

Logical abilities and behavior in strategic-form games*

Ralph C. Bayer[†] & Ludovic Renou[‡]

May 19, 2016

Abstract

This paper investigates the impact of cognitive abilities and preferences on behavior in games. We first measure subjects' logical abilities and social preferences and then evaluate their influence on behavior in strategic-form games. For this purpose, we estimate how logical abilities and preferences relate to the probability of a subject being of a specific behavioral type. We find that better logical abilities increase the likelihood of sophisticated behavior. This shift is stronger for subjects who are classified as "selfish". However, not even the most logically able, selfish subjects are identified as Nash types.

Keywords: logical 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 and are grateful to two anonymous referees and the editor for helpful comments.

[†]Department of Economics, 10 Pulteney Street, Adelaide 5005, Australia. Phone: +61 (0)8 8313 4666, Fax: +61 (0)8 8223 1460. ralph.bayer@adelaide.edu.au

[‡]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom. lrenou@essex.ac.uk

1 Introduction

This paper investigates experimentally how individual characteristics, in particular logical abilities, relate to individual behavior in strategic-form games. A large body of evidence documents systematic behavioral departures from selfish Nash play, particularly in environments with limited learning opportunities. (See Camerer (2003) and references therein.) It is still relatively unclear what drives particular behaviors. There are almost unlimited potential causes for why a player might deviate from choosing strategies prescribed by Nash equilibrium under purely selfish preferences. Bounded rationality, other-regarding preferences, lack of belief in the rationality of the opponent and the belief that the opponent might have other-regarding preferences are just a few possible causes. This paper is mainly concerned with the impact of cognitive abilities on the level of strategic sophistication shown in strategic-form games, while controlling for social preferences and beliefs about the other player.

Our experiment has two essential components. The first component measures 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).¹ 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. Moreover, we vary what a subject knows about the cognitive abilities and social preferences of the opponent, such that we can control for beliefs.

We use the choices made by the subjects in the twelve strategic-form games to identify different behavioral types, which are more or less strategically sophisticated. 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 *behavioral* types. Although we consider fifteen possible behavioral types, only six of them are of significant importance in the sample. An *Altruistic* type identifies the strategy profile that maximizes the sum of his

¹See also Weber (2001).

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 logical 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 score of two or more has the ability to put himself in the position of the opponent (strategic thinking), while a subject with a score 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 7%, 2%, 14.2%, 59.9%, 15.4% and 1.6%, respectively. Most subjects (about 75%) 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 *Omniscient* (rational expectation) or *Nash 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.² 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.

²In 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 and attributed the result to the particular games chosen.

Finally, we turn our attention to the relation between logical abilities and behavioral types. Intuitively, one might expect the behavior of individuals with better cognitive abilities to conform more closely with solution concepts such as iterated deletion of strictly dominated strategies or Nash equilibrium. We find evidence for cognitively more able players playing more sophisticated strategies. Surprisingly, subjects do not take into account information they receive about the cognitive ability and social preferences of their opponents.

The general pattern we observe is as follows: the higher the index of logical abilities, the more sophisticated (i.e., $L1$, $L2$ or $D1$) is the behavioral type. In particular, we find that the best predictor for play is the $L2$ type if a subject has a logical ability of three or more, while the best predictor is the $L1$ type if a subject has a logical ability of two or less. Logical abilities do have a significant impact on play. However, the impact of logical 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 among 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 $Nash$ type would achieve less than 90% of the rational expectation payoff.³ In terms of outcomes, the subjects' behavior is quite sophisticated after all, and subjects with better logical abilities are more likely to be of the types that achieve higher payoffs.

These results are consistent with the findings by Burnham et al. (2009) and Braas-Garza et al. (2012), who show that certain measures of cognitive abilities are correlated with sophistication in beauty contests. In addition to those findings, we are able to look deeper into the forces driving this relationship. We do not find differences in behavior regardless of the information received about the cognitive abilities and social preferences of the opponent. This implies that the impact of cognitive abilities on strategic sophistication is likely to be direct, rather than indirect. Better cognitive abilities itself lead to more sophisticated strategies. The alternative explanation that better cognitive abilities imply optimistic beliefs about the rationality of others which in turn cause more sophisticated strategies is not

³ $L0$ types achieve even less with about 80%.

supported by our data. In that case, varying the information about the opponent’s logical abilities and dictator game choices should change the relationship between logical abilities and play in the strategic form games, which we did not observe.

This paper contributes to the large literature on iterative reasoning in games, which includes 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 important contributions.⁴ 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 in strategic-form games. The use of these games, however, suffers from the lack of control for beliefs about the rationality of others and for social preferences.⁵ A novelty of our experiment is that we are able to provide these controls with a computerized version of the red-hat puzzles. The second strand of literature our paper contributes 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, 2012; Burchardi and Penczynski, 2014).⁶ Finally, our paper contributes to the large literature on how cognitive abilities correlate with behavioral biases, risk and time and social preferences (for recent papers see Oechssler et al., 2009; Dohmen et al., 2010; Benjamin et al., 2013).

2 Experimental Design

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 and were

⁴We refer the reader to chapter 5 of Camerer (2003) for a survey of earlier literature on this topic.

⁵Gneezy 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.

⁶Beyond 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.

recruited using ORSEE (Greiner, 2015). For each of our four conditions (which we will explain later), we conducted two sessions resulting in between 36 and 42 subjects per condition. 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 can be found in Appendix F. The experiment consisted of three stages. In the first two stages we measured logical abilities and got an indication of subjects' social preferences, while the third stage collected the subjects' play in normal-form games. The elicitation of social preferences could not be distinguished by the subjects from the normal-form games part, as it took place within the same computer interface.

2.1 A measure of logical abilities

This section presents the first stage of our experiment, the red-hat puzzles, and describes our measure of logical ability. Along with a measure of “social preferences,” this measure of logical ability is at the heart of our analysis of behavior in strategic-form games.

