Ashton T. Sperry-Taylor

Bounded Rationality in the Centipede Game[1]

ABSTRACT

Normative game theory unsatisfactorily explains rational behavior. Real people do not behave as predicted, and what is prescribed as rational behavior is normally unattainable in real-life. The problem is that current normative analysis does not account for people's cognitive limitations – their bounded rationality. However, this paper develops an account of bounded rationality that explains the rationality of more realistic behavior. I focus on the Centipede Game, in which boundedly rational players explore and test others' immediate behavior, until they can apply limited backward induction. The result is that the game has a solution in the form of a subjective Nash equilibrium, which boundedly rational players can possibly realize.

1. INTRODUCTION

Normative game theory unsatisfactorily explains rational behavior. Real people do not behave as predicted, and what is prescribed as rational behavior is normally unattainable in real-life (McKelvey and Palfrey 1992; Stahl and Wilson 1995; Hedden and Zhang 2002; Camerer 2003; Meijering et al. 2010). The problem is that current normative analysis does not account for people's cognitive limitations – their bounded rationality. Reinhard Selten submits that current normative analysis

> is largely based on an unrealistic picture of human decision making. Economic agents are portrayed as fully rational Bayesian maximizers of subjective utility. This view of economics is not based on empirical evidence, but rather on the simultaneous axiomization of utility and subjective probability. (2001, 13)

Gerd Gigerenzer submits that

> Humans and animals make inferences about unknown features of their world under constraints of limited time, limited knowledge, and limited computational capacities.

Episteme 8.3 (2011): 262–280
DOI: 10.3366/epi.2011.0021

www.eupjournals.com/epi

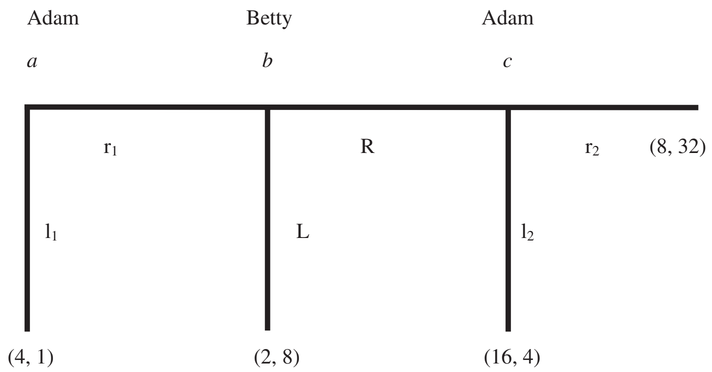

Figure 1. The Three-Stage Centipede Game. The numbers in the parentheses represent monetary payoffs. Adam's is always the left number while Betty's is the right.

> Models of rational decision making in economics, cognitive science, biology, and other fields, in contrast, tend to ignore these constraints and treat the mind as a Laplacean superintelligence equipped with unlimited resources of time, information, and computational might. (2001, 37)

Many game theorists agree with these sentiments, as if normative game theory is now recorded in the annals of failed-analysis. It receives an 'A' for effort, 'F' for accuracy.

Nonetheless, this paper develops an account of bounded rationality that explains the rationality of the behavior found in experimental research. I shall revive normative game theory by removing its idealized account of rational agents, and by extending its methodology to include agents' limited cognitive capacities. For depth, my focus is on one specific game of strategy, the Centipede Game.

Section 2 presents the traditional normative analysis of the Centipede Game, followed by a brief account of the experimental research in section 3. Sections 4 and 5 develop an account of bounded rationality that explains the possible rationality of the behavior found in experimental research. This paper then compares my account to level-*k* theory in section 6, explaining how my account compliments it. My account begins to unify the normative and behavioral disciplines.

2. THE CENTIPEDE GAME

Following Bicchieri and Antonelli (1995), the Centipede Game is a noncooperative, sequential game, depicted as a tree in figure 1, with two players, Adam and Betty. Its importance is that it models the behavior found in any dynamic relationship between individuals or parties in which there are growing potential gains from sequential exchanges, but also increasing temptations for one to capitalize on those gains and end the relationship (Camerer 2003, 218). A noncooperative game precludes binding agreements. A sequential game considers the dynamics of a

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

game: the order in which players move and the kind of information from which they strategize.

Suppose Adam and Betty are both expected payoff maximizers. They also have common knowledge of each other's rationality and of the structure of the game (Adam knows that Betty knows that Adam knows that Betty knows, and so on ad infinitum). Thus, *at least* the following propositions are true concerning Adam and Betty's epistemic states:

(1) Adam is rational at stage *c*.
(2) Betty is rational at stage *b* and knows (1).
(3) Adam is rational at stage *a* and knows (2).

We can *justify* that Adam will play down at his first opportunity by backward induction, where players predict moves at the game's last stage, and then use these predictions to predict moves at the game's second to last stage. Reasoning in this fashion from the game's end to its start, each player settles on a strategy for moves throughout the game. We begin at stage *c* and reason our way backward to stage *a*. Suppose stage *c* is reached. The following proposition is true given (1):

(i) If stage *c* is reached, Adam will play l_2.

Betty knows (1) is true and thus knows Adam is rational at stage *c*. She knows that Adam will play l_2 if stage *c* is reached. She therefore will choose L at stage *b* since she is rational. The following proposition is true given (2):

(ii) If stage *b* is reached, Betty will play L.

