Reasoning About Rationality

A. Bjorndahl[a,*], J.Y. Halpern[b], R. Pass[b]

[a]*Dept. of Philosophy, Carnegie Mellon University, Pittsburgh, PA, USA 15213*
[b]*Dept. of Computer Science, Cornell University, Ithaca, NY, USA 14853*

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 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, standard foundations of many solution concepts require not only that each player

*Corresponding author
Email addresses: `abjorn@andrew.cmu.edu` (A. Bjorndahl), `halpern@cs.cornell.edu` (J.Y. Halpern), `rafael@cs.cornell.edu` (R. Pass)

is rational, but also that each player believes that her opponents are rational, believes that her opponents believe that *their* opponents are rational, and so on. Formal logic furnishes a powerful and versatile framework for representing such complex epistemic reasoning; 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.

Axiomatization is a way of distilling the mathematical properties of a given class of models into a few core principles from which all the rest follow deductively. This provides valuable insight into the logical assumptions implicit in these models. Moreover, in abstracting away from the full mathematical structure, axiom systems serve to illuminate commonalities between different kinds of models and provide logical tools to reason systematically about them.

Consider the standard KD45 axiomatization of the basic modal language of belief (see Sections 2.1 and 3). Intuitively, the axioms of this system represent the foundational properties of a certain conception of belief: the core principles from which all and only the true statements about belief follow. It is standard to employ probability measures to model belief (see Section 2.2), and KD45 axiomatizes certain classes of such models. But this system also axiomatizes classes of models that employ binary relations rather than probability measures to interpret belief, thereby establishing a precise mathematical connection between such so-called "qualitative" models and their "quantitative" or "probabilistic" counterparts.

The main axiomatization result we present in this paper extends KD45, and is defined with respect to a logical language expressive enough to talk about rationality as well as belief. We take as our point of departure axioms for rationality in the sense of expected utility maximization first articulated in [1]. 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 [2] for a discussion of all the decision rules mentioned in this paper). Though the full range of decision rules we consider is quite wide, our axiomatization results for the corresponding notions of rationality are all essentially the same, revealing a common thread running through all these formalizations of rational decision making. Of course, the nature of any axiomatization depends on the expressivity of the underlying logical language, a point we return to throughout the paper, particularly in Section 3.2. In essence, we work with a language that is rich enough to capture rationality and belief without being so rich as to *require* probabilistic models for its interpretation. As such, our results hold in both quantitative and qualitative frameworks (see Sections 2.2 and 3). 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, having developed these logics for reasoning about different notions of rationality, 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 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 $(x_1, \ldots, x_{i-1}, x_i', x_{i+1}, \ldots, x_n)$.

A **(normal-form) game** is a tuple $\Gamma = (I, (\Sigma_i)_{i \in I}, (u_i)_{i \in I})$ where $I = \{1, \ldots, n\}$ is the set of *players*, Σ_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 Γ, 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 Φ of primitive propositions, let $\mathcal{L}_B(\Phi)$ be the language consisting of those formulas obtained by closing off the set Φ of primitive propositions under logical negation ($\neg$), conjunction ($\wedge$), and the unary modalities B_i (for $i \in I$), where $B_i\phi$ is read "player i believes ϕ". We define other connectives such as disjunction ($\vee$), the material conditional ($\rightarrow$), and the biconditional ($\leftrightarrow$) as usual in terms of $\wedge$ and $\neg$, and write $\widehat{B}_i\phi$ for $\neg B_i \neg\phi$ ("player i consider it possible that ϕ"). We also write

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

for "everyone believes ϕ" and its k-fold iteration, respectively. Let

$$\Phi_\Gamma := \{play_i(\sigma_i) \,:\, i \in I, \, \sigma_i \in \Sigma_i\},$$

where we read $play_i(\sigma_i)$ as "player i is playing strategy σ_i"; we write

$$\begin{aligned} play(\sigma) &\equiv play_1(\sigma_1) \wedge \cdots \wedge play_n(\sigma_n) \text{ and} \\ play_{-i}(\sigma_{-i}) &\equiv \bigwedge_{j \neq i} play_j(\sigma_j) \end{aligned}$$

for "the players are playing according to the strategy profile σ" 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 ω and each player i is a probability measure on the set of all worlds, used to interpret player i's beliefs at ω. In the case of a language like $\mathcal{L}_B(\Phi_\Gamma)$, we also must take care to interpret the primitive propositions appropriately.

A **(countable) Γ-structure** is a tuple $M = (\Omega, (Pr_i)_{i\in I}, s)$ satisfying the following conditions:

(C1) Ω is a nonempty, countable set;

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

(C3) if $Pr_i(\omega') \neq Pr_i(\omega)$, then $Pr_i(\omega)(\omega') = 0$;

(C4) $s : \Omega \rightarrow \Sigma$ is such that if $s(\omega') \neq s(\omega)$, then $Pr_i(\omega)(\omega') = 0$.