A remark is in order. There is a large number of cognitive ability tests: the Wechsler Adult Intelligence Test, the Stanford-Binet Intelligence Test, the Woodcock-Johnson Tests of Achievement, the Bender-Gestalt Test, just to name a few. These tests typically measure general cognitive abilities by summing across tests of several specific aptitudes, usually verbal aptitude, quantitative aptitude, and sometimes technical aptitude. We decided not adopt any of these general cognitive tests for a number of reasons. First and foremost, we wanted a *specific* test of cognitive ability able to measure the ability of subjects to perform complex chains of reasoning, in particular counter-factual reasoning. We were less interested in the literacy or visual aptitudes of our subjects. Secondly, our test is simple to administer. Thirdly, the scale of our measure ranges from 0 to 4, while the scale of most general cognitive ability tests ranges from 0 to 250. Such a wide scale would have been problematic for our econometric analysis (see section 2.3). Moreover, our test is easy to interpret: a higher score means a higher ability to perform complex chains of reasoning. This is not the case with general cognitive ability tests. Finally, the logical puzzle we are using requires a stepwise

counter-factual reasoning that is very similar to that used in finding Nash equilibria in games that are dominance-solvable.

2.1.1 The Red-Hat Puzzle

We first present a logical puzzle, which is the basis for our measure of logical ability. 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.” Individuals then learn the answers of the other individuals 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).⁷

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.” However, if Ann says “I have a red hat” right away, then Bob can infer that Ann cannot have seen any other red hats. So Bob knows he has got a white hat. If Bob sees two red hats, then he will be able to deduct his hat color in three steps by applying the logic of how to determine your hat color when you see one red hat to the two people who wear red hats. If they can figure out that they have red hats by their second answer, he knows that they cannot have seen another red hat and his hat is white. If they cannot, then he knows that they must have seen another, i.e. his, red hat.

In general, the greater the number of red hats an individual observes, the more complex is

⁷The 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). Weber (2001) ran experiments based on a Bayesian Game derived from the logical puzzle.

the reasoning needed to logically infer the color of one’s own hat. An individual needs $m + 1$ iterations to figure out his hat color, where m denotes the number of red hats this individual observes. In Bayer and Renou (2007) we showed that this puzzle is both difficult (the median subject can only solve puzzles with up to one observed red hat) and discriminatory (more than 10 percent were able to solve all puzzles including those where they observed three red hats).

We now describe our experimental protocol. In our experiment, a human subject was paired with three computers, which were acting as “players.” Computers were programmed to choose the logically correct answers in each round (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.” In a given puzzle, the first time a subject had to choose an answer, 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.⁸ This stopping procedure is necessary to avoid logical inconsistencies. To see this, 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.”

⁸In a given round, announcing a hat color was correct only if it was actually possible to infer the hat color at this given round.

However, this 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).⁹

2.1.2 Our measure

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 making. Firstly, as expected, the more iterations are required to correctly infer one’s hat color, the lower is the percentage of correct answers. Secondly, and somewhat surprisingly, there is no significant difference between solving a puzzle requiring three iterations and one requiring four iterations (26.6% vs. 26.3%, $p > 0.92$, test of proportions), while there is a significant difference between solving a puzzle requiring two iterations and one requiring three or four ($p < 0.001$ for all pairwise tests of proportions). This suggests that if an individual can perform three steps of iterative reasoning, then she can also do four steps.

⁹For our subjects losing AUD 2.5 with a single wrong decision was quite a strong incentive. A student job was paying about ten Dollars per hour and the median hourly wage in South Australia was about AUD 20 at the time of our experiments. We also conjecture that the loss frame increases effort through loss-aversion.

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.

[Table 1 about here]

We now construct a simple index of logical 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 $n = m$). 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 $(3, 4)$.

The index we adopt is an integer ranging from 0 (all puzzles incorrectly solved) to 4 (all puzzles correctly solved). The index is supposed to tell us how many steps of counter-factual reasoning a person can do. An index of i predicts that a subject solves all puzzles requiring i or less iterations and fails all others. For each subject, we simply ascribe the index that minimizes the number of wrong predictions.¹⁰

[Table 2 about here]

Table 2 shows a few examples of how our index works. The table reports which of the seven puzzles (requiring one to four steps) a person solves correctly (c) or wrongly (w) and shows the index assigned as well as how many prediction errors the assigned index leads to. The puzzles are identified by the number of red hats seen and the total number of red hats (m, n) as explained above. A subject has a pattern of answers consistent with the index if he has solved all puzzles requiring i or less iterations and failed all others. Remarkably, 78.6% of the subjects exhibit a consistent pattern of answers. The first example in the table above shows a person who has a pattern consistent with a person who can do three steps but not four and is therefore assigned an index of three. Since the pattern is fully consistent there are no prediction errors. The second example shows a person who gets all step one and step two puzzles right but fails to solve all higher puzzles except one step four puzzle. If we assign

¹⁰If there are multiple minimizers, we chose the smallest.

this person an index of two, then there is one prediction error (shown in bold), as a step two person is not predicted to solve a level four problem. It is easy to check that any other index value would lead to more than one prediction error. Our procedure would interpret the one correct step-four outcome as luck. In the second example in the table above our procedure assigns a mistake in a low level puzzle to bad luck. The final example shows how our decision to choose the lower index number if two at the same time minimize the prediction errors turns out. Here the subject gets all step one and two puzzles and one of the two step three puzzles right. Both, assigning an index of two or three lead to one prediction error. In this case we are conservative and assign an index of two.

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 90 percent, 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.

[Figure 1 about here]

Looking at the distribution, it catches the eye that the frequency of four step individuals is higher than that of three step subjects, which cause the distribution to be bimodal. There are two potential reasons for this. Firstly, an index value of four steps means that the subject can do *at least* four steps of iteration, as we do not have any puzzles that require more steps of iterative reasoning. So the real distribution of iterative abilities could be unimodal, as the subjects with the index value of four in fact should be distributed over the bins for iteration abilities of four, five, six, up to infinity, which we cannot distinguish. Secondly, there is also a possibility that some subjects categorized as level four have been able to infer the structure of the problem from their counterfactual reasoning. In other words, some people who are able to counter-factually reason a few steps are able to extrapolate the pattern for more difficult tasks without actually requiring to iterate further. This effect can either be due to learning from easier puzzles, which is possible but difficult due to the lack of feedback in our design, or from introspection. From a companion paper Bayer and Renou (2009), where we investigate the impact of feedback, order of puzzles, repetition and one-shot play on performance, it can be inferred that introspection is likely to be the dominant force.

In light of the results, it makes sense to think of alternative ways of characterizing the cognitive abilities of our subjects. First, one could divide subjects into those who can put themselves into the shoes of someone else and those who cannot. All subjects with a score higher than one fall into the former category and those with a score of one into the latter. A second natural way of grouping subjects are the following three groups: a) no counter-factual reasoning (score of one), b) some counterfactual reasoning (score of two) and c) understanding the structure of the problem (scores of greater than two). This way of combining levels three and four makes sense, since about 74 percent of the subjects that were able to solve the problems of level three were also able to solve the most difficult puzzle. We will make use of both of these alternative categorizations, when we investigate the link between cognitive abilities and play in strategic-form games.

