

How Many Levels Do Players Reason? Observational Challenges and a Solution*

[Extended Abstract] *

Adam Brandenburger
NYU Stern School of Business
New York University
adam.brandenburger@stern.nyu.edu

Alex Danieli
alex.daniels85@gmail.com

Amanda Friedenberg
Department of Economics
Arizona State University
amanda.friedenberg@asu.edu

Keywords

Epistemic Game Theory, Bounded Reasoning, Level- m Reasoning, Context of the Game

1. INTRODUCTION

Interactive reasoning is an important aspect of how players behave. To determine whether a particular course of action is good or bad, Ann may need to form a theory about Bob's play of the game. In forming such a theory, she may reason that Bob is 'strategically sophisticated' — if so, she may reason that Bob forms a belief about her own play to determine whether a particular strategy is good or bad for him. That is, Ann may want to form a second-order theory about Bob's play of the game. Of course, Ann may then reason that Bob uses a second-order theory to choose his strategy. In this case, Ann may want to form a third-order theory about Bob's play of the game. And so on.

How many levels of reasoning do players undertake? We address this question for the case where the players' processes of reasoning are not observable. Instead, the researcher observes only the behavior of the players — or, perhaps, only the outcome of the game. We ask: Can the researcher use this information to identify — or provide bounds on — levels of reasoning?

2. A MOTIVATING EXAMPLE

Figure 1 depicts the game of Battle-of-the-Sexes with an Outside Option: Ann can either choose to exercise an outside option or choose to play Battle-of-the-Sexes with Bob.

*This abstract was prepared for TARK 2015. We thank Jessy Hsieh, Terri Kneeland, Xiaomin Li, and Elliot Lipnowski for valuable comments. We also thank audiences at the University of Bristol, the 2013 Asian Meeting of the Econometric Society, National University of Singapore, and the 2013 Workshop on Logical Structure of Correlated Information Change for important input. Adam Brandenburger thanks the NYU Stern School of Business for financial support. Amanda Friedenberg thanks NSF SES-1358008 for financial support and the UCL Economics Department for its unbounded hospitality.

The standard argument is that, if the players are 'strategically sophisticated,' then Ann will choose *In-U* and Bob will choose *L*: The strategy *In-D* is dominated for Ann by the outside option. Thus, if Ann does play *In*, Bob should reason that she will play *U*, since *In-U* is undominated. In this case, his best response is to play *L*. Ann should understand that Bob will reason this way and expect Bob to play *L*. With this expectation, Ann should play *In-U*.




Figure 1: Battle-of-the-Sexes with an Outside Option

Thus, there appears to be a clear prediction: Ann will play *In-U* and Bob will play *L*. Yet, in the lab, *Out* is played with significant frequency. (See, inter alia, [4] and [2].) One might then draw the conclusion that there is limited reasoning: If Bob engages in one level of reasoning, he may choose either *L* or *R* depending on what he believes about Ann's play. So, if Ann engages in two levels of reasoning, she might choose to play *Out*. It is only if Ann reasons three (or more) levels that she would not play *Out*.

But, in fact, this behavior need not be an artifact of limited reasoning. There is another reason why Ann might choose to play *Out* — one based on the idea that there is a "context" to the game. In particular, suppose that it is commonly understood that "Bob is a bully" and, so, whenever a Battle-of-the-Sexes game is played, he attempts to go for

his best option and plays R . To be specific, suppose:

- Bully-1: at each information set, Ann believes that Bob plays R ,
- Bully-2: at each information set, Bob believes “Bully-1,”
- Bully-3: at each information set, Ann believes “Bully-2.”
- ⋮

If this is the context under which the game is played, Ann may play *Out*, even if she reasons three levels. In fact, she may play *Out*, even if she engages in common reasoning about the play of the game.

