

DISCUSSION PAPER SERIES

DP16683

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

Isabelle Brocas and Juan D Carrillo

PUBLIC ECONOMICS

CEPR

Dynamic coordination in efficient and fair strategies: a developmental perspective

Isabelle Brocas and Juan D Carrillo

Discussion Paper DP16683
Published 31 October 2021
Submitted 18 October 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Isabelle Brocas and Juan D Carrillo

Dynamic coordination in efficient and fair strategies: a developmental perspective

Abstract

We study in the laboratory the behavior of children and adolescents (ages 7 to 16) in two repeated coordination games, the stag hunt and battle of the sexes. Coordinating on the efficient and fair long run outcome (EFO) requires participants to share intentions and beliefs. This exercise is arguably complex in the battle of the sexes, as it requires taking turns between the two static Nash equilibria, hence coordinating the strategies. By contrast, in the stag hunt it only requires repeating the action that leads to the Pareto efficient outcome, hence coordinating the actions. We obtain four main findings. First, for both games, we show a significant and remarkably stable increase in the ability to coordinate on the EFO with age. Second, the majority of participants in all ages adhere to one of a small number of relatively simple strategies. Third, jointly profitable outcomes are more prevalent in the stag hunt than in the battle of the sexes. Last, behavior improves between the first and second supergame. This evidence suggests that we gradually learn how to share intentions and beliefs, an ability that we train rapidly and export to new interactions, but that is limited by game complexity.

JEL Classification: C73, C91

Keywords: developmental decision-making, coordination, repeated games

Isabelle Brocas - brocas@usc.edu
University of Southern California and CEPR

Juan D Carrillo - juandc@usc.edu
University of Southern California and CEPR

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

Isabelle Brocas

*University of Southern California
and CEPR*

Juan D. Carrillo

*University of Southern California
and CEPR*

October 2021

Abstract

We study in the laboratory the behavior of children and adolescents (ages 7 to 16) in two repeated coordination games, the stag hunt and battle of the sexes. Coordinating on the efficient and fair long run outcome (EFO) requires participants to share intentions and beliefs. This exercise is arguably complex in the battle of the sexes, as it requires taking turns between the two static Nash equilibria, hence coordinating the strategies. By contrast, in the stag hunt it only requires repeating the action that leads to the Pareto efficient outcome, hence coordinating the actions. We obtain four main findings. First, for both games, we show a significant and remarkably stable increase in the ability to coordinate on the EFO with age. Second, the majority of participants in all ages adhere to one of a small number of relatively simple strategies. Third, jointly profitable outcomes are more prevalent in the stag hunt than in the battle of the sexes. Last, behavior improves between the first and second supergame. This evidence suggests that we gradually learn how to share intentions and beliefs, an ability that we train rapidly and export to new interactions, but that is limited by game complexity.

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. The study was conducted with the University of Southern California IRB approval UP-12-00528. We gratefully acknowledge the financial support of the National Science Foundation grant SES-1851915. Address for correspondence: Juan D. Carrillo, Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089, USA, <juandc@usc.edu>.

1 Introduction

Strategic interactions often feature collaborative outcomes that may be socially desirable but require individual sacrifices. Game theory has successfully modeled 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. The prisoner’s dilemma features a collaborative solution that is not a Nash equilibrium and requires *cooperation*, a collective endeavor, to be achieved. The stag hunt and battle of the sexes exhibit two Nash equilibria and require *coordination*, an agreement to reach a common goal, to be selected. When these games are played repeatedly with the same partner, people can leverage observations of play to assess partners’ intentions and use histories to keep them in check. Unfortunately, the usefulness of theory is very 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, and a myriad of strategies can support each equilibrium outcome. 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 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 suitable game to study the tension between the short term gain of deviation vs. the long term gain of cooperation.¹ On the other hand, there is a surprisingly small experimental literature on repeated coordination games, where the static version has multiple Nash equilibria (McKelvey and Palfrey, 2001; Ioannou and Romero, 2014a; Mathevet and Romero, 2014). A major common finding in these papers is the impressive ability to achieve long run coordination on fair and Pareto efficient equilibria in games where coordination in the one-shot game is infrequent. The result is interesting in the stag hunt game, where coordination on the most desirable outcome necessitates both players to repeatedly choose the high-risk action. The finding is even more impressive in the battle of the sexes, where fairness and Pareto efficiency require individuals to alternate between the two equilibria, each preferred by a different player.

Intuitively, coordination requires a significant ability to share mutual intentions. Not only both players must intend to alternate but they also need to trust that the other will. This likely builds on three critical elements: (i) theory of mind (ToM), the ability to take the perspective of others to assess their motivations and beliefs; (ii) abstract logical reasoning, the ability to think deductively, recursively and in a counterfactual manner about possible long-term goals of other players; and (iii) reciprocal beliefs, a higher form of ToM whereby each player understands and commits to the

¹See Dal Bó and Fréchet (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 Fréchet (2018) for a detailed survey.

long-term shared goal. Because these abilities develop gradually, we conjecture that performance in coordination games will improve with age.

To address this question, we approach coordination games from a developmental perspective. More precisely, we investigate the behavior of children and adolescents (ages 7 to 16) from a single school in two repeated coordination games: stag hunt and battle of the sexes. Our participants play 2 supergames of each game, with a large number of stages (24). Since such design has never been implemented in children and teens,² we put special emphasis in developing a methodology adequate for all age groups, as well as novel story lines with attractive graphical interfaces (see [section 2](#)). The study builds on current knowledge on the development of ToM and logical thinking to form hypotheses (see [section 3](#)). We test several predictions regarding the developmental trajectory of strategic choices and the differences across games. To do so, we devise a set of strategies of potential empirical relevance, analogous but not identical to those typically discussed in the prisoner’s dilemma literature. We consider basic strategies that take the form of a heuristic or fixed rule, and also elaborate strategies that make use of the history of play. One important issue concerns the evolution of behavior. We hypothesize that young participants employ simple strategies and the transition to more sophisticated ones occurs at an older age.

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). Naturally, such changes have large payoff consequences. Because young children mis-coordinate often, their payoffs are smaller than that of their older peers. 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 concrete and easy to follow strategies are. With such strategies, most groups either achieve coordination in the efficient and fair long run outcome early in the supergame or never at all. Even though the set of strategies is similar in all age groups, the proportion of participants that uses each of them evolves with age. Consequently, the changes with age in the payoffs secured by our participants is driven by the frequency with which each strategy is selected. Third, while the trajectory is similar in both games, the levels 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 the action (both play stag in every period) while battle of the sexes necessitates coordinating the strategy (alternate between the two Nash equilibria), which is arguably more challenging. Second, centration (the tendency of young children to focus on salient features, in this case, their own payoff) pushes towards coordination in stag hunt (always play stag) but not in the battle of sexes (always play the favorite action). In any case, by the second supergame

²The only exception is our companion paper on the alternating dictator game ([Brocas et al., 2018](#)). Strictly speaking, it is not a repeated game of imperfect information, although it shares many similarities with the repeated prisoner’s dilemma. More generally, the majority of experimental studies focus on individual decision making paradigms (see [Sutter et al. \(2019\)](#) and [List et al. \(2021\)](#) for detailed surveys).

coordination is remarkably high in stag hunt, even among our youngest participants. 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 supergame. All this suggests a capacity to learn and adapt very rapidly (after only one exposure) and to bring any lesson to the next game with a new partner.

The paper contributes to the growing literature that evaluates the developmental trajectory of decision-making and strategic thinking and how it is impacted by cognitive development (Sher et al., 2014; Brocas and Carrillo, 2020b, 2021; Fe et al., 2020). Existing theories of cognitive development predict that young children choose very self-centered strategies that disregard partners (Piaget et al., 1967; Donaldson, 1982; Crain, 2015). Choices gradually evolve towards strategies that take into account their partner (Perner, 1991; Wellman et al., 2001; Wellman and Liu, 2004) and build on increasingly complex logic (Rafetseder et al., 2013; De Neys and Everaerts, 2008). By analyzing the strategies revealed through play, we can assess the centration tendencies of participants and their level of sophistication. The resulting trajectory of behavior can then be compared to the theoretical predictions to assess whether behavior tracks cognitive development.

More broadly, our study also relates to the large literature that investigates which strategies support desirable outcomes in repeated interactions. This issue is central for a myriad of fields beyond economics, such as evolutionary biology, anthropology, philosophy and psychology. Researchers in those areas approach this question using classical game theory but also other angles: evolutionary game theory (Smith and Price, 1973; Weibull, 1997), reinforcement learning (Sutton and Barto, 1998; Roth and Erev, 1995), theories that allow evolution to operate on the learning mechanisms (McLoone and Smead, 2014), and genetic algorithms that capture heritability and mutations (Browning and Colman, 2004). Their models provide a theoretical framework to study the dynamics of play between fixed rational players, traits (strategies) that compete for dominance and change constantly, or players with limited information who adapt behavior through trial and error. They offer normative benchmarks that help evaluate the behavior of organisms and their evolutionary trajectory. Our study into the developmental trajectory of dynamic coordination has implications in this broad context. First, it documents for the first time what type of strategies children of different ages use, and compares them to game theoretical predictions. This is important to assess the level of rationality of individuals at different developmental stages. Second, it investigates learning within age groups both between games and between supergames of the same game. Because players must infer the intentions of their partners, observations of play offers critical opportunities to learn, which may or may not be leveraged efficiently.

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 payoffs by game and age-group, and [section 5](#) details the classification of participants according to their best-fitted strategy in each supergame. The contribution of demographic variables is investigated through regression analysis in [section 6](#). [Section 7](#) compares the results in two benchmark adult populations, USC undergraduates and teachers from the school of our participants. Concluding remarks are presented in [section 8](#).

