Dynamic coordination in efficient and fair outcomes: a developmental perspective *

Isabelle Brocas
University of Southern California
IAST and CEPR

Juan D. Carrillo
University of Southern California
IAST and CEPR

December 2024

Abstract

We study in the laboratory age-related changes in the behavior of children and adolescents (ages 7 to 16) in two repeated coordination games, stag hunt and battle of the sexes. In these games, sharing intention and beliefs helps participants reach the efficient and fair long run outcome (EFO). In stag hunt, it means coordinating on the Pareto superior Nash equilibrium, hence a coordination of actions. In battle of the sexes, the exercise is arguably more complex as it requires taking turns between the two static Nash equilibria, hence a coordination of strategies. We find in both games a significant and remarkably stable increase in the ability to coordinate on the EFO with age. At the same time, the majority of participants in all ages adhere to one of a small number of relatively simple strategies. EFO is more prevalent in stag hunt and in the second supergame. This evidence suggests that children gradually learn how to share intentions and beliefs, an ability that can be exported to new interactions, but that is limited by game complexity. More generally, it suggests that dynamic cooperation is not instinctive or innate but rather reflective and acquired.

Keywords: developmental decision-making, coordination, repeated games.
JEL Classification: C73, C91.

*We are grateful to members of the Los Angeles Behavioral Economics Laboratory (LABEL) for their insights and comments and to the Lycée International de Los Angeles (LILA) –in particular Emmanuelle Acker, Nordine Bouriche and Anneli Harvey– for their help and support running the experiments in their school. We thank the seminar participants at University of Arkansas, University of Pittsburgh, New York University and the MiddExLab Virtual Seminar at Middlebury College, and the Editor and two referees of this journal. The study was conducted with the University of Southern California IRB approval UP-12-00528. We gratefully acknowledge financial support of the National Science Foundation (NSF) grant SES-2315770, and IAST funding from the French National Research Agency (ANR) under grant ANR-17-EURE-0010. Address for correspondence: Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089, USA, <brocas@usc.edu> and <juandc@usc.edu>.

1 Introduction

Strategic interactions often feature collaborative outcomes that may be socially desirable but require individual sacrifices. Game theory models multi-person interactions through abstract games that capture the main ingredients of collaborative relationships. These include the prisoner's dilemma, stag hunt and battle of the sexes, to name a few. The theoretical predictions in the one-shot version of these games are sharp. Unfortunately, the usefulness of theory is often limited in repeated games since, according to the folk theorem, every individually rational payoff can be sustained in equilibrium if the game is repeated with sufficiently high probability. At the same time, collaboration takes its true meaning. Repeating the cooperative outcome, selecting the Pareto efficient Nash equilibrium of the stage game, and alternating between the two Nash equilibria are both individually and socially desirable outcomes in repeated prisoner's dilemma, stag hunt and battle of the sexes, respectively. Are people capable of leveraging repetition to reach these outcomes? Laboratory experiments can be particularly useful to answer this question.

In the last decade, there has been a rapid development of the experimental literature on the repeated prisoner's dilemma, arguably the most natural game to study the tension between short term gain of deviation v. long term gain of cooperation.[1] By contrast, there is a surprisingly small experimental literature on repeated coordination games, where the static version has multiple Nash equilibria (McKelvey and Palfrey, 2001; Kuzmics et al., 2014; Mathevet and Romero, 2014). A common finding is the impressive ability to achieve long run coordination on fair (equal payoffs) and Pareto efficient outcomes in symmetric, two-person games where coordination in the one-shot game is infrequent.

An explanation that has received significant traction is that beneficial cooperation in humans is intuitive and spontaneous (Rand et al., 2012). Under this view, such innate behavior should be observed from an early age. Alternatively, one may view coordination as a reflective choice that is facilitated by the sharing of intentions and beliefs. This builds on three critical elements: (i) theory of mind (ToM), (ii) abstract logical reasoning, and (iii) reciprocal beliefs. Because these abilities develop gradually (see e.g., Wellman et al. (2001); Royzman et al. (2003); Rafetseder et al.

[1] See Dal Bó and Fréchette (2011); Camera et al. (2012); Fudenberg et al. (2012); Friedman and Oprea (2012); Romero and Rosokha (2018) for some representative examples out of a long list, and Dal Bó and Fréchette (2018) for a detailed survey.

(2013)), one would conjecture that coordination should improve steadily with age. Improvements should also be a function of the difficulty to coordinate effectively. Overall, a developmental approach to dynamic cooperation may help answering the question of whether coordination is instinctive and innate or reflective and acquired.

In this paper, we investigate in the laboratory age-related changes in behavior of children and adolescents (ages 7 to 16) in two repeated coordination games: stag hunt and battle of the sexes. Individuals play 2 supergames of each game, with 24 stages each. The study puts special emphasis in developing a methodology adequate for all ages, which requires novel story lines with attractive graphical interfaces.

Our main results can be summarized as follows. First, we observe a significant and remarkably steady improvement in coordination from young childhood (grades 2-3) to late adolescence (grades 8-10) and into young adulthood (college undergraduates). While this result may be unsurprising for some readers, one should notice that, in our past experience, we sometimes observe limited (e.g., Brocas and Carrillo (2021)) or no (e.g., Brocas and Carrillo (2022b)) improvements with age in games of strategy. These changes have large payoff consequences. Second, we find that the vast majority of participants (with the exception of the youngest subjects in their first supergame) adhere to one of a small number of identifiable strategies, generally avoiding the most intricate ones. In other words, excess complexity, which is often empirically suboptimal in repeated games, is not in the toolkit of our subjects, while simple strategies (some optimal and some not) are. While the set of strategies is similar in all grade-groups, the proportion of participants that uses each of them changes with age. Consequently, the changes with age in the payoffs secured by our participants is driven by the relative frequency with which each strategy is selected. Third, while the trajectory is similar in both games, the levels of coordination are not. Coordination is more prevalent in stag hunt than in battle of the sexes. There are two complementary reasons for such difference. First, stag hunt necessitates coordinating actions (both play stag in every period) while battle of the sexes necessitates coordinating strategies (alternate between the two Nash equilibria), which is arguably more challenging. Second, the natural tendency for young children to focus on the most salient features of the game (their own payoff) leads towards coordination in stag hunt (always play stag) but not in the battle of sexes (always play the favorite action). Finally, we also observe an ability to learn and signal. The decision in the first round of a supergame has a large impact on the likelihood to coordinate, and we observe significant improvements between the first and second supergames. All this suggests a capacity to learn and adapt very rapidly (after one exposure) and to bring any lesson to the next game with a new partner.

The paper contributes to the growing experimental economics literature that studies strategic choice in children. The seminal research focused on traditional games such as ultimatum, trust or public goods (see e.g., Harbaugh and Krause (2000); Harbaugh et al. (2003); Peters et al. (2004)). More recent studies analyze the developmental trajectory of strategic decision making and how it is impacted by cognitive development (see e.g., Sher et al. (2014); Brosig-Koch et al. (2015); Czermak et al. (2016); Brocas and Carrillo (2021); Fe et al. (2022)).[2] To our knowledge, Brocas et al. (2017) is the only paper in experimental economics that has formally studied the change in behavior from childhood to adulthood in a long repeated game, more precisely a 16 round alternating dictator game. The paper studies the incentives to initiate cooperation, reciprocate and forgive in dynamic relationships. As discussed above, we are interested in the related (but different) issue of dynamic *coordination* in games with multiple static Nash equilibria. Relatedly, Blake et al. (2015) and Bašić et al. (2021) concentrate on narrow grade-groups and study cooperation and reciprocity in short finite versions of the repeated prisoner's dilemma. Finally, one should not forget the extensive literature in developmental psychology that addresses related questions. Closest to our work are Grueneisen and Tomasello (2017, 2019) who propose a highly innovative design to study coordination in the chicken game. Again, the authors focus on a narrow range of young children (5 to 8 years old). They are mainly interested in determining which communication rules and strategies (threats, protests, compromises, etc.) lead to long run coordination. Finally, the sustained increase in performance with age matches the behavior in more complex games with asymmetric information (Brocas and Carrillo, 2022a).

The paper is organized as follows. Design and procedures are detailed in section 2 while theory and hypothesis are presented in section 3. Section 4 reports the descriptive analysis of choices and payoff, and section 5 details the classification of participants according to their best-fitted strategy. The contribution of demographic variables is investigated through regression analysis in section 6. Concluding remarks are presented in section 7. Appendix A compares our two benchmark adult populations (USC undergraduates and school teachers), Appendix B contains additional regression analyses, and Appendix C reports the instructions of the game.

[2] We refer the reader to Sutter et al. (2019) and List et al. (2023) for detailed surveys of the experimental economics literature on children and adolescents.

2 Experimental design

We investigate the behavior of children and adolescents (7 to 16 years old) in repeated games of complete and imperfect information. To account for the challenges inherent to the study of this age group, we follow the guidelines proposed by Brocas and Carrillo (2020b) and we develop a graphical version of existing games.[3]

Participants. Our main population consists of 220 school-age individuals from grades 2 to 10 at Lycée International of Los Angeles (LILA), a private school in Los Angeles. We ran 28 sessions that lasted no more than one school period (50 minutes) in June 2017. Sessions were conducted in a classroom at the school using PC tablets and the tasks were programmed with the software 'Multistage Games'. Sessions had 8, 10 or 12 participants. For each session, we tried to have a mix of male and female participants from the same grade, but for logistic reasons, we sometimes had to mix subjects of two consecutive grades (always from the same grade-group). High schoolers from grades 9, 11 and 12 did not participate in the study because they were taking or preparing for national exams. The majority of students at LILA are Americans and Europeans from caucasian families of upper-middle socio-economic status. A homogenous population allows us to make meaningful age comparisons. Indeed, in our previous research we have shown that variations in economic or demographic characteristics are associated with differences in performance in some games (Brocas and Carrillo, 2021) but not in others (Brañas Garza et al., 2023). Given the difficulty to obtain a large sample, avoiding a mix of participants from different schools reduces the number of confounding effects on the developmental trajectory. On the other hand, the pool is not representative of the US population, which limits the generalizability of our findings.

For comparison, we recruited 70 undergraduates at the University of Southern California (USC) and ran 6 sessions using *identical* procedures. With some exceptions (e.g., Cobo-Reyes et al. (2020)), studies with children do not include an adult population. We believe a control group is important to establish a behavioral benchmark and argue that USC is a reasonable match.[4] At the same time, the existing

[3]The relevant principles for this experiment are: (i) simplify the procedures given the participants' limited attention; (ii) offer age-appropriate incentives; (iii) present the task in a simple, graphical and attractive way; and (iv) include, if possible, a benchmark adult comparison group.

[4]After high school, a large fraction of students from LILA go to well-ranked colleges in North America and Europe, including USC and universities in the UC system.

differences between these two populations (nationality, family background, size of peer group, etc.) must be acknowledged.

Finally, we also had the (unexpected) opportunity to conduct the experiment with 30 teachers at LILA. We report a comparison of our two adult populations in Appendix A. Table 1 reports a summary of our 320 participants. For the analysis, we group our school-age participants into four naturally clustered grade-groups: grades 2-3, 4-5, 6-7 and 8-10. Standard errors in the regressions are clustered at the matched pair level.

	LILA								USC	Teachers
Grade	2nd	3rd	4th	5th	6th	7th	8th	10th	U	T
Age	7-8	8-9	9-10	10-11	11-12	12-13	13-14	15-16	18-23	n/a
# subjects	33	21	24	30	24	24	33	31	70	30

Table 1: Summary of participants

Tasks. The experiment had two tasks always performed in the same order.

In the first, every participant played the same four rounds of a three-option dictator game, with different payoff combinations and different partners. Recipients in the four rounds were four anonymous individuals from the same session. To avoid cross-contamination, the aggregate outcome of this short task was only communicated to participants at the end of the experiment and, as announced during the experiment, the identity of dictators and recipients was never communicated. We treated the dictator game as separate and independent from the task in this study. We did not expect or had any hypothesis on a possible relationship between the two, so we did not analyze them in combination. The results of the first task are reported in Brocas and Carrillo (2020a).[5]

After a break, participants moved to the main task with new anonymous partners. It consisted of two repeated battle of the sexes (**BoS**) and two repeated stag hunt (**SH**) supergames with symmetric payoffs. Table 2 shows the normal-form representation of the stage game. We report both the actions as chosen by the participants –{red, green} and {in, out}– as well as the notation adopted in our analysis –$\{M_i, Y_i\}$ and $\{I_i, O_i\}$– as explained in section 3.1. These two games share many

[5]Given the significant difficulty to access a population of children, we sometimes conduct two separate experiments in the same session.

similarities. In particular, both are symmetric 2×2 games with two pure strategy Nash equilibria in the stage game. At the same time, and as discussed later, dynamic coordination is expected to be simpler when there is a Pareto superior static Nash equilibrium (**SH**) than when there is not (**BoS**).

