak{d}$ that suffice to guarantee the existence of equilibria are of interest, and potentially admit a logical characterization. Such questions are the subject of ongoing research.

Yet another issue involves understanding the implications for computability of using various decision rules. In Section 3, we observed that the axiom systems $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ are finite extensions of the $\mathsf{KD45}$ system and thus trivially decidable. Thus, we can, for example, compute whether a formula is a logical consequence of rationality in any given axiom system $\mathsf{GL}_\Gamma^{\mathfrak{d}}$. But there is arguably a more interesting question as far as decidability goes. Up to now we have considered decision rules as functions defined with respect to some fixed game. But rules like expected utility maximization, maximin, or minimax regret can be applied in all games in a uniform way. To capture this, define a **decision paradigm** to be a function that maps each game $\Gamma$ to a decision rule in $\Gamma$. Suppose that we are given decision paradigms $\mathcal{D}_i$ for each player associating with each game $\Gamma$ a decision rule $\mathcal{D}_i(\Gamma)$ for that player in $\Gamma$. We might want to know, given the profile $\mathcal{D} = (\mathcal{D}_i)_{i \in I}$, whether the mapping

$$\Gamma \mapsto \mathsf{GL}_\Gamma^{\mathcal{D}(\Gamma)}$$

is decidable; in other words, given as input a game $\Gamma$, can we effectively determine whether a formula belongs to the axiom system $\mathsf{GL}_\Gamma^{\mathcal{D}(\Gamma)}$? For each game $\Gamma$, this requires determining membership in the sets $\mathcal{S}_i^+(\sigma_i)$ and $\mathcal{S}_i^-(\sigma_i)$, which are defined by existential quantification over simplices $\Delta(\Sigma_{-i})$, subject to constraints based on the decision rules $\mathcal{D}_i(\Gamma)$. In the case of familiar decision paradigms like maximin or expected utility maximization, computing the sets $\mathcal{S}_i^+(\sigma_i)$ and $\mathcal{S}_i^-(\sigma_i)$ boils down to solving systems of linear inequalities. In general, however, we must impose certain computability requirements on the decision paradigms in order to be able to decide whether a formula is an instance of an axiom. To take an extreme example: we could define $\mathcal{D}_i(\Gamma)$ depending on whether the number of players in $\Gamma$ lies in the halting set.

This kind of example suggests that we want to be more restrictive in the form that $\mathcal{D}_i$ can take. In particular, we may be able to get more traction on this problem if we restrict attention to decision paradigms that can be expressed in some limited language. All the standard decision rules—maximizing expected utility, minimax regret, maximin, and so on—are of the form "choose strategy

$\sigma_i$ only if $\gamma$", where $\gamma$ is a collection of constraints expressible in some simple language involving quantification over strategies, linear inequalities, etc. We believe that by identifying appropriate languages and limiting the constraints that can be used to define decision paradigms to those expressible in these languages, we may well be able to establish general decidability results that apply to decision paradigms rather than merely decision rules.

Finally, through the introduction $\mathfrak{D}$-rationality in this paper, we hope to initiate a broader research program investigating the representational power of endogenous decision rules. Belief update in strategic scenarios is widely recognized as a foundational issue in modern game theory; the additional structure of decision rules associated to each state allows a player to learn not just about her opponents' strategy choices and beliefs, but about the mechanism by which they make decisions under uncertainty. As we have already suggested, this kind of belief update is particularly relevant in a dynamic setting. Thus, a natural extension of the present work would be to formulate an extensive-form version of $\mathfrak{D}$-rationalizability and investigate its relationship with standard extensive-form solution concepts and methods of belief update.

# A Proofs

## A.1 Axiomatizing $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$

**Theorem A.1:** $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ *is a sound axiomatization of the language* $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ *with respect to the class of all finite* $\Gamma$*-structures.*

**Proof:** Soundness of the axioms and rules of $\mathsf{KD45}$ can be proved as usual. It therefore suffices to show that **G1**–**G6** are valid in all finite $\Gamma$-structures.

Fix a finite $\Gamma$-structure $M = (\Omega, (Pr_i)_{i \in I}, s)$. Soundness of **G1** and **G2** is an immediate consequence of the fact that $s$ is a (total) function. Soundness of **G3** is a straightforward consequence of condition (C4), while soundness of **G4** follows easily from the combination of conditions (C3) and (C4).

