The development of randomization and deceptive behavior in mixed strategy games [*]

Isabelle Brocas
University of Southern California and CEPR

Juan D. Carrillo
University of Southern California and CEPR

August 2021

Abstract

We study the foundations for the development of optimal randomization in mixed strategy games. We consider a population of children and adolescents (7 to 16 years old) and study in the laboratory their behavior in a non-zero sum, hide-and-seek game with a unique interior mixed strategy equilibrium where each location has a known but different value. The vast majority of participants favor the high-value location not only as seekers (as predicted by theory) but also as hiders (in contradiction with theory). The behavior is extremely similar across all ages, and also similar to that of the college students control adult group. We also study the use of cheap talk (potentially deceptive) messages in this game. Hiders are excessively truthful in the messages they send while seekers have a slight tendency to (correctly) believe hiders. In general, however, messages have a small impact on outcomes. The results point to a powerful (erroneous) heuristic thinking in two-person randomization settings that does not get corrected, even partially, with age.

Keywords: developmental decision making, laboratory experiment, mixed strategy, randomization.

JEL Classification: C72, C93.

[*]We thank the members of the Los Angeles Behavioral Economics Laboratory (LABEL) for their insights and to Chris Crabbe for his outstanding programming skills. We are grateful to the staff of the Lycée International de Los Angeles –in particular Emmanuelle Acker, Nordine Bouriche, Mathieu Mondange and Anneli Harvey– for their help and support in running the experiment in their schools. The study was conducted with the University of Southern California IRB approval UP-12-00528. We also acknowledge the financial support of the National Science Foundation grant SES-1851915. 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

The ability to make choices that are unpredictable to other people is an essential component of optimal decision making in a myriad of strategic contexts where players have conflicting interests (sports, politics, corporations, etc.). Von Newman's minimax theorem for zero-sum games and, more generally, the mixed strategy Nash equilibrium are cornerstones of game theory. The two pillars of the theory state that payoffs should be equal across strategies and choices should be serially uncorrelated. Unfortunately, the beauty and crispness of these theoretical predictions contrasts with the empirical behavior. As demonstrated in the early experimental literature, departures are frequent and with no obvious systematic patterns across games (see e.g. Rapoport and Boebel (1992); Mookherjee and Sopher (1994); Ochs (1995) and the review in Camerer (2003) chapter 3).[1] The subsequent literature has taken two main routes. Some researchers have used existing behavioral models, such as quantal response equilibrium and cognitive hierarchy, to explain the observed deviations (McKelvey and Palfrey, 1995; Goeree et al., 2003; Crawford and Iriberri, 2007). A second strand has investigated differences in behavior between regular subjects and expert randomizers (Walker and Wooders, 2001; Chiappori et al., 2002; Palacios-Huerta, 2003).

In this paper, we follow a third avenue. We provide a foundational study of the ability of individuals to play mixed strategies by investigating the evolution of behavior from a very young age into adulthood. Indeed, while the literature has analyzed unpredictability and optimal alternation of actions by children in individual decision making contexts (Rabinowitz et al., 1989; Towse and Mclachlan, 1999), randomization by children in a game of strategy has not been studied.[2]

A mixed strategy game is an interesting candidate for a developmental study. The game is easy to understand. Optimal play is neither trivial nor extraordinarily difficult. Also, deviations by adults are frequent–yet not omnipresent–and follow systematic patterns. Finally, and despite some advancements (Goeree and Holt, 2001), the fundamental reasons for the difficulty to play optimally are still imperfectly understood. Studying the evolution with age can help determine what makes these apparently simple situations so elusive.[3]

[1] In one of the earliest experiments, O'Neill (1987) concludes that mixed strategy Nash equilibrium cannot be rejected, but Brown and Rosenthal (1990) challenged the conclusion. McCabe et al. (2000) also finds support for the theory in their experiment.

[2] Despite some recent interest in experimental economics to study decision making by children and adolescents, the literature has predominantly focused on individual decision making (see the recent surveys by Sutter et al. (2019) and List et al. (2018)). Some notable exceptions of game theoretic experiments with children include Murnighan and Saxon (1998); Harbaugh and Krause (2000); Sher et al. (2014); Czermak et al. (2016); Chen et al. (2016); Fe et al. (2020); Brocas and Carrillo (2020b, 2021).

[3] We also modified the game to facilitate understanding and learning, which we thought would increase

There are competing hypotheses regarding the developmental trajectory. First, the ability to randomize optimally may be acquired, but only slowly and imperfectly with age. This would mean that young children perform worse than adults, and that the proportion of equilibrium behavior has a positive trend. Alternatively, randomization may be equally difficult at all developmental stages, with no significant improvement with age. Finally, randomization could be an innate ability lost during childhood or adolescence. While the last hypothesis is the least natural, it is worth noting that changes occurring in the brain during adolescence involve losing neural connections deemed non-essential (Giedd, 2008). During this pruning period, adolescents lose skills that have not been used frequently as well as the ability to develop new skills (e.g., speaking a language without an accent). It is therefore possible that abilities that have not been reinforced fade during pruning episodes. Finally, we also incorporate a message game and analyze if the possibility to "trick" the opponent has a significant effect on outcomes.

To study these questions, we propose a novel setting that combines and extends properties from previous games. We also propose a novel experimental design, accessible to young individuals (7- to 16-year-old in our main population and as young as 5-year-old in our shortened version). We develop a non-constant sum, three-action, hide-and-seek game, where locations have known but *different* values.[4] This small modification has important implications. Since action probabilities are determined in a way that the opponent is indifferent between their actions, in equilibrium the seeker should choose more frequently the high-payoff location whereas the hider should choose less frequently the high-payoff location. This means that, the high-payoff location–presumably focal for both roles–coincides with the equilibrium prediction in one role (seekers) but not in the other (hiders). Focality, however, is purely based on payoffs rather than on psychological considerations as in Rubinstein et al. (1997) for example. Notice that a similar property of intuitive vs. counterintuitive randomization was pointed out in Goeree and Holt (2001)'s abstract mixed strategy game. We study if a similar tendency towards the intuitive randomization is also present in our game for some (or all) of our age groups. Importantly, since our participants play in both roles and obtain feedback, we determine whether practical experience affects their dynamic behavior. For the design, we propose a story-based narrative with a graphical presentation of choices and feedback.

The experiment consists of three tasks. In the basic hide-and-seek game (task 1), participants play 12 rounds with the same partner alternating roles and without feedback.

equilibrium compliance by adults. In particular, the presentation was dramatically simplified, participants gained experience in both roles, and they obtained feedback either between or within tasks.

[4]Otherwise, standard hide-and-seek rules apply: when the hider chooses the location of value n, the payoffs of hider and seeker are 0 and n if the seekers finds the location and n and 0 otherwise.

This allows us to infer the frequency of each action for each individual and contrast the empirical choices with the theoretical predictions. In the first extension (task 2), participants play the very same game with a new opponent, except that hiders send a cheap talk message regarding the location chosen. Seekers observe the message before making their choice. Given the opposing interests, the hider's message should be uninformative, and therefore ignored by the seeker. In the second extension (task 3), participants play 12 rounds in one role, then switch and play 12 rounds in the other role, always with the same partner. Participants receive feedback after each round, so we can determine if the dynamics of choices exhibit some trends.

We obtain the following results. With respect to choices (tasks 1, 2 and 3), we notice that every participant consistently randomizes between the alternatives. Almost no subject chooses the high-value option less than 25% or more than 75% of the time in any role. While randomization may seem obvious for educated adults, it could, in principle, elude young children. We show that this is not the case. At the same time, the vast majority of individuals favor the high-value location not only as seekers (as predicted by theory) but also as hiders (in contradiction with theory). The behavior is extremely similar across all age groups, from 2nd graders to 10th graders (and even kindergartners in our shortened version). It is also similar across the three tasks. Adults play somewhat closer to theory as seekers but their behavior is equally opposite to equilibrium as hiders. Overall, there is a powerful tendency to randomize, with an added extra weight on the high-value location.

Regarding the message game (task 2), all our participants also realize the importance of alternating between truthful and deceptive messages. Hiders are excessively truthful, especially in the younger age groups while seekers have a slight tendency to (correctly) believe hiders. In general, messages have a very minor impact on payoffs. Indeed, seekers are somewhat gullible but hiders fail to exploit this bias. On the other hand, hiders are mostly truthful but so incorrect in their hiding patterns that their bias in messages is, in comparison, of very little importance.

Finally, the dynamic choice with feedback (task 3) reveals that choices are serially correlated. More precisely, we observe an excessive tendency to alternate between locations in both roles, especially by participants in 7th grade and below. Again, this violation has a very limited effect on payoffs.

To sum up, we observe some minor differences across ages in the provision and interpretation of messages as well as in the reaction to feedback. As for choices, individuals in all age groups show intent but also inability to be unpredictable. They all randomize but they do so in a suboptimal way: seekers favor the high-value location less than they should and hiders favor the high-value location even though they should not. These deviations have very significant payoff consequences. The similarity in behavior at all ages is striking. We are not aware of other developmental game theoretic studies where the mistakes committed are so remarkably similar in direction and magnitude from early childhood all the way to the peak cognitive performance age. Indeed, with experience and the development of cognitive faculties, individuals typically improve performance in strategic situations involving high-order reasoning (even though they sometimes reach a ceiling during adolescence, as for example in Brocas and Carrillo (2021)). Our results point to a fundamental impediment (perhaps biological) to randomize optimally. It also suggests the existence of a powerful erroneous heuristic that governs behavior. This default choice is ingrained from a very young age. According to our results, it is difficult to overcome even after experiencing both roles, obtaining feedback, encountering different partners and, more generally, becoming older and therefore more cognitively mature individuals. The behavioral models we discuss (social efficiency and noisy introspection) offer some partial but not fully satisfactory explanations for the observed departures.

The paper is organized as follows. In section 2, we lay out the experimental design and procedures, with special emphasis on how we adapt it to our population. In section 3, we present the basic theory and predictions. In section 4, we analyze randomization strategies and payoffs by age. In section 5, we discuss the use and interpretation of deceptive messages. In section 6, we investigate dynamic trends and serial correlation. In section 7, we offer some remarks and directions for future research. Additional analyses and the choices of children from kindergarten and first grade can be found in the Appendix.

2 Experimental design

We conduct a variant of a two-person, hide-and-seek game with children and adolescents. Working with young participants poses a number of methodological challenges. In Brocas and Carrillo (2020a), we provide a review of the main difficulties as well as some general guidelines to address them. We closely follow these principles in this paper.[5]

Participants. We recruited 228 children and adolescents from 2nd to 10th grade (7 to 16 years old) studying at the Lycée International de Los Angeles (LILA), a French-English bilingual private school in Los Angeles. For comparison, we ran the same experiment with 34 young adult students (A) at the University of Southern California. Table 1 reports the distribution of our 262 participants by grade and approximate age. We also ran a shortened version of the experiment with 42 Kindergartners and 34 first graders from the

[5]In short, it is important (i) to adapt the length and procedures to a population with limited attention, (ii) to offer age-appropriate incentives, (iii) to present the task in a way that subjects are not required to possess strong analytical skills to participate, and (iv) to describe the characteristics of the children population and include a benchmark adult comparison whenever possible.

same school. The analysis of this population can be found in Appendix B.

	LILA								USC
Grade	2	3	4	5	6	7	8	10	A
Age	7-8	8-9	9-10	10-11	11-12	12-13	13-14	15-16	18+
# subjects	23	26	30	27	20	40	34	28	34

Table 1: Participants by grade and location.

Most experiments with children do not include an adult population. We believe it is important to include an adult control group using *identical* procedures to establish a behavioral benchmark (Brocas and Carrillo, 2020a). However, it is also key to recognize its limitations. In our case, most students at LILA are from caucasian families of upper-middle socioeconomic status. After high school, they go to well-ranked colleges in Europe and North America (including USC and schools in the UC system). It is therefore an imperfect but reasonable match for the adult population, despite some differences in individual characteristics, such as nationality, family background or peer group size.

Notice also that our population is homogenous. As shown in previous research (Brocas and Carrillo, 2021; Charness et al., 2019), demographic and socioeconomic factors have a strong influence on the behavior of children. Pooling participants of different ages from different schools is likely to introduce confounds that mask any developmental trajectory. We avoid this problem by recruiting children from the same school, with similar social and economic backgrounds, and who follow the same curriculum.

Games. The experiment consisted of two games always performed in the same order. We started with two rounds of a graphical version of the two-person beauty contest, a game with a strictly dominant strategy which we called the “Army game”. After a short break, we ran the hide-and-seek game, which we called “Farmers and Pirates”. In this paper we focus exclusively on the second game and relegate the analysis of the first game to Brocas and Carrillo (2020b). Given that the two games focus on very different paradigms, we have no reason to expect that outcomes in one game could affect choices in the other. In any case, and to avoid any possible cross-contamination, we randomly and anonymously rematched subjects between the two games.

General procedures. We ran 28 sessions at LILA and 3 sessions at USC, each with 8 to 12 subjects. Interactions were computerized through touchscreen PC tablets. Tasks were programmed with the open source software ‘Multistage Games.’[6] Sessions with school-age students were run in a classroom at the school whereas adult sessions were run at the Los

[6]Downloading instruction can be found at `http://multistage.ssel.caltech.edu`.

Angeles Behavioral Economics Laboratory at USC. Procedures were identical, except for payments, as explained below. For school-age sessions, we tried to have males and females from the same grade, but for logistic reasons sometimes had to mix participants of two consecutive grades. Students from grades 9, 11 and 12 did not participate in the study because they were taking or preparing for French or US national exams.

Farmers and Pirates. Participants played a three-action ($n \in \{3, 4, 8\}$), non-zero sum, hide-and-seek game. Table 2 provides its normal-form representation.

		Seeker		
		3	4	8
	3	(0,3)	(3,0)	(3,0)
Hider	4	(4,0)	(0,4)	(4,0)
	8	(8,0)	(8,0)	(0,8)

Table 2: Normal form representation of the hide-and-seek game

In studies with children and adolescents, it is crucial that differences in behavior reflect as much as possible developmental differences in strategic thinking as opposed to developmental differences in attention, numerical skills, mathematical skills, or capacity to understand rules (Brocas and Carrillo, 2020a). As a result, it is imperative to avoid the abstract, formal presentations typical in experiments with adults such as, for example, the normal form representation described in Table 2. We therefore developed for our experiment a simple, graphical interface with an engaging story, which is accessible and appealing to children as young as 5 years of age. We call it the "Farmers and Pirates" game. Participants played three variants of that game, called Tasks 1, 2 and 3.

