Varieties of strategic reasoning in games

Eric Pacuit

Department of Philosophy
University of Maryland, College Park
epacuit@umd.edu
http://pacuit.org

A crucial assumption underlying any game-theoretic analysis is that there is *common knowledge* that all the players are *rational*. Rationality here is understood in the decision-theoretic sense: The players' choices are *optimal* according to some choice rule (such as maximizing subjective expected utility). Recent work in *epistemic game theory* has focused on developing sophisticated mathematical models to study the implications of assuming that all the players are rational and this is commonly known (or commonly believed). [1] However, 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 of a game. It is also important to describe how the players were able to arrive at this crucial state of information. [2] There is now 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 (see [17] for a discussion).

There is a second reason why it is important to develop formal models of the players' process of deliberation in game situations. A number of researchers have questioned the usefulness of models that make explicit assumptions about the players' *higher-order* beliefs in game situations. In the end, we are 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*

[1] See, for example, [20, 9], and the references therein, for a survey of this literature.

[2] This general point about common knowledge was already appreciated by David Lewis when he first formulated his notion of common knowledge [14]. See [6] for an illuminating discussion and a reconstruction of Lewis' notion of common knowledge, with applications to game theory.

> *opponent's behavior are irrelevant, and can be ignored in the analysis by integrating them out of the joint opinion.*" [12, 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. [13]). Part of the difficulty in answering this question comes from experimental work suggesting that humans do not necessarily take into account even second-order beliefs (e.g., a belief about their opponents' beliefs) in game situations (see, for example, [11] and [16, 15]). Of course, this is a *descriptive* issue, and it is very much open how such observations should be incorporated into a general theory of rational deliberation in games (cf. [5]).

In this paper, I will compare and contrast different models of rational deliberation in games. The goal is not only to develop a comprehensive comparison of the different frameworks, but also to show how explicitly representing the players' process of deliberation can shed some light on the role that higher-order information plays in the analysis of 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. Indeed, Robert Aumann and Jacques Dreze [1, pg. 81] note that: "[T]he 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".

The different models of rational deliberation in games that I will discuss in this paper include:

1. John Harsanyi's "tracing procedure": The goal of the tracing procedure is to identify a unique Nash equilibrium in any finite strategic game. Harsanyi thought of the tracing procedure as "a mathematical formalization of the process by which rational players coordinate their choices of strategies" [10].
2. Brian Skyrms' model of "dynamic deliberation": Players deliberate by calculating expected utility and then use this new information to recalculate their probabilities about what they are going to do and what they expect their opponents are going to do [21].
3. Robin Cubitt and Robert Sugden's "reasoning-based expected utility": An iterative procedure for solving strategic games that builds on David Lewis's "common modes of reasoning" [7, 8].
4. Johan van Benthem *et col.*'s "dynamic logic of games": The goal is to characterize different solution concepts as fixed points of iterated "(virtual) rationality announcements" [2, 3, 19].

Although the details of these 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 frameworks mentioned above differ in how they represent the players' "state of indecision" during the deliberation process:

1. Harsanyi does not explicitly represent the players' uncertainty about their strategy choices during his tracing procedure. Instead, the idea is to analyze a continuum of games starting from the original game and a common prior[3], in which the payoffs are replaced by expected payoffs.
2. Skyrms' model represents the players' beliefs during deliberation as a probability measure on i's set of strategies. The interpretation is that it is the *mixed strategy* that the players would adopt if deliberation ended at the current stage.
3. Cubitt and Sugden assume that the players *categorize* their strategies S_i in a game as a pair of sets (S_i^+, S_i^-) where elements of S_i^+ are the "admissible" actions, elements of S_i^- are deemed "irrational", and strategies such that $s \notin S_i^+ \cup S_i^-$ (if any) are not (yet) categorized.
4. van Benthem and colleagues use *relational models*, familiar from the computer science and philosophy literature (see [18] for an overview), to describe what the players know and believe about their own choices, their opponents' choices and their *higher-order* beliefs.

Each of the above models are intended to be a "snapshot" of the players' beliefs about their own choices and their opponents' choices during the deliberational process. The second aspect of the models of deliberation in games are the dynamical rules that are intended to describe how the players' beliefs evolve during deliberation. At each stage of the deliberation, the players determine which of their available strategies are "optimal" and which ones they ought to avoid. Typically, it is assumed that the players are guided by some decision-theoretic choice rule, such as maximizing expected utility or avoiding dominated strategies. Using the information about the players' own choices and what they expect their opponents to do, the current snapshot of the players' beliefs is *transformed* according to some fixed dynamical rule. Different types of transformations represent how confident the players are in their assessment of which of the available choices are "rational". The picture to keep in mind is:

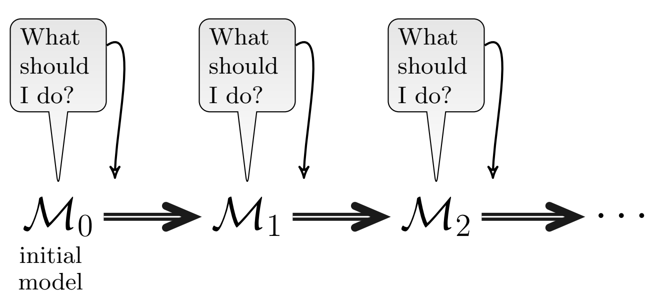

[3] Harsanyi in collaboration with Reinhard Selten offer an algorithm for constructing a common prior given the structure of the game.