Now suppose that $\omega \models \mathit{play}_i(\sigma_i) \wedge \mathit{RAT}_i$; then $\sigma_i \in \mathfrak{d}_i(\mu_{i,\omega})$. Set $S := \mathit{supp}(\mu_{i,\omega})$, and observe that $S \in \mathcal{S}_i^+(\sigma_i)$. For each $\sigma_{-i} \in S$, it is easy to see that $\omega \models \widehat{B}_i\, \mathit{play}_{-i}(\sigma_{-i})$; therefore, we must have

$$\omega \models \bigwedge_{\sigma_{-i} \in S} \widehat{B}_i\, \mathit{play}_{-i}(\sigma_{-i}).$$

Similarly, for each $\sigma_{-i} \notin S$, we have $\omega \models \neg\widehat{B}_i\, \mathit{play}_{-i}(\sigma_{-i})$, so

$$\omega \models \bigwedge_{\sigma_{-i} \notin S} \neg\widehat{B}_i\, \mathit{play}_{-i}(\sigma_{-i}).$$

In other words, $\omega \models \delta_{i,S}$; this establishes the soundness of **G5**.

Finally, suppose that $\omega \models play_i(\sigma_i) \wedge \neg RAT_i$, which implies that $\sigma_i \notin \mathfrak{d}_i(\mu_{i,\omega})$. Set $S \coloneqq supp(\mu_{i,\omega})$; then $S \in \mathcal{S}_i^-(\sigma_i)$ and, as above, we have $\omega \models \delta_{i,S}$, which establishes soundness of **G6**. ∎

We prove completeness by what is essentially the *canonical model method*, a standard method for proving completeness of modal logics (see, e.g., [Fagin, Halpern, Moses, and Vardi 1995] or any standard text on modal logic). Of course, the full canonical model is not finite (for $n > 1$), so we modify the construction by restricting attention to finite sub-languages. More precisely, given a formula $\varphi \in \mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$, we identify a finite sub-language of $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ such that the corresponding canonical model refutes $\varphi$ just in case $\mathsf{GL}_\Gamma^{\mathfrak{d}} \not\vdash \varphi$. This technique is sometimes called *filtration*.

Fix a formula $\varphi \in \mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$. Let $Sub_\Gamma(\varphi)$ denote the collection of all subformulas of $\varphi$ together with all subformulas of instances of the axiom schemes **G1** through **G6**. Define

$$Sub_\Gamma^+(\varphi) \coloneqq Sub_\Gamma(\varphi) \cup \{\neg\psi \,:\, \psi \in Sub_\Gamma(\varphi)\}.$$

Note that there are only finitely many instances of **G1** through **G6**, and therefore $Sub_\Gamma^+(\varphi)$ is finite.

Let $\Omega^\varphi$ be the collection of all maximal, consistent (with respect to $\mathsf{GL}_\Gamma^{\mathfrak{d}}$) subsets of $Sub_\Gamma^+(\varphi)$. Clearly $\Omega^\varphi$ is a finite set. Given $X \subseteq \mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$, set

$$X^{B_i} \coloneqq \{\psi \,:\, B_i\psi \in X\};$$

and, for $F \in \Omega^\varphi$, define

$$\mathrm{Bel}_i(F) \coloneqq \{G \in \Omega^\varphi \,:\, G \supseteq F^{B_i} \text{ and } G^{B_i} = F^{B_i}\}.$$

For each $i \in I$ and each $F \in \Omega^\varphi$, we will define a probability measure $Pr_i^\varphi(F)$ on $\Omega^\varphi$ such that the support of this measure is precisely $\mathrm{Bel}_i(F)$. Loosely speaking, $\mathrm{Bel}_i(F)$ is the set of all $G \in \Omega^\varphi$ that are compatible with the beliefs of player $i$ in $F$. More precisely, $G \in \mathrm{Bel}_i(F)$ if and only if:

(a) $B_i\psi \in F$ implies $\psi \in G$, and

(b) $B_i\psi \in F$ iff $B_i\psi \in G$.

