

Dynamic Models of Rational Deliberation in Games

Eric Pacuit^(✉)

Department of Philosophy, University of Maryland, College Park, USA
epacuit@umd.edu

Abstract. There is a growing body of literature that analyzes games in terms of the “process of deliberation” that leads the players to select their component of a rational outcome. Although the details of the various models of deliberation in games are different, they share a common line of thought: The rational outcomes of a game are arrived at through a process in which each player settles on an optimal choice given her evolving beliefs about her own choices and the choices of her opponents. The goal is to describe deliberation in terms of a sequence of belief changes about what the players are doing or what their opponents may be thinking. The central question is: What are the update mechanisms that match different game-theoretic analyses? The general conclusion is that the rational outcomes of a game depend not only on the structure of the game, but also on the players’ initial beliefs, which dynamical rule is being used by the players to update their inclinations (in general, different players may be using different rules), and what exactly is commonly known about the process of deliberation.

Keywords: Epistemic game theory · Dynamic epistemic logic · Belief revision

1 Introduction and 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, player i ’s strategy describes the action that player i will choose whenever she is required to make a decision according to the rules of the game.

Traditionally, game theorists have focused on identifying profiles of strategies that constitute an “equilibrium” (e.g., the Nash equilibrium and its refinements). A typical game-theoretic analysis runs as follows: Given a game G , there is an associated *solution space* S_G describing all the possible *outcomes* of G . In a one-shot game (called a **strategic game**; see Sect. 2.1 for details), this is the set of all tuples of strategies (a tuple of strategies, one for each player, is called a **strategy profile**)¹. Abstractly, a **solution** for a game G is a subset of the

¹ The assumption is that once each player settles on a strategy, this identifies a unique outcome of the game. This is a simplifying assumption that can be dropped if necessary. However, for this chapter, it is simpler to follow standard practice and identify the set of “outcomes” of a game with the set of all tuples of actions.

solution space S_G . The subset of S_G identified by a solution concept is intended to represent the “rational outcomes” of the game G .

Suppose that $S \subseteq S_G$ is a solution for the game G . The elements of S are privileged outcomes of G , but what, exactly, distinguishes them from the other outcomes in S_G ? The standard approach is to require that for each profile in S , players should not have an incentive to deviate from their prescribed strategy, given that the other players follow their own prescribed strategies. This is an *internal* constraint on the elements of a solution set since it requires that the strategies in a profile are related to each other in a particular way. This chapter takes a different perspective on the above question by imposing a different constraint on the profiles in S : Each player’s prescribed strategy should be “optimal” given her *beliefs* about what the other players are going to do. This constraint is *external* since it refers to the players’ “beliefs”, which are typically not part of the mathematical representation of the game.

It is not hard to think of situations in which the internal and external constraints on solution concepts discussed above are not jointly satisfied. The point is that players may have very good reasons to believe that the other players are choosing certain strategies, and so, they choose an optimal strategy based on these beliefs. There is no reason to expect that the resulting choices will satisfy the above internal constraint unless one makes strong assumptions about how the players’ beliefs are related². The external constraint on solution concepts can be made more precise by taking a “Bayesian” perspective on game theory [46]: In a game-theoretic situation, as in any situation of choice, the *rational choice* for a player is the one that maximizes expected utility with respect to a (subjective) probability measure over the other players’ strategy choices. A sophisticated literature has developed around this simple idea: it focuses on characterizing solution concepts in terms of what the players know and believe about the other players’ strategy choices and beliefs (see, for example, [7, 22, 26, 60] and [63] for a textbook presentation).

In this chapter, I shift the focus from beliefs about the other players’ choices to the underlying *processes* that lead (rational) players to adopt certain strategies. An early formulation of this idea can be found in John C. Harsanyi’s seminal paper [42], in which he introduced the *tracing procedure* to select an equilibrium in any finite game:

The n players will find the solution s of a giving game G through an intellectual process of *convergent expectations*, to be called the *solution process*....During this process, they will continually and systematically modify [their] expectations—until, at the end of this process, their expectations will come to converge on one particular equilibrium point s in the game G .
(original italics) [42, pg. 71]

The goal of the tracing procedure is to identify a unique Nash equilibrium in any finite strategic game. The idea is to define a continuum of games in

² For example, one can assume that each player *knows* which strategies the other players are going to choose. Robert Aumann and Adam Brandenburger use this assumption to provide an epistemic characterization of the Nash equilibrium [10].

such a way that each of the games has a unique Nash equilibrium. The tracing procedure identifies a path through this space of games ending at a unique Nash equilibrium in the original game. Harsanyi thought of this procedure as “being a mathematical formalization of the process by which rational players coordinate their choices of strategies.”

Harsanyi, in collaboration with Reinhard Selten [43], turned these basic ideas into a beautiful theory of equilibrium selection. This theory is now part of the standard education for any game theorist. Nonetheless, it is not at all clear that this theory of equilibrium selection is best interpreted as a formalization of the players’ processes of “rational deliberation” in game situations (see [72, pgs. 154–158] for a discussion of this point). In this chapter, I will critically discuss three recent frameworks in which the players’ process of “rational deliberation” takes center stage:

1. Brian Skyrms’ model of “dynamic deliberation,” in which 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 [72].
2. Robin Cubitt and Robert Sugden’s recent contribution that develops a “reasoning-based expected utility” procedure for solving games (building on David Lewis’ “common modes of reasoning”) [31, 33].
3. Johan van Benthem *et col.*’s analysis of solution concepts as fixed-points of iterated “(virtual) rationality announcements” [3, 15, 18, 20, 21].

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 in this chapter is to provide a (biased) overview of some key technical and conceptual issues that arise when developing mathematical models of players deliberating about what to do in a game situation.

2 Background

I assume that the reader is familiar with the basics of game theory (see [52] and [2] for concise discussions of the key concepts, definitions and theorems) and formal models of knowledge and belief (see [19, 57] for details). In this section, I introduce some key definitions in order to fix notation.

2.1 Strategic Games

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 $i \in N$, $u_i : \prod_{i \in N} S_i \rightarrow \mathbb{R}$ is a utility function assigning real numbers to each outcome of the game (i.e., tuples consisting of the choices for

each player). Strategic games represent situations in which each player makes a single decision, and all the players make their decisions simultaneously. If $s \in \prod_{i \in N} S_i$ is a strategy profile, then write s_i for the i th component of s and s_{-i} for the sequence consisting of all components of s except for s_i (let S_{-i} denote all such sequences of strategies).

Recall from the introduction that the *solution space* S_G for a game G is the set of all outcomes of G . Since we identify the outcomes of a game with the set of strategy profiles, we have $S_G = \prod_{i \in N} S_i$. This means that a “solution” to a strategic game is a distinguished set of strategy profiles. In the remainder of this section, I will define some standard game- and decision-theoretic notions that will be used throughout this chapter.

Mixed Strategies. Let $\Delta(X)$ denote the set of probability measures over the finite³ 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). Mixed strategies are incorporated into a game-theoretic analysis as follows. Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a finite strategic game. The **mixed extension** of G is the strategic game in which the strategies for player i are the mixed strategies in G (i.e., $\Delta(S_i)$), and the utility for player i (denoted U_i) of the joint mixed strategy $m \in \prod_{i \in N} \Delta(S_i)$ is calculated in the obvious way (let $m(s) = m_1(s_1) \cdot m_2(s_2) \cdot \dots \cdot m_n(s_n)$ for $s \in \prod_{i \in N} S_i$):

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

Thus, the solution space of a mixed extension of the game G is the set $\prod_{i \in N} \Delta(S_i)$.

Mixed strategies play an important role in many game-theoretic analyses. However, the interpretation of mixed strategies is controversial, as Ariel Rubinstein notes: “We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes” [69, pg. 913]. For the purposes of this chapter, I will assume that players choose only pure strategies. Mixed strategies do play a role in Sect. 3, where they describe each players’ beliefs about what they will do (at the end of deliberation).

Nash Equilibrium. The most well-known and extensively studied solution concept is the **Nash equilibrium**. Let $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a finite strategic game. A mixed strategy profile $m = (m_1, \dots, m_n) \in \prod_{i \in N} \Delta(S_i)$ is a Nash equilibrium provided for all $i \in N$,

$$U_i(m_1, \dots, m_i, \dots, m_n) \geq U_i(m_1, \dots, m'_i, \dots, m_n), \quad \text{for all } m'_i \in \Delta(S_i).$$

This definition is an example of the *internal* constraint on solutions discussed in the introduction. Despite its prominence in the game theory literature, the

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

Nash equilibrium faces many foundational problems [68]. For example, there are theoretical concerns about what the players need to *know* in order to play their component of a Nash equilibrium [10, 62]; questions about how players choose among multiple Nash equilibria; and many experiments purporting to demonstrate game-theoretic situations in which the player’s choices do not form a Nash equilibrium. Nash equilibrium does not play an important role in this chapter. I focus, instead, on the outcomes of a game that can be reached through a process of “rational deliberation”.