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 \(2020c\)](#) and we develop a graphical version of existing games.³

Participants. Our main population consists of 220 school-age subjects from 2nd to 10th grade at the Lycée International of Los Angeles (LILA), a French-English bilingual private school in Los Angeles. We ran 28 sessions that lasted no more than one school period (50 minutes). Sessions were conducted in a classroom at the school using touchscreen PC tablets and the tasks were programmed with the open source software ‘Multistage Games’.⁴ Sessions had 8, 10 or 12 participants. For each session, we tried to have male and female participants from the same grade, but for logistical reasons, we sometimes had to mix subjects of two consecutive grades. High schoolers from 9th, 11th and 12th grade did not participate in the study because they were taking or preparing for french or US national exams during this period.

The majority of students at LILA are Americans and Europeans from caucasian families of upper-middle socio-economic status. Even though this pool is not representative of the US population, it is homogenous. This allows us to make meaningful age comparisons. Indeed, variations in economic or demographic characteristics have been associated with differences in performance in strategic games ([Charness et al., 2019](#); [Brocas and Carrillo, 2021](#)). By avoiding a mix of participants from different schools, we limit the effect of confounds on the developmental trajectory.

For comparison, we recruited 70 undergraduates at the University of Southern California and ran 6 sessions at Los Angeles Behavioral Economics Laboratory (LABEL) using *identical* procedures. Participants were recruited from the LABEL subject pool. The majority of studies with children do not perform the same experiment with an adult control population. Instead, they compare their findings with results obtained in related paradigms with adults. We believe it is valuable to include an adult control group to establish a behavioral benchmark. This is especially important in an experiment like ours, where the literature on adult behavior is limited and the procedures are different. At the same time, it is important to acknowledge that, for obvious reasons, not all characteristics of the control group match those of the population under study.⁵

Finally, we also had the opportunity to conduct the experiment with 30 teachers at LILA. We conducted 2 sessions with 10 and 12 participants who teach children in pre-K to 5th grade and 1 session with 8 participants who teach children in 6th to 12th grade. We report a comparison of our two adult populations in [section 7](#). A summary of our 330 participants is reported in [Table 1](#).

³The 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.

⁴Software instructions can be found at <http://multistage.ssel.caltech.edu>.

⁵After high school, many students from LILA go to high-ranked colleges in North America and Europe, including USC. While it is a reasonable match for the USC population, some potential differences (such as nationality, family background and size of peer group) exist between the two populations.

	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.

The first task consisted of four trials of a three-options dictator game, with different payoff combinations. Outcomes of this short task were not communicated to participants until the end of the experiment. The results obtained in that task are reported in [Brocas and Carrillo \(2020a\)](#).

After a break, participants moved to the main task, which is the focus of this paper. 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 with the selected payoff values. 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 sections 3.2 and 3.3. We chose these two games because they share many 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 extensively discussed later on, dynamic coordination is expected to be simpler when there is a Pareto superior static Nash equilibrium (**SH**) than when there is not (**BoS**).

	BoS		SH	
	red (Y_2)	green (M_2)	in (I_2)	out (O_2)
red (M_1)	(5,3)	(1,1)	(3,3)	(1,2)
green (Y_1)	(1,1)	(3,5)	(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 (integers between 1 and 5). 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 of the four supergames was identical. Subjects were randomly and anonymously matched with a partner and played 24 rounds of the game with the same partner and with feedback after each round. At the end of the supergame, total payoffs were displayed.

⁶Formally, the basin of attraction is 1/2, so that neither equilibrium is risk dominant. Studies have shown that with these payoff values neither choice is overwhelmingly favored by adults in the one-shot version of the game ([Dal Bó et al., 2020](#)).

New partners were randomly and anonymously drawn and a new supergame was played. In one-half of the sessions participants played two **BoS** supergames followed by two **SH** supergames, whereas in the other half, they played two **SH** supergames followed by two **BoS** supergames.⁷ To avoid end-of-period effects, we did not announce the length of the supergames (the instructions said “you will play many rounds with the same partner”). However, we used the same 24 round length in all four supergames, so some subjects could potentially become aware of it.⁸

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 children as young as 7 years of age. This ruled out the payoff matrices presented in [Table 2](#) as well as other formal representations commonly employed in experimental economics.

For **BoS**, we developed a novel story called the “find the balance” game. In this game, each of the two matched participants in a group was either the red player or the green player. 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 of the information provided. The role (here, 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 (in this example, 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”. In this other game, a blue player and a yellow player possessed a blue and a yellow star, respectively. Each had to decide whether to place their star on or outside a common carpet. Placing the star outside the carpet gave a player 2 points, irrespectively of the other player’s behavior. Placing it on the carpet gave the individual 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.

⁷Length of the experiment is a major constraint in developmental studies due to the limited attention span of participants ([Brocas and Carrillo, 2020c](#)). We believe our supergames are long enough (24 rounds) to allow us to grasp the intentions of participants. On the other hand, 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.

⁸After multiple tests, we concluded that it was preferable to have deterministic unannounced ending rather than the traditional random termination rule. This is less rigorous but much more natural and significantly easier to explain to young children. We think it is not problematic in general, but especially in games where the existence of equilibrium cooperative strategies do not depend on the discount factor. Finally, notice that in our games, if subjects succeed in coordinating in the fair and efficient long run equilibrium, a unilateral deviation in the last period cannot improve the payoff of an individual. This further decreases the relevance of the last period effect.




Figure 1: Experimental design of **BoS**

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 aloud instructions is included in Appendix A.

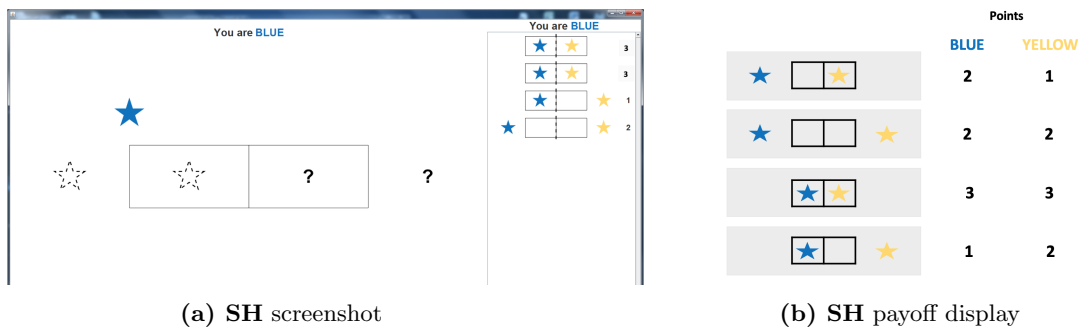


Figure 2: Experimental design of **SH**

Payoffs. During the experiment, subjects accumulated points. Following the methodology described in Brocas and Carrillo (2020c), we used different mediums of payment for different ages. The objective was to equalize, to the best of our ability, the *value of rewards* across age groups instead of equalizing the rewards themselves. USC students, teachers and participants in grades 6 to 10 earned points that were converted into money at a \$0.04 per point conversion rate, and paid immediately after the experiment. USC students and teachers were also paid a \$7 show-up fee, to correct for differences in opportunity cost of time. USC students were paid in cash. Teachers and school-age participants were paid with an Amazon giftcard, since the school does not allow money transactions on premises. 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).

For children in grades 2 to 5 we set up a shop with 20 to 25 pre-screened, age appropriate toys

and stationary.⁹ 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 and, for the youngest participants, we explicitly stated that more points would result in more toys. 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 a scarce resource: no child had excess points after choosing all the toys they liked.¹⁰ 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”.

For the analysis, we group our school-age participants into four naturally clustered age-groups: grades 2-3 (ages 7-9, 54 participants), grades 4-5 (ages 9-11, 54 participants), grades 6-7 (ages 11-13, 48 participants), grades 8-10 (ages 13-16, 64 participants). The control populations include USC students (ages 18-23, 70 participants - U) and LILA school teachers (30 participants - T). Unless otherwise noted, when comparing aggregate choices we perform two-sided t-tests of mean differences. Standard errors are clustered at the individual level whenever appropriate and we use a p-value of 0.05 as the benchmark threshold for statistical significance.

3 Theory and Hypotheses

3.1 Cognitive complexity of coordination games

Playing at equilibrium in one-shot games is cognitively complex, and necessitates a number of abilities that develop gradually during childhood and adolescence. 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 usually exhibit centration, the tendency to focus attention only on one salient dimension of a problem (Piaget et al., 1967; Donaldson, 1982; Crain, 2015). As a result, they mostly pay attention to their own play and payoffs in a game of strategy. With the development of theory of mind—the mental capacity to understand other people’s behavior, intentions and beliefs—children become gradually able to recognize strategic implications (Perner, 1991; Wellman et al., 2001; Wellman and Liu, 2004). Changes in these abilities are profound during elementary school. However, game theoretic paradigms also require logical abilities. A player must use the information regarding their partner recursively to determine a best response. A player must also rely on counterfactual thinking to assess best responses in different contingencies. Hypothetical and counterfactual thinking are known to develop mostly throughout middle school (Piaget, 1972; Rafetseder et al., 2013; De Neys and Everaerts, 2008).

⁹These included gel pens, bracelets, erasers, figurines, die-cast cars and trading cards for younger kids, and apps, calculators and earbuds for older kids. Children, however, could choose any item they wanted.

