# The centipede game at school: does developing backward induction logic drive behavior? *

**Isabelle Brocas**
*University of Southern California and CEPR*

**Juan D. Carrillo**
*University of Southern California and CEPR*

August 2022

**Abstract**

Adults do not play the Nash equilibrium in the well known centipede game. While Palacios-Huerta and Volij (2009) argued that behavior results from the failure of backward induction logic, Levitt et al. (2011) found that players who know how to backward induct still do not play Nash. Here, we ask children and adolescents (ages 8 to 16) to play the centipede game in the laboratory and we leverage knowledge about developing abilities to assess the contribution of backward induction logic. In line with the literature, we find that the ability to perform backward induction increases with age. However, it predicts behavior only in elementary school children: those with advanced logical abilities over-apply their skills. Starting in middle school, students who reason logically know that the unraveling argument should not be applied blindly. They utilize Theory-of-Mind (ToM) abilities to form beliefs about others' play and (optimally) refrain from stopping immediately. Their behavior is in line with the deviations observed in adults. Interestingly, developing ToM leads to a gradual decrease in stopping stages with age, which is accompanied by a decrease in payoffs with age. The results indicate that ToM is the key contributor of behavior that helps departing from backward induction when beneficial.

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

JEL Classification: C72, C90.

*We thank the members of the Los Angeles Behavioral Economics Laboratory (LABEL) for their insights and to Chris Crabbe for his outstanding programming skills. We are grateful to the staff of the Lycée International de Los Angeles (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 grant SES-1851915.

# 1 Introduction

Behavior in strategic games results from the combination of cognition, inferences about others and intrinsic preferences. Under the assumption that people are motivated by rewards, can correctly anticipate what others will do and possess cognitive skills that support flawless logic, game theory provides normative predictions of behavior. However, we often observe departures from these predictions, resulting sometimes in inferior but sometimes in superior outcomes. The centipede game is one of the best examples of a game where Nash 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 (in which case the player takes the largest pile, leave the smallest one to the partner and end the game) or to pass (in which case the partner faces a similar decision with larger piles of money). Equilibrium theory relies on backward induction arguments and prescribes stopping at each stage. This obviously prevents players from taking advantage of jointly increasing rewards.

Since its introduction by Rosenthal (1981), the centipede game has fascinated theorists, who describe it as a prominent paradox of backward induction (see e.g., Aumann (1992); Reny (1992); Ben-Porath (1997)). Even though the theoretical prediction of the backward induction equilibrium is sharp and simple, there are numerous intuitive reasons why participants may decide to pass initially. These include limited cognition and bounded rationality (inability to perform backward induction), non-selfish motivations (preferences for fairness or cooperation) and inferences about mutual uncertainty (about the players' cognitive ability and preferences). 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, 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 the backward induction prediction: few participants stop immediately. Naturally, 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 from Nash predictions indicate that people may not be able to apply backward induction logic, likely the result of limited cognition. This hypothesis has been investigated in two studies. Palacios-Huerta and Volij (2009) argued that experts (chess Grandmasters) are more likely to play the backward induction equilibrium than the regular population. Levitt et al. (2011) did not replicate the results and found, instead, no relationship between ability to backward induct (measured in a dominance solvable task) and the decision to stop immediately. The question of whether limited cognitive skills underlie behavior in the centipede game is therefore not fully answered. It is an important question because behavior in other strategic settings is believed to largely depend on cognition (Gill and Prowse, 2016; Proto et al., 2019, forthcoming; Fe et al., 2020).

Our study leverages age-related changes in both cognition and Theory-of-Mind (ToM) –the ability to read the rival's intentions and form beliefs– to clarify the contribution of these factors to the observed behavior in the centipede game. More specifically, we recruit children and teenagers to participate in an artefactual field experiment (Harrison and List, 2004) and we study the change in behavior with age in a version of the centipede game. Our hypotheses rely on known developmental changes in cognition and ToM abilities throughout development. We run 5 rounds of a centipede game with linearly increasing payoffs for the player who stops and constant payoffs for the one who does not.[1] 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 of the ability to backward induct. Therefore, contrasting behavior in rounds 1 and 2 allows us to disentangle between passing due to limited cognition and passing for other strategic motives at the different ages. Studying the evolution of behavior across long games (rounds 1, 3, 4 and 5) allows us to determine how age and the gradual development of abilities affects learning over the course of the experiment.

