



**University of
Leicester**

DEPARTMENT OF ECONOMICS

**Cognitive abilities and behavior in
strategic-form games.**

Ralph-C. Bayer, University of Adelaide, AUS

Ludovic Renou, University of Leicester, UK

**Working Paper No. 11/16
January 2011**

Cognitive abilities and behavior in strategic-form games.*

Ralph-C. Bayer[†] & Ludovic Renou[‡]

January 18, 2011

Abstract

This paper investigates the relation between cognitive abilities and behavior in strategic-form games with the help of a novel experiment. The design allows us first to measure the cognitive abilities of subjects without confound and then to evaluate their impact on behavior in strategic-form games. We find that subjects with better cognitive abilities show more sophisticated behavior and make better use of information on cognitive abilities and preferences of opponents. Although we do not find evidence for Nash behavior, observed behavior is remarkably sophisticated, as almost 80% of subjects behave near optimal and outperform Nash behavior with respect to expected pay-offs.

Keywords: cognitive ability, behaviors, strategic-form games, experiments, preferences, sophistication.

JEL Classification Numbers: C70, C91.

*We thank the Economic Design Network Australia and the University of Adelaide for funding.

[†]School of Economics, University of Adelaide, Nexus 10, Adelaide 5005, Australia. Phone: +61 (0)8 8303 4666, Fax: +61 (0)8 8223 1460. ralph.bayer@adelaide.edu.au

[‡]Department of Economics, Astley Clarke Building, University of Leicester, University Road, Leicester LE1 7RH, United Kingdom. l.renou@le.ac.uk

1 Introduction

Most strategic situations require individuals to make complex chains of reasoning, especially counter-factual reasoning of the type: “If I were my opponent, I would not play action A because action B gives a strictly higher payoff regardless of the action I actually choose,” or “If my opponent conjectures that I play each of my actions with equal probability, then he would play A and, thus, I should play B because it is my best-reply to A .” Quite naturally, we can imagine arbitrarily complex chains of reasoning of this sort. The cognitive ability of a person, i.e., his ability to make complex chains of reasoning, is therefore likely to influence his behavior in strategic situations. This paper addresses this issue. More precisely, the purpose of this paper is to measure the ability of individuals to perform complex chains of reasoning and to relate the measured cognitive abilities to behavior in strategic-form games.

A main motivation for this work is to better understand reasons for humans deviating from theoretically postulated behavior, such as iterated deletion of strictly dominated strategies or Nash equilibrium. If, for example, the choice of an individual does not coincide with iterated deletion of strictly dominated strategies and yet her cognitive abilities suggest that she is able to draw the necessary logical inferences, then we might conclude that there is a lack of common beliefs in rationality. After all, a common belief in rationality implies that choices are consistent with iterated deletion of strictly dominated strategies. Ultimately, such a finding might call for solution concepts relaxing the assumption of common (or even mutual) beliefs in rationality.

With this objective in mind, we design a novel series of experiments, which allow us to address the issues raised above. Our experiments have two essential components. The first component tests the ability of individuals to perform complex chains of reasoning, in particular counter-factual reasoning. We employ a computerized variant of the red-hat puzzle, as introduced in Bayer and Renou (2007). This computerized variant of the red-hat puzzle

is ideal to measure the cognitive abilities of individuals as it neatly controls for confounding factors, such as other-regarding preferences or beliefs about the cognitive abilities of others. The second component of the experiments presents the subjects with twelve strategic-form games. Ten of these are dominance solvable, for which the logic for iterated deletion of strictly dominated strategies closely parallels the logic underlying the red-hat puzzles. We add two dictator games, which are designed as a control for preferences.

We use the choice made by the subjects in the twelve strategic-form games to identify different behavioral types. Following recent studies, such as Stahl and Wilson (1994; 1995), Costa-Gomes et al. (2001), Costa-Gomes and Weizsäcker (2008), we estimate a mixture density model, where each subject's behavior is determined, possibly with errors, by one of a rich set of possible *behavioral* types. Although we consider a rich set of fifteen possible behavioral types, only six of them are of significant importance in the sample. We now describe these six types. An *Altruistic* type identifies the strategy profile that maximizes the sum of his own and the opponent's payoff and chooses the corresponding action. An *Optimistic* type identifies the strategy profile that maximizes his payoff and takes the corresponding action (i.e., "max max" behavior). An *L0* player uniformly randomizes over his actions. The more sophisticated *L1* type (*Naive* in Stahl and Wilson; 1994 and 1995) best replies to *L0*. An *L2* type is even more sophisticated and best replies to *L1*. Lastly, a *D1* player does one round of deletion of strictly dominated strategies and best replies to a uniform distribution over the opponent's remaining strategies.

We now summarize the main results of the paper. Firstly, from the red-hat puzzles, we derive an index of cognitive ability for each subject. The index ranges from zero to four and can be interpreted as a subject's depth of reasoning. In particular, a subject with a cognitive ability of two or more has the ability to put himself in the position of the opponent (strategic thinking), while a subject with a cognitive ability of one or less cannot. In the sample of 154 subjects, the percentage of subjects showing an iteration depth of

zero, one, two, three and four is 0%, 52.6%, 24.03%, 5.84% and 17.53%, respectively. About one-half of the subjects can think strategically.

Secondly, the econometric analysis of behavioral types gives a relatively simple description of subjects' behaviors. As already alluded to, the most informative statistical model identifies six behavioral types, *Altruistic*, *Optimistic*, *L0*, *L1*, *L2* and *D1* with percentages of 6.50%, 1.95%, 11.69%, 62.99%, 14.94% and 1.95%, respectively. Most subjects (about 77%) are best classified as *L1* or *L2* types. *L1* is the modal type. This is in line with a result from Costa-Gomes and Weizsäcker (2008), who found that *L1* was the best predictor for play in strategic-form games. We find no statistical evidence for the presence of *Sophisticated* (rational expectation) or *Equilibrium* types. Some previous studies, e.g., Stahl and Wilson (1995), Haruvy et al. (1999) and Costa-Gomes et al. (2001), found some evidence for a small fraction of equilibrium types.¹ In accordance with most of the studies mentioned above, we also find some *L2*, *D1* and *Optimistic* types. Unlike most previous studies, we find strong evidence for the existence of some *Altruistic* types.

Finally, we turn our attention to the relation between cognitive abilities and behavioral types. The general pattern is as follows: the higher the index of cognitive abilities, the more sophisticated (i.e., *L1*, *L2* or *D1*) is the behavioral type. However, the impact of cognitive abilities is not as extreme as one might have expected. Even the subjects with very good cognitive ability are not behaving as *L3* or *Nash* types, which are the most sophisticated types we looked for. Yet, the observed *L1* and *L2* types' expected payoffs (given the distribution of play observed in the sample) are very high. The expected payoff of an *L2* type is about 99% of that a player with rational expectations, the ideal of game theory, would achieve. Our most prevalent type, *L1*, still achieves 96% of the rational expectation payoff. In contrast, a

¹In their econometric analysis with information about search patterns, Costas-Gomez et al. found no evidences for equilibrium types. In sharp contrast with our study and most previous studies, Rey-Biel (2009) found that a large fraction of subjects (70%) played in agreement with Nash behavior.

Nash type would achieve less than 90% of the rational expectation payoff.² In terms of outcomes, the subject's behavior is quite sophisticated after all.

We also observe that subjects with a cognitive ability of two or more react to information on their opponent, while subjects with a cognitive ability of one do not. Moreover, subjects with a cognitive ability of two or more only behave as *Altruistic* types if they are informed about opponents' play in the dictator games. Also, subjects with a cognitive ability of two or more are more likely to behave as *L2* types if information is provided about the cognitive abilities of the opponents. Remarkably, all in all, about 80% of our subjects show quite sophisticated behavior despite of an environment that is quite unfamiliar to them (our subjects had not participated in prior experiments).

This paper contributes to the large literature on iterative reasoning in games e.g., McKelvey and Palfrey (1992), Beard and Beil (1994), Nagel (1995), Ho et al. (1998), Goeree and Holt (2004), Huyck et al. (2002), Cabrera et al. (2007), to name just a few.³ A recurring feature of many of these studies is the use of games solvable by iterated deletion of strictly or weakly dominated strategies. In these studies, the ability of individuals to perform complex chains of reasoning is associated with their ability to iteratively delete dominated strategies. Centipede games (e.g., McKelvey and Palfrey, 1992) and beauty contest games (introduced to the literature by Nagel, 1995) are two of the most commonly used games in that literature. Typically, the use of these games suffers from the lack of control for beliefs about the rationality of others and for social preferences.⁴ The main novelty of our ex-

²*L0* types achieve even less with about 80%.

³We refer the reader to chapter 5 of Camerer (2003) for a survey of earlier literature on this topic.