¹⁰In our experience, incentives are key to retain the attention of children. However, it is also important to avoid excessively high variance in payoffs to make sure that no child feels unhappy. Our payoff-calibration emphasized the value of earning tokens, ensured an enjoyable experience for everyone, and was still in line with the incentives in the adult literature. Notice also that children are familiar with accumulating points (or tickets) and subsequently exchanging them for rewards since it is commonly employed in fairs and arcade rooms.

Coordination in the one-shot versions of the stag hunt and battle of the sexes requires an extra level of sophistication. A player must also model the intention of the other player to target one equilibrium, understanding that the other player is making a similar reasoning. This requires an extraordinary ability to share mutual beliefs, a higher-level form of theory of mind. Coordination in these one-shot games is challenging among adults (Camerer, 2011). We should therefore not expect children to succeed either.

At the same time, and as discussed in the introduction, the previous literature has demonstrated that adults coordinate remarkably well their choices on Pareto optimal equilibria in the repeated version of these two games (McKelvey and Palfrey, 2001; Ioannou and Romero, 2014a; Mathevet and Romero, 2014). This is particularly interesting given the little predictive power of the standard theory, as epitomized by the folk theorem.¹¹ It also contrasts with other dynamic games (for example, the repeated samaritan’s dilemma or the more frequently studied repeated prisoner’s dilemma) where empirical behavior is highly heterogeneous (Dal Bó and Fréchette, 2011, 2018). The behavior of adults thus provides a stark template for comparison. It also suggests the possibility that repetition and observation of past play facilitate the establishment of mutual beliefs and expectations.

Developmental psychologists have reported that children as young as 2 years old exhibit collaborative strategies in simple tasks (Brownell et al., 2006; Warneken and Tomasello, 2007). This ability seems to develop in parallel to other-regarding concerns (Gräfenhain et al., 2009) and theory of mind (Grueneisen et al., 2015), and it is facilitated by communication (Siposova et al., 2018; Wyman et al., 2013). It is however unclear whether repetition alone can help children share mutual beliefs about efficient coordination and whether this ability is innate or develops gradually.

To understand the developing ability to coordinate, it is also important to exploit differences across games. Comparing behavior in repeated versions of stag hunt and battle of the sexes may reveal intrinsic difficulties or abilities to reach jointly beneficial outcomes. For instance, let us 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 the alternation between the two static Nash equilibria: (M_1, Y_2) , (Y_1, M_2) , (M_1, Y_2) , etc.¹² In **SH**, EFO entails the repetition of the static Pareto superior Nash equilibrium: (I_1, I_2) , (I_1, I_2) , (I_1, I_2) , etc. While many different strategies may lead to such outcomes, one thing is clear: 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. In that respect, it is also key that the supergame is played with a fixed partner.¹³

¹¹There are some interesting behavioral learning models, although their performance—especially in **BoS**—is not fully satisfactory, as discussed in the above mentioned research.

¹²Naturally, other strategies where players coordinate half the time on (M_1, Y_2) and the other half on (Y_1, M_2) would result in the same payoff. However, we hypothesized (and empirically verified) that such strategies were unlikely to be played in our game, so we did not consider them in our analysis.

¹³Not surprisingly, McKelvey and Palfrey (2001) finds very significant differences in behavior when the repeated coordination games are played with fixed vs. random partners. There is also an interesting literature on cooperative behavior in dynamic prisoner’s dilemma games with random partners (Camera and Casari, 2009; Camera et al., 2013). For our games, fixing partners is important as it allows us to study the dynamics of coordinated strategies

The key issue is therefore how participants devise strategies, at every age, to reach the EFO. Unlike in the prisoner’s dilemma paradigm, the existing literature in coordination games does not provide clear guidelines regarding which basic strategies may capture best the thinking process of individuals. Our goal is not to provide an exhaustive taxonomy of strategies in our games, but to discuss plausible dynamic behaviors. To do so, we focus on simple and intuitive strategies that we think *could* be employed by our participants in the repeated **BoS** and **SH** supergames.¹⁴

3.2 Battle of the Sexes

In **BoS**, we define actions symmetrically for both players, that is, relative to their most and least favorite actions (not by the color choices). We denote by M_i^t the choice by player i in round t of ‘my favorite’ action (‘red’ for red player and ‘green’ for green player). Similarly, Y_i^t 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) . The EFO consists in alternating between the two, which 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 such 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 choice of the partner in the previous round
(5) TRIG	grim-trigger: coordinate on EFO and play M_i^t forever after any deviation
(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, M) depending on the partner’s behavior, while (6) seeks to repeatedly coordinate in one of the static Nash equilibrium by either exploiting the partner or giving-in. The remaining strategies capture a variety of strategic behaviors: (7) is similar to (5) except that it forgives after one period,

(such as EFO) and the development of mutually shared beliefs within the group.

¹⁴Notice that, in the discussion below, we leave the subject’s initial choice unspecified for some strategies. Needless to say, this makes the strategies incomplete. We deliberately adopt this approach for reasons related to our empirical analysis, as discussed in [section 5](#).

(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.

3.3 Stag Hunt

SH is a symmetric game with symmetric payoffs. We denote by I_i^t and O_i^t the ‘stag’ and ‘hare’ choices by player i in round t (‘in’ and ‘out’ in our game). The static Nash equilibria are (I_1, I_2) and (O_1, O_2) (see Table 2). The existence of a Pareto superior static Nash equilibrium (I_1, I_2) implies that coordination in the EFO of the dynamic game is arguably simpler and more intuitive than in **BoS**. As in section 3.2, we discuss some 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: coordinate on EFO and play O_i^t forever after any deviation
(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 4: Some simple strategies in **SH**

The strategies are similar (though not identical) to those in **BoS**, and their behavioral interpretation also differs in some cases. For example, ALT is plausible but less natural than in **BoS** since alternation between static Nash equilibria decreases the payoff of both players. 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. The remaining strategies capture different ways to instill coordination on (I, I) while sanctioning more or less harshly deviations to O .

3.4 Hypotheses

Centration, Theory of Mind and logical thinking—which develop dramatically during our window of observation—are likely to affect behavior significantly. Young children have a hard time considering the perspective of other people and realizing that the decisions of their partners depend on their goals. They also have a limited capacity to perform multiple steps of reasoning, and therefore anticipate how partners will react to their choices. These two traits are known to develop with age. Centration decreases fast and practically disappears around age 7 whereas strategic thinking improves gradually all the way into young adulthood (Miller, 2002).

While this suggests that behavior should be (weakly) closer to equilibrium as participants get older, its significance is not fully clear. For one thing, and as discussed above, the predictive power

of theory is very limited. If very different paths can be rationalized as equilibrium behavior, it becomes difficult to provide an objective metric for deviations.¹⁵ For another, (weak) improvements can take very different forms. In our past research, we have noted strategic games where improvements are constant and sustained (Brocas et al., 2018), steep with a ceiling in middle school (Brocas and Carrillo, 2021) or negligible in the entire window of observation (Brocas and Carrillo, Forthcoming), among others. With these considerations in mind, we next provide some hypotheses about the expected evolution in behavior of our participants.

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

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

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

Hypothesis H4. *Game comparison.* *Participants 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 (heuristically playing always the same action) with strategies that take into consideration the choice of others. Relatedly, **H2** predicts that older participants will adopt a more complex system of rewards, punishments, and forgiveness via strategies that subtly depend on the choice history of both players.¹⁶ **H3** predicts that the combination of the previous hypotheses will result in more frequent and more efficient coordination, as well as higher gains for older participants compared to their younger peers. 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 hypotheses are supported by the data. If there is a choice trajectory, we will also determine whether behavior changes monotonically, changes in steps, or stabilizes after a certain age.

4 Descriptive analysis

We first analyze the aggregate behavior in each age group averaged over the 24 rounds of each supergame. The top graphs of Figure 3 report information regarding **BoS**. They display the average proportion of M_i choices by individuals (left), the average proportion of (M_i, Y_j) outcomes by pairs of individuals (center), and the average individual payoff (right) in each supergame (BoS1 and BoS2) and each age group. The bottom graphs of Figure 3 report information regarding **SH**.

¹⁵Naturally, one can choose one particular equilibrium and determine the deviations from it, but the choice of such equilibrium may be ad-hoc.

¹⁶Our definition of complexity is admittedly vague. However, it serves the purpose of informally classifying strategies in Tables 3 and 4 in three categories: low (1-2-3), medium (4-5-6) and high (7-8-9) complexity. For a recent formal analysis of rule complexity in individual decision making, see e.g. Oprea (2020).

They display the average proportion of I_i choices by individuals (left), the average proportion of (I_1, I_2) outcomes by pairs of individuals (center), and the average individual payoff (right) in each supergame (SH1 and SH2) and each age group.