While there has been an increased interest in studying decision making by children and adolescents in experimental economics, the majority (though certainly not the entirety) of the literature focuses on individual decision making problems (see Sutter et al. (2019) and List et al. (2021) for excellent surveys). The study of games of strategy has revealed some interesting developmental trajectories that track the development of preferences (see e.g., Murnighan and Saxon (1998); Harbaugh and Krause (2000)), reasoning and ToM (see e.g., Sher et al. (2014); Czermak et al. (2016)) and largely correlates with measures of cognitive ability (Fe et al., 2020). 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). Finally, children show gradual improvements in their ability to foresee future events and input them in current calculations although, in some cases, such ability never matures fully or plateaus (Brocas and Carrillo, 2021).

The first and major finding is that the amount of passing the first time the game is played decreases monotonically with age, from elementary school to young adulthood. Differences in passing are 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), finding and playing the equilibrium are intimately related. Children either find and play the equilibrium in both the long and short games (rounds 1 and 2), or they do not play it in either game. In other words, there is a positive relationship between

[1] In section 2, we discuss why we choose this version of the centipede.

passing in both games in elementary school, supporting the hypothesis that cognition drives behavior. 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 in Levitt et al. (2011) and the majority of the experimental literature, 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, a consequence of the increased development of ToM.

The second and related finding is the monotonic decrease in payoffs with age which, to our knowledge, is a first in a game of strategy. A large fraction of elementary children do not backward induct and stop late in both the short and long versions of the game. This behavior is consistent with their immature reward processing and their tendency to focus on salient rewards. As they age, logic and ToM abilities develop making participants more prone to stop immediately in the short game and pass for a few periods only in the long game, with the corresponding decrease in payoffs. Adolescent still pass more than adults, likely due to the combination of their still less developed ToM ability, age-related reward processing, and anticipation of passing by the rival.

Our third finding relates to the change in behavior during the experiment. Participants of all ages stop earlier as they play more rounds. This tendency is a natural result of the asymmetry of incentives: failure in one round pushes subjects to preempt their rival in the next whereas success is unlikely to trigger postponement. However, it does not result in complete unraveling, at least within the five rounds of our experiment; 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 implies that choices (and therefore payoffs) are very similar across ages by the end of the experiment. Still, abilities revealed in the first two rounds continue to play a role: participants who demonstrate a capacity to backward induct in the second round continue to stop earlier than others while participants who stop very late in the first two rounds continue to stop later than others.

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 as a function of age in the first round of the long and short versions of the game. In section 4, we investigate the evolution in behavior during the five rounds of the experiment. Concluding remarks and alleys for future research are collected in section 5.

## 2 Experimental 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). Because working with young participants presents important methodological challenges, we follow the guidelines developed in Brocas and Carrillo (2020a).[2] In particular, we propose a graphical, story-based version of the original game to address those challenges.

*Population.* The experiment was conducted with 315 school-age students from 3rd to 10th grade at the Lycée International de Los Angeles (LILA), a private school in Los Angeles.[3] We also included a control adult population (A) consisting of 72 college students from the University of Southern California (USC).

With some exceptions (see e.g., Cobo-Reyes et al. (2020)), experiments with children and adolescents typically do not feature an adult population and, instead, rely on prior research for a comparison. We believe it is key to include an adult control group that follows *identical* procedures to establish a behavioral benchmark (Brocas and Carrillo, 2020a). This is especially important when procedures are slightly modified, as it is the case here. Ideally, the control population should also be as similar as possible to the treatment population of children to allow for meaningful comparisons. We argue that a private university like USC is a reasonable match for the LILA population.[4]

We follow List (2020) to assess the generalizability of our results. In our case, we argue that the preferences, goals and beliefs of children at LILA in this game are no different from those of children in the general population. Given our efforts to make the game graphical and accessible, it is also equally natural for all populations. On the other hand, because they are raised in a private school featuring small classroom size and individualized attention, the cognitive development of LILA students at a given age is likely to be, on average, superior to that of peers in less favored schools. This could make them more able to find the equilibrium compared to other children (as we found in Brocas and Carrillo (2021) for example). The same arguments also apply to our control population, which is more educated and cognitively developed than the general young adult population. Overall, while our participants have no reason to comprehend and respond to the game in a qualitatively different fashion from other populations, quantitative predictions should

[2]In a nutshell, the principles are: (i) adapt the length and procedures to a population with limited attention span; (ii) offer age-appropriate incentives (possibly different at different ages); (iii) present the task in a way that subjects are not required to possess strong analytical skills to participate (e.g., graphical interfaces and simple instructions); (iv) understand, describe and compare the children population, and (v) include a benchmark adult comparison group whenever possible.

[3]High schoolers from 11th and 12th grade did not participate in the study because they were taking or preparing for national exams during this period.

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

be taken with a grain of salt. In particular, the developmental trajectory may be delayed in other populations.

