Strategic Reasoning in Games

Eric Pacuit

Department of Philosophy
University of Maryland, College Park, MD, 20742, USA
epacuit@umd.edu,
http://ai.stanford.edu/~epacuit

1 Introduction

1.1 General motivation

Strategies are the basic objects of study in a game-theoretic model. The standard interpretation is that a strategy represents a player's *general plan of action.* That is, a strategy for player i describes the action that player i will choose (allowing players to *randomize* their choices) whenever she is required to make a decision according to the rules of the game.

Traditional game theory has focused on identifying profiles of strategies that constitute an "equilibrium" (typical examples include the Nash equilibrium and its refinements or the various notions of iterated dominance). These "equilibrium profiles" are intended to represent the "rational outcomes" of a strategic interactive situation. Much of this work has led game-theorists to rethink the basic view that strategies are simply instructions that "rational" players follow when playing the game. For example, in a well-known paper on the foundations of game theory, Ariel Rubinstein writes:

> I aim to endorse the view that equilibrium strategy describes a player's plan of action, as well as those considerations which support the optimality of his plan (i.e. preconceived ideas concerning the other players' plans) rather than being merely a description of a "plan of action."
> [1, pg. 910]

This broader understanding of a strategy raises a number of interesting foundational questions. It is beyond the scope of this paper to discuss all of these issues (see [1] for an illuminating discussion, including how to interpret *mixed strategies* and strategies in sequential games, and broader issues on how to interpret game-theoretic models). In this paper, I focus on a key foundational issue that stems from the broader understanding of a strategy described in the above quote: How do the (rational or not-so rational) players decide what to do in a strategic situation?

This question directs our analysis to aspects of a strategic interactive situation that are not typically covered by standard game-theoretic models. Much of the work in game theory is focused on identifying the rational *outcomes* of an game-theoretic situation. This is in line with the standard view of a strategy as simply a description of what the players (should) *do*. However, as noted above, a broader understanding of a strategy also includes an important *epistemic* component describing the players opinions about why a particular strategy is a "good" plan of action. This naturally shifts the focus to the underlying *process* of deliberation that leads (rational) players to adopt certain plans of action.

The overall theme of this paper is that the players' "reasoning process" should be made explicit in game-theoretic models. This is in line with much of the recent work on the "epistemic foundations of solution concepts" [2, 3]. I will not attempt to survey this vast literature here (see [4] for a discussion and links to the relevant work). The key idea is to explicitly describe the "informational context" of a game situation (what the players think each other will do, think each other thinks each other will do, and so on) and then *derive* the fact that the players' choices adhere to a given solution concept from an epistemic property (eg., common belief of rationality). But, a question still remains: How do the players arrive at a particular informational context?

1.2 This paper

In this paper, I am not interested in strategies *per se*, but rather, the process of "rational deliberation" that leads players to adopt are particular "plan of action". This has both a normative component (What are the normative principles that guide the players' decision making?) and a descriptive component (Which psychological phenomena best explain discrepancies between predicted and observed behavior in game situations?).

The main challenge is to find the right balance between descriptive accuracy and normative relevance. While this is true for all theories of individual decision making and reasoning, focusing on game situations raises a number of compelling issues. Robert Aumann and Jacques Dreze [2, pg. 81] adeptly summarize one of the most pressing issues when they write: "the fundamental insight of game theory [is] that a rational player must take into account that the players reason about each other in deciding how to play". Exactly how the players (should) incorporate the fact that they are interacting with other (actively reasoning) agents into their own decision making process is the subject of much debate.

A number of frameworks set out to model this "rational deliberation" process in game situations. Key examples include:

1. John C. Harsanyi's *tracing procedure* [5]. The goal of the tracing procedure is to identify a unique Nash equilibrium in a strategic game. The idea is to analyze a continuum of games starting with the original game where the payoffs are replaced by expected payoffs. In fact, Harsanyi himself thought of this procedure as "being a mathematical formalization of the process by which rational players coordinate their choices of strategies."
2. Brian Skyrms' model of "dynamic deliberation" where players deliberate by calculating their expected utility and then use this new information to recalculate their probabilities about the states of the world and their expected utilities [6].

3. Ken Binmore's analysis of rational decision making where players are represented as Turing machines which can *compute* their rational choice(s) [7].
4. Robin Cubitt and Robert Sugden's recent contribution developing a "reasoning-based expected utility" procedure for solving games (building on David Lewis' "common modes of reasoning") [8, 9].
5. Johan van Benthem et col.'s analysis of solution concepts as fixed-points of iterated "(virtual) rationality announcements" [10–14].

Although the details of these frameworks are quite different, they share a common line of thought: In contrast to classical game theory, solution concepts are no longer the basic object of study. Instead, the "rational solutions" of a game are arrived at through a process of "rational deliberation". My goal here is not to provide a comprehensive survey of these different approaches. Instead, I will provide a (biased) overview of some key technical and conceptual issues that arise when developing fine-grained models of players deliberating about what to do in a game situation. In particular, I focus on the last two frameworks listed above in order to connect with recent work on the logical models of belief change over time [15–17].

1.3 Structure of the paper

A general theory of rational deliberation for game and decision theory is a big topic. It is beyond the scope of this article to discuss the many different aspects and competing perspectives on such a theory. The reader is referred to Brian Skyrms' seminal book [6] (especially Chapter 7) for a broader discussion. This paper focuses on four questions concerning the role of (rational) deliberation in a game-theoretic analysis of an interactive situation. Before discussing these questions, in Section 2, I introduce the key notions and definitions from game theory and dynamic logics of knowledge and belief. The goal is not to be comprehensive here, but to provide enough background to fill in some of the technical details discussed in the later sections. The remainder of the paper will focus on four general questions concerning deliberation in games:

What constraints should be imposed on a game model that describe beliefs that are formed through a process of rational deliberation about what the other players might do or are thinking? Section 3 discusses some natural properties and a puzzling consequence.

Do considerations about the players' rational process of deliberation lead to a new type of solution concept? In Section 4, I discuss an iterative procedure recently proposed by Cubitt and Sugden [9] that is intended to describe outcomes that arise when Bayesian deliberators interact.

What type of process can be used to generate *a game model?* Drawing on the extensive literature on dynamics logics of knowledge and belief, I briefly address this question in Section 5.

How do the players' beliefs (about what each other is going to do) change during a play of the game? Section 6 briefly discusses belief change in *extensive games*, where the sequential nature of the game situation is explicitly represented.

Finally, in Section 7, I conclude with some general remarks about strategic reasoning in specific types of game situations.

2 Background

2.1 Strategic games

I assume that the reader is familiar with the basics of game theory (see [18] and [19] for concise discussions of the key concepts, definitions and the relevant textbooks). I introduce some key definitions below in order to fix notation.

A **strategic game** is a tuple $\langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ where N is a (finite) set of players, for each $i \in N$, S_i is a finite set (elements of which are called actions or strategies), and for each $i \in N$, $u_i : \Pi_{i \in N} S_i \to \mathbb{R}$ is a utility function assigning real numbers to each outcome of the game (i.e., tuples consisting on the choices for each player). [1] Strategic games represent situations where the players make a single simultaneous decision. A **strategy profile** is a sequence of strategy choices for each player. If σ is a strategy profile, then I write σ_i for the ith component of σ and σ_{-i} for the sequence consisting of all components of σ except for σ_i (let S_{-i} denote all such sequences of strategies). I make use of the following standard game-theoretic notions (see [19] for an extended discussion of these and related definitions):

A strategy $s \in S_i$ **strictly dominates** strategy $s' \in S_i$ provided,

$$\forall s_{-i} \in S_{-i} \quad u_i(s, s_{-i}) > u_i(s', s_{-i})$$

A strategy $s \in S_i$ **weakly dominates** strategy $s' \in S_i$ provided,

$$\forall s_{-i} \in S_{-i} \quad u_i(s, s_{-i}) \geq u_i(s', s_{-i}) \quad \text{and} \quad \exists s_{-i} \in S_{-i} \quad u_i(s, s_{-i}) > u_i(s', s_{-i})$$

More generally, I say s strictly/weakly dominates s' **with respect to a set** $X \subseteq S_{-i}$ if I replace S_{-i} with X in the above definitions. Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ and $G' = \langle N, \{S'_i\}_{i \in N}, \{u'_i\}_{i \in N} \rangle$ is are strategic games. I say G' is a **restriction** of G provided for each $i \in N$, $S'_i \subseteq S_i$ and u'_i is the restriction of u_i to $\Pi_{i \in N} S'_i \subseteq \Pi_{i \in N} S_i$. To simplify notation, write G_i for the set of strategies for player i. I also defined a notion of **reduction** between restrictions H and H' of G: write $H \longrightarrow_S H'$ whenever $H \neq H'$, H' is a restriction of H and

$$\forall i \in N, \forall s_i \in H_i \setminus H'_i \ \exists s'_i \in H_i \ s_i \text{ is strictly dominated in } H \text{ by } s'_i$$

[1] For simplicity, the outcomes of a game are identified with the choices of each player. This assumption is not crucial for what follows.

So, if $H \longrightarrow_S H'$, then H' is the result of removing some of the strictly dominated strategies from H. We can iterate this process of removing strictly dominated strategies. Formally, I say H is the result of iteratively removing strictly dominated strategies (IESDS) provided $G \longrightarrow_S^* H$, where $\longrightarrow^*$ is the reflexive transitive closure of a relation $\longrightarrow$.

The above definition can be easily adapted to other choice rules, such as weak dominance. Let $\longrightarrow_W$ denote the relation between games defined as above replacing strict dominance with weak dominance. [2] Furthermore, the above definition can be easily adapted for so-called *mixed extensions of strategic games*. Let $\Delta(X)$ denote the set of probability measures over the finite[3] set X. A **mixed strategy** for player i, is an element $m_i \in \Delta(S_i)$. If $m_i \in \Delta(S_i)$ assigns probability 1 to an element $s_i \in S_i$, then m_i is called a **pure strategy** (in such a case, I write s_i for m_i). The utility for player i of the joint mixed strategy $m \in \Pi_{i\in N}\Delta(S_i)$ is calculated in the obvious way (let $m(s) = m_1(s_1)\cdot m_2(s_2)\cdots m_n(s_n)$ for $s \in \Pi_{i\in N}S_i$):

$$u_i(m) = \sum_{s\in\Pi_{i\in N}S_i} m(s)\cdot u_i(s)$$

The above definitions of strict and weak dominance and iterated removal of strictly/weakly dominated strategies can be readily adapted to the mixed extensions of a strategic game.