BoS

	red (Y_2)	green (M_2)
red (M_1)	(5,3)	(1,1)
green (Y_1)	(1,1)	(3,5)

SH

	in (I_2)	out (O_2)
in (I_1)	(3,3)	(1,2)
out (O_1)	(2,1)	(2,2)

Table 2: Normal form representation of the battle of the sexes and stag hunt games

We avoided null payoffs and, for mathematical ease, considered simple numerical values. In **BoS**, we made sure that mis-coordination was sufficiently costly compared to coordination in the least desirable equilibrium (1 vs. 3) and that coordination in the least desirable equilibrium was also sufficiently costly compared to coordination in the most desirable equilibrium (3 vs. 5). In **SH**, we made the least risky strategy riskless for expositional ease. We also made sure that the efficient equilibrium was not overly rewarding to avoid salience effects.[6]

The structure in each supergame was identical. Subjects were randomly and anonymously matched with a partner and played 24 rounds of the game with the same partner and feedback after each round. At the end of the supergame, total payoffs were displayed. New partners were randomly and anonymously drawn and a new supergame was played. For each age-group, participants played two **BoS** supergames followed by two **SH** supergames in approximately half the sessions, and two **SH** followed by two **BoS** in the remaining sessions.[7] To highlight that the horizon was long, we did not announce the number of rounds (the instructions said: "you will play many rounds with the same partner"). However, we used the same length (24 rounds) in all four supergames, so some subjects could potentially become aware of it. It is important to remember that, unlike in the prisoner's

[6]Formally, the basin of attraction is 1/2. Dal Bó et al. (2021) have shown that with these payoff values neither choice is overwhelmingly favored by adults in the one-shot version of the game.

[7]Length of the experiment is a major constraint in developmental studies due to the limited attention span of participants (Brocas and Carrillo, 2020b). In the absence of constraints on time and attention, we would have ideally liked to run more supergames to study learning patterns and to better disentangle between individual strategies that result in identical outcomes.

dilemma where a finite repetition results in a unique Subgame Perfect Equilibrium, in our coordination games the limit perfect "folk theorem" holds: any feasible and individually rational payoff vector of the stage-game is achievable in the finitely repeated game as the time horizon gets sufficiently large (Benoit and Krishna, 1985).

We were very concerned with the possibility that differences in behavior across ages could reflect differences in task comprehension. It was therefore of paramount importance to provide a simple, graphical interface and a story accessible and appealing to young children. This ruled out the payoff matrices presented in Table 2 as well as other formal representations commonly employed.

For **BoS**, we developed a novel story called the "find the balance" game. Each of the two matched participants in a group was assigned a role, red or green. The red player possessed a red scale and the green player possessed a green scale. They also possessed one ball each that they had to simultaneously place in one of the scales. If both participants placed their balls on the same scale, the scale was balanced. The owner of the scale earned 5 points and the other player earned 3 points. If they placed their balls on different scales, the scales would be unbalanced, and players earned 1 point each.

Figure 1a provides a screenshot. The role (red) was displayed at the top. The player had to place the ball on the red or the green scale by tapping on the corresponding dotted circle. The "?" sign described the possible choices of the other player. The right-side of the screen displayed the history of the supergame (here, the first 4 rounds), including the choices of both players and the points earned by the player in each round. This panel filled up in real time as the supergame progressed. For reference, a screen in the front of the room displayed the payoffs of both individuals for each combination of choices, as represented in Figure 1b. This information remained visible during the 48 rounds of the two **BoS** supergames.

For **SH**, we developed a novel story called "risky stars". A blue and a yellow player possessed a blue and a yellow star. Each decided whether to place their star on or outside a common carpet. Placing the star outside the carpet gave 2 points. Placing it on the carpet gave the player 3 points if the other player also placed the star on the carpet and 1 point if the other player placed the star outside the carpet.

Figure 2a provides a screenshot of the **SH** supergame, with the carpet represented by a rectangle, and the right panel describing the history of the first four choices. Just like before, a screen in the front of the room displayed the payoffs of both players for each combination of choices (Figure 2b). A transcript of the read

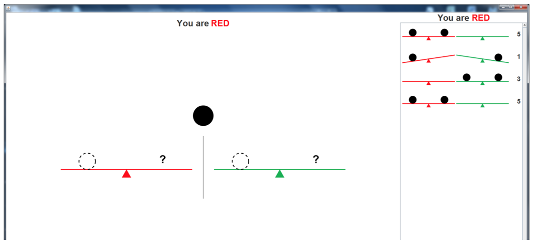

(a) BoS screenshot

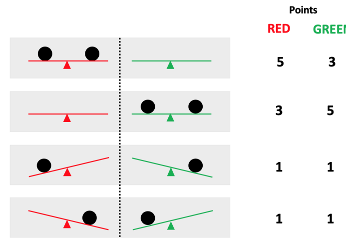

(b) BoS payoff display

Figure 1: Experimental design of **BoS**

aloud instructions is included in Appendix C.

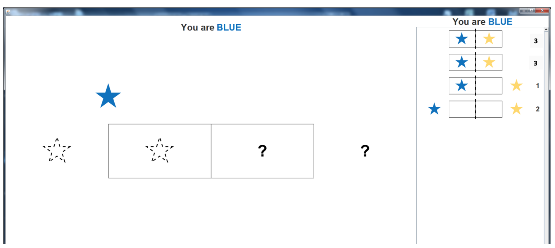

(a) SH screenshot

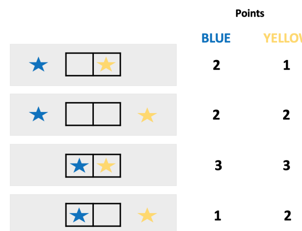

(b) SH payoff display

Figure 2: Experimental design of **SH**

Payoffs. During the experiment, subjects accumulated points. Following Brocas and Carrillo (2020b), we used different mediums of payment for different ages. The objective was to try and equalize the *value of rewards* across grade-groups instead of equalizing the rewards themselves.[8] USC students, teachers and participants in grades 6 to 10 earned points that were converted into money at $0.04 per point, and paid after the experiment in cash (USC) or with an e-giftcard (school students and teachers, since the school does not allow money on premises). USC students and teachers were also paid a $7 show-up fee, to correct for differences in opportunity cost of time. Average payoffs for this section of the experiment (not including show-up fees) were $12.52 (USC), $12.02 (teachers) and $11.74 (grades 6 to 10).

[8]Money is usually the most adequate medium of payment precisely because it is valued most similarly by participants. However, this is not the case when age is a factor. Young children strictly prefer desirable objects for their immediate enjoyment rather than the equivalent amount of money, which they appreciate, but it is likely to be administered by the parents.

For children in grades 2 to 5 we set up a shop with 20 to 25 pre-screened, age appropriate toys and stationary (bracelets, erasers, figurines, apps, earbuds, etc.). Different toys were worth different point prices. Before the experiment, children were taken to the shop and showed the toys they were playing for. They were also instructed about the point price of each toy. At the end of the experiment, participants learned their point earnings and were accompanied to the shop to exchange points for toys. We made sure that every child earned enough tokens to obtain at least three toys. At the same time, points were scarce and valuable.[9] We spent an average of $4 in toys per child. At the end of the experiment, we also collected demographic information consisting of "gender", "grade", and "number of siblings".

3 Theory and Hypotheses

3.1 The cognitive complexity of coordination games

Playing at equilibrium in one-shot games is cognitively complex, and necessitates a number of abilities that develop gradually. It requires a player to realize that another person is involved in the game, and to model the ability and best interest of that other player. Young children appear egocentric (Piaget et al., 1967; Crain, 2015) and mostly pay attention to their own play and payoffs in a strategic interaction. With the development of theory of mind during elementary school –the mental capacity to understand other people's behavior, intentions and beliefs– children become gradually able to recognize strategic implications (Wellman et al., 2001). However, game theoretic paradigms also require hypothetical and counterfactual thinking abilities, which are known to develop throughout middle school (Rafetseder et al., 2013). Finally, one-shot coordination games also requires modeling the intention of the other player to target one equilibrium. This necessitates an extraordinary ability to share mutual beliefs, a higher-level form of theory of mind. Coordination in these one-shot games is challenging even among educated adults (Camerer, 2011). Therefore, we should not expect children to succeed either.

At the same time, the previous literature has demonstrated that adults coordinate remarkably well their choices on Pareto optimal equilibria in the repeated version of these games (McKelvey and Palfrey, 2001; Kuzmics et al., 2014; Math-

[9] While incentives are key to retain the attention of children, it is also important to avoid excessively high variance in payoffs to make sure that no child feels unhappy.

evet and Romero, 2014). This is particularly interesting given the little predictive power of the standard theory (Benoit and Krishna, 1985). It contrasts with other repeated games (for example, the prisoner's dilemma) where empirical behavior is highly heterogeneous (Dal Bó and Fréchette, 2018). The behavior of adults thus provides a stark template for comparison.

To understand the developing ability to coordinate, it is also important to exploit differences across games. In **BoS**, we define actions symmetrically for both players (relative to their most and least favorite actions). We denote by M_i the choice by player i of 'my' favorite action ('red' for red player and 'green' for green player). Similarly, Y_i is the choice of 'your' favorite action ('green' for red player and 'red' for green player). From Table 2, the static Nash equilibria are therefore (M_1, Y_2) and (Y_1, M_2). For **SH**, we denote by I_i and O_i the 'stag' and 'hare' choices by player i ('in' and 'out'). From Table 2, the static Nash equilibria are (I_1, I_2) and (O_1, O_2).

Let us now focus on the Subgame Perfect Equilibrium that is Pareto optimal and gives equal payoff to both players, which we will thereafter refer to as the *Efficient and Fair Outcome* (EFO). In **BoS**, EFO entails alternating between the two static Nash equilibria: (M_1, Y_2) at t, (Y_1, M_2) at $t + 1$, etc.[10] In **SH**, EFO entails the repetition of the static Pareto superior Nash equilibrium: (I_1, I_2) at t, $t + 1$, etc. Notice that to reach EFO, participants need to coordinate their strategies in **BoS** whereas they only need to coordinate their actions in **SH**, with the former being arguably more challenging than the latter.[11]

The issue is to determine how participants of different age devise strategies to reach EFO. In our empirical analysis (section 5), we discuss some possibilities.

3.2 Hypotheses

Centration, theory-of-mind and logical thinking–three features that change significantly during our window of observation–are likely to affect behavior. The tendency to focus mostly on features that affect oneself decreases gradually until age

[10] Any strategy where players coordinate half the time on each static equilibrium would result in the same payoff. We hypothesized (and empirically verified) that such strategies would not be played in our game, so we did not consider them in our analysis.

[11] Not surprisingly, McKelvey and Palfrey (2001) finds very significant differences when coordination games are played with random partners (see also Camera and Casari (2009); Camera et al. (2013) on prisoner's dilemma with random partners). Fixing partners is important in our game as it allows us to study the development of mutually shared beliefs.

11 (Miller, 2002). Also, while children master the most basic false belief ToM tasks by age 5 and some 6 years old children can already understand second-order false beliefs (Grueneisen et al., 2015), the more subtle aspects of this general ability continue to develop throughout adolescence (Royzman et al., 2003). Finally, strategic thinking improves all the way into adulthood and, as witnessed in numerous economic experiments, sometimes never matures fully.

These findings suggest that behavior is likely to be closer to equilibrium as participants get older. However, its exact significance is not fully clear. For one thing, the predictive power of theory is limited. If very different paths can be rationalized as equilibrium behavior, it becomes difficult to provide an objective metric for deviations. For another, the existing research suggests that we can observe very different age-related changes in behavior and payoffs depending on the structure of the game. With these considerations in mind, we next provide some hypotheses about the evolution in behavior of our participants.

Hypothesis H1. *As they get older, participants make less frequently decisions that ignore the behavior of others and more frequently decisions that respond to and prompt their cooperation.*

Hypothesis H2. *As they get older, participants employ more complex strategies.*

Hypothesis H3. *As they get older, participants are more successful in reaching the EFO and obtain higher payoffs.*

Hypothesis H4. *Participants of all ages are more likely to reach the EFO in* **SH** *than in* **BoS**.

According to **H1**, we expect that older participants will replace self-centered behavior with strategies that take into consideration the choice of others. **H2** predicts that older participants will adopt a more complex system of strategies, where "complexity" will be defined later but it roughly consists in decisions that depend on more dimensions of the game. **H3** argues that the combination of the previous hypotheses will result in more frequent and more efficient coordination, and therefore higher gains, for older participants. Finally, since EFO requires participants to coordinate their strategy in **BoS** and coordinate their action in **SH**, **H4** predicts a higher rate of success in the latter game than in the former. Our empirical analysis will study whether these natural age-related and game-related hypotheses are supported by the data.

4 Descriptive analysis

We first report in Figure 3 some aggregate statistics in each grade-group and each supergame, averaged over the 24 rounds, for **BoS** (top) and **SH** (bottom). The left graphs display the proportion of M_i and I_i choices. The center graphs display the average proportion of (M_i, Y_j) and (I_1, I_2) outcomes by pairs of individuals. The right graphs display the average individual payoff, with a range going from random choice (2.5 in **BoS** and 2.0 in **SH**) to EFO (4.0 in **BoS** and 3.0 in **SH**).