Iteratively Removing Strategies. A strategy $s \in S_i$ **strictly dominates** strategy $s' \in S_i$ provided that

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

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

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

More generally, the strategy s strictly/weakly dominates s' **with respect to a set** $X \subseteq S_{-i}$ if S_{-i} is replaced with X in the above definitions⁴. Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ and $G' = \langle N, \{S'_i\}_{i \in N}, \{u'_i\}_{i \in N} \rangle$ are strategic games. The game G' is a **restriction** of G provided that for each $i \in N$, $S'_i \subseteq S_i$ and u'_i is the restriction of u_i to $\prod_{i \in N} S'_i$. To simplify notation, write G_i for the set of strategies for player i in game G . Strict and weak dominance can be used to *reduce* a strategic game. Write $H \rightarrow_{SD} H'$ whenever $H \neq H'$, H' is a restriction of H and

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

So, if $H \rightarrow_{SD} H'$, 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, H is the result of iteratively removing strictly dominated strategies (IESDS) provided that $G \rightarrow_{SD}^* H$, where \rightarrow^* is the reflexive transitive closure⁵ of a relation \rightarrow .

The above definition can be easily adapted to other choice rules, such as weak dominance. Let \rightarrow_{WD} denote the relation between games defined as above using weak dominance instead of strict dominance⁶. Furthermore, the above

⁴ Furthermore, the definitions of strict and weak dominance can be extended so that strategies may be strictly/weakly dominated by *mixed strategies*. This is important for the epistemic analysis of iterative removal of strictly/weakly dominated strategies. However, for my purposes in this chapter, I can stick with the simpler definition in terms of pure strategies.

⁵ The reflexive transitive closure of a relation R is the smallest relation R^* containing R that is reflexive and transitive.

⁶ Some interesting issues arise here: It is well-known that, unlike with strict dominance, different orders in which weakly dominated strategies are removed can lead to different outcomes. Let us set aside these issues in this chapter.

definition of iterated removal of strictly/weakly dominated strategies can be readily adapted to the mixed extensions of a strategic game.

There are a number of ways to interpret the iterative process of removing strategies, defined above. The first is that it is an algorithm that a game theorist can use to find an equilibrium in a game. The second interpretation views the successive steps of the removal process as corresponding to the players' *higher-order beliefs* (i.e., player i believes that player j believes that player i believes that...that player i will not play such-and-such strategy). Finally, the third interpretation is that the iterative process of removing strategies tracks the “back-and-forth reasoning” players engage in as they decide what to do in a game situation (i.e., if player i does not play such-and-such a strategy, then player j will not play such-and-such a strategy, and so on).

Bayesian Rationality. In this chapter, I am interested not only in solutions to a game, but also what the players *believe* about the outcomes of a game. Let $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a strategic game. A probability measure $\pi \in \Delta(S_{-i})$ is called a **conjecture** for player i . The **expected utility** of $s \in S_i$ for player i with respect to $\pi \in \Delta(S_{-i})$ is:

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

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

One conclusion to draw from the discussion in this section is that much can be said about the issues raised in this chapter using standard game-theoretic notions. Indeed, it is standard for a game theorist to distinguish between the *ex ante* and *ex interim* stages of decision making⁷. In the former, the players have not yet decided what strategy they will choose, while, in the latter, the players know their own choices but not their opponents'. However, the *process* by which the players form their beliefs in the *ex interim* stage is typically not discussed. The frameworks discussed in the remainder of this chapter are focused on making this process explicit.

2.2 Game Models

A **game model** describes a particular play of the game *and* what the players think about the other players. That is, a game model represents an “informational context” of a given play of the game. This includes the “knowledge” the players have about the game situation *and* what they think about the other

⁷ There is also an *ex post* analysis when all choices are “out in the open,” and the only remaining uncertainties are about what the other players are thinking.

players' choices and beliefs. Researchers interested in the foundations of decision theory, epistemic and doxastic logic and formal epistemology have developed many different formal models to describe the variety of informational attitudes important for assessing decision maker's choices in a decision- or game-theoretic situation. See [19] for an overview and pointers to the relevant literature. In this section, I present the details of a logical framework that can be used to reason about the informational context of a game.

Syntactic issues do not play an important role in this chapter. Nonetheless, I will give the definition of truth for a relevant formal language, as it makes for a smoother transition from the game theory literature to the literature on dynamic epistemic logic and iterated belief change discussed in Sect. 5.1. Consult [19, 57, 58] for a discussion of the standard logical questions about axiomatics, definability, decidability of the satisfiability problem, and so on.

Epistemic-Plausibility Models. Variants of the models presented in this section have been studied extensively by logicians [13, 17, 19], game theorists [23], philosophers [51, 74] and computer scientists [25, 48]. The models are intended to describe what the players know and believe about an outcome of the game.

The first component of an epistemic-plausibility model is a nonempty set W of **states** (also called **worlds**). Each state in a game model will be associated with an outcome of a game G via a function σ , called the **outcome map**. So, for a state w , $\sigma(w)$ is the element of S_G realized at state w . Let $\sigma_i(w)$ denote the i th component of $\sigma(w)$ (so, $\sigma_i(w)$ is the strategy played by i at state w). The atomic propositions are intended to describe different aspects of the the 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. There are a number of ways to make this precise. Perhaps the simplest is to introduce, for each player i and strategy $a \in S_i$, an atomic proposition $play_i(a)$ intended to mean “player i is playing strategy a .” For a game $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$, let $\text{At}(G) = \{play_i(a) \mid i \in N \text{ and } a \in S_i\}$ be the set of atomic propositions for the game G .

There are two additional components to an epistemic-plausibility model. The first is a set of equivalence relations \sim_i , one for each player. The intended reading of $w \sim_i v$ is that “everything that i knows at w is true at v ”. Alternatively, I will say that “player i does not have enough information to distinguish state w from state v .”

The second component is a *plausibility ordering* for each player: 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 $X \subseteq W$, set $\text{Min}_{\preceq_i}(X) = \{v \in X \mid v \preceq_i w \text{ for all } w \in X\}$, 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”).

Putting everything together, the definition of an epistemic-plausibility model is as follows:

Definition 1. *Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a strategic game. An **epistemic-plausibility model** for G is a tuple $\mathcal{M} = \langle W, \{\sim_i\}_{i \in \mathcal{A}}, \{\preceq_i\}_{i \in \mathcal{A}}, \sigma \rangle$ where $W \neq \emptyset$; for each $i \in \mathcal{A}$, $\sim_i \subseteq W \times W$ 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 $v \sim_i u$); for each $i \in \mathcal{A}$, \preceq_i is a well-founded (every non-empty set of states has a minimal element)⁸ reflexive and transitive relation on W ; and σ is an outcome map. In addition, the following two conditions are imposed for all $w, v \in W$:*

1. *if $w \preceq_i v$ then $w \sim_i v$ (plausibility implies possibility), and*
2. *if $w \sim_i v$ then either $w \preceq_i v$ or $v \preceq_i w$ (locally-connected). ◁*

Models without plausibility relations are called **epistemic models**.

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, $w \not\preceq_i v$ and $v \not\preceq_i w$. Thus, I have the following equivalence: $w \sim_i v$ iff $w \preceq_i v$ or $v \preceq_i w$. 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 $v \preceq_i w$.

For each strategic game G , let $\mathcal{L}_{KB}(G)$ be the set of sentences generated by the following grammar⁹:

$$\varphi := \text{play}_i(a) \mid \neg\varphi \mid \varphi \wedge \psi \mid B_i^\varphi\psi \mid K_i\varphi$$

where $i \in N$ and $\text{play}_i(a) \in \text{At}(G)$. The additional propositional connectives ($\rightarrow, \leftrightarrow, \vee$) are defined as usual and the dual of K_i , denoted L_i , is defined as follows: $L_i\varphi := \neg K_i\neg\varphi$. The intended interpretation of $K_i\varphi$ is “agent i knows that φ ”¹⁰. The intended interpretation of $B_i^\varphi\psi$ is “agent i believes ψ under the supposition that φ is true”.

Truth for formulas in $\mathcal{L}_{KB}(G)$ is defined as usual. 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 $\text{Min}_{\preceq_i}([w]_i \cap X)$ is nonempty if $[w]_i \cap X \neq \emptyset$.

⁸ Well-foundedness is only needed to ensure that for any set X , $\text{Min}_{\preceq_i}(X)$ 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.

⁹ There are other natural modal operators that can. See [57] for an overview and pointers to the relevant literature.

¹⁰ This is the standard interpretation of $K_i\varphi$ in the game theory literature. Whether this captures any of the many different definitions of knowledge found in the epistemology literature is debatable. A better reading of $K_i\varphi$ is “given all of the available evidence and everything i has observed, agent i is informed that φ is true”.