2.2 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 Dollars 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, in order to minimize the rate of possible mistakes. As we wanted to exclude other motivations than social preferences (such as reciprocity or conformity), subjects played one game as dictator and one game as a receiver with different partners and without feedback.

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 subject 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 , while the own payoff decreases.

Our choice of allocations was guided by the idea that a person who prefers C is willing to sacrifice quite a bit of her own payoff in return for social efficiency. A person who chooses A tends to dislike advantageous 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*.

2.3 Behavior in games

The last stage of our experiment 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 had a unique Nash equilibrium in pure strategies.¹¹ 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 Appendix C.

Six games and their respective transposes constitute 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.¹²

A subject was initially seeing his own payoff table, but not the payoff table of his opponent. The subject had to wait for 15 seconds before being able to see the payoff of his opponent; she had to click on a button to do so.¹³

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

To summarize, we obtain a measure of logical ability with the red-hat puzzle. To control for social preferences, we use dictator games. We randomly rematched subjects in between

¹¹In all dominance solvable games, a strictly dominated strategy was dominated by a *pure* strategy.

¹²We found no order effects.

¹³We included the delay and the option to view the payoff of one's opponent to gather some additional information, e.g., whether a player with a dominant strategy viewed the payoff of his opponent. We found no significant effect, however.

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 consists of four different conditions, which differ in the level of information given to subjects in stage three. In the *fullinfo* condition, 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 cognitive ability of the opponent, whereas the choice in the dictator game is a signal about the opponent's preferences. In the *red-hat* condition, only the number of correctly solved RHPs is provided. In the *dictator* condition, the subjects were only told the opponent's choice in the dictator game, while in the *noinfo* condition no information was provided.

3 Hypothesis and experimental results

In what follows we will present the results of our experiments. Before we get into the details, let's talk about hypotheses. Our general goal of the paper is to find out how cognitive abilities as measured by the RHP shape play in normal form games. At the same time we want to control for social preferences and use our different treatments to draw further inferences about what drives play in the normal-form games. One can hypothesize that Nash equilibrium (our standard solution concept) is the process that drives our data. Then the following hypothesis should hold.

Hypothesis 1 *Observed behavior is largely consistent with Nash equilibrium.*

If the hypothesis above holds then one might investigate where observed deviations result from. Candidates are that a subject cannot locate a selfish Nash equilibrium or has social preferences. Since we have information on subjects' cognitive abilities and on their preference we can test the following hypotheses if required.

Hypothesis 2 *Subjects with higher cognitive ability play Nash more often.*

Hypothesis 3 *Subjects who have chosen the selfish option in the dictator game should play selfish Nash more often.*

Now suppose that subjects are boundedly rational and choose their actions based on a mixture of their own cognitive ability, preferences and their beliefs about the ability and the preferences of their opponents. Then we can provide some alternative hypotheses. In what follows, when we speak of sophistication then we mean the level in a level-K style model. First, we would expect that the own ability is one driver of displayed strategic sophistication. Somebody who cannot perform more than one step of counterfactual reasoning, for example, cannot perform the logic required for consciously behaving as a level 2 player in a level-K model.

Hypothesis 4 *Subjects with better cognitive abilities are associated with more sophisticated behavioral types.*

Note that this does not imply that two players with different cognitive abilities will have to behave differently or that players with the same cognitive abilities have to behave identically. Two highly cognitively able players, for example, could have different beliefs about the cognitive abilities of others and as a result one could behave as level 1 and the other as level 2. If subject's beliefs about the other player are an important determining factor of how people behave, then providing a signal about the opponents cognitive abilities should have the effect proposed in the next hypothesis.

Hypothesis 5 *If subjects know the cognitive abilities of the opponent, then on average the sophistication increases with the abilities of the opponent.*

Similarly, if beliefs about the preferences are a driver, then the observed types should depend on a signal about them. Observe that level-K is typically formulated in a way that is implicitly assumes that payoffs in a game represent preferences over outcomes. Subjects with social preferences are then typically misclassified as less sophisticated than they actually are. Often their behavior will be picked up by level zero types or by other non-strategic types. The following hypothesis arises from this logic.

Hypothesis 6 *If subjects observe a signal about the preferences of the opponent that deviates from selfishness, then they are associated more often with non-strategic types.*

3.1 Nash play

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 , under the assumption that preferences are indeed given by G .

[Table 3 about here]

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 any of their actions with equal probability, 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.

From a mechanism design perspective, the poor predictive power of Nash equilibrium suggests that Nash implementation is problematic. Furthermore, the relatively low compliance rate in games with a strictly dominant strategy suggests that implementation of social choice functions in dominant strategies is not even a guarantee. Moreover, even if we restrict our attention to subjects having chosen the action B , i.e., the selfish action, in the dictator game, the aggregate compliance rate with strict dominance is still less than 80%. This seems

to indicate that preferences alone cannot explain why the dominant strategies are not played. Generally, Nash equilibrium is a poor descriptor of observed play in the normal form games.

Finding 1 *Play in normal-form games is not guided by selfish Nash equilibrium.*

This result implies that we have to reject 1 and a further investigation of Hypotheses 2 and 3 is not necessary.

3.2 Information conditions

The treatment manipulation of our study was designed to shed some light on how subjects decide which strategy to play in normal-form games. We varied the amount of information a subject has about the cognitive abilities and social preferences of the opponent. In the four conditions we either gave no information ($N = 36$), information about how many RHPs the opponent solved ($N = 38$), information which dictator-game option the opponent chose ($N = 38$), or both ($N = 42$). A person, who forms conjectures about what the opponent is likely to do in a game in order to best-respond to it, should take the information provided into account. On the other hand, a subject who uses a fixed rule (i.e. playing like a level-K type or a non-strategic type), would disregard such information.

If subjects used the information we would expect a statistical association between the treatment and the distribution of strategies played in a particular game. The distributions of actions taken in the different games look extremely similar across treatments. We conducted Fisher Exact tests on the contingency tables for all 12 games. Only in a single game we could reject the null-hypothesis that there is no association. One successful test out of 12 is well within the realm of false positives. And indeed, using Fisher's method of combining the tests, we find that an association across all games is extremely unlikely ($p = .66$). The table below shows the p-values for the individual and the combined tests.¹⁴

[Table 4 about here]

¹⁴Fisher's combined probability test assumes that the individual tests are independent. This assumption might be violated in our case, since the same subjects play all the games. However, if the p-values of individual tests are correlated, then the test is biased towards rejecting the null, which strengthens our argument.