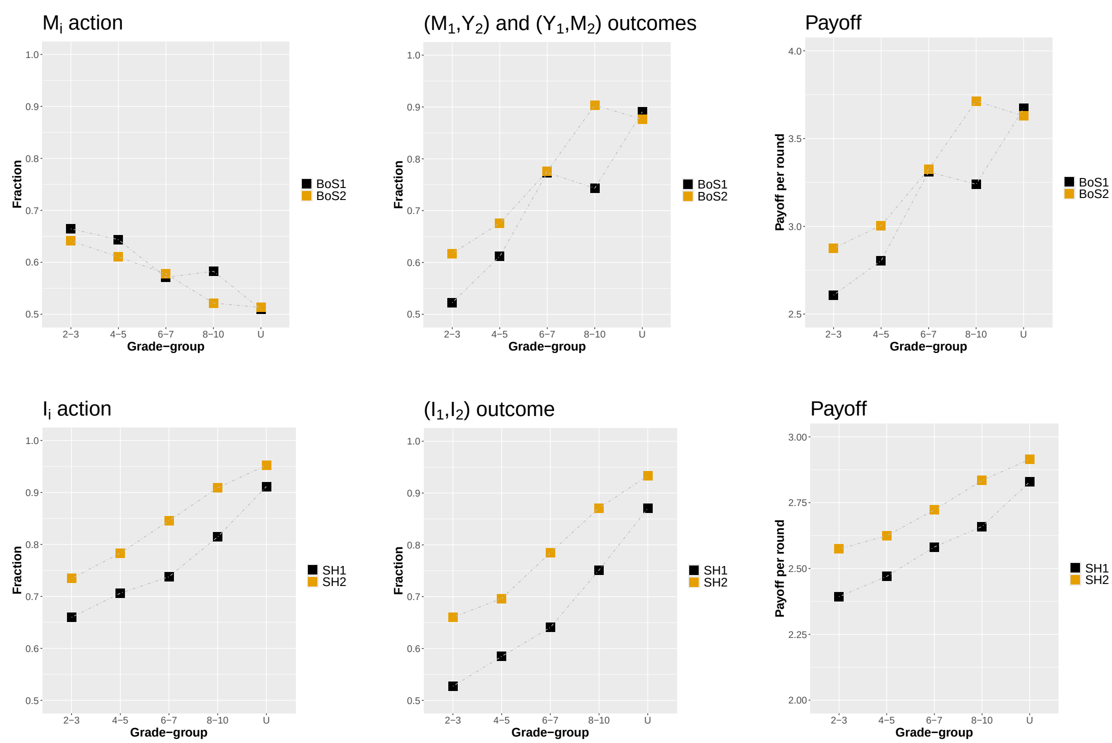

Figure 3: Aggregate behavior and earnings in **BoS** (top) and **SH** (bottom)

Figure 3 is illustrative of the main results that will be emphasized all along the paper. First, and in line with earlier literature, the adult control group achieves almost perfect coordination on one of the Nash equilibria in **BoS** and on the Pareto efficient equilibrium in **SH**, in both supergames (between 87.6% and 93.3% of coordination). They leave very little money on the table (earnings between 90.8% and 97.0% of the group maximum), and therefore provide a sharp template for comparison. Second, there is a significant and remarkably steady improvement in behavior with age. Coordination in (M_i, Y_j) and (I_1, I_2) in our school-age population starts poorly and steadily increases to levels similar to those in the adult population. The increase in coordination in **BoS** with age is due in part to the decreased tendency to choose the favorite option (top left graph). It is also reflected in the payoffs obtained by our participants. Third, there is improvement between the first and second supergames among school-age participants, with the exception of 6-7 in **BoS**. This indicates that participants learn and leverage their experience to improve their strategies and payoffs. There is also an improvement, though less dramatic, between BoS1 and BoS2. Finally, the behavior of individuals is strongly correlated across supergames (Pearson Correlation Coefficient, PCC = 0.64 for $\Pr(M_i)$ and PCC = 0.62 for $\Pr(I_i)$, $p < 0.0001$). Consequently, payoffs are also highly correlated across supergames (PCC = 0.43 in **BoS** and 0.44 in **SH**, $p < 0.0001$).

We next study the dynamics of outcomes. Figure 4 presents the change in the proportion of groups that achieve coordination on (M_i, Y_j) in **BoS** (top) and on (I_1, I_2) in **SH** (bottom) from rounds 1 to 24 in each supergame and grade-group.

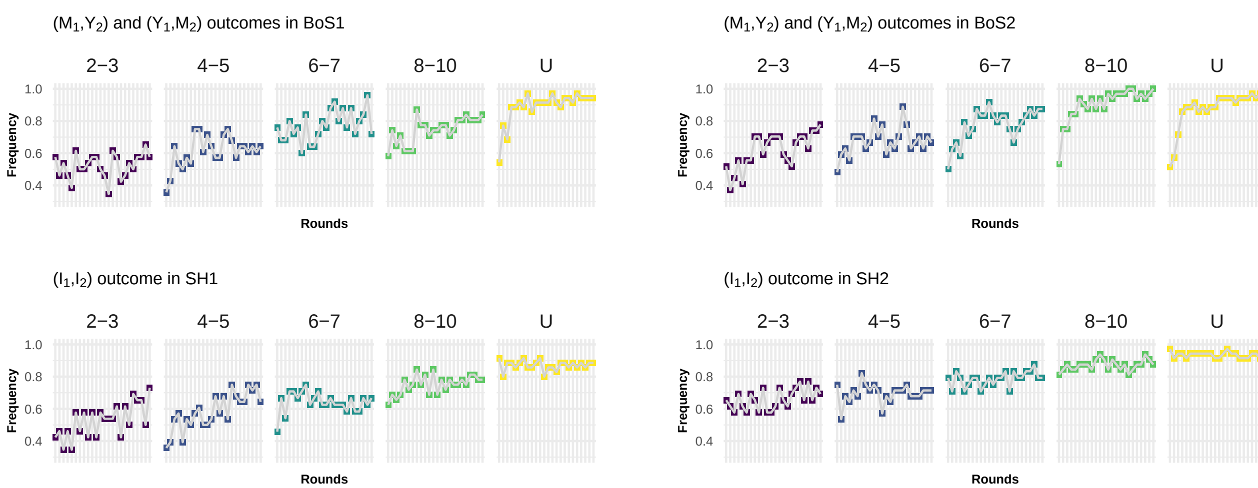

Figure 4: Coordination over rounds in **BoS** (top) and **SH** (bottom)

Not surprisingly given Figure 3 (middle), we find a sustained increase in the level of coordination across grade-groups. However, there is limited evidence of increased coordination across rounds within grade-groups. In **BoS**, we observe no significant trend in the younger grade-groups (augmented Dickey-Fuller test, $p > 0.05$, although this may be due to a lack of statistical power). There is a positive trend in the first supergame for 8-10 and U, which is driven by the first few rounds. Indeed, initial miscoordination is frequent but it is often solved quickly. In **SH**, coordination increases over time for our younger participants (2-3 and 4-5) in the first supergame (augmented Dickey-Fuller test, $p < 0.05$). Coordination in our older school-age participants and control group starts at a high level and remains constant.

The next question is whether the willingness and ability to coordinate lead to convergence to EFO. Figure 5 reports the distribution of rounds at which convergence to EFO is reached by grade-group and supergame. For each group, we determine the round after which the outcome coincides with EFO for the remaining of the supergame. The bar at the extreme left corresponds to the fraction of groups that coordinate from the outset whereas the bar at the extreme right corresponds to the fraction of groups that never coordinate.

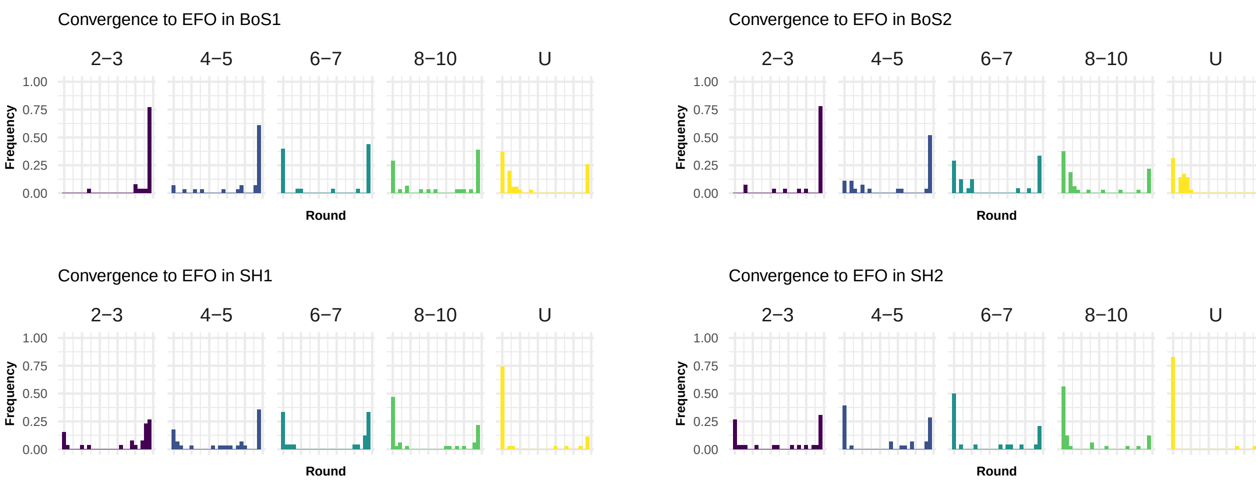

Figure 5: Round of convergence to EFO in **BoS** (top) and **SH** (bottom)

Stage of convergence is bimodal both in **BoS** and in **SH**: groups either coordinate on EFO in an early round or they do not coordinate at all. As participants get older, the fraction of early coordination grows and that of late or no coordination shrinks. Early coordination within a grade-group is more frequent in **SH** than in **BoS**, especially in the second supergame. Also, for groups that manage to coordinate early, it often takes a few more rounds of miscoordination in **BoS** than in **SH** and in SH1 than in SH2. This difference between **BoS** and **SH** is particularly noticeable in our adult control group. Bimodality is consistent with the lack of evidence of an overall improvement in coordination within supergames, as reported in Figure 4.

Finally, it is also instructive to take a closer look at the payoffs obtained by our participants. Figure 6 reports average earnings over the 24 rounds by supergame and grade-group. A dot represents the payoff of a pair of subjects, with the diameter being proportional to the number of pairs with that combination of earnings. For visual ease, we always report in the x-axis the player in the pair with highest gains. The set of attainable average payoffs is delimited by the gray segments.

In **BoS**, groups are more likely to reach a payoff close to $(4, 4)$, the average earnings of the EFO, as they age. This is a strong indication that older school-age participants are better at coordinating their strategy than their younger peers. Systematic miscoordination (payoffs close to (1,1)) and asymmetric outcomes in the frontier set (where one subjects always play M and the other does not) are common in younger children. Behavior in **SH** is less heterogeneous than in **BoS**. A significant fraction of payoffs are concentrated around (3,3), the EFO, especially in the older grade-groups. Behavior is mostly symmetric and payoffs in the neighborhood of $(2, 2)$ (the other static Nash equilibrium) accounts for the second largest fraction of choices. There are extremely few instances of players persevering in I when the partner has decided to play O (area below and to the left of (2,2)).

To sum up, the descriptive analysis has delivered the following conclusions: (i) there is a sustained increase in (M_i, Y_j) and (I_1, I_2) outcomes with age; (ii) age-related improvements in coordination translate into higher earnings; (iii) (M_i, Y_j) and (I_1, I_2) outcomes are more frequent and earnings are higher in the second supergame; (iv) most pairs coordinate on the EFO either early in the supergame or not at all; and (v) as participants get older, the fraction of early coordination on EFO grows and that of late or no coordination shrinks. More generally, the fact that participants become more likely to coordinate (on a static Nash equilibrium) and, most importantly, that they coordinate better on the EFO with age, provides support for **H3**. Nevertheless, only increased converge with age to EFO in **BoS** demonstrates that they move away from centered behavior as they grow older. Indeed, and as noted earlier, centration promotes efficient behavior in **SH** but it does not in **BoS**. This transition is evidence in favor of **H1**.

