Steps of reasoning in children and adolescents *

Isabelle Brocas
University of Southern California and CEPR

Juan D. Carrillo
University of Southern California and CEPR

April 2019

Abstract

We develop a novel graphical paradigm of a strict dominance solvable game to study for the first time the developmental trajectory of steps of reasoning between 8 years old and adulthood. Most participants play the equilibrium action either always or only when they have a dominant strategy. Although age is a determinant of equilibrium choice, some very young participants display an innate ability to play at equilibrium. Finally, the proportion of equilibrium play increases significantly between until 5th grade and stabilizes afterwards, suggesting that the contribution of age to equilibrium play vanishes early in life.

Keywords: developmental decision-making, backward induction, steps of dominance.

JEL Classification: C91, D83.

*We thank Sobhana Atluri, the members of the Los Angeles Behavioral Economics Laboratory (LABEL) and seminar participants at USC and Stanford for their insights and comments. We are grateful to the staff of the Lycée International de Los Angeles (LILA) and Thomas Starr King (KING) –in particular Emmanuelle Acker, Nordine Bouriche, Adriana Díaz, Mathieu Mondange and Anneli Harvey– for their help and support running the experiment in their schools. The study was conducted with the University of Southern California IRB approval UP-12-00528. Address for correspondence: Isabelle Brocas, Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089, USA, <brocas@usc.edu>.

1 Introduction

At which age are individuals capable of selecting rational, forward-looking, optimal decisions in multi-person games of strategy? The goal of this paper is to provide a first answer in a simple, well-defined game theoretic setting. Existing research on adults documents significant disparity in depths of reasoning across individuals,[1] and it is difficult to overemphasize the importance of determining the *reasons* for this observed heterogeneity. One argument stressed in recent studies is the positive correlation between cognitive skills and strategic sophistication in economic choices.[2] We believe that the process through which sophistication develops is hiding the key to understand the reasons for heterogeneity. Therefore, observing how the ability to think strategically evolves with age should provide invaluable evidence to identify general patterns in the acquisition of sophistication, differences across individual trajectories, and causal mechanisms. Is strategic thinking innate? Acquired gradually? The result of experience or repeated exposure? Impacted by the environment in which we grow?

Surprisingly and despite some exceptions (Sher et al., 2014; Czermak et al., 2016; Chen et al., 2016; Brocas et al., 2017; Brocas and Carrillo, 2018; Fe and Gill, 2018), the experimental literature has devoted little attention to the evolution of behavior in strategic games from childhood to adulthood.[3] Existing studies point to behavioral age differences but leave several key questions unanswered. First, evidence is often reported on a snapshot of the developmental trajectory. Second, paradigms are often complex and require the aggregation of many abilities to achieve rational play, thereby introducing a confound between analytical ability, forward looking behavior, beliefs about others' choices and payoff maximizing considerations. This article is the first to directly test the developmental trajectory from childhood to adulthood of strategic sophistication in strict-dominance solvable games.

Dominance solvable games are particularly appealing because steps of reasoning offer a natural algorithm to solve them. Sophistication is defined naturally as the number of steps of reasoning that a subject is able to implement to get closer to the Nash equilibrium. Our objective is to investigate the relationship between that form of sophistication and age. Notice that we are not interested in situations where more steps of reasoning do not move the individual closer to Nash equilibrium (as for example in the 11-20 money request

[1]See e.g., Costa-Gomes et al. (2001); Johnson et al. (2002); Costa-Gomes and Crawford (2006); Brañas-Garza et al. (2011); Arad and Rubinstein (2012); Brocas et al. (2014, 2018) out of a long list. See also the survey by Crawford et al. (2013).

[2]See e.g., Brañas-Garza et al. (2012); Gill and Prowse (2016); Proto et al. (Forthcoming).

[3]Indeed, and as emphasized in Sutter et al. (2019) and List et al. (2018), most economic experiments with children and adolescents focus on individual decision-making paradigms (rationality of choices, time preferences, risk preferences and social preferences).

game of Arad and Rubinstein (2012)) and/or do not result in higher empirical payoffs (as for example in the traveler's dilemma game of Capra et al. (1999)). Instead, we want a setting in which (i) steps of reasoning provide the algorithm to play Nash, (ii) subjects with higher levels of reasoning invariably play closer to Nash and (iii) they obtain higher payoffs. This allows us to rank unambiguously the sophistication of participants.[4]

For this, we need to design a paradigm such that the ability to recursively think about others' behavior simultaneously facilitates the formulation of the (theoretical) equilibrium and the (empirical) payoff-maximizing strategy. We confront two major challenges for the age evolution analysis to be feasible. First, the problem needs to be sufficiently simple and transparent that young children can understand it (without being trivial for high schoolers and young adults). This consideration precludes the use of some standard paradigms, such as the two-person p-beauty contest (Costa-Gomes and Crawford, 2006). This type of games is intuitive for game theorists, replete with interesting and testable properties, but excessively intangible for the minds of children. Second and related, it is critical to minimize the abstract and formal structure of the game. These aspects may lead cognitive ability to confound with mathematical or analytical skills (a competence that is expected to develop during adolescence and facilitated by extra years of schooling). They make standard normal-form representations (Costa-Gomes et al., 2001; Brocas et al., 2018) also unsuitable for our study.[5]

To minimize these concerns, we propose a novel graphical interface where subjects possess three objects with three attributes each: a shape, a color and a letter. Their goal is to select an object with a certain characteristic, which depends on the object selected by another player in the game. This is true for all but one player, who must simply match a feature of a specific single object. This player's decision constitutes the starting point of the iteration process, and the problem of the other players can be iteratively solved by successive elimination, with a maximum of three steps of reasoning.

To analyze the developmental trajectory of behavior in our paradigm, we recruited three populations of children and adolescents. The main experiment reported here involves a population of children and adolescents (8 to 18 years old) recruited at a single private school in Los Angeles as well as a control young adult population from USC. We also recruited younger children (5 to 8 years old) from that same school and implemented a simpler version of the same task. Last, we recruited a population of middle schoolers (11 to 14 years old) from a single public school also in Los Angeles to assess the potential

[4]Fe and Gill (2018) provides an interesting study of a simplified 11-20 money request game with 5 to 12 years old children.

[5]The difficulty to understand a game when it is presented abstractly has been recognized (Chou et al., 2009; Cason and Plott, 2014). We believe it is exacerbated in the case of children.

impact of school characteristics on sophistication.

Our analysis yields three main results. First, the vast majority of participants either play always at equilibrium or they play at equilibrium only when they have a dominant strategy. There are few random players, and virtually no one exhibits an "intermediate" level of reasoning (i.e., plays at equilibrium when it requires two steps of reasoning but not when it requires three steps). This is in sharp contrast with the existing literature that emphasizes large heterogeneity in levels of reasoning and abundance of intermediate types (Costa-Gomes et al., 2001; Johnson et al., 2002; Costa-Gomes and Crawford, 2006; Brañas-Garza et al., 2011; Brocas et al., 2014; Gill and Prowse, 2016; Brocas et al., 2018). Second, although age is a significant determinant of equilibrium thinking, there is also an important innate ability component, and the change with age is not as steep as one might expect. Indeed, one-quarter of our 8 years old subjects play consistently at equilibrium. At the same time, one-third of our 17 years old subjects deviate significantly from it (only recognize a dominant strategy or follow a non-discernible pattern). Third and related, the amount of equilibrium play increases very significantly between 3rd and 5th grade and stabilizes afterwards. In other words, the contribution of age to equilibrium behavior vanishes relatively early in life (between 12 and 13 years old). Our data reveals important predictors of performance. We find that female participants and subjects with a self-reported preference for science subjects perform significantly better. Finally, differences across schools and across tracks within schools are also associated with differences in sophistication.

2 Experiment

2.1 Design and procedures

Participants. We report the results of three experiments, referred to as MAIN, KING and YOUNG. They correspond to the same paradigm but with three different populations.

- MAIN. Our main population consisted of 234 school-age participants studying at the Lycée International de Los Angeles (LILA), a French-English bilingual private school in Los Angeles. We enrolled students from 3rd (34), 4th (17), 5th (20), 6th (41), 7th (27), 8th (18), 9th (25), 10th (23) and 11th (29) grades. For comparison we also recruited 60 students at the University of Southern California (USC). We ran 20 sessions with school-age students in a classroom at LILA and 5 sessions with undergraduates at the Los Angeles Behavioral Economics Laboratory (LABEL) at USC. Sessions had 9, 12 or 15 participants and followed *identical* procedures. For each school-age session, we tried to have subjects from the same grade, but for logistical reasons we sometimes had to mix subjects of two consecutive grades.

Notice that most studies with children do not recruit an adult population. We believe it is important to include an adult control group to establish a behavioral benchmark, *even if the comparison is imperfect.* In our case, the majority of students at LILA are from caucasian families of upper-middle socioeconomic status. After graduating, they typically attend well-ranked colleges in Europe and North America (including USC). Despite some differences (nationality, family background, size of peer group, etc.), the two populations match 'reasonably' well.[6]

- KING. We ran the same experiment with 199 middle school children from the Film and Media Magnet at Thomas Starr King Middle School, a low-income public school in Los Angeles, located less than a mile away from LILA. We used the school's classification criteria to group participants into categories, corresponding to their academic achievements.

