

Cognitive ability and learning to play equilibrium: A level- k analysis *

David Gill [†]
Oxford University
Department of Economics

Victoria Prowse [‡]
Cornell University
Department of Economics

June 27th, 2012

Abstract

In this paper we investigate how cognitive ability influences behavior, success and the evolution of play towards Nash equilibrium in repeated strategic interactions. We study behavior in a p -beauty contest experiment and find striking differences according to cognitive ability: more cognitively able subjects choose numbers closer to equilibrium, converge more frequently to equilibrium play and earn more even as behavior approaches the equilibrium prediction. To understand better how subjects with different cognitive abilities learn differently, we estimate a structural model of learning based on level- k reasoning. We find a systematic positive relationship between cognitive ability and levels; furthermore, the average level of more cognitively able subjects responds positively to the cognitive ability of their opponents, while the average level of less cognitively able subjects does not respond at all. Our results suggest that, in strategic environments, higher cognitive ability translates into better analytic reasoning and a better ‘theory of mind’.

Keywords: Cognitive ability; Bounded rationality; Learning; Convergence; Level- k ; Non-equilibrium behavior; Beauty contest; Repeated games; Structural modeling; Theory of mind; Intelligence; Raven test.

JEL Classification: C92; C73; D83.

*We thank the Economic Science Laboratory at the University of Arizona for hosting our experiment. Prowse gratefully acknowledges financial support from a Zvi Meitar-Oxford University Social Sciences Research Grant.

[†]david.gill@economics.ox.ac.uk

[‡]prowse@cornell.edu

1 Introduction

Little is known empirically about how boundedly-rational agents choose and learn in strategic environments. In this paper, we aim to discover how cognitive ability, measured using a leading test of analytic intelligence, influences behavior, success and the evolution of play towards Nash equilibrium in repeated strategic interactions. Despite well-documented differences in cognitive ability in the population, to the best of our knowledge we are the first to study how cognitive ability affects how people learn to play equilibrium. In our laboratory experiment, we find that more cognitively able subjects choose numbers closer to equilibrium, converge more frequently to equilibrium play and earn more on average even as behavior approaches the equilibrium prediction. To help gain insight into the micro-processes that drive these differences, and thus to understand better how subjects with different cognitive abilities learn differently, we estimate a structural model of learning based on level- k reasoning that fits the observed data well. The model allows subjects' levels to vary in their own cognitive ability and that of their opponents: we find a systematic positive relationship between cognitive ability and levels; furthermore, the average level of more cognitively able subjects responds positively to the cognitive ability of their opponents, while the average level of less cognitively able subjects does not respond at all. Methodologically, our structural analysis builds on existing level- k mixture-of-types models estimated using Maximum Likelihood, including Stahl and Wilson (1995), Ho et al. (1998), Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006), Crawford and Iriberry (2007a), Crawford and Iriberry (2007b) and Costa-Gomes and Weizsäcker (2008).

Cognitive ability correlates with a wide array of preferences, behavioral biases and economic outcomes (e.g., time preferences, Benjamin et al., forthcoming; risk aversion, Dohmen et al., 2010; conservatism in updating probabilities, Oechssler et al., 2009; anchoring, Bergman et al., 2010; and labor market outcomes, Heckman et al., 2006). Furthermore, a small but burgeoning literature is starting to find a link between cognitive ability and behavior in strategic games that are played only once. In the beauty contest, Burnham et al. (2009) and Brañas-Garza et al. (forthcoming) find that subjects with higher cognitive ability choose lower numbers, while Agranov et al. (2011) find that, when subjects are given time to think about their choices, higher cognitive ability subjects' choices fall more with thinking time.¹ In related dominance-solvable and guessing games, working memory (Rydval et al., 2009) and depths of reasoning in the red hat puzzle (Bayer and Renou, 2012) correlate with behavior. Cognitive ability also influences behavior in public good games (Millet and Dewitte, 2007) and in the Prisoner's Dilemma (Burks et al., 2009).

However, to the best of our knowledge, we are the first to investigate how cognitive ability influences learning in strategic environments.² We study strategic behavior in a p -beauty contest:

¹Burnham et al. (2009) have 656 participants and a single winner, and use a test of cognitive ability based on analogies, number series, and logical series. With fewer than 200 participants and groups of 24, Brañas-Garza et al. (forthcoming) find no effect of Raven test scores, but do find an effect of Cognitive Reflection Test scores. Agranov et al. (2011) find an effect of cognitive ability measured in a Bayesian updating task, but not as measured by a numeracy test or the Cognitive Reflection Test.

²Schnusenberg and Gallo (2011) run a three-round beauty contest with one winner per round and no monetary incentives (the prize was a small in-class grade improvement), and find that scores in Frederick (2005)'s three-question Cognitive Reflection Test affect choices in the first round only. In repeated centipede games with rematching, and without measuring cognitive ability, Palacios-Huerta and Volij (2009) and Ho and Su (2012) consider how different subject pools behave differently.

three subjects simultaneously choose an integer between 0 and 100 inclusive, and the subject whose choice is closest to 70% of the average of the three numbers wins \$6. In the unique Nash equilibrium, all subjects choose 0; however, the game is ideally suited to study learning since best responses to non-equilibrium choices are often above 0, but with repetition behavior tends to move towards equilibrium.³ Real-world parallels include timing games in financial and labor markets.⁴

In our experiment, we first measure cognitive ability using the 60 question non-verbal Raven test. We classify each subject as either of ‘high cognitive ability’ or of ‘low cognitive ability’ according to whether her test score is in the top or bottom half of the distribution of scores in her session. Our 510 subjects then play the p -beauty contest ten times with the same opponents and with feedback. In ‘own-matched’ groups, all three members are of the same cognitive ability type. In ‘cross-matched’ groups, the three members are of mixed ability (either two high ability and one low ability subject, or vice-versa). Subjects find out their own cognitive ability type as well as the ability type of the other two group members.⁵ This matching protocol is designed to allow us to discover how behavior varies with subjects’ own cognitive ability and with the cognitive ability of their opponents.

We find striking differences by cognitive ability. On average, high cognitive ability subjects choose lower numbers and converge more frequently to equilibrium and close-to-equilibrium play. In the final two rounds, all three group members choose the equilibrium action 37% of the time in own-matched high ability groups, but only 15% of the time in cross-matched groups and 5% of the time in own-matched low ability groups; similar results hold for measures of close-to-equilibrium play. High cognitive ability subjects are also more successful: in cross-matched groups and across all ten rounds, high cognitive ability subjects earn \$3.56 more on average than low ability subjects (by construction, high ability subjects in own-matched high ability groups earn the same on average as low cognitive ability subjects in own-matched low ability groups). Furthermore, the difference in earnings becomes bigger over rounds, even though the average choices of high and low cognitive ability subjects in cross-matched groups remain similar: high ability subjects seem to ‘learn’ better how to play the game than do low ability subjects.

To shed light on the behavioral mechanisms that underlie these differences, we estimate a structural level- k mixture-of-types model of learning. The level- k model (Stahl and Wilson, 1995; Nagel, 1995) is a powerful tool for analyzing boundedly-rational non-equilibrium behavior and reasoning in strategic interactions. Level-0 types behave in some random fashion, level-1 types best respond to the choices of level-0 types, level-2 types best respond to the choices of level-1 types, and so on. Structural level- k mixture-of-types models have been applied successfully to study behavior in, for example, guessing games (Costa-Gomes and Crawford, 2006), coordination games (Costa-Gomes et al., 2009) and auctions (Crawford and Iriberri, 2007a), and to analyze

³In games such as the p -beauty contest in which actions are strategic complements, the theoretical convergence properties of various learning processes are relatively well understood, but less is known empirically about how behavior evolves towards equilibrium (Chen and Gazzale, 2004).

⁴During a bubble or in a job market, there is an advantage to trading or making job offers a little earlier than competitors, but moving too early is costly (in terms of lost profit on the upward wave of the bubble or missing out on new information about job candidates). Roth and Xing (1994) provide evidence of slow unraveling of the timing of offers in entry-level professional job markets.

⁵As far as possible the instructions use neutral language, and so do not refer to ‘high ability’, ‘low ability’, ‘winning’, and so on. Controlling for a subject’s own test score, we find no evidence that the allocation to cognitive ability type influences behavior.

the role of beliefs in normal-form games (Costa-Gomes and Weizsäcker, 2008).⁶ Following Nagel (1995), Stahl (1996) and Duffy and Nagel (1997), we assume that level-0 types “follow the crowd” in the sense that they copy the average group behavior from the previous round,⁷ and we incorporate a form of rule learning (Stahl, 1996) by including types who switch up one level during the course of the experiment.

Our structural model of learning fits the observed data well. Simulations using the estimated parameters match closely the observed paths of average choices and earnings, and the simulated choices fit well the amount of convergence to equilibrium and close-to-equilibrium play found in the data. When estimated using only the data from the first eight rounds, the model continues to perform well out-of-sample in the final two rounds. Rule learning plays an important role in explaining subjects’ choices: for instance, when we remove rule learners from the model we can no longer fit the pattern of increasing difference in earnings over rounds between high and low cognitive ability subjects.

We find a systematic relationship between subjects’ cognitive ability and their level- k types. In particular, we summarize the estimated proportions of learner types in a single statistic measuring the average level- k choice rule that subjects follow, and find that the average level of high cognitive ability subjects is higher than that of low ability subjects. The result continues to hold when we consider only cross-matched subjects, even though crossed-matched high ability subjects face a lower number of high ability opponents on average than do cross-matched low ability subjects. We also find a difference by cognitive ability in how subjects respond to the cognitive ability of their opponents: the average level of high cognitive ability subjects responds positively to the cognitive ability of their opponents, while the average level of low cognitive ability subjects does not respond at all.⁸

The estimates of the structural model’s parameters also allow us to simulate the earnings of each level- k type, given the estimated distribution of types of their opponents. The analysis shows that subjects are constrained in their levels below those that are optimal, but that high cognitive ability subjects are closer to the optimum: on average, own-matched high ability subjects leave \$2.64 on the table compared to the payoff-maximizing type, own-matched low ability subjects leave \$4.10 on the table, cross-matched high ability subjects leave \$2.37 on the table, and cross-matched low ability subjects leave \$5.72 on the table.

