

Why do children pass in the centipede game?

Cognitive limitations v. risk calculations *

Isabelle Brocas

*University of Southern California
IAST and CEPR*

Juan D. Carrillo

*University of Southern California
IAST and CEPR*

January 2025

Abstract

Children and adolescents from 8 to 16 years old play the centipede game in the laboratory, where non-equilibrium behavior (passing) can occur for two reasons: an inability to backward induct (cognitive limitation) or a decision to best respond to the empirical risk and take a measured chance (behavioral sophistication). We find that logical abilities develop gradually. While young participants are (as expected) least likely to perform backward induction, those who do, tend to over-estimate the ability of their peers to behave similarly. With age, participants gradually learn to think strategically and to best respond to their beliefs about others. Overall, the centipede game is an ideal test case for studying the development of abilities, as it disentangles the causes for passing in young children and in teenagers. Interestingly, shrewdness does not transform into earnings, and we document for the first time a game of strategy where payoffs monotonically *decrease* with age.

Keywords: developmental decision-making, centipede game, backward induction, risk-taking.

JEL Classification: C72, C90.

*We thank the Editor and referees for helpful comments, the members of the Los Angeles Behavioral Economics Laboratory (LABEL) for their insights, and Chris Crabbe for his outstanding programming skills. We are grateful to the staff of the Lycée International de Los Angeles (LILA) –in particular Emmanuelle Acker, Nordine Bouriche, Mathieu Mondange and Anneli Harvey– for their help and support running the experiment 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 (NSF) grant SES-2315770 and the French National Research Agency (ANR) grant ANR-17-EURE-0010 (Investissements d’Avenir program).

1 Introduction

Behavior in strategic games results from the combination of cognitive ability, inferences about others and preferences. Game theory provides normative predictions of behavior under selfish preferences and perfect foresight. However, we often observe departures from these predictions. The centipede game is an excellent example of a game where theoretical predictions yield low payoffs and are counterintuitive to most individuals. In a typical centipede game, two players face two escalating piles of money and take a finite number of turns deciding whether to stop or to pass. Backward induction prescribes stopping at each stage, which prevents players from taking advantage of the jointly increasing rewards.

Since its introduction by [Rosenthal \(1981\)](#), the centipede game has fascinated theorists, who describe it as a paradox of backward induction ([Aumann, 1992](#); [Reny, 1992](#); [Ben-Porath, 1997](#)). Authors have discussed several intuitive reasons why participants may decide to pass initially: limited cognition, social preferences, and inferences on mutual uncertainty. It is worth noting that once the first player has passed in their first opportunity, the second player cannot rely on backward induction to predict the rival's subsequent moves. This provides an argument for the second player to pass in their first opportunity which, in turn, offers a motivation for the first player to pass in the first place. More generally, under some assumptions, the theoretical literature has shown that passing may be consistent with certain definitions of rationality ([Reny, 1992](#)).

The experimental literature, initiated with [McKelvey and Palfrey \(1992\)](#), overwhelmingly sides with the intuitive expectation and against backward induction: few participants stop immediately... but they do not continue all the way to the end, either. As in most robust paradoxes, deviations are attenuated but not eliminated with experience, and they depend on specific elements, such as game length, payoff manipulations and number of players ([Krockow et al., 2016](#)). The persistent departures could indicate that people are not able to apply backward induction logic. This hypothesis has been investigated in two studies. [Palacios-Huerta and Volij \(2009\)](#) argue that experts (chess Grandmasters) are more likely to play the equilibrium than students. [Levitt et al. \(2011\)](#) did not share that result. Also, they found no relationship between ability to backward induct (measured in a challenging dominance solvable task) and the decision to stop immediately. Their result suggests that cognitive limitations are not the main driver of early passing. This is in contrast to other strategic settings where equilibrium behavior largely depends on cognitive ability ([Gill and Prowse, 2016](#); [Fe et al., 2022](#)).

Our study adopts an ontogenic approach. It leverages documented age-related changes

in both cognition and Theory-of-Mind (ToM) –the ability to read the rival’s intentions and form beliefs– to better disentangle the contribution of each of these two factors to choices in the centipede game. We recruit children and teenagers and study the change in behavior with age in the centipede game played over five rounds. Our hypotheses rely on known developmental changes in cognition and ToM abilities throughout development. Rounds 1, 3, 4 and 5 have ten stages whereas round 2 has only four stages. The short four-stage game serves as a diagnostic tool for the ability to backward induct. Contrasting behavior in rounds 1 and 2 allows us to disentangle between passing due to limited cognition and passing for strategic motives at different ages. Studying the evolution of behavior across long games allows us to determine how age and the gradual development of abilities affects learning over the course of the experiment.

The first and major finding is that the frequency of passing on the first attempt in both the long and the short games decreases monotonically with age, from elementary school through young adulthood. This trajectory is due to age-related differences in *both* the ability to perform backward induction and the ability to read the rival’s intentions. For elementary school children (ages 8 to 11), backward induction ability drives behavior. Children either find and play the equilibrium in both the long and short games (the minority), or they do not play it in either game (the majority). However, this correlation disappears starting in middle school. While the fraction of participants who stops immediately in the short game increases significantly with age, this choice does not predict their behavior in the long game. It suggests that, just like the experts in [Levitt et al. \(2011\)](#), passing for children at age 11 and above is often a deliberate decision for strategic considerations. As they grow, participants become more able to evaluate the empirical risk associated with passing and fine tune their stopping time.

The second (related) finding is the systematic decrease in payoffs with age (from elementary school to adulthood) both in the long and in the short game. To our knowledge, this is a first in a game of strategy. The astute reader may find it unsurprising, given that the higher payoffs of our younger participants are due to their larger deviations. One should notice, however, that such argument could also apply to other games where joint deviations increase the payoffs of players (e.g., prisoner’s dilemma), and yet it has never been documented previously. More generally, high ability is usually associated with a better empirical reading of the situation, which results in higher payoffs ([Proto et al., 2019, 2022](#)). Under such definition, one would expect more passing and larger gains as individuals get older, at least in the long game.

Our third finding relates to the choice dynamics. Participants of all ages stop earlier

as they play more rounds, which is rather natural given the asymmetry of incentives: ‘losing’ one round pushes subjects to preempt their rival in the next whereas ‘winning’ is unlikely to trigger deferral. However, it does not result in complete unraveling; by the end of the game, the majority of participants still stop between the second and fourth stage. Interestingly, the change is more pronounced for the younger participants. It results in choices and payoffs being very similar across ages by round 5.

There has been an increased recent interest in studying decision making by youngsters in experimental economics (see [Sutter et al. \(2019\)](#) and [List et al. \(2021\)](#) for excellent surveys). Games theoretical studies have revealed some interesting developmental trajectories that track the development of preferences ([Murnighan and Saxon, 1998](#); [Harbaugh and Krause, 2000](#)), reasoning and ToM ([Sher et al., 2014](#); [Czermak et al., 2016](#)) and largely correlates with measures of cognitive ability ([Fe et al., 2022](#)). Also, while children master the most basic false belief ToM tasks by age 5, the more general ToM ability continues to develop throughout adolescence ([Royzman et al., 2003](#)). Perhaps closest to this study is our recent research on games of strategy in children and adolescents. In [Brocas and Carrillo \(2018, 2020c\)](#), we consider a population of *very young* children (preschoolers to first graders, 4 to 7 years old) and study their ability to perform backward induction in very simple two-person games played in pairs or against the experimenter. In [Brocas and Carrillo \(2020b, 2021\)](#), we study children of a wide age range like in the present paper. We analyze the development of behavior in simultaneous two-person beauty contest and dominance solvable games, respectively, using novel, graphical versions adapted to that population. These games are similar to our four-stage centipede in that solving them requires a small number of steps of reasoning. The papers show a gradual improvement with age in the participants’ ability to anticipate future events and input them in current calculations. This ability typically plateaus around middle school. The novelty of the current project lies in using the long and short centipede games to distinguish between two sources of non-equilibrium behavior: limited cognition and behavioral sophistication. Additionally, it examines how individuals who grasp equilibrium theory develop the capacity to deviate from it for payoff-improving motives. Finally, it is worth emphasizing that, although the behavior of our participants changes with age, the two stylized comparative statics emphasized by [Ho and Su \(2013\)](#) regarding violations of backward induction in centipede games hold generally: in all our age groups, players violate backward induction less in a game with fewer stages and they stop earlier as they play the same game repeatedly.

The paper is organized as follows. In [section 2](#), we detail our population and discuss our design choices. In [section 3](#), we report the choices and payoffs in the first round of the

long and short versions of the game, with age as the main treatment factor. In [section 4](#), we investigate the evolution in behavior during the five rounds of the experiment. In [section 5](#) we present the results of a second experiment with a younger population. Concluding remarks and alleys for future research are collected in [section 6](#).