⁴Gneezy et al. (2010) and Dufwenberg et al. (2010) use a version of the game "Nim" to study if and how humans learn backward induction. Since players have (weakly) dominant strategies, this zero-sum game can be used to measure the depth of iterative reasoning in humans if one accepts the auxiliary hypothesis that it is common knowledge that players do not play weakly dominated strategies.

periment is that we are able to provide these controls. The second strand of literature our paper is related to is that on behavioral types and behavior in strategic-form games (see, among others, Stahl and Wilson, 1994; 1995; Costa-Gomes et al., 2001; Camerer et al., 2004; Costa-Gomes and Crawford, 2006; Rey-Biel, 2009; Arad and Rubinstein, 2010 and Burchardi and Penczynski, 2010).⁵

To summarize, the main objective of the paper is to understand how the ability to perform complex chains of reasoning relates to behavior in strategic-form games, controlling for other-regarding preferences. The experimental design consists of three distinct stages. The first stage provides a measure of the ability of subjects to perform counter-factual reasoning. We use a variant of the red-hat puzzle to obtain such a measure. The second stage controls for other-regarding preferences by presenting the subjects with dictator games. Finally, the last stage consists of twelve strategic-form games. The paper is organized around these three stages. After briefly describing our subject sample in Section 2, Section 3 presents the red-hat puzzles, its computerized experimental design and initial results. Section 4 describes the dictator games. Section 5 presents the twelve strategic-form games, the postulated behavioral types and our econometric analysis. Section 6 relates the estimated behavioral types to cognitive abilities and preferences, while Section 7 concludes.

2 Subject sample and procedure

The experiment was conducted at AdLab, the Adelaide Laboratory for Experimental Economics at the University of Adelaide in Australia. We used Urs Fischbacher's (2007) experimental software Z-tree. A total of 154 subjects participated in eight sessions; all sessions took place over three days:

⁵Beyond the types typically considered, we also included types that capture other-regarding preference (e.g., Fehr and Schmidt (1999) or Charness and Rabin (2002)) and others we deemed promising.

26, 27 and 28 of November 2008. For each of our four treatments (which we will explain later), we conducted two sessions resulting in between 36 and 42 subjects per treatment. A session lasted for about 90 minutes and on average subjects earned about 20.4 Australian Dollars. Earnings ranged from AUD 8 to AUD 36. The 154 participants were mostly students from the University of Adelaide and the University of South Australia. All subjects had no previous experience with participation in economic experiments. The experimental instructions used are available on the author's webpage.

3 Measuring cognitive abilities

This section presents the first stage of our experiment, the red-hat puzzles, and describes our measure of the ability of subjects to perform complex chains of reasoning (for short, cognitive ability). Along with a measure of "social preferences," this measure of cognitive ability is at the heart of our analysis of behavior in strategic-form games.

3.1 Experimental design

We first present a simple puzzle, which is the basis for the first stage of our experiment. Each of N individuals has either a red hat or a white hat, observes the hat color of others, but cannot observe the color of his own hat. Along comes a trusted referee, who declares that "at least one individual has a red hat on the head." The referee then asks the following question: "What is your hat color?" All individuals simultaneously choose an answer out of "I can't possibly know," "I have a red hat," or "I have a white hat." Players then learn the answers of the other players and are asked again what their hat color is. This process is repeated until all individuals have inferred their hat color. This problem is known as the red-hat puzzle (henceforth, RHP).⁶

⁶The same problem is also known as the "Dirty Faces Game." For an alternative exposition see Fudenberg and Tirole (1991, pp. 544-548) or Osbourne and Rubinstein (1994, p. 71).

Suppose that an individual, say Bob, sees that no one else has a red hat. Since it is commonly known that there is at least one red hat, he must conclude that he has a red hat. Now suppose that Bob sees one red hat, say on Ann's head. He should answer "I cannot possibly know" the first time the question is asked, as he sees another red hat. However, if Bob can iterate his reasoning further, he can infer his hat color from Ann's answer. If Ann says "I cannot possibly know," Bob must realize that he has a red hat. For otherwise, Ann would have known that she has a red hat right away and should have answered "I have a red hat." In general, the greater the number of red hats an individual observes, the more complex is the reasoning needed to logically infer the color of one's own hat. An individual needs $m + 1$ iteration steps to figure out his hat color, where m denotes the number of red hats this individual observes.

It is important to note that the above logic for an individual correctly inferring his hat color relies on some crucial assumptions. Firstly, even if an individual has unlimited cognitive abilities, he also needs to know that the answers of the other individuals are logically correct. To see this, suppose that there is a unique red hat and Bob observes this red hat. Bob can only correctly infer his hat color (white) if the individual, who wears the red hat, answers the first question accurately with "I have a red hat." It follows that any experimental design using the red-hat puzzle to measure the cognitive ability of humans has to make sure that each subject knows that the answers of other subjects are logically correct. Secondly, the event "There is at least one red hat" must be common knowledge. Thirdly, subjects must have incentives to correctly infer their hat color and to truthfully report their logical inferences.

We now describe our experimental protocol, and how it addresses the difficulties discussed above. In our experiment, a human subject was paired with three computers, which were acting as "players." Pairing a subject with computers has several advantages given our objective. Firstly, we can reasonably assume that subjects have no concerns for the eventual "pay-

offs” of computers. Secondly, we can ensure that a subject knows that the computers’ answers are logically correct by a) programming the computer-players to choose the logically correct answers and b) communicating this credibly to the subject. Accordingly, computers were programmed to choose the logically correct answers in each round of questions (see below), and the instructions emphasized this point heavily. Additionally, subjects were told (and constantly reminded with an on-screen message) that there was at least one red hat.

Subjects were asked to infer their hat color from the information given to them. At any point when they were asked, they had three possible answers to choose from: “I have a WHITE hat with certainty,” “I have a RED hat with certainty,” and “I cannot possibly know.” The first time a subject had to choose an answer within a puzzle the information a subject had was the hat color of the three computer-players (along with the fact that there was at least one red hat). In any subsequent round within the puzzle, the information a subject had was the complete history of all answers of all players (the computers’ and his) in all previous rounds. Similarly, the initial information a computer-player had was the hat color of the two other computers and the human subject and, subsequently, the complete history of answers. Whenever a computer had to answer, the computer’s answer was the (unique) logically correct answer inferred from their information and history. Before subjects started the experiment, they had to answer some control questions testing their understanding of the instructions.

A RHP was stopped after either a wrong answer by the human or a correct announcement of the hat color.⁷ This stopping procedure is necessary to avoid logical inconsistencies. Suppose there is only one red hat, which is worn by the human subject. The subject initially observes three white hats. Now, if the subject (wrongly) answers “I cannot possibly know,” then computers should logically infer and, if allowed, answer “I have a red hat.” However, this

⁷In a given round, announcing a hat color was correct only if it was actually possible to infer the hat color at this given round.

contradicts what the subject observes. Although the computers in this case would chose the logically correct answers, we would have lost control over how a subject interprets this inconsistency. We believe that the observation of contradicting computer announcements and physical reality would have led subjects to believe that the computers were not properly programmed or that our claim that the computers are logically correct was based on deception.

Since each individual was paired with three computers, we had seven possible distinct logical situations. A logical situation was determined by the number of red hats a subject saw and whether the subject had a red or white hat herself. The more red hats a subject was observing, the more steps (iterations) were required to correctly infer the hat color.

Subjects played all seven situations in increasing order of difficulty and without any feedback in between. For any mistake, we deducted 2.5 Australian Dollars (AUD) from the subjects start-up amount of AUD 20.00. As any mistake immediately terminates a RHP, subjects payout from the first stage of the experiment could be calculated as AUD 2.5 for the show-up fee plus AUD 2.5 per puzzle correctly solved. We used the frame of “deduction per mistake” in order to make the incentive to think hard at every step of the game as strong as possible (given our financial means).⁸

3.2 A measure of cognitive ability

This section briefly examines the determinants to correctly infer one’s hat color in red-hat puzzles. All our results are consistent with a prior experiment we ran and thus refer the reader to our companion paper, Bayer and Renou (2009), for an in-depth analysis. Table 1 reports the percentage of subjects over the entire sample who correctly solved a puzzle (as a function of the number of steps needed to solve it). Some remarks are worth mak-

⁸For our subjects losing AUD 2.5 with a single wrong decision is quite a strong incentive. A student job pays about ten Dollars per hour (the median hourly wage in South Australia is about AUD 20). We also conjecture that the loss frame increases effort through loss-aversion.

ing. Firstly, as expected, the more iterations are required to correctly infer ones hat color, the lower is the percentage of correct answers. Secondly, and somewhat surprisingly, there is almost no difference between solving a puzzle requiring three iterations and one requiring four iterations (26.6% vs. 26.3%), while there is a significant difference between solving a puzzle requiring two iterations and one requiring three or four. This suggests that if an individual can perform three steps of iterative reasoning, then she can also do four steps. Our econometric analysis confirms this empirical observation. Thirdly, and reassuringly, all subjects were able to solve the easiest puzzle, i.e., to state that they had a red hat when they were observing three white hats.

steps	1	2	3	4
solved in %	100.00	55.5	26.6	26.3

Table 1: Correctly solved puzzles by iteration steps required

We now consider the determinants of correctly inferring ones hat color. The dependent variable is “correct,” a dichotomous variable indicating whether a subject had correctly inferred his hat color in a given puzzle. Our econometric model is a probit model with subject specific random effects. Table 2 reports the marginal effects averaged over the whole sample.