Finding 2 *Subject's play in the normal-form games does not differ across information conditions.*

This finding in itself is no proof yet that Hypotheses 5 or 6, which link information to behavior have to be rejected. It is still possible that information-dependent behavior in the treatments where information was provided aggregates to similar over-all behavior as that observed in the treatment without information. However, comparing the behavior of players who observe different signals about the cognitive abilities in the strategic-form games, does not show significant differences either. We take only subjects that have observed the performance of their opponents in the RHP ($N = 80$). Then we compare the distribution of behavior for subjects that saw scores that correspond to LO measures of 1,2,3 and 4. A Fisher exact test does not find a difference in any game and combining the games provides a combined p -value of 0.99.¹⁵ Given this result we have to reject Hypothesis 5. This finding is in contrast to Agranov et al. (2012), who find that manipulating the beliefs about cognitive abilities of others has an impact on play in guessing games. Gill and Prowse (2015) find that in the 11-20 game only subjects who have a high cognitive ability themselves adapt their behavior according to the information they have about the opponent.

Similarly, we do not find differences in behavior depending on the signal a subject receives about the social preferences of the opponent. Comparing play of informed subjects in the games conditional on which action the opponent had chosen in the dictator game does not show any significant differences for any game. Combining the tests using Fishers method yields $p = 0.28$. Therefore, we have to reject Hypothesis 6.

3.3 Behavioral types

So far we have established that neither Nash equilibrium, nor the information received about rationality and social preferences of the opponent are driving behavior. This implies that it is likely that subjects behave according to some fixed decision rules. In what follows, we are postulating some potential types and then via a mixture model estimate the fraction of the

¹⁵Other ways of slicing the sample do not change this result. If one. e.g. compares the distributions of players who faced subjects that could not put themselves in the shoes of others ($LO \leq 1$) and those who could ($LO > 1$) yields $p = 0.91$.

population behaving according to the different types.

We postulate that a subject can be one of fifteen possible behavioral types, labeled *Altruistic*, *Pessimistic*, *Optimistic*, *L0*, *L1*, *L2*, *L3*, *D1*, *D2*, *Equilibrium*, *Omniscient*, *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 maximum 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 conjecture. *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 this conjecture. *Equilibrium* conjectures that the opponent plays a Nash equilibrium strategy and best replies to this conjecture. *Omniscient* 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 four additional types that we ex-ante deemed promising. *Regret* minimizes his maximum 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.¹⁶

3.3.1 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, y_i) consists of the play in the twelve

¹⁶For 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 $u_i(a_i, a_j) \geq u_j(a_i, a_j)$.

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 , together with the outcome of the logical ability test, the choice made in the dictator game, the treatment played, and personal characteristics (age, degree, gender, etc). As already explained, we hypothesize that each subject has a behavioral type $\theta \in \Theta$, which determines a specific play in all games. (See Appendix D.) We assume that conditional on observables y , behavioral types are independently and identically drawn from the distribution p^y . Two important conditioning observables are the logical ability index and the choice made in the dictator games.

We further assume that in each game k , conditional on observables y , subject i of type θ plays the action $x_i^k(\theta)$ consistent with his type with probability $(1 - e_\theta^y)$ and plays any one of the remaining two actions with probability $e_\theta^y/2$. Note that if $e_\theta^y = 2/3$, 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:

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

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

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

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

$$\sum_{i \text{ with observables } y} \ln \left(\sum_\theta p^y(\theta) (1 - e_\theta^y)^{\bar{x}_i(\theta)} (e_\theta^y/2)^{12 - \bar{x}_i(\theta)} \right).$$

To maximize the log-likelihood $l((\cdot, \cdot))$ with respect to (p^y, e^y) , 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).

The confidence intervals are obtained by bootstrapping, as analytical standard errors are not available.¹⁷ For model selection, we use three different criteria for fit and discriminatory

¹⁷For the bootstrap, we re-sampled the original data with replacement, with the number of observations

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. We unfortunately cannot use Likelihood-ratio based test like the Vuong test, which was used in Kawagoe and Takizawa (2012), e.g., since due to the use of the EM algorithm we cannot be sure about the distribution of the test statistic. The higher the entropy, the less informative is the model.¹⁸

3.3.2 Econometric results

Table 5 reports our estimation results for our most preferred models for the whole sample and for sub-samples of subjects with cognitive abilities of one and above on. Models were selected as follows.¹⁹ For each logical ability score, we started with the full model containing all our fifteen types and eliminated the types with inconsistent error estimates.²⁰ After removing the inconsistent types, we had an intermediate model. For the final model selection, we then proceeded in two steps.

[Table 5 about here]

First, we removed the types with estimated error rates of close to two-thirds and replace these types by the type $L0$. Recall that for any game a type predicts that one of the three actions is played. If the error probability (i.e. the probability to choose any of the two alternative strategies with probability $e/2$) for a type is equal to $2/3$, then a player of this type mixes over the three alternatives with equal probabilities. This is just the behavior of a $L0$ type. So the types with estimated error rates of around $2/3$ all pick up the same behavior. Consequently, we created and estimated a pseudo-type $L0$ that always plays the middle action B , but has error rate e . Since the estimated error rate for the new type was equal to the number of observations in the original data, and re-performed the estimation. After 2000 re-samplings, we used the resulting distributions of estimates to determine the bootstrap confidence intervals.

¹⁸El-Gamal and Grether (1995) regard values below .38, which was the highest value they observed in their study, as good.

¹⁹All estimations are available upon request.

²⁰Estimated 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.

indeed close to two-thirds, we are confident that we successfully identify the $L0$ types in the sample.²¹ This procedure allows us to actually pick up $L0$ types, which a variety of studies are not able to do.

Second, we checked if reducing the model further improved parsimony without worsening the fit. Models were compared in terms of their ANE, AIC and BIC scores. Models reported in Table 5 dominate all other models but one in all criteria, and dominate the remaining model in two out of three criteria. Note that the selected models have relatively low ANEs, which confirms their excellent discriminatory power.²²

Table 5 shows the type estimation for the whole sample and also separately for subjects with an ability score of one and those of more than one. Note that in our puzzles an ability of more than one required putting oneself in the shoes of a computer, while a score of one did not. We first note that no subject is best described as the *Nash* type. This is consistent with the observations made in Section 3.1. 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 59.9% 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 (15.4%) are subjects classified as type $L2$.²³ 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 $L1$.²⁴ Next, we need to mention that we found a large number of subjects (more

²¹The five estimates we obtain for the different models all lie between 0.622 and 0.706. All bootstrapped confidence intervals contain $2/3$.