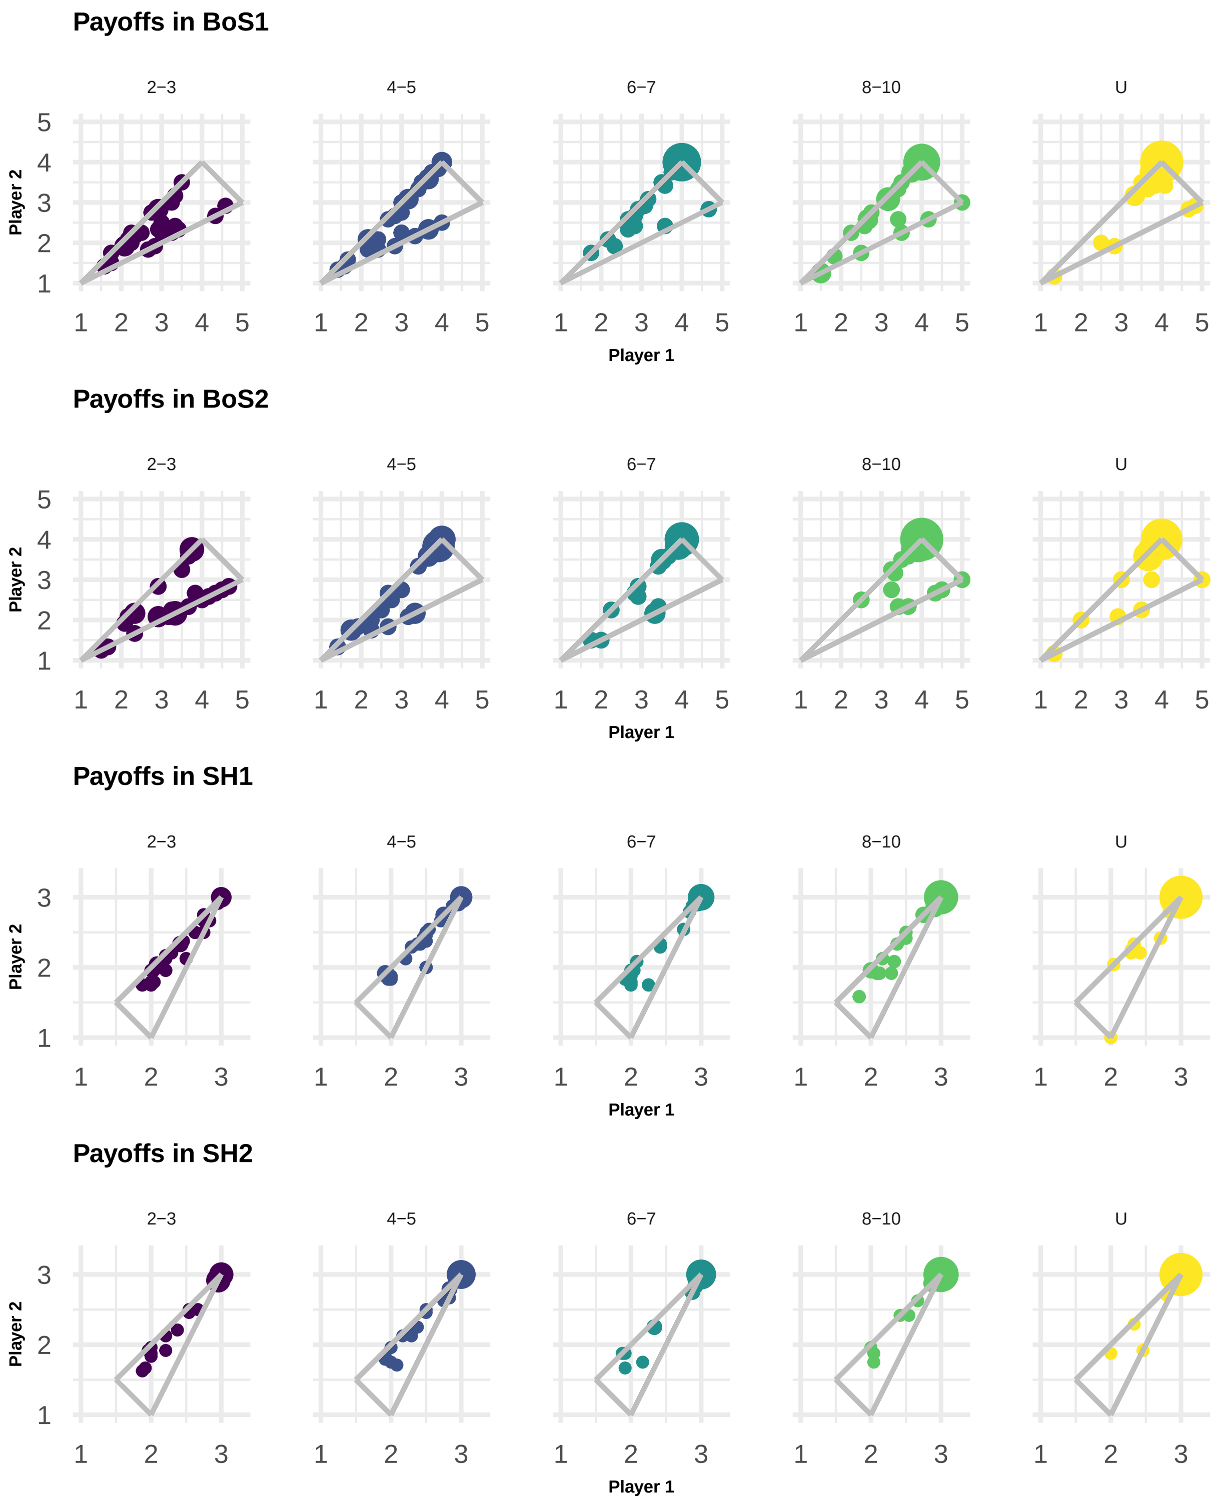

Figure 6: Average earnings of each pair of individuals in **BoS** and **SH**

5 Individual strategies

Descriptive analysis provides an overall picture of the main trends. However, strategic behavior in game theoretic paradigms is usually heterogenous. We investigate this heterogeneity and classify our participants according to their strategy in each supergame. Our methodology is as follows. For each supergame, we study the 20 choices observed in rounds 5 to 24. We conjecture that choices in rounds 1 to 3 are short term explorations and therefore ignore them. We use the outcome in round 4 as the anchor (or initial condition) for the strategy, and consider a number of potential strategies. We then assign to each player the strategy for which the number of deviations in rounds 5 to 24 is smallest, provided it is no greater than 3. If the number of deviations is the same for two or more strategies, we classify the subject at the intersection. Players exhibiting more than 3 deviations from all the strategies are unclassified.[12] In our empirical analysis, we will leave the subject's initial choice unspecified for some strategies. This is not theoretically rigorous but is consistent with using round 4 as the anchor of the empirical analysis.

Note that the behavior of a player in a supergame may be compatible with several strategies. With a large number of supergames, one could disentangle between different strategies by studying the behavior against different partners, adapting some of the sophisticated techniques developed in the repeated prisoner's dilemma literature.[13] Unfortunately, such analysis is not possible with only two supergames. On the other hand, studying overlaps of strategies will be informative about the behavior of a participant, as we will discuss below. We also did not find adequate to impose the same strategy across supergames. While such approach is suitable for studying steady-state behavior when the number of supergames is large, it seems inappropriate with only two supergames and a very real possibility that individuals deliberately change strategies between them.

Finally, and unlike in the prisoner's dilemma, the literature in coordination games does not provide clear guidelines regarding which strategies capture best the decision process of individuals. Our goal here is not to provide an exhaustive taxonomy, but

[12]Results are similar if we use rounds 3 or 5 as anchor (instead of 4) and/or if we allow 2 or 4 deviations (instead of 3). These robustness checks are generally omitted for brevity (except when noted) but they are available upon request.

[13]For example, Camera et al. (2012) trades-off goodness of fit and number of strategies, Aoyagi and Fréchette (2009) conduct Maximum Likelihood Estimation of best fitting strategies and Romero and Rosokha (2019) perform a direct elicitation of strategies.

to discuss some plausible alternatives. We focus on 'simple' strategies that may be 'focal' and *could* be employed by our participants.[14]

5.1 Strategies in Battle of the Sexes

5.1.1 Choosing strategies

Remember that EFO in **BoS** consists in alternating between the two Nash Equilibria: $(M_1^t, Y_2^t), (Y_1^{t+1}, M_2^{t+1}), (M_1^{t+2}, Y_2^{t+2})$, etc. (where we use the superpscript t for round t). This results in an average per-round payoff of 4 for each player. We are interested in strategies that can help sustain EFO, but also in strategies that may be chosen intuitively even though they do not result in EFO. Table 3 reports some possible strategies from simplest to most sophisticated.

strategy	description
(1) ME	play always M_i^t
(2) YOU	play always Y_i^t
(3) ALT	alternate between M_i^t and Y_i^t
(4) TFT	tit-for-tat: replicate the action of the partner in the previous round
(5) TRIG	grim-trigger: play the action consistent with EFO if all past outcomes are consistent with EFO and play M_i^t forever otherwise
(6) REV	reverse tit-for-tat: reverse the choice of the partner in the previous round
(7) FORG	forgiving trigger: play M_i^t unless the last round outcome was (M_i^{t-1}, Y_j^{t-1})
(8) TEACH	play Y_i^t unless the last round outcome was (Y_i^{t-1}, M_j^{t-1})
(9) TEST	play M_i^t unless the last round outcome was (Y_i^{t-1}, M_j^{t-1})

Table 3: Some simple strategies in **BoS**

Strategies (1)-(2)-(3) can be played by naïve players with little understanding of the partners' incentives as well as by strategic players whose objective is to reach an equilibrium (insisting on the best possible for themselves, agreeing on the best for the partner, or targeting the EFO, respectively). Strategies (4)-(5) are typical in other games and may result in EFO but also collapse into (M_1, M_2) depending on the partner's behavior, while (6) seeks to repeatedly coordinate in the same static Nash equilibrium (either always exploiting the partner or always giving-in). The

[14] Simplicity and focality are important characteristics of strategies in the coordination literature on static games (Cooper and Weber, 2020) as well as repeated games (Kuzmics et al., 2014).

remaining strategies capture a variety of strategic behaviors: (7) is similar to (5) except that it forgives after one period, (8) attempts to teach EFO by playing Y after a deviation, and (9) is the opposite of (8) and similar to (6), in that it attempts to exploit partners but gives in to selfish ones.

5.1.2 Empirical behavior

Classifying the maximum number of individuals with the fewest number of strategies can be delicate and subjective. In our case, however, it turned out to be relatively uncontroversial. Indeed, with only four strategies–ME, ALT, TFT and TEST–we can account for the choices of 72.8% and 80.3% in BoS1 and BoS2, respectively.[15] Furthermore, including all five remaining strategies would classify only an additional 3 participants in BoS1 and 3 participants in BoS2. We therefore decided to not include those strategies. As noted before, different strategies may lead to the same choices. Figure 7 provides a Venn diagram describing the overlap between strategies.

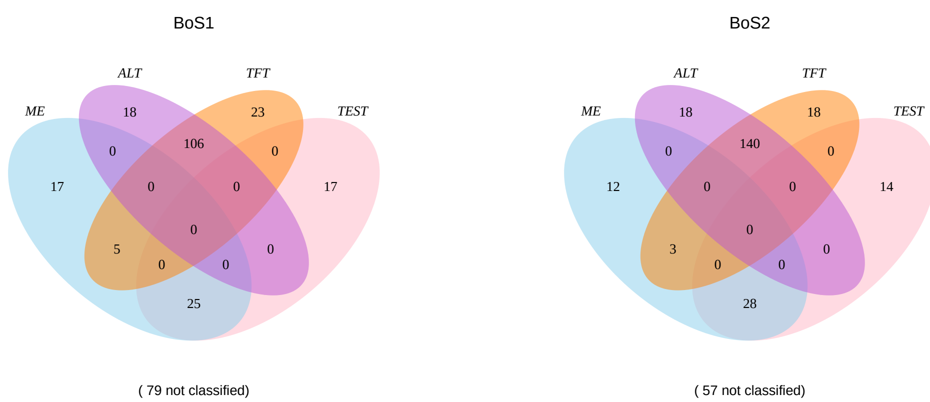

Figure 7: Overlap between strategies in BoS1 (left) and BoS2 (right)

The strategies that overlap the most are ALT/TFT. Players in that intersection belong to groups that successfully coordinate on EFO, since they alternate actions *and* replicate the past action of the partner. These participants differ from individuals classified only as ALT and individuals classified only as TFT. The former are participants who have suffered deviations of their partners. They tended not to punish those deviations and continued alternating, so they typically do not reach EFO.

[15] Allowing instead a maximum of 2 and 4 deviations respectively would classify 64.5% and 77.9% participants in BoS1 and 76.6% and 85.9% participants in BoS2.

The latter are participants who punished deviations more systematically, revealing a strategic tendency. They reach EFO only when they are matched with someone who does not deviate often from EFO, which is also infrequent.

Strategies ME/TFT as well as ME/TEST also overlap. The former are observed in a small number of players who are potentially strategic but they have faced a partner who always chose M. By contrast, the latter are observed in players who have always played M *despite* facing a partner who sometimes played Y. In this respect, the intersections of the Venn diagram are very revealing as they partially separate different motives for identical behaviors. Indeed, just like ME/TEST, ME comprises individuals who acted selfishly while facing potential cooperators. Finally, TEST captures individuals who give in to a selfish partner (consistently play Y against M) or exhibit long streaks of identical behavior. It is an interesting and relatively sophisticated strategy played by a significant number of participants. Figure 8 presents the distribution of strategies in each grade-group.