Task 1. Participants were randomly and anonymously matched in pairs. Each individual in a pair was assigned the role of "Farmer" or "Pirate". The Farmer secretly chose one of 3 islands to grow a flower. The three islands differed in n, the number of flowers that would grow in that island, with $n \in \{3, 4, 8\}$. After the Farmer had made his choice, the Pirate had to decide which island to raid. Whenever the Pirate chose the island where the Farmer planted the flower, he would steal the n flowers. Otherwise, the Farmer would keep the n flowers. The number of flowers that each island could grow was common knowledge and displayed on the screen at all times. Participants played 12 rounds of Task 1 with the same partner, *alternating roles and without feedback* between rounds. At the end of Task 1, participants obtained a summary of information with the number of times they had won as a Farmer and as a Pirate and the total number of points (flowers) accumulated.[7]

[7] We cannot overemphasize the importance of an interactive interface with an intuitive design and an

Figures 1a and 1b present screenshots of Task 1 from the perspective of the Farmer and the Pirate, respectively.[8] In this example, the Farmer planted the flower in island 4 (flower image) and the Pirate incorrectly raided island 3 (red circle).[9] Figure 1c describes the feedback at the end of Task 1.[10]

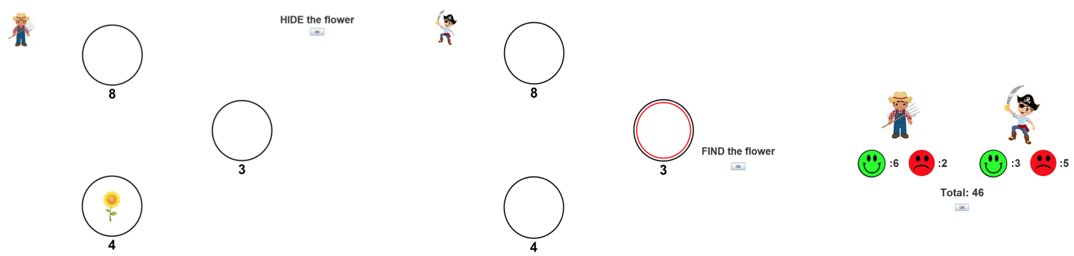

(a) Task 1. Farmer **(b)** Task 1. Pirate **(c)** Task 1. Feedback

Figure 1: Screenshots of Task 1 in the Farmers and Pirates game

Task 2. After receiving the summary feedback, we randomly formed new groups and each participant in a pair was again assigned the role of Farmer or Pirate. Participants would then play Task 2, which was identical to Task 1 with one exception: after choosing the island where the flower was grown, the Farmer had to send a cheap talk message to the Pirate, which consisted in highlighting one island. This message was displayed on the screen of the Pirate, who could decide to look for the flowers in the island pointed by the Farmer or anywhere else. Since we did not want to attach any negative connotation to a misrepresentation of the choice, we explicitly told participants that, as Farmers, they could highlight the island where they actually grew the flowers or any other island. Similarly, as Pirates, they could look for the flowers in the island highlighted by the Farmer or in

attractive story both to facilitate comprehension and maximize attention. In the instructions (Appendix C), we included two practice rounds (called "pretend matches") to familiarize the participants with the rules. We believe that most subjects understood the rules from the outset and virtually everyone figured them out after a few rounds of Task 1. In our informal conversations after the experiment, no children seemed confused. This is rather expected given that we framed the interaction literally as a hide-and-seek game, a recreational activity that most participants have played since toddlerhood.

[8]Unlike Rubinstein et al. (1997), we were not interested in psychological considerations for the hiding and seeking strategies. We avoided focal locations using a circular display with islands located equidistantly. We randomly switched the locations of islands (e.g., at 1, 5 and 9 o'clock vs. at 2, 6 and 10 o'clock, etc.) as well as the location of the high- and low-value islands.

[9]To implement their choice, participants had to tap with their finger on the desired island. They could (and often would) change their mind and switch islands as many times as they wanted. A decision was final only after pressing the "OK" button.

[10]It corresponds to a test session with 16 instead of 12 rounds. Success is captured with a green smiling face and failure with a red frowning face.

any other one.[11] Figures 2a and 2b present screenshots of Task 2 from the perspective of the Farmer and the Pirate. In this example, the Farmer planted the flower in island 3 (flower image) and highlighted island 4 (blue circle). The Pirate observed that island 4 was highlighted (blue circle) but correctly raided island 3 instead (red circle).

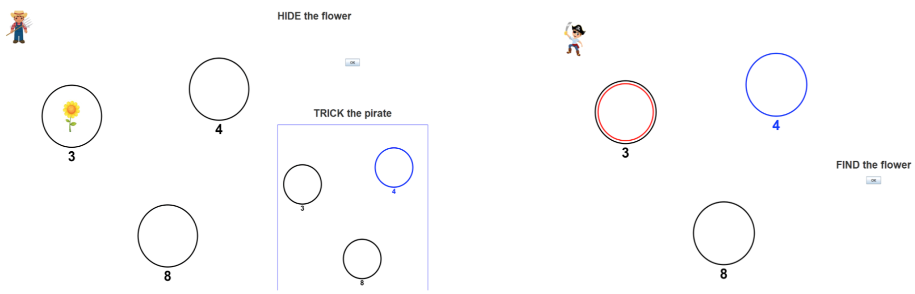

(a) Task 2. Farmer

(b) Task 2. Pirate

Figure 2: Screenshots of Task 2 in the Farmers and Pirates game

Task 3. After receiving the summary feedback from Task 2 (which took the same form as the feedback in Task 1), participants moved to Task 3, where new groups were randomly formed and each participant in a pair was again assigned the role of Farmer or Pirate. Task 3 was identical to Task 1 with three exceptions. First, participants played 24 rounds instead of 12. Second, instead of alternating roles, they played 12 rounds in one role then switched and played 12 rounds in the other role, always with the same partner. Third, they received feedback after each round instead of only at the end of the task: their choice of island and whether they won or not.[12] This information would fill up as rounds progressed. Figures 3a and 3b present screenshots of Task 3 from the perspective of the Farmer and the Pirate. In this example, participants are currently in the fourth round. The Farmer has successfully hidden the plants in islands 8 and 3 in the first two rounds whereas the Pirate has successfully raided island 4 in the third round.[13]

Summing up, Task 1 studies the basic problem of randomization in asymmetric hide-and-seek games, Task 2 adds the effect of cheap talk messages, whereas Task 3 incorporates feedback. A sample of instructions is included in Appendix C.

Notice that our set up does not provide as much data as some traditional experiments (e.g., Ochs (1995)) where mixed strategies are elicited in each round (instead of inferred

[11]In fact, the Farmer's decision was framed as "TRICK the pirate" to make clear that the message could but did not have to be truthful.

[12]While it is typically preferable to introduce only one change at a time, we believed that feedback in Task 3 would be more salient if participants did not alternate roles.

[13]To avoid overwhelming participants with information, we decided not to present the choice of the other person, and only reported success (green smiling face) or failure (red frowning face).

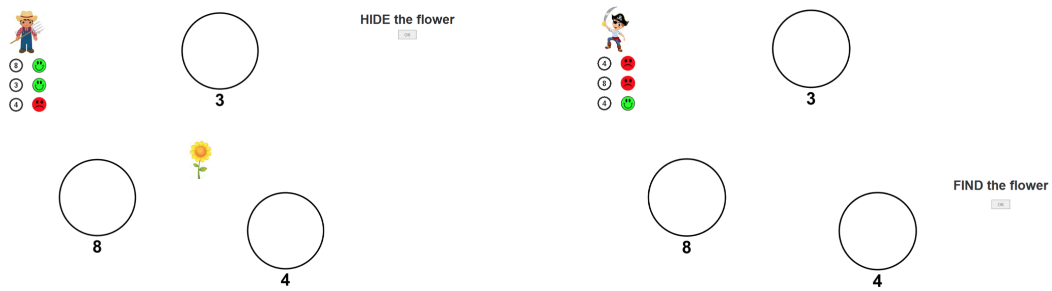

(a) Task 3. Farmer

(b) Task 3. Pirate

Figure 3: Screenshots of Task 3 in the Farmers and Pirates game

from the behavior across rounds) and participants play significantly more times. This loss of valuable information decreases the statistical power of our analysis but it is a necessary compromise given the characteristics of our population.[14]

Payments. During the experiment, subjects accumulated one point per flower. Rather than equalizing payment across ages, we calibrated the incentives to account for differences in marginal value of money (see Brocas and Carrillo (2020a)). To this purpose, we implemented three different payment rules. USC and LILA students from grade 6 and above had points converted into money at the rate of $0.08 and $0.06 per point respectively, paid immediately at the end of the experiment with an amazon e-giftcard. USC students also received a $5 show-up fee. For elementary school participants (grades 2 to 5 as well as K and 1 in the shortened version), we set up a shop with 20 to 25 pre-screened, age appropriate toys and stationary.[15] Different toys had different point prices. Before the experiment, children were taken to the shop and showed the toys they were playing for. They were instructed about the point price 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.[16] Changing the medium of payment is unusual in economics. However, we think it is key in order to try and homogenize within our population the value of the rewards

[14]Despite the new, highly creative approaches to elicit probabilistic choices (Fragiadakis et al., 2019), we believe they are still excessively hard to grasp for young participants.

[15]These included gel pens, bracelets, erasers, figurines, die-cast cars, trading cards, fidget spinners, apps, calculators and earbuds among others.

[16]The procedure emphasizes the importance of accumulating points while making the experience enjoyable. Children are familiar with the method of accumulating points that are subsequently exchanged for rewards since it is commonly employed in fairs and arcade rooms. We spent an average of $4 in toys per child, which is arguably more than usual in economic experiments with elementary school children.

instead of the rewards themselves.

Other information. The Farmers and Pirates game lasted around 60 minutes and the entire experiment around 75 minutes. Average earnings for the Farmers and Pirates game were \$8.1 (LILA) and \$10.5 (USC), given the 33% increase in conversion rate. Total earnings averaged \$11.3 (LILA) and \$16.2 (USC), without counting the \$5 show-up fee at USC. At the end of the experiment, we asked participants to report their gender. For the data analysis, we will group the school-age students in 4 grade categories, 2-3 (ages 7 to 9, 49 subjects), 4-5 (ages 9 to 11, 57 subjects), 6-7 (ages 11 to 13, 60 subjects), 8-10 (ages 13 to 16, 62 subjects), and denote by A (34 subjects) our young adult control group. For the statistical analysis, we perform t-tests for comparisons across age groups and tasks, as well as tests of comparisons of proportions. Unless otherwise noted, tests are implemented with False Discovery Rate (FDR) correction for multiple comparisons whenever appropriate. Last, in the regression analysis, we compute clustered bootstrap standard errors to correct for within subject correlations.

3 Theory and predictions

Farmers and Pirates is a 3×3 asymmetric, hide-and-seek game, in which the action h of the hider (Farmer) is to hide the prize in one of 3 possible locations $h = n \in \{a, b, c\}$ and the action s of the seeker (Pirate) is to seek the prize in one of these 3 locations $s = n' \in \{a, b, c\}$. The payoff of hider and seeker are $(n, 0)$ if $n \neq n'$ and $(0, n)$ if $n = n'$.

Let $h_n \equiv \Pr[h = n]$ and $s_n \equiv \Pr[s = n]$ be the probabilities of hiding in n and seeking in n, respectively. Given the conflict of interests, there is a unique Nash equilibrium $(h_a^*, h_b^*, h_c^*, s_a^*, s_b^*, s_c^*)$ in mixed strategies. In an equilibrium, the hider chooses h_n^* so that the seeker is indifferent between looking in any of the three locations and the seeker chooses s_n^* so that the hider is indifferent between hiding in any of the three locations. Since seeking in n has an expected payoff of $n\, h_n$ for the seeker and hiding in n has an expected payoff of $n(1 - s_n)$ for the hider, an interior mixed strategy Nash equilibrium satisfies the following system of equations:[17]

$$a\, h_a^* = b\, h_b^*; \quad a\, h_a^* = c\, h_c^*; \quad h_a^* + h_b^* + h_c^* = 1; \quad h_n^* > 0 \;\; \forall n$$

$$a(1 - s_a^*) = b(1 - s_b^*); \quad a(1 - s_a^*) = c(1 - s_c^*); \quad s_a^* + s_b^* + s_c^* = 1; \quad s_n^* > 0 \;\; \forall n$$

The solution has the following immediate general consequence:

$$\text{if} \quad a > b > c \quad \text{then} \quad h_a^* < h_b^* < h_c^* \quad \text{and} \quad s_a^* > s_b^* > s_c^*. \tag{1}$$

[17] Under certain conditions, we can obtain corner solutions. We choose the parameters in our experiment to ensure that the theoretical solution is interior for all three locations. This ensures that deviations from equilibrium can go in either direction.