2 The game

2.1 Design and procedures

The paper studies the behavior of children and adolescents in the well-known centipede game, which was first introduced by [Rosenthal \(1981\)](#) and first studied in the laboratory by [McKelvey and Palfrey \(1992\)](#). Since working with young participants presents important methodological challenges, we have developed some methodological guidelines in [Brocas and Carrillo \(2020a\)](#) which we closely follow in this paper. Most notably, we develop a graphical, story-based version of the game that is appealing to children and adolescents.

Population. Experiment 1, the main treatment in the paper, was conducted with 315 school-age students from grades 3 to 10 at Lycée International de Los Angeles (LILA), a private school in Los Angeles.¹ We also included a control adult population (A) consisting of 72 college students from the University of Southern California (USC). A control group that follows an *identical* protocol is important to establish a behavioral benchmark ([Brocas and Carrillo, 2020a](#)), especially when procedures are slightly different than the standard ones, as it is the case here. [Table 1](#) summarizes the participants by grade and age.

Grade	LILA								USC
	3	4	5	6	7	8	9	10	A
Age	8-9	9-10	10-11	11-12	12-13	13-14	14-15	15-16	18-23
# indiv.	53	40	31	54	67	22	14	34	72

Table 1: Summary of participants by grade and population

We then conducted Experiment 2, a simplified version of the game involving 150 kindergarten, first, and second-grade students from the same school. For now, we will focus on the procedures and results of Experiment 1, with the specifics of Experiment 2 deferred to [section 5](#).

¹High schoolers from grades 11 and 12 did not participate in the study because they were taking or preparing for national exams during this period.

Procedures. We ran 27 and 6 sessions of Experiment 1 at LILA and USC with 8 to 14 participants each. Sessions at LILA were run in classrooms during school hours with individual partitions to preserve anonymity. Sessions at USC were run at Los Angeles Behavioral Economics Laboratory (LABEL) in the Department of Economics at USC. For each school-age session, we tried to have male and female participants from the same grade, but for logistic reasons we sometimes had to mix participants from two consecutive grades. Procedures were identical in all cases, except for payments as explained below.

The experiment consisted of two games programmed in ‘oTree’ (Chen et al., 2016) and implemented on touchscreen PC tablets through a wireless closed network. We started with a third-party dictator game. After a short break, we moved to the centipede game. Findings of the dictator game are discussed in a different article (Brocas and Carrillo, 2023). The two games are sufficiently different that we are not concerned about cross contamination. Nevertheless, to avoid any potential issues, we always performed the two games in the same order (dictator followed by centipede) with random and anonymous re-matching of subjects between the two games. Most importantly, we did not announce any result regarding the first game until the second game was finished.

Centipede Game. In developmental game theoretic studies, it is key to provide a simple, graphical interface and a story which is sound, accessible and relatable to young participants. This is all the more important when the age span is large. With this goal in mind, we developed our version of the Centipede Game, called “Going Down the Street”. Figure 1 presents a screenshot of the game from the perspective of player 2 at the second decision node. In our narrative, the blue and orange players walk together down a street and must decide in which house they enter. When arriving at a house, the player whose color matches that of the house decides for both whether to enter or continue to the next house. If they reach the last house, there is no possibility of continuing. Whenever they enter a house, players collect their color coded points and the game ends. (Appendix A provides the full set of instructions).

Rounds and stages. Participants were matched in pairs, assigned a role as player 1 or player 2 (blue or orange in our game) and played in Round 1 (R1) the ten-stage Centipede game described in Figure 1. Then they were randomly rematched, kept the same role and played in Round 2 (R2) a four-stage version of the same game, with only the first four houses. After that, they played Rounds 3, 4 and 5 (R3, R4, R5) of the original ten-stage version with alternating roles and random re-matching between rounds.²

²Thus, a player had one role in R1, R2 and R4 and the other role in R3 and R5.




Figure 1: Screenshot of “Going down the street” game

Contrasting behavior in R1 and R2 allows us to disentangle between different motives for early passing: while R2 serves as a diagnostic test for the ability to backward induct, sophisticated players may pass in R1 as a gamble to increase their expected payoff. Including R3, R4 and R5 helps us study whether and how quickly participants in the different age groups adjust their behavior.

Note that pairing individuals by age introduces a potential confound. Older participants may exhibit different behaviors due to an enhanced capacity for strategic thinking, a recognition that their peers might act differently, or a combination of both factors.

Remark 1. Our setting uses different parameters compared to the original experiment (McKelvey and Palfrey, 1992) and some of the subsequent literature.³ Ten (instead six) stages enriches the number of passing options. A constant (instead of increasing) payoff for the “loser” removes the possibility that both participants strictly win by passing which, in turn, reduces the role of social preferences as a driver of behavior.⁴ It also simplifies the calculations necessary to find the backward induction equilibrium. A linearly (instead of exponentially) increasing payoff for the “winner” ensures that the empirical variance in payments, to which children are particularly sensitive, is large but not massive. Our

³Researchers have studied experimentally many variants of the game, including length (McKelvey and Palfrey, 1992), payoff structure (Fey et al., 1996), incentives (Parco et al., 2002), number of players (Rapoport et al., 2003), game presentation (Nagel and Tang, 1998), and elicitation method (García-Pola et al., 2020a). Qualitative properties of the empirical behavior are usually robust to such modifications (see Krockow et al. (2016) for a survey).

⁴Even though some versions of social preferences à la Charness and Rabin (2002) could, in principle, rationalize passing, they are unlikely to be empirically relevant given the parameters we have chosen. It is also worth noting that García-Pola et al. (2020b) find no support for preference-based models in explaining passing behavior in a variety of centipede games.

modifications are expected to affect the results quantitatively but not qualitatively. We conjecture that the first feature increases the incentives to pass in the early stages for strategic considerations while the second and third feature decrease those incentives. A key feature is that all participants –including our control group– play the same game.

Remark 2. Several other games, besides our four-stage centipede, could be used to diagnose the participants’ ability to backward induct: (i) a multi-stage individual decision making task with intertwined moves by nature (Bone et al., 2009); (ii) a simplified version of the “game of 21” (Dufwenberg et al., 2010); (iii) a child-friendly graphical game that can be solved using iterated elimination of strictly dominated strategies (Brocas and Carrillo, 2021)); or (iv) the same ten-stage centipede game played against a robot with a pre-announced strategy. While each game type has its strengths, we found that a short version of the centipede game offers clear advantages: it closely resembles the long centipede game, is straightforward to understand, and, crucially, does not require new instructions. These features are particularly desirable in our setting because age significantly impacts the ability of participants to maintain concentration and handle extensive information from multiple games. At the same time, our game choice has an important limitation that we will discuss when we present our first result.

Payments and duration. Following the guidelines discussed in Brocas and Carrillo (2020a), we used different mediums of payment for different ages. This comes at added effort but we view it as a key choice. Indeed, an optimal incentive system must equalize as best as possible the *value of rewards* across individuals, not the rewards themselves. Money is usually the most adequate medium of payment precisely because it is valued most similarly by participants. However, this is not the case when age is a factor. Young children prefer objects to money, which they understand and appreciate, but it provides a lower immediate gratification. Participants were paid for all 5 rounds of the experiment. School-age students from grade 6 and above and control adults earned \$0.01 per point in each round paid immediately at the end of the experiment in cash (USC) or with an amazon giftcard (LILA, where cash transfers are not allowed). For elementary school students, we set up a shop with 20 to 25 pre-screened, age-appropriate toys and stationery that children find appealing (bracelets, erasers, figurines, die-cast cars, etc.).⁵ Before the experiment, we took the children to the shop, showed the toys they were playing for and

⁵While the age cut-off between toys and money is arbitrary, it coincides with school practices: only after 6th grade the school offers amazon giftcards as prizes for performance in intra-school activities (math puzzles, art shows, literature competitions, etc.).

explained their point prices. At the end of the experiment, subjects learned their total point earnings and were accompanied to the shop to exchange points for toys.⁶

The game studied in this paper lasted around 30 minutes. The entire experiment never exceeded one school period (50 minutes). Average monetary earnings in the Centipede Game were \$8.25 (LILA grade 6 and above) and \$7.75 (USC). Participants also earned \$2 to \$6 in the other game, and there was a \$5 show-up fee paid only to the adult population to correct for differences in the opportunity cost of time. For children in grades 3 to 5, point prices were calibrated in a way that all children obtained at least three toys, although there was large variance in number and type. We spent on average \$5 per child in toys, which is considerably higher than most experiments with elementary school children.

Clustering. To increase statistical power, some of our analysis clusters the school-age participants into three age-groups: grades 3-4-5 (**C1**, ages 8-11, 124 participants), grades 6-7 (**C2**, ages 11-13, 121 participants) and grades 8-9-10 (**C3**, ages 13-16, 70 participants).⁷ The control adult population consists of USC undergraduates (**C4**, ages 18-23, 72 participants) and it is included only for relevant comparisons. Regressions use either the age in months of the participants or dummies for age-group. 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. We use a p-value of 0.05 as the benchmark threshold for statistical significance, and we apply Holm p-value correction to account for multiple comparisons whenever appropriate.

2.2 Assumptions and Hypotheses

The following assumptions are implicitly built in our experimental design.