An influential interpretation of a mixed strategy for player i is that it represents the *beliefs* of i's opponents about what she will do. This is nicely explained in Robert Aumann's influential paper [20] (see, especially, Section 6 for a discussion of this interpretation of mixed strategies):

> An important feather of our approach is that it does not require explicit randomization on the part of the players. Each player always chooses a definite pure strategy, with no attempt to randomize; the probabilistic nature of the strategies relfects the uncertainties of other players about his choice. [20, pg. 3]

Given this "epistemic interpretation" of mixed strategies, call a probability measure probability measure $\pi \in \Delta(S_{-i})$ a **conjecture** for player i. Let $G = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ be a strategic game. The **expected utility** of a strategy $s \in S_i$ for agent i with respect to a conjecture $\pi \in \Delta(S_{-i})$ is:

$$E_\pi(s) = \sum_{\sigma_{-i}\in S_{-i}} \pi(\sigma_{-i})\cdot u_i(s,\sigma_{-i})$$

A strategy $s \in S_i$ maximizes expected utility with respect to $\pi \in \Delta(S_{-i})$, denoted $MEU(s,\pi)$, if for all $s' \in S_i$, $E_\pi(s) \geq E_\pi(s')$.

[2] There are some interesting issues that arise here: it is well-known that, unlike with strict dominance, the order in which weakly dominated strategies are removed the different orders in which weakly dominated strategies are removed can lead to different outcomes. Let us set aside these issues for this paper.

[3] Recall that I am restricting attention to finite strategic games.

Maximizing expected utility and weak/strong dominance are fundamentally connected, as is shown by the following well-known Lemmas (see [21, Appendix B] for proofs).

Lemma 1. *A strategy $s \in S_i$ is not strictly dominated iff s does not maximize expected utility with respect to any conjecture $\pi \in \Delta(S_{-i})$.*

A conjecture $\pi \in \Delta(S_{-i})$ is said to be **full support** provided $supp(\pi) = \{\sigma_{-i} \in S_{-i} \mid \pi(\sigma_{-i}) > 0\} = S_{-i}$.

Lemma 2. *A strategy $s \in S_i$ is not weakly dominated iff s does not maximize expected utility with respect to any full support conjecture $\pi \in \Delta(S_{-i})$.*

2.2 Game models

A **game model** represents the "informational context" of a given play of the game. The informational context of a game describes the players' *hard* and *soft* information about the possible outcomes of the game. This includes the "knowledge" the players have about the game situation *and* their opinions about the choices and beliefs of the other players. Researchers interested in the foundation of decision theory, epistemic and doxastic logic and formal epistemology have developed many different formal models that can describe the many varieties of informational attitudes important for assessing the choice of a *rational* agent in a decision- or game-theoretic situation. I do not have the space here to provide a complete survey of this literature — see [22, 17] for an overview and pointers to the relevant literature.

Two main types of models have been used in the game theory literature to describe the players' beliefs (and other informational attitudes) in a game situation: *type spaces* [23, 24] and the so-called *Aumann-* or *Kripke-structures* [25, 26]. Although these two approaches have much in common, there are some important differences, but they are not crucial for this paper. A second, more fundamental, distinction is between "quantitative" structures, representing "graded" attitudes (typically via probability distributions), and "qualitative" structures representing "all-out" attitudes. In this section, I present the details of three logical frameworks that have been used to describe the knowledge beliefs of the players in a game situation.

Syntactic issues do not play an important role in this paper. Nonetheless, I will present both the semantics and the relevant formal language (giving the definition of truth). Standard logical questions about axiomatics, definability, decidability of the satisfiability problem, etc. can and have been studied for all of the frameworks discussed in this paper (see [26, 27]). I do not discuss any of these issues here; however, I present the details of the logical systems as it makes for a smoother transition from the game theory literature to the literature on dynamic epistemic logic and iterated belief change.

An immediate question when presenting any formal language for reasoning about games is what do the atomic propositions represent? Given a (strategic or extensive) game G, a **strategy profile** is a sequence of strategy choices for each player. Let S be the set of strategy profiles (recall that S_i denotes the strategies for player i and S_{-i} for the set of tuples of strategies for all players except i, see Section 2.1 for more information). The set S represents the possible outcomes in a game G. Each state in a game model will be associated with an outcome in S via a map σ (so for a state w, $\nu(w)$ is the outcome of S realized at state w). The atomic propositions are intended to describe different aspects of the different outcomes of a game. For example, they could describe the specific action chosen by a player or the utility assigned to the outcome by a given player. I will abuse notation and use the same symbol σ to represent the valuation function assigning states to atomic propositions (so for an atomic proposition At, $\sigma(p)$ is the set of states where p is true). Another convention I will follow is to use lower case letters to represent strategies in a game and the corresponding uppercase letter for the (atomic) proposition expressing that the relevant player choose that action. So if i has a strategy $u \in S_i$, then U is intended to mean that player i is playing action u.

Epistemic models I start with models that have been extensively studied by philosophical logicians [17], computer scientists [26] and game theorists [25].

Definition 1 (Epistemic Model). *An* **epistemic model** *for a game G with players N is a tuple $\langle W, \{\sim_i\}_{i\in N}, \sigma\rangle$ where $W \neq \emptyset$, for each $i \in N$, $\sim_i \subseteq W \times W$ is an equivalence relation (each $\sim_i$ is* reflexive*: for each $w \in W$, $w \sim_i w$;* transitive*: for each $w, v, u \in W$, if $w \sim_i v$ and $v \sim_i u$ then $w \sim_i u$; and* Euclidean*: for each $w, v, u \in W$, if $w \sim_i v$ and $w \sim_i u$, then $v \sim_i u$), and σ is the outcome map.*

The intended reading of $w \sim_i v$ is that "at state w, agent i cannot *rule-out* state v (according to i's current observations and any other evidence she has about the game situation)." Alternatively, I can say that "agent i does not have enough information to distinguish state w from state v."

A simple propositional modal language is often used to describe the agent's knowledge at states in an epistemic model. Formally, let $\mathcal{L}_K$ be the (smallest) set of sentences generated by the following grammar:

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi$$

where $p \in \mathsf{At}$ (the set of atomic propositions). The additional propositional connectives ($\rightarrow, \leftrightarrow, \vee$) are defined as usual and the dual of K, denoted L, is defined as follows: $L_i\varphi \ := \neg K_i \neg\varphi$. The intended interpretation of $K_i\varphi$ is that "given all of the available evidence and everything i has observed, agent i is informed that φ is true". Alternatively, following the standard usage in the epistemic logic and game theory literature, I can say "agent i knows that φ".

Truth is defined in a standard way. The main idea behind the definition of kowledge is very simple: a player knows that some proposition φ is true when φ is true in all the possibilities that the player has not yet "ruled-out".

Definition 2 (Truth for $\mathcal{L}_K$). *Let $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, \sigma\rangle$ be an epistemic model for a game G. For each $w \in W$, φ* **is true at state** *w, denoted $\mathcal{M}, w \models \varphi$, is defined by induction on the structure of φ:*

- $\mathcal{M}, w \models p$ *iff* $w \in \sigma(p)$
- $\mathcal{M}, w \models \neg\varphi$ *iff* $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ *iff* $\mathcal{M}, w \models \varphi$ *and* $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models K_i\varphi$ *iff for all* $v \in W$, *if* $w \sim_i v$ *then* $\mathcal{M}, v \models \varphi$ $\lhd$

Epistemic models represents the informational context of a given game in terms of possible configurations of states of the game and the hard information that the agents have about them. The function σ assigns to each possible world a unique state of the game in which every ground, i.e. non-informational fact is either true or false. If $\sigma(w) = \sigma(w')$ then the two worlds w, w' will agree on all the ground facts (i.e., what actions the players will choose) but, crucially, the agents may have different information in them. So, elements of W are *richer*, than the elements of S.

Epistemic-Plausibility Models Originally used as a semantics for counterfactuals (cf. [28]), *plausibility models* have been extensively used by logicians [29, 17, 30], game theorists [31] and computer scientists [32, 33] to represent rational agents' (all-out) beliefs. Thus, I take for granted that they provide natural models of (multiagent) beliefs and focus on how they can be used to represent "rational deliberation" in a game situation.

The main idea is to endow epistemic ranges with a *plausibility ordering* for each agent: a pre-order (reflexive and transitive) $w \preceq_i v$ that says "agent i considers world w at least as plausible as v." As a convenient notation, for $X \subseteq W$, I set $Min_{\preceq_i}(X) = \{v \in X \mid v \preceq_i w \text{ for all } w \in X\}$, the set of minimal elements of X according to $\preceq_i$. This is the subset of X that agent i considers the "most plausible". Thus, while the $\sim_i$ partitions the set of possible worlds according to i's "hard information", the plausibility ordering $\preceq_i$ represents which of the possible worlds agent i considers more likely (i.e., it represents i's "soft information").

Definition 3 (Epistemic-Plausibility Models). *An* **epistemic-plausibility model** *for a game G with players N is a tuple $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, \{\preceq_i\}_{i\in N}, \sigma\rangle$ where $\langle W, \{\sim_i\}_{i\in N}, \sigma\rangle$ is an epistemic model for G and, for each $i \in N$, $\preceq_i$ is a well-founded*[4] *reflexive and transitive relation on W satisfying, for all $w, v \in W$:*

1. plausibility implies possibility*: if $w \preceq_i v$ then $w \sim_i v$.*
2. locally-connected*: if $w \sim_i v$ then either $w \preceq_i v$ or $v \preceq_i w$.*

Remark 1. Note that if $w \not\sim_i v$ then, since $\sim_i$ is symmetric, I also have $v \not\sim_i w$, and so by property 1, $w \not\succeq_i v$ and $v \not\succeq_i w$. Thus, I have the following equivalence: $w \sim_i v$ iff $w \succeq_i v$ or $v \succeq_i w$. In what follows, unless otherwise stated, I will assume that $\sim_i$ is defined as follows: $w \sim_i v$ iff $w \preceq_i v$ or $v \preceq_i w$.

[4] Well-foundedness is only needed to ensure that for any set X, $Min_{\preceq_i}(X)$ is nonempty. This is important only when W is infinite – and there are ways around this in current logics. Moreover, the condition of connectedness can also be lifted, but I use it here for convenience.

One modal language that can be interpreted on these models extends the basic epistemic language $\mathcal{L}_K$ with a conditional belief operator: Let $\mathcal{L}_{KB}$ be the (smallest) set of sentences generated by the following grammar:

$$\varphi \ := \ p \mid \neg\varphi \mid \varphi \wedge \psi \mid B_i^{\varphi}\psi \mid [\preceq_i]\varphi \mid K_i\varphi$$

where $p \in \mathsf{At}$ (the set of atomic propositions) and $i \in N$. The same conventions apply as above with the additional convention that I write $B_i\varphi$ for $B^{\top}\varphi$. The intended interpretation of $B_i^{\varphi}\psi$ is "agent i believes ψ under the supposition that φ is true" and $[\preceq_i]\varphi$ is "the agent *robustly* believes that φ is true".