Solving the system, we get as the interior solution:

$$\begin{array}{lll} h_a^* = \dfrac{\frac{1}{a}}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}; & h_b^* = \dfrac{\frac{1}{b}}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}; & h_c^* = \dfrac{\frac{1}{c}}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}} \\ s_a^* = \dfrac{\frac{1}{b}+\frac{1}{c}-\frac{1}{a}}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}; & s_b^* = \dfrac{\frac{1}{a}+\frac{1}{c}-\frac{1}{b}}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}; & s_c^* = \dfrac{\frac{1}{a}+\frac{1}{b}-\frac{1}{c}}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}} \end{array} \tag{2}$$

Expected payoffs of hider and seeker are $E_h^* = \dfrac{2}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}$ and $E_s^* = \dfrac{1}{\frac{1}{a}+\frac{1}{b}+\frac{1}{c}}$.

x As in every game with a full support unique mixed strategy Nash equilibrium, players randomize between actions so as to make the opponent indifferent between alternatives. In our setting, it implies that the seeker must look more frequently in the high value location whereas the hider must (perhaps counterintuitively) select more frequently the low value location (see Eq. 1). With three locations, the hider has twice as many winning locations as the seeker. Since equilibrium payoffs are identical across locations, the payoff of hiders should, in equilibrium, be twice as a high as the payoff of seekers. Finally, our problem is built in a way that the strategy of the hider determines the size of the prize whereas the strategy of the seeker determines the sharing of the prize. Therefore, if individuals do not follow the theoretical probability mixture, aggregate payoffs may be increased or decreased depending on whether the person who deviates is the hider or the seeker.[18]

Our experimental treatment corresponds to $a = 8$, $b = 4$, $c = 3$. Using Eq. 2, the mixed-strategy Nash equilibrium is interior and given by:

$$s_8^* = \frac{11}{17},\ s_4^* = \frac{5}{17},\ s_3^* = \frac{1}{17},\ \text{ and } h_8^* = \frac{3}{17},\ h_4^* = \frac{6}{17},\ h_3^* = \frac{8}{17},\ \text{ with } E_s^* = \frac{24}{17},\ \ E_h^* = \frac{48}{17}$$

Under choice independence, the expected equilibrium probabilities of obtaining a positive payoff are:

$$p_s^* = \sum_n s_n^*\, h_n^* = \frac{71}{289} \simeq 0.246 \quad \text{and} \quad p_h^* = 1 - \sum_n s_n^*\, h_n^* = \frac{218}{289} \simeq 0.754$$

These are point estimates based on the parameters of our design. We do not expect participants to play exactly at equilibrium. In fact, given previous knowledge of experimental results on mixed strategies games, we anticipate that deviations will occur. Instead of testing the theory strictly speaking, we test two types of properties. First, we consider general qualitative properties implied by theory which we expect participants to follow, at

[18] For example, if the hider chooses the high (low) location excessively often, the sum of both players' payoffs will be higher (lower) than in equilibrium.

least in some age groups. Second, we formulate hypothesis on the expected deviations from theory. We summarize below our predictions regarding choices, messages and feedback.

Prediction P1. Choices. In Tasks 1, 2 and 3:

(i) $h_8 < h_4 < h_3,\ s_8 > s_4 > s_3,\ E_h \simeq 2\,E_s$ and $p_h \simeq 3\,p_s$.

(ii) Choices in each role are positively correlated across tasks.

(iii) Choices in each task are negatively correlated across roles.

According to **P1**(i), participants should hide most often in the low-value location and seek most often in the high-value location. Also, payoffs should be significantly higher (by a factor of two) for hiders than for seekers. Finally, since participants seek more frequently but hide less frequently in the high-value location, the probability of success for seekers must, in equilibrium, be lower than under uniform choice ($p_s^* < 1/3$ and $p_h^* > 2/3$). In our case, an unsuccessful raid should be around three times as likely as a successful one. **P1**(ii) states that while adding messages (Task 2) and feedback (Task 3) may have some effect on the choices of individuals, it should not fundamentally change their randomization patterns (that is, each individual will follow similar principles in all three tasks). Finally, **P1**(iii) implies that while we expect many participants to play as predicted by theory (low h_8 and high s_8), we also anticipate that individuals who deviate will play the opposite (high h_8 and low s_8), both of which would result in a negative correlation of choices by hiders and seekers.

Next, for Task 2, let us call $m = n \in \{a, b, c\}$ the hider's (truthful or untruthful) message m stating that the chosen location is n, and $m_n \equiv \Pr[m = n]$ the probability that the hider sends message n. We denote by $h_n^{n'} \equiv \Pr[h = n \,|\, m = n']$ the conditional probability that the hider chooses location n after sending message n' and by $s_n^{n'} \equiv \Pr[s = n \,|\, m = n']$ the conditional probability that the seeker looks in location n after receiving message n'. Unlike the original cheap talk literature where agents have partial alignment of preferences (see Crawford and Sobel (1982) and the experiments by Cai and Wang (2006) and Wang et al. (2010)), individuals have opposing interests in our game. As a result, the message of the hider should be uninformative and ignored by the seeker. We summarize this idea with the following prediction.

Prediction P2. Messages. In Task 2: $h_n^{n'} = h_n\ \forall\ n, n'$ and $s_n^{n'} = s_n\ \forall\ n, n'$.

It is important to highlight that the theory does not constrain the frequency of each message m_n. **P2** only states that the message sent by the hider must not convey any information that the seeker could use to his own advantage. A testable implication of this result is that, conditional on the location chosen, truthful messages are neither more nor less frequent than deceptive ones: $h_n^n = h_n^{n'}$ when $n \neq n'$. Similarly, trusted messages are neither more nor less frequent than untrusted ones: $s_n^n = s_n^{n'}$ when $n \neq n'$.

The second pillar of the theory states that choices should be serially uncorrelated. This prediction cannot be tested in Tasks 1 and 2. Since participants do not observe the past choices of their rival, the entire behavior reduces to the choice of a mixed strategy profile. In Task 3, participants observe past choices and outcomes so a test of dynamic behavior is possible. Generally speaking, the theory predicts that location choices in round t' should not give any relevant information to participants regarding location choices in round t ($> t'$), since it could be used to the other party's advantage. Similarly, the outcome in round t' (which, for each role, can be summarized as "win at t'" or $w^{t'}$ and "lose at t'" or $l^{t'}$), should not be indicative of the behavior in t ($> t'$). There are many ways in which individuals can fail to be dynamically unpredictable. Given our limited dataset (12 observations per role for each individual), we cannot reliably test for all possible forms of serial correlation. We therefore decided to focus on two salient properties of choices and outcomes. Denote by $h^t_{n,n'} \equiv \Pr[h^t = n \,|\, h^{t-1} = n']$ the conditional probability that the hider chooses location n in round t after choosing location n' in the immediately preceding round (and the same for seekers with $s^t_{n,n'}$). Denote also $h^t_{n,n'}(w) \equiv \Pr[h^t = n \,|\, h^{t-1} = n', w^{t-1}]$ and $h^t_{n,n'}(l) \equiv \Pr[h^t = n \,|\, h^{t-1} = n', l^{t-1}]$ the same probabilities, except that we also condition on whether the hider won or lost in round $t-1$ (and, again, the same for seekers with $s^t_{n,n'}(w)$ and $s^t_{n,n'}(l)$). We have the following predictions.

Prediction P3. Feedback. In Task 3:
(i) $h^t_{n,n} = h^t_{n,n'}$ and $s^t_{n,n} = s^t_{n,n'}$ $\forall\, t, n, n'$.
(ii) $h^t_{n,n}(w) = h^t_{n,n}(l)$ and $s^t_{n,n}(w) = s^t_{n,n}(l)$ $\forall\, t, n$.

According to **P3**(i), *choices* in a given round should not be indicative of choices in the following one. Existing findings suggest that individuals alternate excessively between options when they try to be unpredictable both in individual decision making problems (Bar-Hillel and Wagenaar, 1991; Camerer, 1995) and in mixed strategy games (Brown and Rosenthal, 1990; Camerer, 2003). If that bias is also present in our setting, it will translate into insufficient repetition of locations both as hiders ($h^t_{n,n} < h^t_{n,n'}$) and as seekers ($s^t_{n,n} < s^t_{n,n'}$). Such departures, if correctly interpreted, could be exploited by a rival. Similarly, **P3**(ii) indicates that *outcomes* should not affect the choices made immediately after. In our setting, and to focus the discussion, we consider choice repetition, and predict that it should not depend on whether the last choice resulted in success or failure.

Overall and despite the discussed data limitations, our experimental setup provides a rich set of testable predictions in terms of choices, messages and feedback. In section 4, we study and compare the choices and payoffs of participants in both roles, across ages and across tasks (**P1**). We then analyze in section 5 the marginal effect of message provision (hider) and message interpretation (seeker) on choices and payoffs (**P2**). We finally investigate in section 6 how feedback affects the dynamic ability to remain unpredictable and/or exploit the predictability of the rival (**P3**). Each of these predictions is studied in the context of age, the main variable of interest in our experiment.

4 Choices

4.1 Decisions to hide and seek: aggregate behavior

From now on, we adopt the terminology used in the experiment, namely "Farmer" and "Pirate", to refer to the roles of hider and seeker, respectively. The first step in our analysis consists in determining aggregate hide and seek choice frequencies. Figure 4 describes for each school-age group the average choices of Farmers (h_n) and Pirates (s_n) in Tasks 1, 2 and 3. For comparison, we include the behavior of the control adult population (A), as well as the theoretical predictions h_n^* and s_n^* (Th).

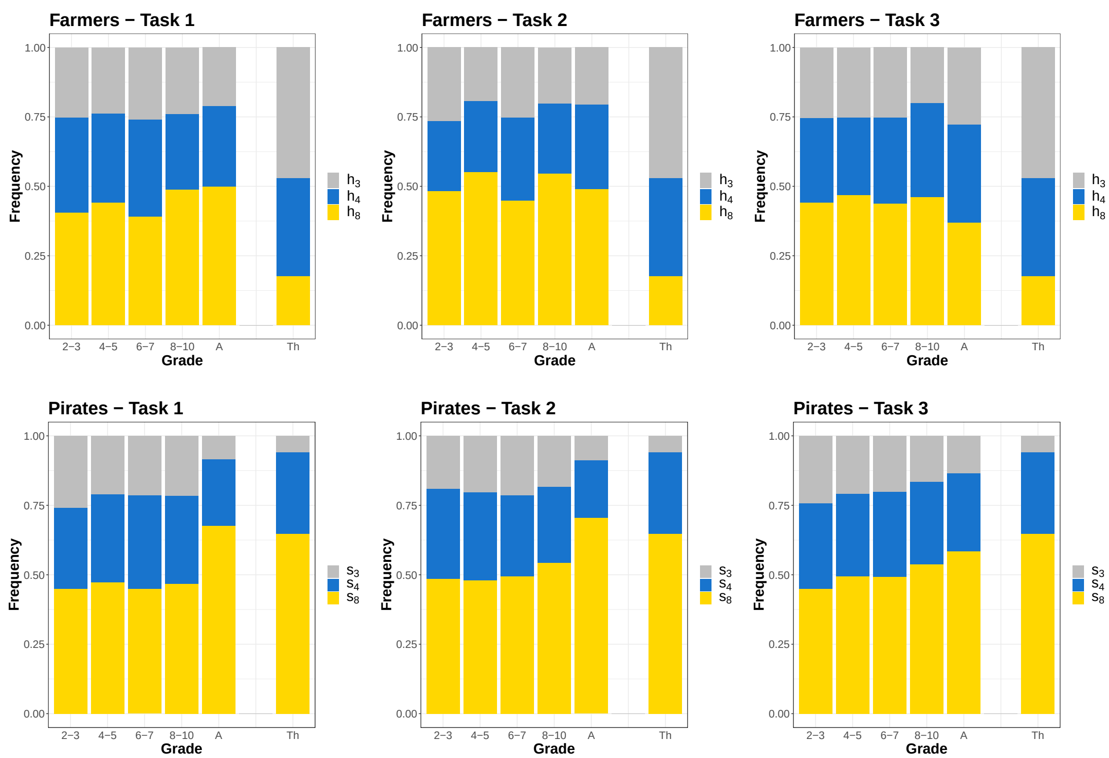

Figure 4: Distribution of decisions as Farmer (top) and Pirate (bottom) in Tasks 1, 2 and 3 by age group.

For the analysis below, we focus the comparison on the choice of location $n = 8$, the most salient one. Behavior as Farmers is similar across all school-age groups in all three tasks. Their behavior is also very similar to that of adults with some minor exceptions (higher choice of h_8 by 4-5 and 8-10 age groups in Task 3, p = 0.016 and p = 0.016). The likelihood that participants choose $n = 8$ is consistently higher than if they behaved randomly ($h_8 > 1/3$, $p < 0.05$) and consequently also significantly higher than predicted by theory ($h_8 > h_8^* = 3/17$, $p < 0.001$) in all tasks and all age groups, including adults.

Behavior as Pirates is also very similar across all school-age groups in all tasks (only 8-10 choose s_8 more than 2-3 in Task 3, p = 0.016). However, s_8 is significantly lower in all school-age groups than in the adult population in all tasks ($p < 0.02$), with the exception of 8-10 in Task 3. All school-age participants choose $n = 8$ in all tasks more often than if they behaved randomly ($s_8 > 1/3$, $p < 0.001$) but less than predicted by theory ($s_8 < s_8^* = 11/17$, $p < 0.001$). Adults choose $n = 8$ as predicted by theory in Tasks 1 and 2 and less than predicted by theory in Task 3.

Comparing the behavior of participants across tasks, we notice very small and non-systematic differences in the school-age population both as Farmers (higher choice of h_8 in Task 2 for 4/5 and 8/10) and as Pirates (lower choice of s_8 in Task 1 for 8/10). Adults behave similarly in Tasks 1 and 2 and differently in Task 3 under both roles (closer to theory as Farmers and farther away as Pirates, $p < 0.02$).

To sum up, although school-age children respond to the payoff asymmetries (h_n and s_n are different from 1/3), **P1**(i) is not supported by the data: subjects respond *incorrectly* as Farmers ($h_8^* < 1/3 < h_8$) and *correctly but insufficiently* as Pirates ($1/3 < s_8 < s_8^*$). We found no differences across school-age subjects in either role and small differences when compared to adults (only in the Pirate role). Behavior of the school-age population is very similar across tasks, suggesting that adding messages and feedback does not affect their choices substantially. Feedback affects marginally and unsystematically the behavior of adults. Overall, there is a robust heuristic tendency to favor the high yield location ($n = 8$) in both roles. At the same time, even our youngest participants realize the importance of unpredictability, as reflected by the absence of an overwhelmingly favored option.

4.2 Payoffs

It is important to determine whether the documented deviations have significant payoff consequences. Table 3 (left) presents the average per-round payoffs by task (and combining all three) in each age group. Table 3 (right) describes the payoffs of an individual who would best respond to the empirical distribution of choices in his age group. We include the location $[h_n]$ and $[s_n]$ pertaining to the optimal best response strategy. The final column reports the predictions of the theory, as developed in section 3.