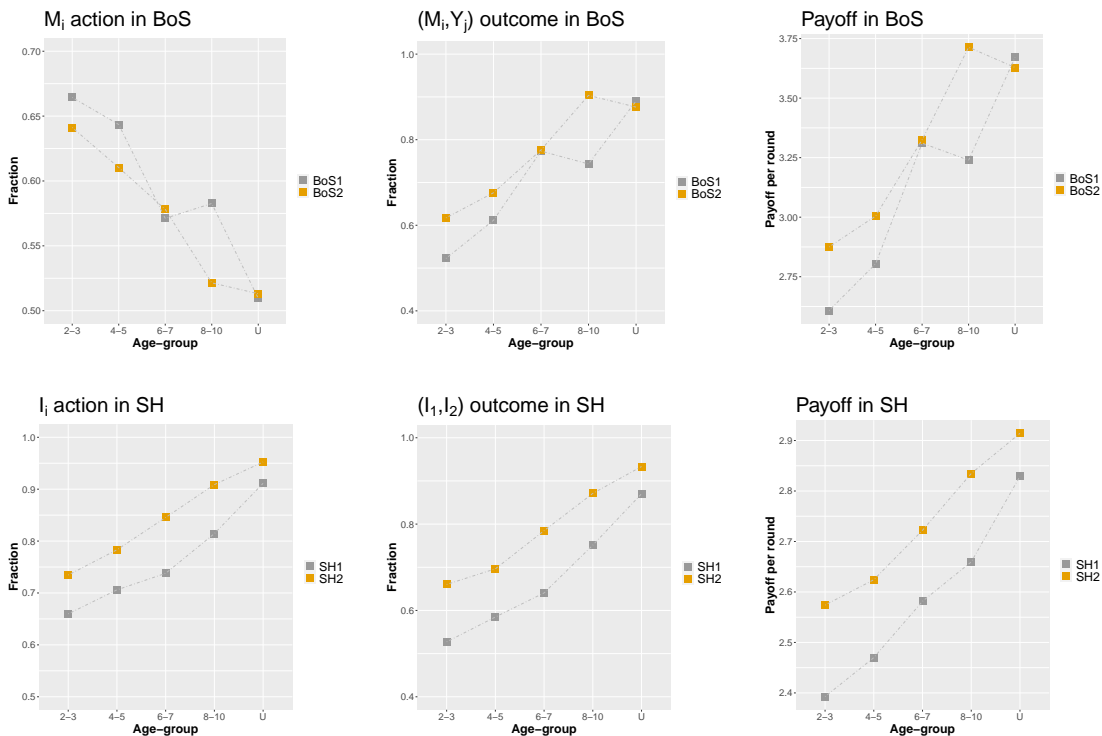


Figure 3: Aggregate behavior and earnings in **BoS** (top) and **SH** (bottom)

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 two 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_i, I_j) in our school-age population starts poorly (52.2% and 52.7% in the first supergames for our 2-3 age-group) and steadily increases to levels similar to those in the adult population (90.3% and 87.1% in the second supergames for our 8-10 age-group). 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. In particular, the behavior

of age groups 2-3 and 8-10 in SH2 is very similar to the behavior of age groups 6-7 and U in SH1, respectively. There is also an improvement, although less dramatic, between BoS1 and BoS2. Finally, the behavior of individuals is highly 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, their payoffs are also highly correlated across supergames ($PCC = 0.43$ in **BoS** and $PCC = 0.44$ in **SH**, $p < 0.0001$). These initial findings will be further investigated in our regression analysis of section 6.1.

We next study the dynamics of outcomes. Figure 4 presents the evolution of the average proportion of groups that achieve coordination on (M_i, Y_j) in **BoS** (top) and on (I_i, I_j) in **SH** (bottom) from round 1 to round 24 in each supergame and age group.

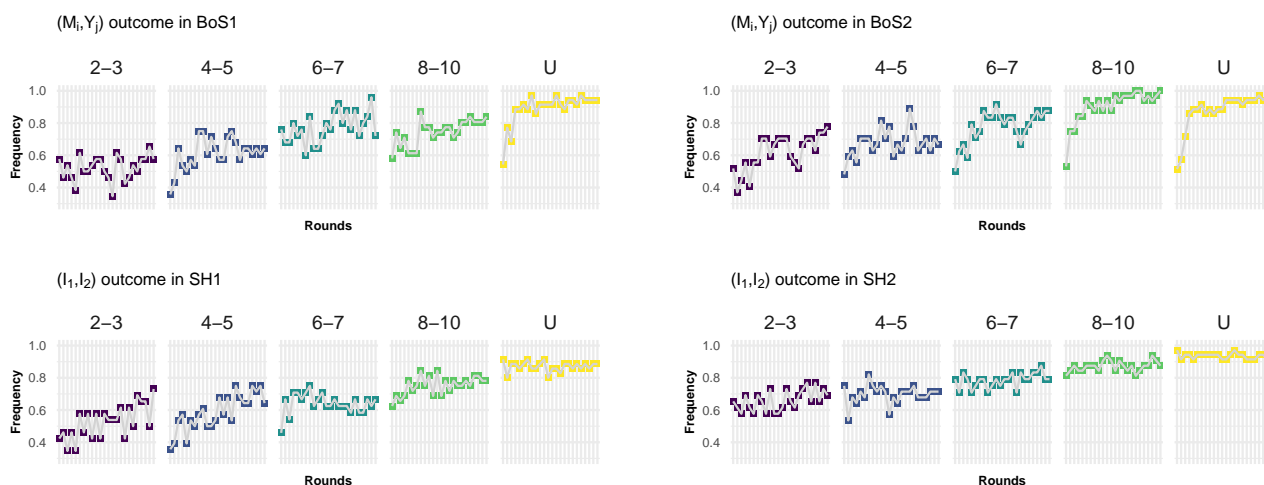


Figure 4: Coordination over rounds in **BoS** (top) and **SH** (bottom)

Not surprisingly given Figure 3 (middle), we find a sustained increase in the level of coordination across age groups. However, there is little evidence of increased coordination across rounds within age groups. In **BoS**, we observe no significant trend in the younger age groups (augmented Dickey-Fuller test, $p > 0.05$). There is a positive trend in the first supergame for age-groups 8-10 and U, although it is entirely 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.

Choices made in the first round of each supergame also provide interesting information regarding the intentions of participants. Table 5 presents the average proportion of outcomes by supergame and age group.

Interestingly, the outcome (Y_1, Y_2) in the first round of **BoS** is infrequent in our younger age groups but it gains some popularity in 8-10 (and even more in the control group). Intuitively,

Age	BoS			SH		
	(M_i, Y_j)	(M_1, M_2)	(Y_1, Y_2)	(I_1, I_2)	(O_1, O_2)	(I_i, O_j)
2-3 [1]	.58	.31	.12	.43	.19	.38
2-3 [2]	.52	.48	.00	.65	.04	.31
4-5 [1]	.35	.54	.11	.36	.18	.46
4-5 [2]	.48	.52	.00	.75	.04	.21
6-7 [1]	.76	.20	.04	.46	.00	.54
6-7 [2]	.50	.42	.08	.79	.04	.17
8-10 [1]	.58	.26	.16	.63	.03	.34
8-10 [2]	.53	.19	.28	.81	.00	.19
A [1]	.54	.11	.34	.91	.00	.09
A [2]	.52	.14	.34	.97	.00	.03

[supergame in brackets]

Table 5: Proportion of behavior in the first round of the supergame by age group

playing the preferred action of the partner may be a way to signal an intention to cooperate for mutual advantage. Despite limited statistical power, a chi-square test of homogeneity confirms that differences exist in the distribution of the first round’s outcome across school-age groups (4-5 vs. 6-7 in BoS1, 2-3 vs. 8-10 in BoS2 and 4-5 vs. 8-10 in BoS2). We also note that undergraduates start with (I_1, I_2) more frequently than all school-age children in **SH**. There is also a significant shift from (I_i, O_j) to (I_1, I_2) between SH1 and SH2, especially in age groups 4-5 and 6-7 (p-value < 0.02). It suggests that participants exhibit an increased willingness to coordinate right from the outset after experiencing just one supergame.

The convergence process is indicative of the ability of participants to share mutual beliefs about their intentions. The next question is whether the willingness and ability to coordinate lead to convergence to EFO, that is to repeat (I_1, I_2) in **SH** and to alternate between the two static Nash equilibria in **BoS**. Figure 5 reports the distribution of rounds at which convergence to EFO is reached by age group and supergame. More precisely, 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.¹⁷

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. We do not observe any noticeable difference across supergames but there are some differences between the two games. Indeed, early

¹⁷Alternatively, we could determine the longest streak consistent with EFO to better classify individuals who, for example, deviate only at the very end. This method, however, would pool groups that coordinate at the beginning and finish in a punishment phase with groups that take time to coordinate but eventually succeed. We therefore preferred the other approach.

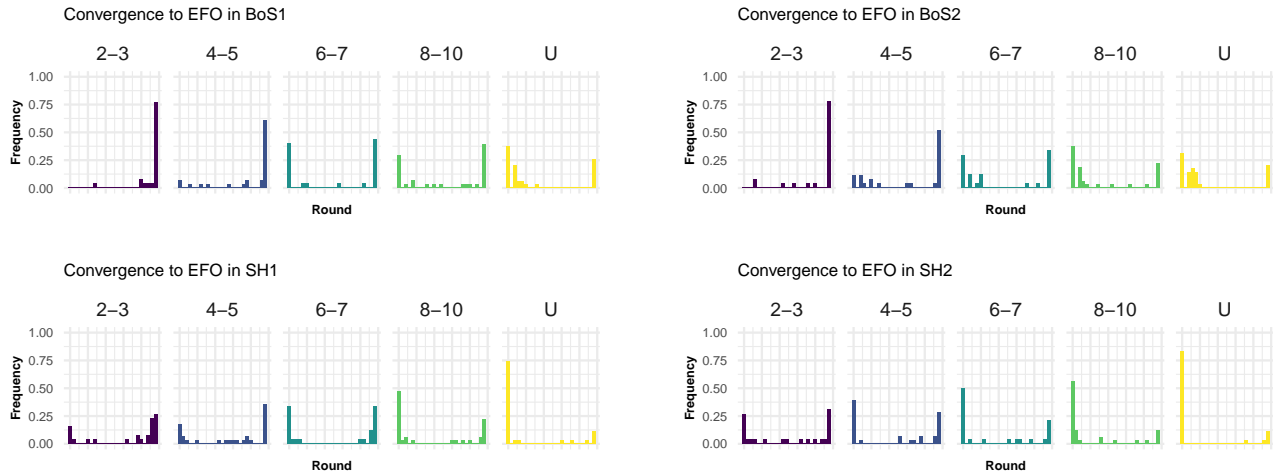


Figure 5: Round of convergence to EFO in **BoS** (top) and **SH** (bottom)


coordination within an age group is more frequent in **SH** than in **BoS**, as reflected in a higher left mode, 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**. This difference is particularly noticeable in our adult control group. Note that 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 age 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. The set of attainable 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 quite remarkable, and 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 age groups. Behavior is mostly symmetric, with payoffs near the ones corresponding to the other static Nash equilibrium $(2, 2)$ accounting for the second largest fraction of choices. There are almost no instances of players persevering in I when the partner has decided to play O (bottom left quadrant 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


Payoffs in BoS1



Payoffs in BoS2



Payoffs in SH1



Payoffs in SH2




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

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. 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**. But only the fact that participants tend to converge to EFO more often with age in **BoS** demonstrates that they move away from centered behavior with age. Indeed, and as noted earlier, saliency promotes efficient behavior in **SH** but it does not in **BoS**. This transition is evidence in favor of **H1**.

