Measuring the Depth of Iteration in Humans

Bayer, R-C. [1], **L. Renou** [2]

[1] School of Economics, The University of Adelaide, Australia
[2] School of Economics, The University of Leicester, United Kingdom
Email: ralph.bayer@adelaide.edu.au

Keywords: *Experimental Economics, Bounded Rationality, Iteration Ability*

EXTENDED ABSTRACT

Economic analysis of situations with strategic interaction is based on equilibrium concepts such as Nash Equilbrium and its refinements like Subgame Perfection for dynamic games, Bayesian Nash Equilibrium for games of incomplete information and Sequential Equilibrium for dynamic games of incomplete information. In reality (and in field and laboratory experiments) we regularly observe individuals deviating from equilibrium play. It is important to understand why individuals fail to play according to standard equilibrium predictions. A better understanding of the reasons for deviations can help to develop new theories that better predict human behaviour. This is highly desirable if economics is viewed as a positive social science. A better understanding of human behaviour is also helpful for normative economics, as the desirability of policy measures depends on the predicted reactions of individuals.

Behavioural economists have identified a variety of reasons why individuals fail to behave according to traditional equilibrium concepts. The main factors have been identified to be social preferences, biases in belief formation, deviations from expected utility maximisation, present biased preferences and bounded rationality. This paper is part of a larger research project, which aims to experimentally separate and quantify the contribution of bounded rationality to the occurrence of equilibrium deviation. The rationality concept underlying game theoretical equilibria is based on individuals iterating towards the equilibrium.

For an individual to choose a strategy which is part of a Nash equilibrium this individual needs to be able to perform a certain number (depending on the nature of the game) iterations and also needs to know that the other players are able to perform the necessary number of steps. In addition an individual player needs to know that the other players know about his ability to iterate, and so on. So rationality (here the ability to perform the necessary number of iteration steps) has to be common knowledge.

A first step in understanding the impact of limited iteration ability on deviations from equilibrium play is to properly measure the number of steps individuals can commonly perform. Earlier studies used interactive games (such as beauty contests, centipede games, the email game, or the dirty-faces game) to make inferences about the iteration ability of individuals. These approaches all have a severe disadvantage in common. As these games are interactive, observed behaviour cannot properly be attributed to a certain iteration depth, as the behaviour is not only dependent on the own iteration ability but also on the beliefs about the iteration ability of the other players. In other words: observing behaviour that deviates from equilibrium play can be the cause of either limited iteration ability, or the belief that the other players' iteration ability is limited.

We propose a novel design to separate the effect of limited iteration ability from that of strategic considerations stemming from the beliefs about the rationality of the others. We use a variant of the dirty-faces game (Littlewood, 1953) and implement it in the laboratory. The main innovation is the introduction of computer players, which always behave rationally. By informing the participants that the other players are rational computer players, which are not making any mistakes, we ensure common knowledge of rationality. Consequently, the failure to play according to the equilibrium prediction must be the result of limited iteration ability. We vary the parameters in the game within subjects, such that the number of iteration steps necessary to arrive at the equilibrium prediction varies from one to four. This enables us to find the number of iteration steps an individual is able to perform.

Previous studies, where deviation from equilibrium is a confound of limited iteration ability and the lack of common knowledge of rationality, have found that the frequency of correct play drops drastically between one and two steps of iteration. We find that removing the doubts about the rationality of the other increases the average iteration depth between two and three steps.

1. INTRODUCTION

It has long been argued that real economic agents are not the flawless iterators we assume to inhabit the world described by economic theory. While the hyper-rationality assumption provides a useful benchmark and parsimoniously explains some actual behaviour observed in the real world, experimental economics has uncovered many cases where individuals systematically deviate from the predictions obtained from standard assumptions.

One possible explanation for people's failure to play Nash is that they are not logically omniscient. Logical omniscience requires that an individual can deduct all logical consequences from her knowledge. We certainly cannot expect any human to possess logical omniscience. Let us consider a concrete example of the logical omniscience problem. Consider a mathematician, as an agent, who fully accepts the axioms of set theory. Now if set theory decides the Goldbach conjecture (every even number is the product of two primes), then the mere fact that the mathematician is assumed to be logically omniscient implies that she knows that Golbach's conjecture is actually not a conjecture but a theorem! Basically, the assumption of logical omniscience in traditional game theory implies that anyone who is told the axioms of set theory should be able to prove or disprove Goldbach's conjecture.