Condition (a) just says that everything player $i$ believes in $F$ is true in $G$; condition (b) says that player $i$'s beliefs in $F$ are the same as her beliefs in $G$, which is reasonable in light of the fact that we are working in a system with positive and negative introspection. (In the full canonical model, (b) follows from (a); here we must impose this condition explicitly because of the way the language has been restricted.)

Since our aim is to define probability measures with the sets $\mathrm{Bel}_i(F)$ as their supports, we must show that these sets are never empty.

**Lemma A.2:** *For each $i \in I$ and $F \in \Omega^\varphi$,* $\mathrm{Bel}_i(F) \neq \emptyset$.

**Proof:** Given $F \in \Omega^\varphi$, set

$$\Lambda \coloneqq \{\psi \,:\, B_i\psi \in F\} \cup \{B_i\psi \,:\, B_i\psi \in F\} \cup \{\neg B_i\psi \,:\, \neg B_i\psi \in F\}.$$

It is easy to see that $\Lambda \subset Sub^+_\Gamma(\varphi)$. In addition, we show that $\Lambda$ is consistent. For suppose not; then

$$\begin{aligned} \mathsf{GL}^\mathfrak{d}_\Gamma \vdash \neg \bigwedge_{\xi\in\Lambda} \xi \quad &\Longrightarrow \quad \mathsf{GL}^\mathfrak{d}_\Gamma \vdash B_i \neg \bigwedge_{\xi\in\Lambda} \xi \\ &\Longrightarrow \quad \mathsf{GL}^\mathfrak{d}_\Gamma \vdash \neg B_i \bigwedge_{\xi\in\Lambda} \xi \\ &\Longrightarrow \quad \mathsf{GL}^\mathfrak{d}_\Gamma \vdash \neg \bigwedge_{\xi\in\Lambda} B_i\xi, \end{aligned}$$

which is a contradiction, since each formula $B_i\xi$ with $\xi \in \Lambda$ is logically equivalent to a formula in $F$, and $F$ is consistent.

From this we can conclude that there exists a $G \in \Omega^\varphi$ such that $G \supseteq \Lambda$. It follows immediately that $G \supseteq F^{B_i}$ and that $G^{B_i} \supseteq F^{B_i}$. Moreover, if $\psi \in G^{B_i}$, then $B_i\psi \in G$ and so certainly $\neg B_i\psi \notin G$, from which it follows that $\neg B_i\psi \notin \Lambda$ and thus $\neg B_i\psi \notin F$. Maximality of $F$ then guarantees that $B_i\psi \in F$, whence $\psi \in F^{B_i}$, and so $G^{B_i} \subseteq F^{B_i}$. This establishes that $G \in \mathrm{Bel}_i(F)$, as desired. ∎

In the classical canonical model construction, it is sufficient to define $Pr_i^\varphi(F)$ to be the uniform distribution on $\mathrm{Bel}_i(F)$. In the present context, however, we need to be more careful, since $Pr_i^\varphi(F)$ is used not only to interpret the belief modalities $B_i$, but also the primitive propositions $RAT_i$. In essence, we must define $Pr_i^\varphi(F)$ in a manner that agrees with whether or not player $i$ is best responding to her beliefs at $F$; not surprisingly, this is precisely where the axiom schemes **G5** and **G6** come into play. At the same time, we have to define $Pr_i^\varphi$ on $\Omega^\varphi$ in a systematic way so as to preserve the introspection condition (C3). What follows is a formalization of this basic recipe, for which several more lemmas and definitions are needed.

**Lemma A.3:** *Let $F \in \Omega^\varphi$. If $\widehat{B}_i\psi \in F$, then there exists a $G \in \mathrm{Bel}_i(F)$ with $\psi \in G$; if $\neg\widehat{B}_i\psi \in F$, then for all $G \in \mathrm{Bel}_i(F)$ we have $\psi \notin G$.*

**Proof:** First suppose that $\widehat{B}_i\psi \in F$, and set

$$\Lambda \coloneqq \{\chi \,:\, B_i\chi \in F\} \cup \{B_i\chi \,:\, B_i\chi \in F\} \cup \{\neg B_i\chi \,:\, \neg B_i\chi \in F\}.$$