Conditions (C1) and (C2) set the stage to interpret player i's beliefs at ω by the measure $Pr_i(\omega)$. Condition (C3) then ensures that at each world ω, 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 ω a strategy profile $s(\omega)$ in game Γ—intuitively, the strategy that each player is playing at ω—and moreover, that each player is sure of her own strategy.

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

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

Thus, the primitive propositions are interpreted in the obvious way using the strategy function (s_i denotes the ith component function of s), the Boolean connectives are interpreted classically, and the formula $B_i\phi$ holds at all and only those worlds ω at which $Pr_i(\omega)$ assigns probability 1 to ϕ. Dually, it is easy to check that

$$[\![\hat{B}_i\phi]\!]_M \quad := \quad \{\omega \in \Omega \,:\, Pr_i(\omega)([\![\phi]\!]_M) > 0\}.$$

As is standard, we often write $(M, \omega) \models \phi$ or just $\omega \models \phi$ to indicate that $\omega \in [\![\varphi]\!]_M$. Similarly, we write $M \models \phi$, and say that ϕ *is valid in* M, if $[\![\phi]\!]_M = \Omega$, and when $(M, \omega) \not\models \phi$, we say that M *refutes* ϕ *(at* ω*)*, or just that ω refutes φ.

Γ-structures come equipped with full-fledged probability measures, but this may seem like overkill since the belief modalities are interpreted only by reference to events of probability 1. One way to make this intuition precise is by defining a simpler kind of model: a **(countable) qualitative Γ-structure** is a tuple $M = (\Omega, (\mathcal{S}_i)_{i \in I}, s)$ satisfying the following conditions:

(C1′) Ω is a nonempty, countable set;

(C2′) $\mathcal{S}_i$ associates with each $\omega \in \Omega$ a subset $\mathcal{S}_i(\omega) \subseteq \Omega$;

(C3′) if $\mathcal{S}_i(\omega') \neq \mathcal{S}_i(\omega)$, then $\omega' \notin \mathcal{S}_i(\omega)$;

(C4′) $s : \Omega \to \Sigma$ is such that if $s(\omega') \neq s(\omega)$, then $\omega' \notin \mathcal{S}_i(\omega)$.

Think of $\mathcal{S}_i(\omega)$ as the set of worlds agent i considers possible at ω. Define an interpretation as before except replacing the clause for the belief modalities with the following:

$$[\![B_i\phi]\!]_M \quad := \quad \{\omega \in \Omega \,:\, \mathcal{S}_i(\omega) \subseteq [\![\phi]\!]_M\}.$$

It is not hard to see that every Γ-structure induces a qualitative Γ-structure by replacing each probability measure with its support: that is, by setting $\mathcal{S}_i(\omega) = \{\omega' \in \Omega : Pr_i(\omega)(\omega') \neq 0\}$. It is also easy to see that this transformation preserves the truth value of all formulas at all worlds. Despite this equivalence, we focus in this paper on the richer class of Γ-structures. This is because, as we will see in the next section, many of the notions of rationality we seek to represent in our logic depend crucially on quantitative information about the players' probabilistic beliefs; this information simply is not available in qualitive Γ-structures.

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 Γ and a probability measure μ on Σ_{-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 μ:

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

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 Γ) with a function $\mathfrak{d}_i^{eu}$ that takes as input a belief μ on Σ_{-i} and returns as output the set of strategies $\sigma_i \in \Sigma_i$ that maximize player i's expected utility given μ:

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

Let $\Delta(\Sigma_{-i})$ denote the set of all probability measures on Σ_{-i}. Then, generalizing the above, we define a **decision rule for player i (in Γ)** 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 σ_i is a best response for player i to the belief μ 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 Γ-structures; roughly speaking, for each world ω, we can define the set of $\mathfrak{d}_i$-best responses for player i at ω, 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 ω. Formally, given a Γ-structure M, for each player i and each world ω, the probability measure $Pr_i(\omega)$ induces a probability measure $\mu_{i,\omega}$ defined on Σ_{-i} as follows:

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

Mathematically, $\mu_{i,\omega}$ is the *pushforward* of $Pr_i(\omega)$ by s_{-i}; it captures what are sometimes called the *first-order beliefs* of player i at ω. Since $Pr_i(\omega)$ is interpreted as representing player i's beliefs at ω, 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 ω** to be $\mathfrak{d}_i(\mu_{i,\omega})$. We say that player i is **$\mathfrak{d}_i$-rational at ω** just in case $s_i(\omega) \in \mathfrak{d}_i(\mu_{i,\omega})$.