Some games require only analytic reasoning: for example, Dufwenberg et al. (2010) and Gneezy et al. (2010) study how players in Race games learn their less-than-obvious dominant strategy. In contrast, the beauty contest requires analytic reasoning, in order to deduce how best to respond to beliefs about how others will choose, as well as the ability to judge well the thinking of others, in order to predict accurately how others will in fact behave. According to Coricelli and Nagel (2009), ‘theory of mind’ is “the ability to think about others’ thoughts and mental states to predict their intentions and actions”, and playing the beauty contest against

⁶The closely-related cognitive hierarchy model (Camerer et al., 2004) has been used to study behavior in, e.g., zero-sum betting games (Brocas et al., 2011), telecoms markets (Goldfarb and Xiao, 2011) and the Lowest Unique Positive Integer game used by the Swedish national lottery (Östling et al., 2011). See Crawford et al. (forthcoming) for a comprehensive survey of applications of level- k and cognitive hierarchy models.

⁷Thus, the group-specific history influences the behavior of every type via the impact on the level-0 type.

⁸Agranov et al. (2012) also find that levels are endogenous: in a one-shot beauty contest, inexperienced undergraduates shift to higher level- k types on average when they play against graduates with some experience of the game.

humans rather than a computer “activated areas commonly associated with theory of mind or mentalizing-thinking about other people’s minds.” Ohtsubo and Rapoport (2006) argue that “one of the most important uses of the theory-of-mind ability is the strategic reasoning used to outwit or manipulate others.” Our results suggest that high scores on a purely analytic test of intelligence translate in a strategic environment into better analytic reasoning and a better theory of mind. High cognitive ability subjects are not only more successful, but are also better able to predict how cognitive ability affects how their opponents behave: high ability subjects adjust to the cognitive ability of their opponents, while low ability subjects do not on average. The earnings simulations reported above show that, by not adjusting their level upward, the low ability subjects lose out more when the cognitive ability of their opponents goes up: high ability subjects leave \$0.27 more on the table, while low ability subjects leave \$1.62 more on the table.

Our finding that people of higher cognitive ability perform better and learn faster in strategic interactions is important for understanding how boundedly-rational people play games in the real world and for interpreting observed heterogeneity in learning processes, but also raises potentially far-reaching practical and ethical questions. For instance: How much protection should public policy afford to the less cognitively able when they operate in markets, especially newer markets in which some participants have price-setting power? How can the design of institutions and mechanisms take into account the impact of bounded rationality on how agents learn to behave in the strategic environment implied by the rules of the institution or mechanism? Is redistribution appropriate to correct for differences in outcomes when people of different cognitive abilities interact repeatedly? Our results are also relevant when interpreting close-to-equilibrium behavior: even if average behavior mimics equilibrium play quite closely after some period of learning, low cognitive ability agents might nonetheless be earning substantially less than their high cognitive ability counterparts, with potentially significant implications for fairness and efficiency.

The paper proceeds as follows: Section 2 describes the experimental design; Section 3 reports descriptive statistics and reduced form regression results; Section 4 presents the structural analysis; and Section 5 concludes.

2 Experimental design

We ran 22 experimental sessions at the University of Arizona’s Experimental Science Laboratory (ESL), all conducted on weekdays between November 2010 and March 2011 and lasting approximately 75 minutes. 18 or 24 student subjects participated in each session, with 510 participants in total.⁹ The participants were drawn from the ESL subject pool, which is managed using a bespoke online recruitment system. We excluded any graduate students in economics. Seating positions were randomized. The experimental instructions (Appendix C) were provided to each subject on their computer screen and were read aloud to the subjects. Questions were answered privately. Each subject was paid a show-up fee of \$5.00 and earned an average of a further \$20.00 during the experiment (all payments were in U.S. dollars). Subjects were paid privately

⁹All sessions were run during the Fall or Spring Semesters. We aimed to have 24 subjects per session, but ran sessions of 18 when fewer than 24 showed up. Before running these sessions, we also ran one pilot session without any monetary incentives and two sessions with a different form of the p -beauty contest (where the target was 90% of the mean of the choices) whose results are not reported here.

in cash. The experiment was programmed in z-Tree (Fischbacher, 2007).

2.1 The test of cognitive ability

Each session consisted of a test of cognitive ability followed by repeated play of the p -beauty contest (Nagel, 1995). In more detail, each session started with a 30 minute computerized test of cognitive ability using Raven’s Progressive Matrices. The Raven test, which consists of non-verbal multiple choice questions, is recognized as a leading measure of analytic intelligence (Carpenter et al., 1990; Gray and Thompson, 2004, Box 1, p. 472).¹⁰ In economics, Raven test scores have been found to correlate positively with fewer Bayesian updating errors (Charness et al., 2011) and with more accurate beliefs (Burks et al., 2009). Each question asks the subjects to identify the missing element that completes a visual pattern. We used the Standard Progressive Matrices Plus version of the Raven test (the level of difficulty lies between that of the Standard Progressive Matrices and the Advanced Progressive Matrices), which consists of 60 questions split into 5 parts of increasing difficulty, labeled A-E, with 12 questions in each. We gave the subjects 3 minutes for each of parts A and B (which are easier than parts C-E) and 8 minutes for each of parts C, D and E. Within each part subjects could move back and forth between the 12 questions in that part and, time permitting, they were allowed to change their previous answers.


We did not provide any monetary incentives for completing the Raven test. This is conventional in the psychology and psychometric literatures and avoids the possibility that income effects might spill over from the test to behavior in the p -beauty contest. We did, however, tell the subjects that we would inform them privately of their own score at very end of the session. Figure 1(a) shows how the Raven test scores of our subjects were distributed. The mean test score was 40.7, with individual scores ranging from 12 to 58.

2.2 Subject matching

Following the Raven test, each subject was determined to be either (i) of ‘high cognitive ability’ if her test score was in the top half of the distribution of scores of the subjects in her session or (ii) of ‘low cognitive ability’ if her test score was in the bottom half of the distribution in her session. Subjects scoring exactly the session median were allocated to ability types so as to ensure an equal number of subjects of each ability type for that session. Figure 1(b) illustrates the densities of the Raven test scores by cognitive ability type. On average, the Raven test score of high ability subjects was 7.0 higher than that of low ability subjects. The region of overlap arises because the median Raven test score varied over sessions, from 37.0 to 42.5. We found no time trend in the median test score over sessions (2-sided $p = 0.557$).

Subjects were then put into groups of 3. In ‘own-matched’ sessions, all 3 members of a group were of the same ability type. 180 subjects participated in own-matched sessions, giving 30

¹⁰Carpenter et al. (1990) define analytic intelligence (also sometimes called fluid intelligence, as opposed to crystallized intelligence) as “the ability to reason and solve problems involving new information, without relying extensively on an explicit base of declarative knowledge derived from either schooling or previous experience”, and show that Raven test scores discriminate according to the ability to use abstract reasoning and correlate highly with scores on other complex cognitive tasks. Raven et al. (2000, SPM25-SPM37) survey the extensive literature that studies: (i) correlations between Raven test scores and scholastic aptitude and achievement tests and other measures of cognitive ability; and (ii) the internal consistency and test-retest reliability of the Raven test.



(a) Histogram of Raven test scores.

(b) Smoothed densities of Raven test scores.

Figure 1: Histogram and densities of Raven test scores

groups of 3 high cognitive ability subjects (‘own-matched high ability groups’) and 30 groups of 3 low cognitive ability subjects (‘own-matched low ability groups’). In ‘cross-matched’ sessions, the 3 members of a group were of mixed ability (half the groups were made up of 2 high cognitive ability subjects matched with 1 low cognitive ability subject and the other half were made up of 2 low cognitive ability subjects matched with 1 high cognitive ability subject). 330 subjects participated in cross-matched sessions, giving 110 ‘cross-matched groups’. Conditional on a subject’s cognitive ability type, the allocation to groups was random. The subject matching implies that, on average, high ability subjects face $4/3$ more high ability opponents in own-matched groups than in cross-matched groups, while low ability subjects face $4/3$ more high ability opponents in cross-matched groups than in own-matched groups.

We informed each subject of her own cognitive ability type as well as the cognitive ability type of the 2 other members of her group. The instructions (Appendix C) did not use the terms ‘high ability’ or ‘low ability’; instead, we referred more neutrally to the top and bottom half of the test scores of all participants in the room. Controlling for a subject’s own test score, we find no evidence that the allocation to cognitive ability type *per se* influenced behavior or earnings in the experiment: Appendix A provides the details of this analysis.

2.3 p -beauty contest game

Each group of 3 then played 10 rounds of the p -beauty contest (Nagel, 1995) with $p = 0.7$ and without rematching (no rematching allows us to treat behavior across groups as independent). In particular, in every round each group member privately chose an integer between 0 and 100 inclusive (the subjects typed their chosen number into a box rather than selecting it from an on-screen grid). The group member whose chosen number was closest to 70% of the mean of the 3 numbers chosen by the group members (the ‘target’) was paid \$6.00 and the other group members received nothing. In the case of ties, the \$6.00 was split equally among the winners. To keep the language as neutral as possible, the instructions (Appendix C) did not use terms such as ‘prize’, ‘winner’, ‘loser’, ‘ties’ or ‘target’. The unique Nash equilibrium is for all players

to choose 0.¹¹

Before the start of the first round, we described the number of rounds, the rules of the game and the information feedback the subjects would receive at the end of each round. At the end of every round, each group member was informed of: (i) the numbers chosen by the group members; (ii) the mean of the 3 chosen numbers; (iii) 70% of the mean (the target); (iv) which group member(s)' number(s) was (were) closest to the target; and (v) how much each group member was paid for the round.¹² While deciding on their choice of number and also when receiving feedback, the subjects could see a reminder of the rules and of the cognitive ability type of each member of their group. All interactions were anonymous, but subjects were given labels (X, Y or Z) which were held fixed for the 10 rounds; hence each subject could link the choices in their group to particular opponents whose cognitive ability type they knew.