Table 1 summarizes the participants by grade and age in our sample.

| | LILA | | | | | | | | USC |
|---|---|---|---|---|---|---|---|---|---|
| Grade | 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

*Procedures.* We ran 27 and 6 sessions 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. The findings of the third-party dictator game are discussed in a different article (Brocas and Carrillo, 2020b). 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 appealing to children and adolescents (Brocas and Carrillo, 2020b). This is all the more important when the age span is large, as in our study (8 to 16 years old). 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 the second player at the second decision node. In our narrative, the blue and orange players walk together down the 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 payoff and the round ends. (Appendix A provides the full set of instructions).

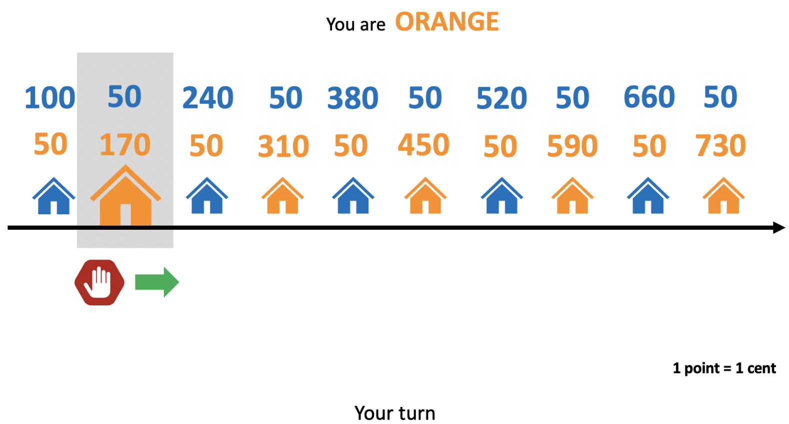


**Figure 1:** Screenshot of "Going down the street" game

*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.[5]

Contrasting behavior in R1 and R2 allows us to disentangle between different motives for early passing. Indeed, while R2 serves as a diagnostic test for the ability to perform backward induction, R1 helps reveal other drivers of behavior. In particular, participants familiar with backward induction but who consider also other motives are less likely to pass in the first stage of R2 than in the first stage of R1. Also, because players receive feedback and have the opportunity to adjust their strategies, including R3, R4 and R5 helps us study whether and how quickly different age groups make these adjustments.

*Remark.* Our setting uses different parameters compared to the original experiment (McKelvey and Palfrey, 1992): four and ten (instead of four and six) stages, a constant (instead of increasing) payoff for the participant who does not stop, and a linearly (instead of exponentially) increasing payoff for the participant who stops.[6] Opting for ten stages increases the incentives to pass in the first stages for strategic considerations. A constant payoff for the "loser" removes the possibility that both participants strictly win by passing

[5]This means that a player had one role in R1, R2 and R4 and the other role in R3 and R5.

[6]Researchers have studied experimentally multiple variants of the game, including (but not limited to) length (McKelvey and Palfrey, 1992), payoff structure (Fey et al., 1996), incentives (Parco et al., 2002), number of players (Rapoport et al., 2003) and game presentation (Nagel and Tang, 1998). The qualitative properties of the empirical behavior are usually robust to such modifications (see Krockow et al. (2016) for a detailed survey).

several times, which makes social preferences an unlikely driver of behavior. Finally, choosing a linearly increasing payoff for the "winner" –as in the original theory– ensures that the empirical variance in payments, to which children are particularly sensitive, is high but not extremely high. Given these changes, it is especially important to include a control group.

*Payments and duration.* Following the guidelines discussed in Brocas and Carrillo (2020a), we used different mediums of payment for different ages. This comes at a substantial added effort for the experimenter but it is, in our view, a key choice. The goal is to equalize, to the best of our ability, the *value of rewards* across age groups instead of equalizing the rewards themselves.[7] School-age students from 6th grade and above and control adults earned $0.01 per point paid immediately at the end of the experiment in cash (at USC) or with an amazon e-giftcard (at LILA, given that cash transfers are not allowed in the school).[8] For elementary school students (grades 3 to 5), we set up a shop with 20 to 25 pre-screened, age-appropriate toys and stationery that children find attractive (bracelets, erasers, figurines, die-cast cars, trading cards, apps, calculators, earbuds, scented pens, squishies, etc.). Before the experiment, we took the children to the shop, showed the toys they were playing for and explained the point prices of each toy. At the end of the experiment, subjects learned their point earnings and were accompanied to the shop to exchange points for toys.[9]

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 6th graders 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 control adult population to correct for differences in the opportunity cost of time. As for elementary school children, point prices were calibrated in a way that all children obtained at least three toys, although there was significant 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 the statistical power, some of our analysis groups the school-age participants into three naturally clustered 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

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

[8]Children in grade 6 and above are familiar with amazon giftcards, which are offered by the school as prizes for performance in intra-school activities (math puzzles, art shows, literature competitions, etc.).