Since we wish to reason formally about rationality in games, we expand our logical 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

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

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

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

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 Γ-structures. For example, as shown by Tan and Werlang **(year?)** and Brandenburger and Dekel **(year?)**, given a game Γ, a strategy σ_i is *rationalizable* if and only if there exists a Γ-structure M and a state ω 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). Thus, the languages $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ we have defined are expressive enough not only to capture any given notion of rationality, but also to characterize rationalizability with a countable set of formulas (the infinitary nature of this characterization can be removed by adding a *common belief* modality, though we do not pursue this theme here).

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 Σ_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]

[1]In this sense, we can think of a decision rule as a *knowledge-based program* [5]. Knowledge-

Example 1. In order to help solidify the framework just introduced, we present a simple example. Consider the standard Bach-or-Stravinsky game Γ_{BoS} [7], 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.

	Bach	Stravinsky
Bach	3,2	0,0
Stravinsky	0,0	2,3

Table 1: Bach or Stravinsky?

We now describe a Γ_{BoS}-structure in which we can reason about the beliefs and rationality of the players. Of course, there are many such Γ_{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 ω_1 and ω_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 μ_{1,ω_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 her expected utility on choosing Stravinsky is 1. Thus,

$$\omega_0 \models RAT_1^{\mathfrak{d}_1^{eu}} \leftrightarrow play_1(\text{Bach}).$$

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 ω_0. □

3. Axiomatization

Although reasoning formally about beliefs and rationality has long been recognized as important to game theory (see, e.g., [8]), 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.

Recall that an axiom system consists of a collection of *axiom schemes* (which are just families of formulas) and *rules of inference* (which give a way of deriving

based programs have previously been used to characterize rationality and solution concepts [6].

new formulas from a collection of formulas already proved). An axiom system AX is a *sound* axiomatization of a given language $\mathcal{L}$ with respect to a class of models $\mathcal{M}$ if every formula in $\mathcal{L}$ that is provable from AX is valid in every model in $\mathcal{M}$. Conversely, AX is a *complete* axiomatization of $\mathcal{L}$ with respect to $\mathcal{M}$ if every formula in $\mathcal{L}$ that is valid in every model in $\mathcal{M}$ is provable in AX. In our case, $\mathcal{M}$ will consist of the countable Γ-structures.

A sound and complete axiomatization is a way of understanding a class of models by condensing their various features into a more digestable form: a (short) list of core principles from which all other truths follow deductively. Consider the following axiom schemes and rule of inference:

K. $B_i(\phi \rightarrow \psi) \rightarrow (B_i\phi \rightarrow B_i\psi)$,

D. $B_i\phi \rightarrow \hat{B}_i\phi$,

4. $B_i\phi \rightarrow B_iB_i\phi$,

5. $\neg B_i\phi \rightarrow B_i\neg B_i\phi$,

N. From ϕ deduce $B_i\phi$.

Adding these to any axiomatization of classical propositional logic produces the system called $\mathsf{KD45}$, which is sound and complete with respect to several important classes of models for belief. The axioms help us understand the nature of belief in these models; they also make plain the implicational relations between the various properties of belief. For instance, the requirement that agents be "introspective" about their beliefs—that is, sure of their own beliefs—is revealed to be a conjunction of two related but logically independent principles: a *positive* and a *negative* version of introspection as captured by axioms **4** and **5**, respectively.

All the axiomatizations we present in this paper are extensions of the basic $\mathsf{KD45}$ system. The logics we are concerned with have been specialized to express statements about a given game Γ, and have been extended to include the primitive propositions RAT_i, interpreted as rationality with respect to a given decision rule $\mathfrak{d}_i$. Each of these features requires additional axioms.

Fix a game Γ and a profile of decision rules $\mathfrak{d}$. To begin, we encode the fact that each player must play exactly one strategy:

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

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

Next come introspection conditions for both strategies and rationality: each player is certain of her strategy and of whether or not she is rational:

G3. $play_i(\sigma_i) \leftrightarrow B_i\, play_i(\sigma_i)$,[2]

[2]In fact, it is not hard to see (in light of **G1** and **G2**) that the reverse implication

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

G1–**G4** are easy to state and understand. What remains are the axioms that capture the specific nature of rationality as expressible in the language; **G4** identifies it as the kind of thing players are always sure about, but is silent about any further properties it might have. In order to state the final axioms of our system, we require some preliminary definitions. Given a subset $S \subseteq \Sigma_{-i}$, let

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

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

Given a measure μ on a countable set X, let

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

the **support** of μ. 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 μ with $supp(\mu) = S$ and with respect to which σ_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 μ with $supp(\mu) = S$ and $\sigma_i \notin \mathfrak{d}_i(\mu)$. Consider the following axiom schemes:

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

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

Intuitively, **G5** says that if RAT_i holds and player i is playing σ_i, then player i must consider possible a collection of strategy profiles[3] on which she could place a probability that would justify her playing σ_i. **G6** is interpreted analogously. Since exactly one of the formulas $\delta_{i,S}$ holds for each player i, one might also read these axioms "in reverse": if i considers all and only the strategy profiles in S possible and is playing σ_i, then **G5** and **G6** taken together either (a) demand that player i rational, (b) demand that player i is irrational, or (c) make no demands, depending (respectively) on whether

(a) $S \in \mathcal{S}_i^+(\sigma_i) \setminus \mathcal{S}_i^-(\sigma_i)$,

---

$B_i\, play_i(\sigma_i) \rightarrow play_i(\sigma_i)$ is redundant here. We state this axiom scheme as a biconditional merely to emphasize the introspective nature of this principle. We thank an anonymous reviewer for pointing this redundancy out to us.

[3]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 σ_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.

(b) $S \in \mathcal{S}_i^-(\sigma_i) \setminus \mathcal{S}_i^+(\sigma_i)$, or

(c) $S \in \mathcal{S}_i^+(\sigma_i) \cap \mathcal{S}_i^-(\sigma_i)$.

Example 2. We define a simple game for the purpose of elucidating axioms **G5** and **G6**. In this example we fix the interpretation of rationality for all players as expected utility maximization. Consider a game Γ with $I = \{1, 2\}$, $\Sigma_1 = \{T, B\}$, $\Sigma_2 = \{L, M, R\}$, and utility functions $u_1 = u_2$ defined as shown in Table 2.

	L	M	R
T	1	0	3
B	2	1	2

Table 2: Utility values for outcomes in the game Γ.

For each player i and each available strategy $\sigma_i \in \Sigma_i$, we must determine membership in the corresponding sets $\mathcal{S}_i^+(\sigma_i)$ and $\mathcal{S}_i^-(\sigma_i)$; we consider a handful of particular examples in order to illustrate the general process.

First we observe that $\{L\} \in \mathcal{S}_1^-(T) \setminus \mathcal{S}_1^+(T)$; this is because the only probability measure that has support equal to $\{L\}$ is one for which T is not a best response. Next we check that $\{L, R\} \in \mathcal{S}_1^+(T) \cap \mathcal{S}_1^-(T)$; this follows from, for example, the fact that T is a best response to beliefs that place probability $\frac{1}{4}$ on L and $\frac{3}{4}$ on R, whereas T is not a best response to beliefs that place probability $\frac{3}{4}$ on L and $\frac{1}{4}$ on R. For a final example, it is not hard to see that $\mathcal{S}_2^+(M) = \mathcal{S}_2^-(R) = \emptyset$, since no matter what beliefs player 2 has, M cannot be a best response and R must be a best response.

One can check that

$$\mathcal{S}_2^+(L) = \{\{B\}\}$$

and

$$\mathcal{S}_1^-(B) = \{\{R\}, \{L, R\}, \{M, R\}, \{L, M, R\}\}.$$

It follows that the formulas

$$(play_2(L) \wedge RAT_2) \rightarrow \delta_{2,\{B\}}$$

and

$$(play_1(B) \wedge \neg RAT_1) \rightarrow (\delta_{1,\{R\}} \vee \delta_{1,\{L,R\}} \vee \delta_{1,\{M,R\}} \vee \delta_{1,\{L,M,R\}})$$

are instances of **G5** and **G6**, respectively. The first formula is equivalent to

$$(play_2(L) \wedge RAT_2) \rightarrow B_2\, play_1(B),$$

so it expresses the fact that if player 2 is rationally playing L, then he must be sure that player 1 is playing B. The second formula is equivalent to

$$(play_1(B) \wedge \neg RAT_1) \rightarrow \hat{B}_1\, play_2(R),$$

which says that if player 1 is irrationally playing B, then she must consider it possible that player 2 is playing R. □

Notice that player i's actual (i.e., quantitative) beliefs are not fully specified in **G5** or **G6**. This may be somewhat surprising, given that the semantic interpretation of RAT_i depends (in principle) on the quantitative probabilities that constitute player i's beliefs: how can it be that a notion of rationality induced by a quantitative decision rule is fully captured by a pair of axioms that make no reference to probabilities? Of course, for decision rules that are purely *qualitative*—i.e., defined solely in terms what each player considers possible, rather than the probabilitic weights they assign to the different possibilities (as, for example, with maximin)—this is a non-issue. The mismatch arises precisely when rationality depends on the quantitative nature of belief, because the belief modalities themselves do not (see Section 2.2, especially the definition of a qualitative Γ-structure).