The subjects had 90 seconds to make their choice in each round. The subjects were told that if they made their choice early, they would still have to wait for the full 90 seconds. If a subject failed to make a choice within 90 seconds, a flashing request prompted an immediate choice. At the end of each round, the subjects could see the feedback information described above for a period of 30 seconds before the next round began.

3 Reduced form results

In this section, we report descriptive statistics and reduced form regression results in order to describe how average behavior, the evolution of play over rounds and the group-by-group dynamics of convergence towards equilibrium vary with cognitive ability. In Section 4, we go on to build and estimate a structural model of learning that aims to explain the main features of the patterns that we describe below.

¹¹The game is discrete, and hence this is not true for all values of p (López, 2001). However, it is relatively straightforward to show for our $p = 0.7$. A proof starts from the observation that the highest chosen number can never win or tie unless all 3 players choose that number. Suppose that $x_i \geq x_j$ and $x_i > x_h$. Let t be the target. If $x_h \geq t$, clearly i cannot win or tie. If $x_h < t$, x_h is closer to the target than x_i if and only if $(x_i - t) - (t - x_h) > 0$. When $p = 0.7$, this difference equals $\frac{1}{15}(8x_i - 7x_j + 8x_h) > 0$. Now suppose a Nash equilibrium exists in which the players do not always all choose zero. Let x^{max} be the highest number that is ever played in equilibrium. x^{max} cannot be a best response. From the observation, the payoff from $x_i = x^{max}$ is strictly positive if and only if $x_j = x_h = x^{max}$ (giving a tie). If that can happen with strictly positive probability, deviating to any lower number is profitable (from the observation giving a win when $x_j = x_h = x^{max}$). If not, deviating to the lowest number ever chosen by your opponents is profitable (from the observation giving a tie when at least one opponent chooses that number).

¹²In terms of information feedback, our design is close to that of Nagel (1995), where the whole distribution of choices was revealed at the end of each round (but with much larger groups and just 4 repetitions).

3.1 Behavior and earnings in the first round

We start by reporting briefly how behavior and earnings vary with cognitive ability in the first round. The mean choice of high cognitive ability subjects is 42.6 while that of low cognitive ability subjects is 45.5. The difference of 2.9 is not quite statistically significant at conventional levels (regressing p -beauty contest choices on cognitive ability type gives a 2-sided $p = 0.106$). To study differences in earnings by cognitive ability we only look at subjects in cross-matched groups: low cognitive ability subjects in own-matched low ability groups must by construction earn as much on average as high cognitive ability subjects in own-matched high ability groups. In the first round, cross-matched high ability subjects earn \$0.11 more than cross-matched low ability subjects (recall that mean earnings per round are \$2.00), but the difference is not statistically significant (2-sided $p = 0.723$). In conclusion, we find weak evidence that high cognitive ability subjects choose lower numbers in the first round, but find no statistically significant evidence that these lower choices translate into greater earnings.

3.2 Behavior across all 10 rounds

We now consider how behavior varies with cognitive ability across all rounds of the experiment, describing both average behavior across all 10 rounds and the evolution of behavior during the course of the experiment. We first study the behavior of all subjects, and then focus on own-matched subjects and cross-matched subjects separately. We study differences in earnings only for cross-matched subjects: as noted in Section 3.1, low cognitive ability subjects in own-matched low ability groups must by construction earn as much on average as high cognitive ability subjects in own-matched high ability groups (i.e., \$2.00 on average in every round).

3.2.1 Behavior of all subjects

Across all 10 rounds of the p -beauty contest, the mean choice of high cognitive ability subjects is 18.1 while that of low cognitive ability subjects is 20.6. Regressing p -beauty contest choices on cognitive ability type, we find that the difference of 2.5 is highly statistically significant (2-sided $p = 0.004$).¹³ Figure 2 shows the round-by-round evolution of mean choices for high and low cognitive ability subjects. We can see that average behavior moves towards the Nash equilibrium for both types. Despite the similarity in average behavior towards the end of the 10 rounds, we find that the high ability subjects earn significantly more than low ability subjects (Section 3.2.3) and converge to a greater degree to Nash equilibrium (Section 3.3.2).

¹³Except in Section 3.1, all regressions cluster by group to allow for within-group non-independence after the first round. The subject matching implies that high cognitive ability subjects face a slightly higher number of high cognitive ability opponents on average than do low ability subjects: in Sections 3.1 and 3.2, the statistical significances of our regression results are robust to controlling for the cognitive ability of opponents by including a variable measuring the proportion (0, 0.5 or 1) of high ability opponents.




Figure 2: Round-by-round means of p -beauty contest choices (all subjects).

3.2.2 Behavior of own-matched subjects

The difference in average behavior between high and low cognitive ability subjects is more pronounced for own-matched subjects than for the sample as a whole. Across all 10 rounds, the mean choice of own-matched high cognitive ability subjects is 6.3 lower than that of own-matched low ability subjects (2-sided $p = 0.000$). Figure 3 shows round-by-round mean choices for own-matched subjects only. In rounds 1-5, choices of own-matched high cognitive ability subjects are on average 8.2 lower than choices of own-matched low ability subjects (2-sided $p = 0.000$). In rounds 6-10, own-matched high cognitive ability subjects' choices remain lower, by an average of 4.4 (2-sided $p = 0.034$). By the final round, the mean choice of own-matched high cognitive ability subjects falls to 4.5, while that of own-matched low ability subjects falls to 7.4.




Figure 3: Round-by-round means of p -beauty contest choices of own-matched subjects.


3.2.3 Behavior and earnings of cross-matched subjects

Across all 10 rounds, the mean choice of cross-matched high cognitive ability subjects is only 0.4 lower than the mean choice of cross-matched low ability subjects, and the difference is not statistically significant (2-sided $p = 0.653$). Figure 4(a) shows round-by-round mean choices for cross-matched subjects only. Cross-matched high and low cognitive ability subjects behave similarly on average in both the first and second halves of the experiment: there is no statistically significant difference in behavior in either rounds 1-5 ($p = 0.181$) or in rounds 6-10 ($p = 0.330$).

Despite the similarity in the average behavior of cross-matched high and low cognitive ability subjects, we find that cross-matched high cognitive ability subjects are more successful in the sense that they earn significantly more money in the experiment. Over the 10 rounds cross-matched high ability subjects earn \$3.56 more than cross-matched low ability subjects, and the difference is highly statistically significant (2-sided $p = 0.007$).

Figure 4(b) shows how the earnings of cross-matched subjects evolve during the course of the experiment. We can see that (i) cross-matched high ability subjects earn more than cross-matched low ability subjects in both the first and second halves of the experiment; and that (ii) the earnings difference becomes bigger in the second half. In rounds 1-5, cross-matched high cognitive ability subjects earn \$0.27 more per round than cross-matched low ability subjects, although the difference is not quite statistically significant (2-sided $p = 0.143$). In rounds 6-10, cross-matched high cognitive ability subjects earn \$0.44 more per round than cross-matched low ability subjects, and the difference is highly statistically significant (2-sided $p = 0.003$). By the final round, the difference in earnings rises to \$0.79.

Thus earnings diverge in the second half of the experiment, even though average behavior remains similar. Somehow, cross-matched high cognitive ability subjects ‘learn’ better how to play the game than do cross-matched low cognitive ability subjects. An important aim of the structural analysis described in Section 4 is to explain the mechanism driving these results.



(a) Round-by-round means of choices.

(b) Round-by-round means of earnings.

Figure 4: p -beauty contest choices and earnings of cross-matched subjects.

3.3 Group-by-group convergence towards equilibrium


Section 3.2 provides a good overview of how average behavior evolves over time. However, this average overview masks a significant amount of group-by-group variation in exactly how play evolves towards equilibrium. In order to better understand how the dynamics of the learning process vary with cognitive ability, in this section we look in more detail at the process of convergence towards Nash equilibrium. We start by providing a visual description of the way in which some individual groups succeed or fail to converge towards equilibrium; we then present results about the proportion of groups that converge and how the degree of convergence varies with group composition.

3.3.1 Visual description of group-by-group behavior


Figures 5(a)-5(d) and 6(a)-6(d) give an overview of the group-by-group variation in the dynamics of convergence. Each of the 8 figures shows, for a specific group, how the choices of the 3 group members change round-by-round. The figures also show the winning choice or choices in each round.

Figures 5(a)-5(c) show examples of 3 groups which slowly converge towards equilibrium. No group ever reaches equilibrium in Figures 5(a)-5(c). However, some groups do successfully converge all the way to equilibrium. Figure 5(d) shows an own-matched high ability group in which convergence to equilibrium is almost immediate: by the fourth round all 3 group members choose 0, and all 3 then stick to the equilibrium choice for the remainder of the experiment. Figure 6(a) shows a cross-matched group in which behavior also converges to equilibrium, although not as fast. As we will see shortly in Section 3.3.2, convergence is much more common when all 3 members of the group are of high cognitive ability.


Convergence towards equilibrium is not the only pattern that we observe. Figures 6(b) and 6(c) show groups in which behavior does not move discernibly towards equilibrium. In both cases, high choices by some of the subjects seem to disturb the learning process. However, Figure 6(d) shows an own-matched high ability group in which such a high choice seems hardly to affect the learning process at all.




(a) Own-matched high ability group.




(b) Own-matched low ability group.



(c) Cross-matched group with 1 high ability subject.




(d) Own-matched high ability group.




Notes: Ties are represented by overlapping markers. In cross-matched groups, the subject represented by the solid line is of the minority cognitive ability type.

Figure 5: Examples of group-by-group behavior: groups 1-4 of 8.




(a) Cross-matched group with 1 high ability subject. (b) Cross-matched group with 1 low ability subject.



(c) Own-matched low ability group.

(d) Own-matched high ability group.



Notes: Ties are represented by overlapping markers. In cross-matched groups, the subject represented by the solid line is of the minority cognitive ability type.