- YOUNG. We ran a simplified version of the experiment with 117 younger children from LILA in grades K, 1^{st} and 2^{nd}.

We will primarily focus on the MAIN population (294 participants). The results of KING (199) and YOUNG (117) are discussed in section 4.

Tasks. The experiment had two tasks implemented on PC tablets, programmed on 'Multi-stage Games' and always performed in the same order. The first task consisted of two trials of a "lying game," where subjects privately rolled a dice and were rewarded according to the number they reported. The findings of this project are discussed in a different article. The second task is the focus of this article. It consisted of 18 trials of a three-person simultaneous strict dominance solvable game.

We designed a simple, graphical interface, which was both accessible and appealing to children as young as 8 years of age. As noted earlier, it was of paramount importance that differences in behavior reflected as much as possible developmental differences in *cognitive abilities* rather than developmental differences in *mathematical skills*, *abstract thinking* and *attention.* This ruled out payoff matrices and other abstract, formal presentations standard in the literature. We ended up with a simple paradigm in which subjects were matched in groups of three and assigned a role as player 1, player 2 or player 3, from now one referred to as role 1, role 2 and role 3. Each player in the group had three objects, and each object had three attributes: a shape (square, triangle or circle), a color (red, blue or yellow) and a letter (A, B or C). Players had to simultaneously select one object.

[6] In any case, we strongly believe that an (imperfect) benchmark is superior to no benchmark.

Role 1 would win if the object he chose matched a given attribute of the object chosen by role 2. Similarly, role 2 would win if the object he chose matched a given attribute of the object chosen by role 3. Finally, role 3 would win if the object he chose matched a given attribute of an extra object. The attributes to be matched were different for different roles and specified by the experimenter. All options and objectives of players were common knowledge and displayed on the computer screen. Figure 1 provides a screenshot of the game as seen by role 2.

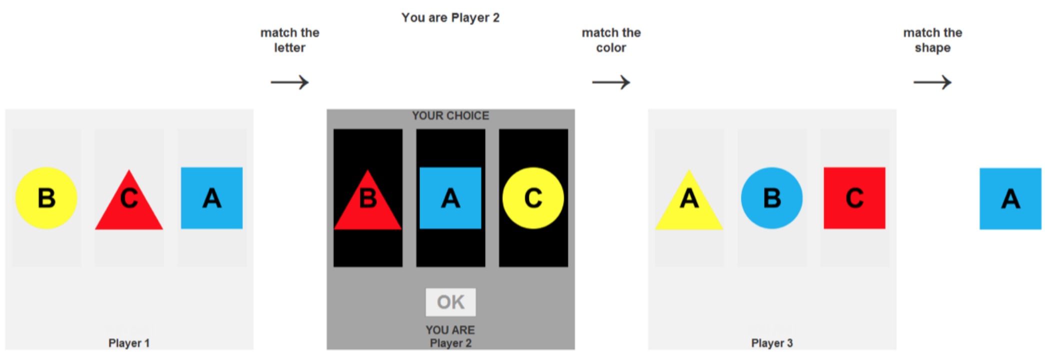

Figure 1: Screenshot of the game (as seen by role 2)

The game can be easily solved with an inductive argument starting from role 3. In the example of Figure 1, role 3 has to match the shape of the outside object, so he wins if he chooses the red square C. Conditional on that choice, role 2 wins if he chooses the red triangle B and, conditional on that choice, role 1 wins if he chooses the yellow circle B.

Participants played 6 trials in each role r ($\in \{1, 2, 3\}$) for a total of 18 trials, with anonymous partners randomly drawn after each trial and without feedback. After each trial, participants changed roles and the software changed the shapes, colors, letters and attributes to be matched. To ensure comprehension, we implemented a quiz before the 18 incentivized trials.[7]

We used 9 combinations of objects for the first 9 trials and we repeated them for the following 9 trials. This allowed us to study learning by comparing the choices in the first and second half of the experiment. Also, to distinguish between equilibrium reasoning and chance, we deliberately introduced focal, non-equilibrium objects in roles 1 and 2 of all 18 trials.[8] A transcript of the instructions and quiz is included in Appendix B.

[7]Subjects had to answer 4 questions. If they missed one or more, a warning sign would appear stating "not all answers are correct, please try again". The experiment started only after all subjects in a session had completed the quiz correctly.

[8]In the example of Figure 1, the blue square A was the extra object to the right of the screen, and it

Payoffs. Subjects accumulated points. We implemented three different conversions depending on the population. LILA students from grade 6^{th} and above, KING students and USC subjects had points converted into money paid immediately at the end of the experiment in cash (USC) or with an amazon e-giftcard (LILA and KING, where cash transfers on premises are not allowed). USC subjects accumulated \$0.50 and \$0.20 per successful and unsuccessful trial respectively, and a \$5 show-up fee. LILA and KING subjects accumulated \$0.40 and \$0.20 per successful and unsuccessful trial with no show-up fee.[9] The experiment lasted between 40 and 50 minutes. Average earnings (on the entire experiment and not including show-up fees) were \$15.2 (USC), \$11.2 (LILA) and \$10.8 (KING).

For elementary school subjects at LILA (grades K to 5^{th}), we set up a shop with 20 to 25 pre-screened, age-appropriate toys and stationery (bracelets, erasers, figurines, die-cast cars, trading cards, apps, calculators, earbuds, fidget spinners, etc.). Participants accumulated 40 points and 20 points per successful and unsuccessful trial respectively, and each toy had a different point price. Before the experiment, children were taken to the shop and showed the toys they were playing for. They were instructed about the prices of each toy and, for the youngest subjects, we explicitly stated that more points would result in more toys. At the end of the experiment, subjects learned their point earnings and were accompanied to the shop to exchange points for toys. We made sure that every child earned enough points to obtain at least three toys. At the same time, no child had excess points after choosing all the toys they liked.[10]

Questionnaire. We collected demographic information consisting of "gender", "age", "grade", "number of siblings" and "favorite subject at school".

2.2 Theory and hypotheses

Consistent with the experimental literature on dominance solvable games reviewed in the introduction, we expect that participants would differ in their ability to iteratively eliminate dominated strategies. More precisely, we anticipate to find four types of individuals: R (subjects who always play randomly), D_0 (subjects who play at equilibrium only if

was also in the choice set of roles 1 and 2. However, no player should, in equilibrium, select it. Appendix A3 reports heuristic rules based on focal objects.

[9] Incentives were calibrated to account for differences in marginal value of money and opportunity cost of time. We provided a positive payment for unsuccessful trials to artificially reduce variance and ensure a pleasant experience of our school-age participants.

[10] The procedure emphasizes the importance of accumulating points while making the experience enjoyable. At this age, a toy is also a significantly more attractive reward than money. Most children are familiar with this method of accumulating points or tickets that are subsequently exchanged for rewards since it is commonly employed in arcade rooms and fairs.

they have a dominant strategy), D_1 (subjects who play at equilibrium when they have a dominant strategy and can best respond to a D_0 type), and D_2 (subjects who can play as D_0 and D_1, as well as best respond to D_1). These types map well into *nested* levels of strategic sophistication from lowest (R) to highest (D_2). The game however does not allow us to distinguish between levels D_2 and above (D_3, Nash, etc.).

The predicted behavior is simple. R plays the equilibrium strategy 1/3 of the time in all roles; D_0 always plays the equilibrium strategy in role 3 and 1/3 of the time in roles 1 and 2; D_1 always plays the equilibrium strategy in roles 2 and 3 and 1/3 of the time in role 1; and D_2 always plays the equilibrium strategy. An immediate implication of this behavioral theory is that we should never observe a subject playing the equilibrium strategy significantly more often in an earlier role than in a later role.

We formulate the following hypotheses.

Hypothesis 1 *The behavior of the vast majority of individuals at all ages is consistent with one of the four types: R, D_0, D_1 or D_2.*

Hypothesis 1 states that the model that has proved successful to describe the reasoning process of adults in dominance-solvable games is expected to capture also the behavior of children and adolescents.

Hypothesis 2 *There are few or no D_2 types in our youngest school-age subjects. There are no R and few or no D_0 types in our oldest school-age and adult subjects. There are D_1 types in all ages.*

Informally, the idea behind Hypothesis 2 is that the game is hard to solve for young children, very easy to solve for young adults, and with a steep learning curve. Also, and in line with the existing literature, a fraction of subjects is expected to perform a positive but limited number of steps of reasoning (D_1).

Hypothesis 3 *There is a gradual and strictly monotonic shift in types with age, from lowest to highest level of sophistication (R to D_0 to D_1 to D_2).*

According to Hypothesis 3, sophistication increases steadily with age. We also anticipate higher sophistication in adults than in our oldest school-age students.

While we think that Hypotheses 1, 2 and 3 are natural, they carry important implications for developmental decision-making. Indeed, validating them would show that cognitive reasoning is quantitatively different but qualitatively identical across ages, with no milestones necessary to succeed in understanding dominance solvable games. It would also imply that, over time, we develop abilities that facilitate performing steps of reasoning. More specifically, they allow us to *smoothly* move from random behavior, to one step of reasoning, then two steps and finally three steps.