Adam knows (2) is true and thus knows Betty is rational at stage *b*, and that she knows Adam is rational at stage *c*. He knows that Betty will play L if stage *b* is reached. He therefore will choose l_1 at stage *a* since he is rational. The following proposition is true given (3):

(iii) Adam will play l_1 at stage *a*.

Adam's strategy choice (l_1l_2) is a best response to Betty's strategy choice (L), and Betty's strategy choice (L) is a best response to Adam's strategy choice (l_1l_2). A Nash equilibrium is a profile of strategies, one for each player, where each strategy is a best response – one has no incentive to unilaterally switch one's strategy – to the other strategy. The Nash equilibrium (l_1l_2, L) is the game's solution entailed by the truths of (1), (2), and (3).

The following section introduces the experimental research on this game of strategy. Notably, subjects in experiments do not behave according to this normative analysis, generally for a greater payoff.

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

3. EXPERIMENTAL RESEARCH ON THE CENTIPEDE GAME

Experimental research contravenes what is predicted by traditional normative analysis (McKelvey and Palfrey 1992). Students ranging from Caltech undergraduates to Pasadena City College undergraduates have been subjects in experimental Centipede Games. They have a wide range of analytical ability as illustrated by their SAT-M scores. But the results are quite stable: subjects rarely play "down" at their first opportunity, and many play "across" to the last stage. The following table aggregates results from four sessions with subjects from Pasadena City College, and two sessions with subjects from Caltech (Camerer 2003, 219).

Game Type	Trials	1	2	3	4	5	6
Four Stages	1–5	0.06	0.32	0.57	0.75		
	6–10	0.08	0.49	0.75	0.82		
High Stakes	1–5	0.08	0.46	0.60	0.80		
	6–10	0.22	0.41	0.74	0.50		
Six Stages	1–5	0.00	0.06	0.18	0.43	0.75	0.81
	6–10	0.01	0.07	0.25	0.65	0.70	0.90

The first column presents two four-stage games, with the second involving greater payoffs (four-fold), and a six-stage game with its first four stages having the same payoffs as the original four-stage game. The second column compares the first five trials to the second five trials. The remaining columns present the percentage of subjects who play "down" at that stage. There are two general results one can extrapolate from the data:

Result 1: Subjects are likelier to revert to their Nash strategy as the game proceeds. Subjects begin to apply backward induction as they approach the game's end. Subjects can apply backward induction to a limited number of stages.

Result 2: Subjects are likelier to revert to their Nash strategy in games with fewer stages. Hence, players deviate from backward induction less often in the four-stage version than in the six-stage version. The simpler the game, the easier the subjects find it to apply backward induction.

This research shows that subjects in these experiments generally exhibit partial, but not complete, rejection of the Nash strategy. Subjects therefore apply limited, but not complete, backward induction.

Thus, normative game theory poorly accounts for players – very smart ones, in fact – rejecting their Nash strategy at their first opportunity.[2] The next section explains why and provides an account of bounded rationality that unifies our traditional normative analysis with the experimental results.

4. AN ACCOUNT OF BOUNDED RATIONALITY

The use of backward induction in section 2 to justify Adam and Betty's Nash strategies involves interdependent decision-making. David Lewis explains that in any problem of interdependent decision-making,

> We may acquire … or correct or corroborate whatever expectations we already have, by putting ourselves in the other fellow's shoes, to the best of our ability. If I know what you believe about the matters of fact that determine the likely effects of your alternative actions, and if I know your preferences among possible outcomes and I know that you possess a modicum of practical rationality, then I can replicate your practical reasoning to figure out what you will probably do, so that I can act appropriately. (1969/2002, 27)

By placing "one's self" in the other's shoes, I replicate your practical reasoning to better strategize accordingly. Adam and Betty's interdependent decision-making involves replicating each other's practical reasoning – their beliefs and strategies – which we can symbolize with a language containing primitive propositions R_i ($i =$ Adam, Betty), meaning "player i is rational." Thus, R_A means "Adam is rational," and R_B means "Betty is rational." There is one epistemic modality K_i operating on any proposition p, where $K_i p$ means "player i knows proposition p."

Experimental research contravenes what is predicted by normative analysis, because it is problematic to ascribe common knowledge of rationality to real people. Real people in strategic interactions do not consider indefinite iterations of 'knows'. Suppose Adam and Betty have common knowledge of each other's rationality, which entails the following epistemic states:

Adam knows:	Betty knows:
R_B	R_A
$K_B R_A$	$K_A R_B$
$K_B K_A R_B$	$K_A K_B R_A$
.	.
.	.
.	.

Their epistemic states go on indefinitely, since having common knowledge of each other's rationality goes on indefinitely. This is unrealistic to ascribe to players with limited cognitive abilities. By removing the common knowledge assumption, we limit the number of iterations of 'knows' that a player can consider, which means we limit the players' abilities to replicate each other's practical reasoning. This is the beginning of an account of bounded rationality.[3]

However, as Cristina Bicchieri (1992, 1993) explains, we can remove the common knowledge assumption and still justify the backward induction solution to the Centipede Game if each player knows what the other knows at subsequent stages. For example, only the following limited interactive knowledge of each other's rationality is required for the backward induction solution to obtain:

Adam knows:	Betty knows:
R_B	R_A
$K_B R_A$	

Combine this with their knowledge of the game depicted by figure 1, and it follows that:

(1) Adam is rational at stage *c*.
(2) Betty is rational at stage *b* and knows (1).
(3) Adam is rational at stage *a* and knows (2).