Definition 2 (Truth for $\mathcal{L}_{KB}(G)$). Given an epistemic-plausibility model $\mathcal{M} = \langle W, \{\sim_i\}_{i \in \mathcal{A}}, \{\preceq_i\}_{i \in \mathcal{A}}, \sigma \rangle$. Truth for formulas from $\mathcal{L}_{KB}(G)$ is defined recursively:

- $\mathcal{M}, w \models \text{play}_i(a)$ iff $\sigma_i(w) = a$
- $\mathcal{M}, w \models \neg\varphi$ 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 $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models B_i^\varphi\psi$ iff for all $v \in \text{Min}_{\preceq_i}([w]_i \cap \llbracket \varphi \rrbracket_{\mathcal{M}})$, $\mathcal{M}, v \models \psi$

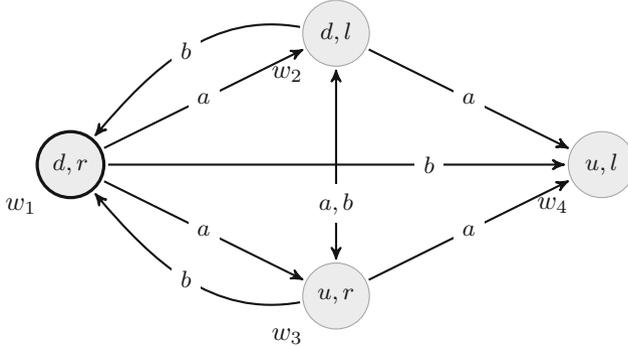
Thus, i believes ψ conditional on φ , $B_i^\varphi\psi$, if i 's most plausible φ -worlds (i.e., the states satisfying φ that i has not ruled out and considers most plausible) all satisfy ψ . Full belief is defined as follows: $B_i\varphi := B_i^\top\varphi$. Then, the definition of plain belief is:

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

I illustrate the above definition with the following coordination game:

Bob													
	<table style="border-collapse: collapse;"> <tr> <td></td> <td style="text-align: center;">l</td> <td style="text-align: center;">r</td> </tr> <tr> <td style="text-align: center;">Ann</td> <td style="border: 1px solid black; padding: 5px;">u</td> <td style="border: 1px solid black; padding: 5px;">d</td> </tr> <tr> <td></td> <td style="border: 1px solid black; padding: 5px;">3, 3</td> <td style="border: 1px solid black; padding: 5px;">0, 0</td> </tr> <tr> <td></td> <td style="border: 1px solid black; padding: 5px;">d</td> <td style="border: 1px solid black; padding: 5px;">1, 1</td> </tr> </table>		l	r	Ann	u	d		3, 3	0, 0		d	1, 1
	l	r											
Ann	u	d											
	3, 3	0, 0											
	d	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 a strategy): I draw an i -labeled arrow from v to w if $w \preceq_i v$ (to keep minimize the clutter, I do not include all arrows; the remaining arrows can be inferred by reflexivity and 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 :

$$B_a(\text{play}_a(u) \wedge \text{play}_b(l)) \wedge B_b(\text{play}_a(u) \wedge \text{play}_b(l))$$

is true at all states. However, Ann and Bob do have different *conditional* beliefs. On the one hand, Ann believes that their choices are independent; thus, she believes that $play_b(l)$ is true even under the supposition that $play_a(d)$ is true (i.e., she continues to believe that Bob will play l even if she decides to play d). On the other hand, Bob believes that their choices are somehow correlated; thus, under the supposition that $play_b(r)$ 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. [49]).

Common Knowledge and Belief. States in an epistemic-plausibility 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. 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 [40] and social interactions [28]. In this section, I briefly recount the standard definition of common knowledge¹¹.

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} :

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

where $X \subseteq N$ is a finite set. 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 [6], the idea is to define common knowledge of φ as the following iteration of everyone knows operators:

$$\varphi \wedge K_N\varphi \wedge K_NK_N\varphi \wedge K_NK_NK_N\varphi \wedge \dots$$

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, a modal operator $C_G\varphi$ with the intended meaning " φ is common knowledge among the group G " must be added to our modal language. Formally:

Definition 3 (Interpretation of C_G). Let $\mathcal{M} = \langle W, \{\sim_i\}_{i \in \mathcal{A}}, V \rangle$ be an epistemic model¹² and $w \in W$. The truth of formulas of the form $C_X\varphi$ is:

$$\mathcal{M}, w \models C_X\varphi \text{ iff for all } v \in W, \text{ if } wR_X^C v \text{ then } \mathcal{M}, v \models \varphi$$

where $R_X^C := (\bigcup_{i \in X} \sim_i)^*$ is the reflexive transitive closure of $\bigcup_{i \in X} \sim_i$.

¹¹ 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 [34, 36, 57, 76].

¹² The same definition will, of course, hold for epistemic-plausibility and epistemic-probability models.

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, $\mathcal{M}, w \models C_X\varphi$ iff every finite path for X from w ends with a state satisfying φ .

This approach to defining common knowledge can be viewed as a recipe for defining common (robust) belief. For example, suppose that wR_i^Bv iff $v \in \text{Min}_{\succeq_i}(\{w\}_i)$, and define R_G^B to be the transitive closure¹³ 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 in [55].

3 Reasoning to an Equilibrium

Brian Skyrms presents a model of the players' process of deliberation in a game in his important book *The Dynamics of Rational Deliberation* [72]. In this section, I introduce and discuss this model of deliberation, though the reader is referred to [72] for a full discussion (see, also, [1, 45] for analyses of this model).

To simplify the exposition, I restrict attention to a two-person finite strategic game. Everything discussed below can be extended to situations with more than two players¹⁴ and to extensive games¹⁵. Suppose that $G = \langle \{a, b\}, \{S_a, S_b\}, \{u_a, u_b\} \rangle$ is a strategic game in which $S_a = \{s_1, \dots, s_n\}$ and $S_b = \{t_1, \dots, t_m\}$ are the players' strategies, and u_a and u_b are utility functions. In the simplest case, deliberation is trivial: Each player calculates the expected utility given her belief about what her opponent is going to do and then chooses the action that maximizes these expected utilities. One of Skyrms' key insights is that this calculation may be informative to the players, and if a player believes that there is any possibility that the process of deliberation may ultimately lead her to a different decision, then she will not act until her deliberation process has reached a stable state¹⁶.

Deliberation is understood as an iterative process that modifies the players' opinions about the strategies that they will choose (at the end of the deliberation). For each player, a **state of indecision** is a probability measure on that player's set of strategies—i.e., an element of $\Delta(S_i)$ for $i = a, b$. Note that each state of indecision is a *mixed strategy*. However, the interpretation of the mixed strategies differs from the one discussed in Sect. 2.1. In this model, the interpretation is that the state of indecision for a player i at any given stage of the deliberation process is the mixed strategy that player i would choose if the player stopped deliberating. It is the players' states of indecision that evolve during the deliberation process.

¹³ Since beliefs need not be factive, I do not force R_G^B to be reflexive.

¹⁴ However, see [1] for interesting new issues that arise with more than two players.

¹⁵ See [72], pgs. 44 – 52 and Chap. 5.

¹⁶ See [72], Chap. 4.

Let $p_a \in \Delta(S_a)$ and $p_b \in \Delta(S_b)$ be states of indecision for a and b , respectively, and assume that the states of indecision are *common knowledge*. One consequence of this assumption is that the players can calculate the expected utilities of their strategies (using their opponent's state of indecision). For example, for $s_j \in S_a$, we have

$$EU_a(s_j) = \sum_{t_k \in S_b} p_b(t_k) u_i(s_j, t_k),$$

and similarly for b . The **status quo** is the expected utility of the current state of indecision:

$$SQ_a = \sum_{s_j \in S_a} p_a(s_j) \cdot EU_a(s_j) \quad SQ_b = \sum_{t_k \in S_b} p_b(t_k) \cdot EU_b(t_k).$$

Once the expected utilities are calculated, the players modify their states of indecision so that they believe more strongly that they will choose strategies with higher expected utility than the status quo. Players can use various rules to update their states of indecision accordingly. In general, any dynamical rule can be used so long as the rule *seeks the good* in the following sense:

1. The rule raises the probability of a strategy only if that strategy has expected utility greater than the status quo.
2. The rule raises the sum of the probabilities of all strategies with expected utility greater than the status quo (if any).

Deliberation reaches a **fixed-point** when the dynamical rule no longer changes the state of indecision. It is not hard to see that all dynamical rules that seek the good have, as fixed-points, states of indecision in which the expected utility of the status quo is maximal. To illustrate Skyrms' model of deliberation with an example, I give the details of one of the rules discussed in [72]:

Nash dynamics. The **covetability** of a strategy s for player i is calculated as follows: $cov_i(s) = \max(EU_i(s) - SQ_i, 0)$. Then, **Nash dynamics** transform a probability $p \in \Delta(S_i)$ into a new probability $p' \in \Delta(S_i)$ as follows. For each $s \in S_i$:

$$p'(s) = \frac{k \cdot p(s) + cov_i(s)}{k + \sum_{s \in S_i} cov(s)},$$

where $k > 0$ is the “index of caution” (the higher the k , the more slowly the decision maker raises the probability of strategies that have higher expected utility than the status quo).