Assumption 1 *Participants who can (cannot) perform backward induction will (will not) stop immediately in R2.*

Assumption 2 *Participants who can perform backward induction will not necessarily stop immediately in R1.*

⁶The procedure emphasizes the importance of accumulating points while making the experience enjoyable. Children at this age are familiar with this method of accumulating points that are subsequently exchanged for rewards since it is commonly employed in arcade rooms and fairs.

⁷While it could be argued that 8th graders should be grouped with the other middle schoolers, we chose otherwise mainly to reach a similar sample size in all school-age groups. Results are similar (though statistical significance is affected) if we consider other grouping methods.

Assumption 3 *More participants can find the backward induction equilibrium in R2 than in R1.*

According to Assumption 1, we expect the four-stage game to be sufficiently simple so that no individual will consciously deviate from the equilibrium as a strategic attempt to increase their payoff. At the same time, we expect individuals who perform non-strategic probabilistic expectations to be lured by the magnitude of the payoffs in the latter stages.

According to Assumption 2, and based on a vast body of research on the centipede game in adults, we expect that a large fraction of forward-looking individuals who can perform backward induction will not stop immediately with the intention to maximize expected earnings. Naturally, one-half of these individuals will end up with a low payoff, hence the calculated risk.

Finally, while the backward induction logic is identical in both games, R2 requires fewer steps of reasoning than R1. Given that the literature strongly supports the idea that individuals are heterogeneous in the number of steps of reasoning they can perform, we conjecture in Assumption 3 that finding the equilibrium is more common in R2 than in R1. On the basis of these assumptions, we formulate the following hypotheses.

Hypothesis 1 *Stopping (and consequently earnings) changes monotonically with age:*

- (a) *In R1, older participants stop later and earn more than younger participants.*
- (b) *In R2, older participants stop earlier and earn less than younger participants.*

Hypothesis 2 *Within each age group, there is no correlation between early stopping in R1 and R2.*

Hypothesis 3 *Participants of all ages stop earlier as they play more rounds of the game.*

There are two possible reasons for passing: inability to backward induct (cognitive limitation) or decision to best respond to the empirical risk and take a measured chance (behavioral sophistication). Under Assumption 1, the second motive is not present in R2. Indeed, forward looking players will expect rivals to stop immediately if they do not stop themselves, so passing is driven by cognitive limitations. Since the cognitive ability to backward induct increases with age (Brocas and Carrillo, 2021), we expect less passing in R2 as individuals get older (**H1b**). By contrast, strategic passing in the initial stages of R1 is empirically optimal (Krockow et al., 2016). Under Assumption 2, we expect that sophisticated players will realize this. As a result, older participants who read better the

situation than their younger, more naïve counterparts will be more likely to pass early on in R1 (**H1a**). Overall, we expect average earnings to increase with age in R1 due to a better empirical read of the game by older participants and to decrease with age in R2, with the group of younger players financially benefitting from their lower cognitive ability.

H2 addresses more directly whether the strategy of early passing in the long game is due to behavioral sophistication or cognitive limitation. Since, the first effect is not present in the short game, the correlation between choices in R1 and R2 can address it. There are arguments for both hypotheses in the literature on adults (Palacios-Huerta and Volij, 2009; Levitt et al., 2011). When it comes to individuals who are developing their cognitive abilities, we expect under Assumption 2 that passing in R1 will often reflect empirically motivated strategic sophistication and therefore will *not* be followed by passing in R2. Thus, as summarized in **H2**, we do not expect a correlation in any of our age-groups between early passing in R1 and early passing in R2.

Finally, **H3** predicts a behavioral change towards the equilibrium in all ages as the game is played repeatedly. In this game, half the participants ‘lose’ in a given round, which provides incentives to stop earlier the next time. The other half of participants, who ‘win’ that round, need to form a counterfactual belief of what could have happened had they waited another turn. We expect that this asymmetric learning structure will imply earlier stopping over the course of the experiment.

3 Initial choices in long and short games

3.1 Choice and earnings as a function of age

We first study stopping behavior by age-group in R1 and R2. We compute the proportion of pairs who stopped at any given stage, with the understanding that only one of the individuals in the pair had a choice in each stage. For visual ease, we group the R1 stopping decision in four categories: stages 1-2 (first chance for each player), 3-4 (second chance), 5-6 (third chance) and 7-10 (fourth and fifth chance). The results are presented in Figure 2. We also include the average stopping stage at the top.

Round 1. In line with the existing literature on adults (and in contrast with the backward induction equilibrium prediction), we observe very significant levels of passing in all age groups of R1, with an average stopping between the fourth and fifth stage. Participants become more inclined to stop at the intermediate stages (3 to 6) with age rather than at the later stages (7 to 10), while stopping in the participants’ first opportunity (1 and 2) occurs between 14.8% and 25.8% of the time. The stopping distributions are




Figure 2: Stopping decision in Round 1 (left) and Round 2 (right)

significantly different between **C1** and **C4** and between **C3** and **C4** (chi-squared tests, $p = 0.021$ and $p = 0.014$ respectively). The average stopping time is also slightly decreasing (4.95 to 4.58), although differences are not statistically significant. Generally speaking, aggregate behavior in R1 is similar across age groups.

Round 2. The trend is much sharper in R2, where participants stop significantly earlier as they grow older. The average stopping stage moves from 2.11 in **C1** to 1.28 in **C4** (**C1** is different from all age groups, $p < 0.0003$, and **C2** is different from **C4**, $p = 0.013$). This reflects an increased understanding of the strategic forces at play: as they age, participants learn to anticipate how the game will unravel if they do not stop immediately and to apply backward induction reasoning.

An important caveat comes with this conclusion. In **C1**, the conditional probability of passing by player 2 in stage 2 is 0.61 which means that the best response of a risk-neutral player 1 to the aggregate empirical behavior is to pass in stage 1. In other words, our setting cannot rule out that, against Assumption 1, some shrewd elementary school children pass in stage 1 of R2 as a payoff maximizing strategy. On the other hand, if the larger levels of passing in stage 1 by the youngest children were due to their being highly sophisticated, then their peers would not pass in stage 2, thereby invalidating their reasons for the initial behavior. Overall, while we believe that the observed behavior in R2 reflects an increased level of sophistication with age, it is key to acknowledge this potential confound.⁸ In section 5, we compare the behavior of these individuals with that of even younger players (Kindergartners, first and second graders) in the short centipede.

As most behavioral theories would predict, immediate stopping is much less likely in

⁸We are thankful to a reviewer for stressing this very important point. In retrospect, a cleaner diagnostic of backward induction test would have been a three-stage centipede game.

R1 than in R2, and this difference increases with age: 8% v. 21% in **C1**, 5% v. 36% in **C2**, 11% v. 60% in **C3**, and 6% v. 72% in **C4**. The behavior of adults in R2 is consistent with [García-Pola et al. \(2020b\)](#) and confirms that social preferences is an unlikely explanation for early passing, at least under the parameters adopted in our version of the game.

Although the descriptive analysis is revealing, it is important to control for other factors, including role and gender. To do so, we focus on the school age population and conduct OLS regressions of the number of stages before stopping (columns 1 and 2) as a function of *Age* (in months), gender (*Male* = 1), the interaction between age and gender, and the participant’s role (*Player2* = 1). To construct our independent variable, we group together two consecutive stages (1-2, 3-4, etc.), to correct for the fact that players 1 and 2 can only stop at the odd and even stages, respectively. We also run OLS regressions of the payoff of the participant who stops as a function of the same variables (columns 3 and 4), this time without grouping two stages together.⁹ We do not include our control undergraduate students in these regressions to not bias the age coefficients of the regression in their direction. The results are presented in [Table 2](#).

	Stages before stop		Payoff of winner	
	R1	R2	R1	R2
<i>Age</i>	-0.012*	-0.005**	-1.651*	-0.642**
	(0.006)	(0.001)	(0.748)	(0.200)
<i>Male</i>	-2.036*	-0.185	-285.1*	-25.9
	(1.015)	(0.278)	(142.1)	(38.9)
<i>Age</i> × <i>Male</i>	0.009	0.001	1.290	0.071
	(0.007)	(0.002)	(0.968)	(0.266)
<i>Player2</i>	-0.469*	-0.292***	4.320	29.2***
	(0.180)	(0.050)	(25.2)	(6.98)
<i>const.</i>	5.015***	2.015***	662.1***	242.0***
	(0.789)	(0.214)	(110.5)	(29.9)
Adj. R ²	0.138	0.250	0.092	0.206
# obs.	158	158	158	158

* p < 0.05, ** p < 0.01, *** p < 0.001.

Table 2: OLS of stopping stage and payoffs in R1 and R2 in school-age population

The stopping stage significantly decreases with age in R1 ($p = 0.029$), and the effect is stronger in R2 ($p = 0.002$). Males stop earlier than females in R1 but not in R2. It suggests no gender differences in the ability to backward induct (in R2) while gender-related motives are at play in R1. Player 2 chooses less frequently to pass than player 1.