3 Results

3.1 Aggregate choice

Figure 2 reports the average number of equilibrium choices by grade and role.

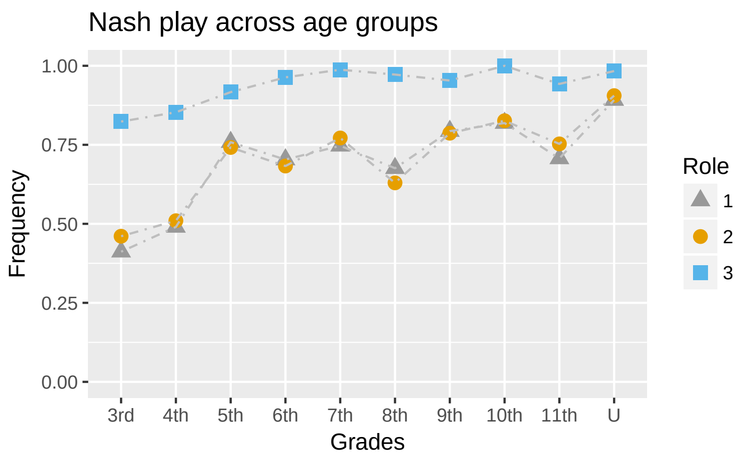

Figure 2: Proportion of equilibrium choices by grade and role

Within grades, equilibrium play is significantly higher in role 3 than in the other roles for all grades. By contrast, and to our surprise, equilibrium play is not significantly different between roles 1 and 2 for any grade. The behavior across grades also unveils interesting patterns. Equilibrium behavior in roles 1 and 2 is the same in 3$^{\text{rd}}$ and 4$^{\text{th}}$ grade, increases significantly in 5$^{\text{th}}$ grade, and remains constant afterwards (there is a dip in 8$^{\text{th}}$ grade, though it is not statistically significant). Compared to middle school, the USC population plays Nash only marginally more often. Finally, in all grades and roles the probability of equilibrium behavior is above 0.33. Therefore, the best response to the empirical behavior is to play the equilibrium strategy for all roles and grades. In other words, deviations from Nash cannot be explained by non-equilibrium behavior as a best response to the empirical strategy of others.

3.2 Individual analysis

Although aggregate behavior is instructive, patterns of choice at the individual level are more revealing. According to section 3.1, playing the equilibrium action is also the payoff-maximizing, best response strategy to the empirical behavior of the population. The fraction of individuals who play the equilibrium action in 17 or 18 trials out of the 18 trials of the game in each grade are: 3$^{\text{rd}}$ (0.24), 4$^{\text{th}}$ (0.12), 5$^{\text{th}}$ (0.60), 6$^{\text{th}}$ (0.46), 7$^{\text{th}}$ (0.59), 8$^{\text{th}}$ (0.39), 9$^{\text{th}}$ (0.64), 10$^{\text{th}}$ (0.78), 11$^{\text{th}}$ (0.59), USC (0.85).

Non-equilibrium players are frequent except for 10$^{\text{th}}$ graders and USC students. The patterns confirm also the aggregate analysis: we observe low levels of equilibrium compliance in 3$^{\text{rd}}$ and 4$^{\text{th}}$ grade, a statistical increase at 5$^{\text{th}}$ grade and a stabilization thereafter (with a statistically not significant dip in 6$^{\text{th}}$ and 8$^{\text{th}}$ grades and a statistically significant peak in 10$^{\text{th}}$ grade and USC).

Given the significant proportion of individuals who do not play the equilibrium strategy, we next classify subjects into types. We use a very simple method. We label the behavior in role r "equilibrium" if the subject played the equilibrium action 5 or 6 times (out of 6) and "random" if the subject played the equilibrium action 0, 1 or 2 times (out of 6). We then use the theory developed in section 2.2 to classify individuals into R (random in all roles), D_0 (random in roles 1 and 2 and equilibrium in role 3), D_1 (random in role 1 and equilibrium in roles 2 and 3), and D_2 (equilibrium in all roles). The remaining subjects are classified as O, for "Other".[11] Figure 3 reports the proportion of subjects by grade who are classified under each type, from most (bottom) to least (top) sophisticated.

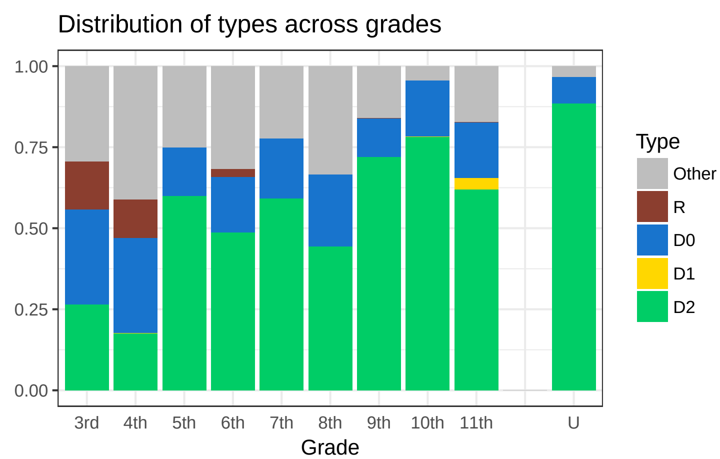

Figure 3: Proportion of subjects by type and grade

In strong support of Hypothesis 1, our theoretical model provides a very solid behavioral template. Indeed, the choice of 76% of LILA students and 97% of USC students can be accounted for by one of the four types described in section 2.2. The proportion of subjects who do not fit in one of these types (O) decreases with age, although it is

[11] Alternatively, we could structurally estimate types using maximum likelihood methods. This is superior only when data abounds and when subjects do not fall neatly into types. Given our data, our simple classification is more revealing.

statistically smaller only for 10^{th} graders.

Result 1 *Hypothesis 1 is supported by the data. The majority of subjects behave consistently with one of the four types of our behavioral model* (R, D_0, D_1, D_2).

However, there are very few R individuals and only one D_1 subject in the entire sample, that is, all the classified individuals are either D_0 (players who can only recognize a dominant strategy) or D_2 (equilibrium players). This is consistent with the result in Figure 2, which highlighted that aggregate equilibrium performance within a grade is very similar in roles 1 and 2. It is in sharp contrast with our Hypothesis 2. Indeed, we expected that subjects would learn gradually with age to perform more and more steps of reasoning. Instead, they either recognize only a dominant strategy or all the steps of reasoning. It is also radically different from the experimental literature that emphasizes large heterogeneity in steps of reasoning in adults. Admittedly, and for the purpose of being accessible to young children, our setting is simpler than most existing games. However, it is also devoid of an analytical framework. We conjecture that part of the reason why some individuals perform some but not all the steps of reasoning in traditional dominance-solvable games is because of the complexity of the formal presentation. In other words, it is possible (and worthy of further investigation) that intermediate levels of reasoning are the result of limitations in the ability to understand the structure of a game and/or mechanically compute all the required steps as opposed to a limitation in the ability to transform knowledge into equilibrium behavior.

Also against our Hypothesis 2, 26% of our 3^{rd} graders are classified as D_2 (21% play the equilibrium strategy in all 18 trials). Conversely, 34% of 11^{th} graders are classified as D_0 or O (21% play less than 4 out of 12 times the equilibrium strategy in roles 1 and 2). These two population are significantly different from 0.

Result 2 *Hypothesis 2 is not supported by the data. Subjects either recognize only a dominant strategy or always play at equilibrium. Also, some very young players display an innate ability to play always at equilibrium while some young adults are unable to perform two-steps of dominance.*

Finally, and in partial support of Hypothesis 3, we notice a weakly monotonic (but not gradual) increase in strategic types with age. Participants in 3^{rd} and 4^{th} grade are mainly type R and D_0 whereas participants in 5^{th} grade and above are predominantly D_2, with only small differences after the 5^{th} grade. However, this classification of proportions by grade is not the most adequate for statistical tests. Therefore, at this stage we refrain from making definitive assertions on the evolution of equilibrium behavior with age. A more in-depth study of this question is performed in the regression analysis of section 3.3.

Taken together, our findings suggest that the ability to solve dominance-solvable games develops differentially. While this ability is acquired instinctively by some young children, it eludes some educated young adults.[12] It also seems that the developmental trajectory plateaus (or at least decelerates) at a relatively young age, around 5$^{\text{th}}$ grade.

3.3 Regression analysis

To better understand the determinants of equilibrium behavior, we perform a series of OLS regressions at the individual level. We only consider the 234 school-age students, to avoid biasing the results with the undergraduate population. In columns 1, 2 and 3 of Table 1, the dependent variable is the percentage of equilibrium choices of each participant in role r ($\in \{1,2,3\}$). Our main independent variable is the *Age* in months of the participant at the date of the experiment. We include a dummy control variable for favorite topic at school ($STEM$ = 1) to account for analytical inclination. STEM refers to a reported preferences for Mathematics, Sciences or Technology. Consistent with the curriculum of the school, the other categories offered were English, French, History/Geography and Arts/Music, which we globally refer to as 'Arts & Humanities'. We also add demographic dummy variables for gender ($Male$ = 1) and whether the participant has siblings ($Siblings$ = 1).