Our regression confirms our initial observation: the more steps required to solve a puzzle, the less likely it is that an individual solves it ($p < 0.01$, Wald test). The probability of solving a puzzle requiring three steps is 0.358 lower than the probability of solving puzzles requiring two steps (the reference). Moreover, as already suggested, there is no statistical difference between solving a puzzle requiring three steps and a puzzle requiring four steps. We also note that the probability of correctly solving puzzles of different difficulties is highly correlated within subjects ($\rho = 0.62$). The high correlation within subjects testifies of a large heterogeneity in cognitive abilities across subjects, which cannot be explained by demographics. This also

Probability of solving puzzle (<i>estimated baseline</i>)		
Dependent variables	Marginal effect	SE
Iteration steps necessary (baseline: <i>2 steps</i>)		
<i>3 steps</i>	-0.358**	(0.048)
<i>4 steps</i>	-0.362**	(0.048)
High school maths	-0.016	(0.109)
Control questions ok	0.204**	(0.072)
Age dummies	not sig.	
Male	0.166	(0.202)
Session dummies	not sig.	
Course dummies (baseline: <i>Arts</i>)		
<i>Commerce and Finance</i>	0.103	(0.204)
<i>Economics</i>	-0.177	(0.117)
<i>Engineering</i>	0.219	(0.208)
<i>Law</i>	-0.047	(0.192)
<i>Medicine</i>	0.060	(0.235)
<i>Science</i>	0.479*	(0.229)
<i>Other</i>	0.080	(0.218)
<i>Not studying</i>	-0.118	(0.201)
<i>N</i>	924	
Log Likelihood	-435.646	
χ^2_{22} , Wald test	105.12**	
ρ	0.615**	(0.052)

** p<0.01, * p<0.05

Table 2: Random-effect panel estimation of the probability of a correct solution

indicates that there is some consistency within subjects: correctly solving a certain puzzle increases the likelihood of solving another puzzle.

We now construct a simple measure of cognitive ability so as to quantify the number of steps of iterative reasoning an individual is able to perform. Each puzzle can be parameterized by a pair (m, n) , where m is the number of red hats a subject observes and n is the actual number of red hats ($n = m + 1$ or $n = m$). Remember that if a subject observes m red hats, $m + 1$ iterations

are required to correctly infer his hat color: the higher m is, the more complex is the chain of reasoning. In our experiment, we had seven possible situations: $(0, 1)$, $(1, 1)$, $(1, 2)$, $(2, 2)$, $(2, 3)$, $(3, 3)$ and $(3, 4)$.

A perfect measure would obtain if a subject solving a puzzle requiring m iterations had actually solved all puzzles requiring $m' \leq m$ iterations. In this idealized situation, the measure of cognitive ability of a subject would be given by the number of iterations required to solve the most difficult puzzle the subject can solve. For instance, if a subject had correctly solved the puzzles $(0, 1)$, $(1, 1)$, $(1, 2)$ and failed the puzzles $(2, 2)$, $(2, 3)$, $(3, 3)$ and $(3, 4)$, his measure of cognitive ability would be 2. Such a measure is called a perfect Guttman scale (Guttman 1944; 1950) in the psychological literature.

It is, however, unreasonable to expect all subjects to exhibit a consistent pattern of answers.⁹ For instance, a subject might fail to solve a puzzle of a given complexity for reasons (e.g., trembles, mistakes, inattention) unrelated to his ability to perform chains of reasoning. Alternatively, a subject might correctly solve a puzzle by sheer luck. Quite surprisingly, 78.6% of the subjects nonetheless exhibit a consistent pattern in our sample. Still, we have to deal with the inconsistent patterns of answers.

The measure we adopt has the following features. First, it is an integer ranging from 0 (all puzzles incorrectly solved) to 4 (all puzzles correctly solved). Moreover, it makes a unique prediction about the puzzles a subject can solve. For instance, if the measure of cognitive ability of an individual is 2, it predicts that the individual solves the puzzles $(0, 1)$, $(1, 1)$ and $(1, 2)$ and fails all others. Second, it maximizes the number of correct predictions. To clarify this last point, consider the following situation (“Y” stands for “correctly solved” and “N” for “failed”):

$(0, 1)$	$(1, 1)$	$(1, 2)$	$(2, 2)$	$(2, 3)$	$(3, 3)$	$(3, 4)$
Y	Y	N	Y	Y	Y	Y

⁹A pattern of answers is consistent if there exists a threshold m^* such that a subject correctly solves all puzzles requiring $m \leq m^*$ iterations and fails all other puzzles.

The subject has correctly solved all puzzles but (1, 2). If we ascribe the measure 4 to the individual, we predict that he can solve all puzzles and thus make one prediction error. If, however, we ascribe any other number, we make at least two wrong predictions since we would predict that he cannot solve puzzles (3, 3) and (3, 4). Thus, we attribute the failure to solve puzzle (1, 2) to a mistake. Alternatively, if an individual had solved puzzles (0, 1) and (3, 4) and failed all others, his measure would be one. The fact that the individual has solved the puzzle (3, 4) would be attributed to sheer luck.

Two additional remarks are worth making. First, if an individual has a consistent pattern of answers, then our measure coincides with a perfect Guttman scale. Second, two different numbers can correctly predict the same number of outcomes. In all such situations, we decided to follow the more conservative approach of assigning the *lowest* number maximizing the number of correct predictions. Nonetheless, all our results are robust to this choice. The two resulting measures are highly correlated ($\rho > 0.86$). A standard criterion for a good measure and the existence of a one-dimensional hierarchical scale is that the coefficient of reproducibility (i.e. the fraction of correctly predicted outcomes) is above 0.9, Guttman (1950). Our measure satisfies this criterion, as it predicts 91.5 percent of the outcomes correctly. Figure 1 shows the distribution of our measure across subjects.

4 Social preferences

In the second stage of the experiment, subjects were presented with a dictator game, designed to provide a crude measure of social preferences. The dictator could implement three different allocations: $A = (4.5, 4.5)$, $B = (6.0, 1.5)$ and $C = (3.0, 9.0)$, expressed in Australian Dollar and with the first coordinate representing the monetary payoff of the dictator. Subjects had to choose between A , B or C and then to confirm their choices, so as to minimize the rate of possible mistakes. We divided the subjects into two equal groups, \mathcal{G}_1 and \mathcal{G}_2 . Initially, subjects from \mathcal{G}_1 were anonymously and randomly paired with

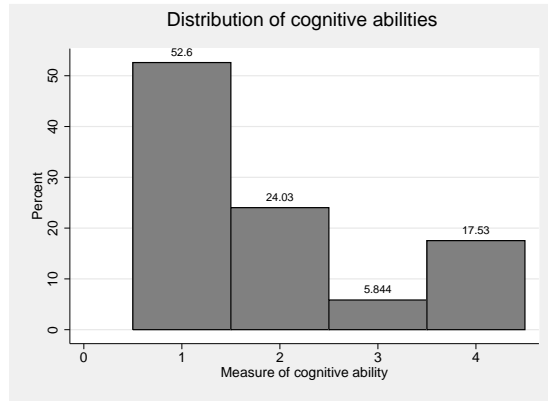


Figure 1: Empirical distribution of the measure of cognitive ability

subjects from \mathcal{G}_2 and assigned the role of dictators. Subsequently, subjects from \mathcal{G}_2 were assigned the role of dictators and anonymously paired with a member of \mathcal{G}_1 different from the initial match. No feedback was provided.

A selfish dictator would choose option B . For a person to choose A or C , other-regarding preferences are necessary. For a subject to choose A over B , the concern for the other subject must be large enough such that the increase in the payoff of the other subjects compensates for the loss in one's own payoff. For a player to prefer C over the other two options, the concern for the other subjects must be even greater. Also, note that the total surplus increases from B to A and from A to C .

Our choice of allocations was guided by the idea that a person who prefers C is keen on social efficiency and has strong social preferences. A person who chooses A tends to dislike inequity, whereas a person who chooses B is relatively selfish. For our purpose, such an admittedly crude classification of subjects is sufficient.

In our sample of 154 subjects, 27% have chosen alternative A , 50% alternative B , and 23%, alternative C .

5 Behavior in games

The last stage of our experiments consisted of twelve strategic-form games. Each player had three pure strategies in each game. All but two games were dominance solvable and each game has a unique Nash equilibrium in pure strategies.¹⁰ More precisely, in three games, one round of deletion led to the subject's equilibrium strategy (i.e., the subject had a dominant strategy). Two and three rounds of deletion were required in three games each, whereas there was one game requiring four rounds. The games are presented in the Appendix.

Six games and their respective transposed constitutes our twelve games; most of them taken from Costa-Gomes and Weizsäcker (2008). Moreover, all games were presented to the subjects as if they were row players. Thus, subjects were actually playing the same twelve games and had the same view of each game, i.e., the view of a row player. We randomized the order of the games so as to control for possible order effects. We had four different sequences of games, in which a subject never played a game and its transposed in succession.¹¹

A subject was initially seeing his own payoff, but not the payoff of his opponent. He had to wait for 15 seconds before being able to see the payoff of his opponent; he had to click on a button to do so.¹²

The monetary payoff from the second and third stage of the experiment consisted of the payoff obtained in two randomly chosen games out of the fourteen games played in phases two and three (twelve strategic-form games,

¹⁰In all dominance solvable games, a strictly dominated strategy was dominated by a *pure* strategy.

¹¹We found no order effects.

¹²We included the delay and the option to view the payoff of one's opponent to gather some additional information. For instance, a selfish player who has a dominant strategy does not have to view the payoff of his opponent. Yet, mere curiosity might induce a player to view the payoff of his opponent even if it is not needed; the delay was included to discourage curious subjects to view the payoff of their opponents. Unfortunately, this did not help us.

one dictator game as the dictator and one as the receiver). Subjects were informed about this in the instructions.