Greater iterations of 'knows' are unnecessary. Propositions (1), (2), and (3) are sufficient to obtain the backward induction solution – just follow the argument in section 2. However, true bounded rationality, the type that realistically explains the results found in behavioral experiments, must be of a limited degree *insufficient* to apply backward induction to solve the game.

A theory of rationality can either express *standards* of rationality that apply to players' acts, or *procedures* that players apply to choose acts. Take the principle of maximizing one's expected payoffs. As a standard, the principle states that an act is rational only if it maximizes one's expected payoffs. This principle, however, does not provide a procedure to maximize one's expected payoffs. Applying a procedure depends on the nature of the game. For instance, backward induction applies to sequential games but not to strategic games where players strategize and act simultaneously. Any standard of rationality can change depending on what is possible for the players' actions. Therefore, a standard of rationality can be sensitive to the players' strategic interactions and their cognitive abilities. This is helpful because it allows rationality to be attainable, adjusting to the imposed limitations. It follows that in strategic situations with boundedly rational players, any principle of rationality can cease to be a standard but remain a *goal* of rationality to which players aspire.

I thus propose that boundedly rational players *aspire* to maximize their expected payoffs, and applying the procedure in which they maximize their expected payoffs depends on their limitation to replicate each other's practical reasoning. This account of bounded rationality specifies conditions insufficient to apply backward induction to traditionally solve the sequential game.

Players can follow two basic approaches to fulfilling their goal in maximizing their expected payoffs. They can approach the goal indirectly by refraining from a present action of positive payoff for a future action with a greater than present

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

payoff. Or they can approach the goal directly by acting on a present positive payoff instead of pursuing a future action with a greater than present payoff. For example, players act on a present payoff in the ideal Centipede Game with common knowledge. The presence of common knowledge guarantees the backward induction solution. Players therefore maximize their payoffs by playing "down" at their first opportunity.

Players will change their approach to fulfilling their goal in maximizing their expected payoffs depending on their ability to replicate each other's practical reasoning. In cases of bounded rationality, the players will have a basic conception of each other's rationality and standing in the game, but not to a level sufficient to reason backward towards a specific equilibrium. They must therefore employ a more general procedure in attaining their goal to maximize their expected payoffs. They will begin by eschewing the current positive payoff for the future opportunities of greater payoff, since their insufficient level of interactive knowledge will not establish concrete beliefs about how the other will behave imminently. Thus, future opportunities of greater payoff will appear readily available, which would not be the case if they could apply backward induction. The dearth of concrete beliefs about each other's imminent actions, combined with their desire to maximize their expected payoffs, jointly entail that Adam and Betty will *explore* and *test* each other's immediate actions. However, players will continue this course of action until their limited interactive knowledge catches up with the remaining number of stages in the finite sequential game. When the players' level of interactive knowledge – the number of iterations of 'knows' – matches the remaining number of stages in their game, they will be able to apply backward induction at their current stage. The players will then have concrete beliefs about the others' available strategies. The players will then apply backward induction and continue accordingly. Players therefore switch their approach, since they will take the current payoff instead of a greater future payoff, since that future opportunity will not occur. The complete course of action requires less cognitive ability, or less approximation of ideal rationality, than traditionally applying backward induction, since the latter will require a greater level of interactive knowledge to infer that only the initial payoff is attainable. I am assuming for now that players have roughly the same cognitive abilities. My account, however, does naturally extend to cover asymmetries in abilities (I explain below).

I shall now show in the next section how this account applies to our Centipede Game. The account will explain why more realistic players do not play their Nash strategies, and will explain the possible rationality of said behavior.

5. BOUNDED RATIONALITY IN THE CENTIPEDE GAME

The Centipede Game can have any number of stages, but let us analyze one version with five stages as depicted by figure 2.

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

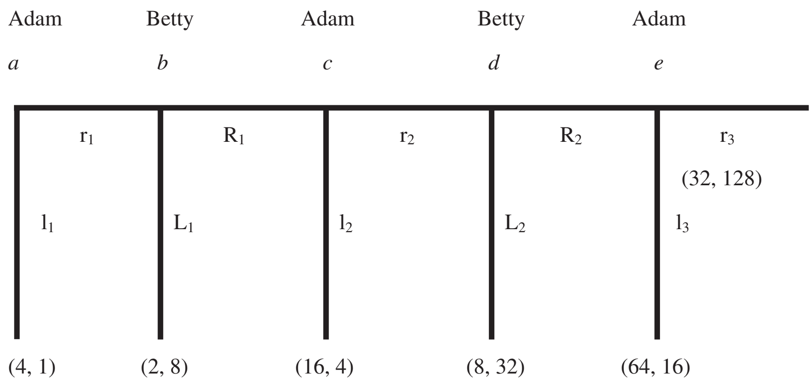

Figure 2. The Five-Stage Centipede Game. The numbers in the parentheses represent monetary payoffs. Adam's is always the left number while Betty's is the right.