Assume for contradiction that $\Lambda \cup \{\psi\}$ is inconsistent. We then have

$$\mathsf{GL}^\mathfrak{d}_\Gamma \vdash \bigwedge_{\xi\in\Lambda} \xi \to \neg\psi,$$

from which it follows that

$$\mathsf{GL}_{\Gamma}^{\mathfrak{d}} \vdash \bigwedge_{\xi \in \Lambda} B_i \xi \to B_i \neg \psi. \tag{1}$$

As observed in Lemma A.2, each $B_i \xi$ is equivalent to a formula in $F$, and therefore (1) implies that $B_i \neg \psi \in F$, contradicting our assumption that $\widehat{B}_i \psi \in F$. Thus $\Lambda \cup \{\psi\}$ is consistent, and so can be extended to some $G \in \Omega^{\varphi}$; moreover, as we saw in Lemma A.2, $G \in \mathrm{Bel}_i(F)$. This proves the first statement of the Lemma. The second statement follows immediately from the definition of $\mathrm{Bel}_i(F)$: if $\neg \widehat{B}_i \psi \in F$, then also $B_i \neg \psi \in F$, and so for all $G \in \mathrm{Bel}_i(F)$ we have $\neg \psi \in G$, whence $\psi \notin G$. ∎

For each $\sigma_{-i} \in \Sigma_{-i}$, define

$$\mathrm{Bel}_i(F; \sigma_{-i}) \coloneqq \{G \in \mathrm{Bel}_i(F) \, : \, \mathit{play}_{-i}(\sigma_{-i}) \in G\}.$$

Given $F \in \Omega^{\varphi}$, it is easy to see, using **G1** and **G2**, that there is a unique $\sigma_i \in \Sigma_i$ with $\mathit{play}_i(\sigma_i) \in F$. If, in addition, $RAT_i \in F$, then by **G5** we know that for some $S \in \mathcal{S}_i^+(\sigma_i)$, $\delta_{i,S} \in F$ (or, in the case where $\mathcal{S}_i^+(\sigma_i) = \emptyset$, we know that no such $F$ exists). Otherwise, if $RAT_i \notin F$, then by **G6** we know that for some $S \in \mathcal{S}_i^-(\sigma_i)$, $\delta_{i,S} \in F$ (or again, in the case where $\mathcal{S}_i^-(\sigma_i) = \emptyset$, that no such $F$ exists). Thus, for each $i \in I$, there is a unique set $S_i(F) \subseteq \Sigma_{-i}$ such that $\delta_{i,S_i(F)} \in F$, and moreover, $S_i(F) \in \mathcal{S}_i^+(\sigma_i)$ if $RAT_i \in F$, and $S_i(F) \in \mathcal{S}_i^-(\sigma_i)$ if $RAT_i \notin F$.

**Lemma A.4:** *The collection* $\{\mathrm{Bel}_i(F; \sigma_{-i}) \, : \, \sigma_{-i} \in \Sigma_{-i}\}$ *partitions* $\mathrm{Bel}_i(F)$*; moreover,* $\mathrm{Bel}_i(F; \sigma_{-i}) \neq \emptyset$ *if and only if* $\sigma_{-i} \in S_i(F)$*.*

**Proof:** The first statement is a straightforward consequence of **G1** and **G2**, while the second is an immediate corollary of Lemma A.3 together with the fact that

$$\delta_{i,S_i(F)} \equiv \bigwedge_{\sigma_{-i} \in S_i(F)} \widehat{B}_i \, \mathit{play}_{-i}(\sigma_{-i}) \; \wedge \bigwedge_{\sigma_{-i} \notin S_i(F)} \neg \widehat{B}_i \, \mathit{play}_{-i}(\sigma_{-i}) \in F. \;\blacksquare$$

For each $\sigma_i \in \Sigma_i$ and $S \in \mathcal{S}_i^+(\sigma_i)$, let $\mu^+_{\sigma_i,S}$ be a fixed probability measure witnessing the fact that $S \in \mathcal{S}_i^+(\sigma_i)$; that is, $\mathit{supp}(\mu^+_{\sigma_i,S}) = S$ and $\sigma_i \in \mathfrak{d}_i(\mu^+_{\sigma_i,S})$. Likewise, for each $\sigma_i \in \Sigma_i$ and $S \in \mathcal{S}_i^-(\sigma_i)$, let $\mu^-_{\sigma_i,S}$ be a fixed probability measure witnessing the fact that $S \in \mathcal{S}_i^-(\sigma_i)$.