5.1 Treatments, matching and information.

To summarize, we obtain a measure of the ability to perform complex chains of reasoning with the red-hat puzzle. To control for social preferences, we use dictator games. We randomly rematched subjects in between stage two and three of the experiment and informed subjects of that procedure. No feedback was provided during stage three (i.e., the strategic-form games).

The experiment consist of four different treatments, which differ in the level of information given to subjects in stage three. In the *fullinfo* treatment, subjects were told which option the opponent had chosen in the dictator game and how many RHPs the opponent had correctly solved. The number of RHPs solved is a signal about the depth of reasoning of the opponent, whereas the choice in the dictator game is a signal about the opponent's preferences. In the *red-hat* treatment, only the number of correctly solved RHPs is provided. In the *dictator* treatment, the subjects were only told the opponent's choice in the dictator game, while in the *noinfo* treatment no information was provided.

5.2 Nash behavior

This section briefly comments on the overall compliance with the concept of Nash equilibrium. To be precise, let G be one of the twelve strategic-form games subjects were presented with. (Remember that all games had a unique Nash equilibrium in pure strategies.) We say that a subject plays according to the concept of Nash equilibrium if he plays the equilibrium strategy of G .¹³

¹³We need to be cautious here. Let G be one of the twelve strategic-form games of the experiment and $u_i : \times A_i \rightarrow \mathbb{R}$ the payoff function in game G . The “experimentally postulated” preferences for player i is \succsim_i , where $(a_i, a_{-i}) \succsim_i (a'_i, a'_{-i})$ if and only if $u_i(a_i, a_{-i}) \geq u_i(a'_i, a'_{-i})$. However, if player i 's true preferences \succsim_i^* differ from the “exper-

<i>Solvable in</i>	<i>1 round</i>	<i>2 rounds</i>	<i>3 rounds</i>	<i>4 rounds</i>	∞ <i>rounds</i>
<i>Nash behavior</i>	66.67%	36.80%	32.03%	32.47%	15.58%

Table 3: Fraction of Nash behavior as a function of rounds of deletion

Table 3 reports the percentage of subjects' behavior consistent with the concept of Nash equilibrium, as a function of number of rounds of iterated deletion of strictly dominated strategies required. The percentages are surprisingly low. In games with a strictly dominant strategy, subjects chose the Nash strategy about two-thirds of the time, but the fraction drops to only about one-third for games where two, three or four rounds of iteration are necessary. In games that are not dominance solvable, the percentage drops even further. Tests of proportions reveal that the fraction of Nash behavior is higher in games with a strictly dominant strategy than in all other games ($p < 0.01$ for all pairwise comparisons) and lower in games, which are not dominance solvable, than in all other games ($p < 0.01$ for all comparisons). Comparing the percentage of Nash behavior to one third, i.e., as if subjects play each of their actions with equal probability ($L0$ type), shows that for games with a strictly dominant strategy, the percentage is significantly higher ($p < 0.01$, binomial distribution test). Strikingly, for games that are not dominance solvable, the percentage of Nash behavior is significantly lower than one-third ($p < 0.01$). For the remaining games, the percentages of Nash behavior are not significantly different from one-third.

Overall, the low agreement with Nash behavior and the fact that the percentage of Nash behavior is significantly lower than one-third in the most difficult games suggest the following. Firstly, Nash behavior seems to be a poor predictor for the behavior of the subjects. Secondly, subjects are likely to use other heuristics than Nash behavior. In what follows, we consider a

imentally postulated" preferences \succsim_i , he might still play an equilibrium strategy of the game induced by his true preferences and yet would appear to us, the experimentalists, as if he is not playing according to the concept of Nash equilibrium.

rich set of behaviors, including Nash behavior, and estimate the likelihood for the observed play to be consistent with our behavioral types.

5.3 Behavioral types

So far, we have devoted our analysis to the solution concept of Nash equilibrium. As already alluded to, a subject's preferences can coincide with the "experimentally postulated" preferences, have unlimited cognitive abilities and yet not play a Nash equilibrium. For instance, the subject might be doubtful about the rationality of his opponent or might have specific conjectures about his opponent, so that deviations from Nash behavior do *not* imply bounded rationality or other-regarding preferences. After all, if a subject expects his opponent to play a certain action, even if it is strictly dominated, it is optimal to best-reply to it.

To address this issue, we postulate that a subject can be one of fifteen possible behavioral types, labeled *Altruistic*, *Pessimistic*, *Optimistic*, *L0*, *L1*, *L2*, *L3*, *D1*, *D2*, *Equilibrium*, *Sophisticated*, *Regret*, *Best-reply to Altruistic*, *Inequity Aversion*, and *Efficiency*. *Altruistic* takes the decision corresponding to the profile of decisions that maximizes the sum of his own and opponent's payoffs. *Pessimistic* takes the decision that maximizes his minimal payoff over his opponent's decision. *Optimistic* takes the decision that maximizes his maximal payoff over his opponent's decision. *L0* uniformly randomizes over his actions. *L1* conjectures that his opponent plays each of his actions with equal probability and best-plies to this conjecture, i.e., *L1* best replies to *L0*. *L2* best replies to *L1*, i.e., he conjectures that his opponent plays according to *L1* and best replies to this conjecture. Similarly, *L3* best replies to *L2*. *D1* does one round of deletion of strictly dominated strategies, conjectures that his opponent plays each of the remaining strategies with equal probability, and best replies to this conjectures. *D2* does two rounds of deletion of strictly dominated strategies, conjectures that his opponent plays each of the strategies remaining with equal probability, and best replies to his conjectures. *Equilibrium* conjectures that the opponent plays a Nash equilib-

rium strategy and best replies to this conjecture. *Sophisticated* best replies to the empirical distribution of played actions in each game. The ten behavioral types we have described so far are borrowed from Nagel (1995), Stahl and Wilson (1995) and Costas-Gomez et al. (2001). We have also considered three additional types that we ex-ante deemed promising. *Regret* minimizes his maximal regret over the decision of his opponent. *Best-reply to Altruistic* best replies to *Altruistic*. *Inequity Aversion* and *Efficiency* take equilibrium decisions, but with other-regarding preferences à la Fehr and Schmidt (1999) for inequity aversion and à la Charness and Rabin (2002) for efficiency.¹⁴

5.3.1 The statistical model

This section presents the statistical model adopted for the estimation of types. For each subject $i \in \{1, \dots, n\}$, the vector of observables x_i consists of the play in the twelve different games, i.e., $x_i := (x_i^1, \dots, x_i^{12})$ where $x_i^k \in \{A, B, C\}$ is the action played by player i in game k . Furthermore, as already explained, we hypothesize that each subject i has a behavioral type $\theta \in \Theta$, independently and identically drawn from the distribution $p \in \Delta(\Theta)$: each behavioral type determines a particular play in all games. For instance, in game 3, *Pessimistic* plays B , while *Altruistic* plays A . The aims of our statistical model is two-fold. Firstly, we want to estimate the distribution p from the observables. Secondly, we want to assign to each subject his most likely type (so that we can relate the type of a subject to his measure of cognitive ability and social preferences).

We assume that in each game k , subject i of type θ plays the action $x_i^k(\theta)$ consistent with his type with probability $(1 - e_\theta)$ and plays any one of the remaining two actions with probability $e_\theta/2$. Note that e_θ is type-specific, but it does neither depend on the game played nor on the identity

¹⁴For inequity aversion, we have used the payoff function $u_i(a_i, a_j) - 1/2 \max(u_j(a_i, a_j) - u_i(a_i, a_j), 0) - 1/4 \max(u_i(a_i, a_j) - u_j(a_i, a_j), 0)$ where u is the material/experimental payoff, while for efficiency, we have used $u_i(a_i, a_j) + 0.5u_j(a_i, a_j)$ if $u_j(a_i, a_j) > u_i(a_i, a_j)$ and $0.5u_i(a_i, a_j) + 0.5u_j(a_i, a_j)$ if $u_i(a_i, a_j) \geq u_j(a_i, a_j)$.

of a subject. It is also worth noting that if $e_\theta = 2/3$, a subject uniformly randomizes over his three actions. This simple observation makes it possible to identify the $L0$ type.