In this case a richer language may well be of interest; we explore this point further as it pertains to expected utility maximization in Section 3.2. Nonetheless, the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ is an important one for the analysis of rationality in games despite (and, in some senses, because of) its expressive limitations. For one thing, as we have already noted, it is expressive enough to characterize iterative rationality. Moreover, for those concerned with purely qualitative decision rules $\mathfrak{d}$, the simpler class of qualitative Γ-structures are natural models to work with, and while it is easy to see that $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ is interpretable in such structures, a richer language that could express quantitive beliefs would not be. At a high level, what axioms **G5** and **G6** offer us is a qualitative perspective on arbitrary decision rules. That these principles *must* hold is perhaps not terribly surprising; on the other hand, that they constitute *all* that can be said about rationality in this framework is quite notable.

Let $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ be the axiom system that results from adding **G1**–**G6** to the $\mathsf{KD45}$ axioms and rules of inference. These axioms completely characterize the logical properties of rationality as expressible in the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$.

Theorem 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 Γ-structures.*[4]

Corollary 2. *Let $\mathfrak{d}$ be a profile of qualitative decision rules. Then $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ is a sound and complete axiomatization of the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ with respect to the class of all qualitative Γ-structures.*

It bears emphasizing that $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ is parametrized by two variables: the underlying game Γ 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 Γ 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

[4]Our proof shows that completeness holds even with respect to finite Γ-structures. That is, we do not get any extra axioms if we restrict to finite Γ-structures.

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 [2] 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 Γ. A **(countable) lower Γ-structure** is a tuple $M = (\Omega, (\mathcal{PR}_i)_{i\in I}, s)$ satisfying the following conditions:

(L1) Ω is a nonempty, countable set;

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

(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:

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

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:

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

Other rules, such as minimax expected regret [9], can also easily be defined in this setting.

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

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

Then we can define $\mathfrak{d}_i$**-rationality** for player i at ω 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

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

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

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

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 $supp(\mathcal{P}) = S$ 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 $supp(\mathcal{P}) = S$ 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. *$\mathsf{GL}_\Gamma^{\mathfrak{d}}$ is a sound and complete axiomatization of the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ with respect to the class of all lower Γ-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^α for each $\alpha \in [0,1]$, where $B_i^\alpha \phi$ is interpreted as saying that player i assigns probability α to ϕ. 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 Γ it is rational only for player i to play σ_i if she assigns probability $\frac{1}{2}$ to σ_{-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 Σ_{-i} is finite and therefore has only finitely-many subsets. By contrast, in the language with belief modalities B_i^α 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 Γ and consider the language of *linear likelihood inequalities* defined by the grammar

$$\phi ::= p \,|\, \neg\phi \,|\, \phi \wedge \psi \,|\, a_1\ell_i(\phi_1) + \cdots + a_k\ell_i(\phi_k) \geq b,$$

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* $\ell_i(\phi)$ are meant to be read as "the probability of ϕ according to player i", and a likelihood formula $a_1\ell_i(\phi_1)+\cdots+a_k\ell_i(\phi_k) \geq b$ should be thought of as asserting the corresponding inequality. More precisely, we interpret such formulas in a Γ-structure M as follows:

$$[\![a_1\ell_i(\phi_1) + \cdots + a_k\ell_i(\phi_k) \geq b]\!]_M \quad := \quad \{\omega \in \Omega \,:\, \sum_{j=1}^{k} a_j Pr_i(\omega)([\![\phi_j]\!]_M) \geq b\}.$$

For example, the formula $\ell_i(\phi) \geq 1$ says that player i assigns probability at least (and therefore exactly) 1 to ϕ, while the formula $(\ell_i(\phi) \geq \frac{1}{2}) \wedge (\ell_i(\neg\phi) \geq \frac{1}{2})$ says that player i assigns probability $\frac{1}{2}$ to ϕ. See [2] for a thorough discussion of this and related logics; a sound and complete axiomatization is given in [10].

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 σ_i is no less than the expected value of playing σ_i', then it is easy to see that

$$RAT_i^{\mathfrak{d}_i^{eu}} \leftrightarrow \bigvee_{\sigma_i\in\Sigma_i} \Big(play_i(\sigma_i) \;\wedge \bigwedge_{\sigma_i'\in\Sigma_i} \sigma_i \succeq \sigma_i'\Big) \tag{1}$$

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

It is not hard to see that a variety of other decision rules (such as maximin and minimax regret) can be likewise captured in this language, so the corresponding notions of rationality can be axiomatized in a similar way. However, these axiomatizations are not particularly informative: the axioms of this language capture general principles of probabilities, leaving the specific nature and properties of rationality completely (and opaquely) wrapped up in definitions as in (1). One might as well point out that the foundational axioms of set theory "axiomatize" rationality in the sense that expected utility maximization can be defined in this general mathematical setting. Languages rich enough to define the notions of rationality we care about give no specific insight into these notions beyond the definitions themselves. Depending on the language, of course, these definitions may be interesting or informative in their own right. For instance, in recent work, Lorini **(year?)** shows that a simple logic capable of expressing knowledge, belief, and "strong belief" can be used to define strong and weak dominance, as well as various solution concepts based on iterative deletion of dominated strategies. Part of the value of Lorini's work therefore lies in the relationship it exposes between the foundational notions of knowledge and (strong) belief and these solution concepts. Our contribution, by contrast, lies not in exploring the relationship between rationality and more "primitive" concepts, but in treating rationality itself as a primitive concept and analyzing its properties.