From section 4.1 we know that Pirates play relatively close to equilibrium in all school-

	Empirical†					Best Response§					Theory
	2-3	4-5	6-7	8-10	A	2-3	4-5	6-7	8-10	A	
Task 1											
Farmer	3.15	3.13	3.03	3.18	2.76	4.41	4.21	4.40	4.26	3.04	2.82
	(0.17)	(0.16)	(0.18)	(0.19)	(0.27)	[h_8]	[h_8]	[h_8]	[h_8]	[h_4]	—
Pirate	2.11	2.49	2.24	2.57	3.02	3.24	3.53	3.13	3.91	4.00	1.41
	(0.19)	(0.17)	(0.17)	(.020)	(0.33)	[s_8]	[s_8]	[s_8]	[s_8]	[s_8]	—
Task 2											
Farmer	3.06	3.53	3.41	3.08	2.54	4.11	4.16	4.04	3.66	3.18	2.82
	(0.20)	(0.16)	(0.18)	(0.19)	(0.26)	[h_8]	[h_8]	[h_8]	[h_8]	[h_4]	—
Pirate	2.61	2.49	2.21	2.84	3.21	3.86	4.42	3.60	4.37	3.92	1.41
	(0.20)	(0.18)	(0.19)	(0.19)	(0.35)	[s_8]	[s_8]	[s_8]	[s_8]	[s_8]	—
Task 3											
Farmer	3.43	3.26	3.18	3.21	3.09	4.41	4.05	4.06	3.70	3.31	2.82
	(0.12)	(0.12)	(0.10)	(0.12)	(0.12)	[h_8]	[h_8]	[h_8]	[h_8]	[h_8]	—
Pirate	2.10	2.36	2.33	2.42	2.12	3.54	3.74	3.51	3.70	2.96	1.41
	(0.11)	(0.13)	(0.12)	(0.12)	(0.16)	[s_8]	[s_8]	[s_8]	[s_8]	[s_8]	—
All											
Farmer	3.27	3.29	3.20	3.17	2.87	4.34	4.12	4.14	3.83	3.21	2.82
	(0.10)	(0.08)	(0.09)	(0.09)	(0.14)						
Pirate	2.23	2.42	2.28	2.56	2.62	3.55	3.86	3.44	3.92	3.46	1.41
	(0.09)	(0.09)	(0.09)	(0.10)	(0.18)						

† (standard errors in parenthesis); §[best response location in brackets]

Table 3: Empirical and best response per-round payoff of Farmers and Pirates

age groups. This means that payoffs of Farmers will be close to theoretical predictions, independently of their behavior. This is precisely what we observe in Table 3. Participants, however, leave some money on the table: Farmers *would* have earned significantly higher gains (up to 56.4% more, $p < 0.001$) if they had best responded by setting $h_8 = 1$.

Outcomes are very different for Pirates. We know that Farmers play the opposite of the theoretical predictions, with excessive number of choices of the high-value location. Participants benefit from such departures when playing as Pirates with very significant gains (between 48.9% and 101.4% payoff increases relative to theory, $p < 0.001$). Such gains would be even higher (up to 213% increase) had they always chosen location $n = 8$.[19]

To sum up, the behavior of Farmers determines the size of the pie whereas the behavior

[19] It is interesting that our adults make the least money: their higher winnings as Pirates due to their better exploitation of Farmers does not compensate for their smaller gains as Farmers. Those differences are not statistically significant after FDR corrections.

of Pirates determines the relative gains of players. The excessive proportion h_8 by Farmers implies that total gains are very significantly above equilibrium. The close to correct proportion s_8 by Pirates means that the latter recoup a large fraction of those benefits. Overall and against **P1**(i), Pirates obtain not only a larger amount but also a larger share of profits than predicted by theory (38.0% to 55.8% instead of 33%). In Appendix A1, we discuss the proportion of rounds that individuals win in each role. We show that, also against **P1**(i), Pirates win significantly more often than predicted by theory.

4.3 Individual behavior across tasks and roles

While the results in section 4.1 suggest that aggregate choices are similar across tasks, it is possible that different individuals treat tasks differently. We therefore explore behavior at a more disaggregate level. For each individual i, we determine the proportion of choices in location $n = 8$ as Farmers (h_8^i) and as Pirates (s_8^i) in each task. We then compute the Pearson Correlation Coefficient (PCC) between each pair of tasks within each age group. The results are presented in Table 4.

	2-3		4-5		6-7		8-10		A	
Tasks	h_8^i	s_8^i	h_8^i	s_8^i	h_8^i	s_8^i	h_8^i	s_8^i	h_8^i	s_8^i
1 & 2	0.40**	0.39**	0.50***	0.39**	0.49***	0.34**	0.52***	0.57***	0.61***	0.26
1 & 3	0.43**	0.05	0.38**	0.28*	0.29*	0.58***	0.32*	0.35**	0.29	0.40*
2 & 3	0.33*	0.28	0.32*	0.41**	0.36**	0.45***	0.37**	0.52***	0.42*	0.47**

* p = 0.05; ** p = 0.01; *** p = 0.001

Table 4: Correlation of choices across tasks by age group

In support of **P1**(ii), Table 4 shows a high and statistically significant correlation of individual behavior across tasks in both roles for participants in age groups 4-5, 6-7 and 8-10 and, to a lesser extent, in our youngest and control groups. Overall, messages and feedback have only a small impact on choices.

We next study the correlation across roles. Indeed, while we have shown mistakes in both roles, we do not know if they are committed by the same or different individuals. The former would result in a negative correlation across roles while the latter would result in a positive one. We present in Figure 5 (left) a heatmap with the choices of our participants. We combine all tasks and represent each individual i by (h_8^i, s_8^i), the proportion of choices in location $n = 8$ (out of the 24) as a Farmer (horizontal axis) and as a Pirate (vertical axis). The theoretical prediction (h_8^*, s_8^*) is represented by a red **x**. Figure 5 (right) depicts the cumulative distribution functions as Farmers and Pirates.

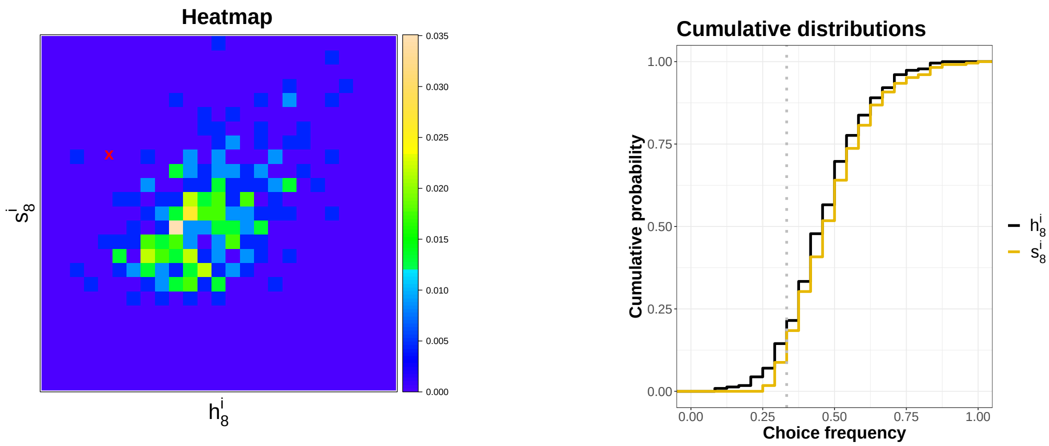

Figure 5: Individual fractions of h_8^i and s_8^i across tasks: heatmap (left) and c.d.f. (right)

In contradiction to **P1**(iii), we notice a positive (0.47) and statistically significant ($p < 0.0001$) correlation across roles. This is best illustrated in Figure 5 (left), where we observe a concentration of choices around the 45° line. No subject behaves as predicted by theory: they all miss at least one dimension, and more often both dimensions of the equilibrium randomization. At the same time, all individuals try to be unpredictable and frequently change locations. This is best reflected in Figure 5 (right), where almost no individual chooses $n = 8$ less than 25% or more than 75% of the time in either role (despite the fact that $h_8^* = 0.18$). Finally, we performed a Wilcoxon signed-rank test for matched pairs and found that the c.d.f. of h_8^i and s_8^i are significantly different from each other ($p = 0.022$), with Pirates seeking in $n = 8$ more often than Farmers hiding in that location.

4.4 Determinants of behavior

We finally run OLS regressions at the individual level to identify the determinants of choice in each role, and to assess their overall contribution to behavior. As in Figure 5, we consider the proportion of choices by an individual in location $n = 8$ across all three tasks (out of 24 decisions), and study separately their behavior as a Pirate (s_8^i) and as a Farmer (h_8^i). We include as explanatory variables the age-group (2-3 is the omitted variable), Task (Task 1 is the omitted variable) and gender ($Male = 1$). In a second step, we also include the subject's behavior in the other role. The results are presented in Table 5.

The regressions confirm previous results. Behavior in both roles is constant in our school-age population, except for a small increase in s_8^i and h_8^i in the older school group 8-10. The adult population plays closer to theory as Pirates and no different from children as

	Pirate		Farmer	
	s_8^i	s_8^i	h_8^i	h_8^i
h_8^i	—	0.212***	—	—
		(0.045)		
s_8^i	—	—	—	0.244***
				(0.052)
4-5	0.007	-0.001	0.036	0.034
	(0.025)	(0.024)	(0.028)	(0.027)
6-7	0.005	0.010	-0.024	-0.025
	(0.024)	(0.024)	(0.027)	(0.025)
8-10	0.055*	0.043	0.056*	0.042
	(0.028)	(0.026)	(0.028)	(0.026)
A	0.191***	0.189***	0.008	-0.038
	(0.034)	(0.034)	(0.039)	(0.040)
Task 2	0.040**	0.026	0.064***	0.054***
	(0.015)	(0.016)	(0.015)	(0.015)
Task 3	0.019	0.019	0.001	-0.004
	(0.014)	(0.014)	(0.015)	(0.014)
Male	0.110***	0.097***	0.062**	0.035
	(0.018)	(0.017)	(0.020)	(0.020)
Const.	0.390***	0.307***	0.393***	0.297***
	(0.020)	(0.028)	(0.022)	(0.028)
Adj. R^2	0.146	0.188	0.049	0.097
# obs.	786	786	786	786
# clusters	262	262	262	262

(clustered bootstrap standard errors in parenthesis)
* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 5: OLS regressions of choice as Pirate and Farmer

Farmers. The strong correlation of choices in both roles is also reflected in the regressions, although the exercise does not allow us to establish causality. Interestingly, males choose $n = 8$ significantly more often than females in both roles, thereby playing closer to theory as Pirates but farther away as Farmers. This could reflect the documented tendency of males to be attracted to socially-efficient outcomes more than females (Andreoni and Vesterlund, 2001) although, as discussed in section 4.5, few participants can be described as social efficiency maximizers. The high-value location is also more prevalent in Task 2, a significant difference that was not apparent in the analysis of section 4.1.

4.5 The underpinnings of suboptimal randomization

As previously mentioned, significant departures from optimal randomization have been discussed since the early experimental literature (O'Neill, 1987; Rapoport and Boebel, 1992; Ochs, 1995). Goeree and Holt (2001) already pointed out the difference between mixing in an "intuitive" and a "counterintuitive" role. However, they argue that subjects in the intuitive role understand and exploit the difficulties faced by their rivals. This conclusion is not supported in our experiment, whether we consider the adult control or any school-age group. Indeed, our participants gain experience in both roles. Therefore, if they realized and actively exploited as Pirates the severe shortcomings of Farmers, they would not incur in such deficient behavior the next time they adopted the role of Farmers. Instead, we observe no significant differences at the aggregate level in behavior across tasks in each role, as well as strong positive correlation of behavior at the individual level across tasks also in each role. The result suggests that the outcome may be based on a powerful heuristic. Such incorrect shortcut reasoning is developed early in life and it is hard to overcome with age, despite the natural development of cognitive abilities. Intuitive presentation, repetition, messages and feedback are also insufficient drivers of optimal behavior. Next, we briefly review two behavioral theories that may account for the observed deviations.

Social efficiency. Since payoffs are substantially above the Nash equilibrium predictions, a natural possibility is that individuals implicitly coordinate on the fair and socially efficient outcome of the dynamic game. This amounts to Farmers always hiding in $n = 8$ and Pirates seeking in that location only one-half of the time.[20] Such behavior necessitates a significant forward looking and self-restraint ability, since the temptation to seek in the high location knowing the hider's behavior is overwhelming. And yet, alternating behavior leading to long run fair and Pareto optimal outcomes is feasible, as it has been documented in related games on adults such as the repeated battle-of-the-sexes (Ioannou and Romero, 2014). On the other hand, individuals in our game do not observe past outcomes within each task, which makes trust and coordination significantly more difficult. Also, social preferences in children are known to evolve from selfish to altruistic (Fehr et al., 2008) and dynamic behavior from myopic to strategic forward looking (Brocas et al., 2017). This would indicate important changes in our window of observation.

To test this theory, we studied whether the behavior of our participants was consistent

[20] As standard in the experimental economics literature, we only take into account the payoffs of participants when determining the socially-efficient outcome. If we include the cost for the experimenter, the game becomes a zero-sum game and every outcome is socially efficient. Naturally, one could adopt an intermediary approach where the cost for the experimenter is positive but smaller than the gain for participants.

with always hiding in $n = 8$ and seeking one-half of the time in $n = 8$. Given the small number of observations at the individual level, we ran binomial tests, which require specifying a probability under the null hypothesis. We set $s_8^0 = 0.5$ and varied h_8^0 between 1 and 0.8. We categorized a participant as socially-efficient in a given role when we did not reject the null hypothesis in that role. However, an individual concerned with social efficiency should presumably play according to this hypothesis in both roles. We therefore categorized a participant as an overall socially-efficient individual if we did not reject the null hypothesis in either role. Table 6 reports the proportion of socially-efficient participants in each age group (by role and overall) when $s_8^0 = 0.5$ and $h_8^0 = 0.8$.