5 Individual strategies

Descriptive analysis allows us to provide an overall picture of the main trends. However, strategic behavior in game theoretic paradigms is usually heterogenous. We next investigate this heterogeneity and we 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 subject to possible random miscoordinations, and we therefore ignore them. We use the outcome in round 4 as the anchor (or initial condition) for the strategy. We consider all the strategies described in [Table 3](#) for **BoS** and [Table 4](#) for **SH**. 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 retained are classified as OTHER. An important remark is in order. The behavior of a player in a supergame may be compatible with several strategies.¹⁸ With a large number of supergames, it would be possible to disentangle between different strategies by studying the behavior against different partners.¹⁹ Unfortunately, this is not possible with only two supergames. We shall however keep this limitation in mind in the analysis and interpretation of the results. In particular, each participant is assigned a combined strategy when several strategies explain their behavior equally well. Last, we decided that analyzing the choice of individuals by imposing the same strategy in the whole experiment was not adequate. While such approach is suitable for studying steady-state behavior when the number of supergames is sufficiently large, it seems inappropriate with only two supergames and a very real possibility that individuals change strategies between the two.

5.1 Strategies in Battle of the Sexes

Classifying the maximum number of individuals with the fewest number of strategies can be delicate. In our case, however, it turned out to be relatively uncontroversial. Indeed, with only four

¹⁸For example, EFO in **BoS** may be the result of two individuals playing, ALT, TFT, TRIG, or FORG.

¹⁹The literature on the repeated prisoner’s dilemma uses a variety of techniques, such as trading-off goodness of fit and number of strategies ([Camera et al., 2012](#)), maximum likelihood estimation of the best fitting strategy ([Aoyagi and Fréchet, 2009](#); [Dal Bó and Fréchet, 2011](#); [Fudenberg et al., 2012](#)) and direct elicitation of strategies ([Romero and Rosokha, 2018, 2019](#); [Dal Bó and Fréchet, 2019](#)).

strategies—ME, ALT, TFT and TEST—we managed to account for the choices of a large fraction of participants (72.8% and 80.3% in BoS1 and BoS2, respectively).²⁰ Furthermore, including all five remaining strategies (YOU, TRIG, REV, FORG and TEACH) would classify only an additional 3 participants in BoS1 and 3 participants in BoS2.²¹ We therefore decided to not include more strategies. As noted before, different strategies may lead to the same choices (depending on the strategy of the partner with whom a player is matched). 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. Of these individuals, 95.3% and 97.1% have 0 deviations from perfect alternation in BoS1 and BoS2, respectively. They also converge instantly to EFO. These participants differ from individuals classified as ALT and individuals classified as TFT. The former are participants who have suffered deviations of their partners. They had a tendency to not punish those deviations and, instead, continued alternating. Their only chance to reach EFO with few deviations is to be matched with someone who also alternates, which is relatively rare in our data.²² 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 unfrequent.²³

²⁰Allowing 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.

²¹In particular, there are no participant in our sample who start playing at EFO and end up punishing and not forgiving a partner's deviation. This absence of grim trigger behavior has been documented in some repeated prisoner's dilemma experiments with adults (see e.g., Camera et al. (2012)).

²²Only 32% and 65% of them coordinate on a static Nash equilibrium more than 16 times in BoS1 and BoS2. Also, only 0% and 25% converge to EFO in less than 5 rounds in their groups in BoS1 and BoS2. This indicates that EFO is quite low for those subjects.

²³Only 27% and 70% of them coordinate on a static Nash equilibrium more than 16 times in BoS1 and BoS2. Also, only 0% and 21% converge to EFO in less than 5 rounds in their groups in BoS1 and BoS2.

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 Venn diagram is revealing as it partially separates different motives for identical choices. 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. Figure 8 presents the distribution of strategies in each age group.

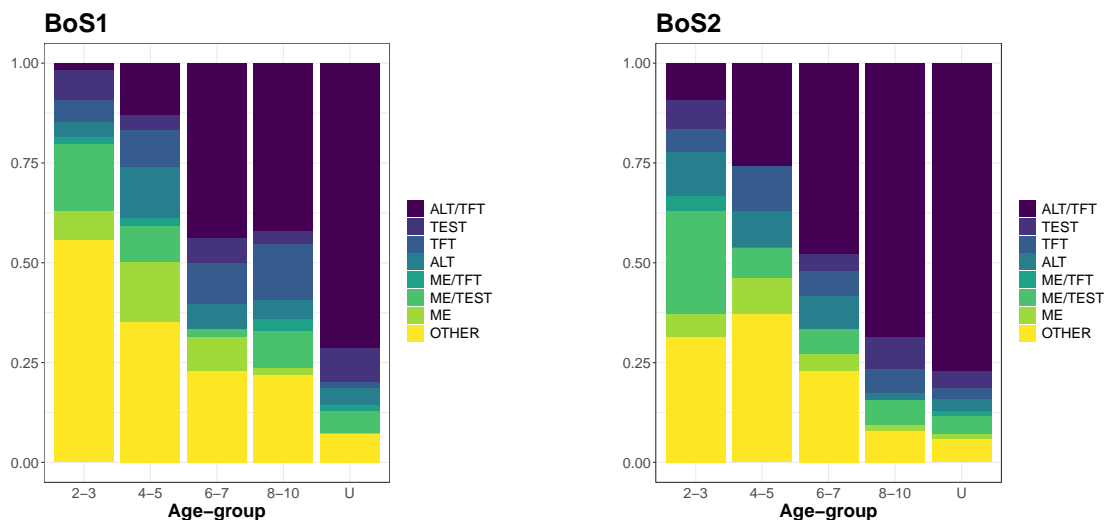


Figure 8: Distribution of strategies across age groups in BoS1 (left) and BoS2 (right)

Perhaps not surprisingly in light of the results in 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, especially among players whose partner attempts to coordinate (ME and ME/TEST), decreases with age. A decreasing trend is also observed among the unclassified players (OTHER). It suggests that a large fraction of our youngest participants have problems devising a strategy. The difficulty to follow some discernible strategy 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 age 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).

Table 6 focuses on the evolution of behavior across supergames. It reports the number of players as a function of their strategy in the first (row) and second (column) supergame.

We obtain several findings. First, most players do not change their strategy between BoS1 and

		BoS2							
		ALT/TFT	TEST	TFT	ALT	ME/TFT	ME/TEST	ME	OTHER
BoS1	ALT/TFT	91	1	5	4	0	1	1	3
	TEST	1	5	1	1	0	2	0	7
	TFT	13	0	3	1	0	0	0	6
	ALT	11	0	1	1	0	0	0	5
	ME/TFT	2	0	1	0	1	1	0	0
	ME/TEST	4	0	0	1	0	14	1	5
	ME	2	0	2	1	1	4	6	1
	OTHER	16	8	5	9	1	6	4	30

Table 6: Individual strategies across supergames in **BoS**

BoS2: some successfully coordinate on EFO (ALT/TFT, 91 subjects), others play selfishly despite opportunities not to do so (ME/TEST or ME, 25 subjects), yet some others behave unpredictably (OTHER, 30 subjects) in both supergames. Second, among participants who play differently, behavior can be considered to improve, at least on average. Indeed, 49 participants coordinate on EFO in BoS2 (ALT/TFT) but not in BoS1 whereas 15 participants do the reverse. It is particularly revealing that among individuals who demonstrably try both times to coordinate on EFO, 24 subjects are successful only the second time (ALT or TFT in BoS1 and ALT/TFT in BoS2) while 9 are successful only the first time (ALT/TFT in BoS1 and ALT or TFT in BoS2). More generally, strategies become more predictable (i.e., more classifiable) in the second supergame.

The choice of strategy has important payoff consequences. Table 7 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).

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

(standard errors in parenthesis)

Table 7: Average payoffs as a function of the strategy in **BoS**

Corroborating previous findings, ALT/TFT are associated to participants who coordinate perfectly on EFO, thereby obtaining a payoff close to 4 on average. Strategies ME/TFT are associated to groups where both partners consistently choose the selfish action, thereby obtaining a payoff close to 1 in every round. Meanwhile, participants classified as ME and especially ME/TEST manage

to exploit their partner, at least in some rounds. However, they still do worse, on average, 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. 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 . Finally, it is worth noting that 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

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 SH2.²⁴ 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 (i.e., where both partners always play I) are classified as IN/TFT. Among these participants, 95.3% and 91.6% have 0 deviations from perfect coordination on EFO in SH1 and SH2, respectively. Similarly, players in groups where both partners always play the Pareto inferior outcome (O, O) 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

²⁴ Again as a robustness check, 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.

latter. Players classified as TFT predominantly coordinate on an equilibrium, but the equilibrium changes over the course of the supergame. 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 resistance from their partner although, as we will see later on, such resistance is usually sparse. Finally, players classified as OUT are those 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 age group.