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 σ_j, but if she plays σ'_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

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

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 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 Γ-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 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., [12]). But in some cases, this too is unsatisfying: for example, "continuing" at the second-last stage of the centipede game (see Example 4) is rational only for a player who believes her 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 Γ 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 Φ_Γ of primitive propositions that we considered earlier by taking

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

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 ω 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 ω a profile of decision rules $r(\omega)$—intuitively, $r_i(\omega) \in \mathfrak{D}_i$ is the rule that player i is using at

ω—and moreover, that each player is sure of her own decision rule. Condition (C6) requires that, at each world ω, 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 ω.

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 [13]; it is reasonable in such games 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 [14], 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 σ_i to be **$\mathfrak{D}$-rationalizable (in Γ)** just in case there exists a $(\Gamma,\mathfrak{D})$-structure in which σ_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, using a straightforward iterated deletion argument, that when the strategy sets are finite, $\mathfrak{D}$-rationalizable strategies must exist.[5] 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:

[5] 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 Σ 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. *For each player i, let $\mathfrak{D}_i \subseteq \mathfrak{D}'_i$. Then if σ_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 Γ and probability measure μ on Σ_{-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 σ_i, where "regret" is interpreted as the difference between the best-case payoff and the actual payoff (with respect to the strategy profiles σ_{-i} that player i considers possible). The minimax regret decision rule $\mathfrak{d}^r_i$ seeks to minimize this value:

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

Example 3. 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:

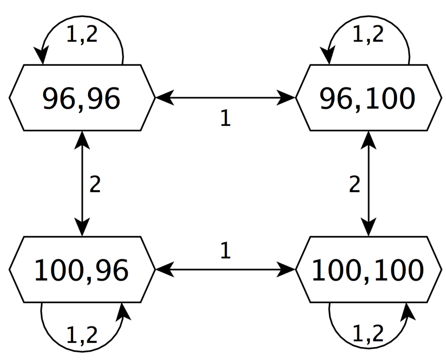

Figure 1: A $(\Gamma, \mathfrak{D})$-structure for the traveler's dilemma

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—we are considering regret, not expected regret—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 when 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 when 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 when 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. Consider the normal-form version of the centipede game [15] 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 minimal stage at which some player quit, 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.

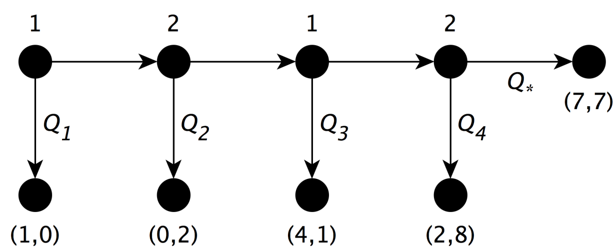

Figure 2: The centipede game

It is well known that in this normal-form version of the centipede game, all pure strategies 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 [16]. By contrast, we now show that when $\mathfrak{D}_i = \{\mathfrak{d}_i^{eu}, \mathfrak{d}_i^r\}$ for $i = 1, 2$, 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 ω 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 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 in which each of these strategies is played and all beliefs are conservative. Consider the structure presented in Figure 3:

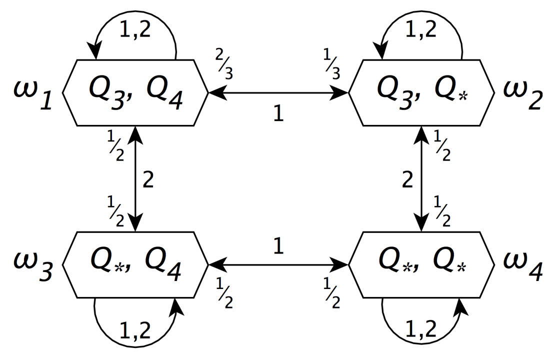

Figure 3: A $(\Gamma, \mathfrak{D})$-structure for the centipede game

As in Figure 1, each of the four states ω_1, ω_2, ω_3, and ω_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 ω_1 and ω_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 ω_3 and ω_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 when 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 ω_3 and ω_4.