Let $F \in \Omega^{\varphi}$, and suppose that $\mathit{play}_i(\sigma_i) \in F$. In light of Lemma A.4, we can define $Pr_i^{\varphi}(F)$ to be the unique probability measure on $\mathrm{Bel}_i(F)$ such that, for all $\sigma_{-i} \in \Sigma_{-i}$,

$$Pr_i^{\varphi}(F)(\mathrm{Bel}_i(F; \sigma_{-i})) = \begin{cases} \mu^+_{\sigma_i,S_i(F)}(\sigma_{-i}) & \text{if } RAT_i \in F \\ \mu^-_{\sigma_i,S_i(F)}(\sigma_{-i}) & \text{if } RAT_i \notin F, \end{cases}$$

and which is uniform within each (nonempty) set $\mathrm{Bel}_i(F;\sigma_{-i})$.

**Proposition A.5:** *$Pr_i^\varphi$ satisfies the following:*

*(a) $Pr_i^\varphi(F)(G) > 0$ iff $G \in \mathrm{Bel}_i(F)$, and*

*(b) $G^{B_i} = F^{B_i}$ implies $Pr_i^\varphi(G) = Pr_i^\varphi(F)$.*

**Proof:**

(a) The forward implication is immediate from the definition. For the reverse implication, suppose that $G \in \mathrm{Bel}_i(F)$; then, by Lemma A.4, we know that $G \in \mathrm{Bel}_i(F;\sigma_{-i})$ for some $\sigma_{-i} \in S_i(F)$, from which it follows that $Pr_i^\varphi(F)(G) > 0$ by definition.

(b) If $G^{B_i} = F^{B_i}$ then $\mathrm{Bel}_i(G) = \mathrm{Bel}_i(F)$. Moreover, axioms **G3** and **G4** guarantee that $play_i(\sigma_i) \in F$ if and only if $play_i(\sigma_i) \in G$, and likewise $RAT_i \in F$ if and only if $RAT_i \in G$. Finally, it is not difficult to see that $S_i(F)$ is completely determined by $\mathrm{Bel}_i(F)$, so $S_i(F) = S_i(G)$. Therefore, by definition of $Pr_i^\varphi$, we can deduce that $Pr_i^\varphi(G) = Pr_i^\varphi(F)$. ∎

Finally, we define a strategy function $s^\varphi : \Omega^\varphi \to \Sigma$ by assigning to each $F \in \Omega^\varphi$ the unique strategy profile $\sigma \in \Sigma$ such that $play(\sigma) \in F$.

**Lemma A.6:** *The tuple $M^\varphi := (\Omega^\varphi, (Pr_i^\varphi)_{i\in I}, s^\varphi)$ is a $\Gamma$-structure.*

**Proof:** Conditions (C1) and (C2) have already been established. By Lemma A.5(b), in order to see that (C3) holds it suffices to observe that $Pr_i^\varphi(F)(G) > 0$ implies that $G^{B_i} = F^{B_i}$, which follows from Lemma A.5(a). This same observation also establishes (C4), since by **G3** we know that $G^{B_i} = F^{B_i}$ implies $s_i^\varphi(G) = s_i^\varphi(F)$. ∎

**Lemma A.7:** *For all formulas $\psi$, for all $F \in \Omega^\varphi$, if $\psi \in Sub^+_\Gamma(\varphi)$ then $F \in [\![\psi]\!]_{M^\varphi}$ if and only if $\psi \in F$.*

**Proof:** The proof proceeds by induction on the structure of $\psi$. We prove here the base cases corresponding to the primitive propositions; the inductive steps can be proved in the standard way (see, e.g., [Fagin, Halpern, Moses, and Vardi 1995]).