As expected, *Age* is a key determinant of equilibrium behavior. Males play at equilibrium less often than females in roles 1 and 2. There is also a strong explanatory power of the self-reported preferred school topic. Indeed, participants who report a preference for STEM play the equilibrium action 20% more often in roles 1 and 2 than those who prefer Arts & Humanities.

The analysis in sections 3.1 and 3.2 suggests an increase in equilibrium behavior with age but also a deceleration after a certain grade. To further investigate these dynamic trends, we conduct a spline regression analysis to estimate the age at which such deceleration occurs for each role. The method consists of running OLS regressions assuming that a kink exists, and in identifying the kink that provides the best R^2-based fit. We include the same controls as above. The spline regressions by role are reported in columns 4, 5 and 6 of Table 1 and represented in Figure 4.

The regressions strongly support our previous conclusions. Performance in all roles (including role 3) increases significantly up to a certain age (around 12 years old), and then stabilizes approximately at 0.75 (roles 1 and 2) and 0.97 (role 3). Fits are virtually identical for roles 1 and 2, the slopes are not statistically different and the knots are estimated to be one month apart.

[12] The fact that most high schoolers and young adults who do not play consistently at equilibrium are classified as D_0 and not R suggests that they have paid attention to the game.

	OLS			Spline OLS		
	Role 1	Role 2	Role 3	Role 1	Role 2	Role 3
Age	0.003***	0.003***	0.001***	—	—	—
	(0.001)	(0.001)	(0.0003)			
Age (spline1)	—	—	—	0.009***	0.007***	0.004***
				(0.002)	(0.002)	(0.001)
Age (spline2)	—	—	—	-0.001	-0.0001	-0.001
				(0.001)	(0.001)	(0.001)
knot (months)	—	—	—	147	148	154
STEM	0.191***	0.224***	0.037	0.196***	0.228***	0.039
	(0.049)	(0.049)	(0.024)	(0.047)	(0.048)	(0.024)
Male	-0.082*	-0.104**	-0.010	-0.086*	-0.108**	-0.014
	(0.047)	(0.048)	(0.023)	(0.046)	(0.047)	(0.023)
Siblings	-0.059	-0.051	-0.019	-0.069	-0.058	-0.022
	(0.053)	(0.054)	(0.027)	(0.052)	(0.053)	(0.026)
Constant	0.240**	0.261**	0.744***	-0.533**	-0.319	0.453***
	(0.121)	(0.123)	(0.060)	(0.245)	(0.246)	(0.107)
Observations	234	234	234	234	234	234
Adj. R^2	0.112	0.123	0.048	0.156	0.147	0.086

(standard errors in parenthesis) * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 1: Regressions of equilibrium choice by role

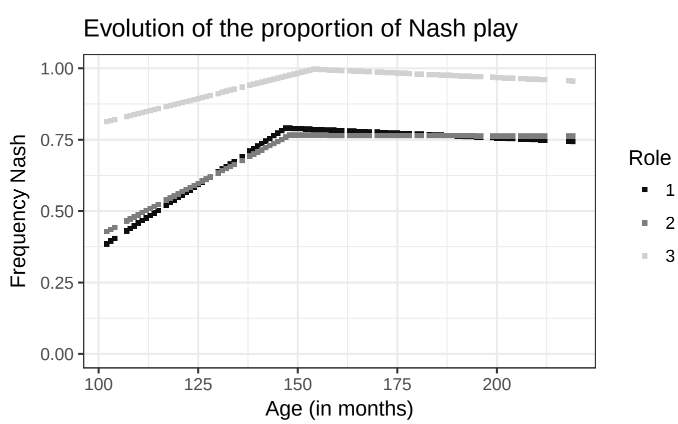

Figure 4: Spline OLS regression

Result 3 *Hypothesis 3 is weakly supported by the data. Equilibrium performance increases with age very significantly during elementary school but it stabilizes in 6^{th} grade.*

3.4 Other analyses

The main analysis suggests that age and a preference for STEM are strong predictors of Nash play at the aggregate level. We show in Appendix A1 that these variables also predict differences in individual performance and type. In Appendix A2, we split the sample between students who report a preference for STEM and those who prefer Arts & Humanities, to isolate the differences in the developmental trajectory of these two groups. For all grades above 5^{th}, students with a STEM preference play Nash more than 75% of the time in roles 1 and 2, while their counterpart never reach that performance, except in 10^{th} grade.

A substantial fraction of participants do not play the equilibrium strategy in roles 1 and 2. As suggested earlier, this cannot be accounted for by a best-response to the empirical behavior because playing Nash is always the payoff-maximizing strategy. It is also implausible that a participant assigned to role 2 deliberately forms the (empirically incorrect) belief that his role 3 counterpart will not pick the matching object, in particular after having played that role. The similarity between behavior in roles 1 and 2 also points to a common reasoning process, one that is unlikely to involve beliefs about others. We investigate in Appendix A3 whether non-equilibrium players follow a discernible strategy. We show that D_0-types consistently pick the focal, non-equilibrium object mentioned in section 2.1. In the absence of a dominant strategy, they are misguided into choosing a suboptimal strategy. This pattern is less pronounced for types R or O.

In Appendix A4, we report a moderate but statistically significant increase in equilibrium behavior between the first and second half of the trials, especially in role 1. A fraction of subjects evolve to a more sophisticated type over time (from D_0 to D_2). This suggests that, despite the absence of feedback, playing in a certain role (e.g., role 3) helps some subjects better understand how to play in other roles (e.g., role 2).

4 Other populations

As mentioned in section 2.1, we studied the same paradigm in two other populations, KING and YOUNG, which are discussed in this section. Table 2 provides a summary of all 610 participants in the three experiments.

Version	simplified			regular										regular		
Population	YOUNG			MAIN										KING		
				LILA									USC			
Grade	K	1	2	3	4	5	6	7	8	9	10	11	U	6	7	8
# subjects	38	37	42	34	17	20	41	27	18	25	23	29	60	110	46	43

Table 2: Summary of participants by population and grade

4.1 Thomas Starr King Middle School (KING)

We had the opportunity to conduct the same experiment with 199 students from the Film & Media Magnet at Thomas Starr King Middle School ("KING").[13] The backgrounds in KING and LILA are extremely different, even though the schools are located less than one mile apart. The majority of students at KING are Latino (about 60%) and have low socioeconomic status (about 85% are economically disadvantaged). A minority of students end up going to college (typically the local community college). Furthermore, differentiation is a core component of public education and KING offers advanced classes to high achieving students as well as special education programs for student with special needs. We use those class criteria to classify students into tracks. The two schools also differ in the curriculum taught (bilingual in LILA vs. monolingual in KING) and in class size (less than 20 students per class at LILA while often reaching 35 at KING, except in special education classes). Table 3 reports the number of participants from KING by grade and track.

	Challenged	Regular	Honors
6^{th}	21	47	42
7^{th}	46	0	0
8^{th}	43	0	0

Table 3: Number of participants at KING by grade and track

Some clarifications are in order. First, the "challenged" track is a mix of children with mild learning disabilities (dyslexia, problems focusing, etc), special needs (english learners) or who are at academic risk (low GPA). The "honors" track comprises children with a GPA significantly higher than their peers. Second, each teacher decided whether to allocate a time for the experiment or not. Among those who allowed their students to participate, we obtained consent from a high fraction of children (65%). We did not

[13] The school hosts two other magnets that were not involved in the study.

have the opportunity to work with teachers in some specific grades and tracks. Finally, we employed the *exact* same protocol as with LILA students from 6^{th} to 11^{th} grade, including classroom layout, instructions, and payment method.

Given the available populations, we conduct two studies separately. First, we analyze the behavior in the three tracks of 6^{th} graders at KING (challenged, regular and honors), and compare them to the LILA 6^{th} graders of our previous sample. Second, we compare the evolution through middle school (6^{th}, 7^{th}, 8^{th}) of LILA students and those in the challenged track at KING.

4.1.1 Comparison of behavior between 6^{th} graders at KING and LILA

Figures 5a and 5b report the proportion of equilibrium choice by track and role, and the distribution of types across tracks.