Suppose Adam and Betty are boundedly rational insofar as they are both rational and both know the structure of the game, and Adam knows that Betty knows these things, and Betty knows that Adam knows these things, but no more. Thus, Adam and Betty are expected payoff maximizers, and in terms of having interactive knowledge of each other's rationality,

Adam knows:	Betty knows:
R_B	R_A
$K_B R_A$	$K_A R_B$

Carefully notice that this level of interactive knowledge is insufficient to apply backward induction for the traditional solution ($l_1 l_2 l_3$, $L_1 L_2$). What *is* sufficient, combined with their knowledge of the game depicted by figure 2, is the following:

(4) Adam is rational at stage *e*.
(5) Betty is rational at stage *d* and knows (4).
(6) Adam is rational at stage *c* and knows (5).
(7) Betty is rational at stage *b* and knows (6).
(8) Adam is rational at stage *a* and knows (7).

Adam and Betty's bounded rationality, in terms of their limited interactive knowledge of each other's rationality, guarantees the truth of propositions (4), (5), and (6), but not (7) and (8). Hence, backward induction is not applicable.

There are other epistemic states that are true about Adam and Betty given their bounded rationality, specifically states about their *lack* of interactive knowledge about each other's rationality. Adam does *not* know that $K_B K_A K_B R_A$, since (8) is false; Betty does *not* know that $K_A K_B R_A$, since (7) is false. Yet Adam and Betty

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

both aspire to maximize their expected payoffs. This is no longer a standard, but a goal of rationality.

It follows that Adam can see that he can either have four dollars in his present stage *a,* or at least double it in future stages. This is also true for Betty. Adam, to explore and test Betty's behavior at stage *b*, believing there is a chance of at least doubling his money, plays r_1 at stage *a*.[4] Betty, when given the opportunity to explore and test Adam's behavior at node *c*, believing there is a chance of at least doubling her money, plays R_1 at stage *b*.[5] Both players follow their desires to maximize their expected payoffs and Adam arrives at stage *c*. However, the game changes pace at this stage. Recall that Adam knows that $K_B R_A$ at stage *e*. Betty knows that R_A at stage *e*. The players' level of interactive knowledge matches the remaining number of stages of this Centipede Game. They can therefore reason backward from stage *e* to their current position – in this case, stage *c* for Adam and stage *d* for Betty. The following is sufficient, combined with their knowledge of the game depicted by figure 2, to apply backward induction at their current position:

(4) Adam is rational at stage *e*.
(5) Betty is rational at stage *d* and knows (4).
(6) Adam is rational at stage *c* and knows (5).

Just follow the backward induction argument in section 2, and the game ends with Adam playing l_2 at stage *c*. Adam's strategy is then ($r_1 l_2 l_3$); Betty's strategy is ($R_1 L_2$).

I have thus far explained the realistic behavior of boundedly rational players in this Centipede Game. This explanation quickly generalizes to any Centipede Game with a finite number of stages. Players with limited interactive knowledge will follow their desire to maximize their expected payoffs until their limited interactive knowledge matches the number of remaining stages. They will then apply backward induction, and the game ends at their current stage.

However, one might point out that players are actually entertaining considerations beyond monetary payoffs. Players are therefore not acting on their desire to maximize their expected payoffs. Indeed, objectors will argue here that players might follow social norms, such as cooperation, when acting in the Centipede Game. Nonetheless, I do not consider social norms influential on the behavior we find in the Centipede Game.

Bicchieri (2006) has the best current account of social norms. A social norm is an informal, possibly unwritten behavioral rule that is enforced, or acted upon, when particular epistemic conditions are satisfied. It is necessary in our Centipede Game that Adam knows that the norm exists and applies in his situation; that Adam has the conditional preference to conform to the norm in his situation; and that Adam believes that Betty expects him to conform to the norm, that Betty prefers him to conform to the norm, and that Betty may sanction his behavior (with punishment, perhaps) if possible. When these conditions are satisfied, Adam's preference relation will change from one that is based on traditional rationality to one that is based on the social norm. For instance, if the norm promotes

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

conditional cooperation, then Adam will prefer to cooperate with Betty unless she does not choose to cooperate. The presence of a social norm changes the players' preference relations, and transforms the original game of strategy into a coordination game with multiple equilibria, one of which is Pareto superior to the others. The presence of the social norm helps players identify and realize the Pareto superior equilibrium.

Bicchieri's account does not apply to sequential games, only strategic games. But we can still gather a good understanding of how the presence of a social norm will affect behavior in the Centipede Game. Consider the following Quasi-Centipede Game in strategic form:[6]

	D_B	A_B
D_A	2, 2	5, 1
A_A	1, 5	4, 4

Both Adam and Betty can either choose to play down or across. If Adam plays down while Betty plays across, then Adam receives four dollars more than Betty. The same is true for Betty when the choices are reversed. If both choose to play across, then both reach the second stage, and we add two dollars to Adam and Betty's pots, which they might divide. Removing the payoffs and replacing them with the ordinal relation of 'B' for best, 'S' for second best, 'T' for third best, and 'W' for worst, we have the following strategic form:

	D_B	A_B
D_A	T, T	B, W
A_A	W, B	S, S

Notice that this game of strategy has one Nash equilibrium (D_A, D_B), which is for both Adam and Betty to play down immediately. This nicely corresponds to the Nash equilibrium of the sequential Centipede Game when backward induction is applied. Adam's ordinal preference relation is: DA > AA > DD > AD. Nonetheless, suppose the epistemic conditions for a social norm to be present are satisfied. Adam would have the following ordinal preference relation based on his conditional preference to cooperate with Betty and play across: AA > DD > DA > AD. This transforms the Quasi-Centipede Game into the following coordination game:

	D_B	A_B
D_A	S, S	T, W
A_A	W, T	B, B

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

which appears as the following when we insert the payoffs:

	D_B	A_B
D_A	4, 4	2, 1
A_A	1, 2	5, 5

Notice that our new coordination game has two Nash equilibria, one (A_A, A_B) Pareto superior to the other (D_A, D_B). Bicchieri argues that the presence of a social norm, in addition to transforming the original game, will coordinate players on a particular equilibrium. It is the Pareto superior equilibrium (A_A, A_B) in this case.

We can repeat the process with a second stage Quasi-Centipede Game:

	D_B	A_B
D_A	4, 4	7, 3
A_A	3, 7	6, 6

which will transform into the following if the same social norm is present:

	D_B	A_B
D_A	6, 6	4, 3
A_A	3, 4	7, 7

Notice again that our new coordination game has two Nash equilibria, one (A_A, A_B) Pareto superior to the other (D_A, D_B). The players will coordinate on (A_A, A_B) in this case. And we can repeat this process for any number of stages. For any number of stages, each game will transform into a coordination game with Nash equilibria (A_A, A_B) and (D_A, D_B), and the players will coordinate on the former equilibrium, which is Pareto superior.

Considering the difference between a repeated Quasi-Centipede Game and a sequential Centipede Game, Bicchieri's theory of social norms predicts that players will either play down immediately at the first stage, or play across to the very last stage in a sequential Centipede Game. Indeed, players will do the latter in the presence of a social norm. However, this prediction contradicts Result 1, which states that subjects are likelier to revert to their Nash strategy as the game progresses. And her prediction also contradicts Result 2, which states that subjects are likelier to revert to their Nash strategy in games with fewer stages. Bicchieri's theory contradicts our experimental results as discussed in section 3. Thus, I propose that a social norm such as conditional cooperation does not account for subjects' behavior in the Centipede Game, at least given Bicchieri's understanding of social norms. A theory of bounded rationality better accounts for the experimental results.

My account, nonetheless, has yet to explain the rationality of this behavior. I argue that it is rational, and will proceed to justify its rationality by providing a solution. A solution is a profile, or set, of strategies with one strategy per player – no more, no less. The solution of choice for ideal games is the Nash equilibrium. It is a profile where each player's strategy is a best response to the others' strategies. A strategy is a best response just in case it maximizes the player's payoff given the profile. A player will not unilaterally switch strategies when the player's strategy is a best response.

Furthermore, a solution can be defined objectively or subjectively. The former approach uses objective facts about strategies and related payoffs. An objective solution is a profile of strategies where each player's strategy maximizes his payoff given the other players' strategies. The latter approach considers players' beliefs and desires, and a subjective solution will adjust to those beliefs and desires (Weirich 2010, 82–3). We can apply these two approaches to the Nash equilibrium. An objective Nash equilibrium is a profile of strategies where each player's strategy is a best response to the others' strategies, in terms of maximizing the game's payoffs independently of players' beliefs. A subjective Nash equilibrium is a profile of strategies where each player's strategy maximizes expected payoffs given the profile (84). The subjective Nash equilibrium is applicable to boundedly rational behavior in the Centipede Game.

Return to our example depicted by figure 2. The strategy profile Adam and Betty realize is ($r_1l_2l_3$, R_1L_2). This profile is a subjective Nash equilibrium and it is the solution to our specific Centipede Game. First, recall that Adam and Betty have limited interactive knowledge of each other's rationality when they begin – limited to the degree that they cannot apply backward induction to determine the other's future behavior. Second, consider Adam and Betty's desires in terms of maximizing payoffs. Both Adam and Betty aspire to maximize their expected payoffs. Their lack of concrete knowledge about each other's imminent behavior, with their desire to maximize their expected payoffs, causes them to naturally reason in terms of weighing the likelihood of obtaining an immediate positive payoff against the likelihood of obtaining a greater future payoff. Indeed, both Adam and Betty can at least double their payoff if they continue to play. These considerations entail that Adam and Betty will *explore* and *test* each other's immediate behavior. Adam will form the (weak, but positive) belief at the start of the game that Betty will likely play R_1 at her first opportunity. Betty has the (weak, but positive) belief at her first opportunity that Adam will likely play r_2 at his second opportunity. They act accordingly, and when they reach the stage where both can apply backward induction, their beliefs about each other's behavior change – Adam knows that Betty will play L_2 at her next opportunity, and Betty knows that Adam will play l_3 at his next opportunity.