⁹In our formulation, the payoff of the participants who do not stop before their rival is constant and equal to 50 independently of the stopping stage, so we can ignore them for our purposes.

This, however, should be interpreted with caution, since it is conditional on being given the same opportunities (which is endogenous to the model). Finally, results are similar (and significance levels of the main variables are identical) if we replace *Age* with *Grade* to control for age heterogeneity within the same class cohort (tables available upon request).

In terms of payoff, this behavioral trajectory results in monotonically decreasing earnings. Notice that the payoff Π of the participant who stops is a linear transformation of the stopping stage t : $\Pi = 30 + 70t$. It is therefore natural that regressions on payoffs and on stopping stage yield very similar findings. Significance levels are not identical only because we group stopping stages in pairs but we consider payoffs separately. In turn, this methodology allows us to unveil an interesting effect of role: player 2 stops, on average, at an earlier *opportunity* (columns 1 and 2) but at a later overall *stage* (columns 3 and 4), hence obtaining a higher payoff conditional on being the player who stops.

In summary, our results support **H1b**, which we interpret as a natural outcome of the developmental trajectory in cognitive abilities (with the noted caveat). However, we find no support for **H1a**: older participants do not stop later in the long game and thus do not achieve the anticipated empirical gains associated with multiple passing.

3.2 Choice across the long and short game

To further investigate the choices of school-age individuals across rounds, we run a Probit regression of the first player’s choice in the first stage of R2 ($\text{Stop}(\text{R2}) = 1$) as a function of the number of stages in which that player chose to pass in R1 (*PassR1*). We use **C2** as the benchmark age group, and include dummies for the other age-groups (*C1* and *C3*) as well as interaction terms, and a dummy for gender. This exercise, however, is imperfect. Indeed, our data is censored since an individual can only make a choice in stage t if the partner passed in stage $t - 1$. Results are reported in [Table 3](#) and sheds some interesting light on the relationship between cognition and behavior.

Repeated passing in R1 is a predictor of passing in the first stage of R2 only in our youngest age group (**C1**). Whenever our elementary school participants realize that the unraveling logic of the game dictates immediate stopping, they apply the argument *indiscriminately in both rounds*. In other words, if they manage to find the equilibrium strategy, they play it, without considering whether their partner will do the same. Foregoing these inferences is payoff-detrimental, particularly in R1. By contrast, our middle- and high-schoolers are more discerning of the situation. For these age groups, passing in the long game is not an indication of limited cognition because it does not predict their

	Stop(R2)
<i>PassR1</i>	-0.256 (.164)
<i>C1</i>	0.756 (.643)
<i>PassR1</i> × <i>C1</i>	-0.998* (.410)
<i>C3</i>	0.490 (.626)
<i>PassR1</i> × <i>C3</i>	-0.035 (.266)
<i>Male</i>	0.396 (.247)
<i>const.</i>	0.105 (.419)
AIC	170.2
# obs.	158

* p < 0.05, ** p < 0.01, *** p < 0.001.

Table 3: Individual Probit of stopping in R2 as a function of choice in R1

behavior in the short game.¹⁰

Overall, the hypothesis that individuals who pass in the long game do so for strategic payoff considerations and therefore such behavior will not be predictive of their choice in the short game (**H2**) is verified for our middle- and high-schoolers but not for our elementary school participants. In this latter group, finding and playing the equilibrium are two intimately related concepts: young children who understand backward induction will typically not realize that others may not be willing or able to apply the Subgame Perfect Equilibrium logic (in Appendix B1, we further study the changes with age in the initial choice of player 1 in R1 and R2). These results, however, should be interpreted with caution. Ideally, a counterbalanced treatment should also be conducted, in which participants first play the four-stage version followed by the ten-stage version. We were unable to implement that variant due to the limited sample size available for this project.

¹⁰As robustness checks, we show that the average stopping for participants in **C1** is significantly lower among individuals who stop immediately in R2 than among those who do not, whereas no differences exist in **C2** and **C3**. Also, adding **C4** to the regression in Table 3 and using a different age group as benchmark does not change the main result that stopping early in R1 predicts stopping in the first stage of R2 only for **C1** (data omitted for brevity but available upon request).

3.3 Summary

Our analysis reveals a clear developmental trajectory. The contrast between behavior in R1 and R2 reveals an interplay between cognitive ability and Theory-of-Mind (ToM), the ability to read the rival’s intentions and form beliefs which develops throughout childhood and adolescence (Royzman et al., 2003). In that respect, we obtain three key findings.

First, many participants in **C1** are lured by the high rewards in late stages and do not stop in either game. Still, 21% of them behave consistently with backward induction logic in R2. Interestingly, these individuals do not anticipate that the majority of their peers will experiment in R1, and *over-apply* their skill. Also, the fact that the least cognitively able participants stop very late suggests that their reward seeking attitude is likely due to a lack of a proper cost-benefit trade-off. This is consistent with the known tendency of children until age 11 to focus on salient features (Miller, 2002).

Second, while some participants in **C2** and **C3** are still unsuccessful at backward inducting in R2 (and suffer losses from it), the behavior of those who succeed does not drive their choice in R1. They have acquired enough ToM abilities to assess empirical risk with some accuracy. This behavior is analogous to Levitt et al. (2011) who show that immediate stopping by chess Grandmasters in the centipede game is not related to their ability to apply backward induction reasoning in other diagnostic tasks.

Third, the behavioral trajectory described above implies that earnings monotonically decrease with age. While the result is expected in R2 (where the game configuration induces forward-looking players to stop immediately), it is surprising in R1 since it prevents older participants to reach solutions that are collectively payoff-improving.

4 Dynamic learning in long games

We next analyze behavior in the long games R1, R3, R4 and R5. We identify systematic differences (or similarities) between age groups and study how experience affects behavior.

4.1 Unraveling within age groups

Figure 3 describes the evolution of the stopping strategy in each age group separately, with the average stopping stage reported at the top.

There is a tendency in all age groups to stop earlier as the experiment progresses. The largest average change occurs between R1 and R3 for all age groups except **C3** (where changes are smoother). By the fifth round, participants stop on average one to two rounds




Figure 3: Evolution of the stopping strategy by age group

earlier than they did in the first round. A t-test reveals that “diff₁₋₅”, the difference between average stopping stage at the beginning and at the end of the experiment (R1 v. R5), is large and highly significant in all age groups: diff₁₋₅ is 1.80 in **C1** ($p < 0.001$), 1.26 in **C2** ($p < 0.001$), 0.94 in **C3** ($p = 0.010$) and 1.47 in **C4** ($p < 0.001$).

This evolution can be best studied with OLS regressions of the stopping stage in each age group as a function of round and gender (R3 being the default round). The results are presented in [Table 4](#).

	Stages before stop			
	C1	C2	C3	C4
<i>R1</i>	1.241*** (0.365)	0.980** (0.295)	0.429 (0.421)	0.683* (0.282)
<i>R4</i>	-0.739* (0.286)	-0.330 (0.240)	-0.581 (0.324)	-0.393 (0.252)
<i>R5</i>	-0.641* (0.307)	-0.283 (0.262)	-0.533 (0.356)	-0.848** (0.324)
<i>Male</i>	-1.543*** (0.395)	-0.231 (0.319)	-0.324 (0.430)	-0.451 (0.362)
<i>const.</i>	4.682*** (0.375)	4.164 (0.249)	4.213*** (0.353)	4.159*** (0.320)
Adj. R ²	0.196	0.072	0.026	0.098
# obs.	248	244	140	140

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

Table 4: OLS Regressions of evolution in stopping stage in each age group

Again, we notice that participants stop earlier in later rounds. The change is most pronounced in **C1**. In **C2**, the change occurs early in the experiment (between R1 and R3). By contrast, changes are not significant in **C3**. An immediate consequence of this behavior is the corresponding decrease in payoffs for participants in all age groups as the

game progresses. We also observe that the gender effect noted in Table 2 –whereby males stop earlier than females– persists but it is significant only in C1.

The dynamics across rounds of our population is consistent with existing findings in adults (McKelvey and Palfrey (1992), Fey et al. (1996), and others). It also connects to earlier defection with experience in repeated prisoner’s dilemma (Embrey et al., 2018). The result is natural. If my rival stops before me, I have incentives to stop earlier the next time. By contrast, if I stop before my rival, I do not know what would have happened had I continued, so I am not necessarily more likely to postpone stopping in future rounds. This asymmetry of incentives about future behavior between “winners” and “losers” results in progressive, though not full, unraveling.

4.2 Age effects

To assess the determinants of stopping within round, we present in Table 5 the same OLS regressions as in Table 2 for Rounds 3, 4 and 5.