Let $[w]_i$ be the equivalence class of w under $\sim_i$. Then local connectedness implies that $\preceq_i$ totally orders $[w]_i$ and well-foundedness implies that $Min_{\succeq_i}([w]_i \cap X)$ is nonempty if $[w]_i \cap X \neq \emptyset$.

Definition 4 (Truth for $\mathcal{L}_{KB}$). *Given an epistemic-plausibility model $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, \{\preceq_i\}_{i\in N}, \sigma\rangle$. The definition of truth for formulas from $\mathcal{L}_K$ is given in Definition 2. The additional belief operators are defined as follows:*

- $\mathcal{M}, w \models B_i^{\varphi}\psi$ *iff for all* $v \in Min_{\preceq_i}([w]_i \cap [\![\varphi]\!]_{\mathcal{M}})$, $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models [\preceq_i]\varphi$ *iff for all* $v \in W$ *if* $v \preceq_i w$ *then* $\mathcal{M}, v \models \varphi$ $\triangleleft$

Thus, ψ is believed conditional on φ, if i's most plausible φ-worlds (i.e., the states satisfying φ that i has not ruled out and considers most plausible) all satisfy ψ. Then, the definition of plain belief (which is defined to be $B^{\top}$) is:

$$\mathcal{M}, w \models B_i\varphi \text{ iff for each } v \in Min_{\preceq_i}([w]_i),\ \mathcal{M}, v \models \varphi$$

To illustrate this definition, consider the following coordination game:

		Bob l	Bob r
Ann	u	$3,3$	$0,0$
Ann	d	$0,0$	$1,1$

The epistemic-plausibility model below describes a possible configuration of *ex ante* beliefs of the players (i.e., before the players have settled on what strategy they will choose): I draw an i-labeled arrow from v to w if $w \preceq_i v$ (to keep the clutter down, I do not include all arrows. The remaining arrows can be inferred by transitivity).

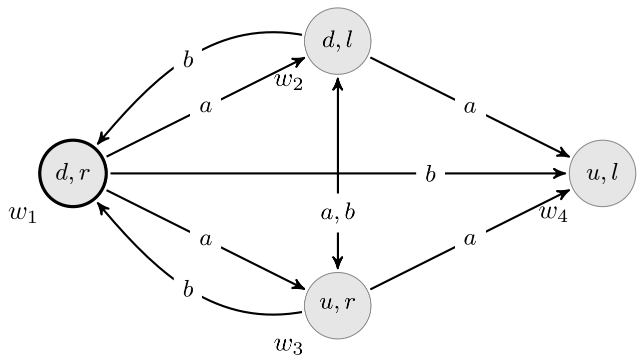

Following the convention discussed in Remark 1, we have $[w_1]_a = [w_1]_b = \{w_1, w_2, w_3, w_4\}$, and so, neither Ann nor Bob knows how the game will end. Furthermore, both Ann and Bob believe that they will coordinate with Ann choosing u and Bob choosing l ($B_a(U \wedge L) \wedge B_b(U \wedge L)$ is true at all the states where U is the proposition expressing Ann chose u, similarly for L and l). However, Ann and Bob do have different *conditional* beliefs. Ann believes that their choices are independent; and so, she believes that L is true even under the supposition that D is true (i.e., she continues to believe Bob will play l even if she decides to play d). On the other hand, Bob believes that their choices are somehow correlated; and so, under the supposition that R is true, Bob believes that Ann will choose d. Conditional beliefs describe an agent's *disposition* to change her beliefs in the presence of (perhaps surprising) evidence (cf. [34]).

Epistemic-Probability Models The logical frameworks introduced in the previous section represent an agent's "full" beliefs. Graded beliefs have also been subjected to sophisticated logical analyses (see, for example, [35–39]).

The dominant approach to formalizing graded beliefs is (subjective) probability theory. A **probability measure** on a set W is a function assigning a positive real number to (some) subsets of W such that $\pi(W) = 1$ and for disjoint subsets $E, F \subseteq W$ (i.e., $E \cap F = \emptyset$) $\pi(E \cup F) = \pi(E) + \pi(F)$. For simplicity, I assume in this section that W is finite. Then, the definition of a probability measure can be simplified: a probability measure on a finite set W is a function $\pi : W \to [0, 1]$ such that for each $E \subseteq W$, $\pi(E) = \sum_{w \in E} \pi(w)$ and $\pi(W) = 1$. Nothing that follows hinges on the assumption that W is finite, but if W is infinite, then there are a number of important mathematical details that add some complexity to the forthcoming definitions.[5] **Conditional probability** is defined in the usual way: $\pi(E \mid F) = \frac{\pi(E \cap F)}{\pi(F)}$ if $\pi(F) > 0$ (if $\pi(F) = 0$, then $\pi(E \mid F)$ is undefined).

The model I study in this section is very close to the epistemic plausibility model of Definition 3 with probability measures in place of plausibility orderings:

Definition 5 (Epistemic-Probability Model). *An* **epistemic probability model** *for a game G with agents N is a tuple $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, \{\pi_i\}_{i \in N}, \sigma\rangle$*

[5] See [40, Chapter 1] for details and pointers to the relevant literature.

where $\langle W, \{\sim_i\}_{i\in N}, \sigma\rangle$ *is an epistemic model and, for each* $i \in N$, π_i *is a probability measure on* W*. I also assume that each* π_i *is* weakly regular[6] *in the sense that for each* $w \in W$, $\pi_i([w]_i) > 0$.

The probability measures π_i represent agent i's prior beliefs about the likelihood of each element of W. Agents then receive private information, represented by the equivalence relations $\sim_i$, and update their initial beliefs with that information. A variety of modal languages have been proposed to reason about graded beliefs. In this section, I focus on a very simple language containing a knowledge modality $K_i\varphi$ ("i is informed that φ is true") and $B_i^q\varphi$ ("i believes φ is true to degree at least q", or "i's degree of belief in φ is at least q") where q is a rational number. More formally, let $\mathcal{L}_{KB}^{prob}$ be the smallest set of formulas generated by the following grammar:

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid B_i^q\varphi \mid K_i\varphi$$

where $q \in \mathbb{Q}$ (the set of rational numbers), $i \in N$ and $p \in \mathsf{At}$.

Definition 6 (Truth for $\mathcal{L}_{KB}^{prob}$). *Suppose that* $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, \{\pi_i\}_{i\in N}, \sigma\rangle$ *is an epistemic-probaiblity model. The definition of truth for formulas from* $\mathcal{L}_K$ *is given in Definition 2. The belief operator is defined as follows:*

– $\mathcal{M}, w \models B_i^q\varphi$ *iff* $\pi_i([\![\varphi]\!]_{\mathcal{M}} \mid [w]_i) \geq q$

where $[\![\varphi]\!]_{\mathcal{M}} = \{w \mid \mathcal{M}, w \models \varphi\}$.

Note that since I assume for each $w \in W$, $\pi_i([w]_i) > 0$, the above definition is always well-defined.

2.3 Hierarchies of beliefs

States in an epistemic (-plausibility/-probability) model not only represent the players beliefs about what their opponents will do, but also their *higher-order* beliefs about what their opponents are thinking. This means that outcomes identified as "rational" in a particular informational context will depend, in part, on the these higher-order beliefs. This raises the following issues:

What role do the higher-order beliefs play in a general theory of rational decision making in game situations? In the end, I are only interested in what (rational) players are going to do. This, in turn, depends only on the agents beliefs about what her opponents are going to do. The players' beliefs about what her opponents are thinking are only relevant because they shape the players first-order beliefs about what her opponents are going to do. Kadane and Larkey nicely explain the issue as follows:

[6] A probability measure is **regular** provided $\pi(E) > 0$ for each (measurable) set E.

> "It is true that a subjective Bayesian will have an opinion not only on his opponent's behavior, but also on his opponent's belief about his own behavior, his opponent's belief about his belief about his opponent's behavior, etc. (He also has opinions about the phase of the moon, tomorrow's weather and the winner of the next Superbowl). *However, in a single-play game, all aspects of his opinion except his opinion about his opponent's behavior are irrelevant, and can be ignored in the analysis by integrating them out of the joint opinion.*" [41, pg. 239, my emphasis]

How much higher-order information should *the players take into account?* The well-known *email game* of Ariel Rubinstein [42] demonstrates that misspecification of arbitrarily high levels of belief can have a great impact on (predicted) strategic behavior. This is an example where (predicted) strategic behavior is *too sensitive* to the players' higher-order beliefs. A theory of rational decision making in game situations need not *require* that the players consider *all* of her higher-order beliefs in her decision making process. The assumption is only that the players recognize that their opponents are "actively reasoning" agents. Precisely "how much" higher-order information *should* be taken into account in such a situation is a very interesting, open question (cf. [43], and Perea's contribution to this volume).

How much higher-order information do *the players take into account?* There is quite a lot of experimental work about whether or not humans take into account even second-order beliefs (eg., belief about her opponents beliefs) in game situations (see, for example, [44, 45]). This is related to a number of issues in cognitive psychology falling under the broad category *the theory of mind* (a famous experiment here is the Wimmer and Perner false-belief task [46]). Of course, this is a *descriptive* question and it is very much open how such observations should be incorporated into a general theory of rational decision making (cf. [47]).

There are no easy answers to the above questions. The lesson I want to draw from the above discussion is that it is important to understand what exactly the higher-order beliefs in the models discussed in the previous section are intended to represent. There are three positions that one can take here: In an epistemic (-plausibility/-probability) model, the higher-order beliefs

1. are an explicit description (perhaps overly precise) of the contents of the players thoughts about her opponents in a game situation;
2. represent the *outcome* of a reasoning process, i.e., the *reasons* rational players can point to in order to justify their choices; or
3. *track* the back-and-forth reasoning that players are engaged in as they deliberate about what to do.