In addition to assuming that the initial states of indecision are common knowledge, it is assumed that each player can *emulate* the other's calculations, and that each player is, in fact, using the same dynamical rule to modify her state of indecision. Given that all of this is common knowledge, the states of indecision resulting from one round of the deliberation process will, again, be common knowledge and the process can continue until a fixed-point is reached.

A simple example will make this more concrete. Consider the following game between two players, Ann (*a*) and Bob (*b*)¹⁷.

		Bob	
		<i>l</i>	<i>r</i>
Ann	<i>u</i>	2, 1	0, 0
	<i>d</i>	0, 0	1, 2

There are two pure Nash equilibria ((u, l) and (d, r)) and one mixed-strategy Nash equilibrium where Ann plays *u* with probability 2/3 and Bob plays *l* with probability 1/3. Suppose that the initial state of indecision is:

$$p_a(u) = 0.2, p_a(d) = 0.8 \quad \text{and} \quad p_b(l) = 0.9, p_b(r) = 0.1.$$

Since both players have access to each other’s state of indecision, they can calculate the expected utilities of each of their strategies:

$$\begin{aligned} EU_a(u) &= 2 \cdot 0.9 + 0 \cdot 0.1 = 1.8 \\ EU_a(d) &= 0 \cdot 0.9 + 1 \cdot 0.1 = 0.1 \\ EU_b(l) &= 1 \cdot 0.2 + 0 \cdot 0.8 = 0.2 \\ EU_b(r) &= 0 \cdot 0.2 + 2 \cdot 0.8 = 1.6 \end{aligned}$$

If the players simply choose the strategy that maximizes their expected utilities, then the outcome of the interaction will be the off-equilibrium profile (u, r) . However, the process of deliberation will pull the players towards an equilibrium. The status quo for each player is:

$$\begin{aligned} SQ_a &= 0.2 \cdot EU_a(u) + 0.8 \cdot EU_a(d) = 0.2 \cdot 1.8 + 0.8 \cdot 0.1 = 0.44 \\ SQ_b &= 0.4 \cdot EU_b(l) + 0.6 \cdot EU_b(r) = 0.9 \cdot 0.2 + 0.1 \cdot 1.6 = 0.34 \end{aligned}$$

The covetabilities for each of the strategies are:

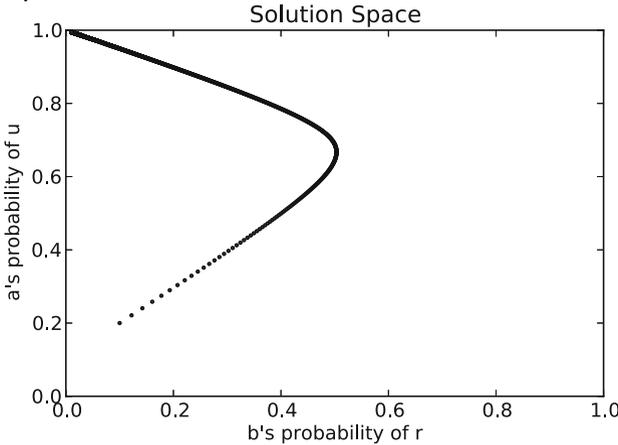
$$\begin{aligned} cov_a(u) &= \max(1.8 - 0.44, 0) = 1.36 \\ cov_a(d) &= \max(0.34 - 0.44, 0) = 0 \\ cov_b(l) &= \max(0.2 - 0.34, 0) = 0 \\ cov_b(r) &= \max(1.6 - 0.34, 0) = 1.26 \end{aligned}$$

Now, the new states of indecision p'_a and p'_b are calculated using Nash dynamics (for simplicity, I assume that the index of caution is $k = 1$):

¹⁷ This game is called the “Battle of the Sexes”. The underlying story is that Ann and Bob are married and are deciding where to go for dinner. Ann would rather eat Indian food than French food, whereas Bob prefers French food to Indian food. They both prefer to eat together rather than separately. The outcome (u, l) is that they go to an Indian restaurant together; (d, r) is the outcome that they go to a French restaurant together; and (u, r) and (d, l) are outcomes where they go to different restaurants.

$$\begin{aligned}
 p'_a(u) &= \frac{p_A(u)+cov_a(u)}{1+(cov_a(u)+cov_a(d))} = \frac{0.2+1.36}{1+1.36} = 0.221183800623 \\
 p'_a(d) &= \frac{p_a(d)+cov_a(d)}{1+(cov_a(u)+cov_a(d))} = \frac{0.8+0}{1+1.36} = 0.778816199377 \\
 p'_b(l) &= \frac{p_b(l)+cov_b(l)}{1+(cov_b(l)+cov_b(r))} = \frac{0.9+0}{1+1.26} = 0.87787748732 \\
 p'_b(r) &= \frac{p_b(r)+cov_b(r)}{1+(cov_b(l)+cov_b(r))} = \frac{0.1+1.26}{1+1.26} = 0.12212251268
 \end{aligned}$$

The new states of indecision are now p'_a and p'_b , and we can continue this process. One can visualize this process by the following graph, in which the x -axis is the probability that Bob will choose r and the y -axis is the probability that Ann will choose u ¹⁸.



The deliberation reaches a fixed-point with Ann and Bob deciding to play their part of the Nash equilibrium (u, l) . In fact, Skyrms shows that under the strong assumptions of common knowledge noted above and assuming that all players use dynamical rules that seek the good, when the process of deliberation reaches a fixed-point, the states of indecision will form a Nash equilibrium¹⁹.

4 Strategic Reasoning as a Solution Concept

A key aspect of the iterative removal of dominated strategies is that at each stage of the process, strategies are identified as either “good” or “bad”. The “good” strategies are those that are not strictly/weakly dominated, while the “bad” ones are weakly/strictly dominated. If the intended interpretation of the iterative procedure that removes weakly/strictly dominated strategies is to represent the players “deliberation” about what they are going to do, then this is a significant assumption. The point is that while a player is deliberating about what to do in a

¹⁸ This graph was produced by a python program with an index of caution $k = 25$ and a satisficing value of 0.01. A satisficing value of 0.01 means that the process stops when the covetabilities fall below 0.01. Contact the author for the code for this simulation.

¹⁹ The outcome may end in a mixed-strategy Nash equilibrium.

game situation, there may be strategies that cannot yet be classified as “good” or “bad”. These are the strategies that the player needs to think about more before deciding how to classify them. Building on this intuition, the *reasoning-based expected utility procedure* of [32] is intended to model the reasoning procedure that a *Bayesian* rational player would follow as she decides what to do in a game.

At each stage of the procedure, strategies are **categorized**. A categorization is a *ternary* partition of the players’ strategies S_i , rather than the usual binary partition in terms of which strategies are strictly/weakly dominated and which are not. The key idea is that during the reasoning process, strategies are accumulated, deleted or neither. Formally, for 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 in $S_i^+ \cup S_i^-$ are classified as “neither accumulated nor 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 let $S_{i,0}^+ = S_{i,0}^- = \emptyset$. Then, for $n \geq 0$, we have:

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

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

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

$$S_{i,n+1}^- \{s_i \in S_{i,n} \mid \text{there is no } \pi \in D_{i,n} \text{ such that } MEU(s_i, \pi)\}.$$

- 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 $\prod_{j \neq i} S_{j,n+1}^+$ and 0 probability to any strategy profile from $\prod_{j \neq i} S_{j,n+1}^-$.

The following example from [32] illustrates this procedure:

		Bob	
		l	r
	u	1, 1	1, 1
Ann	m_1	0, 0	1, 0
	m_2	2, 0	0, 0
	d	0, 2	0, 0

For Bob, strategy l is accumulated since it maximizes expected utility with respect to every probability on Ann’s strategies (note that l weakly dominates r). For Ann, d is deleted, as it does not maximize probability with respect to any probability measure on Bob’s strategies (note that d is *strictly* dominated by u). Thus, we have

$$\begin{aligned}
S_{a,1}^+ &= \emptyset \\
S_{a,1}^- &= \{d\} \\
S_{b,1}^+ &= \{l\} \\
S_{b,1}^- &= \emptyset
\end{aligned}$$

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 Ann playing d . This means that r is accumulated for Bob²⁰ and m_1 is deleted²¹ for Ann:

$$\begin{aligned}
S_{a,2}^+ &= \emptyset \\
S_{a,2}^- &= \{b, m_1\} \\
S_{b,2}^+ &= \{l, r\} \\
S_{b,2}^- &= \emptyset
\end{aligned}$$

At this point, the procedure reaches a fixed-point with Bob accumulating l and r and Ann deleting d and m_1 . The interpretation is that Bob has good reason to play either l or r and, thus, must *pick* one of them. All that Ann was able to conclude is that d and m_1 are not good choices.