Figure 6: Examples of group-by-group behavior: groups 5-8 of 8.

3.3.2 Statistics on group-by-group convergence

Of course, the figures in Section 3.3.1 illustrate the behavior of just a small subset of the 170 groups in our sample. We now study the degree of convergence more systematically, differentiating between own-matched high ability groups, cross-matched groups and own-matched low ability groups.

Table 1 shows the frequency of equilibrium and close-to-equilibrium play in the final 2 rounds of the experiment. The first column shows the proportion of equilibrium play, that is the proportion of group-round observations in which all 3 group members choose 0. Own-matched high ability groups play the equilibrium around 37% of the time in the final 2 rounds, which is statistically significantly more often than for cross-matched groups (15%) and own-matched low ability groups (5%). The second and third columns show that the results extend when we consider close-to-equilibrium play, defined to be the proportion of group-round observations in the final 2 rounds in which the mean choice of the 3 group members is smaller than or equal to 1 (second column) or 2 (third column). Own-matched high ability groups are close to equilibrium statistically significantly more frequently than cross-matched groups, which in turn are close to equilibrium more often than own-matched low ability groups.

	Equilibrium	Group mean ≤ 1	Group mean ≤ 2
<i>Observed proportions:</i>			
Own-matched high ability groups	0.367*** [0.000]	0.500*** [0.000]	0.583*** [0.000]
Cross-matched groups	0.145*** [0.000]	0.282*** [0.000]	0.368*** [0.000]
Own-matched low ability groups	0.050* [0.092]	0.133*** [0.009]	0.217*** [0.001]
<i>Cross-group differences in proportions:</i>			
Own-matched high ability groups vs. own-matched low ability groups	0.317*** [0.001]	0.367*** [0.001]	0.367*** [0.001]
Own-matched high ability groups vs. cross-matched groups	0.221** [0.017]	0.218** [0.027]	0.215** [0.027]
Cross-matched groups vs. own-matched low ability groups	0.095* [0.052]	0.148** [0.027]	0.152** [0.045]

Notes: The first column reports the proportion of group-round observations in rounds 9 and 10 in which all 3 group members choose 0. The second (third) column reports the proportion of group-round observations in rounds 9 and 10 in which the mean choice of the 3 group members is smaller than or equal to 1 (2). p values are shown in brackets and were constructed allowing clustering at the group level. The p values are 1-sided for the null that proportions are equal to 0 (top half of the table) and 2-sided for the null that proportions do not vary between particular group types (bottom half of the table). *, ** and *** denote significance at the 10%, 5% and 1% levels.

Table 1: Proportions of equilibrium and close-to-equilibrium play in rounds 9 and 10.

In Section 4.3.1 we report how well our structural level- k model of learning can match these differences in convergence according to group composition. In particular, we compare the observed proportions of equilibrium and close-to-equilibrium play to those arising from simulated choices constructed using the estimated parameters of the structural model, and we find that the structural model does well in explaining the broad patterns that we find here.

4 Structural analysis

Section 3 presented a detailed reduced form description of how behavior and earnings vary with cognitive ability. In order to shed light on the behavioral *mechanisms* that underlie these differences we turn to a level- k model of learning. Below we describe our empirical model's specification, our estimation strategy and our results.

4.1 Level- k model of learning

We estimate a structural level- k mixture-of-types model of learning using Maximum Likelihood. Choices in the first round serve as the initial conditions. Our level- k model of learning includes nine learner types; and since our model includes rule learners, we distinguish level- k *types* from level- k *choice rules* (described below). Five standard level- k types with $k \in \{0, 1, 2, 3, 4\}$ follow the level- k choice rule in all rounds $r \geq 2$. We also include four rule learner types who switch from following the level- k choice rule in round $r = 2$, with $k \in \{0, 1, 2, 3\}$, to following the level- $(k + 1)$ choice rule in round $r = 10$. For such $Lk - (k + 1)$ rule learners, the probability of choosing the level- $(k + 1)$ choice rule increases linearly over rounds from 0 in round $r = 2$ to 1 in round $r = 10$ (and is thus given by $\frac{r-2}{8}$ in round r), while the complementary probability of choosing the level- k choice rule falls linearly over rounds.¹⁴

Some further notation is necessary to describe the level- k choice rules. Subject $i \in \{1, 2, 3\}$ in group $g \in \{1, 2, \dots, G\}$ is denoted by $ig \in \{1g, 2g, 3g\}$. The choice of subject ig in round $r \in \{1, 2, \dots, 10\}$ is denoted by $x_{ig,r} \in \{0, 1, \dots, 100\}$. The set of choices in round r of the 3 subjects in group g is denoted by $\mathbf{x}_{g,r} \equiv \{x_{1g,r}, x_{2g,r}, x_{3g,r}\}$, with mean choice $\bar{x}_{g,r} \equiv \left(\sum_{i=1}^3 x_{ig,r}\right)/3$.

To model the noise in the choice process, we assume that subjects' choices are independent draws (over rounds and subjects) from discretized and truncated t -distributions.¹⁵ Letting $f(\cdot; \mu, \sigma, \nu)$ be the density of the three-parameter t -distribution with mean μ , scale σ and degrees of freedom ν , the probability of a particular choice x by subject ig when following the level- k choice rule in round $r \geq 2$ given the group-specific choices in the previous round $\mathbf{x}_{g,r-1}$ is

$$\Pr(x|k, \mathbf{x}_{g,r-1}) = (1 - \gamma(r)) \frac{f(x; \mu(k, \mathbf{x}_{g,r-1}), \sigma(\mu, r), \nu)}{\sum_{x=0}^{99} f(x; \mu(k, \mathbf{x}_{g,r-1}), \sigma(\mu, r), \nu)} \mathbf{1}_{x \in \{0, 1, \dots, 99\}} + \gamma(r) \mathbf{1}_{x=100} \quad (1)$$

¹⁴Although the probability of choosing the level- $(k + 1)$ choice rule goes up over time, a $Lk - (k + 1)$ rule learner is allowed to switch back and forth between the level- k and level- $(k + 1)$ choice rules. An alternative specification (AS1) in which rule learners make a once-and-for-all transition to the level- $(k + 1)$ choice rule, and so cannot switch back, fits the data significantly less well (see Table 7 in Appendix B). Note also that in our mixture model framework, we do not model explicitly how rules are chosen; in contrast, Stahl (1996) uses an attraction framework in which a given rule is more attractive the higher its past payoff.

¹⁵Stahl (1996) and Offerman et al. (2002) use truncated normal distributions to model noise in the choice process in structural models of learning. In Table 7 in Appendix B, we show that an alternative specification (AS2) in which the normal distribution replaces the t -distribution fits the data significantly less well.

where $\mu(k, \mathbf{x}_{g,r-1}) = \left(\frac{7}{10}\right)^k \bar{x}_{g,r-1}$, rounded to the nearest integer, and $\gamma(r)$ is the probability of choosing 100 in round r given by the empirical frequency of subjects choosing 100 that we observe in that round in our sample (independently of cognitive ability).¹⁶ Therefore, as in Nagel (1995), Stahl (1996) and Duffy and Nagel (1997), a subject following the level-0 choice rule “follows the crowd” in the sense that she aims to copy the average group behavior from the previous round, while a subject who follows the level- k choice rule for $k > 0$ best responds to level- $(k-1)$ choices in the sense that she aims to hit the current round’s target in her group $t_{g,r} \equiv \frac{7}{10}\bar{x}_{g,r}$ that would result from everybody in the group noiselessly following the level- $(k-1)$ choice rule.¹⁷

For all level- k choice rules the scale parameter σ depends on μ and on the round r in the following way:

$$\sigma(\mu, r) = \exp\left(\alpha + [\mathbf{1}_{\mu=0}, \mathbf{1}_{\mu=1}, \mathbf{1}_{\mu \in \{2,3,4,5\}}] \boldsymbol{\beta} + \delta \left(\frac{r-2}{8}\right)\right). \quad (2)$$

Thus we allow the variance of the discretized choice distribution to vary with μ in a flexible way and to include a round trend.¹⁸ The number of degrees of freedom ν is common to all level- k choice rules.

4.2 Estimation strategy

We estimate probabilities of being the different learner types for own-matched high cognitive ability subjects, cross-matched high ability subjects, cross-matched low ability subjects and own-matched low ability subjects. Thus the parameter vector $\boldsymbol{\theta}$ to be estimated is made up of 38 elements: 32 parameters that measure the probabilities of the different learner types (given 9 learner types) and the 6 parameters of the t -distribution (α , the 3-element vector $\boldsymbol{\beta}$, δ and ν). Letting $k_{ig,r}$ be the level- k choice rule that subject ig follows in round r , the set of level- k choice rules followed by the 3 subjects in group g in round r is denoted by $\mathbf{k}_{g,r} \equiv \{k_{1g,r}, k_{2g,r}, k_{3g,r}\}$, and the set of group g ’s choice rules for every round $r \geq 2$ is denoted by $\mathbf{k}_g \equiv \{\mathbf{k}_{g,2}, \mathbf{k}_{g,3}, \dots, \mathbf{k}_{g,10}\}$. The 32 parameters that measure the probabilities of the different learner types determine the probability of each different combination of choice rules, so $\Pr(\mathbf{k}_g)$ depends on $\boldsymbol{\theta}$. Finally, letting $\mathbf{x}_g \equiv \{\mathbf{x}_{g,2}, \mathbf{x}_{g,3}, \dots, \mathbf{x}_{g,10}\}$ be the set of group g ’s choices for every round $r \geq 2$,

$$\Pr(\mathbf{x}_g | \mathbf{k}_g, \mathbf{x}_{g,1}) = \prod_{r=2}^{10} \Pr(\mathbf{x}_{g,r} | \mathbf{k}_{g,r}, \mathbf{x}_{g,r-1}) = \prod_{r=2}^{10} \prod_{i=1}^3 \Pr(x_{ig,r} | k_{ig,r}, \mathbf{x}_{g,r-1}), \quad (3)$$

¹⁶We observe 67 instances of subjects choosing 100 in our sample, making up 1.3% of observations. Although such extreme choices are a standard feature in beauty contest datasets, their attraction cannot be explained readily by level- k choice rules. Ho et al. (1998) argue that such choices “are probably due to frustration or to misguided attempts to win by singlehandedly raising the mean dramatically.” In Table 7 in Appendix B, we show that an alternative specification (AS3) in which the choices of 100 come from the same t -distribution as for the other choices fits the data significantly less well.