2.4 Common knowledge of rationality

Both game theorists and logicians have extensively discussed different notions of knowledge and belief for a group, such as common knowledge and belief. These notions have played a fundamental role in the analysis of distributed algorithms [48] and social interactions [49]. In this section, I briefly recount the standard definition of common knowledge. [7]

Consider the statement "everyone in group X knows that φ". With finitely many agents, this can be easily defined in the epistemic language $\mathcal{L}_{KB}$:

$$K_X\varphi \quad := \quad \bigwedge_{i\in X} K_i\varphi$$

where $X \subseteq N$. The first nontrivial informational attitude for a group that I study is *common knowledge*. If φ is common knowledge for the group G, then not only does everyone in the group know that φ is true, but this fact is completely transparent to all members of the group. Following [51], the idea is to define common knowledge of φ as the following iteration of everyone knows operators:

$$\varphi \wedge K_N\varphi \wedge K_NK_N\varphi \wedge K_NK_NK_N\varphi \wedge \cdots$$

The above formula is an *infinite* conjunction, and so is not a formula in our epistemic language $\mathcal{L}_{KB}$ (by definition, there can be at most finitely many conjunctions in any formula). In order to express this, I must extend our basic epistemic language with a modal operator $C_G\varphi$ with the intended meaning "φ is common knowledge among the group G". Formally,

Definition 7 (Interpretation of C_G). *Let $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, V\rangle$ be an epistemic model*[8] *and $w \in W$. The truth of formulas of the form $C_X\varphi$ is:*

$$\mathcal{M}, w \models C_X\varphi \quad \textit{iff} \quad \textit{for all } v \in W, \textit{ if } wR^C_Xv \textit{ then } \mathcal{M}, v \models \varphi$$

where $R^C_X := (\bigcup_{i\in X} \sim_i)^*$ is the reflexive transitive closure[9] *of $\bigcup_{i\in X} \sim_i$.*

It is well-known that for any relation R on W, if wR^*v then there is a finite R-path starting at w ending in v. Thus, I have $\mathcal{M}, w \models C_X\varphi$ iff every finite path for X from w ends with a state satisfying φ.

There is an alternative characterization of common knowledge: Call a subset $X \subseteq W$ ***i*-closed** provided for all $w \in W$, if $w \in X$ and $w \sim_i v$ then $v \in X$. If X is i-closed then it is "self-evident" for agent i in the sense that if X obtains (i.e., the current state is in X) then i knows that X obtains (formally, I have $X \subseteq \{w \mid [w]_i \subseteq X\}$). Then,

---

[7] I assume that the formal definition of common knowledge is well-known to the reader. For more information and pointers to the relevant literature, see [50] and [26].

[8] The same definition will of course hold for epistemic-plausibility and epistemic-probability models.

[9] The reflexive transitive closure of a relation R is the smallest relations R^* containing R that is reflexive and transitive.

Fact 1 *Suppose that $\mathcal{M}$ is an epistemic (-plausibility/-probability) model. Then, $\mathcal{M}, w \models C_G\varphi$ iff there is a set $X \subseteq W$ that is i-closed for all $i \in G$ and $X \subseteq [\![\varphi]\!]_{\mathcal{M}}$.*

It is sometimes convenient to use the above characterization common knowledge as the definition of C_G. On epistemic (-plausibility/-probability) models, the two definitions are equivalent.

The approach to defining common knowledge outlined above can be viewed as a recipe for defining common (robust) belief. For example, suppose wR_i^Bv iff $v \in Min_{\preceq_i}([w]_i)$ and define R_G^B to be the transitive closure[10] of $\cup_{i\in G}R_i^B$. Then, **common belief of** φ, denoted $C_G^B\varphi$, is defined in the usual way:

$$\mathcal{M}, w \models C_G^B\varphi \text{ iff for each } v \in W, \text{ if } wR_G^Bv \text{ then } \mathcal{M}, v \models \varphi.$$

A probabilistic variant of common belief was introduced by [52]. It is convenient to give the definition in terms of "self-evident sets".

Definition 8 (Common q-belief). *Call a set $X \subseteq W$ an* **evident q-belief for** *i when $X \subseteq \{w \mid \pi_i(X \mid [w]) \geq q\}$. Then,* **common q-belief for a group** *G, denoted $C_G^q\varphi$, is defined as follows: Let $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, \{\pi_i\}_{i\in N}, \sigma\rangle$ be an epistemic-probability model.*

– *$\mathcal{M}, w \models C_G^q\varphi$ iff there is a set X such that $w \in X$, X is an evident q-belief for all $i \in G$ and $X \subseteq [\![B_i^q\varphi]\!]_{\mathcal{M}}$ for all $i \in G$.*

Common knowledge (and the above variants, common belief and common p-belief) are typically used to make precise the informal game-theoretic assumption that there is "common knowledge of rationality". Rationality here is understand in the decision-theoretic sense: the agents' choices are *optimal* according to some decision-theoretic rule (such as maximizing expected utility if probabilities are available and strict/weak dominance if probabilities are not available). A general theme of this paper is that there is *more* to this basic assumption than the fact that the event where it is common knowledge that all players choose optimally has obtained:

> "It is not just a question of what common knowledge obtains at the moment of truth, but also how common knowledge is preserved, created, or destroyed in the deliberational process which leads up to the moment of truth." [6, pg. 160]

If common knowledge of rationality is to have an "explanatory" role in the analysis of a game-theoretic situation, then it is not enough to simply assume that it has obtained in an informational context. It is also important to describe how the players were able to come to the conclusion that there is common knowledge of rationality. This general point about common knowledge was already appreciated by Lewis when he first formulated his notion of common knowledge [53]. See [8] for an illuminating discussion and a reconstruction of Lewis' notion of common knowledge, with applications to game theory.

[10] Since beliefs need not be factive, I do not force R_G^B to be reflexive.

2.5 Modeling information changes

The simplest type of informational change treats the source of the information as *infallible*. The effect of finding out that φ is true from an infallible source should be clear: *Remove* all states that do not satisfy φ. In the epistemic logic literature this operation is called a *public announcement* [54, 55]. However, calling this an "announcement" is misleading since in this paper I am not modeling any form of "pre-play" communication. The "announcements" are formulas that the players incorporate into the current epistemic state.

Definition 9 (Public Announcement). *Suppose that $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, V\rangle$ is an epistemic model and φ is a formula (in the language $\mathcal{L}_K$). After all the agents find out that φ is true (i.e., φ is* **publicly announced***), the resulting model is $\mathcal{M}^{!\varphi} = \langle W^{!\varphi}, \{\sim_i^{!\varphi}\}, V^{!\varphi}\rangle$ where $W^{!\varphi} = \{w \in W \mid \mathcal{M}, w \models \varphi\}$, $\sim_i^{!\varphi} = \sim_i \cap W^{!\varphi} \times W^{!\varphi}$ for all $i \in N$, and for all $p \in \mathsf{At}$, $V^{!\varphi}(p) = V(p) \cap W^{!\varphi}$.*

The same definition applies *mutatis mutandis* to other models introduced in the previous section[11]. The models $\mathcal{M}$ and $\mathcal{M}^{\varphi}$ describe two different moments in time, with $\mathcal{M}$ describing the current or initial information state of the agents and $\mathcal{M}^{!\varphi}$ the information state *after* the all the agents find out that φ is true. This temporal dimension can also be represented in the logical language with modalities of the form $[!\varphi]\psi$. The intended interpretation of $[!\varphi]\psi$ is "ψ is true after all the agents find out that φ is true", and truth is defined as

– $\mathcal{M}, w \models [!\varphi]\psi$ iff if $\mathcal{M}, w \models \varphi$ then $\mathcal{M}^{!\varphi}, w \models \psi$.

A public announcement is only one type of informative action. For the other transformations discussed in this paper, while the agents do *trust* the source of φ, they do not treat the source as infallible. Perhaps the most ubiquitous policy is *conservative upgrade* ($\uparrow\varphi$), which lets the agent only tentatively accept the incoming information φ by making the best φ-worlds the new minimal set and keeping the old plausibility ordering the same on all other worlds. A second, stronger, operation is *radical upgrade* ($\Uparrow\varphi$) which moves *all* φ worlds before all the $\neg\varphi$ worlds and otherwise keeps the plausibility ordering the same. Before giving the formal definition we need some notation: Given an epistemic-plausibility model $\mathcal{M}$, let $[\![\varphi]\!]_i^w = \{x \mid \mathcal{M}, x \models \varphi\} \cap [w]_i$ denote the set of all φ-worlds that i considers possible and $best_i(\varphi, w) = Min_{\preceq_i}([w]_i \cap \{x \mid \mathcal{M}, x \models \varphi\})$ the best φ-worlds at state w according to agent i.

Definition 10 (Conservative and Radical Upgrade). *Given an epistemic-plausibility model $\mathcal{M} = \langle W, \{\sim_i\}_{i\in N}, \{\preceq_i\}_{i\in N}, \sigma\rangle$ and a formula $\varphi \in \mathcal{L}_{KB}$, the conservative/radical upgrade of $\mathcal{M}$ with φ is the model $\mathcal{M}^{*\varphi} = \langle W^{*\varphi}, \{\sim_i^{*\varphi}\}_{i\in N}, \{\preceq_i^{*\varphi}\}_{i\in N}, V^{*\varphi}\rangle$ with $W^{*\varphi} = W$, for each i, $\sim_i^{*\varphi} = \sim_i$, $V^{*\varphi} = V$ where $* = \uparrow, \Uparrow$. The relations $\preceq_i^{\uparrow\varphi}$ and $\preceq_i^{\Uparrow\varphi}$ are the smallest relations satisfying:*

[11] Of course, the probability measures need to be renormalized.

Conservative Upgrade

1. *If* $v \in best_i(\varphi, w)$ *then* $v \prec_i^{\uparrow\varphi} x$ *for all* $x \in [w]_i$*, and*
2. *for all* $x, y \in [w]_i - best_i(\varphi, w)$*,* $x \preceq_i^{\uparrow\varphi} y$ *iff* $x \preceq_i y$*.*

Radical Upgrade

1. *for all* $x \in [\![\varphi]\!]_i^w$ *and* $y \in [\![\neg\varphi]\!]_i^w$*, set* $x \prec_i^{\Uparrow\varphi} y$*,*
2. *for all* $x, y \in [\![\varphi]\!]_i^w$*, set* $x \preceq_i^{\Uparrow\varphi} y$ *iff* $x \preceq_i y$*, and*
3. *for all* $x, y \in [\![\neg\varphi]\!]_i^w$*, set* $x \preceq_i^{\Uparrow\varphi} y$ *iff* $x \preceq_i y$*.*

As the reader is invited to check, a conservative upgrade is a special case of a radical upgrade: the conservative upgrade of φ at w is the radical upgrade of $best_i(\varphi, w)$. A logical analysis of these operations include formulas of the from $[\uparrow\varphi]\psi$ intended to mean "after everyone conservatively upgrades with φ, ψ is true" and $[\Uparrow\varphi]\psi$ intended to mean "after everyone radically upgrades with φ, ψ is true" . The definition of truth for these formula is as expected:

– $\mathcal{M}, w \models [\uparrow\varphi]\psi$ iff $\mathcal{M}^{\uparrow\varphi}, w \models \psi$
– $\mathcal{M}, w \models [\Uparrow\varphi]\psi$ iff $\mathcal{M}^{\Uparrow\varphi}, w \models \psi$

Note that unlike with public announcements, there is no precondition for these operations.