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

These facts help determine whether Adam and Betty improve their standing if either one unilaterally switches one's strategy. Does Adam improve his payoff if he switches from l_2 at stage *c* to r_2? The answer is 'no'. He knows Betty will play L_2 at stage *d*, so he will switch from his current payoff of sixteen dollars to a future payoff of eight dollars. Does Adam improve his payoff if he switches from r_1 at stage *a* to l_1? The answer, again, is 'no'. He believes Betty will play R_1 at stage *b*, so he will switch from a payoff of sixteen dollars to a payoff of four dollars. Does Betty improve her payoff if she switches from L_2 at stage *d* to R_2? No, she does not, because she knows that Adam will play l_3 at stage *e*. She will switch from her payoff of thirty-two dollars to a future payoff of sixteen dollars. Does Betty improve her payoff if she switches from R_1 at stage *b* to L_1? Again, she does not, because she believes Adam will play r_2 at stage *c*. She will switch from a payoff of thirty-two dollars to a payoff of eight dollars. It follows that both Adam and Betty have no incentive to unilaterally switch strategies, so they realize a subjective Nash equilibrium. If we were considering an objective Nash equilibrium, the answer to these questions is 'yes'. Hence, they do not realize an objective Nash equilibrium.

My argument can generalize to any Centipede Game with a finite number of stages. All my argument requires is that the players do indeed explore and test each other's behavior until they reach a stage where they can apply backward induction. Applying backward induction at that stage, wherever it is located in the game, depends solely on the players' level of interactive knowledge. These conditions jointly ensure a realization of a subjective Nash equilibrium.

My account extends to *asymmetric* abilities to replicate each other's practical reasoning. This entails that Adam and Betty have asymmetric interactive knowledge. Suppose Adam and Betty are boundedly rational insofar as they maximize their expected payoffs, and they have asymmetric interactive knowledge where

Adam knows:	Adam does *not* know:	Betty knows:
R_B	$K_B K_A K_B R_A$	R_A
$K_B R_A$		$K_A R_B$
		$K_A K_B R_A$

Notice that Adam does *not* know that $K_B K_A K_B R_A$; Betty knows that $K_A K_B R_A$. Combine this with their knowledge of the game depicted by figure 2, and the following is true:

(4) Adam is rational at stage *e*.
(5) Betty is rational at stage *d* and knows (4).
(6) Adam is rational at stage *c* and knows (5).

(7) Betty is rational at stage *b* and knows (6).
(8) Adam is rational at stage *a* and does *not* know (7).

Adam can see that he can either have four dollars in his present stage *a*, or at least double that in future stages. Adam, however, does not know that Betty knows that Adam will play l_2 at stage *c*. He therefore plays r_1 at stage *a*. Betty, nevertheless, applies backward induction knowing that Adam will play l_2 at stage *c*. She thus plays L_1 at stage *b*. Adam's strategy is then $(r_1 l_2 l_3)$; Betty's strategy is $(L_1 L_2)$. The reader can verify that this outcome is a subjective Nash equilibrium.

Before concluding this section, consider a more interesting case of asymmetric interactive knowledge. Suppose Adam and Betty have asymmetric interactive knowledge of each other's rationality where

Adam knows:	Betty knows:	Betty does *not* know:
R_B	R_A	$K_A K_B R_A$
$K_B R_A$		$K_A R_B$
$\sim K_B K_A K_B R_A$		

Notice that Adam knows that $\sim K_B K_A K_B R_A$; Betty does *not* know that $K_A K_B R_A$. Combine this with their knowledge of the game depicted by figure 2, and the following is true:

(4) Adam is rational at stage *e*.
(5) Betty is rational at stage *d* and knows (4).
(6) Adam is rational at stage *c* and knows (5).
(7) Betty is rational at stage *b* and does *not* know (6).
(8) Adam is rational at stage *a* and knows (7).

This case illustrates Adam's chance to exploit Betty's lack of knowledge about him. Indeed, any case with asymmetric interactive knowledge where one knows something about one's opponent, and not vice versa, allows one to exploit one's opponent for one's gain. Adam can see that he can either have four dollars in his present stage *a*, or at least double that in future stages. Adam, however, knows that Betty does not know that Adam will play l_2 at stage *c*. He therefore plays r_1 at stage *a*. Indeed, Adam can falsely signal to Betty that he will likely play r_2 at stage *c* by playing r_1 at stage *a*. This action can influence Betty to believe (if she is initially uncertain) that Adam will play r_2 at stage *c*. Or if she already weakly believes this, his action can strengthen her credence. Betty, believing there is a chance of at least doubling her money, plays R_1 at stage *b*. Adam arrives at stage *c*. However, the game changes pace at this stage. Adam knows that $K_B R_A$ at stage *e*. Betty knows that R_A at stage *e*. The players' level of interactive knowledge becomes symmetric and matches the remaining number of stages of this game. They can both apply backward induction at their current position, and the game ends with Adam playing

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

l_2 at stage *c*. Adam's strategy is then $(r_1 l_2 l_3)$; Betty's strategy is $(R_1 L_2)$. The reader can also verify that this outcome is a subjective Nash equilibrium.

I shall now compare my account of bounded rationality to the best current model of bounded rationality, called level-*k* theory. This model's strength is that it is directly connected to the experimental research. I will show, nevertheless, that it has limitations and my account complements it, providing a more complete theory of bounded rationality.