¹⁷If subjects took into account the effect of their own choice on the target, they would choose lower fractions of $\bar{x}_{g,r-1}$; however, in Table 7 in Appendix B, we show that an alternative specification (AS4) in which $\mu(k, \mathbf{x}_{g,r-1})$ is given by this lower fraction fits the data significantly less well.

¹⁸Stahl (1996) also allows the variance of the choice distribution to depend on μ ; Ho et al. (1998) allow a time trend in variances. We need the variance to depend on μ in order to fit the degree of convergence towards equilibrium play that we observe in the data. The time trend is a function of $r-2$ since first round choices serve as the initial conditions and the second round acts as the omitted category.

and the likelihood for group g

$$L_g(\boldsymbol{\theta}|\mathbf{x}_g, \mathbf{x}_{g,1}) = \Pr(\mathbf{x}_g|\mathbf{x}_{g,1}) = \sum_{\mathbf{k}_g} \Pr(\mathbf{k}_g) \Pr(\mathbf{x}_g|\mathbf{k}_g, \mathbf{x}_{g,1}). \quad (4)$$

The sample likelihood is then the product over the G groups of the group likelihoods.

We maximized the sample log likelihood function using a Hessian-based optimization routine. Following Berndt et al. (1974), the Hessian employed in the optimization process was approximated as the sum of outer products of the gradients of the group log likelihoods. The gradients of the group log likelihoods, in turn, were obtained via numerical differentiation. Standard errors were obtained from a Hessian matrix computed using numerical differentiation. We found the optimization problem to be well-behaved. In particular, the optimization routine converged to the same parameter vector for multiple sets of starting values, and the Hessian matrix used to obtain standard errors was never found to be close to singular.¹⁹

4.3 Results

In Section 4.3.1, we start by reporting the estimated proportions of learner types arising from our level- k model of learning, and we provide evidence that our preferred specification fits the observed data well and that rule learning plays an important role in explaining our data. In Section 4.3.2, we then show how the subjects' average level- k choice rule varies with their own cognitive ability and with that of their opponents. Finally, in Section 4.3.3 we simulate the earnings that accrue to different learner types in order to discover which learner types earn the most.

4.3.1 Estimated learner types and model goodness of fit

Table 2 reports the estimated proportions of learner types.²⁰ In Section 4.3.2 below we summarize this information in terms of average level- k choice rules and analyze how the averages vary with own cognitive ability and that of opponents. Instead, our focus here is on the implications of these estimates for behavior in order to see how well our model fits the observed data.

¹⁹According to the model outlined in Section 4.1, choice probabilities conditional on the group playing Nash equilibrium in the previous round are identical across learner types. However, there are enough group-round observations in which Nash equilibrium was not played to identify all the learner type probabilities.

²⁰The estimated parameters of the t -distribution, which complete the estimate of the parameter vector $\boldsymbol{\theta}$, are: $\alpha = -2.691$ (0.057); $\boldsymbol{\beta} = [-2.942$ (0.131), -1.784 (0.098), -1.205 (0.081)]; $\delta = -0.807$ (0.105); and $\nu = 0.777$ (0.044). Standard errors are in parentheses.

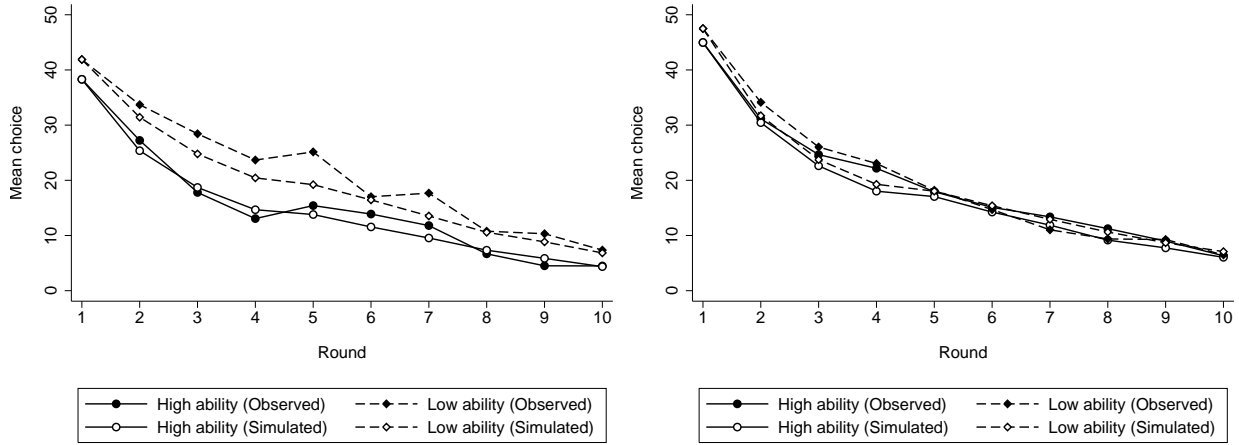
	Own-matched subjects		Cross-matched subjects	
	High ability	Low ability	High ability	Low ability
L0 Type	0.006 (0.016)	0.008 (0.017)	0.006 (0.009)	0.020 (0.019)
L1 Type	0.295*** (0.075)	0.212*** (0.082)	0.168*** (0.053)	0.408*** (0.067)
L2 Type	0.000 —	0.000 —	0.176*** (0.068)	0.280*** (0.069)
L3 Type	0.001 (0.018)	0.000 —	0.066 (0.050)	0.000 —
L4 Type	0.000 —	0.000 —	0.000 —	0.019 (0.024)
L0-1 Rule learner type	0.000 —	0.144*** (0.055)	0.046 (0.029)	0.080* (0.045)
L1-2 Rule learner type	0.243*** (0.085)	0.383*** (0.091)	0.334*** (0.071)	0.067 (0.065)
L2-3 Rule learner type	0.179** (0.081)	0.189*** (0.061)	0.184*** (0.070)	0.084 (0.057)
L3-4 Rule learner type	0.277*** (0.066)	0.064* (0.033)	0.019 (0.031)	0.042 (0.038)
Proportion of rule learners	0.698*** (0.074)	0.780*** (0.084)	0.583*** (0.098)	0.273*** (0.093)

Notes: Standard errors are shown in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% levels.

Table 2: Estimated learner type probabilities.

To give a visual impression of how well the model fits the observed data, we simulate choices over the 10 rounds using the estimated parameters. Figure 7 shows the simulated and observed paths of average behavior for own and cross-matched subjects: we see that the observed and simulated paths match closely (the notes to the figure explain how the simulated paths were constructed).

Figure 8(a) shows the simulated path of earnings for cross-matched subjects (recall that, by construction, own-matched high ability subjects and own-matched low ability subjects must earn \$2 on average in every round). We fit the pattern of divergence in earnings over rounds, although the magnitude of the divergence is not quite as big as that observed in the data (see Figure 4(b) in Section 3.2.3). Our simulations return a divergence in earnings over rounds even though, as in the observed data, high and low ability cross-matched subjects' simulated choices remain similar throughout the experiment. Rule learning is crucial to explaining the pattern of divergence: Table 2 shows that there are 58% of rule learners among cross-matched high ability subjects, but only 27% among cross-matched low ability subjects (the difference is statistically significant with a 2-sided $p = 0.021$), and Figure 8(b) shows that when we re-estimate the model without rule learners we can no longer fit the divergence in earnings over rounds. Further evidence that rule learning plays an important role in explaining subjects' choices more generally comes from the estimates of an alternative specification of the structural model without rule learners: Table 7 in Appendix B shows that this alternative specification (AS5) fits the data significantly less well.




(a) Own-matched subjects.

(b) Cross-matched subjects.

Notes: Simulated choices were obtained using a sequential method. Specifically, for each of the 170 groups in the experimental sample we drew the type of each group member from the appropriate estimated distribution of learner types reported in Table 2. We then simulated the choice of each of the 3 group members in round 2, given the observed behavior, specifically the group average, in the first round. Next, we simulated the choice of each of the 3 group members in round 3, given the average of the *simulated* choices of the group members in round 2. We continued sequentially in this manner to round 10, with simulated choices in each round being based on the group-level average of the simulated choices in the previous round. This procedure was repeated 100 times for each of the 170 groups.

Figure 7: Observed and simulated round-by-round means of choices.



(a) Preferred specification.

(b) Without rule learners.

Notes: Simulated earnings were computed from the simulated choices generated as described in the notes to Figure 7.

Figure 8: Simulated round-by-round means of earnings of cross-matched subjects.

Table 3 shows that the simulated choices match the extent of equilibrium and close-to-equilibrium play that we see in the data quite well. In particular, the simulated choices fit well the broad pattern of increasing convergence as the proportion of high ability subjects in the group goes up.

	Equilibrium	Group mean ≤ 1	Group mean ≤ 2
<i>Own-matched high ability groups:</i>			
Observed	0.367	0.500	0.583
Simulated	0.239	0.396	0.514
<i>Cross-matched groups:</i>			
Observed	0.145	0.282	0.368
Simulated	0.118	0.234	0.347
<i>Own-matched low ability groups:</i>			
Observed	0.050	0.133	0.217
Simulated	0.097	0.210	0.315

Notes: Simulated proportions were computed from the simulated choices generated as described in the notes to Figure 7.

Table 3: Observed and simulated proportions of equilibrium and close-to-equilibrium play in rounds 9 and 10.

Appendix B provides further evidence that our structural level- k model of learning fits well. In particular, Table 7 reports values of log likelihoods and of sums of squared deviations of choices, earnings and convergence statistics computed from simulated choices and shows that our preferred specification fits better than various alternative specifications (some nested and others not), while Table 8 shows that our model continues to perform well out-of-sample.