	Stages before stop			Payoff of winner		
	R3	R4	R5	R3	R4	R5
<i>Age</i>	-0.005 (0.006)	-0.001 (0.004)	-0.003 (0.004)	-0.716 (0.620)	-0.067 (0.545)	-0.367 (0.528)
<i>Male</i>	-2.166* (0.839)	-1.592* (0.777)	-1.784* (0.762)	-303.2* (117.5)	-222.9* (108.7)	-249.7* (106.6)
<i>Age</i> × <i>Male</i>	0.013* (0.006)	0.009 (0.005)	0.010* (0.005)	1.814* (0.805)	1.242 (0.745)	1.463* (0.731)
<i>Player2</i>	0.804*** (0.150)	-0.348* (0.143)	0.334* (0.139)	42.6* (21.0)	21.2 (20.0)	-23.3 (19.5)
<i>const.</i>	2.705*** (0.647)	2.374*** (0.568)	2.353*** (0.561)	408.7*** (90.5)	292.4*** (79.5)	359.4*** (78.5)
Adj. R ²	0.184	0.085	0.067	0.061	0.042	0.043
# obs.	158	158	158	158	158	158

* p < 0.05, ** p < 0.01, *** p < 0.001.

Table 5: OLS of stopping stage and payoffs in R3, R4 and R5 in school-age population

Consistent with the results in Table 4, the OLS regressions of Table 5 show that the effect of age on delaying stopping has dissipated by R3, and remains non-significant in R4 and R5. Consequently, age effects on payoffs also vanish. We also notice that the gender effect observed in Table 2 is still present but it is modulated by age: males delay less the time to stop but the effect is stronger in the youngest ones (significantly in R3 and R5). Again, the effects are very similar if we replace *Age* with *Grade*.

Taken together, **H3** is supported by our analysis: with experience participants stop earlier. The change is stronger in the younger population and, by round 5, there is convergence in behavior, with individuals of all ages stopping early but not immediately. This decrease in “cooperative passing” has detrimental aggregate payoff consequences. In Appendix B2, we further study the behavior of participants in R3, R4 and R5 as a function of their initial choice in R1 and R2.

4.3 Reaction to empirical risk in the long games

As stages progress, participants should realize that passing has both an increased risk (likelihood that the rival stops next) and a higher opportunity cost (difference between current payoff of stopping and payoff if the rival stops). We should therefore observe that participants are each time more likely to stop conditional on reaching a given stage. This pattern would indicate that participants both form logical beliefs about their partners and they themselves reason logically.

To test this hypothesis, we determine for each age group and each stage t , the stopping hazard rate h_t . This is the probability that a participant stops in stage t , p_t , given that such stage has been reached. Formally, $h_t = \frac{p_t}{\sum_{i=t}^T p_i} \in [0, 1]$. Even minimally strategic individuals should exhibit an increasing hazard rate ($h_{t+1} > h_t$). Figure 4 presents the data by age group. Since the number of observations decreases significantly as we move to later stages, we present hazard rates only for stages 1 through 6. We also pool together Rounds 1, 3, 4 and 5 to increase statistical power, even though we realize that behavior changes over rounds, and observations across rounds are collected from the same individuals.

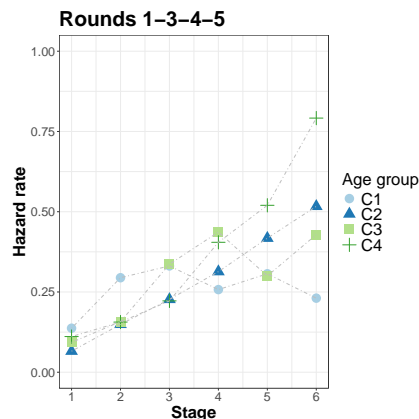


Figure 4: Stopping hazard rate in long games by age group

As in previous experiments, participants in our control group **C4** are increasingly likely to stop as stages advance. The hazard rate’s time series estimated slope is steep ($\hat{s} = 0.124$) and highly significant ($p = 0.008$). Hazard rates also increase steadily in **C2** and **C3**, although the slope is statistically significant in the former ($\hat{s} = 0.089$, $p = 0.009$) but not in the latter ($\hat{s} = 0.067$, $p = 0.132$). **C1**, on the other hand, exhibit a constant hazard rate ($\hat{s} = 0.004$, $p = 0.999$). Behavior of **C2**, **C3** and **C4** conforms to players being strategic. By contrast, behavior in **C1** cannot be reconciled with a theory where participants are minimally strategic, and trade-off the expected costs and benefits of passing at each stage (indeed, the increasing opportunity cost of passing implies that even if players expect rivals to use a constant hazard rate, it is still an optimal best response to use an increasing hazard rate).¹¹

The result reinforces the idea that passing in **C1** is primarily attributable to cognitive limitations (in both the long and short games). This is qualitatively different from the behavior in other age groups, where passing in the long games can be interpreted as a calculated risk-taking strategy.

4.4 Individual analysis

The dynamics of play of our participants are shaped by their cognitive abilities, but also by their assigned roles, which expose them to distinct opportunities to act and learn from feedback. To gain a clearer understanding of individual play and learning dynamics, we separate participants in two groups as a function of the *first* role they take in the game (R1) —blue or orange. Within each group, we examine the rounds in which they act as player 1 (R1 and R4 for blue, R3 and R5 for orange) and categorize participants into four types: (i) WW are players who stop first (and therefore “win”) both times; (ii) WL are risk-taking players, who stop first the first time but wait too long (and “lose”) the second time; (iii) LW are learners, who lose the first time but win the second; finally, (iv) LL are players who pass excessively and lose both rounds. Tables 6 and 7 present descriptive statistics of these four types separated by their role. We highlight in bold the rounds where they choose first.

WL stop significantly later than WW ($p = 0.003$) in R1 and still win, which predisposes them to pass more subsequently. Their success despite more passing slows their learning, preventing them from adapting their play quickly. Conversely, players who initially lose

¹¹The same results hold (though significance is, by construction, weaker) if we look only at the first four stages. Slopes and p-values become: 0.038 ($p = 0.734$) in **C1**, 0.082 ($p = 0.089$) in **C2**, 0.118 ($p = 0.089$) in **C3** and 0.082 ($p = 0.089$) in **C4**.

	Stopping R1	Stopping R3	Stopping R4	Stopping R5	Payoff
WW	4.1	3.6	3.5	4.2	1057 (64.2)
WL	5.6	4.3	–	3.5	898 (38.0)
LW	–	3.2	3.2	3.5	682 (31.6)
LL	–	4.0	–	3.2	527 (34.6)

Table 6: Blue role. Average stopping round by type and conditional on winning, and total payoff by type (standard errors in parenthesis).

tend to adjust more their strategy. By construction, being the player who stops results in higher overall payoffs (all p-values < 0.025). Also, WL achieve higher earnings than LW ($p = 0.002$). Although both groups miss one opportunity, the former exploit the higher aggregate levels of passing in R1 than in R4.

	Stopping R1	Stopping R3	Stopping R4	Stopping R5	Payoff
WW	3.9	3.5	3	3.2	962 (47.7)
WL	5.0	4.8	4	–	938 (46.4)
LW	4.9	–	3.4	3.2	748 (43.0)
LL	5.9	–	3.9	–	631 (42.5)

Table 7: Orange role. Average stopping round by type and conditional on winning, and total payoff by type (standard errors in parenthesis).

A similar pattern emerges for the orange role. WL stop significantly later than WW in R3 ($p = 0.001$). This delay aligns with the one observed in R1 ($p = 0.04$). LL exhibit an even more pronounced version of this behavior. Meanwhile, LW tend to pass more frequently than WW in R1 ($p = 0.07$), but they eventually adapt their behavior by R5. Once again, earnings depend on the number of winning rounds, and WL secure a higher payoff than LW ($p = 0.004$) for the same reasons as before.

The key takeaway from this exercise is that, for both blue and orange roles, there is dynamic adaption. Winning after stopping late in early rounds leads to slower learning and eventual losses in later rounds. Conversely, losing by passing excessively in early rounds foster earlier stopping and improved payoffs later on.

4.5 Summary

Participants stop earlier as the experiment progresses and, by round 5, behavior converges across age groups. Since choice in the first round is more noticeably different in **C1**,

the changes are more abrupt in that group. Still, we observe that the underpinnings of behavior continue to differ. While age groups **C2**, **C3** and **C4** respond to empirical risk in a logical fashion by decreasing the conditional probability of passing as the game progresses, young participants in **C1** do not account for the increased incentives to stop within a round.

5 Experiment 2: the behavior of younger participants

While we have argued that third, fourth and fifth graders (**C1**) are less strategic than their older peers, their behavior in the short game could potentially be explained by a highly sophisticated ability to anticipate and react to the behavior of others. Such argument seems unlikely for a majority of them given the other aspects of their choice (e.g., the group exhibits a constant hazard rate in the long games) but it cannot be ruled out.

This raises the following question: at which age are children strategic enough to solve the four-stage centipede game? This game is significantly more complex than the highly simplified hit-N and dominance games studied for children of this same age in [Brocas and Carrillo \(2018, 2020c\)](#). It is of similar difficulty to the bi-dimensional, two-person beauty contest in [Brocas and Carrillo \(2020b\)](#), which kindergartners have significant problems solving.