Another nice illustration of the burden that logical omniscience puts on the rationality of players is the game of chess. We know from Zermelo's theorem that chess has a unique subgame perfect outcome (one of the players has a winning strategy or can at least force a draw). Logical omniscience implies that the player with the winning strategy (or the strategy that ensures a draw) should be able to realise this and play according to it. If chess players were logically omniscient the game of chess would be boring, as we knew the strategies that will be played and therefore the result beforehand. But how can we expect a human to be able to perform a number of iterations not even the most powerful computers can perform. This is the sort of extreme view that motivates the relaxing of the logical omniscience assumption.

Note that the failure to adhere to a Nash equilibrium strategy is not necessarily the result of an individual's inability to draw all correct inferences from their knowledge. Logical omniscience does not require that humans are selfish. As noted by many authors, some off-equilibrium behaviour can be explained by individuals having social preferences, where not only their own payoff but also that of others influences individual's utility. Models of social preferences are centered around inequality aversion (e.g. Fehr and Schmidt, 1999, or Bolton and Ockenfels. 2000) and may additionally include tastes for social efficiency and for kindness (e.g. Charness and Rabin, 2002). Further deviations from the preferences economist usually assume that can be responsible for unexpected play are deviations from expected utility theory (see Starmer, 2000 for an overview) or from orthodox time preferences with exponential discounting (see Frederick *et al*., 2002 for a review of the literature).

A third potential reason for the deviation from equilibrium play lies in how people form beliefs. By beliefs, we mean probabilistic beliefs - i.e., player i assigns probability $p_i(\omega_{-i})$ to the realization of the state ω_{-i}. In matrix games the beliefs refer to the probability assigned by a player to the other players choosing particular actions. In games with uncertainty beliefs also assign probabilities to possible states of the world, while in dynamic games updating of beliefs comes into play.

Given that there are so many potential reasons why people choose actions in games that are not part of a Nash equilibrium (or one of its refinements and extensions), the task of explaining to which extent the potential reasons contribute to deviations in different settings becomes vitally important for the formulation of new and more accurate theories. This paper is part of a larger research program with this task in mind. As a first step we aim to measure the level of logical omniscience (or equivalently the number of iteration steps) in humans. Existing studies have tried to infer the iteration depth humans have by implementing dominance-solvable games in the laboratory. Dominance-solvable games are games where a finite number of iteration steps lead to the equilibrium strategy. These earlier studies have in common that the number of iteration steps performed by an individual is not readily observable. An observed deviation from equilibrium play can either be due to the inability to iterate or due to the belief that the other player(s) are not able to iterate deeply enough.

To see this, consider the centipede game depicted in Figure 1, which is taken from McKelvey and Palfrey (1992). Centipede games are regularly used to make inferences about the iterative ability of individuals. The subgame perfect Nash equilibrium outcome is player one playing T (take) at the first node. The equilibrium actions at the other information sets are such that each player takes at any information set. A player can determine this outcome by iteration from the back (backward induction). The reasoning goes like this: I know that player 2 will play T at the last node then it is better for player 1 taking at the penultimate node, which means that player 2 should play T at the node before that, and so on.

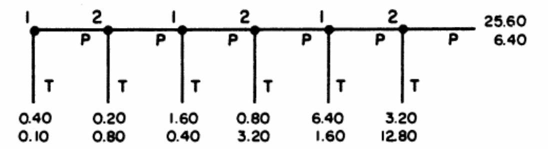

Figure 1: A centipede game

Now one might use the node at which a player chooses T as a measure for the number of steps of iteration a player might be able to do. Playing T at the penultimate is equivalent to one step of iteration, playing T at the node before that to two steps, and so on.