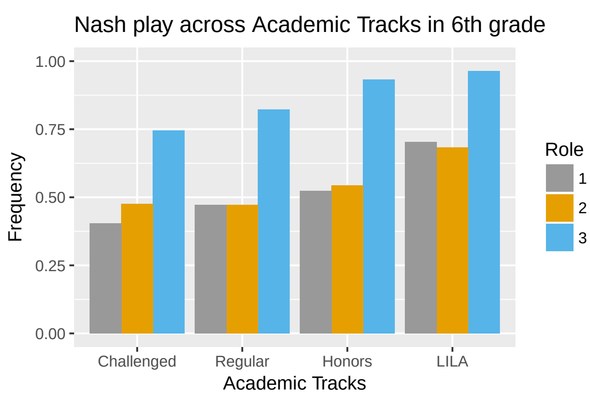

(a) Equilibrium choices by track

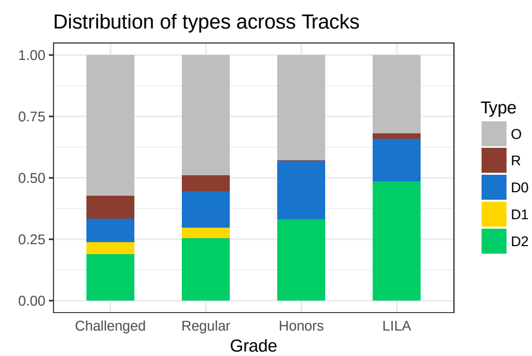

(b) Types by track

Figure 5: Equilibrium behavior and type classification of 6^{th} graders by track

As we can see from Figure 5a, in all tracks the proportion of equilibrium choices is the same in roles 1 and 2 and statistically higher in role 3. Comparing across tracks, we notice that in roles 1 and 2 there are no statistical differences within the KING population and all three tracks perform significantly worse than LILA subjects. In role 3, honors and LILA perform better than challenged and regular.

Figure 5b confirms these results. There are no statistical differences in types across KING tracks. By contrast, the proportion of D_2 is significantly higher in LILA than in the challenged and regular tracks ($p < 0.05$) but not significantly different in LILA than in the honors track. Perfect Nash players are also more prevalent in LILA (37%) than in the honors (14%), regular (9%) or challenged (5%) tracks of KING ($p < 0.04$).

Overall, 6^{th} graders at KING behave more like 3^{rd} or 4^{th} graders at LILA, with honors students, slightly outperforming the other tracks within the school.

We finally conduct a similar OLS regression as in columns 1, 2 and 3 of Table 1, except that we remove the *Age* variable (since all the subjects are 6^{th} graders) and include dummies for tracks, with "challenged" being the omitted variable. The results are summarized in Table 4.

	Role 1	Role 2	Role 3
Regular	0.069	-0.0005	0.079
	(0.093)	(0.093)	(0.049)
Honors	0.123	0.064	0.191***
	(0.095)	(0.095)	(0.050)
LILA	0.312***	0.242**	0.225***
	(0.097)	(0.098)	(0.051)
Male	0.052	0.025	0.033
	(0.060)	(0.060)	(0.031)
STEM	-0.0001	0.095	-0.001
	(0.061)	(0.061)	(0.032)
Siblings	-0.042	-0.033	0.005
	(0.078)	(0.078)	(0.041)
Constant	0.406***	0.438***	0.720***
	(0.113)	(0.113)	(0.059)
Observations	151	151	151
Adj. R^2	0.055	0.039	0.128

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 4: Regressions of equilibrium choices

LILA students perform significantly better in all roles and honors students are closer to equilibrium than the other two tracks in the easiest role 3. In this population, we find no effect of gender, siblings, or school topic preference on performance.

4.1.2 Evolution during middle school between KING and LILA

We next study the evolution during middle school (6^{th} to 8^{th} grade) for the challenged track at KING and compare it to the students at LILA. Figure 6 depicts the proportion of equilibrium choices of these two groups by role and grade.

Corroborating all previous findings, performance in both tracks and all grades is similar in roles 1 and 2 and higher in role 3. There is also no statistical change over the middle school years in either track or role. Finally and also supporting previous results, LILA students perform significantly better in all roles than those in the challenged track at

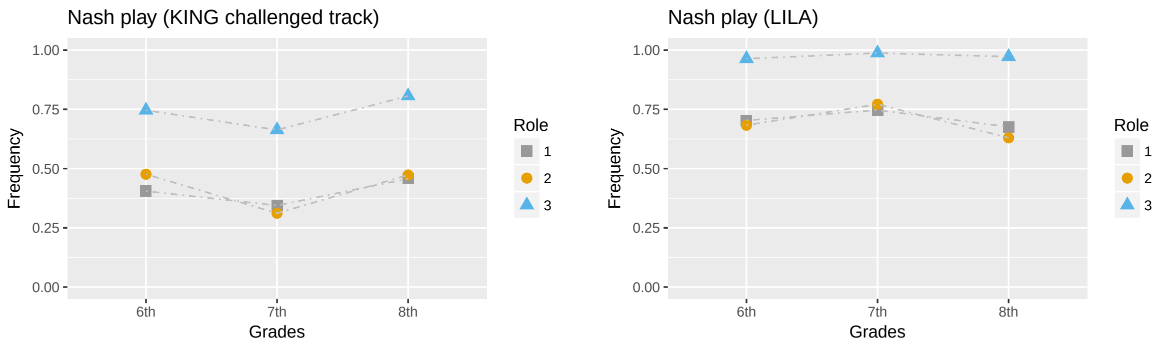

Figure 6: Equilibrium behavior in middle school (KING and LILA)

KING. In particular, the latter perform far from optimal in the simple role 3 (around 75% of equilibrium play) and only slightly better than random in roles 1 and 2.

Figure 7 compares the types in both populations.

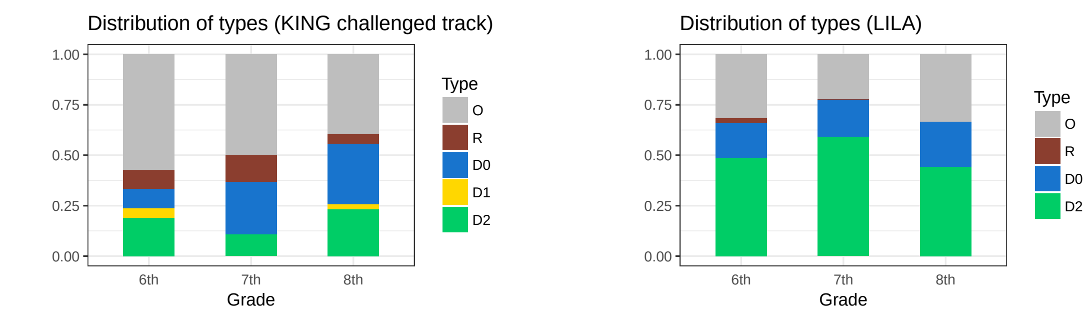

Figure 7: Proportion of KING and LILA middle school subjects by type

As expected, the fraction of D_2 subjects in KING is small. Also, and contrary to LILA, more than 50% of subjects in 6^{th} and 7^{th} grade cannot be classified in one of our four types. This means not only that many subjects play non-equilibrium actions in the simple role 3, but also that they do not play at equilibrium more often in later than in earlier roles, thereby evidencing significant confusion.

Finally, we run OLS regressions to pinpoint the determinants of equilibrium behavior in each role. The variables we use are age, a dummy for school ($LILA = 1$) and the same controls as previously. The results are presented in Table 5.

Within middle school, age is not a determinant factor of equilibrium choice in any role. LILA subjects perform drastically better than KING subjects in the challenged track. Females and those with a preference for STEM also play closer to equilibrium, but only in role 3.

	Role 1	Role 2	Role 3
Age	0.001	-0.002	0.001
	(0.002)	(0.002)	(0.001)
LILA	0.322***	0.301***	0.259***
	(0.048)	(0.051)	(0.034)
Male	0.027	0.012	-0.062*
	(0.047)	(0.050)	(0.033)
STEM	0.044	0.085	0.102***
	(0.049)	(0.052)	(0.034)
Siblings	-0.003	-0.088	0.038
	(0.064)	(0.068)	(0.045)
Constant	0.283	0.796**	0.491**
	(0.341)	(0.361)	(0.239)
Observations	196	196	196
Adj. R^2	0.175	0.156	0.238

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 5: Regressions of equilibrium choices

To sum up, the KING experiment reveals significant differences within grade levels across tracks and schools. In particular, the performance of 6$^{\text{th}}$ graders in roles 1 and 2 ranges between 40% and 70% (section 4.1.1). It suggests that important unobserved environmental factors shape the developmental trajectory of children. It also casts a warning flag to the practice of pooling data from non-homogenous schools. On the other hand, we do not find any improvement over middle school in any school and track, suggesting similar qualitative trajectories in this three-year span (section 4.1.2).

4.2 Young elementary LILA children (YOUNG)

Our initial intention was to run the experiment with all the children at LILA, starting from Kindergarten. However, during pilot testing we realized that the game was overwhelming for some of our youngest participants, due to the amount of information required to explain the game. We, therefore, decided to develop a simplified version for the youngest population at LILA. We recruited 38 subjects from K, 37 subjects from 1$^{\text{st}}$ and 42 subjects from 2$^{\text{nd}}$ grade for a total of 117 children, from now on referred to as the "YOUNG" population. We presented them a simpler game with only two players in each group (the analogue of roles 2 and 3), only two attributes (shape and color), and only 8 trials (4 in each role). Figure 8 presents a screenshot of the game.[14] While the game is simpler, the

[14]Since the youngest participants have trouble reading, we also removed all written instructions on the screen: the roles were "frog" and "owl" instead of "player x" and "player y", and the objectives "match

methods are the same: random and anonymous partners in every trial, roles changed after each trial, no feedback between trials, and focal objects (in the example of Figure 8, the red triangle is the extra object and an object in the frog's choice set).