4.3.2 The impact of cognitive ability on average level- k choice rules

We now analyze how learner types vary with subjects' own cognitive ability and that of their opponents. To do this, we summarize the estimated proportions of learner types in a single statistic measuring the average level- k choice rule followed by the subjects (the notes to Table 4 provide the details of how these averages of choice rules are computed). From Table 4, we can see that, across all rounds, the average level- k choice rule followed by own-matched high ability subjects is 2.08, while for own-matched low ability subjects the average level is 1.55, with the difference statistically significant at the 1% level. In Section 4.3.3, we simulate the earnings that accrue to the different learner types, and show that the higher average level of the own-matched high ability subjects is not fully explained by the higher number of high ability opponents that they face. For cross-matched subjects, the difference is less pronounced but still evident: the average level followed by cross-matched high ability subjects is 1.77, while for cross-matched low ability subjects the average level is 1.54, with the difference significant at the 5% level. The average level followed by cross-matched high ability subjects is higher even though they face a lower number of high ability opponents on average than do cross-matched low ability subjects.

Thus subjects' own cognitive ability type has a significant effect on the average level that they follow: high cognitive ability subjects follow level- k choice rules that are significantly higher on average than those followed by low ability subjects. Further evidence that cognitive ability plays an important role in determining subjects' choice rules comes from estimating an alternative specification of the structural model in which the probabilities of the different learner types are not allowed to vary with subjects' own cognitive ability or with whether subjects are

Own-matched subjects			Cross-matched subjects			Diff.- in-Diff.
High ability	Low ability	Diff.	High ability	Low ability	Diff.	
2.076*** (0.128)	1.554*** (0.098)	0.523*** (0.161)	1.770*** (0.067)	1.542*** (0.078)	0.228** (0.103)	0.295 (0.191)

High ability subjects			Low ability subjects			Diff.- in-Diff.
Own-matched	Cross-matched	Diff.	Cross-matched	Own-matched	Diff.	
2.076*** (0.128)	1.770*** (0.067)	0.307** (0.144)	1.542*** (0.078)	1.554*** (0.098)	-0.012 (0.125)	0.318* (0.191)

Notes: The average choice rule of an Lk type is k , as this type follows the level- k choice rule in every round. The average choice rule of an $Lk - (k + 1)$ rule learner is $k + \frac{1}{2}$. We use the estimated proportions of learner types from Table 2 to compute the average level- k choice rule over all learner types. Tests for the significance of differences and of differences-in-differences are 2-sided. Standard errors are shown in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% levels.

Table 4: Estimated averages of level- k choice rules.

in own-matched or cross-matched groups. Table 7 in Appendix B shows that this alternative specification (AS6) fits the data significantly less well.

We also find significant differences in how average level- k choice rules respond to the cognitive ability of opponents. We noted in Section 2.2 that, on average, high ability subjects face $4/3$ more high ability opponents in own-matched groups than in cross-matched groups, while low ability subjects face $4/3$ more high ability opponents in cross-matched groups than in own-matched groups. The second row of Table 4 shows that high cognitive ability subjects respond to the cognitive ability of their opponents: across all rounds, the average level followed by own-matched high cognitive ability subjects is 0.31 higher than that followed by cross-matched high ability subjects, with the difference significant at the 5% level. Low cognitive ability subjects, on the other hand, do not respond at all to the cognitive ability of their opponents: the average level followed by cross-matched low ability subjects is almost identical to that followed by own-matched low ability subjects; and the difference of -0.01 is not close to being statistically significant. Furthermore, Table 4 shows that the difference-in-difference in how the average levels followed by high and low cognitive ability subjects respond to the cognitive ability of opponents is statistically significant (a 2-sided test shows significance at the 10% level).

Thus subjects' own cognitive ability has a significant effect on how the average level that they follow responds to the cognitive ability of their opponents: the average level- k choice rule followed by high cognitive ability subjects responds significantly to the number of high ability opponents that they face, while the average level followed by low cognitive ability subjects does not respond at all. The earnings simulations in Section 4.3.3 show that, by not adjusting their level upward on average, low cognitive ability subjects lose out more than do high cognitive ability subjects when the cognitive ability of their opponents goes up. Further evidence that cognitive ability plays an important role in determining how subjects' choice rules respond to the cognitive ability of opponents comes from estimating an alternative specification of the structural model in which the probabilities of the different learner types are not allowed to vary

with whether subjects are in own-matched or cross-matched groups (but are allowed to vary with subjects' own cognitive ability). Table 7 in Appendix B shows that this alternative specification (AS7) fits the data significantly less well.

4.3.3 Levels and earnings

Finally, we turn to the question of which learner types earn the most, given the estimated distribution of learner types of their opponents. Table 5 shows how simulated earnings vary according to learner type (the notes to the table describe how the simulated earnings were computed). The differences in earnings by learner type are large: for instance, L2 types earn about 50% more than L1 types across the board. In every case, the 3 best-performing learner types are L3 types, L4 types and L3-4 rule learners, although the order varies. This means that it is always optimal to be a high-level learner type, even for subjects facing a large proportion of low ability opponents who tend to choose higher numbers. In order of decreasing earnings, the order among the other 6 types is always: L2-3 rule learners; L2 types; L1-2 rule learners; L1 types; L1-0 rule learners; and finally L0 types.

	Own-matched subjects		Cross-matched subjects	
	High ability	Low ability	High ability	Low ability
L0 Type	8.31	8.67	8.80	8.25
L1 Type	13.17	15.43	14.65	14.16
L2 Type	19.94	21.89	20.94	20.75
L3 Type	22.04	23.95	24.15 [†]	23.54
L4 Type	22.29	23.47	23.20	23.94 [†]
L0-1 Rule learner type	10.84	11.70	11.86	11.00
L1-2 Rule learner type	16.71	18.20	18.19	17.26
L2-3 Rule learner type	20.59	22.76	22.56	21.62
L3-4 Rule learner type	22.64 [†]	24.10 [†]	24.08	23.77

Notes: [†] denotes the learner type that maximizes earnings. The earnings of a player of a particular learner type were computed based on simulated choices constructed using a method similar to that described in the notes to Figure 7 in which opponents' types were drawn from the appropriate estimated distribution of learner types from Table 2.

Table 5: Simulated mean total earnings (in dollars) by learner type.

Given the estimated proportions of learner types reported in Table 2 and the average level- k choice rules that subjects follow reported in Table 4, it is clear that subjects are constrained in their levels below those that are optimal, given the learner types of their opponents. Of course, unless subjects all play equilibrium, their level- k choice rules will tend to be too low on average; nonetheless, subjects are leaving a substantial amount of money on the table.

On average own-matched high ability subjects earn \$20 across all rounds, and so are leaving \$2.64 on the table compared to the payoff-maximizing L3-4 rule learner type (recall from Table 4 that the average level- k choice rule followed by these subjects is only 2.08). Own-matched low ability subjects also earn \$20 on average, and so are leaving \$4.10 on the table compared to the payoff-maximizing L3-4 rule learner type (from Table 4, their average level is 1.55). Even though

high and low ability own-matched subjects earn the same amount on average by construction, the low ability subjects are further from optimal play: they leave more money on the table and their average level- k choice rule is further from that followed by the payoff-maximizing learner type. Thus, the higher average level of the own-matched high ability subjects is not fully explained by the higher number of high ability opponents that they face.

Cross-matched high ability subjects earn \$21.78 on average, and so are leaving \$2.37 on the table compared to the payoff-maximizing L3 type (their average level is 1.77). Cross-matched low ability subjects earn \$18.22 on average, and so are leaving \$5.72 on the table compared to the payoff-maximizing L4 type (their average level is just 1.54). L1 types are the most common among cross-matched low ability subjects, making up 41% of the total: they earn a striking \$9.78 less on average than the payoff-maximizing learner type.

Finally, we can see that, by not adjusting their level upward on average, the low cognitive ability subjects lose out more when the cognitive ability of their opponents goes up: high ability subjects leave \$0.27 more on the table in own-matched high ability groups than in cross-matched groups (from above, \$2.64 minus \$2.37), while low ability subjects leave \$1.62 more on the table in cross-matched groups than in own-matched low ability groups (from above, \$5.72 minus \$4.10).

5 Conclusion

Cognitive ability varies greatly across the population. Our analysis is the first to show that cognitive ability affects behavior and learning in a repeated strategic environment: for example, our results show that cognitive ability drives observed heterogeneity in choices and earnings, and strongly predicts how quickly groups of agents learn to play equilibrium. A structural level- k model of learning helps to understand the differences in behavior that we discover and reveals that high cognitive ability agents follow significantly higher level- k choice rules than do agents of low cognitive ability.

Our model of learning is portable to other strategic games; thus we hope that future research will identify the extent to which the processes that we bring to light can help explain behavioral differences according to cognitive ability in other strategic settings where learning is important. In other words, the model is testable outside our specific experimental setup. Indeed, many real-world transactions are repeated and strategic. For example, a firm may make a sequence of entry and exit decisions, while firms and individuals repeatedly procure or sell items via auctions. In the context of the labor market, hiring decisions and many aspects of the Principal-Agent relationship are repeated and strategic in nature. Our results suggest that cognitive ability is likely to affect the dynamics of behavior and profits or incomes in these real-world settings: cognitively less able agents may learn more slowly than more able agents, and even as behavior approaches equilibrium low cognitive ability may remain associated with relatively unfavorable economic outcomes.

Valuable extensions of our analysis would investigate the broader ethical and practical implications of our results. There is a wider debate about how clever agents could create environments and mechanisms designed to exploit learning deficiencies (Sobel, 2000, p. 259). In our context, more cognitively able agents may expend resources to ensure that they interact with those less

cognitively able. In addition to being socially wasteful, such efforts have the potential to increase income inequality. We leave it to future research to investigate the merits of interventions, such as training and advice, designed to reduce the earnings gap between high and low cognitive ability agents who interact repeatedly.