²²As robustness tests we also estimated a variety of models for the whole sample using the same model selection criterion. Additionally, we also estimated a model using all data, where error rates could differ not only by type but also by the level of cognitive ability. The resulting ex ante type probabilities were very close to those reported in Table 5. These estimation results are available from the authors on request.

²³In their econometric analysis of behavioral types without information about search patterns, Costa-Gomes et al. (2001) found $L2$ to be the most prevalent type.

²⁴An *Omniscient* type would certainly do better; but according to our most preferred models we did not find this type in the sample.

than 14%) exhibiting *L0* behavior. The remaining subjects were classified as either *Altruistic* (about 7%), *Optimistic* (2%) or *D1* (1.6%).

Finding 3 *Behavior in the strategic-form games is heterogenous and best described by a distribution of player types that does not include a Nash type.*

We now concentrate on the relationships between logical ability and behavior. The separate estimations for subjects with cognitive ability one versus more than one, show that being able to put oneself into the shoes of others (required for $\text{step} > 1$) leads to more sophisticated play in the strategic-form games. The mass of the distribution of types shifts from less sophisticated types such as *L0* and *L1* to more sophisticated types such as *D1* and more prominently *L2*. In particular, a subject with a logical ability greater than one is more than four times as likely to be of type *L2* than one with the logical ability of one.

Next we further separate the sample in order to see if type distributions differ across subjects with different cognitive abilities above one. Figure 2 plots the estimated probabilities to be of a certain behavioral type conditional on a logical ability score of one, two, and three and more (see the additional estimates in Table 7 in Appendix E). Remarkably, a subject with a score of three or more is about twice as likely to be of the *L2* type than a subject with a score of two, and about ten times as likely than a subject with a score of one. Mann Whitney U-Tests confirm the observation that subjects with better logical abilities are more likely to be of the *L2*-type ($p < 0.002$). In fact, conditional on a logical ability index of three or more, the most likely type is *L2*, while *L1* is the most likely type conditional on a logical ability index of two or less.

Also, a subject with an index of one is infinitely more likely to be of type *L0* than a subject with an index of two, and four times more likely than a subject with a score of three or more (Mann Whitney U-Test, $p < 0.001$). These observations are reassuring as *L0* behavior *does not* require to put oneself in the position of the opponent, which subjects with a logical ability index of one cannot do. On the contrary, *L2* behavior *does* require to put oneself in the position of the opponent, which subjects with a logical ability index of two or more can do.

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

Recall that an *Optimistic* chooses the action that maximizes his maximum payoff over his opponent’s decision. In sum, logical abilities have a strong impact on behavior in our games and Hypothesis 4 is confirmed.

Finding 4 *Subjects with higher logical ability are likely to play strategically and more sophisticatedly in the normal-form games.*

[Figure 2 about here]

We now turn our attention to the relationship between the “revealed” preferences in the dictator game, logical abilities, and behavior in strategic-form games. First note that there is no significant relationship between behavior in the dictator game and logical abilities. In an ordered probit regression with the logical ability as dependent variable the dummies for play in the dictator game are all insignificant ($p > 0.7$ and $p > 0.86$). In what follows we look at the association between behavioral type and behavior in the dictator game conditional on a certain level of logical ability. Figures 3 and 4 plot the estimated probabilities to be of a given behavioral type, conditional on the choice made in the dictator game and a logical ability index of one or one and more.²⁵

[Figure 3 about here]

For subjects with a logical ability score of two or more, the choice of actions A or C in the dictator game does not generate statistical differences in the probabilities to be of a certain type. In contrast to this, the choice of action B (the selfish action) does generate statistical differences in the probabilities to be of a certain type when compared with both the choice of A and C (p -values ranging from 0.001 to 0.007 for pair-wise Mann-Whitney U-Tests). For instance, subjects who have chosen the selfish action B in the dictator game are more likely to behave as $L2$ or $D1$ types than their non-selfish counterparts. To understand this observation, recall that Lk and Dk types are defined with respect to “selfish” preferences,

²⁵These probabilities are based on the two models from Table 5 and are the average of the probabilities of all subjects having chosen a certain action in the dictator game. In general it would have been possible to estimate six separate models for all combinations. However, due to the low number of observations for some combinations the estimation would not have been very precise.

i.e., assuming that the monetary payoffs coincide with a subject's preferences. Also, subjects with a logical ability of two or more have the ability to perform counter-factual reasoning, as the behavior of an $L2$ or $D1$ type requires. Both observations explain why the relative frequency of $L2$ and $D1$ types is greater among subjects with a logical ability of two or more, who have chosen the (selfish) alternative B in the dictator game. Moreover, notice that almost all subjects who have made a choice of B in the dictator game are of types $L1$, $L2$, and $D1$. Those types are strategic and obey dominance.

[Figure 4 about here]

For subjects with a logical ability index of one, we can notice that the probability to be of type $L1$ conditional on having chosen action B is statistically different from the probability to be of type $L1$ conditional on having chosen action A or C ($p < 0.034$ vs A , $p < 0.005$ vs C). This suggests that the probability to be of type $L1$ is not driven by preferences. Also, the probability to be of type $L0$ conditional on having chosen action A or C is statistically different (higher) from the probability to be of type $L0$ conditional on having chosen action B ($p < 0.008$). This suggests that the type $L0$ captures not only genuine bounded rationality, but also behaviors driven by differences in preferences. This confirms hypothesis 6.

Finding 5 *Social preferences have an influence on behavior in strategic form games. But differences in social preferences are not responsible for the correlation between logical abilities and strategic sophistication in the strategic-form games.*

4 Concluding Discussion

The paper has analyzed how individual characteristics, in particular logical abilities, relate to behaviors in strategic-form games. We have shown that the level- k model explains the data reasonably well. This study thus reinforces previous studies, e.g., Stahl and Wilson, (1994; 1995); Costa-Gomes et al. (2001); Camerer et al. (2004); Costa-Gomes and Crawford (2006).

The main contribution of the paper, however, is to relate individual characteristics to the likelihood to be of a certain behavioral type. Among other things, we find that subjects

with better logical abilities are more likely to be of type $L2$. Also, subjects with high logical abilities and who have chosen the selfish action in the dictator game are even more likely to be of type $L2$. Conversely, subjects with lower logical abilities or who have not chosen the selfish action in the dictator game are found to be more likely of type $L0$ or $L1$.