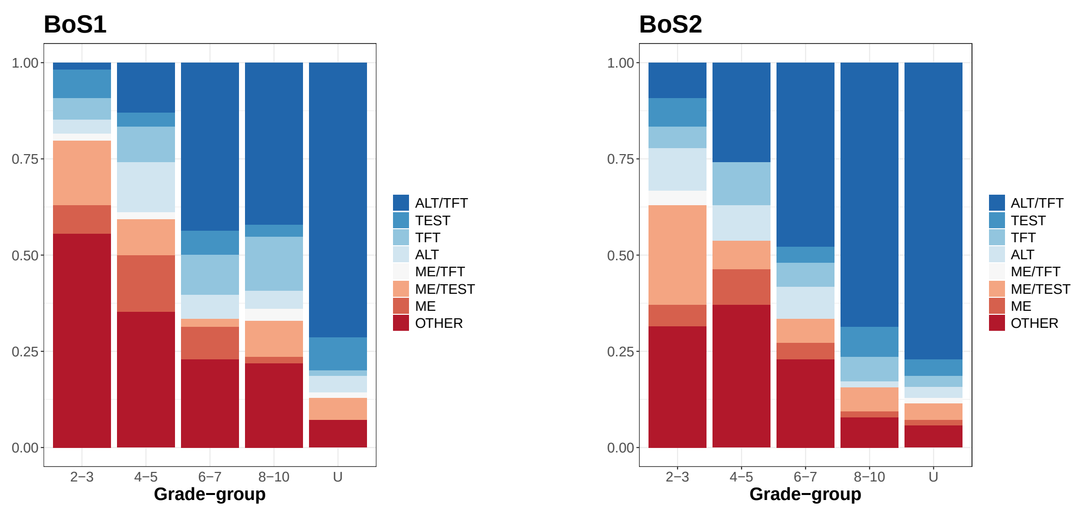

Figure 8: Distribution of strategies across grade-groups in BoS1 and BoS2

Perhaps not surprisingly in light of section 4, coordination in EFO (ALT/TFT) increases very significantly with age and from the first to the second supergame, with numbers ranging from 1.9% in 2-3 BoS1 to 68.8% in 8-10 BoS2. Conversely, selfish behavior (ME and ME/TEST) and unclassified players (OTHER) decrease with age. It suggests that many of our youngest participants have problems devising a strategy. This difficulty diminishes both with age and after experiencing one supergame. Tit-for-tat (TFT), streaks of identical choices (TEST) and alternation that does not lead to EFO (ALT) occur sporadically in all grade-groups. Finally, the improvement in the selection of strategies is also apparent. Indeed, the distribution of strategies in a given group in BoS2 is often similar to the distribution of an older group in BoS1 (2-3 in BoS2 is similar to 4-5 in BoS1 and 8-10 in BoS2 is similar to U in BoS1, chi-square tests, p-values > 0.05).

The choice of strategy has important payoff consequences. Table 4 reports the average per-round payoff in rounds 5 to 24 of each supergame as a function of the strategy employed by the player (as well as the overall per-round payoff "all").

	ALT/TFT	TEST	TFT	ALT	ME/TFT	ME/TEST	ME	OTHER	**all**
BoS1	3.98	3.08	3.18	3.16	1.22	3.35	2.36	2.49	3.21
	(.01)	(.13)	(.12)	(.16)	(.04)	(.26)	(.17)	(.06)	(0.05)
BoS2	3.98	2.91	3.64	3.54	1.20	3.57	2.52	2.39	3.44
	(.01)	(.11)	(.10)	(.13)	(.06)	(.21)	(.29)	(.07)	(0.05)

(standard errors in parenthesis)

Table 4: Average payoffs as a function of the strategy in **BoS**

Corroborating previous findings, ALT/TFT are associated to participants who coordinate on EFO (average payoff close to 4) while ME/TFT are associated to groups where both partners choose the selfish action (average payoff close to 1). Meanwhile, participants classified as ME and especially ME/TEST manage to exploit their partner in some rounds, and yet they still do worse than under joint collaboration. Participants classified as ALT and TFT attempt to reach EFO, but do not succeed in a number of rounds, with the corresponding payoff decrease due to miscoordination, especially in BoS1. Those classified as TEST obtain a similar (if slightly lower) payoff than those under the previous two strategies. However, this occurs through a different channel, namely by giving in and playing with high frequency the partner's preferred equilibrium. Using a non-discernible pattern yields only a slightly lower payoff than the expected payoff of all the players who consistently choose M. Last, but importantly, average payoffs within strategies are quite similar in both supergames. The observed average payoff differences across supergames (3.21 vs. 3.44, t-test, p = 0.001) are mainly driven by changes in the proportion of participants who are classified under the different strategies.

5.2 Strategies in Stag Hunt

5.2.1 Choosing strategies

For **SH**, EFO is the repetition of (I_1^t, I_2^t) for all t. The existence of a Pareto superior static Nash equilibrium (I_1, I_2) implies that coordination is arguably simpler and more intuitive than in **BoS**. We present a set of strategies that may be empirically relevant.

strategy	description
(1) IN	play always I_i^t
(2) OUT	play always O_i^t
(3) ALT	alternate between I_i^t and O_i^t
(4) TFT	tit-for-tat: replicate the choice of the partner in the previous round
(5) TRIG	grim-trigger: play the action consistent with EFO if all past outcomes are consistent with EFO and play O_i^t forever otherwise
(6) REV	reverse tit-for-tat: reverse the choice of the partner in the previous round
(7) FORG	forgiving trigger: play O_i^t unless the last round outcome was (I_i^{t-1}, I_j^{t-1})
(8) PAVLOV	play I_i^t if players coordinated in the last round and O_i^t otherwise
(9) STICK	play I_i^t unless the last round outcome was (O_i^{t-1}, O_j^{t-1})

Table 5: Some simple strategies in **SH**

The strategies are similar to those in **BoS**, but their behavioral interpretation differs in some cases. For example, ALT is less natural than in **BoS**. On the other hand, TRIG is closer in spirit to grim trigger in the prisoner's dilemma, since the punishment outcome is a subgame Perfect equilibrium of the continuation game in **SH** (but not in **BoS**). The remaining strategies capture different ways to instill coordination on (I_1, I_2) while sanctioning more or less harshly deviations to O.

5.2.2 Empirical behavior

We follow the same methodology to study individual strategies in **SH**. Again, with only four strategies, in this case IN, OUT, TFT and ALT, we can classify the behavior of 75.2% and 86.2% of participants in SH1 and SH2 respectively. Including all five remaining strategies would classify 12 more participants in SH1 and 4 more in SH2.[16]

[16] If we instead allowed a maximum of 2 and 4 deviations, we would classify 67.9% and 81.0% of participants in SH1 and 81.0% and 89.3% in SH2.

Figure 9 provides a Venn diagram describing the overlap between strategies.

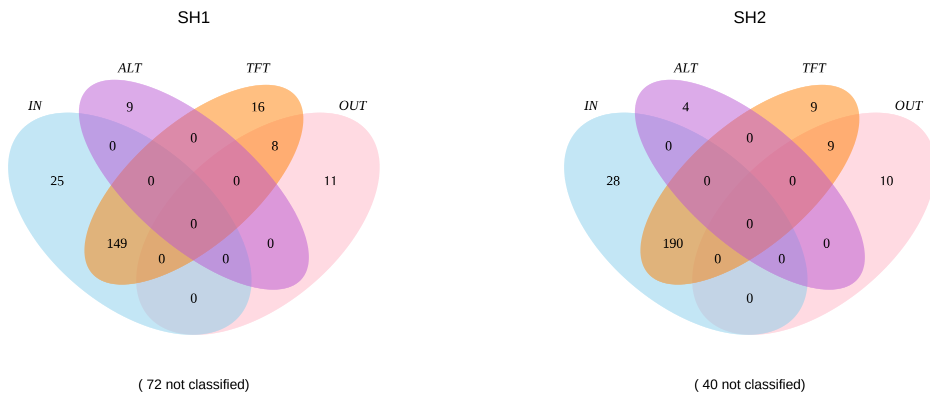

Figure 9: Overlap between strategies in SH1 (left) and SH2 (right)

Players in groups that coordinate on EFO are classified as IN/TFT. whereas players in groups where both partners always play the Pareto inferior outcome are classified as OUT/TFT. These are the only overlapping strategies, with an overwhelming majority in the former and a small but positive number in the latter. TFT are players who often coordinate on an equilibrium, but the equilibrium changes over time. Strategy ALT captures a curious behavior since there is a priori no intuitive reason for such alternation. Players classified as IN are individuals who insist on the potentially superior outcome but face some (transitory) resistance from their partner. Finally, OUT are players who select the safe strategy *despite* being incited by their partner to coordinate on the superior Nash equilibrium. Figure 10 reports the distribution of strategies in each grade-group.

As in **BoS**, there is a significant and sustained increase in EFO with age (IN/TFT) and a general improvement between SH1 and SH2, especially in the younger population. There is also a decrease with age in the proportion of unclassified players (OTHER). Coordination on O (OUT/TFT) is infrequent and spread throughout all grade-groups, whereas choosing I irrespective of the behavior of the partner (IN) is more common in the younger population. More generally, we again notice improvements, with a given grade-group in SH2 behaving like an older group in SH1 (2-3 and 4-5 in SH2 is similar to 8-10 in SH1 while 6-7 and 8-10 in SH2 is similar to U in SH1, chi-square tests, p-values > 0.05). Finally, it is worth stressing out that even our youngest participants play this game remarkably well: almost half the pairs of 2nd and 3rd graders (7 to 9 years old) manage to coordinate perfectly in the Pareto

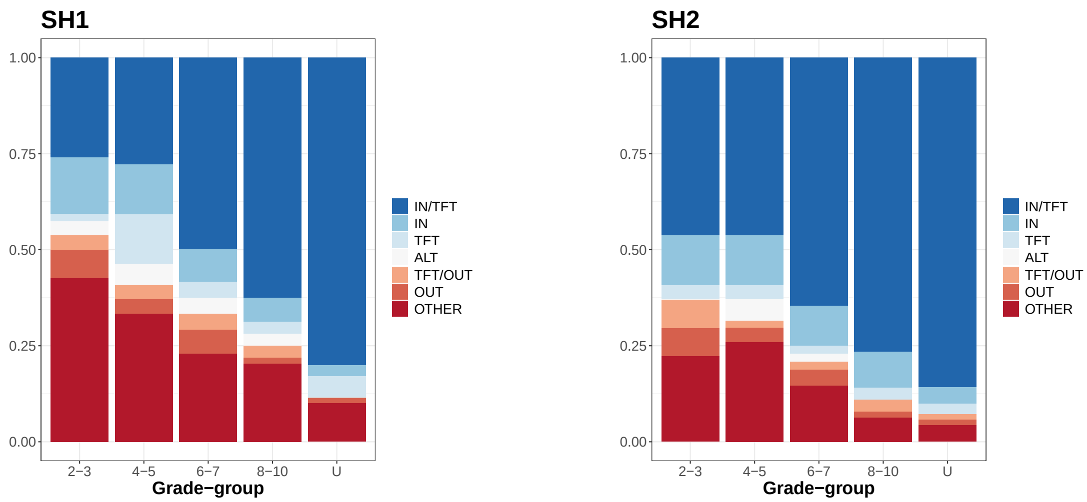

Figure 10: Distribution of strategies across grade-groups in SH1 and SH2

superior equilibrium by the second time they play this game.

Last, Table 6 reports the payoffs of the different strategies in each supergame, averaged over rounds 5 to 24 and with the overall payoff in the last column.

	IN/TFT	IN	TFT	ALT	TFT/OUT	OUT	OTHER	**all**
SH1	2.99	2.54	2.48	2.38	1.92	2.00	2.14	2.63
	(.00)	(.09)	(.08)	(.04)	(.02)	(.01)	(.03)	(0.03)
SH2	2.99	2.75	2.63	2.39	1.97	1.95	2.08	2.76
	(.00)	(.04)	(.06)	(.07)	(.01)	(.01)	(.040)	(0.02)

(standard errors in parenthesis)

Table 6: Average payoffs as a function of the strategy in **SH**

By construction, payoffs of IN/TFT players are close to 3 and payoffs of TFT/OUT and OUT players are close to 2. Participants classified as IN have a relatively high payoff because their partner chooses I in 78% and 88% of the rounds, on average, in SH1 and SH2. Deviations by partners are thus sporadic, which makes it easy to coordinate. By contrast, ALT players incur significant losses because they end up coordinating on (I, I) only between 7 and 10 times. Earnings of unclassified players (OTHER) are not much higher than the earnings of those who plays O. As in **BoS**, the difference in average payoffs between the two supergames reported in the last column (t-test, $p < 0.001$) is mainly driven by the increased proportion of players with strategies compatible with EFO at the expense of unclassified players.