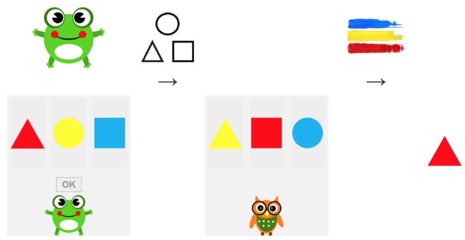

Figure 8: Screenshot for grades K, 1 and 2

We used a similar procedure as previously to classify subjects. The method is bound to be more imprecise since we have fewer observations (eight). Also, we only have the analogue of roles 2 and 3 and therefore we cannot distinguish between two and three steps of reasoning. Table 6 summarizes the classification method for this population.

	role 2	role 3
R'	random	random
D'_0	random	equilibrium
$D'_{1/2}$	equilibrium	equilibrium

random: 0-1 (of 4); equilibrium: 3-4 (of 4)

Table 6: Classification of YOUNG subjects into types

Figure 9a presents the same information as Figure 2, namely the average percentage of equilibrium behavior by grade, except that we add the YOUNG population (first three values for roles 2 and 3). Figure 9b reports the analogue of Figure 3 (proportion of types by grade) in the YOUNG population.

Equilibrium behavior is not significantly different between K and 1^{st} graders in roles 2 and 3, but they are both lower than that of 2^{nd} graders ($p < 0.02$). Choices between 2^{nd} and 3^{rd} graders are not statistically different in either role. However, remember that the protocol and number of observations are different in 2^{nd} and 3^{rd} graders, so the results are not directly comparable.

the shape" or "match the color" were replaced by graphical descriptions. These presentation changes are minor and introduced only to facilitate comprehension.

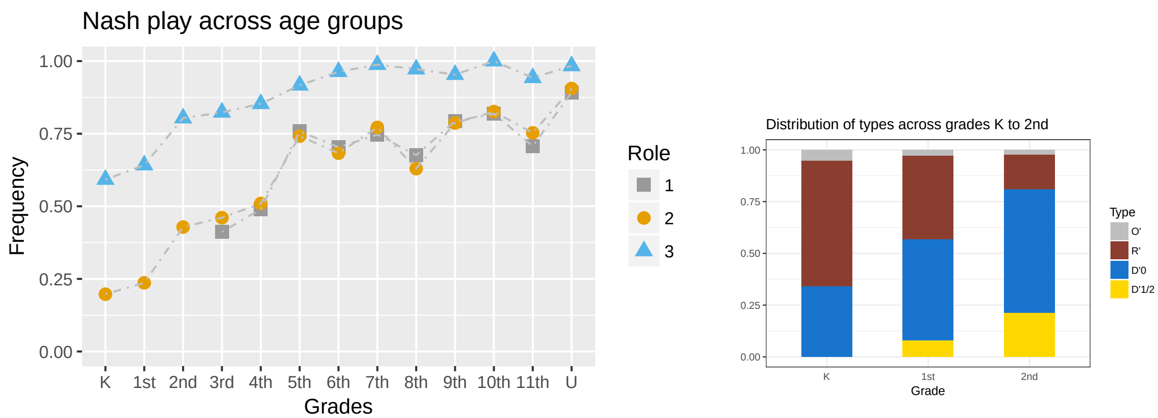

(a) Equilibrium choice by grade **(b)** Types by grade

Figure 9: Equilibrium behavior and type classification of YOUNG subjects

We also notice similar proportions of types between K and 1$^{\text{st}}$: mostly R' and D'_0, with very few equilibrium players (only 1 participant in 1$^{\text{st}}$ grade plays the equilibrium in all 8 trials). By contrast, 9 participants in 2$^{\text{nd}}$ grade (21%) are classified as $D'_{1/2}$, of which 6 participants (14%) play the equilibrium in all 8 trials. This group looks similar to 3$^{\text{rd}}$ graders (again acknowledging the difficulty to compare these two grades). Overall, the data shows severe difficulty in K and 1$^{\text{st}}$ to think beyond a dominant strategy. There is a leap in understanding in 2$^{\text{nd}}$ grade, which seems to stabilize in 3$^{\text{rd}}$ and 4$^{\text{th}}$ grade, followed by another increase in 5$^{\text{th}}$ grade.

Finally, we ran OLS regressions similar to Table 1 with the YOUNG population to better understand the determinants of equilibrium choices in roles 2 and 3. The results can be found in Table 7.

Confirming our previous findings, age is a major determinant of equilibrium behavior in both roles. Contrary to our main population, a preference for STEM is not indicative of equilibrium choice, maybe because some of these participants are too young to have clearly established inclinations.

Overall, the YOUNG experiment shows that the developmental trajectory of steps of reasoning starts early. The gap in performance between one and two steps of reasoning is preserved. Also, the level of sophistication increases with age. While Kindergartners and 1$^{\text{st}}$ graders are rarely able to play at equilibrium when more than one step of reasoning is necessary, the performance of 2$^{\text{nd}}$ graders in the comparable section of the experiment is very similar to that of 3$^{\text{rd}}$ graders.

	OLS	
	Role 2	Role 3
Age	0.011***	0.008***
	(0.003)	(0.003)
Male	0.110*	0.023
	(0.060)	(0.058)
STEM	-0.020	0.070
	(0.068)	(0.065)
Siblings	0.095	-0.010
	(0.067)	(0.064)
Constant	-0.772***	0.017
	(0.243)	(0.233)
Observations	105	105
Adj. R^2	0.150	0.058

* p < 0.1; ** p < 0.05; *** p < 0.01

Table 7: OLS regression of equilibrium choices in roles 2 and 3 for the YOUNG population

5 Conclusion

In this study, most participants either play at equilibrium or recognize only a dominant strategy. It is very unlikely that roles 1 and 2 think that their role 3 counterparts cannot solve the game. Moreover, the similar performance in roles 1 and 2 indicates that mistakes are not due to calculation limitations. If it were the case, we would observe a lower performance in the most calculation-intensive role 1. Overall, the data suggest that both under-performance and absence of intermediate levels result from a purely cognitive limitation rather than beliefs or attention-based calculation mistakes. Therefore, the ubiquitous intermediate levels reported in the literature are likely due to features that are not present in our design. They may stem from valid concerns regarding the ability of others to reach a decision, or from complex designs requiring long iterations towards the equilibrium, which are prone to mechanical (not conceptual) errors.

While equilibrium performance increases with age, there is a substantial innate component: some of our youngest participants play perfectly from the first trial whereas some of our oldest participants do not go beyond one step of reasoning. Even though there is some evidence of learning, repeated exposure is ineffective at bringing participants to play Nash. Note also that, even though dominance solvable games are relatively simple, they require the ability to form hypotheses about what others can do and use this knowledge to formulate the best course of action. This form of reasoning is abstract enough for a young child to miss. Yet, many of them perform remarkably well, echoing neo-piagetian theories of cognitive development that reject the concept of strict stages and emphasize individual differences (Morra et al., 2012).

Performance increases significantly between 8 and 12 years of age and stabilizes afterwards. The developmental trajectory suggests that most of what is needed to solve dominance solvable games is acquired by the end of elementary school. After that, children do not improve their logical abilities. Interestingly, most students acquire complex mathematical skills during adolescence. Our observations suggest that this extra knowledge does not translate into better strategic decision making.

Our data also reveals an unexpected gender difference. We do not have an interpretation for this result because sophistication in our game is a logical ability. It is hard to link this finding to the research on gender, cognition and IQ since that literature provides inconsistent results (Lynn and Irwing, 2004; Reynolds et al., 2008). At the same time, gender differences have been observed in beauty contest games as a function of context and incentives (Cubel and Sanchez-Pages, 2017). It is possible that the school environment promotes female confidence and leads them to engage in more steps of reasoning. Further research on this topic would be very enlightening.

Last, subjects with a self-reported preference for science have a significantly higher level of sophistication. Differences across schools and across tracks are also associated with differences in sophistication. Even though it is impossible to link differences across topic inclination, schools and tracks to differences in cognitive ability with a formal test, we view this result as consistent with studies showing a relationship between cognition and performance in games (Brañas-Garza et al., 2012; Gill and Prowse, 2016; Proto et al., Forthcoming; Fe and Gill, 2018). At the same time, we believe that other factors such as differences in class size, bilingualism, or underlying socio-economic variables may have also played an important role.

References

Ayala Arad and Ariel Rubinstein. The 11-20 money request game: A level-k reasoning study. *American Economic Review*, 102(7):3561–73, 2012.

Pablo Brañas-Garza, Maria Paz Espinosa, and Pedro Rey-Biel. Travelers' types. *Journal of Economic Behavior & Organization*, 78(1-2):25–36, 2011.

Pablo Brañas-Garza, Teresa Garcia-Munoz, and Roberto Hernán González. Cognitive effort in the beauty contest game. *Journal of Economic Behavior & Organization*, 83 (2):254–260, 2012.

Isabelle Brocas and Juan D Carrillo. Iterative dominance in young children: Experimental evidence in simple two-person games. *Journal of Economic Behavior & Organization*, 2018.