2.6 Information change over time

I leave aside games for the moment, and concentrate on the dynamic of information change over time. The general idea is to use single update steps (eg., public announcement, radical or conservative upgrade) and iterate them to according to some predefined "protocol" (describing the sequence of formulas that are presented to the agent(s)). The problem of *iterated revision* has been extensively studied [56–58], and although there are many proposals, there still remain a number of conceptual problems (see [59] for a discussion).

One key issue is this: Suppose that the agent receives a sequence of consistent formulas and uses, for example, radical upgrade to adjust her plausibility orderings. Since the information is consistent, no matter what the order in which she incorporates the information, she will always end up with the same beliefs. However, the different orders can lead to very different *conditional* beliefs, and this, in turn, means that there could be drastic differences in the result of incorporating information that contradicts one of the previous pieces of information.

The main issue that interests use for this paper is the limit behavior of iterated sequence of announcements. That is, what happens to the epistemic -plausibility/-probability models *in the limit*? Do the players' knowledge and beliefs stabilize or keep changing in response to the new information?

An initial observation is that iterated public announcement of any formula φ is a (epistemic -plausibility/-probability) model must stop at a limit model where either φ or its negation is true at all states (see [60] for a discussion and proof). In addition to the limit dynamics of knowledge under public announcements, there is the limit behavior of beliefs under soft announcements (radical/conservative upgrades). See [60] and [14, Section 4] for general discussions. The first observation is that there need not be convergence at all as an update process can oscillate forever (the proofs are repeated here for completeness):

Proposition 1 (Baltag and Smets [60]). *The repeated iteration of a true formula under conservative upgrade need not stabilize.*

Proof. Consider an initial model $\mathcal{M}_1$ with three states w_1, w_2 and w_3 satisfying r, q and p, respectively. Suppose that the agent's plausibility ordering is $w_1 \prec w_2 \prec w_3$. Then the agent believes that r. Consider the formula

$$\varphi := \ p \vee (r \wedge \neg Br) \vee (\neg r \wedge Br)$$

This is true at w_3 in the initial model. Since $\llbracket\varphi\rrbracket^{\mathcal{M}_1} = \{w_3, w_2\}$, we have $\mathcal{M}_1^{\uparrow\varphi} = \mathcal{M}_2$. Furthermore, $\llbracket\varphi\rrbracket^{\mathcal{M}_2} = \{w_3, w_1\}$, so $\mathcal{M}_2^{\uparrow\varphi} = \mathcal{M}_3$. Since, $\mathcal{M}_3$ is the same model as $\mathcal{M}_1$, we have a cycle:

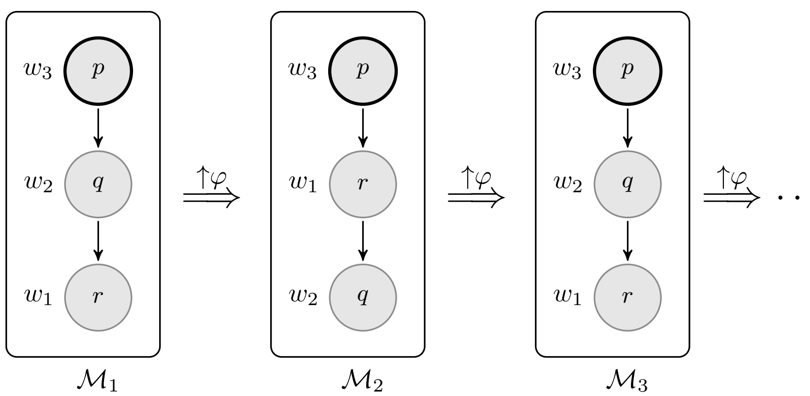

So, repeated conservative upgrades of *true* formulas need not stabilize. A similar phenomena occurs when the agents use radical upgrade to change the model:

Proposition 2 (Baltag and Smets [60]). *The repeated iteration of a true formula under radical upgrade need not stabilize.*

Proof. Consider an initial model $\mathcal{M}_1$ with three states w_1, w_2 and w_3 satisfying r, q and p, respectively. Suppose that the agent's plausibility ordering is $w_1 \prec w_2 \prec w_3$. Then the agent believes that r. Consider the formula

$$\varphi := \ (r \vee (B^{\neg r} q \wedge p) \vee (B^{\neg r} p \wedge q))$$

This is true at w_1 in the initial model. Since $\llbracket\varphi\rrbracket^{\mathcal{M}_1} = \{w_3, w_1\}$, we have $\mathcal{M}_1^{\Uparrow\varphi} = \mathcal{M}_2$. Furthermore, $\llbracket\varphi\rrbracket^{\mathcal{M}_2} = \{w_2, w_1\}$, so $\mathcal{M}_2^{\Uparrow\varphi} = \mathcal{M}_3$. Since, $\mathcal{M}_3$ is the same model as $\mathcal{M}_1$, we have a cycle:

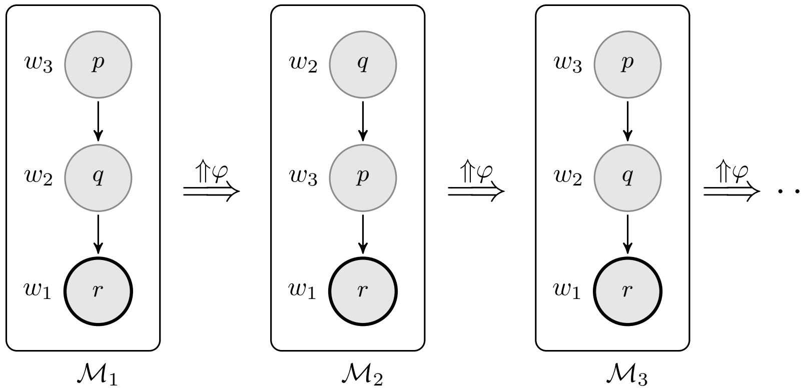

In the above example, the player's conditional beliefs keep changing during the update process. However, the player's first-order beliefs remain fixed throughout the process. In fact, the following is true (consult the cited paper for a proof):

Theorem 2 (Baltag and Smets [60]). *Every iterated sequence of truthful radical upgrades stabilizes all non-conditional beliefs in the limit.*

See [61, 16] for generalizations and broader discussions about the interest of the above theorem.

3 Constraints on game models

The perspective put forward in this paper is that the game models describe the beliefs that the players acquire after a process of "rational" deliberation. This suggests additional constraints on a game model. One natural constraint is that the deliberation was successful in the sense that each player comes to know the choice that they have made:

> **Knowledge of own action**: Players know their own actions in an epistemic (-plausibility/-probability) model provided for all states w and v, if $w \sim_i v$ then $(\sigma(w))_i = (\sigma(v))_i$.

Another natural constraint is that players cannot *rule-out* the fact that their opponents will choose an optimal outcome[12] (in general, there will be more than one "choice-worthy" option for a player). This property was called "privacy of tie breaking" by Cubitt and Sugden [9, pg. 8] and "no extraneous restrictions on beliefs" by Geir Asheim and Martin Dufwenberg [62]. Wlodeck Rabinovich [63] takes this idea even further and argues that from the principle of indifference, players must assign equal probability to all choice-worthy options.

[12] That is, if, according to player i's beliefs, strategy s is optimal for player j, then i cannot rule-out *all* states where player j follows strategy s.

There are a number of ways to make the above idea precise. I need some notation. First, given any probability measure over the strategy profiles $S = \Pi_{i \in N} S_i$, the *expected utility* of action $a \in S_i$ with respect to $\pi \in \Delta(S)$ is

$$E_\pi(a) = \sum_{s_{-i} \in S_{-i}} \pi(a, s_{-i}) \cdot u_i(a, s_{-i})$$

Let $\mathcal{M} = \langle W, \{\sim_i\}_{i \in N}, \{\pi_i\}_{i \in N}, \sigma \rangle$ be an epistemic-model for a game $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$. Each probability measure π_i, generates a probability measure over the strategy profiles as follows: for each $s \in S$, $\widehat{\pi_i}(s) = \pi_i(\sigma^{-1}(s))$. Using this machinery, I can define what it means for an action $a \in S_i$ to be **optimal for agent** i. Formally, for each $a \in S_i$ define the proposition $Opt_i(a)$ as follows:

$$\mathcal{M}, w \models Opt_i(a) \text{ iff for all } a' \in S_i,\ E_{\widehat{\pi_i}(\cdot \mid [w]_i)}(a) \geq E_{\widehat{\pi_i}(\cdot \mid [w]_i)}(a')$$

I can now formally state two constraints on epistemic-probability models. These are natural assumptions to impose if we assume that the epistemic-probability models describe the players knowledge and beliefs as the players deliberate about what to do in a game situation. The first is:

> **Regularity**: For all $i \in N$ and $w \in W$, $\pi_i(w) > 0$.

The idea is that the first step of a reasoning process is that each player i rules-out all the states that are initially considered *impossible* (i.e., assigned probability 0 by π_i). The second is the "privacy of tie breaking" property discussed above:

> **Privacy of tie-breaking**: If $\mathcal{M}, w \models Opt_i(a)$ then for all $j \neq i$, there exists a $v \in [w]_j$ such that $\sigma_i(v) = a$.

The explanation of this property is given above. Cubitt and Sugden show that there is a conflict between these two properties. Consider the following three-person game:

in_3:

	in_2	out_2
in_1	$1, 1, 1$	$1, 1, 1$
out_1	$1, 1, 1$	$0, 1, 1$

out_3:

	in_2	out_2
in_1	$1, 1, 1$	$1, 0, 1$
out_1	$1, 1, 0$	$0, 0, 0$

Proposition 3 (Cubitt and Sugden [64, 9]). *There is no model of the above game satisfying privacy of tie-breaking, regularity knowledge of own choice where there is common knowledge that the players choose optimally.*

I give a high-level sketch of the proof (a more formal proof can be found in [9][13]). Suppose that player 1 considers out_2 possible. Since, in_1 weakly dominates out_1, then in_1 is the only rational choice for player 1 (this follows from Lemma 2). Thus, both players 2 and 3 know that out_1 is impossible. Since player 3 thinks out_1 is impossible, both in_3 and out_3 are rational choices. Again, players 1 and 2 know this, and so by privacy of tie breaking must consider both outcomes possible. Since player 2 considers both in_3 and out_3 possible and in_2 weakly dominates out_2, in_2 must be the only rational choice for player 2. But, players 1 and 3 know this, and so 1 cannot consider out_2 possible.