First consider the primitive proposition $play_i(\sigma_i)$. We have

$$\begin{aligned} F \in [\![play_i(\sigma_i)]\!]_{M^\varphi} \quad &\text{iff} \quad s_i^\varphi(F) = \sigma_i \\ &\text{iff} \quad play_i(\sigma_i) \in F, \end{aligned}$$

as a direct consequence of the definition of $s^\varphi$. Next consider the primitive proposition $RAT_i$; we have

$$\begin{array}{rll} F \in [\![RAT_i]\!]_{M^\varphi} & \text{iff} & s_i^\varphi(F) \in \mathfrak{d}_i(\mu_{i,F}) \\ & \text{iff} & RAT_i \in F, \end{array}$$

the last equivalence being a consequence of the definition of $Pr_i^\varphi$, which ensures that $s_i^\varphi(F)$ is a best response to (the pushforward of) $Pr_i^\varphi(F)$ precisely when $RAT_i \in F$. This completes the proof. ∎

**Theorem A.8:** $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ *is a complete axiomatization of the language* $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ *with respect to the class of all finite* $\Gamma$*-structures.*

**Proof:** Suppose that $\mathsf{GL}_\Gamma^{\mathfrak{d}} \not\vdash \varphi$. Then $\{\neg\varphi\}$ is consistent and so can be extended to a maximal consistent set $F \in \Omega^\varphi$. By Lemma A.7, this implies that $F \notin [\![\varphi]\!]_{M^\varphi}$ and so, in particular, that $M^\varphi \not\models \varphi$, as desired. ∎

## A.2 Belief as lower probability

**Theorem A.9:** $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ *is a sound and complete axiomatization of the language* $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ *with respect to the class of all finite lower* $\Gamma$*-structures.*

**Proof:** The proof given in Appendix A.1 works here as well, modulo the obvious minor alterations in keeping with the generalized definitions given in Section 3.1. In particular, for each $\sigma_i \in \Sigma_i$ and $S \in \mathcal{S}_i^+(\sigma_i)$, we define $\mathcal{P}_{\sigma_i,S}^+$ to be a fixed *set* of probability measures such that $supp(\mathcal{P}_{\sigma_i,S}^+) = S$ and $\sigma_i \in \mathfrak{d}_i(\mathcal{P}_{\sigma_i,S}^+)$; likewise, for each $\sigma_i \in \Sigma_i$ and $S \in \mathcal{S}_i^-(\sigma_i)$, define $\mathcal{P}_{\sigma_i,S}^-$ to be a fixed set of probability measures witnessing the fact that $S \in \mathcal{S}_i^-(\sigma_i)$. Then, given $F \in \Omega^\varphi$ with $play_i(\sigma_i) \in F$, define $\mathcal{PR}_i^\varphi(F)$ as follows: for each $\mu \in \mathcal{P}_{\sigma_i,S}^+ \cup \mathcal{P}_{\sigma_i,S}^-$, let $\tilde{\mu}$ be the unique probability measure on $\mathrm{Bel}_i(F)$ such that, for all $\sigma_{-i} \in \Sigma_{-i}$,

$$\tilde{\mu}(\mathrm{Bel}_i(F;\sigma_{-i})) = \mu(\sigma_{-i}),$$

and which is uniform on each set $\mathrm{Bel}_i(F;\sigma_{-i})$; then set

$$\mathcal{PR}_i^\varphi(F) := \begin{cases} \{\tilde{\mu} \,:\, \mu \in \mathcal{P}_{\sigma_i,S}^+\} & \text{if } RAT_i \in F \\ \{\tilde{\mu} \,:\, \mu \in \mathcal{P}_{\sigma_i,S}^-\} & \text{if } RAT_i \notin F. \end{cases}$$ ∎

## A.3 Axiomatizing $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{D}})$

To obtain a complete axiomatization of this language, it is not sufficient to simply let **G4–G6** range over all decision rules $\mathfrak{d}_i \in \mathfrak{D}_i$ for each player $i$; in general, this system is sound but not complete. Roughly speaking, this is because it is possible for both $play_i(\sigma_i) \wedge RAT_i^{\mathfrak{d}_i}$ and $play_i(\sigma_i) \wedge RAT_i^{\mathfrak{d}'_i}$ to be consistent with $\delta_{i,S}$ for some $S \subseteq \Sigma_{-i}$, yet no measure $\mu$ with $supp(\mu) = S$ is such that $\sigma_i \in \mathfrak{d}_i(\mu) \cap \mathfrak{d}'_i(\mu)$. In this case, the formula