Denote $\bar{x}_i(\theta)$ the number of actions consistent with type θ out of subject i 's profile of actions x_i , i.e., $\bar{x}_i(\theta) := |\{k : x_i^k = x_i^k(\theta)\}|$. It follows that the probability that subject i of type θ plays the sequence x_i of actions is:

$$\Pr(x_i|e_\theta, \theta) = (1 - e_\theta)^{\bar{x}_i(\theta)} (e_\theta/2)^{12 - \bar{x}_i(\theta)}.$$

Thus, the probability that subjects play the profile $x := (x_1, \dots, x_n)$ is:

$$\Pr(x|e, p) = \prod_i \left(\sum_\theta p(\theta) (1 - e_\theta)^{\bar{x}_i(\theta)} (e_\theta/2)^{12 - \bar{x}_i(\theta)} \right),$$

where e is the profile of error rates $(e_\theta)_\theta$. The log-likelihood $l((p, e))$ of (p, e) is thus given by:

$$\sum_i \ln \left(\sum_\theta p(\theta) (1 - e_\theta)^{\bar{x}_i(\theta)} (e_\theta/2)^{12 - \bar{x}_i(\theta)} \right).$$

To maximize the log-likelihood $l((\cdot, \cdot))$ with respect to (p, e) , we implement the EM algorithm of Dempster et al. (1977) (see also Redner and Walker, 1984; Little and Rubin, 1987; El-Gamal and Grether, 1995). More precisely, the t -th iteration of the algorithm is as follows. Given the current estimate p_i^t of the distribution of player i 's type, the M(aximization) step produces the estimates

$$(e_\theta^t)_\theta \in \arg \max_{(e_\theta)_\theta} \sum_i \sum_\theta p_i^t(\theta) \ln((1 - e_\theta)^{\bar{x}_i(\theta)} (e_\theta/2)^{12 - \bar{x}_i(\theta)}).$$

After some algebra, we obtain that the estimate e_θ^t for type θ is

$$e_\theta^t := \frac{\sum_i p_i^t(\theta) [12 - \bar{x}_i(\theta)]}{12 \sum_i p_i^t(\theta)}.$$

Then, the E(xpectation) step defines the consistent estimator for the type distribution p^t as $p^t := (1/n) \sum_{i=1}^n p_i^t$ and computes p_i^{t+1} as the Bayesian posterior of p^t given the estimated error terms $(e_\theta^t)_\theta$:

$$p_i^{t+1}(\theta) = \frac{(1 - e_\theta^t)^{\bar{x}_i(\theta)} (e_\theta^t/2)^{12 - \bar{x}_i(\theta)} p^t(\theta)}{\sum_{\theta \in \Theta} (1 - e_\theta^t)^{\bar{x}_i(\theta)} (e_\theta^t/2)^{12 - \bar{x}_i(\theta)} p^t(\theta)}.$$

Lastly, to initialize the algorithm, fix p_1 in the interior of $\Delta(\Theta)$.¹⁵ The procedure is repeated until convergence. We refer the reader to Redner and Walker (1984) for more details about the EM algorithm and mixture density models (as ours).

The confidence intervals are obtained by bootstrapping, as analytical standard errors are not available.¹⁶ For model selection, we use three different criteria for fit and discriminatory power: the Akaike Information Criterion (AIC) adjusted for small samples, the Bayesian Information Criterion (BIC) and the Average Normalized Entropy (ANE). The lower AIC and BIC are, the better is the model.

Let $p_i \in \Delta(\Theta)$ be the estimated distribution over subject i 's type. Clearly, if p_i is close to the uniform distribution, the estimation is not very informative about the “true” type of subject i . Alternatively, if p_i assigns probability 1 to a particular type, the estimation is very informative. A good measure of the informativeness of p_i is its entropy. More precisely, we use the Average Normalized Entropy (proposed by El-Gamal and Grether, 1995) as a measure of the informativeness of our estimated model. It is given by:

$$ANE = -\frac{1}{N} \sum_{i=1}^N \sum_{\theta=1}^{|\Theta|} p_i(\theta) \log_{|\Theta|} (p_i(\theta)).$$

The ANE takes on values between zero and one. The ANE is zero if all subjects have an estimated probability of one for one type (a perfect

¹⁵We chose the uniform distribution as the initial prior. Yet, our results are robust to this choice.

¹⁶For the bootstrap, we re-sampled the original data with replacement, with the number of observations equal to the number of observations in the original data, and re-performed the estimation. After 1000 re-samplings, we used the resulting distributions of estimates to determine the bootstrap confidence intervals.

classification) and is one if each subject has an estimated distribution that is uniform. The higher the entropy, the less informative is the model.¹⁷

5.3.2 Econometric results

Table 4 reports our estimation results for different models. We started with the full model containing all our types and then eliminated the types with inconsistent error estimates.¹⁸ After removing the inconsistent types, we arrive at Model 1, which includes the same types as in Costa-Gomes et al. (2001). For model selection, we now proceed in two steps. First, we remove the types with estimated error rates of close to two-thirds (i.e., *Nash*, *Pessimistic*, and *L3*) and replace these types by the type *L0*.¹⁹ A test of whether this procedure is valid involves the resulting error rate for the *L0* type in the new model (Model 2, second column in table 4). An error rate of close to two-thirds is required for the new model, Model 2, to pass the test. Since the error rate for the new type is indeed close to two-thirds, we are confident that we successfully identify the *L0* types in the sample.

Second, starting from Model 2, we check if reducing the model further can improve parsimony without worsening the fit. No other reduced model improves the values of both *BIC* and *AIC*. Yet, removing *Sophisticated* and *D2* (Model 3) improves the BIC score but worsens the AIC score. If, furthermore, we consider the average normalized entropy ANE, Model 3 outperforms Model 2 in two of our three criteria for fit and discriminatory power. All other combinations of the types in Model 1 are dominated in all criteria. Accordingly, we choose Model 3 as our preferred model and use it for

¹⁷El-Gamal and Grether (1995) regard values below .38, which was the highest value they observed in their study, as good.

¹⁸Estimated error rates of well above two-thirds imply that subjects played the strategies consistent with their behavioral types with a lower probability than any of the other two alternatives.

¹⁹Concretely, we created a pseudo-type *L0* that plays *B* with the error rate e . Whenever e is equal to two-thirds, this pseudo-type is nothing else than the true type *L0*.

Type	Model 1		Model 2		Model 3	
	p	e	p	e	p	e
Nash	0.008 [0.00; 0.14]	0.767 [0.32; 0.87]				
Altruist	0.068 [0.03; 0.16]	0.212 [0.06; 0.45]	0.071 [0.03; 0.15]	0.219 [0.06; 0.45]	0.070 [0.02; 0.16]	0.214 [0.06; 0.42]
Optimist	0.018 [0.01; 0.08]	0.182 [0.09; 0.98]	0.020 [0.01; 0.10]	0.195 [0.09; 1.00]	0.020 [0.01; 0.12]	0.196 [0.09; 1.00]
Pessimist	0.132 [0.00; 0.18]	0.722 [0.45; 0.84]				
$L1$	0.582 [0.44; 0.68]	0.247 [0.20; 0.29]	0.581 [0.45; 0.68]	0.247 [0.21; 0.29]	0.599 [0.49; 0.70]	0.249 [0.22; 0.28]
$L2$	0.147 [0.07; 0.22]	0.250 [0.17; 0.32]	0.146 [0.07; 0.23]	0.249 [0.17; 0.33]	0.154 [0.08; 0.24]	0.249 [0.18; 0.33]
$L3$	0.006 [0.00; 0.13]	0.616 [0.09; 0.69]				
$D1$	0.015 [0.00; 0.06]	0.109 [0.09; 0.74]	0.015 [0.00; 0.09]	0.109 [0.09; 0.75]	0.016 [0.00; 0.07]	0.111 [0.09; 0.75]
$D2$	0.004 [0.00; 0.16]	0.154 [0.11; 0.79]	0.004 [0.00; 0.16]	0.152 [0.11; 0.77]		
Sophist.	0.019 [0.00; 0.09]	0.229 [0.16; 0.89]	0.020 [0.00; 0.11]	0.233 [0.16; 0.91]		
$L0$			0.143 [0.01; 0.18]	0.685 [0.50; 0.91]	0.142 [0.01; 0.20]	0.678 [0.51; 0.84]
$\ln(\mathcal{L})$	-1423.627		-1425.704		-1431.824	
AIC	2888.090		2883.146		2886.577	
BIC	2942.957		2926.962		2919.053	
ANE	0.201		0.203		0.200	

Table 4: EM estimation of types

assigning types to individual subjects.²⁰

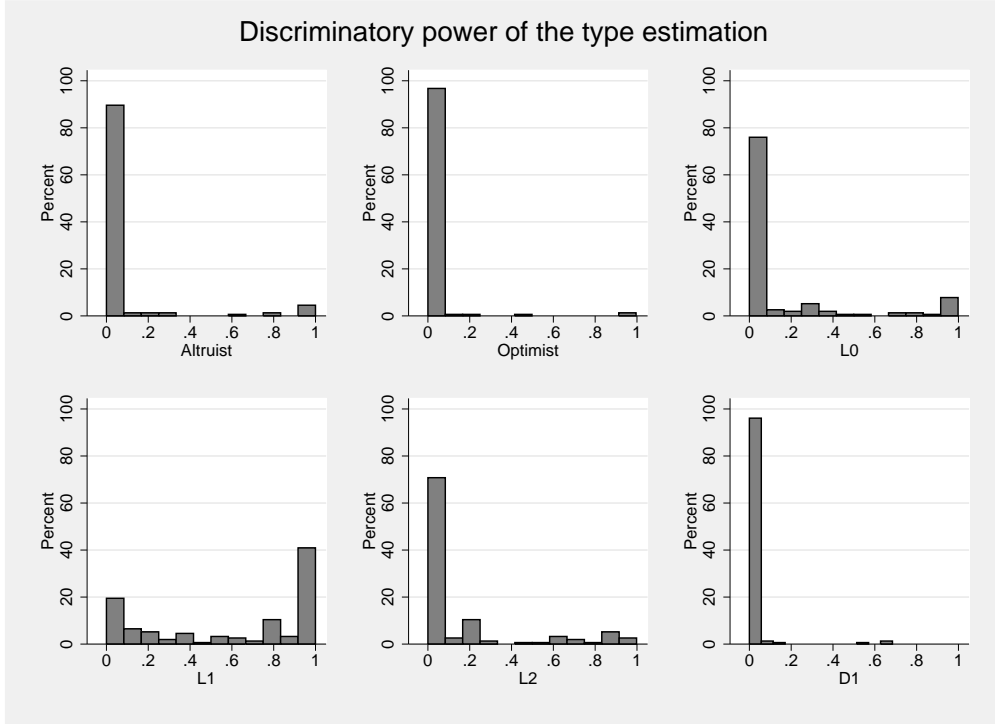


Figure 2: Posterior probabilities for subjects being of a certain type

Figure 2 demonstrates further that Model 3 gives a crisp separation of types. Figure 2 plots the empirical distribution across subjects of the estimated posterior probability $p(\theta|x_i)$ of type θ for the six types in Model 3. We observe that for most types, the empirical distribution is concentrated around 0 or 1, which confirms the excellent discriminatory power of our model. This is further corroborated by the relatively low ANE.