The general message is that players may not be able to identify a *unique* rational strategy by strategic reasoning alone (as represented by the iterative procedure given above). There are two reasons why this may happen. First, a player may accumulate more than one strategy, and so the player must “pick”²² one of them. This is what happened with Bob. Given the observation that Ann will not choose b , both of the choices l and r give the same payoff, and so Bob must *pick* one of them²³. Second, players may not have enough information to identify the “rational” choices. Without any information about which of l or r Bob will pick, Ann cannot come to a conclusion about which of u or m_2 she should choose. Thus, neither of these strategies can be accumulated. Ann and Bob face very different decision problems. No matter which choice Bob ends up picking, his choice will be rational (given his belief that Ann will not choose irrationally). However, since Ann lacks a probability over how Bob will pick, she cannot identify a rational choice.

5 Reasoning to a Game Model

The game models introduced in Sect. 2.2 can be used to describe the informational context of a game. A natural question from the perspective of this

²⁰ If Bob assigns probability 0 to Ann playing d , then the strategies l and r give exactly the same payoffs.

²¹ The only probability measures such that m_1 maximizes expected utility are the ones that assign probability 1 to Bob playing r .

²² See [56] for an interesting discussion of “picking” and “choosing” in decision theory.

²³ Of course, Bob may think it is possible that Ann is irrational, and so she could choose the strictly dominated strategy d . Then, depending on how likely Bob thinks it is that Ann will choose irrationally, l may be the only rational choice for him. In this chapter, we set aside such considerations.

chapter is: How do the players arrive at a particular informational context? In this section, I introduce different operations that transform epistemic-plausibility models. These operations are intended to represent different ways a rational agent’s knowledge and beliefs can change over time. Then, I show how to use these operations to describe how the players’ knowledge and beliefs change as they each deliberate about what they are going to do in a game situation.

5.1 Modeling Information Changes

The simplest type of informational change treats the source of the information as *infallible*. The effect of finding out from an infallible source that φ is true should be clear: *Remove* all states that do not satisfy φ . In the epistemic logic literature, this operation is called a *public announcement* [37, 65]. However, calling this an “announcement” is misleading since, in this chapter, I am not modeling any form of “pre-play” communication. The “announcements” are formulas that the players incorporate into the current epistemic state.

Definition 4 (Public Announcement). *Suppose that $\mathcal{M} = \langle W, \{\sim_i\}_{i \in \mathcal{A}}, V \rangle$ 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 $\mathcal{M}^{1\varphi} = \langle W^{1\varphi}, \{\sim_i^{1\varphi}\}, V^{1\varphi} \rangle$, where $W^{1\varphi} = \{w \in W \mid \mathcal{M}, w \models \varphi\}$; $\sim_i^{1\varphi} = \sim_i \cap W^{1\varphi} \times W^{1\varphi}$ for all $i \in \mathcal{A}$; and $\sigma^{1\varphi}$ is the restriction of σ to $W^{1\varphi}$.*

The models \mathcal{M} and $\mathcal{M}^{1\varphi}$ describe two different moments in time, with \mathcal{M} describing the current or initial information state of the agents and $\mathcal{M}^{1\varphi}$ the information state *after* 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 $\mathcal{M}^{1\varphi}, w \models \psi$].

A public announcement is only one type of informative action. For the other transformations discussed in this chapter, while the agents do *trust* the source of φ , they do not treat it 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 operation is *radical upgrade* ($\uparrow\uparrow\varphi$), which moves *all* the 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 $\llbracket\varphi\rrbracket_i^w = \{x \mid \mathcal{M}, x \models \varphi\} \cap [w]_i$ denote the set of all φ -worlds that i considers possible at state w and $best_i(\varphi, w) = Min_{\preceq_i}(\llbracket\varphi\rrbracket_i^w)$ be the best φ -worlds at state w , according to agent i .

Definition 5 (Conservative and Radical Upgrade). *Given an epistemic-plausibility model $\mathcal{M} = \langle W, \{\sim_i\}_{i \in \mathcal{A}}, \{\preceq_i\}_{i \in \mathcal{A}}, \sigma \rangle$ and a formula $\varphi \in \mathcal{L}_{KB}$, the*

conservative/radical upgrade of \mathcal{M} with φ is the model $\mathcal{M}^{*\varphi} = \langle W^{*\varphi}, \{\sim_i^{*\varphi}\}_{i \in N}, \{\preceq_i^{*\varphi}\}_{i \in N}, V^{*\varphi} \rangle$ 10:44 AM 9/11/2015 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 $\preceq_i^{\uparrow\uparrow\varphi}$ are the smallest relations satisfying:

Conservative Upgrade

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

Radical Upgrade

1. for all $x \in \llbracket \varphi \rrbracket_i^w$ and $y \in \llbracket \neg \varphi \rrbracket_i^w$, set $x \prec_i^{\uparrow\varphi} y$;
2. for all $x, y \in \llbracket \varphi \rrbracket_i^w$, set $x \preceq_i^{\uparrow\varphi} y$ iff $x \preceq_i y$; and
3. for all $x, y \in \llbracket \neg \varphi \rrbracket_i^w$, 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 $\text{best}_i(\varphi, w)$. A logical analysis of these operations includes formulas of the form $\llbracket \uparrow \varphi \rrbracket \psi$ intended to mean “after everyone conservatively upgrades with φ , ψ is true” and $\llbracket \uparrow\uparrow \varphi \rrbracket \psi$ intended to mean “after everyone radically upgrades with φ , ψ is true”. The definition of truth for these formula is as expected:

- $\mathcal{M}, w \models \llbracket \uparrow \varphi \rrbracket \psi$ iff $\mathcal{M}^{\uparrow\varphi}, w \models \psi$
- $\mathcal{M}, w \models \llbracket \uparrow\uparrow \varphi \rrbracket \psi$ iff $\mathcal{M}^{\uparrow\uparrow\varphi}, w \models \psi$

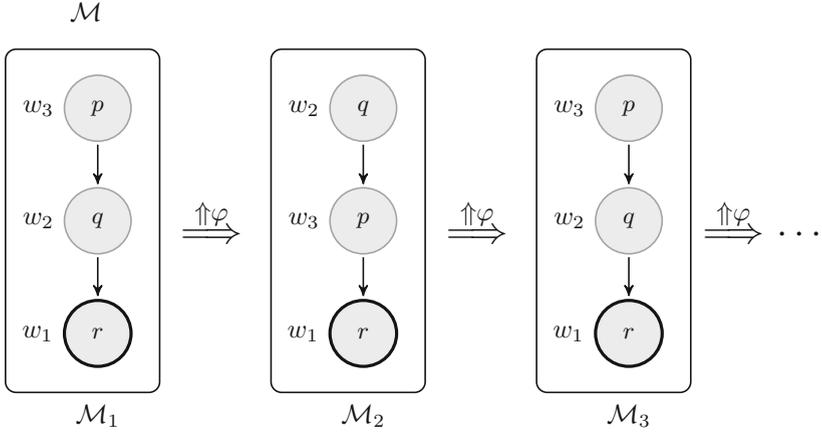
The main issue of interest in this chapter is the limit behavior of iterated sequences of announcements. That is, what happens to the epistemic-plausibility models *in the limit*? Do the players’ knowledge and beliefs stabilize or keep changing in response to the new information?

An initial observation is that iterated public announcement of any formula φ in an epistemic-plausibility model must stop at a limit model where either φ or its negation is true at all states (see [14] 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 [14] and [21, Sect. 4] for general discussions. I conclude this brief introduction to dynamic logics of knowledge and beliefs with an example of the type of dynamics that can arise.

Let \mathcal{M}_1 be an initial epistemic-plausibility model (for a single agent) 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

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

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 non-conditional beliefs remain fixed throughout the process. In fact, Baltag and Smets have shown that every iterated sequence of truthful radical upgrades stabilizes all non-conditional beliefs in the limit [14]. See [12, 38, 58] for generalizations and broader discussions about the issues raised in this section.

5.2 Rational Belief Change During Deliberation

This section looks at the operations that transform the informational context of a game *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 (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 their 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

$$\mathcal{M}_0 \xrightarrow{\tau(D_0)} \mathcal{M}_1 \xrightarrow{\tau(D_1)} \mathcal{M}_2 \xrightarrow{\tau(D_2)} \dots \xrightarrow{\tau(D_n)} \mathcal{M}_{n+1} \implies \dots$$

where each D_i is some proposition describing the “rational” options and τ is a model transformer (e.g., public announcement, radical or conservative upgrade).

Two questions are important for the analysis of this process. First, what type of transformations are the players using? Second, where do the propositions D_i come from?

Here is a baseline result from [18]. Consider a propositional formula Rat_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 1 (van Benthem [18]). *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 $\bigwedge_i \text{Rat}_i$ for the players stabilize at a submodel whose domain contains w .*

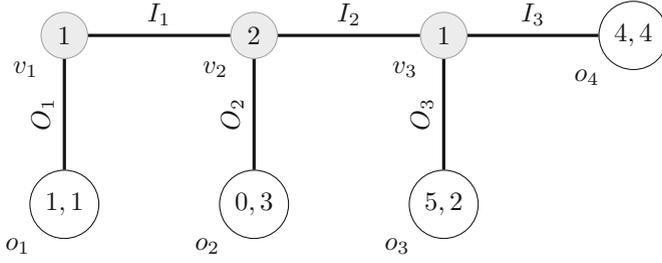
This theorem gives a precise sense of how the process of iteratively removing strictly dominated strategies can be viewed as a process of deliberation (cf. the discussion in Sect. 2.1). See [4] for a generalization of this theorem focusing on arbitrary “optimality” propositions satisfying a monotonicity property and arbitrary games. A related analysis can be found in [59], which provides an in-depth study of the upgrade mechanisms that match game-theoretic analyses.