One potential useful application for our results could be mechanism design theory. As argued in Crawford et al. (2009), adopting a structural non-equilibrium view, in particular the level- k model, in mechanism design problems should lead to more effective designs in concrete problems. Yet, an essential difficulty for policy makers remains: how to estimate the distributions of Lk types in the targeted population? Crawford et al. (2009) assume that the distribution of Lk types is known to the policy makers, but it has to be estimated in practice. We do provide such an estimate conditional on two characteristics: logical abilities as measured in the red-hat puzzles, and concerns for others as expressed in the dictator games. Since logical abilities correlate with education and years of schooling and the concern for others may correlate with occupation, subjects studied or religion, this provides a possible route for a better estimation of the distribution of Lk types in a population. Unfortunately, in a recent paper Georganas et al. (2015) have demonstrated that type distributions might not be stable across games. We regard our study as a first step of what we believe to be an important area of future research.

References

- Agranov, M., E. Potamites, A. Schotter, and C. Tergiman (2012). Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior* 75(2), 449 – 463.
- Arad, A. and A. Rubinstein (2012). The 11-20 money request game: A level-k reasoning study. *American Economic Review* 102(7), 3561–73.
- 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.
- Benjamin, D. J., S. A. Brown, and J. M. Shapiro (2013). Who is behavioral? cognitive ability and anomalous preferences. *Journal of the European Economic Association* 11(6), 1231–1255.
- Braas-Garza, P., T. Garca-Muoz, and R. H. Gonzalez (2012). Cognitive effort in the beauty contest game. *Journal of Economic Behavior & Organization* 83(2), 254 – 260.
- Burchardi, K. B. and S. P. Penczynski (2014). Out of your mind: Eliciting individual reasoning in one shot games. *Games and Economic Behavior* 84(0), 39 – 57.
- Burnham, T. C., D. Cesarini, M. Johannesson, P. Lichtenstein, and B. Wallace (2009, Oct). Higher cognitive ability is associated with lower entries in a p-beauty contest. *Journal of Economic Behavior & Organization* 72(1), 171175.
- 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.
- Crawford, V. P., T. Kugler, Z. Neeman, and A. Pauzner (2009). Behaviorally optimal auction design: Examples and observations. *Journal of the European Economic Association* 7(2-3), 377–387.
- 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.
- Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review* 100(3), 1238–60.
- 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.
- Georganas, S., P. J. Healy, and R. A. Weber (2015). On the persistence of strategic sophistication. *Journal of Economic Theory* 159, Part A, 369 – 400.
- Gill, D. and V. Prowse (2015). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy* forthcoming.
- 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.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association* 1(1), 114–125.
- 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.
- Kawagoe, T. and H. Takizawa (2012, May). Level-k analysis of experimental centipede games. *Journal of Economic Behavior & Organization* 82(2-3), 548–566.

- 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.
- Oechssler, J., A. Roider, and P. W. Schmitz (2009). Cognitive abilities and behavioral biases. *Journal of Economic Behavior & Organization* 72(1), 147 – 152.
- 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.
- Weber, R. A. (2001). Behavior and learning in the dirty faces game. *Experimental Economics* 4(3), 229–242.

A Tables

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

Table 1: Correctly solved puzzles by iteration steps required

(0, 1)	(1, 1)	(1, 2)	(2, 2)	(2, 3)	(3, 3)	(3, 4)	index	errors
c	c	c	c	c	w	w	3	0
c	c	c	w	w	c	w	2	1
c	c	w	c	c	w	w	3	1
c	c	c	c	w	w	w	2	1

Table 2: Examples of how the index is assigned

<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

<i>Game</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>10</i>	<i>11</i>	<i>12</i>	<i>combined</i>
<i>p-value</i>	.82	.33	.80	.72	.35	.83	.03*	.69	.10	.98	.61	.59	.66

Table 4: P-values for Fisher tests of association

	All $N = 154$		1 Step $N = 81$		> 1 Steps $N = 73$	
	p_θ	e_θ	p_θ	e_θ	p_θ	e_θ
Altruist	0.070 [0.025,0.157]	0.217 [0.066,0.438]	0.066 [0.023, 0.227]	0.152 [0.000, 0.517]	0.077 [0.014, 0.169]	0.284 [0.097, 0.626]
Optimist	0.020 [0.007,0.118]	0.194 [0.090,0.499]	0.063 [0.014, 0.151]	0.270 [0.095, 0.488]		
L1	0.599 [0.496,0.705]	0.249 [0.211,0.286]	0.631 [0.479, 0.762]	0.261 [0.203, 0.315]	0.534 [0.370, 0.684]	0.232 [0.186, 0.275]
L2	0.154 [0.076,0.243]	0.249 [0.181,0.333]	0.0630 [0.011, 0.251]	0.169 [0.095, 0.705]	0.275 [0.145, 0.413]	0.244 [0.184, 0.309]
D1	0.016 [0.000,0.097]	0.111 [0.095,0.748]			0.049 [0.000, 0.132]	0.131 [0.097, 0.264]
L0	0.142 [0.007,0.201]	0.685 [0.524,0.836]	0.212 [0.017, 0.302]	0.706 [0.526, 0.866]	0.064 [0.000, 0.137]	0.622 [0.386, 0.742]
$\ln \mathcal{L}$	-1431.82		-784.77		-639.364	
AIC	2868.58		1588.81		1298.16	
BIC	2919.05		1609.09		1317.34	
ANE	0.200		0.164		0.256	

Table 5: Type Estimation

B Figures




Figure 1: Empirical distribution of the measure of logical ability

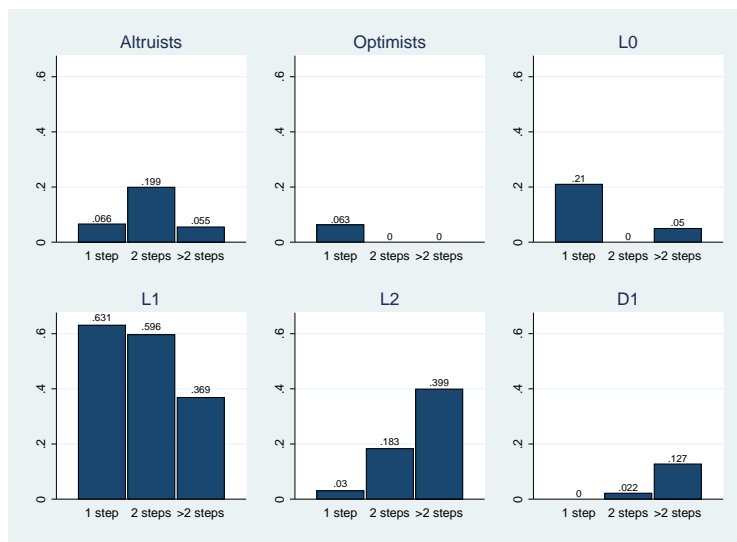


Figure 2: Logical abilities and behavioral types




Figure 3: Choices in dictator games and behavioral types: subjects with an index of two or more