Suppose 1 does not consider out_2 possible. This means that both in_1 and out_1 are rational choices for player 1. By privacy of tie-breaking, both players 2 and 3 consider these outcomes possible. Since in_3 weakly dominates out_3, in_3 is the only rational choice for player 3. Players 1 and 2 know this fact, and so do not consider out_3 possible. Since player 2 does not consider out_3 possible, both in_2 and out_2 are rational choice. By privacy of tie breaking, players 1 and 3 must consider these outcomes possible. In particular, player 1 considers out_2 possible, which is a contradiction.

4 Reasoning to a solution

The previous section points to an interesting tension between assuming that the players reason "strategically" and the standard assumption that there is common knowledge that the players choose optimally. But, what about describing the players' "reasoning process" as a way of "solving" a game? The *reasoning-based expected utility procedure* of [65] is an example of such a solution concept. This procedure is intended to model the reasoning procedure a *Bayesian* rational player would follow as she decides what to do in a game. Let $G = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ be a strategic game. Let $\Delta(X)$ denote the set of probability measures on a set X (we assume X is finite, otherwise assume an appropriate σ-algebra is given). Recall that the expected utility of an action $a \in S_i$ for agent i with respect to some probability $\pi \in \Delta(S_{-i})$ is:

$$E_\pi(a) = \sum_{\sigma_{-i}\in S_{-i}} \pi(\sigma_{-i}) \cdot u_i(a, \sigma_{-i})$$

We say that $a \in S_i$ maximizes expected utility with respect to $\pi \in \Delta(S_{-i})$, denoted $MEU(a, \pi)$, if for all $a' \in S_i$, $E_\pi(a) \geq E_\pi(a')$.

A key notion for this procedure is a **categorization**. A categorization is a *ternary* partition of the players strategies S_i (rather than a binary partition of what is out (or "irrational") and what is in (or "not irrational"). That is, during the reasoning process strategies are accumulated, deleted or neither. Formally, for

[13] Note that the framework used in [9] differs in small but important ways from the epistemic-probability models introduced in this paper. These technical details are not important for the main point I am making here, and, indeed, this argument can be made more formal using epistemic-probability models.

each player i, let $S_i^+ \subseteq S_i$ denote the set of strategies that have been accumulated and $S_i^- \subseteq S_i$ the set of strategies that have been deleted. The innovative aspect of this procedure is that $S_i^+ \cup S - i^-$ need not equal S_i. So, strategies in S_i but not $S_i^+ \cup S_i^-$ are classified as "neither accumulated or deleted". The reasoning-based expected utility procedure proceeds as follows: The procedure is defined by induction. Initially, let $D_{i,0} = \Delta(S_{-i})$, the set of all probability measures over the strategies of i's opponents, and $S_{i,0}^+ = S_{i,0}^- = \emptyset$. Then for $n \geq 1$, we have:

- Accumulate all all strategies for player i that maximize expected utility for *every* probability in D_i. Formally,

$$S_{i,n+1}^+ = \{a \in S_{i,n} \mid MEU(a, \pi) \text{ for all } \pi \in D_{i,n}\}$$

- Delete all strategies that do not maximize probability against *any probability distribution*

$$S_i^- \{a \in S_{i,n} \mid \text{ there is no } \pi \in D_{i,n} \text{ such that } MEU(a, \pi)\}$$

- Keep all probability measures that assign positive probability to opponents playing accumulated strategies and zero probability to deleted strategies. Formally, let $D_{i,n+1}$ be all the probability measures from $D_{i,n}$ that assign positive probability to any strategy profile from $\Pi_{j \neq i} S_{i,n+1}^+$ and 0 probability to any strategy profile from $\Pi_{j \neq i} S_{i,n+1}^-$.

The following example from [65] illustrates this procedure:

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

For Bob, strategy L is accumulated since it maximizes expected utility with respect to every measure on Ann's strategies (note L weakly dominates R). Furthermore, B is deleted as it does not maximize probability with respect to any probability measure on Bob's strategies (note that B is strictly dominated by U). In the next round Ann must consider only probability measures that assign positive probability to Bob playing L and Bob must consider only probability measures assigning probability 0 to B. At the next stage, R is accumulated for Bob as he must consider probability measures that assign 0 to Ann playing B. Similarly, M_1 is deleted for Ann since she must assign positive probability to L. At this point, the procedure stabilizes.

In this example, the procedure completely categorized all strategies for the two players. However, this need not always be true. That is, by reasoning alone the players may not be able to determine whether or not a strategy is deleted or accumulated. To see this, consider the following example:

		Bob	
		l	r
Ann	u	$1,1$	$1,0$
	d	$1,0$	$0,1$

In this game, u is accumulated in the first round since it maximizes expected utility with respect to all probability measures on Bob's strategies. No other strategies are deleted or accumulated. In the second round, nothing changes as Bob can find a probability measure over Ann's choices where r maximizes expected utility. Consult [65] and [9] for an extended discussion of this procedure and how it is related to other well-known solution concepts (such as iteratively removing strictly/weakly dominated strategies).

5 Reasoning to a game model

The game models introduced in Section 2.2 can be used to describe the informational context of a game. A natural question from the perspective of this paper is: How do the players arrive at a particular information context?

We are interested in the operations that transform the informational context *as the players deliberate about what they should do in a game situation.* The main idea is that in each informational context (viewed as describing one stage of the deliberation process), the players determine which options are "*optimal*" and which options the players ought to avoid (which is guided by some choice rule). This leads to a transformation of the informational context as the players adopt the relevant beliefs about the outcome of their *practical reasoning.* The different types of transformation mentioned above then represent how confident the player(s) (or modeler) is (are) in the assessment of which outcomes are rational. In this new informational context, the players again think about what they should do, leading to another transformation. The main question is does this process *stabilize*?

The answer to this question will depend on a number of factors. The general picture is

$$\mathcal{M}_0 \stackrel{\tau(D_0)}{\Longrightarrow} \mathcal{M}_1 \stackrel{\tau(D_1)}{\Longrightarrow} \mathcal{M}_2 \stackrel{\tau(D_2)}{\Longrightarrow} \cdots \stackrel{\tau(D_n)}{\Longrightarrow} \mathcal{M}_{n+1} \Longrightarrow \cdots$$

where each D_i is some proposition describing the "rational" options and τ is a model transformer (eg., public announcement, radical or conservative updgrade). Two questions are important for the analysis of this process. First, what type of transformations are the players using? The second question is where do the propositions D_i come from?

Here is a baseline result from [12]. Consider a propositional formula $\mathsf{WR_i}$ that is intended to mean "i's current action is not strictly dominated in the set of actions that the agent currently considers possible". This is a propositional formula whose valuation changes as the model changes (i.e., as the agent removes possible outcomes that are strictly dominated). An epistemic model is **full** for a game G provided the map σ from states to profiles is onto. That is, all outcomes are initially possible.

Theorem 3 (van Benthem [12]). *The following are equivalent for all states w in an epistemic model that is full for a finite game G:*

1. *The outcome $\sigma(w)$ survives iterated removal of strictly dominated strategies.*
2. *Repeated successive* **public announcements** *of $\bigwedge_i \mathsf{WR_i}$ for the players stabilizes at a submodel whose domain contains w.*

This result was generalize by [66] where they focus on arbitrary "optimality" propositions satisfying a monotonicity property and arbitrary games. A related analysis can be found in [67] which provides an in-depth study of the upgrade mechanisms that match game theoretic analyses.

A key issue in the epistemic foundations of game theory is the epistemic analysis of iterated removal of *weakly* dominated strategies. Many authors have pointed out puzzles surrounding such an analysis [62, 68, 69]. For example, [68] showed (among other things) that "common knowledge of admissibility" may be an inconsistent concept (in the sense that there is a game which does not have a model with a state satisfying "common knowledge of rationality" [68, Example 8, pg. 305]). This is illustrated by the following game:

		Bob	
		l	r
Ann	u	$1, 1$	$1, 0$
	d	$1, 0$	$0, 1$

The key issue is that the assumption that players only play *non-weakly dominated* strategies conflicts with the logic of iteratively removing strategies deemed "irrational". This is closely related to the issue discussed in Section 3. The dynamic analysis of this game from [67] offers a new perspective on this issue.[14] Dynamically, Samuelson's non-existence result corresponds to the fact that a certain upgrade streams does not stabilize. That is, the players are not able to deliberate their way to a stable, common belief in admissibility.

6 Rational belief change during game play

The importance of explicitly modeling belief change over time becomes even more evident when considering *extensive games*. An extensive game makes explicit the

[14] Note that the idea is not to provide a new "epistemic foundation" for the iterated admissibility solution concept. Both [69] and [70] have convincing results here. The goal is to use this solution concept as an illustration of the general dynamic approach to games.

sequential structure of the choices in a game. Formally, an extensive game is a tuple $\langle N, T, \tau, \{u_i\}_{i\in N}\rangle$, where

- N is a finite set of players;
- T is a tree (we assume the reader is familiar with the formal definition of a tree) describing the temporal structure of the game situation: Formally, T consists of a set of nodes partitioned into a set of decision nodes V and outcome nodes O. Let $\rightarrowtail$ denote the immediate successor relation on $V \cup O$ where there are no successors at any outcome node $o \in O$. The edges at a decision node $v \in V$ are each labeled with an action. Let $A(v)$ denote the set of actions available at v;
- τ is a turn function assigning a player to each node $v \in V$ (let $V_i = \{v \in V \mid \tau(v) = i\}$; and
- $u_i : O \to \mathbb{R}$ is the utility function for player i assigning real numbers to outcome nodes.

The following is an example of an extensive game:

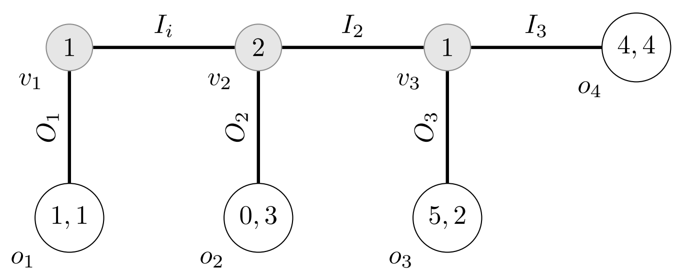

This is an extensive game with $V = \{v_1, v_2, v_3\}$, $O = \{o_1, o_2, o_3, o_4\}$, $\tau(v_1) = \tau(v_3) = Ann$ and $\tau(v_2) = Bob$, and, for example, $u_1(o_2) = 0$ and $u_2(o_2) = 3$. Furthermore, we have, for example, $v_1 \rightarrowtail o_1$ and $v_1 \rightsquigarrow o_4$ with $A(v_1) = \{I_1, O_1\}$.