Expanding the trajectory of behavior in the four-stage game to younger players is instructive as it may help us understand whether the behavior of their slightly older peers is likely to be the result of a naive or a sophisticated strategy. With this objective in mind, we ran a second experiment with 150 additional children: 55 kindergartners (ages 5-6), 49 first graders (ages 6-7) and 46 second graders (ages 7-8) in the same school.

Participants were subject to the same experimental conditions as **C1** with two exceptions: (i) they played four rounds of the four-stage centipede game with alternating roles and random rematching between rounds; (ii) to avoid “big” numbers, the payoffs were modified to (2,1), (1,3), (4,1), (1,5).¹² We did not run the ten-stage version since we believe it is excessively difficult given their early developing cognitive abilities.

[Figure 5](#) describes the evolution of the stopping strategy over the four rounds of the experiment in grades K, 1 and 2, with the average stopping stage reported at the top. To facilitate comparisons, we also report the behavior in R2 of Experiment 1 for grades 3, 4 and 5. [Table 8](#) reports an OLS regression of the stopping stage with age (in months),

¹²Conversion rates were adapted accordingly. Note that, in order to keep the same multiplier, we should have used 3.4, 4.8 and 6.2 instead of 3, 4, and 5 but this was obviously not a reasonable option.

gender and round (1, 2, 3 or 4) as the explanatory variables.




Figure 5: Evolution of stopping strategy in the four-stage centipede for grades K, 1 and 2 (Experiment 2, left) and behavior in R2 for grades 3, 4 and 5 (Experiment 1, right)

	Stages before stop	Payoff of winner
<i>Age</i>	-0.019** (0.006)	-0.019** (0.006)
<i>Male</i>	0.208 (0.187)	0.208 (0.187)
<i>Round</i>	-0.141** (0.001)	-0.141** (0.001)
<i>const.</i>	4.047*** (0.000)	5.047*** (0.000)
Adj. R ²	0.057	0.057
# obs.	300	300

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; ° $p < 0.1$

Table 8: OLS of stopping stage and payoff to winner in Experiment 2

From the descriptive and regression analysis, we obtain the following findings. (i) Second graders tend to stop earlier than kindergartners and first graders. This evolution matches the behavior observed in other games (Brocas and Carrillo, 2020b, 2021), where the ability to think strategically exhibits a significant jump between first and second grade. Also, a fraction of participants in K and first grade miss every strategic aspect of the game, as reflected in their choice to pass in stage 3. This behavior disappears almost entirely by second grade. (ii) In all grades, participants stop earlier as the experiment progresses. The change is largest between rounds 1 and 3 and, by round 4, behavior seems to have stabilized. Although, we would need more rounds to determine the long run strategy of players, there is no evidence that players would eventually converge to immediate stopping. Also, compared to the only instance where third, fourth and fifth graders play the short

centipede (R2), second graders stop around the same time in round 1 (t-test, $p > 0.1$) and earlier in rounds 3 and 4 (t-tests, $p < 0.001$). (iii) With the exception of second graders in round 3, in all other cases the best response to the empirical behavior of rivals is to pass in stage 1. This matches the findings of **C1** in Experiment 1. However, it is hard to conclude that behavioral sophistication is the reason for passing. Indeed, in K and first grade, many children pass as player 1 in stage 3. Also, in all grades, many children pass as player 2 in stage 2 of rounds 2 and 4, even though these were the individuals who acted as player 1 and passed in stage 1 of round 1. Finally, like in Experiment 1, the payoff of the winner is a linear transformation of the stopping stage ($\Pi = 1 + t$). Thus groups that stop later obtain higher aggregate payoffs.

The behavior observed in Experiment 2 naturally sets the stage for the behavior displayed by elementary school children in Experiment 1. It demonstrates that lower levels of rationality and a limited ability to use backward induction, as exhibited by the youngest children, result in higher collective gains in the short game.

6 Conclusion

Behavior in strategic games is driven by cognitive ability, theory of mind, individual motivation and incentives. In the centipede game, each decision to pass is a gamble, and participants might be affected by these prospects. While equilibrium theory prescribes stopping at each stage, numerous experiments show that adults delay stopping. With our payoffs, passing can occur for two main reasons: inability to perform logical inferences or decision to best respond to the empirical risk and take a measured chance. Our study sheds light on the importance of these motives and on how they develop with age.

While young participants are least likely to act logically, those who do, tend to overestimate the ability of their peers to behave similarly. As a consequence, they apply the logic blindly. With age, participants learn to anticipate what others may do and best respond to their beliefs. Starting in middle school, students who reason logically know that the unraveling argument should not be applied blindly. The behavior of middle- and high-schoolers is in line with the literature on the centipede that documents deviations from backward inductions even after experience. It is also consistent with [Levitt et al. \(2011\)](#) who show no correlation between early stopping in the centipede game and ability to backward induct in more challenging, dominance-solvable games.

The intuitive way to approach the game for participants with limited cognition is as a series of gambles leading to increasing rewards, which is what many young participants

do. The inclination of young children to take high risks has been widely documented in the literature (Paulsen et al., 2012; Sutter et al., 2019) and it likely results from their difficulty to simultaneously process rewards and probabilities (Brocas et al., 2019). This leads them to pay disproportionately more attention to rewards. In our case, it makes them not only willing to take high risks, but also to stop at each stage with the same conditional probability. Overall, passing in young children results from cognitive limitations, while it follows a calculated empirical risk in older participants.

More generally, the underlying mechanisms that lead to choice not only are complex but they also change with age. Delayed stopping and constant hazard rate by many young participants is only consistent with salience effects and limited cognition. By contrast, delayed stopping and increasing hazard rate by teenagers more likely reflects impulsivity or underdeveloped risk avoidance mechanisms. Also, the studies referenced above have typically found that female are more risk averse than males, which would predict that females would stop earlier than males in our game. This is not supported by our data. It may be due to the fact that the task does not explicitly describes probabilities and that behavior in our strategic task relates to gender-related differences in competitiveness (Niederle and Vesterlund, 2011). These observations taken together suggest that differences in stopping are driven more by differences in cognition and ToM abilities than by differences in intrinsic risk attitudes.

The study also documents large heterogeneity within age: 21% of very young participants understand the unraveling logic of the short game while 28% of highly educated adults do not. This is consistent with results obtained in other paradigms involving iterative reasoning (Brocas and Carrillo, 2021). It is also interesting to notice that cognitively equipped young children are not able to *not apply* their skill, while their older peers are, which underscores the importance of ToM.

To our knowledge, this is the first experimental study where earnings monotonically decrease over a large age span (8 to 16 years old, and up to 23 years old if we include the control group). Young children obtain, on average, larger earnings in the short game because they are lured by the high rewards in late stages, but they also obtain higher payoffs in the long game. Larger deviations by less cognitively developed participants may seem unsurprising. And yet, we expected that more sophisticated individuals would be able to replicate the behavior of their less sophisticated peers and obtain (at least) their same expected payoff in the long game. This is what typically happens in developmental studies, but not what we find in this paper.

Experience heavily disciplines behavior and, by the fifth round, the distribution of

choices in very similar in all grades. The fact that young children rapidly adapt their decisions indicates that they use feedback as an input for their next choice, and apply inductive logic to what they observe. It would be interesting to refine those ideas and see whether children learn to respond to empirical risk, to backward induct, or both. It would also be interesting to link individual learning to feedback, as we expect children to learn differently if their first rival stops early or takes some chances.

As noted earlier, the centipede game is largely dependent on beliefs. It is therefore an ideal game to combine participants who differ in cognitive abilities but also in abilities to read other players' minds. An interesting alley for future research is to mix children from different grades. This exercise would be instructive for two reasons. First, this approach would allow us to distinguish choices driven by an individual's ability to strategize within the game from choices influenced by beliefs about the rival's choice. Second, it would enable us to assess young children's capacity to behave more maturely in the presence of older peers and older children's willingness to exploit younger rivals who may be presumed less cognitively capable.

References

- Robert J Aumann. Irrationality in game theory. *Economic analysis of markets and games*, pages 214–227, 1992.
- Elchanan Ben-Porath. Rationality, nash equilibrium and backwards induction in perfect-information games. *The Review of Economic Studies*, 64(1):23–46, 1997.
- John Bone, John D Hey, and John Suckling. Do people plan? *Experimental Economics*, 12:12–25, 2009.
- Isabelle Brocas and Juan D Carrillo. The determinants of strategic thinking in preschool children. *Plos one*, 13(5):e0195456, 2018.
- Isabelle Brocas and Juan D Carrillo. Introduction to special issue “understanding cognition and decision-making by children.” studying decision-making in children: challenges and opportunities. *Journal of Economic Behavior & Organization*, 179:777–783, 2020a.
- Isabelle Brocas and Juan D Carrillo. The evolution of choice and learning in the two-person beauty contest game from kindergarten to adulthood. *Games and Economic Behavior*, 120:132–143, 2020b.
- Isabelle Brocas and Juan D Carrillo. Iterative dominance in young children: Experimental evidence in simple two-person games. *Journal of Economic Behavior & Organization*, 179:623–637, 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. Developmental meritocracy. USC Working Paper, 2023.
- Isabelle Brocas, Juan D Carrillo, T Dalton Combs, and Niree Kodaverdian. The development of consistent decision-making across economic domains. *Games and Economic Behavior*, 116:217–240, 2019.
- Gary Charness and Matthew Rabin. Understanding social preferences with simple tests. *The quarterly journal of economics*, 117(3):817–869, 2002.
- Daniel L Chen, Martin Schonger, and Chris Wickens. otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.
- Simon Czermak, Francesco Feri, Daniela Glätzle-Rützler, and Matthias Sutter. How strategic are children and adolescents? experimental evidence from normal-form games. *Journal of Economic Behavior & Organization*, 2016.
- Martin Dufwenberg, Ramya Sundaram, and David J Butler. Epiphany in the game of 21. *Journal of Economic Behavior & Organization*, 75(2):132–143, 2010.