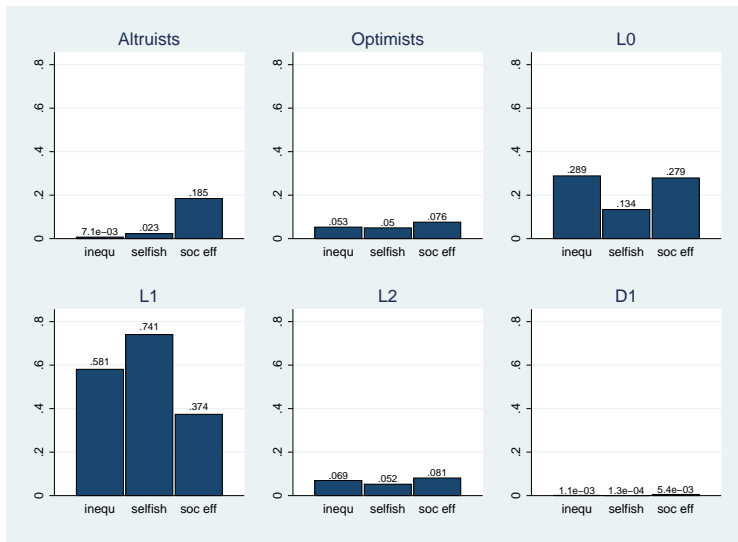


Figure 4: Choices in dictator games and behavioral types: subjects with a score of one

C 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 :(∞,∞)

D Behavioral types

Table D 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 G_n 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)
Omniscient	(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 6: Predicted behaviors in all games

E Additional estimations

	1 Step $N = 81$		2 Steps $N = 37$		> 2 Steps $N = 36$	
	p_θ	e_θ	p_θ	e_θ	p_θ	e_θ
Altruist	0.066 [0.023, 0.227]	0.152 [0.000, 0.517]	0.199 [0.027, 0.353]	0.442 [0.286, 0.581]	0.055 0.022, 0.148	0.211 [0.083, 0.834]
Optimist	0.063 [0.014, 0.151]	0.270 [0.095, 0.488]				
L1	0.631 [0.479, 0.762]	0.261 [0.203, 0.315]	0.596 [0.389, 0.808]	0.229 [0.149, 0.308]	0.337 [0.145, 0.603]	0.218 [0.173, 0.278]
L2	0.0630 [0.011, 0.251]	0.169 [0.095, 0.705]	0.183 [0.049, 0.368]	0.217 [0.106, 0.383]	0.399 [0.174, 0.618]	0.254 [0.205, 0.362]
D1			0.023 [0.000, 0.105]	0.091 [0.086, 0.821]	0.127 [0.000, 0.291]	0.178 [0.092, 0.591]
L0	0.212 [0.017, 0.302]	0.706 [0.526, 0.866]			0.050 [0.000, 0.133]	0.686 [0.168, 0.756]
$\ln \mathcal{L}$	-784.77		-330.68		-307.446	
AIC	1588.81		677.30		636.35	
BIC	1609.09		686.64		647.14	
ANE	0.164		0.177		0.311	

Table 7: Additional Type Estimations

F Sample instructions

Instructions

Thank you for your participation in this experiment. If you read these instructions carefully and act upon them, you can earn real money.

You are not allowed to communicate with other participants during the course of the experiment. If you do not follow this rule you may be excluded from the experiment.

Your task

Your task in this experiment is to determine the colour (red or white) of your hat. You will be paired with 3 computer-players. You will be able to see the colour of the hats of the computer-players, but **not** the colour of your own hat. The computer-players are in a similar situation. They observe your hat colour and the hat colours of their fellow computer-players, but not the colour of their own hat. However, everybody knows (you and the computer-players) that *at least one player has a red hat*. The picture below shows a typical situation:

THIS IS A NEW GAME!!!				
There is at least one RED hat.				
	You	Computer 1	Computer 2	Computer 3
<i>You see this</i>	??	RED	WHITE	WHITE

What can you infer?

- I have a WHITE hat with certainty
- I have a RED hat with certainty
- I can't possibly know

OK

You observe in this case that one of the computer-players has a red hat while the other two have white hats. The question marks “??” indicate that you do not know your hat colour.

You are asked to decide what you can infer from the information you are given. Possible answers are: “I have a WHITE hat with certainty”, “I have a RED hat with certainty”, and “I can’t possibly know”. One of these answers is correct, the two others are wrong. Note that answering “I can’t possibly know” is wrong whenever it is possible to correctly infer the hat colour from the information given. Similarly ticking

“I have a WHITE hat with certainty” or “I have a RED hat with certainty” is only correct if it is actually possible to logically infer that your hat colour is white or red.

The game may end after your initial decision. If the game continues, you will be given the additional information of what the computer-players have inferred from their observation. Recall that the computer-players face the same problem as you do. They can see the hats of all the others but not their own. Therefore, in the above situation, Computer 1 knows that the hats of computers’ 2 and 3 are white, and also knows your hat colour. However, it does not know its own hat colour. Consequently, the computers also have a logically correct answer to the question: what can you (Computer) infer about your hat colour? The computers ALWAYS choose the logically CORRECT answer.

THIS IS THE 2ND DECISION ROUND IN THIS GAME!!!				
There is at least one RED hat.				
	You	Computer 1	Computer 2	Computer 3
<i>You see this</i>	??	RED	WHITE	WHITE
 <i>Announcements in previous Rounds</i>				
	You	Computer 1	Computer 2	Computer 3
<i>Announcements 1</i>	I Can't know	My hat is RED	I Can't know	I Can't know
What can you infer? <input type="radio"/> I have a WHITE hat with certainty <input type="radio"/> I have a RED hat with certainty <input type="radio"/> I can't possibly know				
<input type="button" value="OK"/>				

Above you can see a possible screen for your second decision. You again have to decide what you can infer about your hat’s colour. However, now you have the additional information about what the computers (correctly) announced in the decision round before. After you have made another decision, the game may end or continue. If the game continues, you will again be given the additional information of what the computers inferred from the previous round. This process will go on until you either correctly inferred your hat colour or until you made a mistake.