This reasoning is compelling but not strictly valid. Suppose a researcher observes player one playing P (pass) at the first node. Given the logic above we would conclude that this individual is not able to perform five steps of iteration. Unfortunately, this does not have to be true. We should observe the same behaviour if player one is able to perform five steps of iteration but beliefs that it is very likely that player two is not able to perform four steps of iteration. If player two is not able to perform four steps of iteration he might not play T at the next node. If player one considers this to be sufficiently likely it is best for her – even if she can do five steps of iteration – to play P at the first node and then to take at the third node if player two passes at the second node, which would give her a payoff of 1.60 compared to 0.4 if she takes immediately, what is prescribed by subgame perfection.

Other studies using games with strategic interdependence suffer from the same problem: Deviation from Nash equilibrium can be a confound of limited iterative ability and strategic uncertainty stemming from a lack of common knowledge of rationality. These two effects cannot be cleanly separated. Examples are the dirty- faces game (Weber, 2001 and Bayer and Chan, 2007), centipede games (McKelvey and Palfrey, 1992 and Fey *et al.* 1996), beauty contests (first introduced by Nagel, 1995) and matrix games (Stahl and Wilson, 1994 and Costa-Gomez *et al.*, 2001).

In what follows we propose a novel experimental design that overcomes the problem of non-separable effects. We use a modified version of the dirty-faces game with computerised players and implement it in the laboratory. A human player is paired with a number of logically omniscient computer players (i.e. the computer players are programmed such that they have an unlimited depth of iteration). The human is informed about playing with computers and about the computer's ability to reason flawlessly. This procedure does not only overcome the separation problem described above but also rules out any social influence of social preferences, as the computers do not receive any payoff.

We vary the parameter setting in the games within subjects such that one to four steps of iteration are necessary for the human to play the game according to the equilibrium prediction. With this design we are able to measure how many steps of iteration an individual can actually perform. We find that humans are able to perform more steps of iteration than conjectured using games with where uncertainty about the common knowledge prevails. This indicates that some deviations from equilibrium play that were commonly attributed to limited cognitive abilities are actually the consequence of the individual's doubt about the cognitive ability of other players.

The remainder of this paper is organizes as follows. Section 2 describes the logic of the game used to experimentally measure the depth of iteration in humans. Section 3 presents the experimental design, while our main results are presented in Section 4. Section 5 concludes.

2. THE RED-HAT PUZZLE

The level of logical omniscience is positively related to the number of logical connectives a person can master. For this reason the measurement of the level of logical omniscience requires a variety of tasks, where the successful completion of a particular task can be related to the number of connectives. The tasks must be such that they have different difficulties - from easy (only few connectives are required) to very difficult (many connectives are required). Furthermore, the tasks should be similar; such that potential other factors (like framing) do not influence behaviour differently across tasks.

For this reason we use a game, where a variation of the starting information requires a different number of logical iterations in order to find the correct answer. We use the so called "Red Hat Puzzle" (RHP), which is sometimes also referred to as the "Dirty-Faces Game". This game has been used to investigate iteration depth in humans by Weber (2000) and Bayer and Chan (2007). We follow the exposition of Fagin *et al.* (1995) in order to explain the puzzle.

Consider N agents are playing together. Each of these agents has either a red hat or a white hat. The agent cannot see the colour of their own hat but sees the colour of the hat the other players are wearing. Suppose that some of the agents, say n, have red hats. Along comes a referee, who says that “at least, one of you has a red hat on his head”. Then he asks the following question: “Does any of you know whether you have a red hat on your head?” As long as some people answer “No, I don’t know” the referee keeps repeating his question. Which kind of reasoning does this trigger in the mind of logically omniscient agents?

We can prove that the first $n-1$ times the referee asks the question, logically omniscient agents answer ”No”, but then the n^{th} time, agents with red hats will all answer “Yes”. The proof is by induction on n. For $n=1$, the one with a red hat sees that no one else has a red hat. Since he knows (it is common knowledge) that there is at least one red hat, he concludes that he must be the one. Here just one step of reasoning is required.

Now suppose $n=2$. So there are just two players with a red hat, say players 1 and 2. Each of them answers “No” the first time they are asked because they both see the red hat of the other player, which prevents them from knowing their own hat colour. But the moment 2 says “No”, 1 realises that he must have a red hat herself. The mere fact that 2 could not establish his hat colour before means that he must have seen a red hat, which 1 now knows to be on her head. Thus 1 answers “Yes, it is red” the second time the referee asks. The same reasoning applies to 2, who will also answer “Yes” when asked the second time. For these two players two steps of iterated reasoning are necessary to determine their hat colour.