Lastly, we assign to each subject the behavioral type with the highest estimated posterior probability. More than 73% of subjects have a maximum

²⁰All subsequent findings are robust to this choice. In fact, only one subject's type classification changes (from *Sophisticated* to *L2*) when we switch from Model 2 to Model 3.

posterior probability of at least 0.75 and still half of the subjects have a maximum posterior of more than 0.9. Figure 3 shows the resulting type distribution.

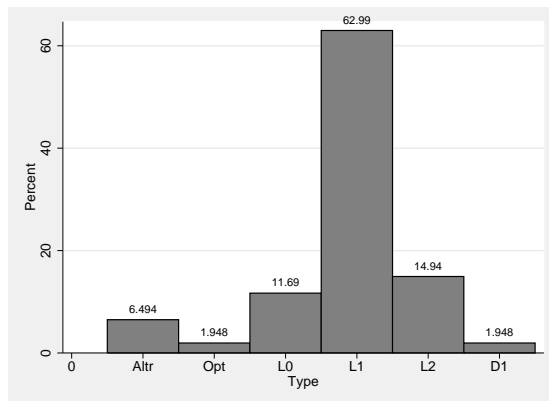


Figure 3: Distribution of behavioral types

We first note that no subject is best described as the *Nash* type. This is largely consistent with the observations made in Section 5.2. Moreover, the distribution over types in our sample is roughly consistent with existing findings. Some of the games in our experiments were first used by Costa-Gomes and Weizsäcker (2008), who concluded that their subjects typically behaved as *L1* types. Our results are consistent with this finding: the most prevalent type in our sample is the *L1* type (roughly 63% of the subjects). Additionally, in an earlier study, Costa-Gomes et al. (2001) equally found *L1* to be the most prevalent type in their econometric analysis of behavioral types with information about search patterns. The next largest group of subjects (15%) are subjects classified as type *L2*.²¹ Note that behaving as an *L2* type is near optimal in our sample, as *L1* is the most prevalent type and *L2* best replies to *L1*.²² We will return to this point in the next section.

²¹In their econometric analysis of behavioral types without information about search patterns, Costa-Gomes et al. (2001) found *L2* to be most prevalent type.

²²A *Sophisticated* type would certainly do better; but according to Model 3 we did not

Next, we need to mention that we found a large number of subjects (almost 12%) exhibiting *L0* behavior. The remaining subjects were classified as either *Altruistic* (about 6%) or *Optimistic* (1%).

6 Cognitive ability, preferences and behavior

This final section combines our previous results in order to analyze the relationships between cognitive abilities, preferences and behavior in strategic-form games.

We first analyze the relationships between cognitive ability and behavior. Prima facie, we would expect subjects with a higher measure of cognitive ability to display a more sophisticated behavior (e.g., *L2*) in games. Figure 4 plots the distribution of cognitive abilities per behavioral type. About 70% of the *L2* types have a measure of cognitive ability of two or above and 72% percent of *L0* types have a measure of one. This observation is reassuring as *L0* behavior *does not* require to put oneself in the position of the opponent, which subjects with a cognitive ability of one cannot do. On the contrary, *L2* behavior *does* require to put oneself in the position of the opponent, which subjects with a cognitive ability of two or more can do. Mann Whitney U-Tests confirm the observation that more sophisticated types have better cognitive abilities. An *L2* type tends to score higher on our cognitive ability measure than an *L1* type and an *L0* type ($p < 0.01$ vs. *L0*, $p < 0.02$ vs. *L1*; one-sided tests). Also, an *L1* type tends to have a slightly better cognitive ability than an *L0* type ($p < 0.09$).

An additional piece of evidence for the impact of cognitive abilities on behaviors in strategic-form games is the fact that all *Optimistic* subjects have a measure of cognitive ability of one.

Finding 1 *There is a positive relationship between cognitive abilities and the level of strategic sophistication in strategic-form games.*

find this type in the sample.

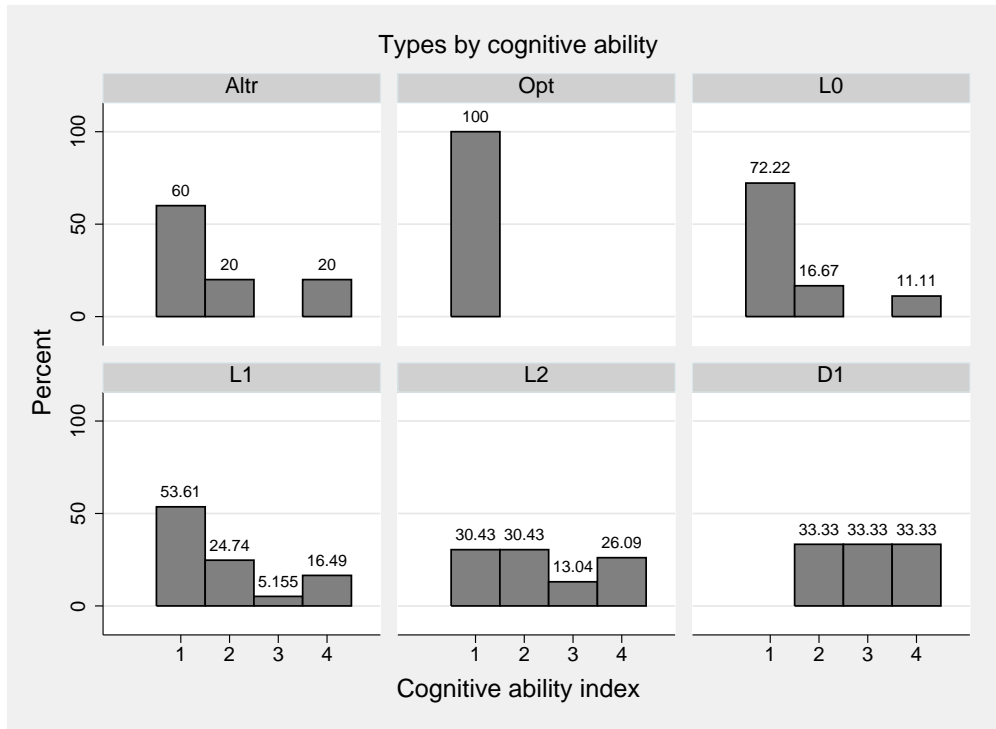


Figure 4: Cognitive abilities of different types

We now turn our attention to the relationship between the “revealed” preferences in the dictator game and behavior in strategic-form games. Figure 5 plots the distributions of behavioral types as a function of the alternative chosen in the dictator game and cognitive abilities.

Independently of the cognitive ability, the distribution of behavioral types differs with the alternative chosen in the dictator game (likelihood-ratio χ^2 -tests, $p < 0.01$ and $p < 0.02$). In particular, the large majority (more than 80%) of *Altruistic* types have chosen alternative *C* in the dictator game, the alternative associated with a taste for social efficiency. Moreover, for the subjects with a cognitive ability of one, this constitutes the main difference. *Altruistic* contributes more than 50% to the χ^2 statistic. However, for subjects with a cognitive ability of two or more, the distributional differences are

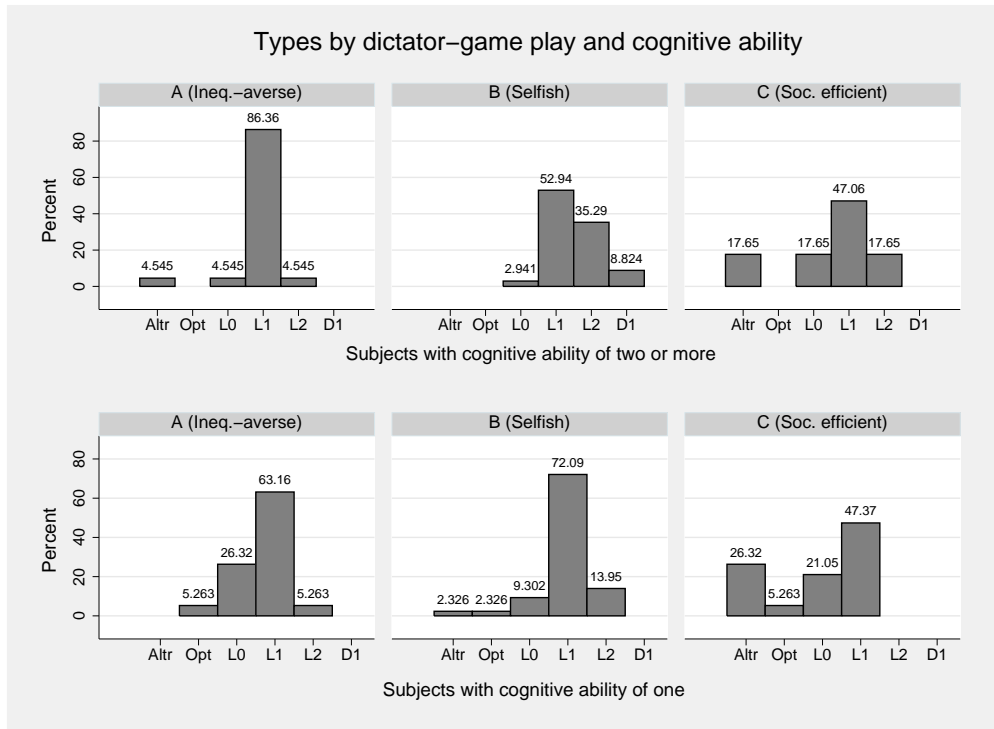


Figure 5: Choices in dictator games and behavioral types

driven by both *Altruistic* and *L2* types. These two types contribute more than 55% to the χ^2 statistic. A switch from alternative *A* or *C* to alternative *B* in the dictator game increases the relative frequency of *L2* types considerably.