A **strategy** for player i in an extensive is a function σ from V_i to nodes such that $v \mapsto \sigma(v)$. Thus, a strategy prescribes a move for player i *at every possible node where i moves.* For example, the function σ with $\sigma(v_1) = O_1$ is $\sigma(v_3) = I_3$ is a strategy for player i even though, by following the strategy, i knows that v_3 will not be reached. The main solution concept for extensive games is the *subgame perfect equilibrium* [71], which is calculated using the "backwards induction (BI) algorithm":

BI Algorithm: *At terminal nodes, players already have their values marked. At non-terminal nodes, once all daughters are marked, the player to move gets her maximal value that occurs on a daughter, while the other, non-active player gets his value on that maximal node.*

On the extensive game given above, the BI algorithm proceeds as follows:

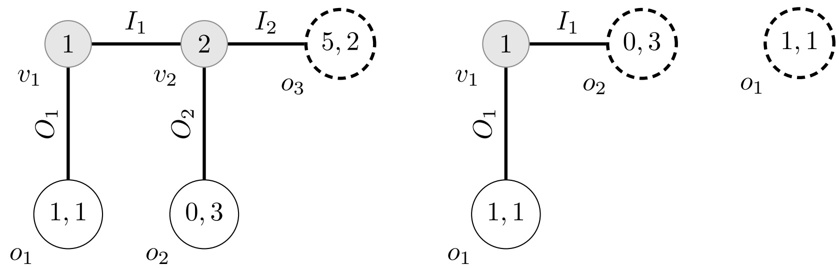

The BI strategy for player 1 is $\sigma(v_1) = O_1$, $\sigma(v_3) = O_3$ and for player 2 is $\sigma(v_2) = O_2$. If both players follow their BI strategy, then the resulting outcome is o_1 ($v_1 \mapsto o_1$ is called the *BI path*).

Suppose that $a, a' \in A(v)$ for some $v \in V_i$. The move a strictly dominates in beliefs move a' provided all of the most plausible nodes reachable by playing a at v are preferred to all the most plausible outcomes reachable by playing a'. Consider an initial epistemic-plausible model consisting of the four outcomes $\{o_1, o_2, o_3, o_4\}$ and both players consider all outcomes equally plausible. Then, at v_2, O_2 is not strictly dominated in beliefs over I_2 since the nodes reachable by I_2 are $\{o_3, o_4\}$, both are equally plausible and player 2 prefers o_4 over o_2, but o_2 over o_3. However, for player 1, since o_3 is preferred to o_4, O_3 strictly dominates in beliefs I_3. Suppose that R is interpreted as no player chooses an action that is strictly dominated in beliefs. Thus, in the initial model model where all four outcomes are equally plausible, the interpretation of R is $\{o_1, o_2, o_3\}$.

$$o_1 \approx o_2 \approx o_3 \approx o_4 \quad \stackrel{\Uparrow R}{\Longrightarrow} \quad o_1 \approx o_2 \approx o_3 \preceq o_4 \quad \stackrel{\Uparrow R}{\Longrightarrow} \quad o_1 \approx o_2 \preceq o_3 \preceq o_4$$

$$\stackrel{\Uparrow R}{\Longleftarrow} \quad o_1 \preceq o_2 \preceq o_3 \preceq o_4$$

This sequence of radical upgrades is intended to represent the "pre-play" deliberation leading to a model where there is common belief that the outcome of the game will be o_1.

The soundness of the deliberation sequence is derived from the assumption that there is common knowledge that the players are "rational" (in the sense, that players will not knowingly choose an option giving them lower payoffs). The basic problem is this: under common knowledge that the players are *rational* (i.e., make the optimal choice when given the chance), player 1 *must* choose O_1 at node v_1. The backwards induction argument for this is based on what the players *would* do if 1 chose I_1. But, if player 1 did in fact choose I_1, then common knowledge of rationality is violated. Thus, the argument does not hold up. Much has been written about this issue and the general relationship between common knowledge of rationality and subgame perfect equilibrium. I do not discuss this literature here (key papers here include [72–76], see [77] for a complete survey of the literature). . The general message in this literature is that players are assumed to be unwaveringly optimistic: no matter what is observed, players maintain the belief that everyone is *rational* at future nodes.

There are many ways to formalize the above intuition that players are "optimistic". I conclude by briefly discussing the approach from [11] since it touches on a number issues raised in this paper. The key idea is to encode the players' strategies as conditional beliefs in an epistemic-plausibility model. For example, consider the following epistemic-plausibility model on the four outcomes of the above extensive game:

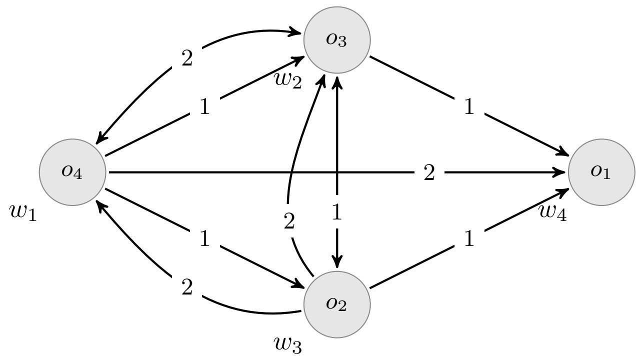

It is assumed that there are atomic propositions for each possible outcome. Formally, suppose that there is an atomic proposition $\mathsf{o_i}$ for each outcome o_i (assume that $\mathsf{o_i}$ is true only at state o_i). The non-terminal nodes $v \in V$ are then identified with the set of outcomes reachable from that node:

$$\mathsf{v} := \bigvee_{v \rightsquigarrow o} \mathsf{o}$$

In the above model, both players 1 and 2 believe that o_1 is the outcome that will be realized, and none of the possible outcomes are initially-ruled out by the players. That is, the model satisfies the "open future" assumption of [11] (none of the players have "hard" information that an outcome is ruled-out). The fact that player 1 is committed to the BI strategy is encoded in the conditional beliefs of the player: both $B_1^{\mathsf{v_1}}\mathsf{o_1}$ and $B_1^{\mathsf{v_3}}\mathsf{o_3}$ are true in the above model. For player 2, $B_2^{\mathsf{v_2}}(\mathsf{o_3} \vee \mathsf{o_4})$ is true in the above model, which implies player 2 plans on choosing action I_2 at node v_2.

The dynamics of actual play is then modeled as a sequence of public announcements (cf. Definition 9). The players' belief change as they learn (irrevocably) which of the nodes in the game are reached. This process produces a sequence of epistemic-plausibility models: For example, a possible sequence of the above game starting with the initial model $\mathcal{M}$ given above is:

$$\mathcal{M} = \mathcal{M}^{!\mathsf{v_1}}; \mathcal{M}^{!\mathsf{v_2}}; \mathcal{M}^{!\mathsf{v_3}}; \mathcal{M}^{!\mathsf{o_4}}$$

The assumption that the players are "incurably optimistic" is represented as follows: no matter what true formula is publicly announced (i.e., no matter how the game proceeds), there is common belief that the players will make a rational choice (when it is their turn to move). Formally, this requires introducing an *arbitrary public announcement operator* [78]: $\mathcal{M}, w \models [\,!\,]\varphi$ provided for all formulas ψ if $\mathcal{M}, w \models \psi$ then $\mathcal{M}, w \models [!\psi]\varphi$. Then, there is *common stable belief* in φ provided $[\,!\,]C^B\varphi$ is true, where $C^B\varphi$ is intended to mean there is common belief in φ (cf. Section 2.4). The key result from [11] is:

Theorem 4 (Baltag, Smets and Zvesper [11]). *Common knowledge of the game structure, of open future and* common stable belief *in dynamic rationality implies common belief in the backward induction outcome.*

Consult the cited paper for the formal details (eg., the definition of dynamic rationality) and the proof.

7 Conclusion: Strategic reasoning in specific games

The analyses in this paper all focused on reasoning in *arbitrary* strategic (and extensive) games. Focusing on *specific* game situations raises new and interesting questions. Here is one example: consider the so-called "hi-low" game:

		Bob	
		l	r
Ann	u	3,3	0,0
	d	0,0	1,1

Since there are no strictly or weakly dominated strategies, it is not hard to find epistemic (-plausibility/-probability) models where there is common knowledge of rationality and any of the outcomes of the game are realized. However, common sense and various empirical observations suggest that *rational* players somehow manage to reason to the *focal point* outcome (u, l):

> "There are these two broad empirical facts about Hi-Lo games, people almost always choose A [Hi] and people with common knowledge of each other's rationality think it is obviously rational to choose A [Hi]."
>
> [79, pg. 42]

What is the underlying dynamics driving players choices in hi-low games? See [80] for an interesting discussion that connects with many of the topics raised in this paper. Perhaps, it involves some form of "group-think" or *team reasoning* [81], or, perhaps, it involves reasoning about which option is *salient* [82]. Each comes with interesting information dynamics that can be uncovered with the logical systems presented in this paper. Finding, a general logic of rational reasoning in such games may be too much to ask:

> "The basic intellectual premise, or working hypothesis, for rational players in this game seems to be the premise that some rule must be used if success is to exceed coincidence, and that the best rule to be found, whatever its rationalization, is consequently a rational rule."
>
> [83, pg. 283, footnote 16]

Nonetheless, focusing on specific instances of strategic interaction raises many new and interesting issues for the dynamic logic analysis of games.