5.3 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 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 describing the temporal structure of the game situation: Formally, T consists of a set of nodes and an immediate successor relation \succrightarrow . Let O denote the set of leaves (nodes without any successors) and V the remaining nodes. 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 . Let \rightsquigarrow be the transitive closure of \succrightarrow .
- τ is a turn function assigning a player to each node $v \in V$ (let $V_i = \{v \in V \mid \tau(v) = i\}$).
- $u_i : O \rightarrow \mathbb{R}$ is the utility function for player i assigning real numbers to outcome nodes.

The following is an example of an extensive game:

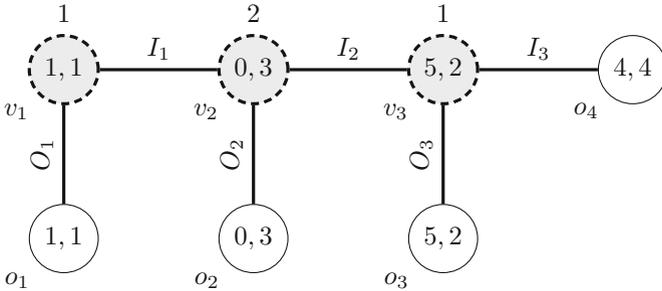


This is an extensive game with $V = \{v_1, v_2, v_3\}$, $O = \{o_1, o_2, o_3, o_4\}$, $\tau(v_1) = \tau(v_3) = 1$ and $\tau(v_2) = 2$, and, for example, $u_1(o_2) = 0$ and $u_2(o_2) = 3$. Furthermore, we have, for example, $v_1 \rightsquigarrow o_1$, $v_1 \rightsquigarrow o_4$, and $A(v_1) = \{I_1, O_1\}$.

A **strategy** for player i in an extensive game 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 $\sigma(v_1) = O_1$ and $\sigma(v_3) = I_3$ 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 “backward induction (BI) algorithm”:

BI Algorithm: At terminal nodes, players already have the nodes marked with their utilities. At a non-terminal node n , once all daughters are marked, the node is marked as follows: determine whose turn it is to move at n and find the daughter d that has the highest utility for that player. Copy the utilities from d onto n .

In the extensive game given above, the BI algorithm leads to the following markings:



The BI strategy for player 1 is $\sigma(v_1) = O_1$, $\sigma(v_3) = O_3$ and for player 2 it 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*).

Much has been written about backward induction and whether it follows from the assumption that there is common knowledge (or common belief) that all players are *rational*²⁴. In the remainder of this section, I explain how epistemic-plausibility models and the model transformations defined above can make this

²⁴ The key papers include [8, 9, 16, 39, 75]. See [61] for a complete survey of the literature.

more precise. The first step is to describe what the players believe about the strategies followed in an extensive game and how these beliefs may change during the play of the game. I sidestep a number of delicate issues in the discussion below (see [24] for a clear exposition). My focus is on the players’ beliefs about which outcome of the game (i.e., the terminal nodes) will be realized.

Suppose that $a, a' \in A(v)$ for some $v \in V_i$. We say **move a strictly dominates move a' in beliefs** (given some epistemic-plausibility model), provided that all of the most plausible outcomes reachable by playing a at v are preferred to all the most plausible outcomes reachable by playing a' . Consider an initial epistemic-plausibility model in which the states are the four outcomes $\{o_1, o_2, o_3, o_4\}$, and both players consider all outcomes equally plausible (I write $w \approx v$ if w and v are equally plausible—i.e., $w \succeq v$ and $v \succeq w$). Then, at v_2 , O_2 is not strictly dominated over I_2 in beliefs 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, since player 1 prefers o_3 to o_4 , O_3 strictly dominates I_3 in beliefs. Suppose that R is interpreted as “no player chooses an action that is strictly dominated in beliefs”. Thus, in the initial model, in which all four outcomes are equally plausible, the interpretation of R is $\{o_1, o_2, o_3\}$. We can now ask what happens to the initial model if this formula R is iteratively updated (for example, using radical upgrade).

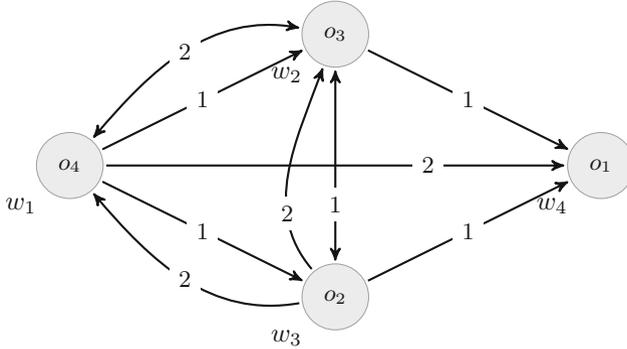
$$\begin{array}{ccc}
 o_1 \approx o_2 \approx o_3 \approx o_4 & \xrightarrow{\uparrow R} & o_1 \approx o_2 \approx o_3 \prec o_4 & \xrightarrow{\uparrow R} & o_1 \approx o_2 \prec o_3 \prec o_4 \\
 & & & & \swarrow \uparrow R \\
 & & & & o_1 \prec o_2 \prec o_3 \prec o_4
 \end{array}$$

This sequence of radical upgrades is intended to represent the “pre-play” deliberation leading to a model in which there is common belief that the outcome of the game will be o_1 . But, what justifies both players deliberating in this way to a common epistemic-plausibility model?

The correctness 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 that will give them lower payoffs). But there is a potential problem: 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 backward induction argument for this is based on what the players *would* do if player 1 chose I_1 . But, if player 1 did, in fact, choose I_1 , then common knowledge of rationality is violated (player 1’s choice would be “irrational”). Thus, it seems that common knowledge of rationality, alone, cannot be used to show that the players will make choices consistent with the backward induction path. An additional assumption about how the players’ beliefs may change during the course of the game is needed. The underlying assumption is that the 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 “unwaveringly optimistic”. I briefly discuss the approach from [15] since it touches on

a number of issues raised in this chapter. 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 \mathbf{o}_i for each outcome o_i (assume that \mathbf{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:

$$v := \bigvee_{v \rightsquigarrow o} \mathbf{o}.$$

In the above model, both players 1 and 2 believe that o_1 is the outcome that will be realized, and the players initially rule out none of the possible outcomes. That is, the model satisfies the “open future” assumption of [15] (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 $B_1^{v_1} \mathbf{o}_1$ and $B_1^{v_3} \mathbf{o}_3$ are true in the above model. For player 2, $B_2^{v_2} (\mathbf{o}_3 \vee \mathbf{o}_4)$ is true in the above model, which implies that player 2 plans to choose action I_2 at node v_2 .

The dynamics of actual play is then modeled as a sequence of public announcements (cf. Definition 4). The players' beliefs 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:

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

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* [11]: $\mathcal{M}, w \models [!]\varphi$ provided that, for all formulas²⁵ ψ , if $\mathcal{M}, w \models \psi$ then $\mathcal{M}, w \models [!\psi]\varphi$. Then, there is *common stable*

²⁵ Strictly speaking, it is all *epistemic* formulas. The important point is to not include formulas with the $[!]$ operator in them.

belief in φ provided that $[!]C^B\varphi$ is true, where $C^B\varphi$ is intended to mean that there is common belief in φ (cf. Sect. 2.2). The key result is:

Theorem 2 (Baltag, Smets and Zvesper [15]). *Common knowledge of the game structure, and of open future and common stable belief in dynamic rationality, together, imply common belief in the backward induction outcome.*

6 Concluding Remarks

This chapter has not focused on strategies *per se*, but, rather, on the process of “rational deliberation” that leads players to adopt a particular “plan of action”. Developing formal models of this process is an important and rich area of research for anyone interested in the foundations of game theory.

The economist’s predilection for equilibria frequently arises from the belief that some underlying dynamic process (often suppressed in formal models) moves a system to a point from which it moves no further.

[22, pg. 1008]

Many readers may have been expecting a formal account of the players’ *practical reasoning* in game-theoretic situations. Instead, this chapter presented three different frameworks in which the “underlying dynamic process” mentioned in the above quote is made explicit. None of the frameworks discussed in this chapter are intended to model the players’ practical reasoning. Rather, they describe deliberation in terms of a sequence of belief changes about what the players are doing or what their opponents may be thinking. This raises an important question: In what sense do the frameworks introduced in this chapter describe the players’ *strategic reasoning*? I will not attempt a complete answer to this question here. Instead, I conclude with brief discussions of two related questions.