Finding 2 *Subjects who have chosen the efficient alternative C in the dictator game are more likely to behave as Altruistic. Conditional on a cognitive ability of two or more, subjects who have chosen the (selfish) alternative B in the dictator game are more likely to behave as L2 type than their non-selfish counterparts.*

To understand this last finding, recall that *Lk* types are defined with respect to “selfish” preferences, i.e., assuming that the monetary payoffs

coincide with a subject's preferences. Also, subjects with a cognitive ability of two or more have the ability to perform counter-factual reasoning, as the behavior of an $L2$ type requires. Both observations explain why the relative frequency of $L2$ types is greater among subjects with a cognitive ability of two or more, who have chosen the (selfish) alternative B in the dictator game.

We now investigate whether there are some treatment effects, i.e., if and how subjects react to information provided to them. We hypothesize that subjects with a cognitive ability of one do not react to information about either the cognitive ability (*red-hat* treatment) or preferences (*dictator* treatment) or both (*fullinfo* treatment) of their opponents, while subjects with a cognitive ability of two or more do react. Intuitively, the information about the number of red-hat puzzles the opponent has solved or about the choice in the dictator game is a signal about the cognitive ability and preferences of the opponent and, thus, about the opponent's view of the strategic situation. To make use of this information, however, requires to be able to think through the opponent's eyes, which subjects with a cognitive ability of two or more can do, while subjects with a cognitive ability of one cannot. Figure 6 shows the distribution of behavioral types per treatment for the subjects with a cognitive ability of one.

A likelihood ratio χ^2 -test shows that there is no significant association between the treatment and the type frequencies ($p > 0.28$). Inspecting the graph and the contribution to the χ^2 statistic reveal that if there is an effect at all, which is not significant though, then it is the near disappearance of *Altruistic* and the increase in the frequency of $L0$ types in the *dictator* treatment.

Finding 3 *Subjects with a cognitive ability of one do not significantly react to information. If at all, information about the dictator games decreases the fraction of Altruistic types and increases the fraction of $L0$ types.*

We now turn to the subjects with a cognitive ability of two or more.

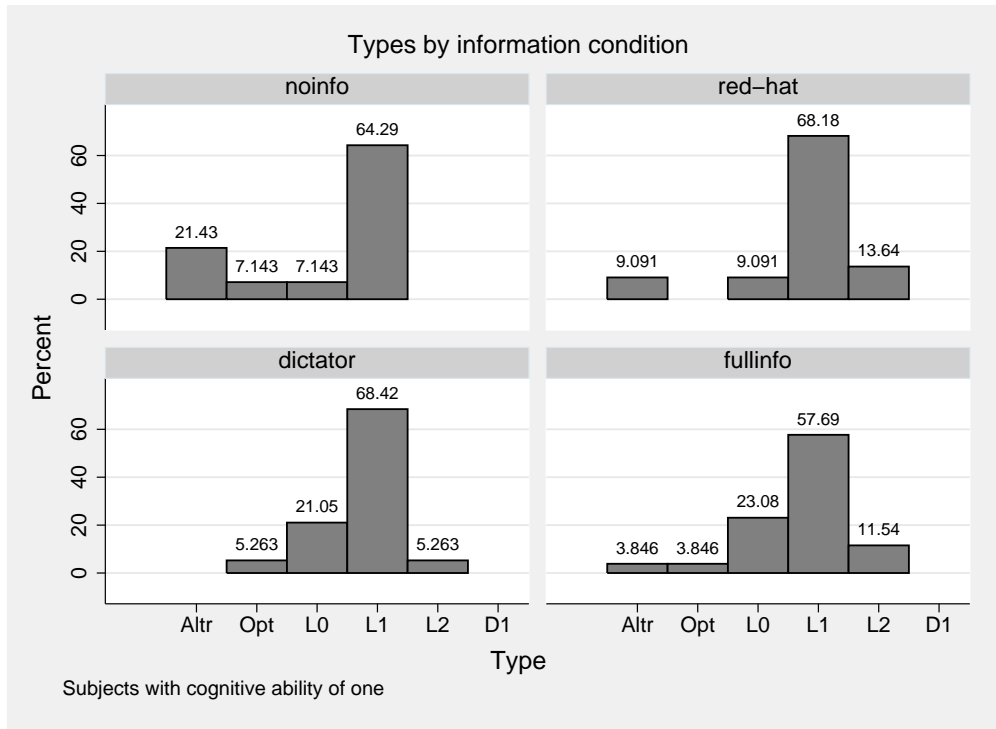


Figure 6: Treatments and behavioral types (subjects with a cognitive ability of one)

Figure 6 shows the distribution of behavioral types per treatment for these subjects. Visual inspection suggests that there is an association between the treatments and the distribution of types. A likelihood ratio χ^2 -test confirms this ($p < 0.065$).

A closer analysis of the types that contribute the most to the likelihood ratio χ^2 score shows that these are *L2* and *Altruistic*. The *Altruistic* score contribution in the *dictator* treatment is about 40% of the total score, whereas the contribution of the *L2* types across all treatments is close to 30%. This implies that the differences in type distributions across treatments are largely driven by the changes in relative frequencies of these two types.

A closer look at Figure 6 is revealing: there is no *Altruistic* type in the




Figure 7: Treatments and behavioral types (subjects with a cognitive ability of two or more)

noinfo and *red-hat* treatments, i.e., the treatments with no information about the choice in the dictator game, while there are *Altruistic* types in the two other treatments. Moreover, in the treatments where information about the dictator games was revealed, three-quarters of the subjects behaving as *Altruistic* have observed their opponents choosing alternative *A* or *C* in the dictator game, i.e., a “non-selfish” alternative. This provides some evidence that *Altruistic* type chose the socially efficient allocation after strategic considerations.²³

Finding 4 *A subject with a cognitive ability of two or more is more likely*

²³All the subjects with higher cognitive ability classified as *Altruists* had chosen a non-selfish option.

to behave as Altruistic if he is informed that his opponent did not choose the “selfish” alternative in the dictator game.

Lastly, we consider the differences between treatments *noinfo*, *red-hat* and *fullinfo*. See Figure 7. From the *noinfo* treatment to the *red-hat* treatment, the fraction of *L2* types increases by about 130%, whereas the increase is even greater from the *dictator* treatment to the *fullinfo* treatment (more than 250%). The increase of the fraction of *L2* types comes at the expense of *L1* types (and also the *L0* types in the *red-hat* treatment). Furthermore, it is illuminating to consider the sub-sample of subjects who have chosen the selfish alternative in the dictator game. In that sub-sample, we are only left with *k*-level types plus a few *D1* types; *Altruistic* and *Optimistic* disappear. The fraction of *L2* types doubles: from 0.25 to 0.5 from the *noinfo* treatment to the *red-hat* treatment and from 0.22 to 0.44 from the *dictator* treatment to the *fullinfo* treatment. A Mann-Whitney U-test confirms that “selfish” subjects with a cognitive ability of two or more behave more sophisticatedly in treatments where information about the cognitive abilities of opponents is provided ($p < 0.05$, one-tailed).

Finding 5 *Subjects with a cognitive ability of two or more behave more sophisticatedly when information about the cognitive abilities of opponents is provided.*

A simple explanation for the shift towards more sophisticated behavior when information about the cognitive abilities of others is given is as follows. Without information, subjects with a cognitive ability of two or more are overconfident about their own cognitive abilities and expect most other subjects to have worse cognitive abilities. With information, however, subjects have to revise their expectation upwards. This explains the significant increase of the number of *L2*-types at the expense of *L1*-types. Moreover, since *L2* best replies to *L1* and *L1* is the most prevalent type in the sample, *L2* is nearly optimal. In fact, compared to *Sophisticated*, the ideal of game theory, *L2*'s loss in payoff (in our sample) is only 1.15%. In other words, *L2*

types make almost 99% of the payoff a *Sophisticated* type would have made, and *L2* is a substantially simpler heuristic. In sharp contrast, behaving as a *Nash* type would have guaranteed no more than 90% of the payoff of a *Sophisticated* type. In our experiment, behaving as an *L2* type proved to be an extremely sophisticated behavior.

7 Conclusion

This paper has analyzed the relation between cognitive abilities and behavior in strategic-form games. To this end, we first measured subjects' cognitive abilities with the help of a computerized version of the red-hat puzzle. A large fraction (47.4% according to our conservative measure) of subjects were able to perform counter-factual reasoning and showed the cognitive ability required to put themselves in the position of their opponents. These subjects satisfied one of the most essential idea of Game Theory, *strategic thinking*.