Different games

After a game has ended, another new game will start. You will be given a set of 7 games.

Payment

You will start with 20 Australian Dollars. For each mistake you make, we will deduct 2.50 Dollars from your account. After the 7 games you will be paid the amount remaining in cash.

Introductory questions

Before you start the actual game we will ask you some questions about the game. These questions will be designed to test if you understand the instructions. Please make sure to read the instruction very carefully, as failing to answer the pre-game questions correctly may lead to exclusion from the experiment.

Questions

Do you have any questions? If yes please raise your hand and we will come and answer them in private.

Instructions

Welcome to our experiment. Please read these instructions carefully. Understanding the instructions is crucial for earning money.

This is an experiment in decision-making. You will be paid for your participation. The exact amount of money you will receive will be determined during the experiment and will depend on your choices. This amount will be added to your total earnings from all task of today's session. In this experiment you will play in 14 situations.

If you have any questions during the experiment, raise your hand and the experimenter will assist you. It is strictly forbidden to talk, exclaim, or to communicate with other participants during the experiment. It is very important for us that you obey these rules. Otherwise the data generated in this session are useless.

The task

You will be asked to choose an action each in nine different situations. In each situation you are randomly paired with another participant. The participant will also have to choose an action. You and the participant you are paired with have to choose the action at the same time, such that both you and your partner cannot condition the choice on the choice of each other. The payoff you receive in a situation depends on your action and on the action of the other participant, as the other participants payoff will depend on his/her and your choice. All payoffs are in Dollars. Recall that we will randomly determine one of the nine situations where you will pay your payoff in Dollars.

The actions you can choose from vary from game to game. The actions the other participant can choose from also vary from game to game. You (and the participant you are paired with in a particular round) may have one (A), two (A or B) or three actions (A, B, or C). The possible choices will be communicated to you in the top left corner of your decision screen. The participant you are paired with in a particular round may also have one, two or three choices (X, Y, Z).

Your payoffs - depending on your choice and the choice of the participant you are paired with - will be shown on the same screen below the information on the available choices.

Further below you can mark your decision and finalize it by clicking the “OK” button.


About 20 seconds after this information is presented to you a button appears. Pressing this button allows you to see the payoff of the participant you are paired with depending on your and his/her choice. The information on your payoffs will remain on the screen no matter if you choose to view the payoff of the other person or not. It is entirely up to you if you want to view the information on the other person’s payoff or not. Note that you can mark and finalize your choice already before this button even appears.

We end these instructions by presenting some screenshots (in chronological order of their possible appearance). In the actual game, the white circles will be replaced by numbers, which represent payoffs in Dollars.

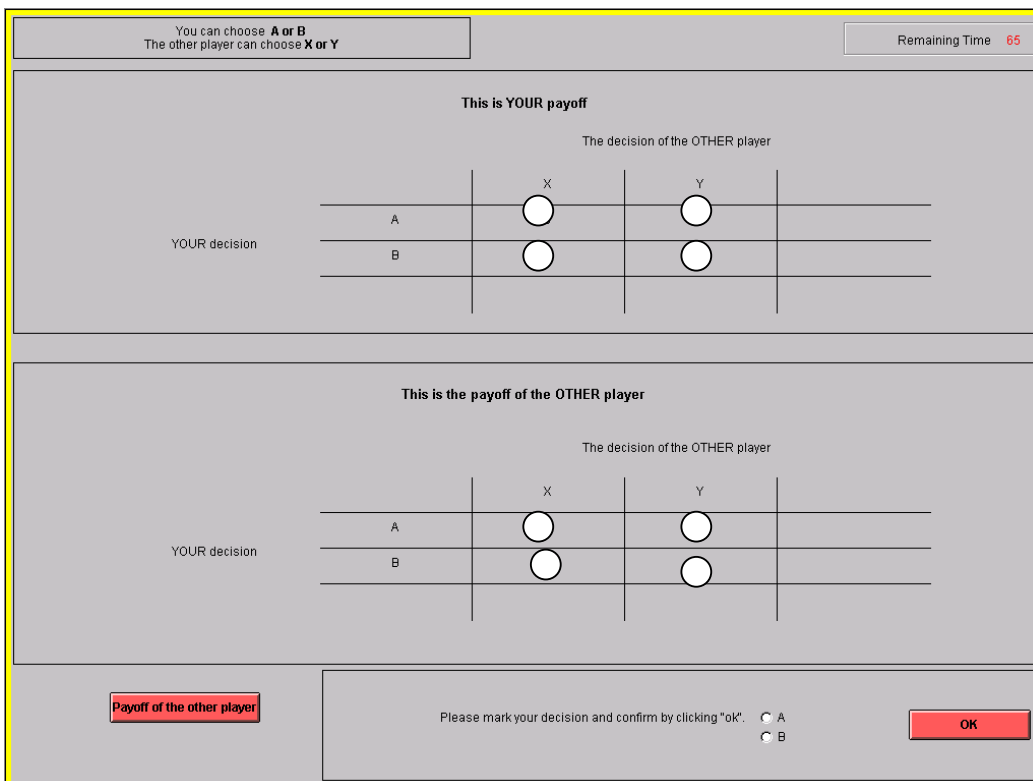
The screenshot shows a game interface with a grey background. At the top left, a box contains the text: "You can choose A or B" and "The other player can choose X or Y". At the top right, a box shows "Remaining Time 100". The main area is titled "This is YOUR payoff" and contains a payoff matrix. The matrix has "YOUR decision" on the left, with options "A" and "B". The top of the matrix is labeled "The decision of the OTHER player" with options "X" and "Y". Each cell in the matrix contains a white circle. Below the matrix, there is a box with the text: "Please mark your decision and confirm by clicking 'ok'." and two radio buttons labeled "A" and "B". To the right of this box is a red "OK" button.

		The decision of the OTHER player	
		X	Y
YOUR decision	A	○	○
	B	○	○

Screenshot1: This is what a screen initially looks like (you can already decide if you want)



Scenshot 2: After twenty seconds the button has appeared. You can click on the button see the other's payoff or not



Scenshot 3: This shows what you see if you click on the button