Appendix

A Effect of allocation to cognitive ability type

Section 2.2 describes how subjects were allocated to cognitive ability type. To test whether the allocation to cognitive ability type *per se* affected behavior or earnings, we regress *p*-beauty contest choices and earnings on cognitive ability type *controlling for a subject's own test score* by including a full set of Raven test score dummies. We can follow this approach since the session median Raven test score varied from 37.0 to 42.5, giving some degree of randomness in the allocation for subjects towards the middle of the cognitive ability range. Figure 1(b) in Section 2.2 shows the densities of Raven test scores for high and low ability subjects separately and thus illustrates the region of overlap in the middle cognitive ability range. In the case of earnings, we consider only subjects in cross-matched groups, since low cognitive ability subjects in own-matched low ability groups must by construction earn as much on average as high cognitive ability subjects in own-matched high ability groups. We run the regressions with and without a control for the cognitive ability type of opponents. As Table 6 illustrates, we find no statistically significant effects of the allocation to cognitive ability type, either for the first round, the second round, the first 2 rounds together, the first 5 rounds or all 10 rounds together. Note also how the signs of the coefficients are unstable.

	Not controlling for opponent type		Controlling for opponent type	
	Choices	Earnings (\$ per round)	Choices	Earnings (\$ per round)
Round 1	1.96 [0.551]	0.29 [0.680]	2.34 [0.480]	0.28 [0.712]
Round 2	-2.46 [0.392]	-0.19 [0.780]	-2.12 [0.466]	-0.30 [0.637]
Rounds 1-2	-0.25 [0.909]	0.05 [0.910]	0.11 [0.961]	-0.01 [0.980]
Rounds 1-5	-0.42 [0.830]	-0.03 [0.912]	0.19 [0.923]	-0.13 [0.612]
Rounds 1-10	0.90 [0.581]	0.15 [0.514]	1.34 [0.400]	0.05 [0.820]

Notes: To isolate the effect of the allocation to cognitive ability type all the regressions include a full set of Raven test score dummies, which control for a subject's own ability. To control for the cognitive ability of opponents in the right-hand-side panel, we include a variable measuring the proportion (0, 0.5 or 1) of high ability opponents. 2-sided *p* values are shown in brackets and were constructed allowing clustering at the group level. *, ** and *** denote significance at the 10%, 5% and 1% levels.

Table 6: Effect of allocation to high cognitive ability type.

B Further goodness of fit analysis

Tables 7 and 8 provide further evidence that our structural level- k model of learning fits well. Table 7 reports values of log likelihoods and of sums of squared deviations of choices, earnings and convergence statistics computed from simulated choices, which together show that our model fits the observed data well. The table also shows that our preferred specification fits better than various alternative specifications: in all cases, a Vuong test (for non-nested model comparisons) or a likelihood ratio test (for nested model comparisons) rejects the null that the preferred and alternative specification fit the data equally well in favor of the preferred specification. We now provide a brief description of each of the alternative specifications:

- AS1: Once a $Lk - (k + 1)$ rule learner switches to the level- $(k + 1)$ choice rule, she never switches back. In each round $r \geq 3$, one-eighth of the $Lk - (k + 1)$ rules learners switch to the level- $(k + 1)$ choice rule; therefore, as in the preferred specification, in round 2 they all follow the level- k choice and by round 10 they all follow the $Lk - (k + 1)$ choice rule.
- AS2: The normal distribution replaces the t -distribution in the model of the choice process.
- AS3: The choices of 100 come from the same t -distribution as for the other choices (instead of the probability of a choice of 100 coming from the round-specific but cognitive ability independent empirical frequency observed in the sample).
- AS4: Subjects following the level- k choice rule for $k > 0$ take into account the effect of their own choice on the target, and understand that subjects following lower level- k choices rules do so as well. For $k > 0$, $\mu(k, \mathbf{x}_{g,r-1})$ is thus given by a lower fraction of $\bar{x}_{g,r-1}$ than in the preferred specification.
- AS5: Rule learner types are not included in the model (so there are just five standard level- k learner types, with $k \in \{0, 1, 2, 3, 4\}$, who follow the level- k choice rule in all rounds).
- AS6: The probabilities of learner types are not allowed to vary with subjects' own cognitive ability or with whether subjects are in own-matched or cross-matched groups.
- AS7: The probabilities of learner types are not allowed to vary with whether subjects are in own-matched or cross-matched groups (but are allowed to vary with subjects' own cognitive ability).

Table 8 shows that our structural level- k model of learning continues to perform well out-of-sample, providing further support for the model. In particular, the table reports how the model performs when we simulate choices for all rounds, but estimate the model's parameters using only the data from rounds 1-8 (3rd and 4th columns), compared to performance using parameters estimated from all the data (1st and 2nd columns). In each case, we provide statistics of fit computed from the simulated choices for all rounds and for rounds 9 and 10 alone.

	Preferred specification	Alternative specifications			
		AS1	AS2	AS3	
<i>Sum of squared deviations of choices:</i>					
Own-matched high ability subjects	22.3	46.2	28.6	51.6	
Own-matched low ability subjects	84.1	213.0	85.7	132.6	
Cross-matched high ability subjects	31.5	93.5	12.1	58.1	
Cross-matched low ability subjects	31.8	60.4	53.9	32.1	
<i>Sum of squared deviations of earnings:</i>					
Cross-matched high ability subjects	0.207	0.214	0.279	0.212	
<i>Sum of squared deviations of convergence statistics:</i>					
Own-matched high ability groups	3.21	0.86	0.45	3.47	
Cross-matched groups	0.35	0.11	0.12	0.48	
Own-matched low ability groups	1.77	3.19	7.90	1.18	
Log likelihood	-14,162	-14,201	-15,330	-14,324	
<i>Test against the preferred specification:</i>					
<i>p</i> value	-	0.013 ^a	0.000 ^b	0.000 ^a	
		Alternative specifications			
		AS4	AS5	AS6	AS7
<i>Sum of squared deviations of choices:</i>					
Own-matched high ability subjects	28.6	29.9	52.4	22.9	
Own-matched low ability subjects	78.8	185.2	187.2	114.6	
Cross-matched high ability subjects	31.3	57.7	17.7	39.7	
Cross-matched low ability subjects	22.4	49.9	31.6	21.5	
<i>Sum of squared deviations of earnings:</i>					
Cross-matched high ability subjects	0.297	0.242	0.489	0.201	
<i>Sum of squared deviations of convergence statistics:</i>					
Own-matched high ability groups	0.78	4.92	15.64	8.12	
Cross-matched groups	0.09	0.59	0.17	0.11	
Own-matched low ability groups	3.19	1.12	4.00	1.28	
Log likelihood	-14,197	-14,221	-14,185	-14,176	
<i>Test against the preferred specification:</i>					
<i>p</i> value	0.060 ^a	0.000 ^b	0.003 ^b	0.024 ^b	

Notes: Descriptions of each alternative specification are in the text of Appendix B. Squared deviations of choices and earnings (in dollars) were computed in each round from the simulated choices generated as described in the notes to Figure 7, and then summed over rounds. By construction, the sum of squared deviations of earnings: (i) is identical for cross-matched high ability and low ability subjects; and (ii) is zero for both own-matched high ability and low ability subjects. Squared deviations of convergence statistics were computed for each of the 3 proportions of equilibrium and close-to-equilibrium play in Table 3, and then summed and multiplied by 100.

^a 2-sided Vuong test.

^b Likelihood ratio test.

Table 7: Goodness of fit: preferred specification vs. alternatives.

	Estimation using all rounds		Estimation using rounds 1-8	
	Goodness of fit:		Goodness of fit:	
	All rounds	Rounds 9 & 10	All rounds	Rounds 9 & 10
<i>Sum of squared deviations of choices:</i>				
Own-matched high ability subjects	22.3	1.84	27.1	1.41
Own-matched low ability subjects	84.1	2.44	91.1	4.14
Cross-matched high ability subjects	31.5	1.61	38.4	4.22
Cross-matched low ability subjects	31.8	0.70	26.4	4.12
<i>Sum of squared deviations of earnings:</i>				
Cross-matched high ability subjects	0.207	0.057	0.351	0.162
<i>Sum of squared deviations of convergence statistics:</i>				
Own-matched high ability groups	-	3.21	-	0.63
Cross-matched groups	-	0.35	-	0.08
Own-matched low ability groups	-	1.77	-	2.03
Log likelihood	-14,162	-	-14,195	-

Notes: The statistics were computed in the same way as described in Table 7, except that the 2nd and 4th columns report the statistics computed only for rounds 9 and 10, while the 3rd and 4th columns were computed from simulated choices based on an estimate of the parameter vector θ obtained using only the data from rounds 1-8. The statistics in the 4th column are therefore out-of-sample quantities.

Table 8: Goodness of fit: in-sample vs. out-of-sample.

C Experimental instructions [Intended for online publication]

Please look at your screen now. I am reading from the instructions displayed on your screen. Please now turn off cell phones and any other electronic devices. These must remain turned off for the duration of this session. Please do not use or place on your desk any personal items, including pens, paper, phones etc. Please do not look into anyone else's booth at any time. Thank you for participating in this experimental session on economic decision-making. You were randomly selected from the Economic Science Laboratory's pool of subjects to be invited to participate in this session. There will be a number of pauses for you to ask questions. During such a pause, please raise your hand if you want to ask a question. Apart from asking questions in this way, you must not communicate with anybody in this room or make any noise.

You will be paid a show-up fee of \$5 together with any money you accumulate during this session. The amount of money you accumulate will depend partly on your actions and partly on the actions of other participants. You will be paid privately in cash at the end of the session. The session is made up of 2 parts. In the first part you will complete a test. Right at the end of the session you will find out your own test score, but you will not be paid for completing the test. I will describe the second part of the session after you have completed the test. Please raise your hand if you have any questions.