Isabelle Brocas, Juan D Carrillo, Stephanie W Wang, and Colin F Camerer. Imperfect choice or imperfect attention? understanding strategic thinking in private information games. *Review of Economic Studies*, 81(3):944–970, 2014.

Isabelle Brocas, Juan D Carrillo, and Niree Kodaverdian. Altruism and strategic giving in children and adolescents. USC Working Paper, 2017.

Isabelle Brocas, Juan D Carrillo, and Ashish Sachdeva. The path to equilibrium in sequential and simultaneous games: A mousetracking study. *Journal of Economic Theory*, 178:246–274, 2018.

C Monica Capra, Jacob K Goeree, Rosario Gomez, and Charles A Holt. Anomalous behavior in a traveler's dilemma? *American Economic Review*, 89(3):678–690, 1999.

Timothy N Cason and Charles R Plott. Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *Journal of Political Economy*, 122(6):1235–1270, 2014.

Jingnan Chen, Daniel Houser, Natalia Montinari, and Marco Piovesan. Beware of popular kids bearing gifts: A framed field experiment. *Journal of Economic Behavior & Organization*, 132:104–120, 2016.

Eileen Chou, Margaret McConnell, Rosemarie Nagel, and Charles R Plott. The control of game form recognition in experiments: Understanding dominant strategy failures in a simple two person "guessing" game. *Experimental Economics*, 12(2):159–179, 2009.

Miguel Costa-Gomes, Vincent P Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235, 2001.

Miguel A Costa-Gomes and Vincent P Crawford. Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5):1737–1768, 2006.

Vincent P Crawford, Miguel A Costa-Gomes, and Nagore Iriberri. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62, 2013.

María Cubel and Santiago Sanchez-Pages. Gender differences and stereotypes in strategic reasoning. *The Economic Journal*, 127(601):728–756, 2017.

Simon Czermak, Francesco Feri, Daniela Glätzle-Rützler, and Matthias Sutter. How strategic are children and adolescents? experimental evidence from normal-form games. *Journal of Economic Behavior & Organization*, 128:265–285, 2016.

Eduardo Fe and David Gill. Cognitive skills and the development of strategic sophistication. *Available at SSRN 3285484*, 2018.

David Gill and Victoria Prowse. Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, 124(6):1619–1676, 2016.

Eric J Johnson, Colin Camerer, Sankar Sen, and Talia Rymon. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47, 2002.

JA List, R Petrie, and A Samek. How experiments with children can inform economics. Working paper, 2018.

Richard Lynn and Paul Irwing. Sex differences on the progressive matrices: A meta-analysis. *Intelligence*, 32(5):481–498, 2004.

Sergio Morra, Camilla Gobbo, Zopito Marini, and Ronald Sheese. *Cognitive development: neo-Piagetian perspectives*. Psychology Press, 2012.

Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, personality, and gains from cooperation in repeated interactions. *Journal of Political Economy*, Forthcoming.

Matthew R Reynolds, Timothy Z Keith, Kristen P Ridley, and Puja G Patel. Sex differences in latent general and broad cognitive abilities for children and youth: Evidence from higher-order mg-macs and mimic models. *Intelligence*, 36(3):236–260, 2008.

Itai Sher, Melissa Koenig, and Aldo Rustichini. Children's strategic theory of mind. *Proceedings of the National Academy of Sciences*, 111(37):13307–13312, 2014.

Matthias Sutter, Claudia Zoller, and Daniela Glätzle-Rützler. Economic behavior of children and adolescents–a first survey of experimental economics results. *European Economic Review*, 111:98–121, 2019.

Appendix A. Supplementary analysis

A1. Other regression analysis

We perform some robustness checks on the evolution of equilibrium behavior with age. More precisely, in column 1 of Table 8, we present a Probit regression where the binary dependent variable is whether the individual played Nash in all 18 trials. In columns 2, 3 and 4, we report multinomial regressions where the dependent variable is the individual's type. We use D_0 as the omitted category and do not perform a regression on D_1 since it contains only one individual. In all the regressions, we include the same independent variables as in the OLS regression of Table 1.

	Probit	Multinomial regression		
	Nash	D_2	R	O
Age	0.012***	0.016***	-0.061**	-0.002
	(0.003)	(0.006)	(0.024)	(0.007)
Male	-0.350*	-0.645*	0.290	-0.079
	(0.183)	(0.385)	(0.864)	(0.418)
STEM	0.587***	1.579***	0.035	0.568
	(0.188)	(0.447)	(1.002)	(0.495)
Siblings	-0.139	-0.459	0.538	-0.425
	(0.205)	(0.461)	(1.170)	(0.498)
Constant	-2.065***	-1.436	5.657*	0.819
	(0.479)	(1.021)	(3.024)	(1.096)
Log Likelihood	-142.7	—	—	—
AIC	—	502.3	502.3	502.3

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 8: Probit and Multinomial regressions

The Probit regression supports the existing findings that age, gender and a preference for STEM are indicative of Nash behavior. The regressions on types also yield similar conclusions to the standard OLS. Compared to individuals who only recognize a dominant strategy (D_0), older subjects are more likely to play at equilibrium (D_2) and less likely to play randomly (R). A preference for STEM (and to a lesser extent females) increases the likelihood of equilibrium play but has no effect on playing randomly or the dominant strategy. There are no significant differences between D_0 and O, reinforcing the idea that although O types typically play better than random, they are not very sophisticated either.

A2. Preference for school subjects: STEM vs. Arts & Humanities

The difference in performance as a function of school preference noted in Tables 1 and 8 is both surprising and interesting. To investigate this effect in more detail, we present in Figure 10, the same information as in Figure 2, namely the proportion of equilibrium choices by grade and role, separately for subjects with a preference for STEM (86 subjects) and with a preference for Arts & Humanities (148 subjects).

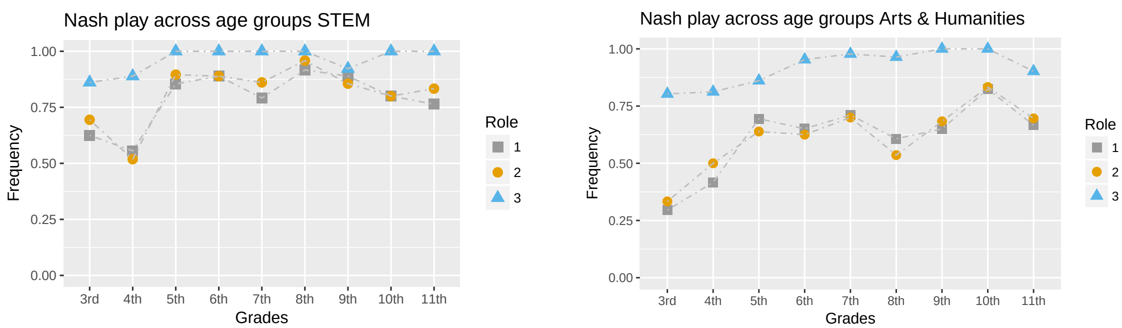

Figure 10: Equilibrium choices by grade, role and favorite school subject

The graph illustrates the difference in performance in roles 1 and 2 across topic preferences. Averaging across grades, subjects who like STEM play 77.9% and 80.4% the equilibrium action in roles 1 and 2 compared to 61.4% and 61.3% for the subjects who prefer Arts & Humanities. These differences are highly significant ($p < 0.001$).

If we run the same OLS regressions as in columns 1, 2 and 3 of Table 1 separately on each subsample, we obtain that age loses significance for subjects who prefer STEM but not for those who prefer Arts & Humanities. This captures the fact that equilibrium performance of the former is always high independent of age, with the exception of 4^{th} graders (regressions omitted for brevity). Overall, we understand that self-reported preferences partly capture intrinsic taste but they also capture self-perceived ability over topics. Also, while we tried to minimize analytical requirements to understand the game, we may not have fully succeeded. With these caveats in mind, the result nevertheless suggests that a scientific inclination is correlated with equilibrium behavior, and this holds independent of age.

A3. The behavior of non-equilibrium players

Our theory has assumed that an individual who does not unveil the logic of equilibrium play in roles 1 and 2 will choose randomly between the three options. At the same time, and as briefly mentioned in section 2.1, we have introduced non-equilibrium focal objects to minimize the likelihood of spurious equilibrium choices. In this section, we briefly study whether non-equilibrium players follow any discernible strategy.