$$play_i(\sigma_i) \wedge RAT_i^{\mathfrak{d}_i} \wedge RAT_i^{\mathfrak{d}'_i} \wedge \delta_{i,S}$$

is not satisfiable, but there is no way to prove its negation from the axioms. However, provided that each set $\mathfrak{D}_i$ is finite, we can deal with this problem by replacing **G5** and **G6** with the following collection of axioms for each player $i$, each strategy $\sigma_i \in \Sigma_i$, and every subset $D \subseteq \mathfrak{D}_i$:

$$\Big(play_i(\sigma_i) \wedge \bigwedge_{\mathfrak{d}_i \in D} RAT_i^{\mathfrak{d}_i} \wedge \bigwedge_{\mathfrak{d}_i \notin D} \neg RAT_i^{\mathfrak{d}_i}\Big) \rightarrow \bigvee_{S \in \mathcal{S}_i^D(\sigma_i)} \delta_{i,S},$$

where $\mathcal{S}_i^D(\sigma_i)$ is the collection of all $S \subseteq \Sigma_{-i}$ such that there exists a probability measure $\mu$ on $S$ such that $supp(\mu) = S$ and for every $\mathfrak{d}_i \in D$, $\sigma_i \in \mathfrak{d}_i(\mu)$, and for every $\mathfrak{d}_i \notin D$, $\sigma_i \notin \mathfrak{d}_i(\mu)$. **G5** and **G6** are special cases occurring when $|\mathfrak{D}_i| = 1$, corresponding to $D = \mathfrak{D}_i$ and $D = \emptyset$, respectively.


# Acknowledgements

Bjorndahl is supported in part by NSF grants IIS-0812045, CCF-1214844, DMS-0852811, and DMS-1161175, and ARO grant W911NF-09-1-0281. Halpern is supported in part by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, by MURI grant FA9550-12-1-0040, and by ARO grant W911NF-09-1-0281. Pass is supported in part by an Alfred P. Sloan Fellowship, a Microsoft Research Faculty Fellowship, NSF Awards CNS-1217821 and CCF-1214844, NSF CAREER Award CCF-0746990, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government. This paper subsumes and extends an earlier conference publication [**?**].


# References


Aumann, R. J. (1999). Interactive epistemology II: probability. *International Journal of Game Theory 28*(3), 301–314.

Battigalli, P. and M. Siniscalchi (2002). Strong beliefs and forward-induction reasoning. *Journal of Economic Theory 106*, 356–391.

Bjorndahl, A., J. Y. Halpern, and R. Pass (2011). Reasoning about justified belief. In *Theoretical Aspects of Rationality and Knowledge: Proc. Thirteenth Conference (TARK 2011)*, pp. 221–227.

Brandenburger, A. and E. Dekel (1987). Rationalizability and correlated equilibria. *Econometrica 55*, 1391–1402.

Capra, M., J. K. Goeree, R. Gomez, and C. A. Holt (1999). Anomalous behavior in a traveler's dilemma. *American Economic Review 89*(3), 678–690.

Fagin, R., J. Y. Halpern, and N. Megiddo (1990). A logic for reasoning about probabilities. *Information and Computation 87*(1/2), 78–128.

Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995). *Reasoning About Knowledge*. Cambridge, Mass.: MIT Press. A slightly revised paperback version was published in 2003.

Halpern, J. Y. (2003). *Reasoning About Uncertainty*. Cambridge, Mass.: MIT Press.

Halpern, J. Y. and Y. Moses (2007). Characterizing solution concepts in games using knowledge-based programs. In *Proc. Twentieth International Joint Conference on Artificial Intelligence (IJCAI '07)*, pp. 1300–1307.

Halpern, J. Y. and R. Pass (2012). Iterated regret minimization: a new solution concept. *Games and Economic Behavior 74*(1), 194–207.

Halpern, J. Y. and R. Pass (2013). Conservative belief and rationality. *Games and Economic Behavior 80*, 186–192.

Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory 139*(1), 242–268.

Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge, Mass.: MIT Press.

Rosenthal, R. W. (1982). Games of perfect information, predatory pricing, and the chain store paradox. *Journal of Economic Theory 25*, 92–100.

Tan, T. and S. Werlang (1988). The Bayesian foundation of solution concepts of games. *Journal of Economic Theory 45*(45), 370–391.