Here is the basic idea: When the game is played in this context, if Ann reasons at least one level, then she must play *Out*. (Ex ante, she expects that Bob will play R , and so *Out* is the unique best choice.) If Bob reasons at least two levels, then he will reason that Ann reasons at least one level, provided he has not observed information that contradicts this hypothesis. Thus, if Bob reasons at least two levels, he must begin the game believing Ann plays *Out*. Conditional upon Ann’s playing *In*, he is forced to reason that Ann does not reason one level. So, Bob can reason two levels and conclude (after observing *In*) that Ann is playing *In-U*, but he can also conclude that Ann is playing *In-D*. Thus, if Bob reasons two levels, he can play either of L or R . It follows that believing that Bob play will R is consistent with three levels of reasoning for Ann.

In light of the above, it appears premature to conclude that observing the play of *Out* indicates that Ann is necessarily only a level m -reasoner, for some $m \leq 2$. This paper focuses on the case where the researcher cannot observe players’ actual beliefs or which beliefs players consider possible. Thus, we seek an answer to our question that is independent of the context of the game.

3. APPROACH AND CHALLENGE

In keeping with what we have just seen, we will need to describe what beliefs players do vs. do not consider possible in a particular game. The device we will use to describe these beliefs is an **epistemic type structure**, denoted by \mathcal{T} . An epistemic type structure will consist of a set of types for each player, where a type for a player will describe that player’s hierarchies of beliefs about the play of the game. Different type structures are associated with different events which are commonly believed. For instance, for Battle-of-the-Sexes with an Outside Option, there is a type structure where the event “Bob is a bully” is commonly believed and there are other type structures where it is not.

Call a pair (Γ, \mathcal{T}) an **epistemic game**. For a given epistemic game (Γ, \mathcal{T}) , we can define the set of strategy-type pairs which are consistent with m levels of reasoning, to be denoted $R^m(\mathcal{T})$. Refer to Figure 2, that illustrates these sets. Specifically:

- (0) The set $R^0(\mathcal{T})$ is the set of all strategy-type pairs. This set captures **level-0 reasoning**, since there is no requirement on reasoning.
- (1) The set $R^1(\mathcal{T})$ is the set of strategy-type pairs where the players’ strategies are optimal given their belief (i.e., type). This set captures **level-1 reasoning**.




Figure 2: Level- m Reasoning

- ⋮
- (m) The set $R^{m+1}(\mathcal{T})$ is the set of strategy-type pairs in $R^m(\mathcal{T})$ where each player reasons that the other players engage in **level- m reasoning**. This set captures **level- $(m + 1)$ reasoning**.
- ⋮
- (∞) The set $R^\infty(\mathcal{T})$ is the set of strategy-type pairs in $R^m(\mathcal{T})$ for all m . This set captures **level- ∞ or common reasoning**.

The sets $R^m(\mathcal{T})$ depend not only on the game Γ but also on the type structure \mathcal{T} . This fits with our informal analysis of Battle-of-the-Sexes with an Outside Option, where the behavior of a level-3 reasoner depended on whether or not the event “Bob is a bully” is commonly understood.

Observe that the sets $R^0(\mathcal{T}), \dots, R^m(\mathcal{T}), \dots$ are decreasing. This reflects the fact that, if players reason at least $(m + 1)$ levels, then they reason at least m levels. As a consequence, we will not be able to identify the minimum number of levels of reasoning by observing behavior alone.¹ Instead, we seek to identify the maximum number of levels of reasoning consistent with observed behavior.

The goal then is to identify when a strategy is consistent with m but not $(m + 1)$ levels of reasoning. Toward this end, we seek to construct an ordered partition of the strategy set, denoted by $\mathcal{L} = \{L^0, L^1, \dots, L^m, \dots, L^\infty\}$, that satisfies the following criteria: When m is finite, $s \in L^m$ if it is

1. consistent with level- m reasoning in some epistemic game (Γ, \mathcal{T}^*) , but
2. inconsistent with level- $(m + 1)$ reasoning in any epistemic game (Γ, \mathcal{T}) .

When m is infinite, $s \in L^m$ if it is consistent with level- ∞ reasoning in some epistemic game (Γ, \mathcal{T}^*) . We refer to players who choose $s \in L^m$ as **Level- m Reasoners** (or **L^m -Reasoners**).

¹As the “level- k ” and “cognitive hierarchy” literatures make clear, it may be possible to identify the minimum number of levels of reasoning by making auxiliary assumptions about behavior or beliefs. See the discussion below.