	2-3	4-5	6-7	8-10	A
Pirates	1.00	0.93	0.92	0.84	0.59
Farmers	0.04	0.14	0.08	0.16	0.08
Overall	0.04	0.12	0.07	0.08	0.03

Table 6: Proportion of socially-efficient individuals

While there is strong support for the social efficiency hypothesis in Pirates, there is much less support for social efficiency in Farmers, and little evidence that participants target social efficiency in both roles (in particular, there is not a single overall socially-efficient individual when $h_8^0 = 1$). The result suggests that the number of individuals purely driven by social efficiency is small. On the other hand, a large fraction of individuals hide in the high-value location more often than predicted by theory and seek in the high-value location less often than predicted by theory ($h_8^i > h_8^*$ and $s_8^i \in (0.5, h_8^*)$, see Figure 5). It is therefore possible that some individuals are mainly concerned with their own payoff but they still put some weight on social efficiency.

Quantal Response Equilibrium (QRE). Noisy introspection has received support in the literature as a possible explanation for the observed deviations in games with a unique mixed strategy equilibrium (McKelvey and Palfrey, 1995). However, subsequent research has also pointed out some problems in other mixed strategy games (McKelvey et al., 2000). We therefore decided to bring this stochastic best response theory to the data.

QRE is a one-parameter model that encompasses both random choice and Nash equilibrium as extreme cases. The econometric strategy consists in running an algorithm that computes a likelihood for each noise parameter and selects the value for which the highest likelihood is returned. A best fit away from no noise signifies that a noisy model of strategic choice outperforms Nash. Table 7 presents the MLE estimated parameter for each age group. Details of the procedure are relegated to Appendix A2.

	2-3	4-5	6-7	8-10	A
MLE	0.23	0.28	0.28	0.34	0.65

Table 7: Maximum Likelihood estimation of noise by age group

Given that QRE offers additional options to fit the data, it is bound to perform better than Nash. Still, there are some objective reasons why QRE is a candidate to explain behavior in our case. First, the set of possible equilibria generated by the QRE model is quite restricted. Observing that the data lie close to predictions provides support for the model. Second, it is possible to quantify the relative likelihood of the best fit compared to Nash, and to run post-hoc inferential tests to assess whether the best fit is consistent with the data in general. Appendix A2 discusses the benefits and limits of the QRE model in light of our estimates.

To sum up, social efficiency has a difficult time fitting the behavior of individuals, mainly because no participant hides in the high value location all the time. As for QRE, it is hard to interpret the meaning of stochastic mistakes when the best response is already stochastic. A third possibility is that participants do not think in terms of social preferences or best response behavior. Instead, they use a very simple heuristic reasoning: they randomize but add an extra weight in the most attractive location, so that when they win, they win "big." More generally, the problem might lie in the unrealistic expectation of finding one theory that explains the deviations of all participants. More likely, different individuals deviate for different reasons, all contributing to the observed discrepancies.

Finally, small variants could help discriminate between theories. For example, it would be interesting to modify this experiment by allowing Pirates to seek in two locations out of three rather than only one. Although, theory predicts the exact same equilibrium behavior for Farmers as in our current game ($h_8^{**} = h_8^*$) and a qualitatively similar behavior for Pirates ($s_8^{**} > s_4^{**} > s_3^{**}$), our conjecture is that behavior in both roles would move closer to theory or even reverse the bias ($h_8 < h_8^{**}$ and $s_8 > s_8^{**}$). Indeed, Farmers may expect Pirates to always search in the high value location and to randomize between the other two locations, in which case they could respond by rarely hiding in $n = 8$. If this conjecture holds, it would also run against the theories discussed above, thereby opening the door to new interpretations.

5 Messages

5.1 The message game

In Task 2, Farmers send a cheap talk message to Pirates regarding the hiding location, and Pirates observe the message before choosing where to seek. Theory predicts that rational individuals should anticipate the conflict of preferences: Farmers should randomize messages in a way that the hiding location remains unpredictable and Pirates should consequently disregard the messages, making this option irrelevant. This argument is only correct if they both play at equilibrium. In this section, we study in more detail the marginal implications of the message game. We first document in Figure 6 the frequency of messages sent by Farmers in the different age groups.

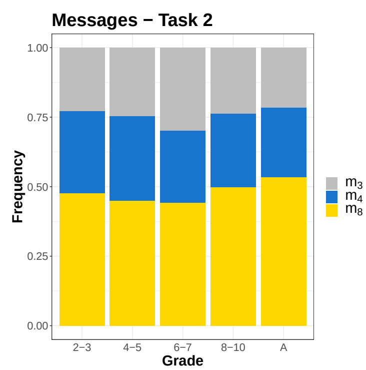

Figure 6: Messages sent by Farmers in each age group.

Subjects in all age groups choose m_8 more frequently than other messages (p < 0.05), with no significant differences between m_3 and m_4. While this difference is interesting, it is important to realize that theory is silent on the relative frequencies of messages.

5.2 Truthfulness and credence

As developed in section 3, the key variable is the message *informativeness* (or lack of). Recall that we have defined $h_n \equiv \Pr[h = n]$ and $h_n^{n'} \equiv \Pr[h = n \,|\, m = n']$ (and analogously for s_n and $s_n^{n'}$). For each individual, we call $t_n \equiv h_n^n - h_n$ the "truthfulness" of the Farmer's message n. To be unpredictable, t_n should be 0. A positive number ($h_n^n > h_n$) reflects excessive truthfulness in location n, whereas a negative number ($h_n^n < h_n$) reflects insufficient truthfulness in location n. Either out-of-equilibrium deviation, would result in a decreased payoff if correctly interpreted by the other party and an increased payoff if misinterpreted. Finally, we summarize by $t = \sum_n m_n \, t_n$, the Farmer's overall truthfulness.

Analogously, for each individual we call $c_n \equiv s_n^n - s_n$ the "credence" of the interpretation of the message received as a Pirate. In equilibrium, c_n should also be 0. A positive number ($s_n^n > s_n$) reflects gullibility, whereas a negative number ($s_n^n < s_n$) reflects skepticism. Again, either deviation could be exploited by the other party. We also summarize by $c = \sum_n m_n c_n$ the Pirate's overall credence.

Figure 7 reports the average truthfulness of the messages sent by Farmers (left) and the average credence of the messages interpreted by Pirates (right) within each age group, computed separately for each location (t_n and c_n) as well as in aggregate (t and c).[21]

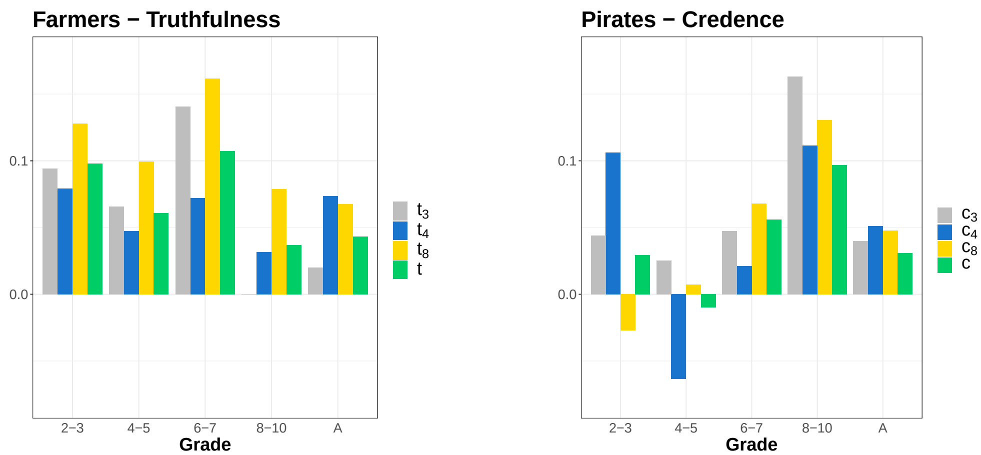

Figure 7: Truthfulness of Farmers (left) and credence of Pirates (right).

Our participants are excessively truthful Farmers in all age groups and for all messages, although given the small number of observations, the effect is significant only for school-age subjects 2-3 to 8-10 in the high-value location t_8 ($p < 0.05$). The aggregate truthfulness is significant in age groups 2-3, 4-5 and 6-7. Credence is typically positive too, although there are some exceptions. On aggregate, there is significant gullibility only in 6-7 and 8-10. Overall, messages are mildly informative so **P2** is not supported by the data. There are also interesting differences in the use and interpretation of messages across school age-groups: 2-3 and 4-5 are excessively truthful, 8-10 are excessively gullible and 6-7 are both truthful and gullible. Finally, there is a positive correlation between truthfulness and credence in the combined school population (PCC $= 0.18$, $p < 0.01$). When we look at each age group separately, the significance persists only in 8-10 (PCC $= 0.29$, $p < 0.05$).

To sum up, the asymmetry between Farmers and Pirates in messages is similar to that of location choices: message interpretation by Pirates best responds to the announcements

[21] Notice that we cannot compute t_n for an individual who never sends message n and c_n for an individual who never receives message n. We therefore omit those individuals from the average of the age group.

of Farmers (correctly believing them) while these announcements do not best respond to the Pirates' interpretation (incorrectly reporting truthfully). Again, the asymmetry is all the more surprising that individuals alternate between roles and therefore gain significant experience on both sides.

In Appendix A3, we compute the monetary consequences of the message game and show that effects are small. Indeed, Pirates are somewhat gullible but Farmers fail to exploit this bias, as they tend to send a truthful message. On the other hand, Farmers are mostly truthful but so incorrect in their hiding behavior that their bias in messages is, in comparison, largely irrelevant.

5.3 Determinants of behavior

Finally, we run OLS regressions at the individual level to identify the determinants of message choice and message interpretation. We focus on the aggregate truthfulness (t) and credence (c) of messages, respectively. The explanatory variables are age-group (2-3 is the omitted variable), gender ($Male = 1$) and the behavior of the individual in the message game when playing in the other role. We present the results in Table 8.

The previous results are confirmed in these regressions. Truthfulness and credence are strongly correlated across individuals, although causality cannot be established. Truthfulness is higher and credence is lower in our youngest school-age students (2-3) than in our oldest ones (8-10), suggesting a different use and interpretation of deception across ages. However, such differences have only minor payoff consequences. Finally, there is an interesting gender effect: males are significantly less truthful than females. This result is reminiscent of studies reporting that males lie more than females in the context of black lies (Dreber and Johannesson, 2008) or when lies benefit both the liar and another person (Erat and Gneezy, 2012). Still, other studies did not replicate those findings in adults (Childs, 2012; Cappelen et al., 2013) or in children and teenagers (Bucciol and Piovesan, 2011; Glätzle-Rützler and Lergetporer, 2015), indicating that small differences in design may affect these particular results.

6 Feedback

There are three main changes in Task 3 relative to Task 1: (i) participants play twice as many rounds (24); (ii) they change roles only once after the first half of the task; and (iii) they observe their rival's decision at the end of each round, thereby allowing them to adapt the strategy to observed choices and outcomes.

	Farmer t	Pirate c
c	0.229**	—
	(0.072)	
t	—	0.168**
		(0.053)
4-5	-0.022	-0.034
	(0.037)	(0.032)
6-7	0.009	0.024
	(0.037)	(0.032)
8-10	-0.077*	0.078*
	(0.037)	(0.031)
A	-0.054	0.011
	(0.043)	(0.037)
Male	-0.050*	0.007
	(0.024)	(0.021)
Const.	0.115***	0.010
	(0.030)	(0.026)
Adj. R^2	0.053	0.064
# obs.	262	262

(standard errors in parenthesis)
* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 8: OLS regressions of truthfulness and credence by Farmers and Pirates

6.1 Feedback from choices

While the theory predicts no significant changes in the randomization strategies when past choices are observable (**P3**), there are numerous reasons why we may expect this prediction not to hold. First, unpredictable randomization is difficult to achieve, even when probabilities are known (Bar-Hillel and Wagenaar, 1991; Camerer, 1995). Also, a subject may deviate assuming that the opponent will misinterpret the deviation and react in the subject's advantage. This is akin to the behavior in the message game, where it is possible that Farmers were excessively truthful (incorrectly) expecting the skepticism of Pirates while Pirates were gullible (correctly) expecting the truthfulness of Farmers. Finally, some of our subjects are relatively young and inexperienced. We may, therefore, observe differences across age groups due to a disparity in their ability to process past information. Unfortunately, our limited dataset prevents us from conducting a full analysis of serial correlation. We will therefore look for the presence (or not) of specific patterns.

Figure 8 reports the likelihood that a given individual chooses location n at date t as a function of their location n' at $t-1$, averaged over rounds 2 to 12 and over all individuals in that age group. The information is presented separately for Farmers ($h^t_{n,n'}$, left) and Pirates ($s^t_{n,n'}$, right). To help the visualization, we group together locations 3 and 4 under the label "-8".

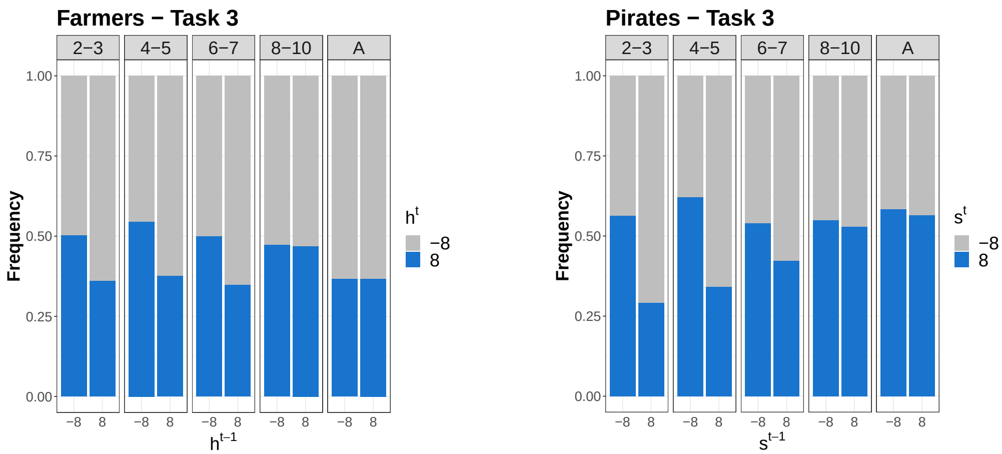

Figure 8: Persistence in location choice of Farmers (left) and Pirates (right) by age group.