The remaining players if logically omniscient will be able to determine that their hats are white after observing the first two players announcing that their hat colour is red in round two. They can recover the reasoning of players 1 and 2 and in a third step of iteration and conclude that players one and two must have seen a white hat on their heads.

Now suppose that $n=3$ – there are three players, 1, 2, and 3 – with a red hat. Player 1 thinks as follows. Assume I do have a white hat. Then by the case of $n=2$, both 2 and 3 should answer “Yes, it's red” the second time. When they do not, 1 realises that his assumption was false and concludes that he has a red hat. He will answer “Yes, it's red” to the third question. The same logic applies to players 2 and 3. They all need three steps of iteration to determine their hat colour. The remaining players can then conclude in round four that their hats are white, as otherwise the three others could not have found out that their hat colour was red. For this they need to be able to perform four steps of iteration. The general argument for $n>3$ proceeds along the same lines.

What is the role of the referee's public announcement? One might wonder why this announcement is informative. After all, each player initially knows that at least one hat is red when $n>1$. Thus, one might conjecture that this announcement is useless, but this is false!!! This announcement provides common knowledge and starts the chain of reasoning as it determines that someone can say I have a red hat if $n=1$. Observe that here common knowledge is required. The mutual knowledge in the cases $n>1$ is not enough.

We implement this game in the laboratory and pair a human player with three computers. Given the discussion above, it is very important to ensure that subjects have common knowledge with respect to the fact that there is at least one red hat if one wants to measure the level of logical omniscience. We ensure this by stressing this point in the experimental instructions and asking control questions to test if the subjects understand this important fact.

There is a second kind of common knowledge necessary in order to cleanly measure the level of logical omniscience. Suppose a logically omniscient player does not know with certainty that all the other players are also logically omniscient. Suppose we have $n=2$ and all players have answered “No” in the first round of questions. Then a player who sees one other red hat can conclude that he must have a red hat *if and only if* he knows that the player with the red hat he sees has answered correctly in the first round. If he has doubts about the cognitive ability of the other player he cannot draw this conclusion with certainty. So then this player will answer with “No”. Consequently, a design that does not ensure that a player is absolutely certain that the other players are logically omniscient does not measure logical omniscience but a combination of logical omniscience and the beliefs of a player about the level of logical omniscience of the others. Then these two effects cannot be cleanly separated. Previous studies using this game (Weber, 2001 and Bayer and Chan, 2007) suffer from this problem.

We therefore introduced logically omniscient computer players. Providing subjects with the information what they see (the hat colours of the other three players, the computers) one of the answers "I can infer that I have a red hat", "I can't possibly know" "I can infer that I have a white hat" is logically correct in the first round. Informing the subjects that a computer always chooses the logically correct answers then ensures that the subjects understand that the computer is logically omniscient. The same is true for the second (and subsequent) rounds of questions. Then, once the answers of the four players (one human and three computers) are made common knowledge, there is again one logically correct answer. The subjects were informed that the computers choose the logically correct answer at any stage.

Note that a mistake of the human player in one round leads to inconsistencies with the computer players. Then a computer player might infer the wrong hat colour. In order to prevent subjects from observing such a situation we end a puzzle immediately once the human player has made a mistake. We also end the puzzle once all players (computers and humans) have correctly inferred their hat colour.

3. EXPERIMENTAL PROCEDURE

In a four player version of the Red-Hat Puzzle (*RHP* thereafter) there are seven logically different situations. The situations differ by the number of red hats a subject sees and whether she has a red or white hat. So there are two situations where the subject sees three red hats (one where she has a white hat herself and one where she has a red hat). There are also two situations each where the subject sees one or two red hats. In the case where the subject does not see any red hats there is only one situation, as the announcement "There is at least one red hat" requires the subject to have a red hat.