6. LEVEL-*K* THEORY

Level-*k* theory currently provides the most successful modeling of bounded rationality. Its success stems from being directly connected to the behavioral research (Nagel 1995; Stahl and Wilson 1995; Stahl 1996; Ho, Camerer, and Weigelt 1998; Costa-Gomes, Crawford, and Broseta 2001; Costa-Gomes and Crawford 2006; Crawford and Iriberri 2007a, 2007b).

Despite variations in application, the general model operates in the following way. Players choose from a set of iteratively defined rules. The rules prescribe an action at every stage of a game of strategy. The chosen rule follows a best-response hierarchy such that the level-*k* rule (for $k = 1, 2, 3, \ldots$) best-responds to the level-(k-1) rule, the level-(k-1) rule best-responds to the level-(k-2) rule, and so forth. Players maximize their subjective expected payoffs by choosing their rules based on their beliefs of which rules the others have chosen.

One general problem with level-*k* theory is that it is unclear how bounded rationality, as a limitation on one replicating another's practical reasoning, actually affects the players' choices of rules or strategies in a game of strategy, such as the Centipede Game. It is unclear because level-*k* theory does not itself refer to anything conceptually important about bounded rationality when modeling experimental behavior. For instance, level-*k* theory posits that Adam and Betty will believe to some degree what rule the other will choose, and then best-respond to that subjective belief. Yet this feature makes no distinction between, say, ideally rational players with false beliefs and boundedly rational players being limited in replicating each other's practical reasoning. Adam and Betty might choose rules that prescribe initially rejecting their Nash strategies in the Centipede Game, and either interpretation of level-*k* theory will predict this about Adam and Betty. Of course, many proponents of level-*k* theory appear to interchange the model and its specific features with the concept of bounded rationality. However, it is only an assumption that a limitation on replicating another's practical reasoning maps onto the essential features of level-*k* theory. This assumption exists *outside* of the model.

Some may characterize my objection as being unfair, but it illustrates the limitations of mathematical models. Level-*k* theory's explanatory power is limited. It will explain the experimental behavior, such as discussed in section 3, by appealing to the rules chosen by the players. A certain distribution of the chosen rules will explain why subjects are likelier to apply backward induction as the game proceeds, or explain why players apply backward induction earlier in simpler Centipede Games than in more complex ones. If the majority of subjects choose rules ranging from L-0 to L-3, then Result 1 (subjects begin to apply backward induction as they approach the game's end) and Result 2 (players deviate from backward induction less often in the four-stage version than in the six-stage version) will follow almost immediately from level-*k* theory.

But this is where the theory's explanation ends. It does not explain how the players' bounded rationality affects their strategy choices, nor does it explain the rationality of their behavior. A complete theory of bounded rationality provides principles of bounded rationality. Principles of bounded rationality will either express procedures that players apply to choose acts, or standards of bounded rationality that apply to players' acts. And a standard of bounded rationality will apply to procedures of bounded rationality. My account complements level-*k* theory by providing a simpler procedure for choosing acts and a standard of bounded rationality. My account extends level-*k* theory to provide a more complete theory of bounded rationality for the Centipede Game.

7. CONCLUSION

My account of bounded rationality explains the two results mentioned in section 2. The first result states that subjects revert to their Nash strategy more often as the game proceeds. Subjects begin to apply backward induction as they approach the game's end. Subjects can therefore apply backward induction to a limited number of stages. My account can accommodate with ease the degree of backward induction found in the Centipede Game. Subjects in experimental settings generally exhibit zero to three levels of interactive knowledge, regardless of the complexity of the game (Ho, Camerer, and Weigelt 1998). My account explains how boundedly rational players who exhibit a fixed degree of interactive knowledge (e.g., one level of interactive knowledge) can apply backward induction to a limited number of stages (e.g., from the last stage to the penultimate stage of the game).

This fact also accounts for the second result, which states that subjects revert to their Nash strategy more often in games with fewer stages. Therefore, players deviate from backward induction less often in the four-stage version compared to the six-stage version. The simpler the game, the easier the subjects find it to apply backward induction. My account explains why rational players who exhibit a fixed degree of interactive knowledge (e.g., one level of interactive knowledge) will apply backward induction sooner in a simpler Centipede Game (the second stage of a three-stage game) than in a more complex version (the penultimate stage of a five-stage game).

Furthermore, regardless of the type of Centipede Game one investigates, there exists a solution in the form of a subjective Nash equilibrium that the players can possibly realize. This solution, as a profile of strategies, depends on the players' degree of bounded rationality and their approach to payoff maximization given the limitation of their interactive knowledge of each other's rationality. The closer the players' bounded rationality approximates ideal rationality, the closer the subjective Nash equilibrium approximates the traditional backward induction solution, which is an objective Nash equilibrium – subjective and objective Nash equilibria coincide in ideal games with ideal players.

My account's greatest benefit, furthermore, is that it provides a general normative account of realistic behavior. It provides standards of rationality for realistic behavior. Where playing one's Nash strategy is a standard for ideal games, the players' *goal* to maximize their expected payoffs remains in nonideal games. The *standard* of playing one's Nash strategy, nonetheless, must conform to the players' cognitive abilities. In realistic cases, the standard is relaxed to playing one's subjective Nash strategy, which depends on one's beliefs and desires. We find here that my account begins to unify a traditional, normative account with the results found in behavioral game theory.