[9]The 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.

13-16, 70 participants).[10] The control adult population consists of USC undergraduates (**C4**, ages 18-23, 72 participants). The regression analyses 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 and we use a p-value of 0.05 as the benchmark threshold for statistical significance.

# 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 (player 1 or player 2) had a choice in each stage. For visual ease, we group the R1 stopping decision in four categories: stages 1-2, 3-4, 5-6, 7-10. The results are presented in Figure 2.

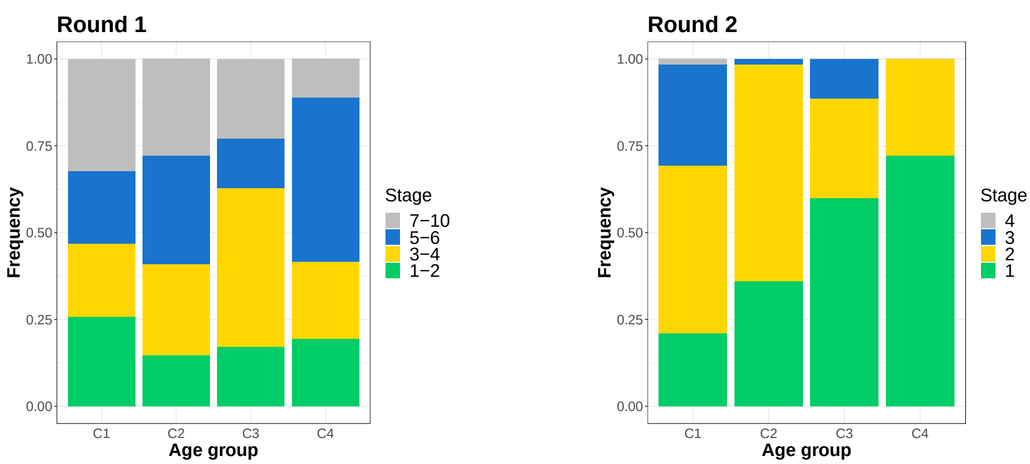


**Figure 2:** Stopping decision in Round 1 (left) and Round 2 (right)

The developmental trajectory is qualitatively similar in R1 compared to R2, featuring a gradual decrease in the tendency to stop at the latest stages. Younger participants act as if they are taking marginally more risks (stop later). In R1, participants become more inclined to stop at the intermediate stages (3-4 and 5-6) with age rather than at the later stages (7-10). The stopping distributions in C1 and C4 as well as C3 and C4 are

[10] While 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 (LILA has recently expanded the size of elementary and middle school, which explains the higher number of participants in those grades). Results are similar (though statistical significance is affected) if we consider other grouping methods.

significantly different (chi-squared tests, p = 0.021 and p = 0.014 respectively). Stopping in the participant's first opportunity (1-2) is roughly constant and rather infrequent in all age-groups. In R2, participants also stop significantly earlier as they grow older, with the average stopping stage moving from 2.11 in **C1** to 1.28 in **C4**, (**C1** different from all age groups, $p < 0.0001$, **C2** different from **C4**, p = 0.013). This also reflects an increased understanding of the strategic forces at play: as they age, participants learn to anticipate how the game will unravel if they don't stop early and to apply backward induction reasoning. Among Player 1 participants, 21% in **C1**, 36% in **C2**, 60% in **C3** and 72% in **C4** stop immediately.

To refine this analysis, we focus on the school age population and we conduct OLS regressions of the number of stages before stopping as a function of the participant's role (*Player2* = 1), *Age* (in months), gender (*Male* = 1) and the interaction between age and gender. 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.[11] We do not include our control undergraduate students in these regressions to not bias the coefficients of the regression in their direction. The results are presented in Table 2.

| | Stages before stop | | Payoff of winner | |
|---|---|---|---|---|
| | R1 | R2 | R1 | R2 |
| *Age* | -0.012* | -0.005** | -1.651* | -0.642** |
| | (0.006) | (0.001) | (0.748) | (0.200) |
| *Male* | -2.036* | -0.185 | -285.1* | -25.9 |
| | (1.015) | (0.278) | (142.1) | (38.9) |
| *Age×Male* | 0.009 | 0.001 | 1.290 | 0.071 |
| | (0.007) | (0.002) | (0.968) | (0.266) |
| *Player2* | -0.469* | -0.292*** | 4.320 | 29.2*** |
| | (0.180) | (0.050) | (25.2) | (6.98) |
| *const.* | 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 regressions strengthen the previous results. Indeed, the stopping stage significantly decreases with age in R1 (p = 0.029). The effect is much stronger in R2 (p = 0.002).

[11] In 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.