We define the options as follows. We call *rational* the option chosen by an equilibrium player, *heuristic* the option chosen by a subject who matches the attribute of the extra object, and *alternative* the option chosen by a subject who does neither of the previous two. By definition, *rational* and *heuristic* coincide for role 3. More importantly, the construction of our "focal objects" is such that *rational* and *heuristic* are always different in roles 1 and 2.[15] This means that in roles 1 and 2, there is always one *rational*, one *heuristic*, and one *alternative* option. In Figure 11, we present the proportion of *rational*, *heuristic*, and *alternative* choices in roles 1 and 2 by subjects classified as D_0, R and O.[16]

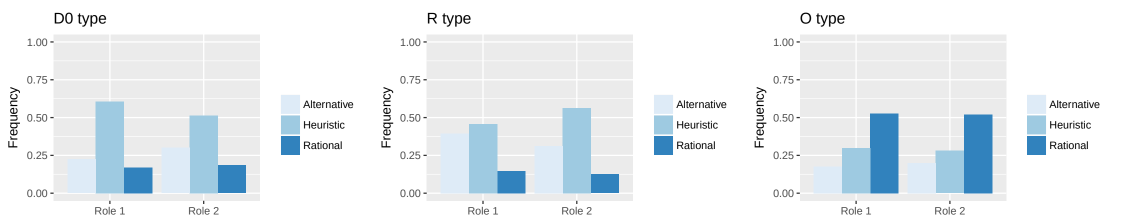

Figure 11: Options chosen in roles 1 and 2

By construction, D_0 and R subjects have chosen the *rational* option at most twice in each role whereas O subjects are likely to have chosen it three or four times (otherwise they would have been classified as one of the other types). It is, therefore, not surprising that the *rational* option is under-represented in D_0 and R and over-represented in O. Interestingly, we find that all the types choose more often the *heuristic* than the *alternative* option in both roles. These differences are significant for D_0 ($p < 0.001$) and to a lesser extent for O in both roles ($p < 0.04$). One interpretation of this finding is that subjects who recognize a dominant strategy and only a dominant strategy (D_0) erroneously apply the same logic to other roles as well. Subjects who are less (R) or more (O) sophisticated than D_0 are less prone to this mistake.[17]

A4. Learning

Many of our participants did not grasp the backward induction logic from the outset. However, after playing a few times in different roles, they may have used their behavior in

[15] For example, in Figure 1 the *heuristic* option in roles 1 and 2 is the blue square A, which coincides with the extra object and is different from the *equilibrium* options.

[16] We do not include D_2 subjects since, by definition, they have chosen 5 or 6 times out of 6 the *rational* option in both roles.

[17] Notice that in all roles and ages a *rational* choice is more likely than a *heuristic* choice (this is obvious in grades 5$^{\text{th}}$ and above but it is also true in 3$^{\text{rd}}$ and 4$^{\text{th}}$ grade). So, for all roles and ages, even an individual who correctly anticipated the empirical behavior of the age group would find it optimal to play the equilibrium strategy.

a certain role to deduce what to do in another role. For example, after playing in role 3, they may have understood the dominant strategy in that role and used it to best respond in role 2. If a significant fraction of subjects are in this category, a classification method based on the entire game may be misleading or incomplete.

To address changes in behavior during the experiment, we present in Figure 12 the fraction of equilibrium choices by grade and role (as in Figure 2), separated between the first and last nine trials of the game.

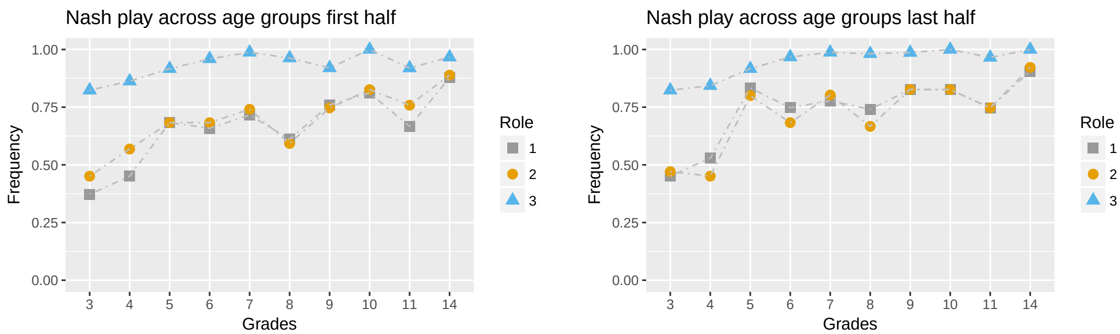

Figure 12: Equilibrium choices in first (left) and last (right) half of trials

We notice a small but sustained increase in equilibrium behavior for roles 1 and 2, with the exception of 4^{th} graders. Averaging across all school-age subjects, individuals play in roles 1, 2 and 3 the equilibrium action 68.5%, 71.6% and 93.6% of the time in the first half of the experiment compared to 75.5%, 74.2% and 95.4% in the second half. The difference is highly significant for role 1 ($p < 0.001$) and marginal for roles 2 ($p = 0.075$) and 3 ($p = 0.052$).

We next perform a similar classification exercise of types as previously, separately in each subsample. With fewer observations, the classification is bound to be more inaccurate. We labeled the behavior in half of the trials of a role "equilibrium" if all 3 observations were consistent with theory and "non-equilibrium" otherwise (0, 1 or 2 out of 3 observations consistent with theory). For each subsample $X \in \{F, L\}$ (where F is the first half and L is the last half), we considered the same types as before: R^X, D_0^X, D_1^X and D_2^X. The remaining subjects are denoted O^X. Table 9 reports the type of subject in the first and last half of the experiment. For this analysis, we focus on LILA population.

The type of two-thirds of our subjects does not change between the first and last half of the trials. Among those who change types, 42% learn to play the equilibrium in the second half (32 subjects) for 8% who play equilibrium in the first half but not the second (6 subjects). Also, 18% improve from R^F to D_0^L (14 subjects) for 8% who decrease from D_0^F to R^L (6 subjects).

	O^L	R^L	D_0^L	D_1^L	D_2^L
O^F	3	1	2	0	7
R^F	2	12	14	0	0
D_0^F	9	6	48	2	13
D_1^F	0	0	2	3	12
D_2^F	3	0	0	3	92

Table 9: Types of school-age subjects in first (F) and last (L) half of trials.

Overall, there is some small evidence of change in behavior during the experiment, and it predominantly takes the form of learning to play closer to equilibrium.

Appendix B. Instructions and Quiz

B1. Instructions

Hi, everyone. Today we are going to play a few games. In all the games, you will earn points that will be placed in your virtual wallet.

[For subjects in grades 6 and above] At the end of the experiment you will be paid 1 cent for each point you obtained with an Amazon gift card. You will get several hundred points, so you will be able to get a nice gift card.

[For subjects in grades 3, 4 and 5] At the end of the experiment we will go to the toy shop and you will be able to buy the toys you like with the points you earned.

In all the games, you will play through the tablets. We ask you to not talk and keep your decisions private.

This game is called the "matching game." In this game, you will be playing many times. Each time, you will be playing in groups of 3. The computer will decide with whom you play and you will not know who that is. If you are player 1, you will see a screen like this.

[SLIDE 2]

At the top of the screen, it tells you are player 1. There are 3 large grey pictures on the screen. Yours is the darkest. On this picture, you can read "YOU ARE Player 1". There are 3 objects on this picture. Each object is a colored shape that is marked with a letter. Shapes, colors and letters are all different. You have to select one object by clicking on it and pressing OK. On your screen, you can also see the objects on Player 2's picture and the objects on Player 3's picture. There is also one object outside the 3 pictures.

If you are Player 2, you will see a screen like this.

[SLIDE 3]

This is the same screen as for Player 1 except that your picture is the darkest one in the middle where you can read "YOU ARE Player 2." If you are Player 3, you will see a screen like this.

[SLIDE 4]

Again, this is the same screen as for Players 1 and 2 except that your picture is the darkest one in the middle where you can read "YOU ARE Player 3." All right, now, how do you obtain points?

[SLIDE 5]

In this game Player 1 has to choose an object that has something in common with the object Player 2 chooses. The arrow between the picture of Players 1 and 2 tells you what they need to have in common. In this example, Player 1 needs to choose an object that has the same letter as the object chosen by Player 2. Now what about Player 2?

[SLIDE 6]

Player 2 has to choose an object that has something in common with the object that Player 3 chooses. The arrow between the picture of Players 2 and 3 tells you what they need to have in common. In this example, Player 2 needs to choose an object that has the same color as the object chosen by Player 1. What about Player 3?

[SLIDE 7]

Player 3 has to choose an object that has something in common with the object that is outside the pictures. In that example Player 3 needs to choose an object that has the same shape as the object outside the pictures. Each time you play, you will know what each player needs to do to win because you all see the same screen. However, when you make a choice, you do not know what objects the others have chosen. If you choose the object that matches what you are asked to match, 40 points will be added to your wallet. If you miss, only 20 points will be added. Is it clear for everyone?

Remember, you will play several times. Sometimes you will be player 1, sometimes you will be player 2 and sometimes you will be player 3. Each time you play, you will play with different people. Also, the shapes, colors, letter and characteristics that you need to match will change. We will not tell you how much you earned each time you played. We will only tell you how many points you have earned in total at the end of the game.

Everybody understands? Let's answer some questions, just to make sure everybody understands. Look at the screen here.

[SLIDE 8]

On your computers are some questions that you have to answer correctly before we start the game. If you need help with the questions, raise your hand and we will come to assist you.

Are you ready to start the game? Remember, you will play several times, always with different partners. The roles, the objects and what you need to match will change. We will tell you at the end how many points you obtained.

Figure 13: Slides 1-8 projected in the main screen

B2. Quiz

Below are the questions that each subject would see on their computer screen. All questions had to be answered correctly before the paid part of the experiment could start.

Figure 14: Quiz included before the paid part of the experiment