Participants in age-groups 2-3, 4-5 and 6-7 are more likely to change the hiding location as Farmers than to repeat the previous one. Formally, $h^t_{8,8} < h^t_{8,-8}$ (which automatically implies $h^t_{-8,-8} < h^t_{-8,8}$). The difference is large (up to 17.0 p.p) and highly significant (p $<$ 0.002). We do not observe the same effect in 8-10 or A. The conclusion is similar–if anything stronger–for Pirates. Our participants in age-groups 2-3, 4-5 and 6-7 are more likely to alternate than to repeat seeking locations ($s^t_{8,8} < s^t_{8,-8}$) with a large (up to 28.0 p.p) and highly significant (p $<$ 0.004) difference. No effect is observed in 8-10 or A. As with some previous results, this non-equilibrium behavior as Pirates is a best response to their behavior as Farmers, while the behavior as Farmers is not a best response to their choice as Pirates. Summing up, against **P3**(i), our young participants alternate excessively between options in their attempt to be unpredictable.[22]

In Appendix A4, we analyze the reaction to outcomes and show that success or failure in a given round has a negligible impact on the subsequent choice of the individual.

[22] When we consider all three locations separately, we obtain similar effects, although the significance is reduced due to the smaller number of observations in each category.

6.2 Regression analysis of location changes across rounds

We also ran Probit regressions of the change in location between two consecutive rounds $t-1$ and t. For the regressors, we included a dummy to capture the outcome in round $t-1$ (*Win = 1*), as well as dummies for location, with $n = 3$ being the omitted variable (*Loc.4* and *Loc.8*). We also included a dummy *FirstRole* to control for order effects (whether the individual started as a Farmer or as a Pirate) and a *Period* variable that takes values 2 to 12 to capture trend effects. Finally, we added the familiar age and gender dummies. The results for Farmers ($h^t_{n,n'}$) and Pirates ($s^t_{n,n'}$) are reported in Table 9.

	Farmer $h^t_{n,n'}$	Pirate $s^t_{n,n'}$
Win	0.070	0.016
	(0.053)	(0.075)
Loc.4	-0.136	-0.155
	(0.080)	(0.094)
Loc.8	-0.455***	-0.822***
	(0.077)	(0.088)
FirstRole	0.029	0.050
	(0.063)	(0.068)
Period	0.016*	0.004
	(0.008)	(0.008)
4-5	-0.125	-0.036
	(0.102)	(0.099)
6-7	-0.131	-0.305**
	(0.102)	(0.105)
8-10	-0.367***	-0.472***
	(0.095)	(0.107)
A	-0.370***	-0.524***
	(0.102)	(0.098)
Male	0.077	-0.105
	(0.064)	(0.064)
Const.	0.730***	1.269***
	(0.104)	(0.134)
AIC	3478	3252
# obs.	2876	2876
# clusters	262	262

(clustered bootstrap st. errors in parenthesis)
* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 9: Probit regressions of location changes by Farmers and Pirates

According to these regressions, and in support of the analysis in Appendix A4, participants do not change their behavior as a function of the payoff obtained in the preceding round. In other words, strategies of the type "win-stay, lose-switch" or "win-switch, lose-stay" are not prevalent in our experiment. We do observe, however, a lower tendency to alternate in location $n = 8$ and by older participants (grade 8 and above as Farmers and grade 6 and above as Pirates). Overall, the results confirm previous findings regarding dynamic behavior: no effect of outcome and a tendency to alternate excessively in the younger population.

7 Conclusion

We have reported the results of a hide-and-seek experiment with children and adolescents. Given the opposite interests, equilibrium behavior requires players to be unpredictable and impervious both to cheap talk messages and feedback. Our graphical paradigm and developmental approach allows us to study foundational aspects of these predictions by testing the evolution of choices from childhood to adulthood. The game features locations with known but different values, which creates a natural way to study the relationship between potential failures of unpredictability and payoff-saliency effects. We found that the vast majority of participants favor the high-value location not only as seekers (as predicted by theory) but also as hiders (in contradiction with theory). Strikingly, we found little to no differences in behavior across age-groups, despite the very significant changes in cognitive capacities in our window of observation.

The lack of improvement contrasts with existing games in the developmental literature and could be explained by a powerful and intuitive (yet erroneous) heuristic thinking. It also sharply contrasts with the close to equilibrium behavior documented in the choices of "expert randomizers", such as professional sport players (Walker and Wooders, 2001; Chiappori et al., 2002; Palacios-Huerta, 2003). It is possible that these activities attract individuals with a better practical understanding of the subtleties behind the mixed strategy Nash equilibrium. Perhaps more likely, their continual exposure to the negative effects of suboptimal randomization with very significant payoff consequences promotes the (conscious or unconscious) development of equilibrium strategies. Finally, it also departs from behavioral studies reporting Nash compliance in chimpanzees playing symmetric and asymmetric matching pennies games (Martin et al., 2014). Contrary to non-human primates, we may have lost or not developed this ability because our environment does not select for unpredictability (Miller, 1997), making us ill-adapted to these games.

We added a cheap talk stage to the original design, and addressed the potential for deception. Even though hiders in our game are excessively truthful and seekers have a slight tendency to believe them, messages do not impact outcomes significantly. Also, the version with feedback addresses learning patterns. In general, both hiders and seekers alternate excessively between locations, especially in the younger population, but once again the impact on outcomes remains small.

Playing a mixed strategy requires putting our decisions in the hands of a mental randomization device. The problem is related to the generation of random numbers, a well-studied issue. Experiments show that sequences exhibit too few repetitions and serially correlated responses (Rabinowitz et al., 1989; Falk and Konold, 1997), with only modest developmental improvements (Towse and Mclachlan, 1999). These limitations have been traced to executive functions (Jahanshahi et al., 2000). Overall, the inability of our participants to mix strategies in the correct proportions is consistent with these individual decision making results. It suggests that *succeeding* at being unpredictable in game theoretical settings relies on an inhibitory mechanism that never fully develops. At the same time, even our youngest participants realize the importance of *trying* to be unpredictable. Indeed, no subject in the entire sample overwhelmingly favors one option over the others.

Neuroscientific research reports that saliency detection is a key attentional process that allows organisms to focus their limited perceptual and attentional resources on what is pertinent (Itti and Koch, 2001; Friston, 2010). However, the process is subject to a bias. Attention should be focused on the features relevant for the behavioral goal (e.g., to select the correct strategy) and salient elements sometimes act as distractors, creating a conflict that must be resolved. Biasing selection towards or away from salient stimuli is implemented by different competing brain areas (Mevorach et al., 2006). Lack of coordination may result in a behavioral bias, which provides a unique and unifying explanation for experimental results where salient options are favored. In our case, payoff-salience is an added difficulty. Not only do the participants not produce responses independent of previous responses but they are attracted by high rewards which, in the case of hiders, lure them away from equilibrium choices.

There is still much to learn about the foundations of suboptimal behavior in mixed strategy games. New experimental variants could help disentangle the reasons for the observed deviations. For example, it would be useful to run both the asymmetric (as in the current paper) and the symmetric hide-and-seek game (the traditional rock-paper-scissors) with the same population. In the symmetric version, we expect that participants from a young age will intuitively understand the optimality of playing each action with probability 1/3 (thereby eliminating the bias due to payoff salience). However, we would still expect deviations due to the difficulty to implement an unpredictable strategy as well as the prospect of exploiting the predictability of the rival.

References

James Andreoni and Lise Vesterlund. Which is the fair sex? gender differences in altruism. *The Quarterly Journal of Economics*, 116(1):293–312, 2001.

Maya Bar-Hillel and Willem A Wagenaar. The perception of randomness. *Advances in applied mathematics*, 12(4):428–454, 1991.

Elizabeth Bonawitz, Stephanie Denison, Alison Gopnik, and Thomas L Griffiths. Win-stay, lose-sample: A simple sequential algorithm for approximating bayesian inference. *Cognitive psychology*, 74:35–65, 2014.

Isabelle Brocas and Juan D Carrillo. Introduction to special issue "understanding cognition and decision-making by children." studying decision-making in children: challenges and opportunities. *Journal of Economic Behavior & Organization*, 179:777–783, 2020a.

Isabelle Brocas and Juan D Carrillo. The evolution of choice and learning in the two-person beauty contest game from kindergarten to adulthood. *Games and Economic Behavior*, 120:132–143, 2020b.

Isabelle Brocas and Juan D Carrillo. Steps of reasoning in children and adolescents. *Journal of Political Economy*, 129(7):2067–2111, 2021.

Isabelle Brocas, Juan D Carrillo, and Niree Kodaverdian. Altruism and strategic giving in children and adolescents. CEPR DP12288, 2017.

James N Brown and Robert W Rosenthal. Testing the minimax hypothesis: a re-examination of o'neill's game experiment. *Econometrica*, pages 1065–1081, 1990.

Alessandro Bucciol and Marco Piovesan. Luck or cheating? a field experiment on honesty with children. *Journal of Economic Psychology*, 32(1):73–78, 2011.

Hongbin Cai and Joseph Tao-Yi Wang. Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36, 2006.

Colin Camerer. Individual decision making. *Handbook of experimental economics*, 1995.

Colin Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.

Alexander W Cappelen, Erik Ø Sørensen, and Bertil Tungodden. When do we lie? *Journal of Economic Behavior & Organization*, 93:258–265, 2013.

Gary Charness, John A List, Aldo Rustichini, Anya Samek, and Jeroen Van De Ven. Theory of mind among disadvantaged children: Evidence from a field experiment. *Journal of Economic Behavior & Organization*, 166:174–194, 2019.

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.

P-A Chiappori, Steven Levitt, and Timothy Groseclose. Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, 92 (4):1138–1151, 2002.

Jason Childs. Gender differences in lying. *Economics Letters*, 114(2):147–149, 2012.

Vincent P Crawford and Nagore Iriberri. Fatal attraction: Salience, naivete, and sophistication in experimental" hide-and-seek" games. *American Economic Review*, 97(5):1731–1750, 2007.

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, pages 1431–1451, 1982.

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*, 2016.

Anna Dreber and Magnus Johannesson. Gender differences in deception. *Economics Letters*, 99 (1):197–199, 2008.

Sanjiv Erat and Uri Gneezy. White lies. *Management Science*, 58(4):723–733, 2012.

Ido Erev and Alvin E Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881, 1998.

Ruma Falk and Clifford Konold. Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, 104(2):301, 1997.

Eduardo Fe, David Gill, and Victoria Prowse. Cognitive skills, strategic sophistication, and life outcomes. Working Paper, 2020.

Ernst Fehr, Helen Bernhard, and Bettina Rockenbach. Egalatarianism in young children. *Nature*, 454(28):1079–1084, 2008.

Daniel Fragiadakis, Ada Kovaliukaite, and David Rojo Arjona. The belief elicitation by superimposition approach. Working Paper, 2019.

Karl Friston. The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11 (2):127–138, 2010.

Jay N Giedd. The teen brain: insights from neuroimaging. *Journal of adolescent health*, 42(4): 335–343, 2008.

Daniela Glätzle-Rützler and Philipp Lergetporer. Lying and age: An experimental study. *Journal of Economic Psychology*, 46:12–25, 2015.

Jacob K Goeree and Charles A Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 91(5):1402–1422, 2001.

Jacob K Goeree, Charles A Holt, and Thomas R Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113, 2003.

William T Harbaugh and Kate Krause. Children's altruism in public good and dictator experiments. *Economic Inquiry*, 38(1):95–109, 2000.

Christos A Ioannou and Julian Romero. A generalized approach to belief learning in repeated games. *Games and Economic Behavior*, 87:178–203, 2014.

Laurent Itti and Christof Koch. Computational modelling of visual attention. *Nature reviews neuroscience*, 2(3):194–203, 2001.

Marjan Jahanshahi, Georg Dirnberger, Rebecca Fuller, and Chris D Frith. The role of the dorsolateral prefrontal cortex in random number generation: a study with positron emission tomography. *Neuroimage*, 12(6):713–725, 2000.

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

Christopher Flynn Martin, Rahul Bhui, Peter Bossaerts, Tetsuro Matsuzawa, Colin Camerer, and Colin Camerer. Experienced chimpanzees behave more game–theoretically than humans in simple competitive interactions. *Sci Rep*, 4:5182, 2014.

Kevin A McCabe, Arijit Mukherji, and David E Runkle. An experimental study of information and mixed-strategy play in the three-person matching-pennies game. *Economic Theory*, 15(2): 421–462, 2000.

Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

Richard D McKelvey, Thomas R Palfrey, and Roberto A Weber. The effects of payoff magnitude and heterogeneity on behavior in 2× 2 games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization*, 42(4):523–548, 2000.

Carmel Mevorach, Glyn W Humphreys, and Lilach Shalev. Opposite biases in salience-based selection for the left and right posterior parietal cortex. *Nature neuroscience*, 9(6):740–742, 2006.

Geoffrey F Miller. 12 protean primates: The evolution of adaptive unpredictability in competition and courtship. *Machiavellian intelligence II: Extensions and evaluations*, page 312, 1997.

Dilip Mookherjee and Barry Sopher. Learning behavior in an experimental matching pennies game. *Games and Economic Behavior*, 7(1):62–91, 1994.

J Keith Murnighan and Michael Scott Saxon. Ultimatum bargaining by children and adults. *Journal of Economic Psychology*, 19(4):415–445, 1998.

Jack Ochs. Games with unique, mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, 10(1):202–217, 1995.

Barry O'Neill. Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the national academy of sciences*, 84(7):2106–2109, 1987.

Ignacio Palacios-Huerta. Professionals play minimax. *The Review of Economic Studies*, 70(2): 395–415, 2003.

F Michael Rabinowitz, William P Dunlap, Malcolm J Grant, and Joseph C Campione. The rules used by children and adults in attempting to generate random numbers. *Journal of Mathematical Psychology*, 33(3):227–287, 1989.

Amnon Rapoport and Richard B Boebel. Mixed strategies in strictly competitive games: A further test of the minimax hypothesis. *Games and Economic Behavior*, 4(2):261–283, 1992.

Ariel Rubinstein, Amos Tversky, and Dana Heller. Naive strategies in competitive games. In *Understanding strategic interaction*, pages 394–402. Springer, 1997.

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.

Mark Steyvers, Michael D Lee, and Eric-Jan Wagenmakers. A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179, 2009.