Interestingly, males stop earlier than females in R1 but not in R2. This suggests that there is no gender differences in the ability to respond purely strategically (in R2) while gender-related motives are at play in R1. Table 2 also reveals that player 2 chooses less frequently to pass than player 1, although this should be interpreted with caution, since it is conditional on being given the same opportunities (which is endogenous to the model).

In terms of payoff, this behavioral trajectory results in a decrease in payoffs from age 8 to age 16: younger children end up stopping later and getting higher payoffs compared to their older peers. Notice that the payoff ($\Pi$) 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 (for comparability between player 1 and 2) whereas we consider payoffs separately. In turn, this methodology allows us to unveil an interesting effect of role (*Player2*): 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.

### 3.2 Choice across rounds

Taking advantage of the fact that participants keep the same role in R1 and R2, we ask whether the behavior of an individual in R1 relates to their behavior in R2. To do this, we run a Probit regression of the first player's choice in the first stage of R2 (Stop(R2) = 1) as a function of the number of stages in which that player chose to pass in R1 (*PassR1*). Again, we focus on the school-age population and we use **C2** as the benchmark age group. We include dummies for the other age-groups (*C1* and *C3*) as well as interaction terms, and a dummy for gender. It is important to acknowledge the imperfect nature of this exercise. Indeed, our data is censored given that an individual can only make a choice in stage $t$ if the partner passed in stage $t-1$. Results are reported in Table 3

Remember that R2 partly serves as a diagnostic test of cognitive ability. In the four-stage version, stopping immediately is both the equilibrium and the optimal strategy assuming that a partner is minimally strategic. In R1, however, while the backward induction equilibrium prescribes stopping at the earliest opportunity, a participant may not assume that their partner (even a strategic one) will choose that strategy. Table 3 sheds some interesting light on the relationship between cognition and behavior.

We observe that 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. They seem to play the equilibrium strategy if they manage to find it, without asking whether their partner is capable of the same logic or willing to apply it always. While foregoing these inferences is not payoff-detrimental in

| | Stop(R2) |
|---|---|
| *PassR1* | -0.256 |
| | (.164) |
| *C1* | 0.756 |
| | (.643) |
| *PassR1* $\times$ *C1* | -0.998* |
| | (.410) |
| *C3* | 0.490 |
| | (.626) |
| *PassR1* $\times$ *C3* | -0.035 |
| | (.266) |
| *Male* | 0.396 |
| | (.247) |
| *const.* | 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

R2, it is in R1.[12]

By contrast, our middle- and high-schoolers are more discerning of the situation. Passing in the long game does not predict their behavior in the short version and is therefore not an indication of limited cognition. Participants who understand the incentives to stop and apply that logic (optimally) in R2 act as if they draw inferences about their partner that guide them away from the equilibrium behavior.[13]

Differences in behavior across games and players can also be addressed by asking whether the number of foregone stopping opportunities in R1 are predicted by age and (our crude measure of) cognitive ability. For that purpose, Table 4 reports an OLS regression of the number of stages before stopping in R1 on the decision to stop immediately in R2, age and gender.

This result shows that participants who are less cognitively equipped (do not stop immediately in R2) tend to delay stopping significantly more in R1. It indicates that these participants are less concerned by or aware of the inherent risks associated with not stopping. The effect is modulated by age and is weaker in older participants, consistent

[12] Ideally, one would like to run another treatment where participants play first the four-stage version and then the ten-stage version.

[13] As 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 would not change the result that stopping early in R1 predicts stopping in the first stage of R2 only for **C1** (data omitted for brevity).

| | Stages before Stop (R1) |
|---|---|
| *Age* | -0.008 |
| | (.004) |
| *Stop(R2)* | -3.200** |
| | (1.101) |
| *Male* | -0.337 |
| | (.410) |
| *Age×Stop(R2)* | 0.016* |
| | (.007) |
| Adj. $R^2$ | 0.159 |
| # obs. | 158 |

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

**Table 4:** OLS of stopping stage in R1 as a function of Stopping in R2

with the fact that behavior in R2 is predicted by behavior in R1 (and conversely) only in young children.

### 3.3 Factors of behavior

Joint behavior in the first round of R1 and R2 may be used to infer drivers of behavior. Arguably, individuals who play Nash and stop immediately in both games strictly apply backward induction. Those who stop immediately in R2 but wait in R1 seem to know both how to apply backward induction logic *and* also to use ToM abilities to infer and react to the empirical behavior of others. Those who wait in both games are likely motivated by seeking high rewards at the expense of some risk, and those who stop in R1 but not in R2 probably miss all strategic aspects of the game. Here, we create the categorical variable to summarize joint behavior in the first round of R1 and R2. Following the order of the description above, categories are denoted *Logic*, *ToM*, *Reward* and *Illogic*. The distribution across age groups are represented in Figure 3.