5.3 Summary

Behavior in **BoS** steeply improves with age, starting with either indiscernible or egocentric strategies in the youngest grade-group, and ending with strategies that support EFO. There is also significant improvement after only one supergame, which suggests fast learning and rapid adaptation to the lessons learned. Patterns are similar in **SH**. In both cases, most participants adhere to one of a small number of strategies. All participants, but especially the younger ones, find it much easier in **SH** than in **BoS** both to avoid an indiscernible strategy and to coordinate their behavior in the EFO. Indeed, the proportion of strategies leading to EFO is 36.6% in BoS1 against 60% in SH1 and 48.3% in BoS2 against 75% in SH2. Differences in both cases are highly significant (tests of comparison of proportions, $p < 0.001$).

The age-related increased ability of participants to coordinate on EFO in **BoS** provides strong support for **H1**. The overall increased performance in both **BoS** and **SH** is evidence in favor of **H3** and the significant differences between games is consistent with **H4**. By contrast, **H2** is not supported by the data. Participants of all ages restrict their attention to relatively simple strategies and avoid complex options (only TEST is selected among the most sophisticated alternatives, that is, options 7-8-9 of each game). A posteriori, this is not surprising. We know since the pioneering work of Axelrod (1985) that excessively sophisticated strategies are neither empirically optimal nor widely common in the population. Our school-age participants, especially the older ones, manage to coordinate on the EFO with simple (though not simplistic) strategies. The data also provides support for rapid learning, a result we had not hypothesized given the short window. The idea that behavior of a given grade-group in the second supergame is similar to that of their older peers in the first supergame suggests that children have an intrinsic ability to "skip developmental stages" when they are exposed, even if briefly, to certain problems.

6 Regression analysis

Last, we perform some regressions to further investigate the effects highlighted above. We report a brief summary of the findings and refer the reader to Appendix

B for the full analysis.

First, we conduct OLS regressions at the choice level to investigate the effect of age on actions, outcomes and earnings in **BoS** and **SH**. Corroborating previous findings, we note that age is a powerful predictor of equilibrium play in all four supergames and that behavior improves in the second supergame. Interestingly, we also find a certain portability across games: participants who start with **SH** perform better when they move to **BoS** than those who start with this more complex game. Second, we conduct Probit regressions at the individual level to study the factors that contribute to the selection of the different strategies. Once again, age is a predictor of EFO. It is also negatively related to the choice of inferior strategies in **BoS** (ME) although not in **SH** (OUT). As in the OLS regressions, playing **SH** first helps participants to play closer to EFO and away from inferior strategies in **BoS**, while the reverse does not have a significant effect. Finally, we also find a strong correlation between individual strategy choices across supergames and games, suggesting that "good" and "bad" decisions are traits that extend within and across games. The correlations hold after controlling for age, which implies that variation in strategic thinking is due to age but also to general cognitive abilities.

7 Conclusion

This paper has studied choice by children and adolescents in repeated **BoS** and **SH**. It has three distinctive features. It is the first to investigate the developmental trajectory of behavior in long repeated coordination games. It proposes a novel methodology and a story line that can be exported to other populations that might find abstract representations challenging. It reports strategies of potential empirical relevance analogous but not identical to those studied in other repeated games.

Both games feature an increase in coordination on EFO with age, marked by a higher choice of reactive strategies capable of supporting EFO at the expense of strategies that ignore the behavior of partners. Performance is also higher in **SH** than in **BoS**, and in supergame 2 than in 1. Natural extensions include other populations and other coordination games.

Research in developmental psychology has reported evidence of collaborative behavior in children. However, studies focus on the communication strategies that facilitate coordination, and they restrict the analysis to narrow age ranges (Grueneisen and Tomasello, 2017, 2019). Instead, we think that studying the evolution in strategies and outcomes over time allows us to determine whether coordination is innate or acquired, intuitive or learned, achievable or not. Our study indicates that the ability to coordinate through repeated interactions develops progressively. Contrary to the "natural heuristic hypothesis" of Rand et al. (2012), we find that the impressive capacity of adults to coordinate in these games is not natural and instinctive.[17] Rather, development acts on traits that are *gradually* expressed. At the same time, all individuals select their strategy from the same small set.

We can see from our data how development operates. As children grow, they learn to form beliefs about the intentions and goals of people they interact with. This opens the door to the development of strategic behavior. In our games, children can use inductive logic, and build a theory of the best course of action from the observation of desirable and undesirable outcomes. This reasoning is simple in **SH**. Indeed, once a child selects the action with highest potential (I) and observes the most desirable outcome (I, I), it is enough to repeat the same action. However, it is more complex in **BoS**. After choosing M and observing one's favorite outcome (M, Y), a child may be tempted to replicate the same action M, which will often lead to miscoordination. To infer the optimality of alternation, it is necessary to keep track of past moves, observe long sequences, and understand the need to sacrifice current payoff to induce collaboration. These conceptual differences are likely the main reason why children coordinate better and faster in **SH** than in **BoS**.

Our study also points to important phenomena that should be studied further. In particular, it is interesting that behavior of a given grade-group in the second supergame is similar to that of their older peers in their first attempt, not only in terms of outcomes and payoffs, but even in the distribution of strategies. Natural questions are whether young children can learn to play like adults and how many supergames it would take. And, importantly, what kind of learning explains these improvements: is it mechanical imitation or the application of some elements of logic they already possess? Would they be able to export what they learn to other games and to other contexts? Finally, the fact that participants perform better in **BoS** if they have already played two rounds of **SH**, but not the other way around, also indicates that some form of skill transfer operates.

[17]The subsequent literature (e.g., Krajbich et al. (2015); Bouwmeester et al. (2017); Alós-Ferrer and Garagnani (2020)) argues that evidence for coordination heavily depends on experimental conditions and data interpretation.

Finally, differences in behavior also point to differences in cognitive capabilities. Some of these differences are age-related but others are due to individual heterogeneity. A recent strand of the experimental literature has shown a positive association between IQ and performance in games of strategy (Brañas-Garza et al., 2012; Gill and Prowse, 2016; Proto et al., 2019, Forthcoming; Fe et al., 2022). In this study, we have highlighted considerable differences within grade-groups. This is consistent with Brocas and Carrillo (2021), where a significant proportion of very young children choose optimally in reasonably sophisticated games while a non-negligible fraction of late teens and adults do not. Such heterogeneity in behavior points to differences in cognition at a given age. It also suggests that differences in genetic makeup and/or unique environmental influences shape cognitive development. These findings inform us about the (highly inter-related) contribution of nature and nurture to human cognitive abilities and behavior. They are important to design interventions, such as school improvements, that optimize human capital accumulation (Cunha et al., 2006; Cunha and Heckman, 2007), and underscore the need to disentangle these different influences to assess the relationship between schooling and life outcomes.

References

Carlos Alós-Ferrer and Michele Garagnani. The cognitive foundations of cooperation. *Journal of Economic Behavior & Organization*, 175:71–85, 2020.

Masaki Aoyagi and Guillaume Fréchette. Collusion as public monitoring becomes noisy: Experimental evidence. *J. of Economic theory*, 144(3):1135–1165, 2009.

Robert M Axelrod. *The evolution of cooperation*. Basic Books, New York, 1985.

Zvonimir Bašić, Parampreet Christopher Bindra, Daniela Glätzle-Rützle, Angelo Romano, Matthias Sutter, and Claudia Zoller. The roots of cooperation. CESifo Working Paper, 2021.

Jean-Pierre Benoit and Vijay Krishna. Finitely repeated games. *Econometrica*, 53 (4):905–22, 1985.

Peter R Blake, David G Rand, Dustin Tingley, and Felix Warneken. The shadow of the future promotes cooperation in a repeated prisoner's dilemma for children. *Scientific reports*, 5(1):14559, 2015.

Samantha Bouwmeester, Peter PJL Verkoeijen, Balazs Aczel, Fernando Barbosa, Laurent Bègue, Pablo Brañas-Garza, et al. Registered replication report: Rand, greene, and nowak (2012). *Perspectives on Psychological Science*, 12(3):527–542, 2017.

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

Pablo Brañas Garza, Isabelle Brocas, and Juan D Carrillo. Developmental meritocracy. mimeo, USC., 2023.

Isabelle Brocas and Juan D Carrillo. The development of social strategic ignorance and other regarding behavior from childhood to adulthood. *Journal of Behavioral and Experimental Economics*, 85:101524, 2020a.

Isabelle Brocas and Juan D Carrillo. Studying decision-making in children: Challenges and opportunities. *Journal of Economic Behavior & Organization*, 179: 777–783, 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 and Juan D Carrillo. Adverse selection and contingent reasoning in preadolescents and teenagers. *Games and Economic Behavior*, 133:331–351, 2022a.

Isabelle Brocas and Juan D Carrillo. The development of randomization and deceptive behavior in mixed strategy games. *Quantitative Economics*, 13(2):825–862, 2022b.

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

Jeannette Brosig-Koch, Timo Heinrich, and Christoph Helbach. Exploring the capability to reason backwards: An experimental study with children, adolescents, and young adults. *European Economic Review*, 74:286–302, 2015.

Gabriele Camera and Marco Casari. Cooperation among strangers under the shadow of the future. *American Economic Review*, 99(3):979–1005, 2009.

Gabriele Camera, Marco Casari, and Maria Bigoni. Cooperative strategies in anonymous economies: An experiment. *Games and Economic Behavior*, 75(2):570–586, 2012.

Gabriele Camera, Marco Casari, and Maria Bigoni. Money and trust among strangers. *PNAS*, 110(37):14889–14893, 2013.

Colin F Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton university press, 2011.

Ramon Cobo-Reyes, Jose J Dominguez, Fernando García-Quero, Brit Grosskopf, Juan A Lacomba, Francisco Lagos, Tracy Xiao Liu, and Graeme Pearce. The development of social preferences. *Journal of Economic Behavior & Organization*, 179:653–666, 2020.

David J Cooper and Roberto A Weber. Recent advances in experimental coordination games. *Handbook of experimental game theory*, pages 149–183, 2020.

W. Crain. *Theories of Development: Concepts and Applications*. Psychology Press, 6 edition, 2015. ISBN 9781317343219.

Flavio Cunha and James Heckman. The technology of skill formation. *American Economic Review*, 97(2):31–47, 2007.

Flavio Cunha, James J Heckman, Lance Lochner, and Dimitriy V Masterov. Interpreting the evidence on life cycle skill formation. *Handbook of the Economics of Education*, 1:697–812, 2006.

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.

Pedro Dal Bó and Guillaume R Fréchette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1): 411–29, 2011.

Pedro Dal Bó and Guillaume R Fréchette. On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1): 60–114, 2018.

Pedro Dal Bó, Guillaume R Fréchette, and Jeongbin Kim. The determinants of efficient behavior in coordination games. *Games and Economic Behavior*, 130: 352–368, 2021.

Eduardo Fe, David Gill, and Victoria Prowse. Cognitive skills, strategic sophistication, and life outcomes. *Journal of Political Economy*, 130(10):2643–2704, 2022.

Daniel Friedman and Ryan Oprea. A continuous dilemma. *American Economic Review*, 102(1):337–63, 2012.

Drew Fudenberg, David G Rand, and Anna Dreber. Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, 102(2): 720–49, 2012.

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

Sebastian Grueneisen and Michael Tomasello. Children coordinate in a recurrent social dilemma by taking turns and along dominance asymmetries. *Developmental psychology*, 53(2):265, 2017.

Sebastian Grueneisen and Michael Tomasello. Children use rules to coordinate in a social dilemma. *Journal of Experimental Child Psychology*, 179:362–374, 2019.

Sebastian Grueneisen, Emily Wyman, and Michael Tomasello. “i know you don’t know i know…” children use second-order false-belief reasoning for peer coordination. *Child development*, 86(1):287–293, 2015.

William Harbaugh, Kate Krause, Steven Liday, and Lise Vesterlund. Trust in children. *Trust and reciprocity: interdisciplinary lessons from experimental research*, pages 302–322, 2003.

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

Ian Krajbich, Björn Bartling, Todd Hare, and Ernst Fehr. Rethinking fast and slow based on a critique of reaction-time reverse inference. *Nature communications*, 6 (1):1–9, 2015.

Christoph Kuzmics, Thomas Palfrey, and Brian W Rogers. Symmetric play in repeated allocation games. *Journal of Economic Theory*, 154:25–67, 2014.

John A List, Ragan Petrie, and Anya Samek. How experiments with children inform economics. *Journal of Economic Literature*, 61(2):504–564, 2023.

Laurent Mathevet and Julian Romero. Predictive repeated game theory: Measures and experiments, 2014.

Richard D McKelvey and Thomas R Palfrey. Playing in the dark: Information, learning, and coordination in repeated games. *Caltech*, 2001.

Patricia H Miller. *Theories of developmental psychology*. Macmillan, 2002.

H Elizabeth Peters, A Sinan Ünür, Jeremy Clark, and William D Schulze. Free-riding and the provision of public goods in the family: a laboratory experiment. *International Economic Review*, 45(1):283–299, 2004.