References

1. Rubinstein, A.: Comments on the interpretaiton of game theory. Econometrica **59**(4) (1991) 909–924
2. Brandenburger, A.: The power of paradox: some recent developments in interactive epistemology. International Journal of Game Theory **35** (2007) 465–492
3. Perea, A.: Epistemic Game Theory: Reasoning and Choice. Cambridge University Press (2012)
4. Pacuit, E., Roy, O.: Epistemic game theory. In Zalta, E.N., ed.: Stanford Encyclopedia of Philosophy. Volume forthcoming. (2012)
5. Harsanyi, J.: The tracing procedure: A bayesian approach to defining a solution for n-person noncooperative games. International Journal of Game Theory **4** (1975) 61 – 94
6. Skyrms, B.: The Dynamics of Rational Deliberation. Harvard University Press (1990)
7. Binmore, K.: Modeling rational players: Part I. Economics and Philosophy **3** (1987) 179 – 214
8. Cubitt, R.P., Sugden, R.: Common Knowledge, Salience and Convention: A Reconstruction of David Lewis' Game Theory. Economics and Philosophy **19**(2) (2003) 175–210
9. Cubitt, R., Sugden, R.: Common reasoning in games: A Lewisian analysis of common knowledge of rationality. CeDEx Discussion Paper (2011)
10. Apt, K., Zvesper, J.: Public announcements in strategic games with arbitrary strategy sets. In: Proceedings of LOFT 2010. (2010)
11. Baltag, A., Smets, S., Zvesper, J.: Keep 'hoping' for rationality: a solution to the backwards induction paradox. Synthese **169** (2009) 301–333
12. van Benthem, J.: Rational dynamics and epistemic logic in games. International Game Theory Review **9**(1) (2007) 13–45
13. van Benthem, J., Gheerbrant, A.: Game solution, epistemic dynamics and fixed-point logics. Fund. Inform. **100** (2010) 1–23
14. van Benthem, J., Pacuit, E., Roy, O.: Towards a theory of play: A logical perspective on games and interaction. Games **2**(1) (2011) 52 – 86
15. Baltag, A., Smets, S.: Group belief dynamics under iterated revision: fixed points and cycles of joint upgrades. In: Proceedings of the Twelfth Conference on Theoretical Aspects of Rationality and Knoweldge. (2009) 41–50
16. Baltag, A., Gierasimczuk, N., Smets, S.: Belief revision as a truth-tracking process. In: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge. TARK XIII, ACM (2011) 187–190
17. van Benthem, J.: Logical Dynamics of Information and Interaction. Cambridge University Press (2011)

18. Leyton-Brown, K., Shoham, Y.: Essentials of Game Theory: A Concise Multidisciplinary Introduction. Morgan & Claypool Publishers (2008)
19. Apt, K.R.: A primer on strategic games. In Apt, K.R., Grädel, E., eds.: Lectures in Game Theory for Computer Scientists. (2011) 1 – 33
20. Aumann, R.: Correlated equilibrium as an expression of bayesian rationality. Econometrica **55**(1) (1987) 1 – 18
21. Pearce, D.G.: Rationalizable strategic behavior and the problem of perfection. Econometrica **52**(4) (1984) 1029 – 1050
22. Pacuit, E.: Dynamic epistemic logic. Philosophy Compass (2012) forthcoming
23. Harsanyi, J.: Games with incomplete information played by 'bayesian' players, I - III. Management Science **14** (1967) 159–182, 320–334, 486–502
24. Siniscalchi, M.: Epistemic game theory: Beliefs and types. In Durlauf, S., Blume, L., eds.: The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke (2008)
25. Aumann, R.: Interactive epistemology I: Knowledge. International Journal of Game Theory **28** (1999) 263–300
26. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. The MIT Press (1995)
27. van Benthem, J.: Modal Logic for Open Minds. Center for the Study of Language and Information (2010)
28. Lewis, D.K.: Counterfactuals. Harvard University Press (1973)
29. van Benthem, J.: Dynamic logic for belief revision. Journal of Applied Non-Classical Logics **14**(2) (2004) 129 – 155
30. Baltag, A., Smets, S.: ESSLLI 2009 course: Dynamic logics for interactive belief revision. Slides available online at `http://alexandru.tiddlyspot.com/#%5B%5BESSLLI09%20COURSE%5D%5D` (2009)
31. Board, O.: Dynamic interactive epistemology. Games and Economic Behavior **49** (2004) 49 – 80
32. Boutilier, C.: Conditional Logics for Default Reasoning and Belief Revision. PhD thesis, University of Toronto (1992)
33. Lamarre, P., Shoham, Y.: Knowledge, certainty, belief and conditionalisation. In: Proceedings of the International Conference on Knoweldge Representation and Reasoning. (1994) 415 – 424
34. Leitgeb, H.: Beliefs in conditionals vs. conditional beliefs. Topoi **26**(1) (2007) 115 – 132
35. Fagin, R., Halpern, J., Megiddo, N.: A logic for reasoning about probabilities. Information and Computation **87**(1) (1990) 78 – 128
36. Fagin, R., Halpern, J.: Reasoning about knowledge and probability. Journal of the ACM **41**(2) (1994) 340 – 367
37. Heifetz, A., Mongin, P.: Probability logic for type spaces. Games and Economic Behavior **35** (2001) 31 – 53
38. Zhou, C.: A complete deductive system for probability logic. Journal of Logic and Computation **19**(6) (2009) 1427 – 1454
39. Goldblatt, R.: Deductive systems for coalgebras over measurable spaces. Journal of Logic and Computation (2008) doi: 10.1093/logcom/exn092
40. Halpern, J.: Reasoning about Uncertainty. The MIT Press (2003)
41. Kadane, J.B., Larkey, P.D.: Subjective probability and the theory of games. Management Science **28**(2) (1982) 113–120
42. Rubinstein, A.: The electronic mail game: A game with almost common knowledge. American Economic Review **79** (1989) 385 – 391

43. Kets, W.: Bounded reasoning and higher-order uncertainty. Working paper (2010)
44. Hedden, T., Zhang, J.: What do you think i think you think? strategic reasoning in matrix games. Cognition **85** (2002) 1 – 36
45. Stahl, D., Wilson, P.W.: On players' models of other players: Theory and experimental evidence. Games and Economic Behavior **10** (1995) 218 – 254
46. Wimmer, H., Perner, J.: Beliefs about belief: Representaiton and constraining function of wrong beliefs in youg children's understanding of deception. Cognition **13** (1983) 103 – 128
47. Verbrugge, R.: Logic and social cognition: The facts matter, and so do computational models. Journal of Philosophical Logic **38**(6) (2009) 649 – 680
48. Halpern, J., Moses, Y.: Knowledge and common knowledge in a distributed environment. Journal of the ACM **37**(3) (1990) 549 – 587
49. Chwe, M.S.Y.: Rational Ritual. Princeton University Press (2001)
50. Vanderschraaf, P., Sillari, G.: Common knowledge. In Zalta, E.N., ed.: The Stanford Encyclopedia of Philosophy. Spring 2009 edn. (2009)
51. Aumann, R.: Agreeing to disagree. Annals of Statistics **4** (1976) 1236 –1239
52. Monderer, D., Samet, D.: Approximating common knowledge with common beliefs. Games and Economic Behavior **1** (1989) 170 – 190
53. Lewis, D.: Convention. Harvard University Press (1969)
54. Plaza, J.: Logics of public communications. In Emrich, M.L., Pfeifer, M.S., Hadzikadic, M., Ras, Z., eds.: Proceedings, 4th International Symposium on Methodologies for Intelligent Systems. (1989) 201–216 (republished as [84])
55. Gerbrandy, J.: Bisimulations on Planet Kripke. PhD thesis, University of Amsterdam (1999)
56. Boutilier, C.: Iterated revision and minimal revision of conditional beliefs. Journal of Philosophical Logic **25** (1996) 262 – 304
57. Darwiche, A., Pearl, J.: On the logic of iterated belief revision. Artificial Intelligence **89** (1997) 1 – 29
58. Nayak, A., Pagnucco, M., Peppas, P.: Dynamic belief revision operators. Artificial Intelligence **146** (2003) 193 – 228
59. Stalnaker, R.: Iterated belief revision. Erkentnis **70** (2009) 189 – 209
60. Baltag, A., Smets, S.: Group belief dynamics under iterated revision: Fixed points and cycles of joint upgrades. In: Proceedings of Theoretical Aspects of Rationality and Knowledge. (2009)
61. Gierasimczuk, N.: Knowing Ones Limits: Logical Analysis of Inductive Inference. PhD thesis, ILLC University of Amsterdam (2011)
62. Asheim, G., Dufwenberg, M.: Admissibility and common belief. Game and Economic Behavior **42** (2003) 208 – 234
63. Rabinowicz, W.: Tortous labyrinth: Noncooperative normal-form games between hyperrational players. In Bicchieri, C., Chiara, M.L.D., eds.: Knowledge, Belief and Strategic Interaction. (1992) 107 – 125
64. Cubitt, R., Sugden, R.: Rationally justifiable play and the theory of non-cooperative games. The Economic Journal **104** (1994) 798 – 803
65. Cubitt, R., Sugden, R.: The reasoning-based expected utility procedure. Games and Economic Behavior (2011) In Press
66. Apt, K.R., Zvesper, J.A.: Public announcements in strategic games with arbitrary strategy sets. CoRR **abs/1012.5173** (2010)
67. Pacuit, E., Roy, O.: A dynamic analysis of interactive rationality. In: Proceedings of the Third International Conference on Logic, Rationality, and Interaction. LORI'11, Springer-Verlag (2011) 244–257

68. Samuelson, L.: Dominated strategies and common knowledge. Game and Economic Behavior **4** (1992) 284 – 313
69. Brandenburger, A., Friedenberg, A., Keisler, H.J.: Admissibility in games. Econometrica **76**(2) (2008) 307–352
70. Halpern, J., Pass, R.: A logical characterization of iterated admissibility. In Heifetz, A., ed.: Proceedings of the Twelfth Conference on Theoretical Aspects of Rationality and Knoweldge. (2009) 146 – 155
71. Selten, R.: Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory **V4**(1) (1975) 25 – 55
72. Aumann, R.: Backward induction and common knowledge of rationality. Game and Economic Behavior **8**(6 - 19) (1995)
73. Stalnaker, R.: Belief revision in games: Forward and backward induction. Mathematical Social Science **36** (1998) 31 – 56
74. Aumann, R.: On the centipede game. Game and Economic Behavior **23** (1998) 97 – 105
75. Halpern, J.Y.: Substantive rationality and backward induction. Games and Economic Behavior **37**(2) (2001) 425–435
76. Battigalli, P., Siniscalchi, M.: Strong belief and forward induction reasoning. Journal of Economic Theory **105** (2002) 356 – 391
77. Perea, A.: Epistemic foundations for backward induction: an overview. In van Benthem, J., Gabbay, D., Löwe, B., eds.: Proceedings of the 7th Augustus de Morgan Workshop, Texts in Logic and Games, Amsterdam University Press (2007) 159 – 193
78. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., De Lima, T.: 'knowable' as 'known after an announcement'. The Review of Symbolic Logic **1**(3) (2008) 305–334
79. Bacharach, M.: Beyond Individual Choices: Teams and Frames in Game Theory. Princeton University Press (2006)
80. Bicchieri, C.: Rationality and Coordination. Cambridge University Press (1993)
81. Sugdent, R.: The logic of team reasoning. Philosophical Explorations **63**(165 - 181) (2003)
82. Bardsley, N., Mehta, J., Starmer, C., Sugden, R.: The nature of salience revisited: Cognitve hierarchy theory versus team reasoning. Economic Journal **120**(543) (2010) 40 – 79
83. Schelling, T.: The Strategy of Conflict. Harvard University Press (1960)
84. Plaza, J.: Logics of public communications. Synthese: Knowledge, Rationality, and Action **158**(2) (2007) 165 – 179