The proportion of players who display both logic and Theory of mind (*ToM*) increases with age while the proportion of players motivated by risky rewards (*Reward*) decreases. By contrast, *Logic* and *Illogic* play are rare in all age groups. The distribution of factors 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 *ToM* types over *ToM* and *Logic* types) is significantly smaller in **C1** (61.5%) compared to **C2**, **C3** and **C4** ($p = 0.026$).

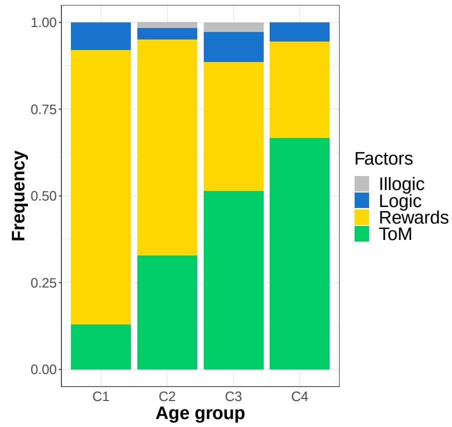


**Figure 3:** Joint behavior in first rounds of R1 and R2

### 3.4 Summary

Our analysis reveals a clear developmental trajectory. In both rounds, older participants stop earlier than their younger peers, but these age-related differences are modulated by the length of the game, with a steeper gradient in the short than in the long game. Participants tend to stop at intermediate stages in the long game and immediately in R2 due to the stronger incentives to undercut the rival.

The short R2 round is instructive as a diagnostic tool for cognitive ability. More cognitive able participants should stop immediately by applying backward induction logic. The contrast between behavior in R1 and R2 reveals an important 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). Because rivals take risks empirically in situations like R1 (but not in R2), it is optimal to depart from strict backward induction logic and factor that belief in the formulation of an empirical best response. The joint behavior in R1 and R2 can be used to determine what combination of cognition, ToM and reward seeking attitude drives a player's behavior. There are 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, a significant fraction (21%) behave consistently with backward induction logic in R2. Interestingly, these individuals fail ToM and do not anticipate that the majority of their peers would experiment in R1. As a result they *over-apply* their skill. Also, the fact that the least cognitively able participants stop very late suggests that their reward seeking attitude is due to a lack of awareness of risks rather than an intrinsic preferences for (calculated) risks. This is consistent with the known tendency of children to focus on salient features, which decreases gradually until age 11 (Miller, 2002). It is also consistent with the fact that young children respond to magnitude of payoffs rather than probabilities over such payoffs in risk tasks, and learn to integrate both only gradually during elementary school (Brocas et al., 2019).

Second, while participants in **C2** and **C3** still respond to rewards and do not all apply logic in R2, the behavior of those who know how to apply backward induction in R2 does not drive their behavior in R1 anymore. They have acquired enough ToM abilities to assess empirical risk with reasonable accuracy. The fact that backward induction abilities in R2 do not drive behavior in R1 among older participants supports the findings in Levitt et al. (2011) who show (in contradiction with Palacios-Huerta and Volij (2009)) that chess Grandmasters do not stop immediately in the centipede game even though they apply backward induction reasoning in other diagnostic tasks. Last, it is worth noting that reward seeking does not disappear during teenage time: participants in **C2** and **C3** are still prone to stop later than adults. However, while reward seeking is best explained by the combination of a lack of cognition and salience effects among our younger participants in **C1**, it likely reflects significant markers of differential behavior during teenage time, such as impulsivity and immature risk avoidance processes (Li, 2017).

Finally, while reward seeking starts as the main driver of behavior, the contributions of cognition and ToM increase with age, in that order. This trajectory results in earnings monotonically decreasing with age both in R1 and R2. A majority of young children and a significant number of adolescents take large risks and get, on average, high rewards. As this proportion decreases and participants settle on passing less, aggregate payoffs shrink.

# 4 Behavior in long games

In this section, we analyze differences in behavior in the long games R1, R3, R4 and R5. We want to both identify systematic differences (or similarities) between age groups and reveal how experience affects behavior.

## 4.1 Unraveling within age groups

Figure 4 describes the evolution of the stopping strategy in each age group separately. Table 5 reports the average stopping stage in R1 and R5 in each age group.

As we can see from Figure 4, there is a noticeable tendency in all age groups to stop earlier as the experiment progresses. The differences between the average stopping stage in the first and last round of the game reported in Table 5 are significant in all age groups (all p-values $< 0.05$). By the time they reach the fifth round, participants stop between one and two rounds earlier than they did in the first round.