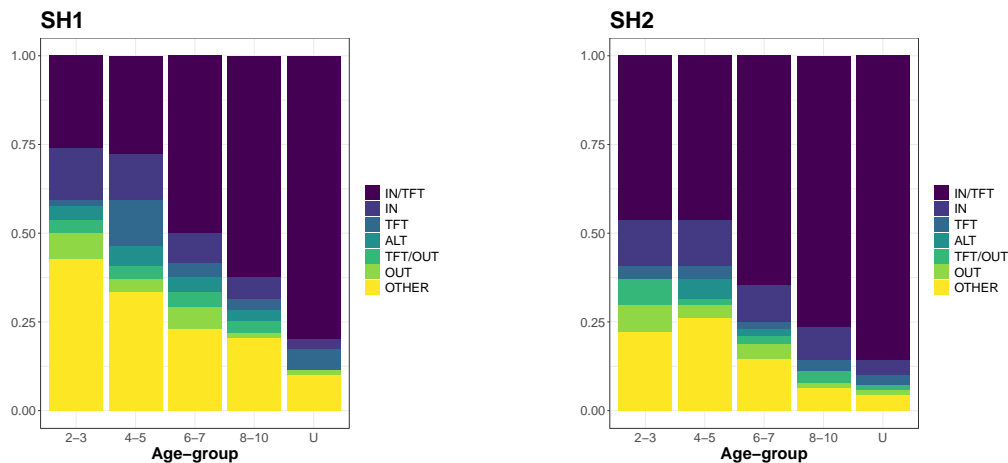


Figure 10: Distribution of strategies across age groups in SH1 (left) and SH2 (right)

As in **BoS**, there is a significant and sustained increase in EFO with age (strategy 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 is infrequent and spread throughout all age groups (OUT/TFT). Choosing I irrespective of the behavior of the partner (IN) is more common in the younger population, although it mostly captures occasional deviations by the partner. More generally, we again notice improvements in the selection of strategies, with a given age 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 emphasizing that even our youngest participants play this game remarkably well: almost half the groups of 2nd and 3rd graders (7 to 9 years old) manage to coordinate perfectly in the Pareto superior equilibrium by the second time they play this game.

Table 8 reports the number of players exhibiting each strategy in the first (row) and second (column) supergame of **SH**.

The vast majority of participants maintain their strategy across supergames. Most seem to target coordination on EFO (IN/TFT or IN, 158 subjects) while a few are unclassified (OTHER, 24 subjects). We observe significant improvements, mostly by participants who switch from OTHER or TFT to IN/TFT (41 participants). However, a few individuals succeeded in coordinating on EFO

		SH2						
		IN/TFT	IN	TFT	ALT	TFT/OUT	OUT	OTHER
SH1	IN/TFT	125	12	3	0	1	1	7
	IN	18	3	2	0	0	0	2
	TFT	9	2	1	1	0	0	3
	ALT	1	1	0	3	1	0	3
	TFT/OUT	2	1	0	0	5	0	0
	OUT	3	0	0	0	2	5	1
	OTHER	32	9	3	0	0	4	24

Table 8: Individual strategies across supergames in **SH**

only in the first supergame. As in **BoS**, the number of participants with a non-discernible behavior decreases between SH1 and SH2 (72 to 40), and we conjecture that it would have stabilized to an even lower number if we had run a few more supergames. Two reinforcing factors may have contributed to this improvement: learning how to play and facing a partner who knows better how to play.

Finally, [Table 9](#) reports the payoffs associated with 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 (.00)	2.54 (.09)	2.48 (.08)	2.38 (.04)	1.92 (.02)	2.00 (.01)	2.14 (.03)	2.63 (0.03)
SH2	2.99 (.00)	2.75 (.04)	2.63 (.06)	2.39 (.07)	1.97 (.01)	1.95 (.01)	2.08 (.040)	2.76 (0.02)

Standard errors in parenthesis

Table 9: Average payoffs as a function of the strategy in **SH**

By construction, the payoffs of participants classified as IN/TFT are close to 3 and the payoffs of participants classified as TFT/OUT and OUT are close to 2. Participants classified as IN have a relatively high payoff, in particular because their partner also choose *I* in 78% and 88% of the rounds, on average, in SH1 and SH2. As explained earlier, deviations by partners are sporadic, which makes it easy to coordinate. Participants classified as ALT incur significant losses. This is the case because, empirically, they end up coordinating on (I, I) only between 7 and 10 times. The per-round payoff of participants with an unidentified strategy (OTHER) is not much higher than the payoff of someone 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 non-discernible or purely selfish strategies in the younger age group, and ending with strategies that largely 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**, with sustained improvements both with age and across supergames. In both cases, most participants adhere to one of a small number of strategies. There is also no sign of stabilization in behavior after a certain age in either game. Interestingly, participants in all age-groups, but especially the younger ones, find it easier in **SH** than in **BoS** both to avoid an undiscernible strategy as well as to coordinate their behavior in the EFO. Indeed, the proportion of strategies leading to EFO is 36.6% in BoS1 against 60% in SH1. These numbers increase to 48.3% in BoS2 and 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** reiterates the strong support for **H1**, and the overall increased performance in both **BoS** and **SH** is evidence in favor of **H3**. The significant differences between games is consistent with **H4**. By contrast, **H2** is not supported. Participants in all age-groups restrict their attention to relatively simple strategies (ALT or TFT) and avoid more complex options (only TEST is selected among options 7-8-9 in both games). A posteriori, this is not excessively surprising. We know since the pioneering work by [Axelrod \(1985\)](#) that excessively sophisticated strategies are neither empirically optimal nor widely common in the general population. Our school-age participants, especially the older ones, manage to coordinate on the EFO with simple (though not simplistic) strategies. Interestingly, the data provides support for learning, a result we had not hypothesized given the short window (only two supergames). This result is important because it suggests that children have an intrinsic ability to “skip developmental stages” when they are exposed, even if briefly, to coordination problems.

6 Regression analysis

This section reports regression analysis. Our adult control group is a benchmark of comparison only. In particular, it is not the culmination of the developmental trend of this specific school population. Therefore, to avoid polluting the trajectory, we include only the 220 school-age participants (grades 2 to 10). Since we are particularly interested in changes in behavior as children develop, 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 (minor) age differences may exist between participants in the same grade.

6.1 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 $Order(SH)$ ($= 1$ 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 ($Male = 1$), whether the participant has one or more siblings ($Siblings = 1$), and the supergame from which the observation is drawn ($BoS2 = 1$ or $SH2 = 1$) to determine potential changes across supergames. Notice that these variables cannot be included in the regressions on outcomes, since these observations are at the group level, and partners change across supergames. We therefore perform regressions on outcomes separately for each supergame ([1] or [2]). The results are reported in Table 10.

	BoS				SH			
	Pr(M_i)	Pr(M_i, Y_j)		Payoff	Pr(I_i)	Pr(I_i, I_j)		Payoff
		[1]	[2]			[1]	[2]	
<i>Grade</i>	-0.015*** (.003)	0.029*** (.008)	0.040*** (.007)	0.119*** (.014)	0.026*** (.005)	0.034** (.013)	0.037** (.013)	0.042*** (.008)
<i>Order(SH)</i>	-0.061*** (.018)	0.120** (.040)	0.072 (.037)	0.336*** (.072)	-0.006 (.027)	-0.044 (.067)	-0.025 (.065)	-0.069 (.040)
<i>Male</i>	0.039* (.018)	—	—	-0.001 (.072)	-0.026 (.027)	—	—	-0.031 (.040)
<i>Siblings</i>	0.020 (.021)	—	—	0.034 (.084)	0.046 (.032)	—	—	0.111* (.046)
<i>BoS2</i>	-0.030 (.018)	—	—	0.256*** (.072)	—	—	—	—
<i>SH2</i>	—	—	—	—	0.088** (.027)	—	—	0.164*** (.040)
<i>const.</i>	0.698*** (.030)	0.332*** (.054)	0.385*** (.049)	2.110*** (.121)	0.567*** (.046)	0.458*** (.089)	0.561*** (.086)	2.256*** (.067)
Adj. R ²	0.068	0.162	0.230	0.191	0.068	0.048	0.056	0.106
# obs.	440	110	110	440	440	110	110	440

Table 10: OLS Regressions of choices, outcomes and payoffs

As expected, age is a very powerful predictor of equilibrium play in both games. Equilibrium outcomes are 2.9 to 4.0 percentage points higher as we move from one grade to the next. Per-round payoff increases per grade are, on average, 0.12 points in **BoS** and 0.04 points in **SH**. 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 highly significant, at least at the 1%-level and often at the 0.01%-level. We observe significant improvements in the second supergame of both **BoS** and **SH**, reinforcing the idea that learning and adapting to the behavioral consequences of choices are key for performance. Also,

participants who start with the arguably simpler stag hunt game perform significantly better when they move to the battle of the sexes than those who start with this more complex game. It indicates portability across games, although it is unidirectional (no order effect is found for **SH**, perhaps because behavior is closer to equilibrium so there is less room for improvement). Finally, we found no systematic effect of gender or siblings on the behavior of our school-age participants.