The difficulty is determined by the number of red hats an individual sees. If we denote the number of red hats seen as r then a subject needs $r+1$ steps of iteration to correctly determine the hat colour. In our treatment the subjects played all seven RHPs. The puzzles where ordered by difficulty. So individuals played the easiest puzzle ($r=0$) first and the two hardest puzzles ($r=3$) last. In between puzzles subjects were given no feedback. So they were not told if they had solved the previous RHP correctly. This prevents the subjects from learning in between puzzles.

The puzzles were programmed in z-tree (Fischbacher, 2007). Before the actual puzzles were started subjects had to answer some control questions in order to ensure that they understood the instructions, the screen layout and that "there is at least one red hat". Instructions and treatments can be obtained from the authors upon request.

In order to provide strong incentives payment was organised as follows. Subjects started with a show-up fee of AUD 17.50. Then for each mistake AUD 2.50 were deducted. Assuming that subjects are loss-averse this setup provides very strong incentives. It took the subjects about an hour to play the seven situations. We conducted five sessions with a total of 94 subjects.

The subjects were mostly students at Adelaide University. Their background varied widely. The degrees these students were enrolled in covered almost the whole spectrum. The years of university education also varied greatly among subjects (from first-year students to PhD students).

4. RESULTS

We now present our main results. We estimate a panel probit model with random effects, where the correct solution of a RHP is the dependent variable. The average marginal effects of the independent variables are summarized in Table 1 below.

We observe a very high correlation within a subject (*rho=0.45*), which shows that the different individuals were differently successful with the RHPs and that the success across situations within a subject is consistent. Another result that validates our results is that all subjects got the situation with one step right. So the way we induced common knowledge about the fact there is at least one red hat must have worked.

We further observe that between two and three steps of iteration occurs a break. While a puzzle with two steps was 33 percentage point less likely to be solved than the puzzle with one step, both three and four step puzzles both were 57 percentage points less likely to be solved correctly than the one-step puzzle. So people that are able to solve the puzzle with three steps are also able to solve the puzzle with four steps. We conjecture that people that manage to do three steps of iteration are able to use induction and therefore have understood the pattern of correct answers.

We controlled for the courses the students were enrolled in. The only difference to the reference group (Economics students) were students from the Medical School. They were on average much more likely to solve a puzzle, which is not a surprise, as the entry requirements are among the highest there.

.**Table 1:** Marginal effects of a panel probit model

Independent Variable	Marginal Effect
Difficulty dummies (one step is the reference	
2 Steps	-0.33**
	(0.00)
3 Steps	-0.57**
	(0.00)
4 Steps	-0.57**
	(0.00)
Course dummies (Economics is the reference group)	
Medicine	0.48**
	(0.00)
All others not significant	
Socioeconomic Background	
Higher Maths	0.18**
	(0.04)
Male	0.18**
	(0.01)
Age dummies	Not sign.
Control questions and decision times	
Control questions correct	0.22**
	(0.01)
Time taken for first decision (in 10 sec)	0.00
	(0.38)
Time taken for critical decision (in 10 sec)	0.06**
	(0.03)
P-Values in parentheses, ** sign. 5%-level	

Also not surprising is the finding that students who studied advanced maths in high school were more successful at solving the puzzles. More puzzling is the finding that males were better at solving the puzzles than females.

Subjects who were able to answer the control questions correctly on average had a much higher chance of solving a puzzle than those who failed answering them correctly. A very interesting result has to do with decision times. We see that the decision time at the first question "What's the colour of you hat?" does not predict wether a subject solves a puzzle correctly or not. However, the decision time at the question which is critical, i.e. when subjects have to switch from "I don't know" to "red" or "white", has an influence. Subjects who thought harder and took more time at this crucial decision were more likely to answer correctly. One additional minute of thinking increased the probability of being correct by 36 percentage points.

We also created a measure for logical omniscience. The measure provides the number of steps an individual is able to perform. We use the number of correct solutions in the seven puzzles and then assign the number of steps an individual was able to perform. For example a person who got all seven RHPs right will have an LO score of 4 as this individual solved all puzzles including the ones where four steps of iteration were required. An individual with three correct answers will be assigned an LO score of 2, as this person was able to solve the three easiest puzzles up to an iteration level of two.