6.1 What Are the Differences and Similarities Between the Different Models of Strategic Reasoning?

The three frameworks presented in this paper offer different perspectives on the standard game-theoretic analysis of strategic situations. To compare and contrast these different formal frameworks, I will illustrate the different perspectives on the following game from [70, Example 8, pg. 305]:

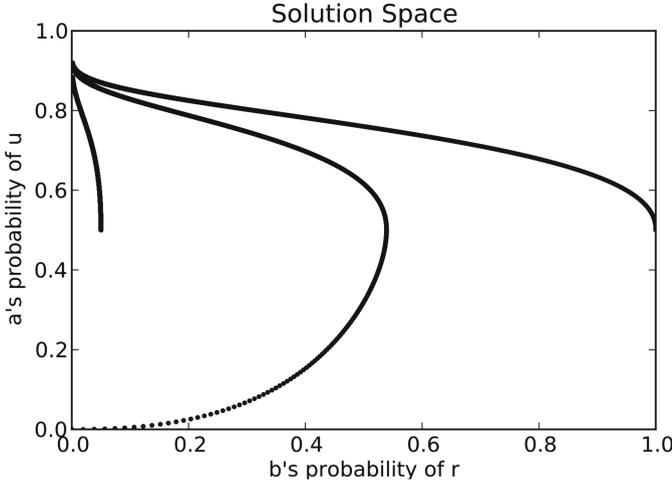
		Bob	
		l	r
Ann	u	1, 1	1, 0
	d	1, 0	0, 1

In the above game, d is weakly dominated by u for Ann. If Bob knows that Ann is rational (in the sense that she will not choose a weakly dominated strategy), then he can rule out option d . In the smaller game, action r is now strictly dominated by l for Bob. If Ann knows that Bob is rational and that Bob knows that she is rational (and so, rules out option d), then she can rule out option r . Assuming that the above reasoning is transparent to both Ann and Bob, it is common knowledge that Ann will play u and Bob will play l . But now, what is the reason for Bob to rule out the possibility that Ann will play d ? He knows that Ann knows that he is going to play l , and both u and d maximize Ann's expected utility with respect to the belief that Bob will play l .

Many authors have pointed out puzzles surrounding an epistemic analysis of iterated removal of weakly dominated strategies [5, 27, 70]. The central issue is that the assumption of common knowledge of rationality seems to conflict with the logic of iteratively removing weakly dominated strategies. The models introduced in this paper each provide a unique perspective on this issue. Note that the idea is not to provide a new "epistemic foundation" for iterated removal of weakly dominated strategies. Both [27] and [41] have convincing results here. Rather, the goal is to offer a different perspective on the existing epistemic analyses.

I start with Skyrms' model of rational deliberation from Sect. 3. There are two Nash equilibria: the pure strategy Nash equilibrium (u, l) and the mixed Nash equilibrium, where Ann plays u and d each with probability 0.5 and Bob plays strategy l . Rational deliberation with any dynamical rule that "seeks the good" (such as the Nash dynamics) is guaranteed to lead the players to one of the two equilibria. However, there is an important difference between the two Nash equilibria from the point of view of rational deliberators. Through deliberation, the players will *almost always* end up at the pure-strategy equilibrium. That is, unless the players start deliberating at the mixed-strategy Nash equilibrium, deliberation will lead the players to the pure-strategy equilibrium. This makes sense since playing u will always give a greater expected utility for Ann than any mixed strategy, as long as there is a chance (no matter how small) that Bob will play r . I can illustrate this point by showing the deliberational path that is generated if the players start from the following states of indecision: (1) Ann is playing d with probability 1 and Bob is playing l with probability 1; (2) Ann is playing u and l with probability 0.5 and Bob is playing l with probability 0.95; and (3) Ann is playing u with probability 0.5 and Bob is playing r with probability 0.5²⁶.

²⁶ These graphs were generated by a python program using a satisficing value of 0.001 and an index of caution of 50. The reason that the simulations stopped before reaching the pure Nash equilibrium is because the simulation is designed so that deliberation ends when the covetabilities fall below the satisficing value.



The second perspective comes from the reasoning-based expected utility procedure discussed in Sect. 4. For Ann, 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. Thus, the procedure stabilizes in the second round without categorizing any of Bob’s strategies or Ann’s strategy d . So, u is identified as a “good” strategy, but d is not classified as a “bad” strategy. Furthermore, neither of Bob’s strategies can be classified as “good” or “bad”.

Finally, I turn to the approach outlined in Sect. 5. An analysis of this game is discussed in [59]. In that paper, it is shown that certain deliberational sequences for the above game do not stabilize. Of course, whether a deliberational sequence stabilizes depends crucially on which model transformations are used. Indeed, a new model transformation, “suspend judgement”, is used in [59] to construct a deliberational sequence that does not stabilize. The general conclusion is that the players may not be able to deliberate their way to an informational context in which there is common knowledge of rationality (where rationality includes the assumption that players do not play weakly dominated strategies).

Each of the different frameworks offers a different perspective on strategic reasoning in games. The perspectives are not *competing*; rather, they highlight different aspects of what it means to reason strategically. However, more work is needed to precisely characterize the similarities and differences between these different models of rational deliberation in games. Such a comprehensive comparison will be left for another paper.

6.2 What Role Do Higher-Order Beliefs Play in a General Theory of Rational Decision Making in Game Situations?

Each model of deliberation discussed in this chapter either implicitly or explicitly made assumptions about the players’ *higher-order* beliefs (see Sect. 2.2). In the end, I am interested only in what (rational) players are going to do. This, in

turn, depends only on what the players believe the other players are going to do. A player's belief about what her opponents are thinking is relevant only because they shape the players first-order beliefs about what her opponents are going to do. Kadane and Larkey explain the issue nicely:

“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.*”

[46, pg. 239, my emphasis]

A theory of rational decision making in game situations need not *require* that a player considers *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. [47, 64]).

There is quite a lot of experimental work about whether or not humans take into account even second-order beliefs (e.g., a belief about their opponents' beliefs) in game situations (see, for example, [30, 44, 73]). It is beyond the scope of this chapter to survey this literature here (see [29] for an excellent overview). 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 deliberation in games (cf. [53, 54, 77]).

* * * * *

A general theory of rational deliberation for game and decision theory is a broad topic. It is beyond the scope of this chapter to discuss the many different aspects and competing perspectives on such a theory²⁷. A completely developed theory will have both a normative component (What are the normative principles that guide the players' thinking about what they should do?) 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: “[T]he fundamental insight of game theory [is] that a rational player must take into account that the players reason

²⁷ Interested readers are referred to [72] (especially Chap. 7), and [35, 50, 67] for broader discussions.

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.

References

1. Alexander, J.M.: Local interactions and the dynamics of rational deliberation. *Philos. Stud.* **147**(1), 103–121 (2010)
2. Apt, K.R.: A primer on strategic games. In: Apt, K.R., Grädel, E. (eds.) *Lectures in Game Theory for Computer Scientists*, pp. 1–33. Cambridge University Press, Cambridge (2011)
3. Apt, K.R., Zvesper, J.A.: Public announcements in strategic games with arbitrary strategy sets. In: *Proceedings of LOFT 2010* (2010)
4. Apt, K.R., Zvesper, J.A.: Public announcements in strategic games with arbitrary strategy sets. *CoRR* (2010). <http://arxiv.org/abs/1012.5173>
5. Asheim, G., Dufwenberg, M.: Admissibility and common belief. *Games Econ. Behav.* **42**, 208–234 (2003)
6. Aumann, R.: Agreeing to disagree. *Ann. Stat.* **4**, 1236–1239 (1976)
7. Aumann, R.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55**(1), 1–18 (1987)
8. Aumann, R.: Backward induction and common knowledge of rationality. *Game Econ. Behav.* **8**, 6–19 (1995)
9. Aumann, R.: On the centipede game. *Game Econ. Behav.* **23**, 97–105 (1998)
10. Aumann, R., Brandenburger, A.: Epistemic conditions for Nash equilibrium. *Econometrica* **63**, 1161–1180 (1995)
11. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., De Lima, T.: ‘Knowable’ as ‘known after an announcement’. *Rev. Symb. Log.* **1**(3), 305–334 (2008)
12. 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*, pp. 187–190, ACM (2011)
13. Baltag, A., Smets, S.: ESSLLI 2009 course: dynamic logics for interactive belief revision (2009). Slides available online at <http://alexandru.tiddlyspot.com/#%5B%5BESSLLI09%20COURSE%5D%5D>
14. 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)
15. Baltag, A., Smets, S., Zvesper, J.A.: Keep ‘hoping’ for rationality: A solution to the backwards induction paradox. *Synthese* **169**, 301–333 (2009)
16. Battigalli, P., Siniscalchi, M.: Strong belief and forward induction reasoning. *J. Econ. Theor.* **105**, 356–391 (2002)
17. van Benthem, J.: Dynamic logic for belief revision. *J. Appl. Non-Class. Log.* **14**(2), 129–155 (2004)
18. van Benthem, J.: Rational dynamics and epistemic logic in games. *Int. Game Theor. Rev.* **9**(1), 13–45 (2007)
19. van Benthem, J.: *Logical Dynamics of Information and Interaction*. Cambridge University Press, Cambridge (2011)
20. van Benthem, J., Gheerbrant, A.: Game solution, epistemic dynamics and fixed-point logics. *Fundam. Inform.* **100**, 1–23 (2010)