The effect is most pronounced in the youngest population who, as we documented previously, stopped the latest in R1. It is followed by the control group. This change in behavior over rounds is consistent with existing findings in adult populations (McKelvey

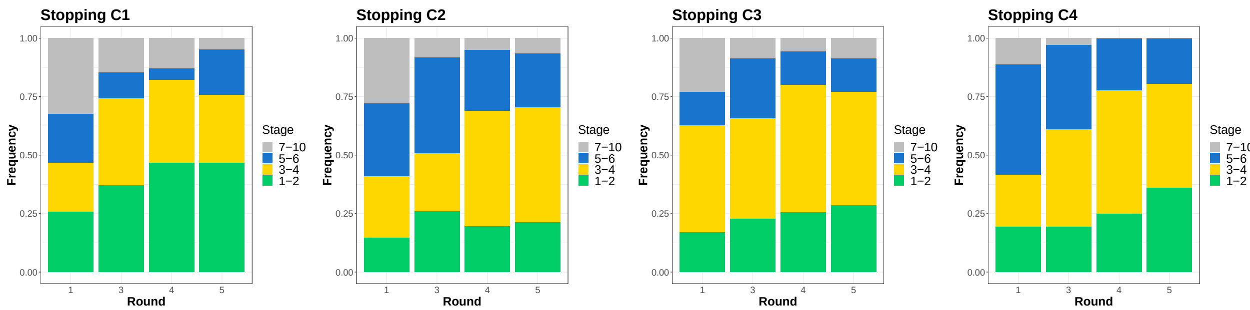


**Figure 4:** Evolution of the stopping strategy by age group

| | **C1** | **C2** | **C3** | **C4** |
|---|---|---|---|---|
| Round 1 | 4.95 | 5.05 | 4.43 | 4.58 |
| | (0.34) | (0.26) | (0.37) | (0.31) |
| Round 5 | 3.15 | 3.79 | 3.49 | 3.11 |
| | (0.32) | (0.38) | (0.43) | (0.41) |
| difference | 1.80*** | 1.26*** | 0.94* | 1.47*** |

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

**Table 5:** Average stopping stage in R1 and R5

and Palfrey, 1992; Fey et al., 1996). It is also quite natural. Participants whose rivals stop before they do receive feedback that incentivizes them to stop earlier in the next round. By contrast, those who stop before their rivals do not know the counterfactual of delaying stopping. This makes them less likely to postpone stopping in future rounds. Such asymmetry of incentives results in progressive –though not full– unraveling.

To further investigate the evolution of play within each age group, we run OLS regressions of the stopping stage as a function of round (R1, R3, R4, R5) and gender, separately for each age group. We use R3 as the default round, to cleanly determine the changes in behavior for the entire length of the experiment. The results are presented in Table 6.

The regressions confirm that the stage at which participants stop decreases over the course of the experiment in all age groups. Once again, the change is most pronounced in **C1**, featuring significant differences between all rounds. In **C2**, the change is most abrupt early in the experiment (between R1 and R3). By contrast, changes are not significant in **C3**. An immediate implication is the corresponding decrease in payoffs for participants in all age groups as the game progresses. We also observe that the gender effect previously noted in R1 and R2, whereby males are more likely to stop earlier, persists but is significant only in **C1**.

| | Stages before stop | | | |
|---|---|---|---|---|
| | **C1** | **C2** | **C3** | **C4** |
| *R1* | 1.241*** | 0.980** | 0.429 | 0.683* |
| | (0.365) | (0.295) | (0.421) | (0.282) |
| *R4* | -0.739* | -0.330 | -0.581 | -0.393 |
| | (0.286) | (0.240) | (0.324) | (0.252) |
| *R5* | -0.641* | -0.283 | -0.533 | -0.848** |
| | (0.307) | (0.262) | (0.356) | (0.324) |
| *Male* | -1.543*** | -0.231 | -0.324 | -0.451 |
| | (0.395) | (0.319) | (0.430) | (0.362) |
| *const.* | 4.682*** | 4.164 | 4.213*** | 4.159*** |
| | (0.375) | (0.249) | (0.353) | (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 6:** OLS Regressions of evolution in stopping stage in each age group

## 4.2 Age effects

To better assess the determinants of stopping and identify potential changes as the experiment progresses, we conduct the same OLS regressions as in Table 2 for Rounds 3, 4 and 5. The results are presented in Table 7.

In Table 6, we emphasized changes in behavior between R1 and R3 and we noted that they were largest for children in elementary, followed by children in middle-school. Consistent with this result, Table 7 shows that the effect of age on delaying stopping has dissipated by Round 3, and remains non-significant in Rounds 4 and 5. Consequently, age effects on payoffs also dissipate. 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).

Taken together, the analysis of the evolution of stopping behavior reveals convergence. While participants start with different behavior at different ages, they adapt over time and end up with a similar strategy, stopping early but not immediately. This decrease in "cooperative passing" has detrimental payoff consequences in all age groups.