Jean Piaget, David Elkind, and Anita Tenzer. *Six psychological studies*. New York, NY: Vintage Books., 1967. ISBN 0394704622.

Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, personality, and gains from cooperation in repeated interactions. *Journal of Political Economy*, 127(3):1351–1390, 2019.

Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, errors and cooperation in repeated interactions. *Review of Economic Studies*, Forthcoming.

Eva Rafetseder, Maria Schwitalla, and Josef Perner. Counterfactual reasoning: From childhood to adulthood. *J. of experimental child psychology*, 114(3):389–404, 2013.

David G Rand, Joshua D Greene, and Martin A Nowak. Spontaneous giving and calculated greed. *Nature*, 489(7416):427–430, 2012.

Julian Romero and Yaroslav Rosokha. Constructing strategies in the indefinitely repeated prisoner's dilemma game. *European Economic Review*, 104:185–219, 2018.

Julian Romero and Yaroslav Rosokha. The evolution of cooperation: The role of costly strategy adjustments. *American Economic Journal: Microeconomics*, 11 (1):299–328, 2019.

Edward B Royzman, Kimberly Wright Cassidy, and Jonathan Baron. "i know, you know": Epistemic egocentrism in children and adults. *Review of General Psychology*, 7(1):38–65, 2003.

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

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

Henry Wellman, David Cross, and Julanne Watson. Meta-analysis of theory-of-mind development: the truth about false belief. *Child development*, 72:655–684, 2001.

Appendix A. Undergraduates vs. teachers

Opportunities to test adults from the general population are useful to put findings obtained with undergraduates into perspective. Here, we report the behavior of a sample of teachers at the school (T). Although procedures are identical, the comparison must be taken with a grain of salt as the sample of teachers is small (30 participants). There is also significant heterogeneity in terms of their age and academic achievement (from BA to PhD) and, as in college students control groups, there are differences in the academic focus of teachers (arts, sciences and humanities). Overall, it is an interesting alternative reference to USC undergraduates. While the latter can be seen as an (imperfect) proxy of what school age students will become right after graduating, teachers are working professionals, significantly older and more experienced. On the other hand, teachers spend a large fraction of their time in the same environment as the school grade-group.

Table 7 presents a comparison of the behavior of teachers (T) and USC undergraduates (U) in the first [1] and second [2] supergame of **BoS** and **SH**. We include the descriptive statistics of the actions, outcomes and payoff variables from section 4 as well as the two polar classes of individual strategies developed in Appendix C (*EFO* and *Inferior*).

			Teachers		USC	
			T [1]	T [2]	U [1]	U [2]
BoS	descriptive	$\Pr(M_i)$	0.63	0.53	0.51	0.51
		$\Pr(M_i, Y_j)$	0.59	0.77	0.77	0.77
		Payoff	2.98	3.57	3.67	3.63
	strategies	*EFO*^{o}	0.25	0.67	0.71	0.77
		Inferior†	0.07	0.06	0.07	0.07
SH	descriptive	$\Pr(I_i)$	0.99	1.00	0.91	0.95
		$\Pr(I_i, I_j)$	0.99	0.99	0.87	0.93
		Payoff	2.98	2.99	2.83	2.91
	strategies	*EFO*‡	1.00	1.00	0.83	0.90
		Inferior§	0.00	0.00	0.01	0.02

^{o}ALT/TFT; †ME, ME/TEST and ME/TFT; ‡IN and IN/TFT; §OUT and TFT/OUT

Table 7: Summary statistics of control populations

Teachers perform worse than USC undergraduates in the first supergame of **BoS**. While they do not follow a myopic ME strategy, they still choose their preferred action too often, fail to coordinate and therefore obtain lower payoffs (all p-values < 0.05). The effect is particularly noticeable in the proportion of EFO (ALT/TFT strategies), suggesting an initial difficulty to understand (or believe that the partner will understand) the joint benefits of alternating. These differences disappear by the second supergame, when teachers and undergraduates become statistically similar.

Results are different in **SH**. Even though undergraduates play this game very well, they are still outperformed by teachers who achieve perfect coordination, thus maximizing earnings. Despite the improvement of undergraduates in the second supergame, differences remain statistically significant in actions, payoffs and strategies (all p-values < 0.05).

We do not have a clear explanation why the battle of the sexes is relatively harder and the stag hunt is relatively easier for teachers than for undergraduates. One could find ex-post reasons for those differences. For example, we could invoke a higher homogeneity in the undergraduate population, therefore a higher capacity to anticipate and mimic the choice of others. And yet, both games are symmetric, which raises the question of why homogeneity would be more conducive to equilibrium in one case than in the other. Additional treatments would be useful to ascertain the robustness of this finding and understand its roots.

Appendix B. Regression analysis

We next report some regression analyses. Our adult control group is a benchmark of comparison (not the culmination of the developmental trend of this school population) so, to avoid polluting the trajectory, we include only the LILA students. We would ideally like to include their age in months. Unfortunately, this information is not available. We therefore include instead the numerical grade as a proxy for age, with the understanding that some age differences may exist between participants in the same grade.

B1. Actions, outcomes and payoffs

We conduct OLS regressions to investigate the effect of age–captured by the numerical variable *Grade* that takes values 2 to 10–on actions, outcomes and earnings. Individual choices are captured by the percentage of times players choose their favorite action ($\Pr(M_i)$) in **BoS** and the risky action ($\Pr(I_i)$) in **SH**, respectively. Outcomes are modeled as the percentage of times a group coordinates on either static Nash equilibria ($\Pr(M_i, Y_j)$) in **BoS** and on the Pareto superior static Nash equilibrium ($\Pr(I_i, I_j)$) in **SH**. Earnings are computed as individual per-round payoffs. We control for order effects by including the dummy variable *1stSH* (= 1 if the participant plays first **SH** and then **BoS**). For individual measures (actions and payoffs), we also include dummy variables to study the effect of gender (*Male* = 1) and whether the participant has one or more siblings (*Siblings* = 1). All the regressions are performed separately for each supergame ([1] and [2]). For individual measures, we also run a regressions with all the observations ([all]) and a supergame dummy (*BoS2* = 1 or *SH2* = 1) to determine potential changes across supergames. Finally, in [all] we also include an interaction term *Grade*1stSH* to study whether order effects are different for participants in different grades. Standard errors are clustered at the matched pair level. The results are reported in Table 8 for **BoS** and in Table 9 for **SH**.

	$\Pr(M_i)$				$\Pr(M_i, Y_j)$		Payoff			
	[1]	[2]	[all]	[all]	[1]	[2]	[1]	[2]	[all]	[all]
Grade	-0.013**	-0.017***	-0.015***	-0.018***	0.035***	0.047***	0.102***	0.137***	0.119***	0.137***
	(.004)	(.004)	(.003)	(.004)	(.009)	(.007)	(.026)	(.021)	(.017)	(.027)
1stSH	-0.071***	-0.051*	-0.061***	-0.097*	0.133**	0.092*	0.396**	0.276*	0.336***	0.523*
	(.021)	(.020)	(.016)	(.038)	(.046)	(.042)	(.137)	(.127)	(.093)	(.223)
*Grade*1stSH*	—	—	—	0.006	—	—	—	—	—	-0.033
				(.005)						(.035)
Male	0.052*	0.026	0.039*	0.038*	—	—	0.035	-0.037	-0.001	0.002
	(.024)	(.028)	(.018)	(.019)			(.100)	(.102)	(.071)	(.071)
Siblings	0.023	0.016	0.020	0.019	—	—	0.050	0.019	0.034	0.040
	(.027)	(.035)	(.022)	(.022)			(.113)	(.112)	(.079)	(.080)
BoS2	—	—	-0.030*	-0.030*	—	—	—	—	0.256**	0.256**
			(.014)	(.014)					(.093)	(.093)
const.	0.683***	0.683***	0.698***	0.719***	0.396***	0.435***	2.150***	2.327***	2.110***	2.003***
	(.031)	(.041)	(.025)	(.036)	(.055)	(.060)	(.176)	(.190)	(.127)	(.178)
Adj. R^2	0.070	0.048	0.068	0.067	0.168	0.246	0.145	0.197	0.191	0.192
# obs.	220	220	440	440	110	110	220	220	440	440

*p<0.05; **p<0.01; ***p<0.001

Table 8: OLS Regressions of choices, outcomes and payoffs in **BoS**

As expected, age is a very powerful predictor of equilibrium play in all four supergames. Equilibrium outcomes are 3.4 to 4.7 percentage points higher as we move from one grade to the next. Accordingly, per-round payoffs increase, on average, 0.12 points in **BoS** and 0.04 points in **SH** from one grade to the next. This is considerable given that participants obtain 4 and 3 points under EFO in **BoS** and **SH** respectively, and that 3 and 2 points are easy to secure (by always choosing Y_i and O_i). All these effects are significant at least at the 1%-level. Marginal effects are similar in both supergames. At the same time, the joint regressions reveal highly significant payoff increases in the second supergame. It reinforces the idea previously documented that our participants learn, adapt and therefore improve earnings between supergames. It is also interesting to observe that participants who start with **SH** perform significantly better when they move to **BoS** than those who start with this more complex game. This indicates a certain portability across games as well as the desirability for learning purposes of moving from an easier to a more difficult game: by first getting used to think strategically in a simple context (let's both choose the action that maximizes our payoffs) participants are then able to also successfully coordinate in

	Pr(I_i)				Pr(I_i, I_j)		Payoff			
	[1]	[2]	[all]	[all]	[1]	[2]	[1]	[2]	[all]	[all]
Grade	0.024*	0.027**	0.026***	0.041***	0.034**	0.037**	0.042**	0.042***	0.042***	0.061***
	(.010)	(.010)	(.007)	(.010)	(.012)	(.012)	(.015)	(.015)	(.010)	(.014)
1stSH	-0.012	-0.001	-0.006	0.148	-0.044	-0.025	-0.086	-0.052	-0.069	0.123
	(.054)	(.053)	(.038)	(.096)	(.067)	(.065)	(.079)	(.078)	(.055)	(.138)
*Grade*1stSH*	—	—	—	-0.027	—	—	—	—	—	-0.034
				(.014)						(.021)
Male	-0.011	-0.042	-0.026	-0.024	—	—	-0.020	-0.043	-0.031	-0.028
	(.038)	(.044)	(.029)	(.029)			(.053)	(.063)	(.041)	(.041)
Siblings	0.044	0.047	0.046	0.050	—	—	0.166*	0.056	0.111*	0.116*
	(.047)	(.053)	(.035)	(.036)			(.072)	(.072)	(.051)	(.051)
SH2	—	—	0.088*	0.088*	—	—	—	—	0.164**	0.164**
			(.037)	(.037)					(.055)	(.055)
const.	0.572***	0.651***	0.567***	0.478***	0.458***	0.561***	2.218***	2.457***	2.256***	2.145***
	(.081)	(.084)	(.059)	(.082)	(.090)	(.096)	(.119)	(.121)	(.087)	(.113)
Adj. R^2	0.032	0.051	0.068	0.080	0.048	0.056	0.082	0.058	0.106	0.114
# obs.	220	220	440	440	110	110	220	220	440	440

*p<0.05; **p<0.01; ***p<0.001

Table 9: OLS Regressions of choices, outcomes and payoffs **SH**

the complex one (let's now take turns choosing our favorite action). Portability, however, is unidirectional and no order effect is found when participants start with **BoS**. This may be partly due to the fact that behavior in **SH** is already close to equilibrium for many participants, so there is small(er) room for improvement. The interaction between grade and order is not statistically significant, which suggests that the order effect is independent of the individual's age. Finally, we found no systematic effect of gender and only a small effect of siblings on behavior and payoffs.

B2. Strategies

Section 5 reported heterogeneity in the choice of strategies within and across grade-groups. Here, we investigate the contribution of several factors on the selection of strategies. For each supergame, we focus on two types of strategies: those leading to *EFO* (ALT/TFT in **BoS** and IN or IN/TFT in **SH**) and those leading to *Inferior* outcomes (ME, ME/TEST or ME/TFT in **BoS** and OUT or TFT/OUT in **SH**). We construct a strategy outcome variable and we classify participants into those who choose such strategy (= 1) and those who do not (= 0). We then conduct a Probit regression of the strategy outcome variable on *Grade*, *1stSH*, *Male* and *Siblings* as well as a variable that captures the group's initial behavior. To wit, we use dummy variables *1st*(M_1, M_2) and *1st*(I_1, I_2) that take value 1 if the first round's outcome is (M_1, M_2) in **BoS** and (I_1, I_2) in **SH**. The idea is that initial choices may serve as anchor or signaling and therefore be conducive of non-cooperative and cooperative behavior, respectively.[18] The results are presented in Table 10.