21. van Benthem, J., Pacuit, E., Roy, O.: Towards a theory of play: A logical perspective on games and interaction. *Games* **2**(1), 52–86 (2011)
22. Bernheim, B.D.: Rationalizable strategic behavior. *Econometrica* **52**(4), 1007–1028 (1984)
23. Board, O.: Dynamic interactive epistemology. *Games Econ. Behav.* **49**, 49–80 (2004)
24. Bonanno, G.: Reasoning about strategies and rational play in dynamic games. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) *Models of Strategic Reasoning*. LNCS, vol. 8972, pp. 34–62. Springer, Heidelberg (2015)
25. Boutilier, C.: Conditional logics for default reasoning and belief revision. Ph.D. thesis, University of Toronto (1992)
26. Brandenburger, A.: The power of paradox: some recent developments in interactive epistemology. *Int. J. Game Theor.* **35**, 465–492 (2007)
27. Brandenburger, A., Friedenberg, A., Keisler, H.J.: Admissibility in games. *Econometrica* **76**(2), 307–352 (2008)
28. Chwe, M.S.-Y.: *Rational Ritual*. Princeton University Press, Princeton (2001)
29. Colman, A.: Cooperation, psychological game theory, and limitations of rationality in social interactions. *Behav. Brain Sci.* **26**, 139–198 (2003)
30. Colman, A.: Depth of strategic reasoning in games. *TRENDS Cogn. Sci.* **7**(1), 2–4 (2003)
31. Cubitt, R.P., Sugden, R.: Common knowledge, salience and convention: A reconstruction of David Lewis’ game theory. *Econ. Philos.* **19**(2), 175–210 (2003)
32. Cubitt, R.P., Sugden, R.: The reasoning-based expected utility procedure. *Games Econ. Behav.* **71**(2), 328–338 (2011)
33. Cubitt, R.P., Sugden, R.: Common reasoning in games: A Lewisian analysis of common knowledge of rationality. *Econ. Philos.* **30**(03), 285–329 (2014)
34. van Ditmarsch, H., van Eijck, J., Verbrugge, R.: Common knowledge and common belief. In: van Eijck, J., Verbrugge, R. (eds.) *Discourses on Social Software*, pp. 99–122. Amsterdam University Press, Amsterdam (2009)
35. Douven, I.: Decision theory and the rationality of further deliberation. *Econ. Philos.* **18**(2), 303–328 (2002)
36. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: *Reasoning about Knowledge*. The MIT Press, Cambridge (1995)
37. Gerbrandy, J.: *Bisimulations on planet Kripke*. Ph.D. thesis, University of Amsterdam (1999)
38. Gierasimczuk, N.: *Knowing one’s limits: Logical analysis of inductive inference*. Ph.D. thesis, Institute for Logic, Language and Information, University of Amsterdam (2011)
39. Halpern, J.: Substantive rationality and backward induction. *Games Econ. Behav.* **37**(2), 425–435 (2001)
40. Halpern, J., Moses, Y.: Knowledge and common knowledge in a distributed environment. *J. ACM* **37**(3), 549–587 (1990)
41. 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 Knowledge*, pp. 146–155 (2009)
42. Harsanyi, J.: The tracing procedure: a Bayesian approach to defining a solution for n -person noncooperative games. *Int. J. Game Theor.* **4**, 61–94 (1975)
43. Harsanyi, J., Selten, R.: *A General Theory of Equilibrium Selection in Games*. The MIT Press, Cambridge (1988)
44. Hedden, T., Zhang, J.: What do you think I think you think? strategic reasoning in matrix games. *Cognition* **85**, 1–36 (2002)

45. Jeffrey, R.: Review of the dynamics of rational deliberation by Brian Skyrms. *Philos. Phenomenol. Res.* **52**(3), 734–737 (1992)
46. Kadane, J.B., Larkey, P.D.: Subjective probability and the theory of games. *Manag. Sci.* **28**(2), 113–120 (1982)
47. Kets, W.: Bounded reasoning and higher-order uncertainty. Working paper (2010)
48. Lamarre, P., Shoham, Y.: Knowledge, certainty, belief and conditionalisation. In: *Proceedings of the International Conference on Knowledge Representation and Reasoning*, pp. 415–424 (1994)
49. Leitgeb, H.: Beliefs in conditionals vs. conditional beliefs. *Topoi* **26**(1), 115–132 (2007)
50. Levi, I.: Feasibility. In: Bicchieri, C., Chiara, L.D. (eds.) *Knowledge, Belief and Strategic Interaction*, pp. 1–20. Cambridge University Press, Cambridge (1992)
51. Lewis, D.K.: *Counterfactuals*. Harvard University Press, Cambridge (1973)
52. Leyton-Brown, K., Shoham, Y.: *Essentials of Game Theory: A Concise Multidisciplinary Introduction*. Morgan & Claypool Publishers, San Rafael (2008)
53. Meijering, B., van Rijn, H., Taatgen, N., Verbrugge, R.: I do know what you think I think: Second-order social reasoning is not that difficult. In: *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*, pp. 1423–1428 (2010)
54. Meijering, B., van Rijn, H., Taatgen, N., Verbrugge, R.: What eye movements can tell about theory of mind in a strategic game. *PLoS ONE* **7**(9), e45961 (2012)
55. Monderer, D., Samet, D.: Approximating common knowledge with common beliefs. *Games Econ. Behav.* **1**, 170–190 (1989)
56. Morgenbesser, S., Ullmann-Margalit, E.: Picking and choosing. *Soc. Res.* **44**(4), 757–785 (1977)
57. Pacuit, E.: Dynamic epistemic logic I: Modeling knowledge and beliefs. *Philos. Compass* **8**(9), 798–814 (2013)
58. Pacuit, E.: Dynamic epistemic logic II: Logics of information change. *Philos. Compass* **8**(9), 815–833 (2013)
59. Pacuit, E., Roy, O.: A dynamic analysis of interactive rationality. In: Ju, S., Lang, J., van Ditmarsch, H. (eds.) *LORI 2011*. LNCS, vol. 6953, pp. 244–257. Springer, Heidelberg (2011)
60. Pearce, D.G.: Rationalizable strategic behavior and the problem of perfection. *Econometrica* **52**(4), 1029–1050 (1984)
61. 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*, pp. 159–193. *Texts in Logic and Games*, Amsterdam University Press (2007)
62. Perea, A.: A one-person doxastic characterization of Nash strategies. *Synthese* **158**, 251–271 (2007)
63. Perea, A.: *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press, Cambridge (2012)
64. Perea, A.: Finite reasoning procedures for dynamic games. In: van Benthem, J., Ghosh, S., Verbrugge, R. (eds.) *Models of Strategic Reasoning*. LNCS, vol. 8972, pp. 63–90. Springer, Heidelberg (2015)
65. Plaza, J.: Logics of public communications. In: Emrich, M.L., Pfeifer, M.S., Hadzikadic, M., Ras, Z.W. (eds.) *Proceedings, 4th International Symposium on Methodologies for Intelligent Systems*, pp. 201–216 (republished as [66]) (1989)
66. Plaza, J.: Logics of public communications. *Synthese* **158**(2), 165–179 (2007)
67. Rabinowicz, W.: Does practical deliberation crowd out self-prediction? *Erkenntnis* **57**, 91–122 (2002)

68. Risse, M.: What is rational about Nash equilibrium? *Synthese* **124**(3), 361–384 (2000)
69. Rubinstein, A.: Comments on the interpretation of game theory. *Econometrica* **59**(4), 909–924 (1991)
70. Samuelson, L.: Dominated strategies and common knowledge. *Games Econ. Behav.* **4**, 284–313 (1992)
71. Selten, R.: Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theor.* **4**(1), 25–55 (1975)
72. Skyrms, B.: *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge (1990)
73. Stahl, D.O., Wilson, P.W.: On players' models of other players: theory and experimental evidence. *Games Econ. Behav.* **10**, 218–254 (1995)
74. Stalnaker, R.: Knowledge, belief, and counterfactual reasoning in games. *Econ. Philos.* **12**, 133–163 (1996)
75. Stalnaker, R.: Belief revision in games: Forward and backward induction. *Math. Soc. Sci.* **36**, 31–56 (1998)
76. Vanderschraaf, P., Sillari, G.: Common knowledge. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*. Spring 2009 edition (2009)
77. Verbrugge, R.: Logic and social cognition: the facts matter, and so do computational models. *J. Philos. Log.* **38**(6), 649–680 (2009)