In a second step, we estimated behavioral types from the choices made by the subjects in the twelve strategic-form games and related it to the subjects' cognitive abilities. Prima facie, the behavior of our subjects did not seem very sophisticated, since we did not find evidence for *Equilibrium* or *Rational expectation* types. Closer inspection revealed that subjects showed a remarkable level of sophistication though. Comparing the subjects without the ability to perform counter-factual reasoning with the subjects with this ability revealed a large shift from *L0* and *L1* types towards *L2* and *D1*. Most remarkably, types *L2* and *D1* behaved near optimal. The expected payoff for these types (given the empirical distribution of choices in the sample) totaled about 99% of the payoff a subject with rational expectations would have been able to achieve. For comparison, playing according to Nash equilibrium would have resulted in a much lower expected payoff of about 88% of the rational expectation payoff.

Subjects, who were able to perform counter-factual reasoning, also made use of information on the opponents' in a sophisticated and profitable man-

ner. Providing these subjects with information on the opponents' cognitive abilities further shifted the type distribution towards the near optimal types of $L2$ and $D1$. This is good news for theories on learning in games. Overall, the level of strategic sophistication of the subjects with the capacity for counter-factual reasoning is well beyond our initial expectation.

A Games

This section presents the 12 games that each subject played. For any $i > 1$ even, game i is the transposed of the game $i - 1$. Games 5 and 7 are taken from Costa-Gomes and v. Weizsäcker (2008), while game 9 is adapted from Costa-Gomes and v. Weizsäcker. In parentheses, we indicate the number of rounds of deletion of strictly dominated strategies required to reach the Nash equilibrium, while we indicate in bold the Nash payoff. Iteration n consists in deleting all pure strategies that are strictly dominated when the opponent's strategy space is given by the strategies not deleted at iteration $n - 1$.

	A	B	C
A	(47,56)	(13,68)	(17,17)
B	(62,37)	(35,45)	(19,21)
C	(46,21)	(20,22)	(12,19)

Game 1 :(1,1)

	A	B	C
A	(82,63)	(44,37)	(14,72)
B	(92,21)	(26,29)	(48,36)
C	(36,17)	(71,41)	(16,63)

Game 3 :(2,1)

	A	B	C
A	(73,80)	(20,85)	(91,12)
B	(45,48)	(64,71)	(27,59)
C	(40,76)	(53,17)	(14,98)

Game 5 :(3,2)

	A	B	C
A	(74,38)	(78,71)	(26,43)
B	(96,12)	(10,89)	(37,25)
C	(15,51)	(83,18)	(39,62)

Game 7 :(2,3)

	A	B	C
A	(30,59)	(34,91)	(96,43)
B	(36,48)	(85,33)	(39,18)
C	(49,86)	(43,14)	(25,55)

Game 9 :(4,3)

	A	B	C
A	(92,41)	(36,26)	(24,22)
B	(43,17)	(70,50)	(40,87)
C	(75,16)	(49,75)	(57,35)

Game 11 :(∞,∞)

B Behavioral types

Table B compactly presents the predicted play of an individual of type θ in any of our twelve games. In each cell of the table, the vector (\cdot, \cdot) represents the predicted play of the row player and the column player (the first element corresponds to the row player). For instance, an *Altruistic* type is predicted to play A in game G7 as a row player and B as a column player. To obtain the predicted play in game Gn for $n > 1$ even, it suffices to consider the play in game $G(n-1)$ and to permute the vector. For instance, in game G4, pessimistic is predicted to play C as a row player and B as a column player.

Types	G1	G3	G5	G7	G9	G11
Altruistic	(A,A)	(A,A)	(A,A)	(A,B)	(A,C)	(A,A)
Pessimistic	(B,B)	(B,C)	(B,B)	(C,A)	(B,A)	(C,B)
Optimistic	(B,B)	(B,C)	(A,C)	(B,B)	(A,B)	(A,C)
L1	(B,B)	(B,C)	(A,A)	(A,B)	(A or B,A)	(C,B)
L2	(B,B)	(B,C)	(A,B)	(C,A)	(C, A or B)	(B,B)
D1	(B,B)	(B,C)	(A,B)	(C,B)	(B, A)	(C,B)
D2	(B,B)	(B,C)	(B,B)	(C,B)	(B, A)	(C,B)
Sophisticated	(B,B)	(B,C)	(A,B)	(A,B)	(B, B)	(B,B)
Minimax regret	(B,B)	(A,B)	(A,A)	(A,B)	(A,A)	(B or C,B)
Best reply to altruistic	(B,B)	(B,C)	(A,B)	(C,B)	(A,B)	(A,A)
Taste for efficiency	(B,B)	(C,B)	(A,A)	(A,B)	(C,A)	(No pure,No pure)
Inequity aversion	(B,B)	(C,B)	(A,A)	(A,B)	(C,A)	(No pure,No pure)
L3	(B,B)	(B,C)	(B,B)	(C,C)	(B or C, A)	(B,C)
NE	(B,B)	(B,C)	(B,B)	(C,C)	(C, A)	(A,A)

Table 5: Predicted behaviors in all games

References

- Arad, A. and A. Rubinstein (2010). The 11-20 money request game: Evaluating the upper bound of k-level reasoning. Eitan Berglas School of Economics, Tel Aviv.
- Bayer, R.-C. and L. Renou (2007). Measuring the depth of iteration in humans. In L. Oxley and D. Kulasiri (Eds.), *MODSIM 2007 International Congress on Modelling and Simulation*, pp. 379–385.
- Bayer, R.-C. and L. Renou (2009). Logical omniscience at the laboratory. University of Leicester, School of Economics Working Paper.
- Beard, T. R. and J. Beil, Richard O. (1994). Do people rely on the self-interested maximization of others? an experimental test. *Management Science* 40(2), 252–262.
- Burchardi, K. B. and S. P. Penczynski (2010, April). Out of your mind: eliciting individual reasoning in one shot games. Unpublished Manuscript.
- Cabrera, S., C. Capra, and R. Gomez (2007). Behavior in one-shot travelers dilemma games: model and experiments with advice. *Spanish Economic Review* 9, 129–152.
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119(3), pp. 861–898.
- Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117(3), 817–869.
- Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica* 69(5), 1193–1235.

- Costa-Gomes, M. and G. Weizsäcker (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies* 75(3), 729–762.
- Costa-Gomes, M. A. and V. P. Crawford (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review* 96(5), p1737 – 1768.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39(1), pp. 1–38.
- Dufwenberg, M., R. Sundaram, and D. J. Butler (2010). Epiphany in the game of 21. *Journal of Economic Behavior & Organization* 75(2), 132 – 143.
- El-Gamal, M. A. and D. M. Grether (1995). Are people bayesian? uncovering behavioral strategies. *Journal of the American Statistical Association* 90(432), 1137–1145.
- Fehr, E. and K. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114(3), 817–868.
- Fischbacher, U. (2007). Z-tree - Zurich toolbox for readymade economic experiments. *Experimental Economics* 10(2), 171–178.
- Fudenberg, D. and J. Tirole (1991). *Game Theory*. Cambridge, Mass.: MIT Press.
- Gneezy, U., A. Rustichini, and A. Vostroknutov (2010). Experience and insight in the race game. *Journal of Economic Behavior & Organization* 75(2), 144 – 155.
- Goeree, J. K. and C. A. Holt (2004, February). A model of noisy introspection. *Games and Economic Behavior* 46(2), 365–382.

- Guttman, L. (1944). A basis for scaling qualitative data. *American Sociological Review* 9(2), pp. 139–150.
- Guttman, L. (1950). The basis for scalogram analysis. In E. A. Suchman, L. C. DeVinney, S. A. Star, and R. M. Williams Jr (Eds.), *Studies in Social Psychology in World War II: The American Soldier*, Volume IV. Wiley New York.
- Haruvy, E., D. O. Stahl, and P. W. Wilson (1999). Evidence for optimistic and pessimistic behavior in normal-form games. *Economics Letters* 63(3), 255 – 259.
- Ho, T.-H., C. Camerer, and K. Weigelt (1998). Iterated dominance and iterated best response in experimental "p-beauty contests". *The American Economic Review* 88(4), 947–969.
- Huyck, J. B. V., J. M. Wildenthal, and R. C. Battalio (2002). Tacit cooperation, strategic uncertainty, and coordination failure: Evidence from repeated dominance solvable games. *Games and Economic Behavior* 38(1), 156 – 175.
- Little, R. J. A. and D. B. Rubin (1987). *Statistical Analysis with Missing Data*. New York: Wiley.
- McKelvey, R. D. and T. D. Palfrey (1992). An experimental study of the centipede game. *Econometrica* 60(4), 803–836.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review* 85(5), 1313–1326.
- Osbourne, M. J. and A. Rubinstein (1994). *A course in game theory*. MIT Press.
- Redner, R. A. and H. F. Walker (1984). Mixture densities, maximum likelihood and the em algorithm. *SIAM Review* 26(2), pp. 195–239.

- Rey-Biel, P. (2009). Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior* 65(2), 572 – 585.
- Stahl, D. O. and P. W. Wilson (1994, December). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization* 25(3), 309–327.
- Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior* 10(1), 218–254.