For the case of a matrix, standard results give that $s \in L^m$ if and only if s is m but not $(m + 1)$ -rationalizable. One might conjecture that the same is true for the tree, where we now take “rationalizability” to mean “extensive-form rationalizability” as in [8] and [1] (or, equivalently, “iterated conditional dominance,” as in [9]). This is not the case. Refer back to Battle-of-the-Sexes with an Outside Option (Figure 1). There, *Out* is consistent with two but not three rounds of extensive-form rationalizability. But, the “Bob is a bully” analysis showed that $Out \in L^\infty$.

The main paper provides a novel procedure which serves to construct the partition \mathcal{L} in a finite number of steps. A major challenge in providing such a procedure is the fact that the definition of L^m makes reference to all type structures. For a given finite tree, there are infinitely many associated type structures and, therefore, searching across all type structures would appear to be an infinite task. The Main Theorem in the paper provides a way to side-step this difficulty for generic games. It characterizes the set of strategies consistent with m levels of reasoning as a property of the game alone — without reference to any type structure. It goes on to show how to implement the procedure in a “simple” manner.

4. UPPER BOUND ON REASONING

The main paper identifies the maximum number of levels of reasoning consistent with observed behavior. The focus on this upper bound is inevitable, absent making auxiliary assumptions on behavior and/or beliefs. (The “level- k ” and “cognitive hierarchy” literatures, e.g., [7], [10], [5], [3], etc., obtain exact identification by imposing such restrictions.) But, there are also good reasons why this upper bound is of interest. To the extent that the researcher is interested in using the number of levels of reasoning as an empirical input, the researcher may care only that the player acts ‘as if’ she is an L^2 -Reasoner — even if she is, in fact, an L^1 -Reasoner or L^0 -Reasoner. A similar argument applies to higher levels of reasoning.

For certain datasets, it may be possible to distinguish an L^2 -Reasoner from an L^1 -Reasoner who acts as if she is an L^2 -Reasoner. With a large dataset, we should expect, on statistical grounds, to see more than just occasional in-sample play of strategies inconsistent with higher levels of reasoning. In the spirit of the “level- k ” and “cognitive hierarchy” literatures (e.g., [7], [10], [5], [3], etc.), the researcher may be able to design an experiment so that, even in small samples, we would expect to see play inconsistent with higher levels of reasoning. For instance, the design in [6] is based on “ring games,” where the payoff to any player depends only on his/her left-hand neighbor’s choice. So, changing payoffs to a player two steps to the left of Ann should not affect her behavior if she is a L^1 -Reasoner, but should affect her behavior if she is a L^2 -Reasoner. As a consequence, varying payoffs across the experimental session should generate distinct behavior for L^1 - and L^2 -Reasoners, thereby allowing the experimenter to identify the L^2 -Reasoner.

5. REFERENCES

[1] P. Battigalli. On Rationalizability in Extensive Games. *Journal of Economic Theory*, 74(1):40–61, 1997.

[2] J. Brandts and C. Holt. Limitations of dominance and forward induction: Experimental evidence. *Economics Letters*, 49(4):391–395, 1995.

[3] C. Camerer, T. Ho, and J. Chong. A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.

[4] R. Cooper, D. DeJong, R. Forsythe, and T. Ross. Forward induction in the battle-of-the-sexes games. *American Economic Review*, 83(5):1303–1316, 1993.

[5] M. Costa-Gomes, V. Crawford, and B. Broseta. Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica*, 69(5):1193–1235, 2001.

[6] T. Kneeland. Testing Behavioral Game Theory: Higher-Order Rationality and Consistent Beliefs. <http://terri.microeconomics.ca>, 2013.

[7] R. Nagel. Unraveling in Guessing Games: An Experimental Study. *The American Economic Review*, 85(5):1313–1326, 1995.

[8] D. Pearce. Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica*, 52(4):1029–1050, 1984.

[9] M. Shimoji and J. Watson. Conditional Dominance, Rationalizability, and Game Forms. *Journal of Economic Theory*, 83(2):161–195, 1998.

[10] D. Stahl and P. Wilson. On Players’ Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.