- Matthew Embrey, Guillaume R Fréchet, and Sevgi Yuksel. Cooperation in the finitely repeated prisoner's dilemma. *The Quarterly Journal of Economics*, 133(1):509–551, 2018.
- Eduardo Fe, David Gill, and Victoria Prowse. Cognitive skills, strategic sophistication, and life outcomes. *Journal of Political Economy*, 130(10):2643–2704, 2022.
- Mark Fey, Richard D McKelvey, and Thomas R Palfrey. An experimental study of constant-sum centipede games. *International Journal of Game Theory*, 25(3):269–287, 1996.
- Bernardo García-Pola, Nagore Iriberri, and Jaromír Kovářík. Hot versus cold behavior in centipede games. *Journal of the Economic Science Association*, 6(2):226–238, 2020a.
- Bernardo García-Pola, Nagore Iriberri, and Jaromír Kovářík. Non-equilibrium play in centipede games. *Games and Economic Behavior*, 120:391–433, 2020b.
- David Gill and Victoria Prowse. Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, 124(6):1619–1676, 2016.
- William T Harbaugh and Kate Krause. Children's altruism in public good and dictator experiments. *Economic Inquiry*, 38(1):95–109, 2000.
- Teck-Hua Ho and Xuanming Su. A dynamic level-k model in sequential games. *Management Science*, 59(2):452–469, 2013.
- Eva M Krockow, Andrew M Colman, and Briony D Pulford. Cooperation in repeated interactions: a systematic review of centipede game experiments, 1992–2016. *European Review of Social Psychology*, 27(1):231–282, 2016.
- Steven D Levitt, John A List, and Sally E Sadoff. Checkmate: Exploring backward induction among chess players. *American Economic Review*, 101(2):975–90, 2011.
- John A List, Ragan Petrie, and Anya Samek. How experiments with children inform economics. Technical report, National Bureau of Economic Research, 2021.
- Richard D McKelvey and Thomas R Palfrey. An experimental study of the centipede game. *Econometrica: Journal of the Econometric Society*, pages 803–836, 1992.
- Patricia H Miller. *Theories of developmental psychology*. Macmillan, 2002.
- J Keith Murnighan and Michael Scott Saxon. Ultimatum bargaining by children and adults. *Journal of Economic Psychology*, 19(4):415–445, 1998.
- Rosemarie Nagel and Fang Fang Tang. Experimental results on the centipede game in normal form: an investigation on learning. *Journal of Mathematical psychology*, 42(2-3): 356–384, 1998.
- Muriel Niederle and Lise Vesterlund. Gender and competition. *Annu. Rev. Econ.*, 3(1): 601–630, 2011.

- Ignacio Palacios-Huerta and Oscar Volij. Field centipedes. *American Economic Review*, 99(4):1619–35, 2009.
- James E Parco, Amnon Rapoport, and William E Stein. Effects of financial incentives on the breakdown of mutual trust. *Psychological Science*, 13(3):292–297, 2002.
- David J Paulsen, Michael L Platt, Scott A Huettel, and Elizabeth M Brannon. From risk-seeking to risk-averse: the development of economic risk preference from childhood to adulthood. *Frontiers in psychology*, 3:313, 2012.
- Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, personality, and gains from cooperation in repeated interactions. *Journal of Political Economy*, 127(3):1351–1390, 2019.
- Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, errors, and cooperation in repeated interactions. *The Review of Economic Studies*, 89(5):2723–2767, 2022.
- Amnon Rapoport, William E Stein, James E Parco, and Thomas E Nicholas. Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior*, 43(2):239–265, 2003.
- Philip J Reny. Rationality in extensive-form games. *Journal of Economic perspectives*, 6(4):103–118, 1992.
- Robert W Rosenthal. Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic theory*, 25(1):92–100, 1981.
- Edward B Royzman, Kimberly Wright Cassidy, and Jonathan Baron. “i know, you know”: Epistemic egocentrism in children and adults. *Review of General Psychology*, 7(1):38–65, 2003.
- Itai Sher, Melissa Koenig, and Aldo Rustichini. Children’s strategic theory of mind. *Proceedings of the National Academy of Sciences*, 111(37):13307–13312, 2014.
- Matthias Sutter, Claudia Zoller, and Daniela Glätzle-Rützler. Economic behavior of children and adolescents—a first survey of experimental economics results. *European Economic Review*, 111:98–121, 2019.

Appendix A. Instructions: LILA (6-10) and USC

GOING DOWN THE STREET GAME

In this game, the computer will pair you with one other student. One of you will be BLUE and the other will be ORANGE. The computer decides who is BLUE and who is ORANGE. In this game, BLUE and ORANGE are walking down a street. On that street, there are blue and orange houses (see Figure 6 for the slides).

[SLIDE 1]

The first house is blue, the second one is orange, the third one is blue again, and so on. If BLUE and ORANGE pass by a BLUE house, BLUE decides whether to stop the game or to continue. If they pass by an orange house, ORANGE decides whether to stop the game or to continue. Because the first house is blue, BLUE always makes the first choice. Once a player stops, the game is over.

When a player decides to stop the game, both BLUE and ORANGE get points. But how many points?

[SLIDE 2]

When BLUE stops, BLUE gets more points the farther down the street he stops: 100, 240, 380, etc. . . In this game, 1 point equals 1 cent, so 1 dollar, 2.40 dollars, 3.80 dollars, etc.

[SLIDE 3]

When ORANGE stops, ORANGE also gets more points the farther down the street he stops: 170, 310, 450, etc.

[SLIDE 4]

Finally, when BLUE stops ORANGE gets only 50 points and when ORANGE stops BLUE gets only 50 points.

[SLIDE 5]

Putting all together, these are the points each person gets depending on who stops where.

[SLIDE 6]

For instance, imagine they reach the fifth house. This is just an example.

- Whose turn is it to decide? [answer: BLUE]
- If BLUE stops, how many points does he get? [answer: 380]
- How many points does ORANGE get? [answer: 50]
- If BLUE continues and ORANGE stops, how many points does ORANGE get? [answer: 450]
- And BLUE? [answer: 50]. Is it clear?

Now, let's look at what you will see on your tablet at the beginning of the game. If you are BLUE, your screen looks like this.

[SLIDE 7]

At the top of the screen, it says you are BLUE. You can see the whole street with the blue and orange houses. You and ORANGE are in front of the first blue house, which is highlighted in grey. Since the house is BLUE, it is your turn to choose. You decide to stop the game by clicking STOP

and then OK or to continue by clicking the green arrow and then OK. If you stop the points you and ORANGE get are highlighted in grey.

If you are ORANGE, your screen looks like this.

[SLIDE 8]

It says you are ORANGE. It is the same as what BLUE saw except that you cannot make any choice, since you are in a BLUE house. If BLUE decides to stop, the points you and BLUE get are highlighted in grey. If BLUE decides to continue, you will both go down the street, reach the next house, which is now ORANGE and it will be your turn to choose. In that case, BLUE will see this screen.

[SLIDE 9]

and ORANGE will see this screen

[SLIDE 10]

Is the game clear? OK, let's play then. Remember you are matched with someone from this room, but you will not know with whom you are playing, and it is not the point to know.

Start CENTIPEDE GAME 1

[At the end] We are going to play another round of this game, except that the street is much shorter this time, with only 4 houses. You will be playing with a different person than before. Are you ready?

Start CENTIPEDE GAME 2

Ok, we are going to play a few more times, again with different partners each time. It is the same game as the first time, with a long street. You will also change color each time (one time blue, one time orange).

Start CENTIPEDE GAME 3 (3 rounds)

We are done. In the next page you will see how many points you got in the "Going Down the Street" Game. These are worth 1 cent. You don't need to record it. The computer takes care of it. We are going to ask you some final questions. Please fill the questionnaire. We will then tell you how much money you made and you will receive your amazon e-giftcard in your email. Thanks for participating.