6.2 Strategies

Section 5 reported heterogeneity in the choice of strategies within and across age 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**). To study the contribution to each type of strategy, we construct a strategy outcome variable and we classify participants into those who choose such strategy (coded as 1) and those who do not (coded as 0). We then conduct a Probit regression of the strategy outcome variable on *Grade*, *Order(SH)*, *Male* and *Siblings* as well as a variable that captures the group’s initial behavior. To wit, we use dummy variables $1st(M_1, M_2)$ in **BoS** and $1st(I_1, I_2)$ in **SH** that take value 1 if the first round’s outcome is (M_1, M_2) and (I_1, I_2) respectively. The rationale is that initial choices will serve as an anchor or a signaling device and therefore be conducive of non-cooperative and cooperative behavior, respectively.²⁵ The results are presented in Table 11.

In support of previous findings, age is a very strong predictor of EFO strategies in both games. 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 Table 10, playing stag hunt first leads participants to play closer to EFO and away from inferior strategies in **BoS**, while it has a negligible effect on **SH**. 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** (even after controlling for all other variables), but only in the first supergame. Conversely, both individuals starting in I_i is a strong predictor of EFO in **SH**.

We conducted ordered probit regressions of changes in strategies between supergames (not reported for brevity but available upon request). They confirm the significant increase in performance documented earlier, with participants moving from strategies not conducive to EFO to strategies conducive to EFO. However, these effects are not modulated by age, except in BOS for the 8-10 age 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

²⁵Remember 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 $1st(I_1, I_2)$.

	BoS				SH			
	<i>EFO</i> ^o		<i>Inferior</i> [†]		<i>EFO</i> [‡]		<i>Inferior</i> [§]	
	[1]	[2]	[1]	[2]	[1]	[2]	[1]	[2]
<i>Grade</i>	0.244*** (.044)	0.297*** (.044)	-0.097* (.044)	-0.165*** (.047)	0.109** (.037)	0.122** (.042)	-0.063 (.052)	-0.094 (.054)
<i>Order(SH)</i>	0.636** (.221)	0.760*** (.204)	-0.689** (.221)	-0.533* (.220)	-0.275 (.191)	-0.131 (.208)	-0.540* (.266)	-0.207 (.259)
<i>Male</i>	0.032 (.213)	-0.102 (.202)	0.603** (.219)	0.119 (.219)	0.184 (.191)	-0.263 (.210)	0.545* (.271)	0.475 (.270)
<i>Siblings</i>	-0.268 (.246)	0.066 (.231)	-0.014 (.246)	-0.253 (.238)	-0.190 (.226)	0.115 (.243)	-0.251 (.283)	-0.286 (.282)
<i>1st(M₁, M₂)</i>	-1.038*** (.286)	-0.222 (.216)	0.506* (.217)	0.242 (.221)	—	—	—	—
<i>1st(I₁, I₂)</i>	—	—	—	—	1.407*** (.196)	1.740*** (.229)	—	—
<i>const.</i>	-2.139*** (.387)	-2.386*** (.391)	-0.580 (.351)	0.168 (.364)	-0.972** (.308)	-1.242*** (.368)	-0.955* (.397)	-0.895* (.394)
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

Note 1. The strategies included are: ^oALT/TFT; [†]ME, ME/TEST and ME/TFT; [‡]IN and IN/TFT; [§]OUT and TFT/OUT

Note 2. Including *1st(I₁, I₂)* in the *Inferior* regressions leads to robustness problems due to lack of observations.

Table 11: Probit Regressions of individual strategies

SH2, and 0.27 between **BoS** and **SH** ($p < 0.001$ in all cases), suggesting that the ability to reach fair and efficient outcomes carries over not only from one supergame to the next, but also to a different game. 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** ($p = 0.015$). This means that the inability to coordinate in one game or supergame is also a trait that extends to other situations. Importantly, these results continue to hold after controlling for age. This suggests that the observed variations in strategic thinking that carry over contexts are not only age-related: variation is also due to differences in general cognitive abilities.

7 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 limited (30 participants). There is also significant heterogeneity in terms of their age and academic achievement (from BA to PhD). 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. At the same time, teachers spend a large fraction of their time in the same environment as the school age group.

Table 12 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 section 6.2 (*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	<i>EFO</i> ^o	0.25	0.67	0.71	0.77
		<i>Inferior</i> [†]	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	<i>EFO</i> [‡]	1.00	1.00	0.83	0.90
		<i>Inferior</i> [§]	0.00	0.00	0.01	0.02

^oALT/TFT; [†]ME, ME/TEST and ME/TFT; [‡]IN and IN/TFT; [§]OUT and TFT/OUT

Table 12: Summary statistics of control populations

Teachers perform significantly 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 striking in the proportion of EFO (ALT/TFT strategies), suggesting an initial difficulty to 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 our undergraduates play this game very well, they are still outperformed by teachers who achieve perfect coordination, thereby maximizing earnings. Despite the improvement of undergraduates in the second supergame, the 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.

8 Conclusion

This paper has investigated the developmental trajectory of the ability to coordinate on ‘desirable’ equilibria in the repeated battle of the sexes (**BoS**) and the repeated stag hunt (**SH**) games. Desirable equilibria have been defined as equilibria that lead to efficient and fair outcomes (EFO): alternation between the two static Nash equilibria in **BoS** and repetition of the Pareto superior static Nash equilibrium in **SH**. The study has three distinctive features. It is the first to investigate the developmental trajectory of behavior in long repeated games. It also proposes a novel methodology and a story line that can be exported to other populations of children and teenagers, to other age groups and, more generally, to populations who might find abstract representations challenging. Last, we construct strategies of potential empirical relevance analogous but not identical to those studied in other repeated games. In our analysis, we pay special attention to the short term learning curve, instead of studying convergence and steady state behavior. We believe this is relevant for situations that occur only from time to time in real life. In both games, we observe a significant increase in the ability to coordinate on EFO with age. This age trajectory is marked by (i) a decreased usage of strategies that ignore the behavior of partners and (ii) an increased usage of reactive strategies capable of supporting EFO. We have also noticed a better performance in **SH** than in **BoS**, and a significant improvement between supergames in both cases. Some final comments are in order.

Even though studies in developmental psychology have reported evidence of collaborative behavior in children, the focus is usually limited to very simple games (Brownell et al., 2006; Grueneisen et al., 2015) and very narrow age ranges (Wyman et al., 2013; Siposova et al., 2018). It is, however, critical to extend this knowledge to more complex interactions, in which reaching desirable outcomes is challenging. Such games allow us to determine how sophisticated strategies are, how learning acts on their selection and whether there is convergence towards desirable outcomes. Studying variations in these aspects across age groups is also important. Indeed, it allows us to assess whether coordination is innate or acquired, intuitive or learned. Our study indicates that, while the ability to coordinate through repeated interactions may improve rapidly through play even among young children, it is not entirely natural and, instead, develops gradually. Therefore, the fact that adults coordinate very well in these games does not result from an instinctive tendency to see the mutually beneficial outcomes. Rather, development acts on traits that are gradually expressed.

Concretely, we can see from our data how development operates. The decrease in self-centered strategies that ignore the behavior of the partner is consistent with the known gradual replacement of centration with ToM. As children grow, they learn to form beliefs about the intentions and goals of people they interact with (Crain, 2015; Perner, 1991). This opens the door to the development

of strategic responses. In the context of our games, children can use inductive logic, an ability that starts developing in early elementary school (Feeney and Heit, 2007; Rhodes et al., 2008), and build a theory of the best course of action from the observation of desirable outcomes. This reasoning is quite simple to implement in **SH**. Indeed, once a child selects the action with highest potential, I , and observes the desirable outcome (I, I) , it is enough to repeat the same action to observe the same outcome again, and establish the causal link between action and outcome. 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 may lead to mis-coordination this time. To infer the optimality of alternation, it is necessary to keep track of past moves, observe longer sequences, and understand the need to sacrifice current payoff in order to induce collaboration. These conceptual differences are likely the main reason why children coordinate both better and faster in **SH** than in **BoS**. Older children may also resort to deductive logic, an ability that develops after inductive logic (Goswami, 2011), and requires building a model and deducing a course of action from it. But here again, finding the model is more complex in **BoS**. It requires thinking in abstract terms, running different counterfactuals to assess what the partner might do and using this recursively to infer what one should do. Eventually, one must also realize that interests across players differ but can be reconciled through alternation. Further remains the question of whether the partner is cognitively capable of reaching the same conclusion, which should eventually lead a player to select a strategy designed to keep the partner on track (e.g., tit-for-tat). This kind of logic develops gradually in middle and high school (Rafetseder et al., 2013). Overall, the differences observed in terms of strategies, but also in terms of time to convergence to coordination reflects the differences in logical abilities. Older children and adults are more successful because they are better equipped to think inductively, deductively and abstractly. And yet, some very young children are successful too.

Our study also points to important phenomena that should be studied further, in particular, learning within and across supergames. Our classification strategy permits only 3 rounds to test a partner before settling on a strategy. The fact that a large fraction of participants are classified based on their behavior in the remaining rounds (with few deviations and despite no prior exposure to the game) indicates that this short test period provides enough time to learn how to best-respond to a partner. It is also interesting that behavior of a given age 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 to reach that outcome. And, importantly, what kind of learning explains these improvements: is it imitation or the application of some elements of logic they already possess? Are they able to export what they learn to other games, to other contexts and to encounters in the more distant future? 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. Perhaps learning how to coordinate on an action is a necessary step to unlock the key of the more complex coordination on a strategy. It is also possible that those who have completed these steps in the correct order have

acquired more durable (and transferable) knowledge and may be more able to achieve coordination in other games.