Similar arguments show that Q_4 is a $\mathfrak{d}_2^{eu}$-best response in states ω_1 and ω_3, and Q_* is a $\mathfrak{d}_2^r$-best response in states ω_2 and ω_4. Thus, we can set

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

to make Figure 3 into a $(\Gamma, \mathfrak{D})$-structure and each of the strategies played in the structure $\mathfrak{D}$-rationalizable 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. $play_i(\sigma_i) \leftrightarrow B_i\, play_i(\sigma_i)$,[6]

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 μ 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 σ_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 σ_i (though she may, coincidentally, have beliefs with respect to which σ_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 6 given in Appendix Appendix A.1.

Theorem 5. $\mathsf{GL}_{\Gamma,\mathfrak{D}}$ *is a sound and complete axiomatization of the language* $\mathcal{L}_B(\Phi_{\Gamma,\mathfrak{D}})$ *with respect to the class of all* $(\Gamma, \mathfrak{D})$*-structures.*

[6] As with **G3**, the reverse implications for **P3** and **P6** are redundant, but we write these axioms as biconditionals to emphasize the introspective nature of strategies and rules in this setting. A similar remark holds for our later soundness and completeness results, although we do not mention it explicitly.

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 σ_i in the support of ℓ_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 Σ_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 Γ to a decision rule in Γ. Suppose that we are given decision paradigms $\mathcal{D}_i$ for each player associating with each game Γ a decision rule $\mathcal{D}_i(\Gamma)$ for that player in Γ. 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 Γ, can we effectively determine whether a formula belongs to the axiom system $\mathsf{GL}_{\Gamma}^{\mathcal{D}(\Gamma)}$? For each game Γ, 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, given a non-computable set $H \subseteq \mathbb{N}$, we could define $\mathcal{D}_i(\Gamma)$ depending on whether the number of players in Γ lies in H.

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 σ_i only if γ", where γ 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 $\mathfrak{D}$-rationalizability in this paper, we hope to initiate a broader research program investigating the role of state-dependent decision rules in solution concepts. 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.

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. A preliminary version of this paper appeared in the Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning, 2014 [17]. This version includes more detailed proofs and a significantly expanded discussion; in addition, most of the material in Section 4 is new, including the definition of $\mathfrak{D}$-rationalizability and all the results and examples pertaining thereto.

Appendix A. Proofs

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

Theorem 6. $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ *is a sound axiomatization of the language $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ with respect to the class of all Γ-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 Γ-structures.

Fix a Γ-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 $\omega \models play_i(\sigma_i) \wedge RAT_i$; then $\sigma_i \in \mathfrak{d}_i(\mu_{i,\omega})$. Set $S \coloneqq supp(\mu_{i,\omega})$, and observe that $S \in \mathcal{S}_i^+(\sigma_i)$. For each $\sigma_{-i} \in S$, it is easy to see that $\omega \models \widehat{B}_i\, play_{-i}(\sigma_{-i})$; therefore, we must have

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

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

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

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**. $\square$

We prove completeness by what is essentially the *canonical model method*, a standard method for proving completeness of modal logics (see, e.g., [5] or any standard text on modal logic). Of course, the full canonical model is not countable (for $n > 1$), so we modify the construction by restricting attention to finite sub-languages. More precisely, given a formula $\phi \in \mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$, we identify a finite sub-language of $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ and a corresponding canonical model that refutes ϕ just in case $\mathsf{GL}_\Gamma^{\mathfrak{d}} \nvdash \phi$. This technique is sometimes called *filtration*; it will actually establish the slightly stronger result mentioned in footnote 4, namely, completeness with respect to the class of all *finite* Γ-structures.

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

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

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

Let Ω^ϕ be the collection of all maximal, consistent (with respect to $\mathsf{GL}_\Gamma^{\mathfrak{d}}$) subsets of $Sub_\Gamma^+(\phi)$. Clearly Ω^ϕ 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 $F \in \Omega^\phi$, define

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

For each $i \in I$ and each $F \in \Omega^\phi$, we will define a probability measure $Pr_i^\phi(F)$ on Ω^ϕ such that the support of this measure is precisely $\mathrm{Bel}_i(F)$. Loosely speaking, $\mathrm{Bel}_i(F)$ is the set of all $G \in \Omega^\phi$ 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 7. *For each $i \in I$ and $F \in \Omega^\phi$,* $\mathrm{Bel}_i(F) \neq \emptyset$.

PROOF. Given $F \in \Omega^\phi$, 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 $\Lambda \subset Sub^+_\Gamma(\phi)$. In addition, we show that Λ 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 $G \in \Omega^\phi$ 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. $\square$

Lemma 8. *Let $F \in \Omega^\phi$. 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

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

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

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

As observed in Lemma 7, each $B_i\xi$ is equivalent to a formula in F, and therefore (A.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 $G \in \Omega^\phi$; moreover, as we saw in Lemma 7, $G \in \mathrm{Bel}_i(F)$. 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 $\psi \notin G$. $\square$

In the classical canonical model construction, it is sufficient to define $Pr^\phi_i(F)$ to be the uniform distribution on $\mathrm{Bel}_i(F)$. In the present context, however, we need to be more careful, since $Pr^\phi_i(F)$ is used not only to interpret the belief modalities B_i, but also the primitive propositions RAT_i. In essence, we must define $Pr_i^\phi(F)$ 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^ϕ on Ω^ϕ 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.