I will now describe the test which makes up the first part of the session. The test is made up of 60 questions, divided into parts A, B, C, D and E. Each of these parts is made up of 12 questions. For every question, there is a pattern with a piece missing and a number of pieces below the pattern. You have to choose which of the pieces below is the right one to complete the pattern. For parts A and B of the test, you will see 6 pieces that might complete the pattern. For parts C, D and E you will see 8 pieces that might complete the pattern. In every case, one and only one of these pieces is the right one to complete the pattern.²¹ For each question, please enter your answer in the column to the right of the pattern. You will score 1 point for every right answer. You will not be penalized for wrong answers. You will have 3 minutes to complete each of parts A and B, and you will have 8 minutes to complete each of parts C, D, and E. During each part, you can move back and forth between the 12 questions in that part and you can change your previous answers. The top right-hand corner of the screen will display the time remaining (in seconds). Before we start the test, please raise your hand if you have any questions. During the test, please raise your hand if you have a problem with your computer. [Subjects complete test]

Your screen is now displaying whether your test score was in the top half of the test scores of all participants in the room or was in the bottom half of the test scores of all participants. [30 second pause] [Example (not read aloud): Your test score was in the top half of the test scores of all participants in the room.] At the end of the session you will find out your own test score.

I will now describe the second and final part of the session. This second part is made up of 10 rounds. You will be anonymously matched into groups of 3 participants. You will stay in the same group for all 10 rounds. In each round, you and your other 2 group members will separately choose a whole number between 0 and 100 (0, 100 or any whole number in between is allowed). The group member whose chosen number is closest to 70% of the average of all

²¹The wording of this description follows the standard Raven test convention.

3 chosen numbers will be paid \$6 for that round and the other 2 group members will be paid nothing. If more than one group member chooses a number which is closest to 70% of the average of all 3 chosen numbers, the \$6 will be split equally among the group members who chose the closest number or numbers. Your total payment will be the sum of your payments in each round together with your show-up fee of \$5. In each round you will have 90 seconds to choose your number. If you choose your number early you will still have to wait until the end of the 90 seconds. The top right-hand corner of the screen will display the time remaining (in seconds). The screen will also include a reminder of the rules.

At the end of each round you will discover: (i) the numbers chosen by all your group members; (ii) the average of all 3 chosen numbers; (iii) what 70% of the average of all 3 chosen numbers was; and (iv) how much each group member will be paid for the round. Please raise your hand if you have any questions.

You will stay in the same group of 3 for all 10 rounds. Each group member has been randomly allocated a label, X, Y or Z. Your screen is now displaying your label and whether the test scores of the members of your group were in the top half or the bottom half of the test scores of all participants in the room. [60 second pause] [Example (not read aloud): You are group member Y. Your test score was in the top half of the test scores of all participants in the room. You have been matched with 2 participants (group member X and group member Z). Group member X was randomly selected from those whose test scores were also in the top half. Group member Z was randomly selected from those whose test scores were in the bottom half.] Please raise your hand if you have any questions. There will be no further opportunities for questions.

[10 rounds of beauty contest with feedback as described in Section 2.3]

[Screen asks subjects to report their gender]

[Screen reports the subject's score in the Raven test]

The session has now finished. Your total cash payment, including the show-up fee, is displayed on your screen. Please remain in your seat until you have been paid. Thank-you for participating.

References

- Agranov, M., Caplin, A., and Tergiman, C.** (2011). The process of choice in guessing games. *Mimeo, Caltech*
- Agranov, M., Potamites, E., Schotter, A., and Tergiman, C.** (2012). Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior*, 75(2): 449–463
- Bayer, R.C. and Renou, L.** (2012). Logical abilities and behavior in strategic-form games. *Mimeo, University of Leicester*
- Benjamin, D.J., Brown, S.A., and Shapiro, J.M.** (forthcoming). Who is behavioral? Cognitive ability and anomalous preferences. *Journal of the European Economic Association*
- Bergman, O., Ellingsen, T., Johannesson, M., and Svensson, C.** (2010). Anchoring and cognitive ability. *Economics Letters*, 107(1): 66–68
- Berndt, E., Hall, B., Hall, R., and Hausman, J.** (1974). Estimation and inference in nonlinear statistical models. *Annals of Economic and Social Measurement*, 3(4): 653–665
- Brañas-Garza, P., García-Muño, T., and Hernán, R.** (forthcoming). Cognitive effort in the beauty contest game. *Journal of Economic Behavior & Organization*
- Brocas, I., Carrillo, J.D., Wang, S.W., and Camerer, C.F.** (2011). Imperfect choice or imperfect attention? Understanding strategic thinking in private information games. *Mimeo, University of Southern California*
- Burks, S.V., Carpenter, J.P., Goette, L., and Rustichini, A.** (2009). Cognitive skills affect economic preferences, strategic behavior, and job attachment. *Proceedings of the National Academy of Sciences*, 106(19): 7745–7750
- Burnham, T., Cesarini, D., Johannesson, M., Lichtenstein, P., and Wallace, B.** (2009). Higher cognitive ability is associated with lower entries in a p -beauty contest. *Journal of Economic Behavior & Organization*, 72(1): 171–175
- Camerer, C.F., Ho, T.H., and Chong, J.K.** (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3): 861–898
- Carpenter, P.A., Just, M.A., and Shell, P.** (1990). What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices test. *Psychological review*, 97(3): 404–413
- Charness, G., Rustichini, A., and van de Ven, J.** (2011). Self-confidence and strategic deterrence. *Mimeo, UCSB*
- Chen, Y. and Gazzale, R.** (2004). When does learning in games generate convergence to Nash equilibria? The role of supermodularity in an experimental setting. *American Economic Review*, 94(5): 1505–1535
- Coricelli, G. and Nagel, R.** (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 106(23): 9163–9168
- Costa-Gomes, M., Crawford, V.P., and Broseta, B.** (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5): 1193–1235
- Costa-Gomes, M.A. and Crawford, V.P.** (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5): 1737–1768
- Costa-Gomes, M.A., Crawford, V.P., and Iriberri, N.** (2009). Comparing models of strategic thinking in Van Huyck, Battalio, and Beil’s coordination games. *Journal of the European Economic Association*, 7(2-3): 365–376
- Costa-Gomes, M.A. and Weizsäcker, G.** (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies*, 75(3): 729–762
- Crawford, V.P., Costa-Gomes, M.A., and Iriberri, N.** (forthcoming). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*
- Crawford, V.P. and Iriberri, N.** (2007a). Level- k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica*, 75(6): 1721–1770
- Crawford, V.P. and Iriberri, N.** (2007b). Fatal attraction: Salience, naïveté, and sophistication in experimental “Hide-and-Seek” games. *American Economic Review*, 97(5): 1731–1750

- Dohmen, T., Falk, A., Huffman, D., and Sunde, U.** (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review*, 100(3): 1238–60
- Duffy, J. and Nagel, R.** (1997). On the robustness of behaviour in experimental ‘beauty contest’ games. *Economic Journal*, 107(445): 1684–1700
- Dufwenberg, M., Sundaram, R., and Butler, D.J.** (2010). Epiphany in the Game of 21. *Journal of Economic Behavior & Organization*, 75(2): 132–143
- Fischbacher, U.** (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2): 171–178
- Frederick, S.** (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4): 25–42
- Gneezy, U., Rustichini, A., and Vostroknutov, A.** (2010). Experience and insight in the Race game. *Journal of Economic Behavior & Organization*, 75(2): 144–155
- Goldfarb, A. and Xiao, M.** (2011). Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. *American Economic Review*, 101(7): 3130–3161
- Gray, J. and Thompson, P.** (2004). Neurobiology of intelligence: Science and ethics. *Nature Reviews Neuroscience*, 5(6): 471–482
- Heckman, J.J., Stixrud, J., and Urzua, S.** (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics*, 24(3): 411–482
- Ho, T.H., Camerer, C., and Weigelt, K.** (1998). Iterated dominance and iterated best response in experimental “*p*-beauty contests”. *American Economic Review*, 88(4): 947–969
- Ho, T.H. and Su, X.** (2012). A dynamic level-*k* model in sequential games. *Mimeo, University of California at Berkeley*
- López, R.** (2001). On *p*-beauty contest integer games. *UPF Economics and Business Working Paper No. 608*
- Millet, K. and Dewitte, S.** (2007). Altruistic behavior as a costly signal of general intelligence. *Journal of Research in Personality*, 41(2): 316–326
- Nagel, R.** (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5): 1313–1326
- Oechssler, J., Roider, A., and Schmitz, P.W.** (2009). Cognitive abilities and behavioral biases. *Journal of Economic Behavior & Organization*, 72(1): 147–152
- Offerman, T., Potters, J., and Sonnemans, J.** (2002). Imitation and belief learning in an oligopoly experiment. *Review of Economic Studies*, 69(4): 973–997
- Ohtsubo, Y. and Rapoport, A.** (2006). Depth of reasoning in strategic form games. *Journal of Socio-economics*, 35(1): 31–47
- Östling, R., Wang, J.T., Chou, E.Y., and Camerer, C.F.** (2011). Testing game theory in the field: Swedish LUPI lottery games. *American Economic Journal: Microeconomics*, 3(3): 1–33
- Palacios-Huerta, I. and Volij, O.** (2009). Field centipedes. *American Economic Review*, 99(4): 1619–1635
- Raven, J., Raven, J.C., and Court, J.H.** (2000). *Manual for Raven’s Progressive Matrices and Vocabulary Scales Section 3: Standard Progressive Matrices*. Pearson, San Antonio, TX
- Roth, A. and Xing, X.** (1994). Jumping the gun: Imperfections and institutions related to the timing of market transactions. *American Economic Review*, 84(4): 992–1044
- Rydval, O., Ortman, A., and Ostadnick, M.** (2009). Three very simple games and what it takes to solve them. *Journal of Economic Behavior & Organization*, 72(1): 589–601
- Schnusenberg, O. and Gallo, A.** (2011). On cognitive ability and learning in a beauty contest. *Journal for Economic Educators*, 11(1): 13–24
- Sobel, J.** (2000). Economists’ models of learning. *Journal of Economic Theory*, 94(2): 241–261
- Stahl, D.O.** (1996). Boundedly rational rule learning in a guessing game. *Games and Economic Behavior*, 16(2): 303–330
- Stahl, D.O. and Wilson, P.W.** (1995). On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1): 218–254