	BoS				SH			
	EFO^o		*Inferior*†		*EFO*‡		*Inferior*§	
	[1]	[2]	[1]	[2]	[1]	[2]	[1]	[2]
Grade	0.244***	0.297***	-0.097*	-0.165***	0.109*	0.122	-0.063	-0.094
	(.050)	(.054)	(.044)	(.048)	(.049)	(.064)	(.059)	(.067)
1stSH	0.636*	0.760**	-0.689**	-0.533*	-0.275	-0.131	-0.540	-0.207
	(.302)	(.279)	(.211)	(.216)	(.242)	(.272)	(.366)	(.328)
Male	0.032	-0.102	0.603**	0.119	0.184	-0.263	0.545*	0.475
	(.228)	(.220)	(.201)	(.221)	(.1982)	(.233)	(.238)	(.338)
Siblings	-0.268	0.066	-0.014	-0.253	-0.190	0.115	-0.251	-0.286
	(.259)	(.216)	(.254)	(.238)	(.193)	(.214)	(.283)	(.304)
1st(M_1, M_2)	-1.038**	-0.222	0.506*	0.242	—	—	—	—
	(.396)	(.216)	(.313)	(.231)				
1st(I_1, I_2)	—	—	—	—	1.407***	1.740***	—	—
					(.248)	(.308)		
const.	-2.139***	-2.386***	-0.580	0.168	-0.972*	-1.242**	-0.955*	-0.895*
	(.433)	(.389)	(.352)	(.366)	(.388)	(.405)	(.473)	(.414)
AIC	197.3	228.2	194.3	189.6	242.6	198.2	123.8	121.6
# obs.	220	220	220	220	220	220	220	220

*p<0.05; **p<0.01; ***p<0.001

Strategies included are: ^{o}ALT/TFT; †ME, ME/TEST and ME/TFT; ‡IN and IN/TFT; §OUT and TFT/OUT

Including *1st*(I_1, I_2) in the *Inferior* regressions leads to robustness problems due to lack of observations.

Table 10: Probit Regressions of individual strategies

Age is a strong predictor of EFO in **BoS** and to a lesser extent in **SH**. It is also negatively related to the choice of inferior strategies in **BoS**. By contrast, age does not explain inferior strategies in **SH**, mainly because OUT-compatible strategies are relatively rare in the population (see Figure 10). As in Tables 8 and 9, playing **SH** first leads participants to play closer to EFO and away from inferior strategies in **BoS**, while the reverse has no effect. There is a small indication that males play more often inferior strategies in the first supergame and no significant effect of siblings. Finally, groups where both individuals choose M_i in the first round are more likely to miscoordinate in **BoS** (after controlling for all other variables), but only in the first supergame. Conversely, both individuals starting in I_i is a very strong predictor of EFO in **SH**.

We conducted ordered probit regressions of changes in strategies between supergames (omitted for brevity but available upon request). They confirm the significant increase in performance, with participants moving from inferior to EFO strategies. However, these

[18]Remember that individual strategies are determined based on choices in rounds 5 to 24, so there is no endogeneity problem with the variables *1st*(M_1, M_2) and *1st*(I_1, I_2).

effects are not modulated by age, except in BOS for the 8-10 grade-group ($p = 0.025$). This suggests a general learning trend from early elementary school until late middle school.

Finally, there is a strong correlation between individual strategy choices across supergames and games. The PCC of *EFO* strategies are: 0.54 between BoS1 and BoS2, 0.42 between SH1 and SH2, and 0.27 between **BoS** and **SH** ($p < 0.001$ in all cases). Similarly, the PCC of *Inferior* strategies are: 0.51 between BoS1 and BoS2, 0.60 between SH1 and SH2 ($p < 0.001$ in both cases) and 0.16 between **BoS** and **SH** ($p = 0.015$). It suggests that both "good" and "bad" decisions are traits that extend within and across games. Importantly, the results continue to hold after controlling for age. Therefore, the observed variations in strategic thinking are due to age but also to general cognitive abilities.

Appendix C. Sample of instructions

This game is called "Find the balance". The computer will decide with whom you play this game. One of you will be "RED" and the other will be "GREEN". The computer also decides who is RED and who is GREEN. If you are RED, your screen looks like this (see Figure 11 for the slides).

[SLIDE 1]

At the top of the screen, it says you are RED. You own the RED scale and the black ball in the middle of your screen. You can see a GREEN scale which belongs to your partner. You need to decide whether to put your ball on the dotted circle of your RED scale or on the dotted circle of your partner's GREEN scale.

If you are GREEN, your screen looks like this.

[SLIDE 2]

It says you are GREEN at the top. You own the GREEN scale and the black ball in the middle of your screen. You can see a RED scale which belongs to your partner. You need to decide whether to put your ball on the dotted circle of your GREEN scale or on the dotted circle of your partner's RED scale. These are the two screens together.

[SLIDE 3]

How do you get points?

[SLIDE 4]

If both balls are put on the RED scale, then player RED gets 5 points and player GREEN gets 3 points. If both balls are put on the GREEN scale, then player RED gets 3 points and player GREEN gets 5 points. If the balls are put on different scales, each player gets 1 point. This information will remain in this screen during your choices.

Now this is very important. You will play many rounds with the same partner. In each round, you will make your choices at the same time. This means that you will not know what your partner did when you make your choice. It is only after both of you have made a choice, that you will both know what each of you did and how many points you got. This will appear in the right column of your screen.

[SLIDE 5]

For example, this is what your screen may look like after 4 rounds. In the first round [explain]. In the second round [explain]. In the third round [explain]. In the fourth round [explain]. You have accumulated a total of [explain] points so far. All right, are you ready to play? The computer will now select partners. When you are ready, make your choices.

[At the end of the 1st game] The game has ended. You can see on your screen the points you have accumulated. We will now mark it down on the record sheet. The computer will now select new partners and you will play the same game again.

[At the end of the 2nd game] This game is finished. Let's move to our next game.

This game is called "Risky stars". As before, the computer will decide with whom you play that game. One of you will be "BLUE" and the other will be "YELLOW". The computer also decides who is BLUE and who is YELLOW. If you are BLUE, your screen looks like this.

[SLIDE 6]

At the top of the screen, it says you are BLUE. In the middle of the screen there is a carpet divided in two. You own the BLUE star. You need to decide whether to put your star on the carpet or outside the carpet. If you are YELLOW, your screen looks like this.

[SLIDE 7]

It says you are YELLOW at the top and you see the same carpet as your partner. You own the YELLOW star and you need to decide whether to put your star on the carpet or outside the carpet. These are the two screens together.

[SLIDE 8]

Now, how do you get points?

[SLIDE 9]

If you put your own star outside the carpet, you get 2 points, no matter what your partner does. If you put it on the carpet, what you get depends on what your partner does. If he also puts his star on the carpet, you both earn 3 points. But if he puts it outside the carpet, then you earn 1 point (while he earns 2 points). This information will remain in this screen during your choices.

As in our first game, you will play many times with the same partner. Each time, you will make your choices at the same time. This means that you will not know what your partner did when you make your choice. It is only after both of you have made a choice, that you will both know what each of you did and how many points you got. This will appear in the right column of your screen.

[SLIDE 10]

For instance, this is what your screen may look like after 4 rounds. In the first round [explain]. In the second round [explain]. In the third round [explain]. In the fourth round [explain]. You have accumulated a total of [explain] points so far. All right, are you ready to play? The computer will now select partners. When you are ready, make your choices.

[At the end of the 1st game]

The game has ended. You can see on your screen the points you have accumulated.

We will now mark it down on the record sheet. The computer will now select new partners and you will play the same game again.

[At the end of the 2nd game] The game has ended. Please answer a few questions and we are done.

[When they have answered the questions] We will now call you one by one and tell you how much money you earned. You can tell your friends how much you got or not. It is totally up to you. You will get today an email from Amazon with an amazon e-giftcard for that amount. Thanks for playing with us.

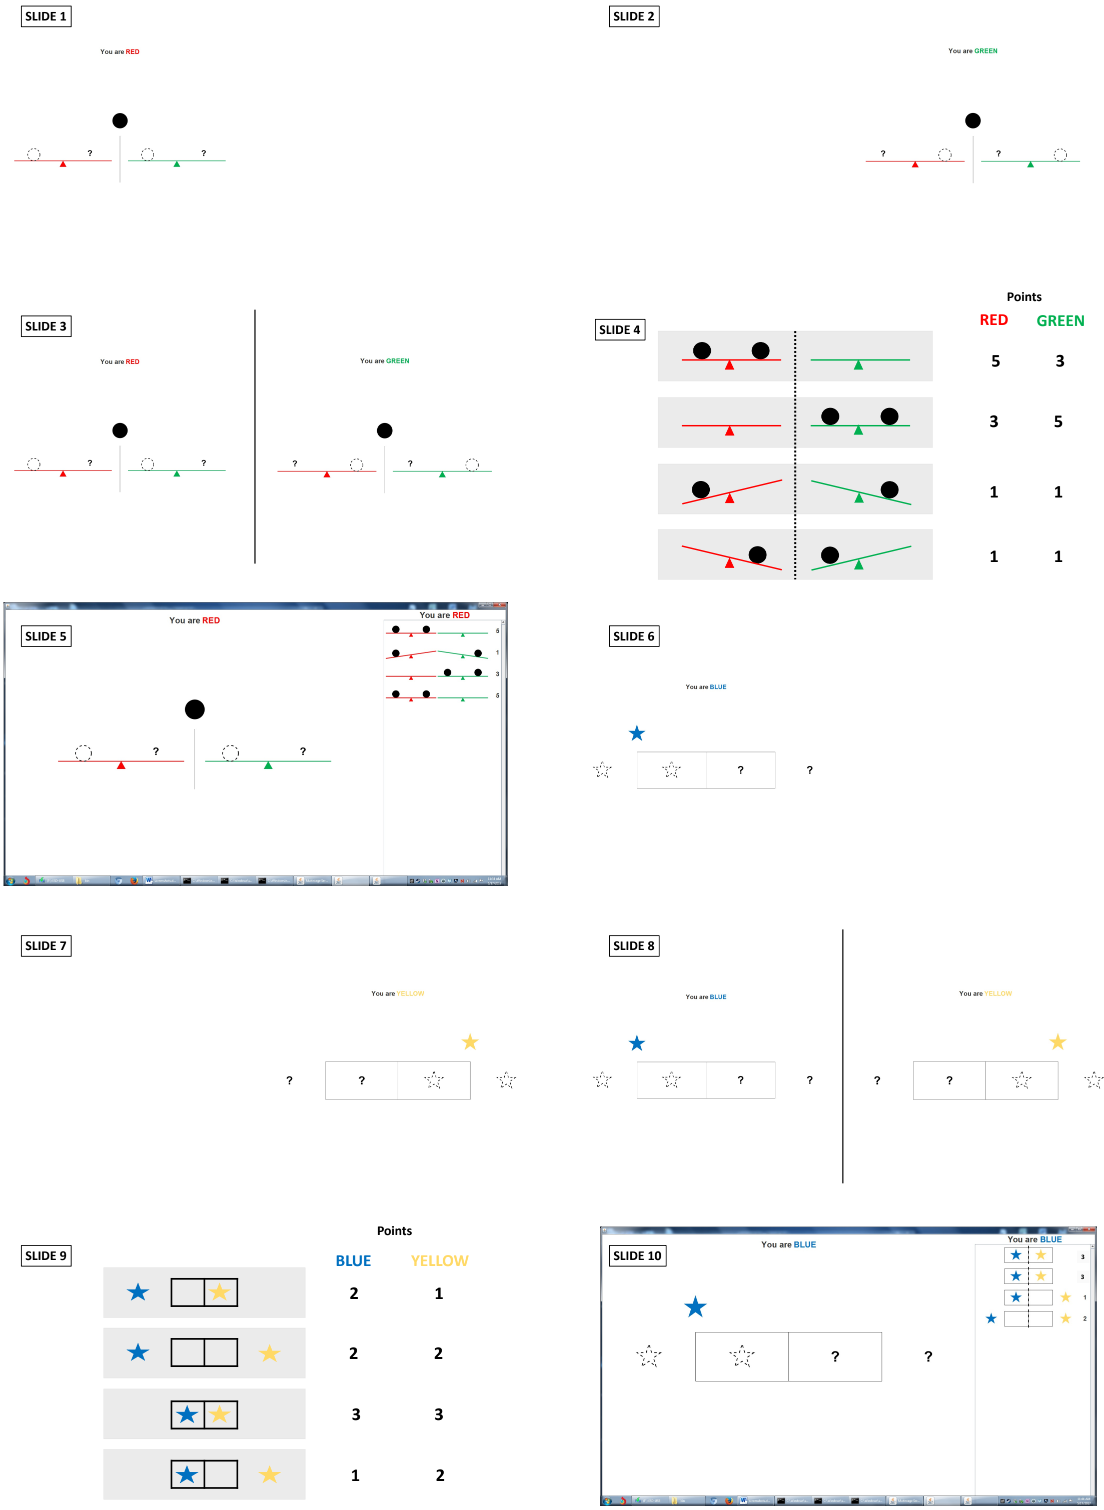

Figure 11: Slides to accompany instructions