Deliberation concludes when the players reach a fixed point in the above process. The goal is not to develop a formal account of the players' *practical reasoning* in game situations. Rather, it 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 different frameworks mentioned above can be compared and contrasted in terms of the specifics of the models: How are the players' beliefs represented during the deliberational process and what are the properties of the operations that are used to transform the players' beliefs? The general conclusion is that each of the frameworks offer a different perspective on strategic reasoning in games. These perspectives are not *competing*; rather, they highlight different aspects of what it means to reason strategically in interactive situations.

A second dimension along which the different models of deliberation can be compared is in terms of the outcomes that can be reached by the deliberational processes and what are the properties of the paths that lead to these outcomes. For instance, the outcomes of both Harsanyi's tracing procedure and Skryms's model of dynamic deliberation are qualitatively similar: Both procedures lead players to choose their component of a Nash equilibrium. However, in Skyrms's model, the rate of convergence depends on the players' initial beliefs in an interesting way; and this, in turn, suggests a more refined analysis of Nash equilibrium [21, pgs. 154 - 158].

Finally, moving away from the mathematical properties of the different frameworks, it is also important to focus on the underlying motivations. Here, the tracing procedure can be singled out as it is focused on identifying a *unique* rational solution to every strategic game. The other frameworks are more "personalistic": The rational outcomes of a game depends not only on the structure of the game, but also on the players' initial beliefs, which dynamical rule is being used by the players (in general, different players may be using different rules), and what exactly is commonly known about the process of deliberation.

I will conclude the paper by carefully comparing the above general approach to modeling strategic reasoning in interactive situations with a recent contribution from the cognitive science literature. In [22], Stuhlmüller and Goodman show how probabilistic programs can explicitly represent conditioning which enables them to describe reasoning about others' reasoning using *nested conditioning*. These probabilistic programs can then be used to capture the "back-and-forth" reasoning that is characteristic of many game-theoretic situations (cf. [4] for an alternative approach).

References

1. Aumann, R., Dreze, J.: Rational expectations in games. American Economic Review, 98(1), 72-86 (2008)
2. Baltag, A., Smets, S., Zvesper, J.A.: Keep 'hoping' for rationality: A solution to the backwards induction paradox. Synthese 169, 301–333 (2009)
3. van Benthem, J.: Logic in Games. The MIT Press (2014)

4. Bergwerff, G., Meijering, B., Szymanik, J., Verbrugge, R, Wierda, S.M.: Computational and algorithmic models of strategies in turn-based games, Proceedings of the 36th Annual Conference of the Cognitive Science Society, 2014. http://jakubszymanik.com/papers/CogSci2014_23.04_final.pdf
5. Colman, A.: Cooperation, psychological game theory, and limitations of rationality in social interactions. Behavioral and Brain Sciences 26, 139 – 198 (2003)
6. Cubitt, R., Sugden, R.: Common knowledge, salience and convention: A reconstruction of David Lewis' game theory. Economics and Philosophy 19(2), 175–210 (2003)
7. Cubitt, R., Sugden, R.: The reasoning-based expected utility procedure. *Games and Economic Behavior*, 71(2), 328 – 338 (2011)
8. Cubitt, R., Sugden, R.: Common reasoning in games: A Lewisian analysis of common knowledge of rationality (2011), ceDEx Discussion Paper
9. Dekel, E., Siniscalchi, M.: Epistemic game theory. In: Handbook of Game Theory. Forthcoming (2014)
10. Harsanyi, J.: The tracing procedure: A Bayesian approach to defining a solution for n-person noncooperative games. International Journal of Game Theory 4, 61 – 94 (1975)
11. Hedden, T., Zhang, J.: What do you think I think you think? Strategic reasoning in matrix games. Cognition 85, 1 – 36 (2002)
12. Kadane, J.B., Larkey, P.D.: Subjective probability and the theory of games. Management Science 28(2), 113 – 120 (1982)
13. Kets, W.: Bounded reasoning and higher-order uncertainty (2012). Available at SSRN: http://ssrn.com/abstract=2116626
14. Lewis, D.: Convention. Harvard University Press (1969)
15. 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), 1 – 8 (2012)
16. Meijering, B., van Rijn, H., Taatgen, N., Verbrugge, R.: I do know what you think I think: Second-order social reasoning. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*, 1423 – 1428 (2010)
17. Pacuit, E.: Strategic reasoning in games. (2014). Forthcoming in *Reasoning about Strategies*, J. van Benthem, S. Ghosh, and R. Verbrugge (eds.).
18. Pacuit, E.: Dynamic epistemic logic I: Modeling knowledge and beliefs. *Philosophy Compass*, 8(9), 798 – 814 (2013)
19. Pacuit, E., Roy, O.: A dynamic analysis of interactive rationality. In: Proceedings of the Third International Conference on Logic, Rationality, and Interaction. pp. 244–257. LORI'11, Springer-Verlag (2011), http://dl.acm.org/citation.cfm?id=2050423.2050441
20. Perea, A.: Epistemic Game Theory: Reasoning and Choice. Cambridge University Press (2012)
21. Skyrms, B.: The Dynamics of Rational Deliberation. Harvard University Press (1990)
22. Stuhlmüller, A., Goodman, N.: Reasoning about reasoning by nested conditioning: Modeling theory of mind with probabilistic programs. Cognitive Systems Research 28, 80 – 99 (2014)