There are two things worth noting here. This LO score assumes that people's behaviour conforms to a perfect Guttmann scale, which means that the performance is weakly decreasing in the difficulty. However, perfect Guttmann scales are hardly ever observed in the real world. Our scale is not perfect either. However, more than 80 percent of our subjects showed monotonous answering patters, which is quite a good level of consistency.

Additionally, for iteration levels two, three and four there were two situations each. The logic in the two puzzles with the same iteration level is slightly different. In one case subjects had to infer from the computers not determining their hat colours yet that their hat colour must be red, while in the other problem they had to infer that their hat colour must be white, as the computers have found out that their hat colour is red. For this reason our measure allows for half steps. Figure 2 shows the distribution of the LO scores.

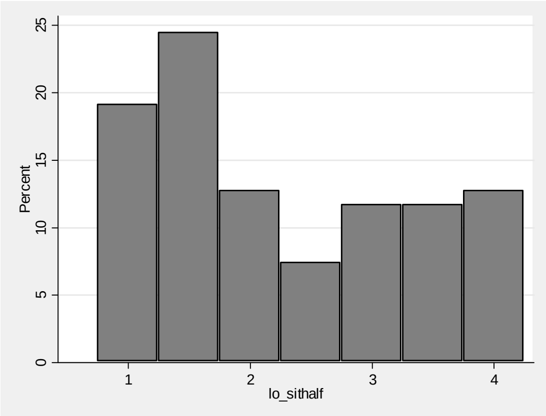

Figure 2: Distribution of LO scores

The median LO score is 2 (mean 2.27). In studies, where doubts about the rationality of the other players were also a factor, subjects usually behave on average as if their iteration ability was between one and two. So around half of an iteration step was lost in these studies due to subjects not relying on the rationality of the other players. The distribution of LO scores shows that surprisingly many subjects were actually able to iterate with depth three or four.

5. CONCLUSION

In this paper we present a novel experimental design, which for the first time allows the measurement of iteration depth in humans without further assumptions about the beliefs individuals have about the rationality of other players. We find that the actual iteration depth is about half a step higher than measured by studies which did not control for the beliefs about the rationality of other players. We also find that individuals who are able to perform three steps of iteration are quite likely to go all the way and are able to perform four steps as well. One, two, three, epiphany.

6. REFERENCES

Bayer, R.-C. and M. Chan (2007), The Dirty Faces Game Revisited, *School of Economics Working Paper*, The University of Adelaide.

Bolton, G. and A. Ockenfels (2000), ERC: A Theory of Equity, Reciprocity, and Competition, *American Economic Review* 90, 166-193.

Charness, G and M. Rabin (2002), Understanding Social Preferences with Simple Tests, *Quarterly Journal of Economics* 117, 817-869.

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

Fagin, R., J. Y. Halpern, Y Moses and M. Y. Vardi (1995), *Reasoning about Knowledge*, MIT Press, Cambridge, USA.

Fehr, E. and K. Schmidt (1999), A Theory of Fairness, Competition, and Cooperation, *Quarterly Journal of Economics* 114, 817-868.

Fey, M, R. D. McKelvey and T. D. Palfrey (1996), An Experimental Study of Constant-Sum Centipede Games, *International Journal of Game Theory* 25, 269-287.

Fischbacher, U.(2007), z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics* 10(2), 171-178.

Frederick, S. G. Loewenstein and T. O'Donoghue (2002), Discounting and Time Preference: A Critical Review, *Journal of Economic Literature* 40(2), 351-401.

Littlewood, J. E. (1953). *A Mathematician's Miscellany*. London: Meuthen & Co. Ltd.

McKelvey, R. D. and T. D. Palfrey (1992), An Experimental Study of the Centipede Game, *Econometrica* 60, 803-836.

Nagel, R. (1995), Unravelling in Guessing Games: An Experimental Study, *American Economic Review* 85, 1313-1326.

Stahl, D. O. and P. W. Wilson (1994), Experimental evidence on players' models of other players, *Journal of Economic Behavior & Organization* 25, 309-327.

Starmer, C. (2000) Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk, *Journal of Economic Literature* 38(2), 332-382.

Weber, R. A. (2001), Behaviour and Learning in the Dirty Faces Game, *Experimental Economics* 4,229-242.