Start QUESTIONNAIRE

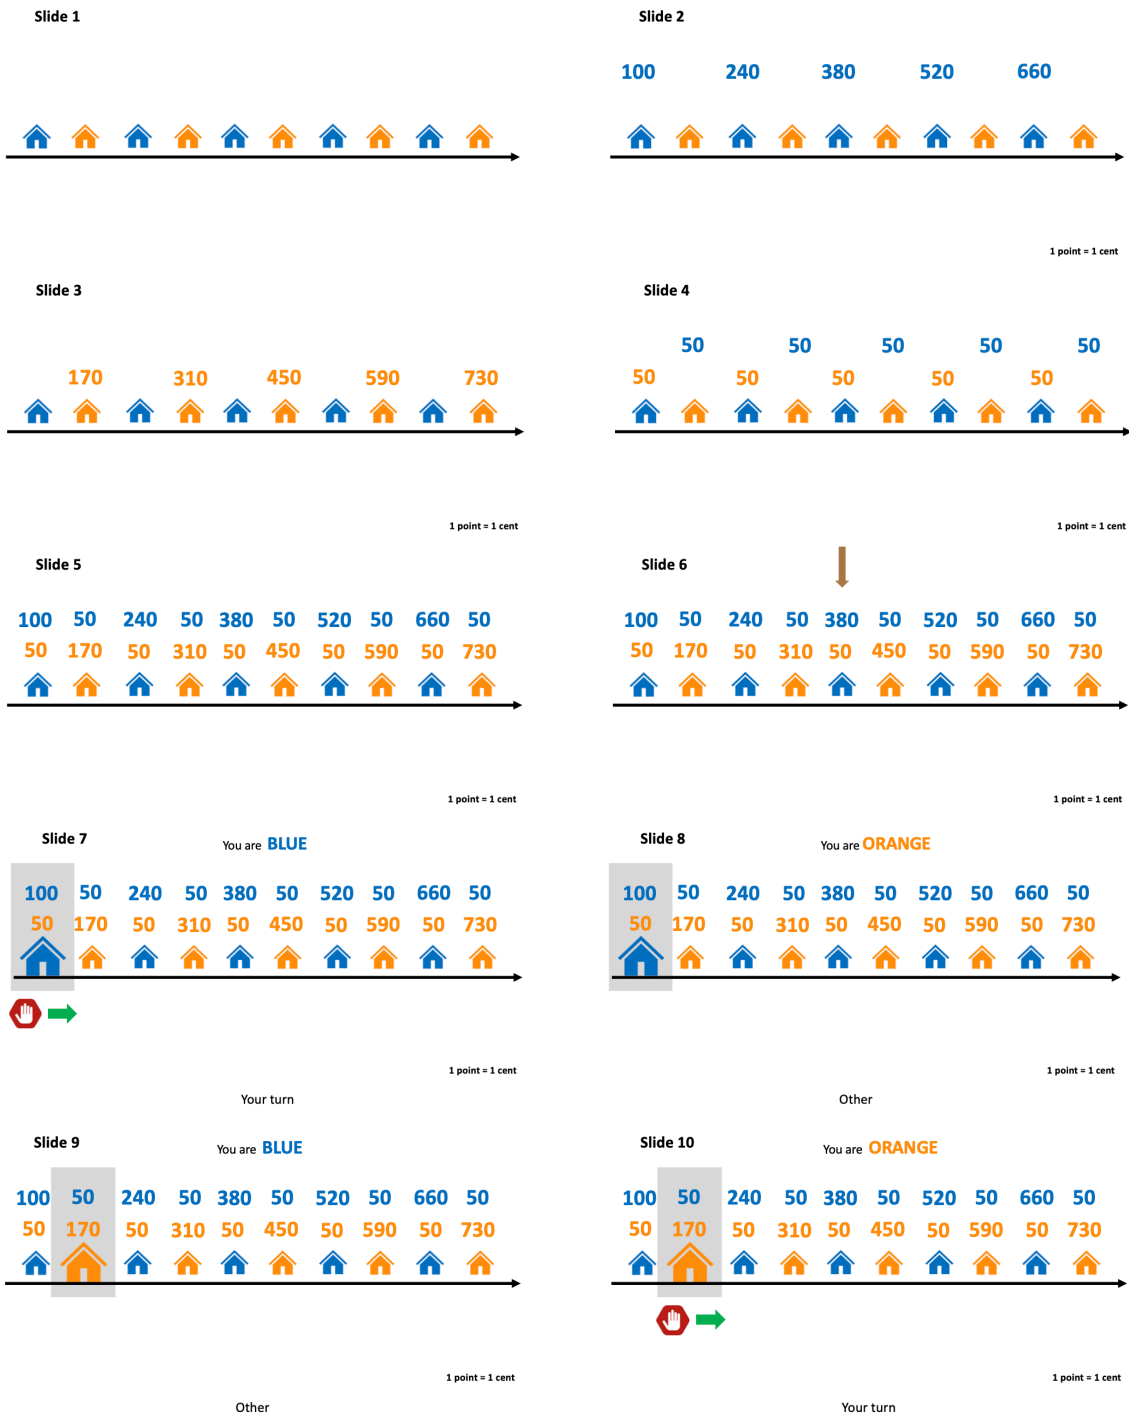


Figure 6: Slides projected on screen for instructions

Appendix B. Additional analysis for Experiment 1

Appendix B1. Factors of behavior in R1 and R2

Since player 1 keeps the same role in R1 and R2, we can use his initial choice in R1 and R2 (stop *S* or continue *C*) to infer drivers of behavior. *SS* corresponds to individuals who strictly apply backward induction. *CS* are participants who try to maximize expected payoffs. Those who stop immediately in R2 but continue in R1 (*CS*) seem to know both how to apply backward induction logic *and* also read the intentions of others and optimally pass when it is payoff improving. Those who wait in both games (*CC*) are likely motivated by seeking high rewards ignoring the risk involved (which is quite significant for the case of R2). Finally, individuals who stop in R1 but not in R2 (*SC*) probably miss all strategic aspects of the game. Figure 7 reports the proportion of player 1 types in each age group.

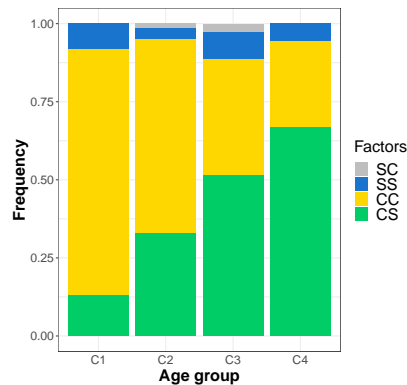


Figure 7: Joint behavior in first round of R1 and R2

The proportion of players who understand backward induction and possess theory of mind (*CS*) increases with age while the proportion of players motivated by risky rewards (*CC*) decreases. By contrast, *SS* and *SC* are infrequent in all age groups. The distribution in **C1** is significantly different from all other age groups ($p < 0.03$). The distribution in **C2** is significantly different from **C4** ($p = 0.007$) and there is no difference between **C3** and **C4**. Also, and consistent with the previous analysis, the proportion of players who pass in the first stage of R1 among those who stop immediately in R2 (formally, the proportion of *CS* types over *CS* and *SS* types) is significantly smaller in **C1** (61.5%) compared to **C2**, **C3** and **C4** altogether (89.9%, $p = 0.026$).¹³

Appendix B2. Factors of behavior in the long games

We have noted that behavior converges with repeated play and that age is not a predictor of behavior from round 3 on. Here, we hypothesize that heterogeneity within age groups may underlie this observation: while age per se does not tell how a participant is likely to behave, their reasoning

¹³This result should not be overemphasized as the number of *SS* observations is small in all age groups.

abilities may. We retain the participants who played first in R1 and R2 and whose behavior was categorized as *SS*, *CS* and *CC* depending on their initial choice in both games (we omitted the 2 participants classified as *SC*). We then run an OLS regression of the stopping stages of these participants in the long games R3 through R5, whenever they were given the opportunity to stop and acted upon it. We use *CC* as the benchmark category. It is worth noting that the number of observations drops significantly compared to previous regression analyses. The result is reported in Table 9.

	Stages before stop		
	R3	R4	R5
<i>SS</i>	-1.389*	-1.895**	-1.219
	(0.635)	(0.616)	(0.712)
<i>CS</i>	-0.117	-0.451	-0.828*
	(0.388)	(0.379)	(0.393)
<i>Male</i>	-0.146	-0.682	-0.700
	(0.375)	(0.355)	(0.383)
<i>const.</i>	3.963***	3.973***	4.391***
	(0.312)	(0.255)	(0.290)
Adj. R ²	0.027	0.116	0.105
# obs.	94	107	88

* p < 0.05, ** p < 0.01, *** p < 0.001.

Table 9: OLS of stopping stage in long games on factors of behavior

Compared to participants who behaved as if they were motivated by rewards in R1 and R2 (*CC* players), both those who played at equilibrium (*SS*) and those who applied ToM reasoning (*CS*) continue to stop earlier in all games. These differences are significant in R3 and R4 for *SS* participants and in R5 for *CS* participants. It indicates that tendencies observed at the beginning of the experiment persist for its entire duration.