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.

John N Towse and Amy Mclachlan. An exploration of random generation among children. *British Journal of Developmental Psychology*, 17(3):363–380, 1999.

Mark Walker and John Wooders. Minimax play at wimbledon. *American Economic Review*, 91 (5):1521–1538, 2001.

Joseph Tao-yi Wang, Michael Spezio, and Colin F Camerer. Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, 100(3):984–1007, 2010.

Appendix A. Additional analyses

A1. Probability of success in each role

As emphasized in section 3, since Farmers should hide predominantly in the low-payoff locations and Pirates should seek predominantly in the high-payoff ones, the former should in equilibrium be successful more often than under uniform random behavior (which would correspond to 2/3) and the latter less often (which would correspond to 1/3). Table 10 reports the empirical probabilities of successful hides and successful raids. We also add the theoretical prediction for comparison.

	Empirical†					Theory
	2-3	4-5	6-7	8-10	A	
Task 1	0.62 / 0.38 (0.03)	0.60 / 0.40 (0.03)	0.60 / 0.40 (0.03)	0.60 / 0.40 (0.03)	0.57 / 0.43 (0.04)	0.75 / 0.25
Task 2	0.59 / 0.42 (0.03)	0.62 / 0.38 (0.02)	0.65 / 0.35 (0.03)	0.58 / 0.42 (0.03)	0.55 /0.45 (0.04)	0.75 / 0.25
Task 3	0.65 / 0.35 (0.02)	0.62 / 0.38 (0.02)	0.62 / 0.38 (0.02)	0.62 / 0.38 (0.02)	0.67 / 0.33 (0.02)	0.75 / 0.25

† (standard errors in parenthesis)

Table 10: Probability of success of Farmer / Pirate by task and age group

We have already emphasized that when Pirates engage in a successful raid they obtain higher payoffs than predicted by theory because Farmers hide excessively often in the high yield locations. According to Table 10, they also win significantly more often than predicted by theory: between 35% and 42% compared to the predicted 25% ($p < 0.001$). The effect is magnified in the adult population. Overall, it is the combination of both excessively high probability of success and excessive payoff whenever the raid is successful that accounts for the large payoff increase of Pirates relative to the prediction of our game.

A2. Quantal Response Equilibrium

In a QRE, best response functions are stochastic. Actions with higher expected payoffs are selected more often, but the best option is not chosen with probability one. For the sake of parsimony and simplicity, we consider the version where players are symmetric, that is, they are all subject to the same stochasticity. Denote by $h_{(n,\lambda)}$ and $s_{(n,\lambda)}$ the probability that a Farmer and a Pirate choose location n in a QRE with error parameter λ. We adopt the standard Logit parametrization, so the mixed strategy QRE $(h_{(n,\lambda)}, s_{(n,\lambda)})$ solves the following system of equations:

$$h_{(n,\lambda)} = \frac{e^{\lambda(1-s_{(n,\lambda)})n}}{\sum_n e^{\lambda(1-s_{(n,\lambda)})n}} \quad \text{and} \quad s_{(n,\lambda)} = \frac{e^{\lambda\, h_{(n,\lambda)}\, n}}{\sum_n e^{\lambda\, h_{(n,\lambda)}\, n}} \qquad n \in \{3,4,8\} \tag{3}$$

with the familiar limit results of uniform random choice when $\lambda = 0$ ($h_{(n,0)} = s_{(n,0)} = 1/3$) and convergence to Nash equilibrium when $\lambda = \infty$ ($h_{(n,\infty)} = h^*_n$ and $s_{(n,\infty)} = s^*_n$). Figure 9 depicts with a green curve all the possible QRE predictions as a function of λ in locations $n = 8$ (left), $n = 4$

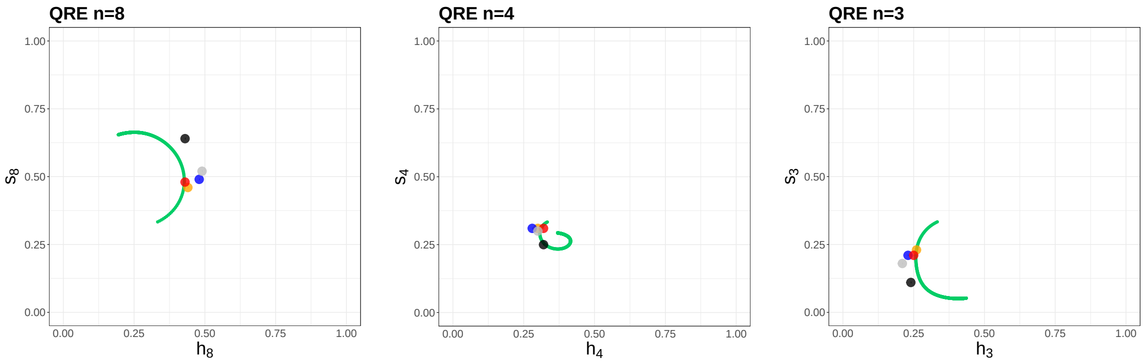

Figure 9: Quantal Response Equilibrium.

(center) and $n = 3$ (right). The endpoints of the curve correspond to Nash equilibrium (h_n^*, s_n^*) and random behavior $(1/3, 1/3)$. The dots represent the average choices in each age group.

Note that the set of equilibria predicted by the QRE model across all possible parameters is relatively small. The fact that the data lie close to the green curve in all three locations (yet far away from either extreme) is suggestive evidence that the model is capturing some aspects of the observed choices (Figure 9). Notice also that behavior would not be captured well by a simpler model where the decision is a convex combination of Nash equilibrium and random choice (without the fixed point element of stochastic best response inherent to QRE). Predicted choice in such model could only lie on the line that connects the endpoints of the green curve.

To provide a measure of fit of the QRE model, we compute an estimate of the noise parameter that best represents the aggregate choice. We implement maximum likelihood estimation (MLE) of the system of equations that characterizes the QRE equilibrium (Eq. 3). For each age group, the algorithm computes a likelihood (based on the trial-by-trial choices) for each possible noise parameter and selects the parameter obtaining the highest likelihood. Since the procedure estimates the parameter on the choices both as Farmers and as Pirates, it is blind to potential "role" effects. The results are reported in Table 7. To assess whether choices are similarly noisy across roles, we use a minimum-distance estimation (MDE) technique. Formally, we compute the euclidian distance between the aggregate distribution of play and the distribution of play predicted by QRE for each possible noise parameter in each role. Our algorithm returns the parameter yielding the smallest distance. We also provide MDE estimates across roles. Table 11 reports the MDE estimates for each age group. For comparison, we also include the MLE estimates previously reported.

	2-3	4-5	6-7	8-10	A
MDE Farmers	0.25	0.26	0.33	0.28	0.36
MDE Pirates	0.23	0.28	0.26	0.34	1.01
MDE Both roles	0.23	0.27	0.27	0.32	0.60
MLE Both roles	0.23	0.28	0.28	0.34	0.65

Table 11: Best noise estimate λ^* by age group and role.

Importantly, Table 11 highlights the congruency of MDE estimates across roles in all the children age groups. This contrasts with the adult population, for which we need significantly different parameters to fit the behavior in each role. Since the same set of subjects play in both roles, the result suggests that, if anything, a model of noisy introspection captures better the behavior of children than that of adults in our population.

A common way to evaluate QRE is to compare the likelihoods attached to Nash equilibrium ($\lambda = \infty$) and to the best fit ($\lambda = \lambda^*$). This requires using our MLE estimates. Simple likelihood ratio computations show that the data favor the best fit by a factor ranging from 64 for the adult group to 6250 for the 2-3 age group. A complementary approach consists in testing the hypothesis that the prediction of the best fit is consistent with the data. To do this, we ran Hotelling's T-squared test, the multivariate counterpart of the t-test, on the distribution of the individual distributions of choices. We tested the hypothesis that the mean distribution was equal to the distribution of choices corresponding to the estimated noise parameter (under both MLE and MDE). The equality hypothesis was not rejected in groups 2-3 and A for both Farmers and Pirates and in group 6-7 for Farmers, under either estimation method. For completeness, we also tested the hypothesis that the mean of this distribution was equal to the choices dictated by Nash theory. It was rejected in all roles and all age groups (all $p < 0.003$). Taken together, the results indicate that QRE performs better than Nash. On the other hand, and as discussed earlier, QRE does not account for all relevant sources of behavioral variations.

A3. Payoffs in Task 2

For Task 2, we can compute the best response of individuals to the empirical behavior of subjects in their age group. It is the same exercise as we did in Task 1 (right columns of Table 3), except that we enlarge it by incorporating the message game. Table 12 summarizes the information. For Farmers, it includes the per-round payoff under best response (BR payoff), the best response hiding and message strategy $[h_n, m_{n'}]$ (BR strat.) and, for comparison, the empirical payoff obtained by individuals in that age group (emp. payoff). For Pirates, it includes the per-round payoff under best response and the best response seeking strategy $[s_n]$ as a function of the message observed $m_{n'}$. Again for comparison, we also compute the payoff if the Pirate always believed the Farmer's message, separated by message (naïve payoff) as well as the overall value (naïve average). Finally, we include the empirical payoff by message (emp. payoff) and overall (emp. average).

Even when we include the option of sending a message, it is still in the Farmers' best interest to hide in location $n = 8$ in all school age groups. Interestingly, the (small) gullibility of Pirates implies that Farmers should lie about it (with the curious exception of our youngest group 2-3 who are the only ones to be skeptical about m_8). Comparing the best response payoffs in Tables 3 and 12, we notice that the message adds only a small value to Farmers in all age groups except 8-10: the best response payoff is only increased marginally when we include this possibility. Since gullibility is high among 8-10 participants, by misleading with message m_3, payoffs in that group could be increased by 37.7% compared to best responding without taking the message into consideration and by 63.6% compared to the payoff empirically obtained.

Pirates, on the other hand, should optimally ignore the message and seek always in the high yield location. Indeed, despite the significant truthfulness of Farmers (as reported in Figure 7), their hiding behavior is so systematically biased towards h_8 that Pirates should exploit this tendency at the expense of any other information conveyed by the message. Notice also that Pirates would lose a non-negligible fraction of earnings if they blindly followed the message of Farmers in locations 3 and 4 (naïve payoff). However, the naïve average in the school-age groups would still be 4.9%

	Farmers				
	2-3	4-5	6-7	8-10	A
BR payoff	4.36	4.19	4.81	5.04	3.61
BR strat.	$[h_8, m_8]$	$[h_8, m_3]$	$[h_8, m_3]$	$[h_8, m_3]$	$[h_8, m_4]$
emp. payoff	3.06	3.53	3.41	3.08	2.54

	Pirates														
	2-3			4-5			6-7			8-10			A		
	m_3	m_4	m_8	m_3	m_4	m_8	m_3	m_4	m_8	m_3	m_4	m_8	m_3	m_4	m_8
BR payoff	2.74	3.12	4.86	3.90	3.69	5.19	2.39	2.80	4.88	3.18	4.04	5.10	2.91	3.45	4.55
BR strat.	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$	$[s_8]$
naïve payoff	1.16	1.33	4.86	0.68	1.23	5.19	1.15	1.57	4.88	0.72	1.13	5.10	0.95	1.33	4.55
naïve average		2.98			2.87			2.98			2.93			2.97	
emp. payoff	2.05	2.16	2.70	2.28	2.18	3.01	1.65	1.37	3.11	1.59	2.32	3.64	2.01	3.10	3.92
emp. average		2.61			2.49			2.21			2.84			3.21	

Table 12: Empirical and best response payoff of Farmer and Pirate by task and age group

to 34.8% higher than their actual gains. Only adults would lose 7.4% of payoffs by following the Farmers' suggestion every time in every location.

A4. Feedback from outcomes

Although choices are informative, our participants may also react to outcomes. Indeed, researchers have documented that payoff-based strategies of the type "win-stay, lose-switch" and other reinforcement learning algorithms are common (Erev and Roth, 1998; Steyvers et al., 2009; Bonawitz et al., 2014). To study a basic version of this possibility, we analyze the differences in the choice immediately after a success or a failure. Formally, we look at the variables $(h^t_{n,n'}(w), h^t_{n,n'}(l))$ for Farmers and $(s^t_{n,n'}(w), s^t_{n,n'}(l))$ for Pirates. There are numerous ways to look at these changes. To avoid an excessively fine partition with few observations in each case, we focus on the frequency that choices are repeated vs. changed, aggregated across all locations. Figure 10 reports the likelihood that a given individual repeats (stay) or changes (switch) location between two consecutive rounds as a function of whether his last choice was a success (w) or a failure (l), averaged over rounds 2 to 12 and over all individuals in that age group. The information is separated between Farmers (left) and Pirates (right).

There are some differences in the likelihood of a location change as a function of the outcome. Farmers in 2-3 and 8-10 switch more frequently after winning ($p < 0.03$) whereas Pirates in 8-10 and A switch less frequently after winning ($p < 0.02$). However, the effects are non-systematic and small (up to 13.5 p.p). Overall, in support of **P3**(ii) we found relatively minor effects of success and failure on the dynamics of choice. By contrast, Figure 10 (right) reinforces the findings of section 6.1 which stated that, independent of the outcome, younger participants switch locations as Pirates more frequently than older ones.

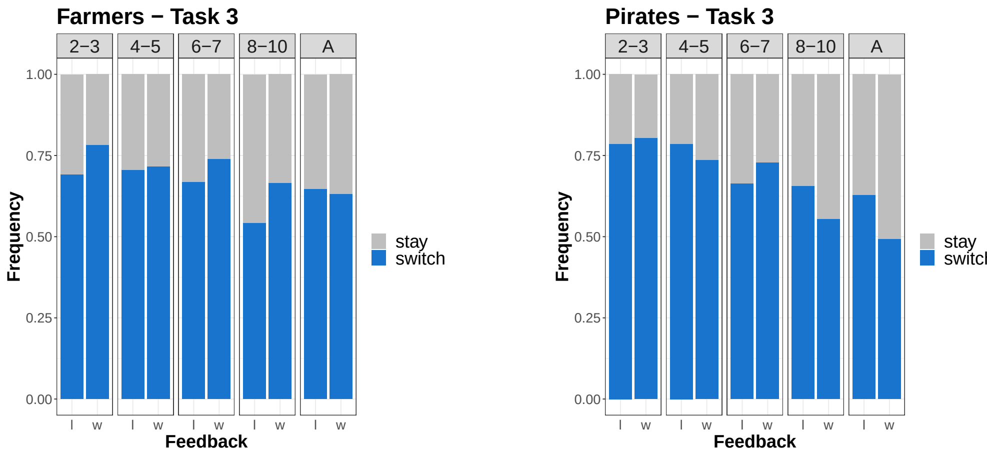

Figure 10: Choice persistence after success and failure.