REFERENCES

Bicchieri, C. 1992. "Knowledge-Dependent Games: Backward Induction." In C. Bicchieri and M. L. D. Chiara (eds.), *Knowledge, Belief, and Strategic Interaction*, pp. 327–43. Cambridge: Cambridge University Press.

Bicchieri, C. 1993. *Rationality and Coordination.* Cambridge: Cambridge University Press.

Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms.* New York: Cambridge University Press.

Bicchieri, C. and G. A. Antonelli.1995. "*Game-Theoretic Axioms for Local Rationality and Bounded Knowledge*." *Journal of Logic, Language, and Information* 4: 145–67.

Camerer, C. F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton: Princeton University Press.

Costa-Gomes, M., V. Crawford, and B. Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica* 69: 1193–235.

Costa-Gomes, M. and V. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review* 96(5): 1737–68.

Crawford, V. and N. Iriberri. 2007a. "Fatal Attraction: Salience, Naïveté, and Sophistication in Experimental 'Hide-and-Seek' Games." *American Economic Review* 97(5): 1731–50.

Crawford, V. and N. Iriberri. 2007b. "Level-*k* Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica* 75: 1721–70.

Gigerenzer, G. 2001. "The Adaptive Toolbox." In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*, pp. 37–50. Cambridge, MA: The MIT Press.

Hedden, T. and J. Zhang. 2002. "What Do You Think I Think You Think? Strategic Reasoning in Matrix Games." *Cognition* 85: 1–36.

Ho, T., C. F. Camerer, and K. Weigelt. 1998. "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests.'" *The American Economic Review* 88(4): 947–69.

https://doi.org/10.3366/epi.2011.0021 Published online by Cambridge University Press

Johnson, E. J., C. Camerer, S. Sen, and T. Rymon. 2002. "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining." *Journal of Economic Theory* 104: 16–47.

Lewis, D. 1969/2002. *Convention: A Philosophical Study*. Oxford: Blackwell Publishers.

McKelvey, R. D. and T. R. Palfrey. 1992. "An Experimental Study of the Centipede Game." *Econometrica* 60: 803–36.

Meijering, B., L. Van Maanen, H. Van Rijn, and R. Verbrugge. 2010. "The Facilitative Effect of Context on Second-Order Social Reasoning." In R. Catrambone and S. Ohlsson (eds.), *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*, pp. 1423–8. Austin, TX: Cognitive Science Society.

Nagel, R. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review* 85(5): 1313–26.

Selten, R. 2001. "What is Bounded Rationality?" In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*, pp. 13–36. Cambridge, MA: The MIT Press.

Smead, R. 2008. "The Evolution of Cooperation in the Centipede Game with Finite Populations." *Philosophy of Science* 75: 157–77.

Stahl, D. O. 1996. "Boundedly Rational Rule Learning in a Guessing Game." *Games and Economic Behavior* 16: 303–30.

Stahl, D. O. and P. W. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10: 218–54.

Weirich, P. 2010. *Collective Rationality: Equilibrium in Cooperative Games*. Oxford: Oxford University Press.

NOTES

1 I would like to thank Paul Weirich for his encouragement and helpful comments on an earlier draft of this paper.

2 I take normative game theory to predict that experimental subjects will select their Nash strategies. An alternative interpretation states, however, that normative game theory treats the behavior of *ideal* players, not real people. Normative game theory therefore will not predict that experimental subjects will select their Nash strategies. Even so, neither interpretation extends normative game theory to real people. And we need a normative theory of realistic human behavior.

3 One can object, however, that particular situations of communication or co-presence will engender common knowledge of features of a game in players with limited cognitive ability. For instance, if players can communicate with each other, or if players can see and hear each other when information about the game is dictated in a public setting, they can then deduce common knowledge of important features of the game. Communication is precluded in the Centipede Game, and administrators conduct their experiments by first publicly reading the game's instructions and allowing subjects to publically ask general questions. Some administrators will also publically announce the results of a concluded round, which subjects could include in their practical reasoning in subsequent rounds of experimentation. These practices are instances of co-presence, and there is no evidence of the subjects having common knowledge of any important feature of the game.

4 One can calculate the threshold for Adam to believe Betty will play R_1 at stage *b*. Adam has the choice of playing l_1 or r_1 at stage *a*. If he plays l_1, the game ends and he receives a payoff of four dollars. His playing r_1 depends on how likely he believes Betty is to play L_1 or R_1 at stage *b*. Adam's expected utility for playing r_1 is then $\$16p + \$2(1\text{-}p)$. To solve for the probability *p*, we let Adam's expected utility equal four dollars: $\$4 = \$16p + \$2(1\text{-}p)$. The probability *p* is one-seventh. Thus, Adam will do better by playing r_1 at stage *a* if he believes Betty will play R_1 at stage *b* with probability greater than one-seventh. Of course, I am assuming two things here. One, Adam is not considering the later stages in the game. Two, his present act at stage *a* does not influence Betty's act at stage *b*.

5 The reasoning is similar to Adam's in note 4.

6 See (Smead 2008) for an evolutionary game theoretic explanation of partially cooperative behavior in the Quasi-Centipede Game.

Ashton T. Sperry-Taylor is a visiting assistant professor of philosophy at the University of Missouri. His principle areas of research include the philosophy of science (broadly construed) and decision and game theory.