The process of convergence at different ages should also be investigated further. One way to approach the problem is through variants of the win-stay-lose-switch reinforcement learning model (Roth and Erev, 1995; Erev and Roth, 1998).²⁶ It would be interesting to fit the behavior of children who play more and longer supergames, and describe their ability to balance exploration and exploitation to move from strategies conducive of inferior outcomes to strategies conducive of EFO. We conjecture that young children would be best modeled as independent learners who mutually ignore each other, consistent with their limited ToM and limited ability to apply abstract logic and recursive arguments. By contrast, older children and adults would be best classified as joint-action learners (Claus and Boutilier, 1998; Busoniu et al., 2008) due to their greater cognitive skills. However, given that differences in the complexity of learning strategies do not necessarily matter to achieve Pareto superior outcomes (Claus and Boutilier, 1998; Kimbrough and Lu, 2005), it is likely that all players will end up achieving EFO when given enough time. And since reinforcement learning models do not properly account for the human ability to share mutual beliefs, some players may even outperform the models, converging much faster to EFO than the model predicts.²⁷

Finally, individual behavior also points to underlying biological mechanisms. Indeed, while population averages and their evolution with age can help us delineate a developmental template, differences within age groups are considerable. A recurrent finding in our studies with homogeneous populations of children and adolescents (Brocas and Carrillo, 2020b, 2021) is that a significant proportion of very young children play optimally (e.g., the dominant strategy) in reasonably sophisticated games while a non-negligible fraction of late adolescent and adults do not (e.g., use a dominated strategy). Such large heterogeneity in behavior points to differences in genetic makeup. This type of studies are therefore informative about the (highly inter-related) contribution of nature and nurture to human cognitive abilities and behavior. This observation is important to design interventions, such as school improvements that optimize human capital accumulation (Cunha et al., 2006; Cunha and Heckman, 2007). In other words, to study the importance of parental environment on the relationship between schooling and life outcomes, one must control for genetic components.

²⁶The more general Experience Weighted Attraction (EWA) model, an adaptive learning model that accounts for foregone payoffs (Camerer and Ho, 1999), may also be leveraged to study learning after properly adapting it to repeated game settings (Ioannou and Romero, 2014a).

²⁷It would be interesting to compare experimental data across ages with the predictions derived in Ioannou and Romero (2014b).

References

- Masaki Aoyagi and Guillaume Fréchet. Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic theory*, 144(3):1135–1165, 2009.
- Robert M Axelrod. *The evolution of cooperation*. Basic Books, New York, 1985.
- Simon Baron-Cohen. The evolution of a theory of mind. In M.C. Corballis and S.E.G. Lea, editors, *The descent of mind: Psychological perspectives on hominid evolution*, pages 261–277. Oxford University Press, 1999.
- 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. The evolution of choice and learning in the two-person beauty contest game from kindergarten to adulthood. *Games and Economic Behavior*, 120: 132–143, 2020b.
- Isabelle Brocas and Juan D Carrillo. Introduction to special issue “understanding cognition and decision making by children.” studying decision-making in children: Challenges and opportunities. *Journal of Economic Behavior & Organization*, 179:777–783, 2020c.
- 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. The development of randomization and deceptive behavior in mixed strategy games. *Quantitative Economics*, Forthcoming.
- Isabelle Brocas, Juan D Carrillo, and Niree Kodaverdian. Altruism and strategic giving in children and adolescents. USC working paper, 2018.
- Celia A Brownell, Geetha B Ramani, and Stephanie Zerwas. Becoming a social partner with peers: Cooperation and social understanding in one-and two-year-olds. *Child development*, 77(4):803–821, 2006.
- Lindsay Browning and Andrew M Colman. Evolution of coordinated alternating reciprocity in repeated dyadic games. *Journal of Theoretical Biology*, 229(4):549–557, 2004.
- Anke F Bullinger, Emily Wyman, Alicia P Melis, and Michael Tomasello. Coordination of chimpanzees (pan troglodytes) in a stag hunt game. *International Journal of Primatology*, 32(6): 1296–1310, 2011.
- Lucian Busoni, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.

- Peter Calow. *Life cycles: an evolutionary approach to the physiology of reproduction, development and ageing*. Springer Science & Business Media, 2012.
- 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. *proceedings of the National Academy of Sciences*, 110(37):14889–14893, 2013.
- Colin Camerer and Teck Hua Ho. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.
- Colin F Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton university press, 2011.
- Gary Charness, John A List, Aldo Rustichini, Anya Samek, and Jeroen Van De Ven. Theory of mind among disadvantaged children: Evidence from a field experiment. *Journal of Economic Behavior & Organization*, 166:174–194, 2019.
- Caroline Claus and Craig Boutilier. The dynamics of reinforcement learning in cooperative multi-agent systems. *AAAI/IAAI*, 1998(746-752):2, 1998.
- 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.
- Pedro Dal Bó and Guillaume R Fréchet. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–29, 2011.
- Pedro Dal Bó and Guillaume R Fréchet. On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1):60–114, 2018.
- Pedro Dal Bó and Guillaume R Fréchet. Strategy choice in the infinitely repeated prisoner’s dilemma. *American Economic Review*, 109(11):3929–52, 2019.
- Pedro Dal Bó, Guillaume R Fréchet, and Jeongbin Kim. The determinants of efficient behavior in coordination games. Working Paper, 2020.

- Wim De Neys and Deborah Everaerts. Developmental trends in everyday conditional reasoning: The retrieval and inhibition interplay. *Journal of Experimental Child Psychology*, 100(4):252–263, 2008.
- Margaret Donaldson. Conservation - What is the question? *British Journal of Psychology*, 23(2):199, 1982.
- Ido Erev and Alvin E Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881, 1998.
- Eduardo Fe, David Gill, and Victoria L Prowse. Cognitive skills, strategic sophistication, and life outcomes. *Strategic Sophistication, and Life Outcomes (November 6, 2020)*, 2020.
- Aidan Feeney and Evan Heit. *Inductive reasoning: Experimental, developmental, and computational approaches*. Cambridge University Press, 2007.
- 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.
- Liane Gabora and Anne Russon. The evolution of intelligence. 2011.
- Usha Goswami. Inductive and deductive reasoning. 2011.
- Maria Gräfenhain, Tanya Behne, Malinda Carpenter, and Michael Tomasello. Young children’s understanding of joint commitments. *Developmental psychology*, 45(5):1430, 2009.
- 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.
- Brian K Hall. *Evolutionary developmental biology*. Springer Science & Business Media, 2012.
- Christos A Ioannou and Julian Romero. A generalized approach to belief learning in repeated games. *Games and Economic Behavior*, 87:178–203, 2014a.
- Christos A Ioannou and Julian Romero. Learning with repeated-game strategies. *Frontiers in neuroscience*, 8:212, 2014b.
- Steven O Kimbrough and Ming Lu. Simple reinforcement learning agents: Pareto beats nash in an algorithmic game theory study. *Information Systems and E-Business Management*, 3(1):1–19, 2005.

- John A List, Ragan Petrie, and Anya Samek. How experiments with children inform economics. Technical report, National Bureau of Economic Research, 2021.
- 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. *California Institute of Technology*, 2001.
- Brian McLoone and Rory Smead. The ontogeny and evolution of human collaboration. *Biology & Philosophy*, 29(4):559–576, 2014.
- Patricia H Miller. *Theories of developmental psychology*. Macmillan, 2002.
- Ryan Oprea. What makes a rule complex? *American economic review*, 110(12):3913–51, 2020.
- Josef Perner. *Understanding the representational mind*. The MIT Press, 1991.
- Jean Piaget. Intellectual evolution from adolescence to adulthood. *Human development*, 15(1):1–12, 1972.
- Jean Piaget, David Elkind, and Anita Tenzer. *Six psychological studies*. New York, NY: Vintage Books., 1967. ISBN 0394704622.
- Eva Rafetseder, Maria Schwitalla, and Josef Perner. Counterfactual reasoning: From childhood to adulthood. *Journal of experimental child psychology*, 114(3):389–404, 2013.
- Marjorie Rhodes, Daniel Brickman, and Susan A Gelman. Sample diversity and premise typicality in inductive reasoning: Evidence for developmental change. *Cognition*, 108(2):543–556, 2008.
- 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.
- Alvin E Roth and Ido Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, 8(1):164–212, 1995.
- 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.
- Barbora Siposova, Michael Tomasello, and Malinda Carpenter. Communicative eye contact signals a commitment to cooperate for young children. *Cognition*, 179:192–201, 2018.
- J Maynard Smith and George R Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.

- 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.
- Richard S Sutton and Andrew G Barto. Reinforcement learning: an introduction mit press. *Cambridge, MA*, 22447, 1998.
- Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(5): 675–691, 2005.
- Felix Warneken and Michael Tomasello. Helping and cooperation at 14 months of age. *Infancy*, 11(3):271–294, 2007.
- Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.
- Henry M Wellman and David Liu. Scaling of theory-of-mind tasks. *Child development*, 75(2): 523–541, 2004.
- Henry M Wellman, David Cross, and Julanne Watson. Meta-analysis of theory-of-mind development: the truth about false belief. *Child development*, 72(3):655–684, 2001.
- Emily Wyman, Hannes Rakoczy, and Mike Tomasello. Non-verbal communication enables children’s coordination in a “stag hunt” game. *European Journal of Developmental Psychology*, 10 (5):597–610, 2013.

Appendix A. 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.


SLIDE 1

You are RED




SLIDE 2

You are GREEN




SLIDE 3


You are RED



You are GREEN



SLIDE 4



Points


RED GREEN

5 3

3 5


1 1

1 1



SLIDE 6

You are BLUE



SLIDE 7

You are YELLOW




SLIDE 8


You are BLUE



You are YELLOW



SLIDE 9



Points

BLUE YELLOW

2 1

2 2

3 3

1 2

SLIDE 10

You are BLUE



You are BLUE



Figure 11: Slides to accompany instructions