Appendix B. Younger children

We had the opportunity to run a shortened version of the experiment with 42 children from Kindergarten (5-6 years old) and 34 children from 1st grade (6-7 years old) from the same school (LILA), called K-1 age group. Given their (very) young age, their attention span is limited. Since the game is quite simple and intuitive, we decided to follow the same experimental procedures but run only Task 1. Table 13 presents a summary of their choices and payoffs obtained as Farmers and Pirates. For ease of comparison we also include the data from age group 2-3, already discussed in the paper.

Farmers				Pirates			
		K-1	2-3			K-1	2-3
Choices	h_8	0.47	0.41	Choices	s_8	0.42	0.45
	h_4	0.28	0.34		s_4	0.31	0.29
	h_3	0.25	0.25		s_3	0.27	0.26
Payoffs	empirical	3.57	3.15	Payoffs	empirical	2.04	2.11
	best response	4.68	4.41		best response	3.73	3.24
	BR strat.	$[h_8]$	$[h_8]$		BR strat.	$[s_8]$	$[s_8]$

PCC $(h_8, s_8) = 0.27$ (p = 0.017) [K-1]; PCC $(h_8, s_8) = 0.05$ (p = 0.72) [2-3]

Table 13: Behavior of young children (K and 1st) in Task 1.

From this table, we notice a remarkable similarity in all aspects of the game between K-1 and 2-3 (who, as we have already extensively discussed, are also similar to the older participants). Indeed, as in the other age groups, Farmers play the opposite of the theory prediction with an excessive choice of h_8. The choice of Pirates matches qualitatively the theoretical predictions, although s_8 is, once again, too low. Since actions of K-1 are not different from 2-3, payoffs are also very similar. The slightly higher choice of h_8 by K-1 implies more gains to distribute whereas the slightly lower choice of s_8 results in a higher share for Farmers than in 2-3. Nevertheless, differences are marginal and best response strategies are identical to all other age groups (h_8 and s_8). Finally, there is also a noticeable positive correlation in the choice of location $n = 8$ in both roles so that, once again, Pirates best respond to their own behavior as Farmers but Farmers do not best respond to their behavior as Pirates. In conclusion, there are clear systematic patterns in the randomization choices of school age individuals and, at the same time, a very remarkable lack of change from 5 to 16 years old.

Appendix C. Sample of instructions (grades 2 to 5)

FARMERS AND PIRATES

This game has three levels. In each level, you will be playing with a different person in the room. The computer will decide who this person is but you will not know and it is not the point to find out. You will be either a "Farmer" or a "Pirate". Farmers grow magic flowers. If you are a Farmer, you will see a screen like this (a copy of the slides is provided in Figure 11):

[SLIDE 1]

Your job is to hide the flower on one of the three islands. But, in some islands flowers grow better than in others. If you hide the magic flower in this island, 8 baby flowers will grow. If you hide it in that other island, 4 baby flowers will grow. And in that other, 3 baby flowers will grow. You choose where to hide the flower. To do so, you just need to tap on the island you want the flower to be hidden in. For example, you can hide it in this island.

[SLIDE 2]

You can change your mind as many times as you want. When you are sure of your choice, you need to press OK to lock your choice in.

Now, if you are a Pirate, your goal is to steal the baby flowers. Your screen will say WAIT when the farmer is hiding the flower. When the farmer has decided where to hide the flower, you will see a screen like this.

[SLIDE 3]

You know that in this island, 8 baby flowers can grow, in this other, 4 baby flowers can grow and in this other, 3 baby flowers can grow. Your job is to guess where the farmer hid the flower. When you think you know, you just tap on the island you think the baby flowers are. A red circle appears showing that island.

[SLIDE 4]

You can change your mind as many times as you want. When you are sure of your choice, you press OK to lock it in. But how do you get points in this game? Well, if the Pirate finds the island where the flower was hidden, he steals the baby flowers! And each baby flower is worth one point. If the Pirate does not find it, the Farmer keeps the baby flowers and, again, each baby flower is worth 1 point. So, for example, if the Farmer hides the flower in the island that gives 4 baby flowers and the Pirate finds it, the Pirate gets the 4 baby flowers and the Farmer gets nothing. If he does not find it, then the Farmer keeps the 4 baby flowers and the Pirate gets nothing [LAUNCH THE GAME].

LEVEL 1

This is level 1 of the Farmers and Pirates game. Now, the computer will match you with a person in the room. In this level, you will alternate roles: one time farmer, one time pirate, one time farmer, one time pirate and you will play many times. You will not know what the other person did each time. Each time you play, you can do the same as before or something different. It is entirely up to you. The computer will tell you at the end of the game how many baby flowers you got. Remember, more baby flowers, more points and more points, more toys in the shop.

Let's start with two pretend matches. These matches do not count for real so you will not get points here. Just use it to practice what you want to do next. If you are a farmer choose one of the island. Once you are done, wait. If you are a pirate wait for the farmer then guess which island has the flower [wait].

Let's play a second time, but now with the opposite role. If you were a farmer last time, now you are a pirate, and if you were a pirate, now you are a farmer. Again, this game is not for real [AFTER THE PRACTICE] Are you ready to play for real now?

[START LEVEL 1]   [AT THE END]

Level 1 is finished. You will now see a screen like this.

[SLIDE 5]

It tells you when you were a farmer, how many times you kept the flowers [point to happy face] and how many times you lost the flowers to the pirate [point to sad face]. It also tells you when you were a pirate, how many times you found the flowers [point to happy face] and how many times you did not find the flowers [point to sad face]. Finally, it also tells you all the points you have won in this level of the game.

LEVEL 2

OK, now we will move to level 2. The game is the same as before except that, now, if you are the Farmer, you can try to trick the Pirate. If you are a Farmer, this is what you will see on your screen:

[SLIDE 6]

Just as before, you decide where to hide the flower by tapping on an island. But now, there is another screen on the right. This screen is for you to tell the Pirate where you hid the flower. You can tell the truth, or you can lie. It is up to you. When you have decided what to tell the Pirate, you just tap on the island you want him to believe the flower is in. A blue circle appears in the island you click.

[SLIDE 7]

Here for example, this is where I hid the flower, and this is where I said I hid the flower. You can change your mind. When you are sure where you want to hide the flower and what you want to tell the pirate, you press OK.

Now if you are a Pirate, you will see a screen that says WAIT while the Farmer makes his choice. Then you will see the same as before but also the island that the Farmer wants you to believe the flower is in. That island is circled in blue.

[SLIDE 8]

You then need again to guess where the flower is. You may choose to believe the farmer or not. To select an island, you tap on it and, just as before, it gets circled in red.

[SLIDE 9]

Just like before, if the Pirate finds the island with the flower, the Pirate gets the baby flowers. If the Pirate does not find the island with the flower, the Farmer keeps the baby flowers.

[LAUNCH THE GAME]

The computer will now match you with another person in the room. Again, you will alternate roles: one time farmer, one time pirate, one time farmer, one time pirate and you will play many times. Whenever you are the Farmer, you can try to trick the pirate or tell the truth. When you are the Pirate, you can believe the farmer or not. Each time you play, you can do the same as before or something different. It is entirely up to you. As before, the computer will tell you only at the end of the game how many baby flowers you got when you were a Farmer and how many you got when you were a Pirate. Are you ready to play?

[START] [AT THE END]

Level 2 is finished. Just like at the end of level 1, you will see a screen like this.

[SLIDE 10]

It tells you how many times you won as a Farmer, how many times you won as a Pirate, and your total number of points.

LEVEL 3

Let's move to level 3. In this level, you cannot trick the Pirate anymore but you can see what the other person did the last time you played, and so will the other person. In level 3, you will play as a farmer or as a pirate several times in a row. Only after several times you will switch role and play in your new role for several times. Suppose you are the farmer and you have been playing already 3 times as a Farmer, you will see a screen like this:

[SLIDE 11]

This screen tells you that you are currently the Farmer. On this side, it tells you in which island you hid the flowers each of the 3 times you already played as a Farmer and whether you won each time. If you are the pirate, you see a screen like this:

[SLIDE 12]

This screen tells you that you are currently the Pirate. On this side, it tells you in which island you looked for the flower each time and whether you won. Do you have any questions? Let's play then. Remember you are not going to change roles every time you will be always the farmer or always the pirate.

[LAUNCH THE GAME] [AFTER HALF THE MATCHES]

Ok. Now you guys are going to switch roles. If you were a farmer you will now be a pirate and if you were a pirate you will now be a farmer. You will play with the same person.

[AT THE END, SHOW THEM FEEDBACK]

Level 3 is finished. You will see a screen like this.

[SLIDE 13]

It tells you how many baby flowers you got as a Farmer, how many you got as a Pirate and the total number of baby flowers. The computer will now sum up all the points. You can exchange these points for toys in the next room. Did you guys have fun today?

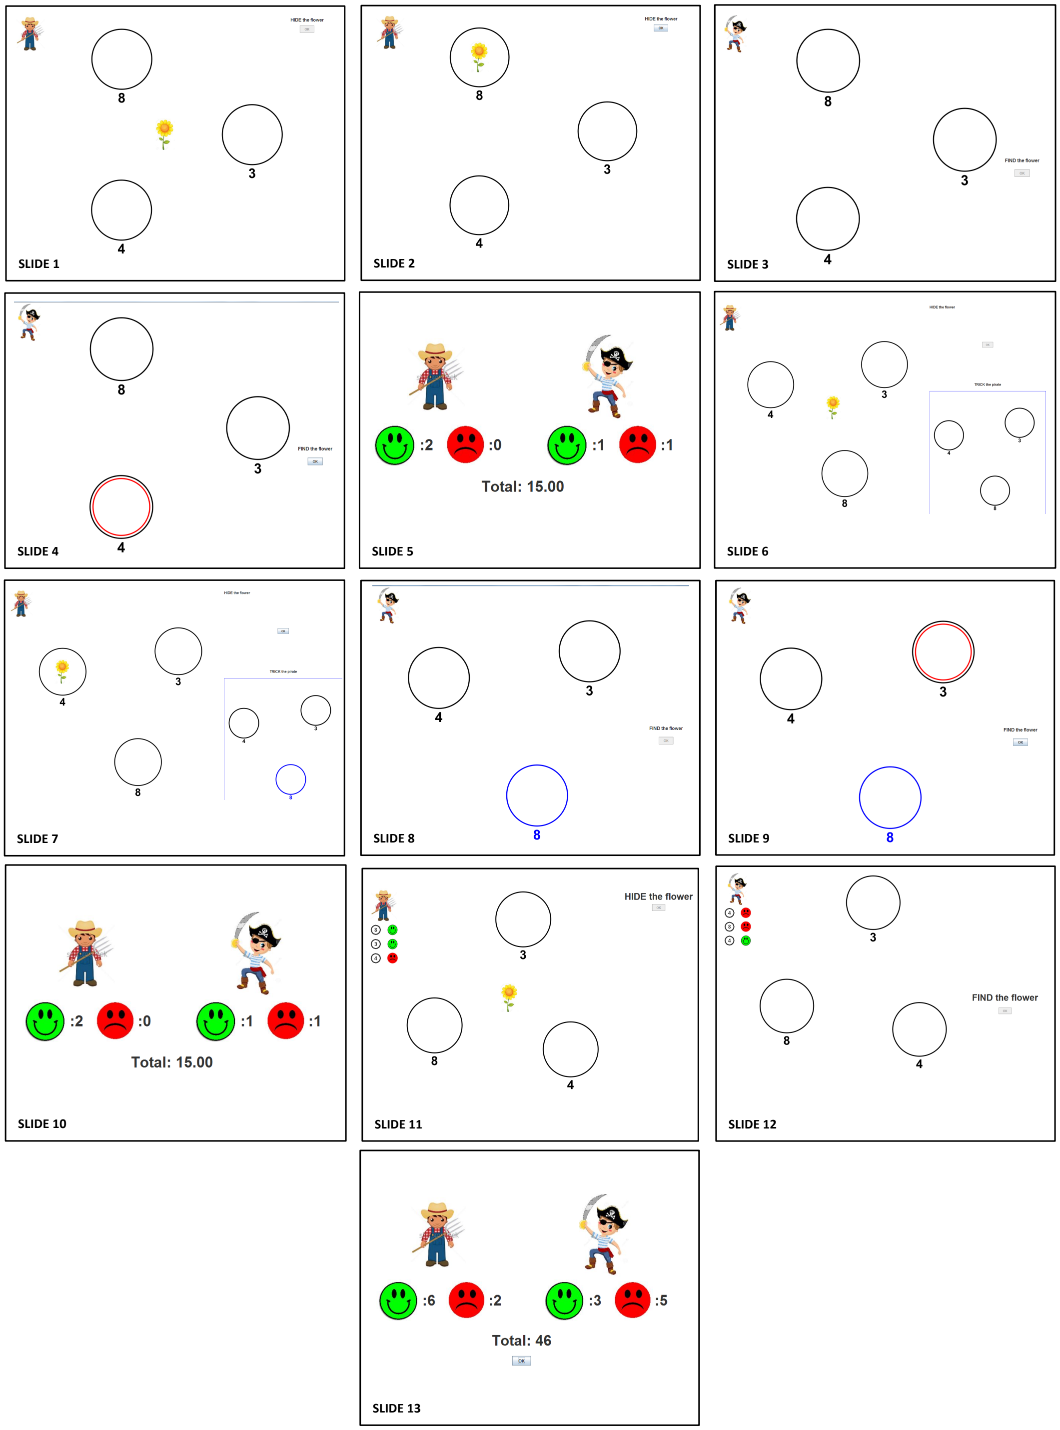

Figure 11: Slides projected on screen during instructions