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 $F \in \Omega^\phi$, 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 9. *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 8 together with the fact that

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

$\square$

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 $F \in \Omega^\phi$, and suppose that $\mathit{play}_i(\sigma_i) \in F$. In light of Lemma 9, we can define $Pr_i^\phi(F)$ to be the unique probability measure on $\mathrm{Bel}_i(F)$ such that, for all $\sigma_{-i} \in \Sigma_{-i}$,

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

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

Proposition 10. Pr_i^ϕ *satisfies the following:*

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

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

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 9, we know that $G \in \mathrm{Bel}_i(F;\sigma_{-i})$ for some $\sigma_{-i} \in S_i(F)$, from which it follows that $Pr_i^\phi(F)(G) > 0$ 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^ϕ, we can deduce that $Pr_i^\phi(G) = Pr_i^\phi(F)$. □

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

Lemma 11. *The tuple* $M^\phi \coloneqq (\Omega^\phi, (Pr_i^\phi)_{i\in I}, s^\phi)$ *is a* Γ*-structure.*

Proof. Conditions (C1) and (C2) have already been established. By Lemma 10(b), in order to see that (C3) holds it suffices to observe that $Pr_i^\phi(F)(G) > 0$ implies that $G^{B_i} = F^{B_i}$, which follows from Lemma 10(a). This same observation also establishes (C4), since by **G3** we know that $G^{B_i} = F^{B_i}$ implies $s_i^\phi(G) = s_i^\phi(F)$. □

Lemma 12. *For all formulas* ψ*, for all* $F \in \Omega^\phi$*, if* $\psi \in Sub_\Gamma^+(\phi)$ *then* $F \in [\![\psi]\!]_{M^\phi}$ *if and only if* $\psi \in F$.

Proof. The proof proceeds by induction on the structure of ψ. We prove here the base cases corresponding to the primitive propositions; the inductive steps can be proved in the standard way (see, e.g., [5]).

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

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

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

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

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

Theorem 13. $\mathsf{GL}_\Gamma^\mathfrak{d}$ *is a complete axiomatization of the language* $\mathcal{L}_B(\Phi_\Gamma^\mathfrak{d})$ *with respect to the class of all finite* Γ*-structures.*

Proof. Suppose that $\mathsf{GL}_\Gamma^\mathfrak{d} \nvdash \phi$. Then $\{\neg\phi\}$ is consistent and so can be extended to a maximal consistent set $F \in \Omega^\phi$. By Lemma 12, this implies that $F \notin [\![\phi]\!]_{M^\phi}$ and so, in particular, that $M^\phi \not\models \phi$, as desired. □

Appendix A.2. Belief as lower probability

Theorem 14. $\mathsf{GL}_\Gamma^{\mathfrak{d}}$ *is a sound and complete axiomatization of the language* $\mathcal{L}_B(\Phi_\Gamma^{\mathfrak{d}})$ *with respect to the class of all lower* Γ*-structures.*

Proof. The proof given in Appendix 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 $F \in \Omega^\phi$ with $play_i(\sigma_i) \in F$, define $\mathcal{PR}_i^\phi(F)$ 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

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

□

Appendix 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 μ 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 μ 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.

References

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

[2] J. Y. Halpern, Reasoning About Uncertainty, MIT Press, Cambridge, MA, 2003.

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

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

[5] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, Reasoning About Knowledge, MIT Press, Cambridge, Mass., 1995, a slightly revised paperback version was published in 2003.

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

[7] M. J. Osborne, A. Rubinstein, A Course in Game Theory, MIT Press, Cambridge, MA, 1994.

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

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

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

[11] E. Lorini, A minimal logic for interactive epistemology, Synthese 3 (2016) 725–755.

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

[13] J. Y. Halpern, R. Pass, Iterated regret minimization: a new solution concept, Games and Economic Behavior 74 (1) (2012) 194–207. doi:10.1016/j.geb.2011.05.012.

[14] M. Capra, J. K. Goeree, R. Gomez, C. A. Holt, Anomalous behavior in a traveler's dilemma, American Economic Review 89 (3) (1999) 678–690. doi:10.1257/aer.89.3.678.

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

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

[17] A. Bjorndahl, J. Y. Halpern, R. Pass, Axiomatizing rationality, in: Principles of Knowledge Representation and Reasoning: Proc. Fourteenth International Conference (KR '14), 2014, pp. 178–187.