## 4.3 Age vs. experience

While only some differences across stages and age groups are statistically significant, an important question is whether they are behaviorally meaningful. In particular, we want to determine if a sophisticated strategic player would behave differently as a function of the round and the partner's age. To answer this question, we compute the best response

| | Stages before stop | | | Payoff of winner | | |
|---|---|---|---|---|---|---|
| | R3 | R4 | R5 | R3 | R4 | R5 |
| *Age* | -0.005 | -0.001 | -0.003 | -0.716 | -0.067 | -0.367 |
| | (0.006) | (0.004) | (0.004) | (0.620) | (0.545) | (0.528) |
| *Male* | -2.166* | -1.592* | -1.784* | -303.2* | -222.9* | -249.7* |
| | (0.839) | (0.777) | (0.762) | (117.5) | (108.7) | (106.6) |
| *Age×Male* | 0.013* | 0.009 | 0.010* | 1.814* | 1.242 | 1.463* |
| | (0.006) | (0.005) | (0.005) | (0.805) | (0.745) | (0.731) |
| *Player2* | 0.804*** | -0.348* | 0.334* | 42.6* | 21.2 | -23.3 |
| | (0.150) | (0.143) | (0.139) | (21.0) | (20.0) | (19.5) |
| *const.* | 2.705*** | 2.374*** | 2.353*** | 408.7*** | 292.4*** | 359.4*** |
| | (0.647) | (0.568) | 0.561) | (90.5) | (79.5) | (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 7:** OLS of stopping stage and payoffs in R3, R4 and R5 in school-age population

strategy of a risk-neutral individual to the empirical behavior of the participants in our experiment. Formally, we consider each round and age group separately, and determine the aggregate empirical likelihood that a player stops at each stage. Given such probabilities, we then compute for the player in the other role the expected payoff of stopping at all possible stages. In order to study the evolution of behavior, we present in Table 8 the information for the first (left) and last (right) round of player 1 (top) and player 2 (bottom), separately for each age group (data for rounds 2, 3, and 4 is omitted for brevity but available). We highlight in **bold** the best response strategy, that is, the stopping stage that maximizes the participant's expected payoff.

| **Pl.1** | Round 1 | | | | | Round 5 | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | 1 | 3 | 5 | 7 | 9 | 1 | 3 | 5 | 7 | 9 |
| **C1** | 100 | 172 | **199** | 186 | 89 | 100 | 122 | **141** | 82 | 71 |
| **C2** | 100 | 207 | **210** | 131 | 67 | 100 | **183** | 116 | 50 | 50 |
| **C3** | 100 | **216** | 111 | 79 | 50 | 100 | **151** | 116 | 81 | 50 |
| **C4** | 100 | 195 | **223** | 95 | 50 | 100 | **169** | 91 | 50 | 50 |

| **Pl.2** | Round 1 | | | | | Round 5 | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | 2 | 4 | 6 | 8 | 10 | 2 | 4 | 6 | 8 | 10 |
| **C1** | 150 | **209** | 192 | 155 | 50 | **130** | 105 | 62 | 50 | 50 |
| **C2** | 156 | **240** | 219 | 50 | 50 | **155** | 151 | 102 | 50 | 50 |
| **C3** | 145 | 187 | **197** | 107 | 50 | **152** | 115 | 90 | 50 | 50 |
| **C4** | 154 | **223** | 103 | 50 | 50 | **128** | 115 | 50 | 50 | 50 |

**Table 8:** Best response payoffs in R1 and R5 by role and age group

There are three important findings. First, except for player 1 in **C1** and **C3**, best response behavior always entails stopping earlier in round 5 than in round 1, resulting in significantly lower expected payoffs. Pooling both players together, the maximum attainable expected payoff at the end of the experiment is between 66.4% and 75.1% the maximum expected payoff at the beginning, depending on the age group. In other words, even a best responder incurs sizable earning losses as the experiment advances. Second, the best response strategy is quite similar in all age groups, with the exception of **C3** in R1 and **C1** in R5. In R1 and in the presence of a **C3** rival, a best responder should stop earlier as Player 1 and later as Player 2 indicating that participants stop sooner as Player 2 and later as Player 1 compared to participants in other age groups. In R5 and in the presence of a **C1** rival, a best responder should stop later as Player 1 indicating that participants stop later as Player 2 compared to participants in the other age groups. Last, best response behavior converges (with the exception of **C1**): player 1 should pass only once and player 2 should stop immediately.

Overall, differences across age groups and rounds are such that a best responder should adopt very similar strategies irrespective of the rival's age but should adjust them as the experiment develops. This reinforces the findings of the previous subsections, that behavior changes more with experience than with age.

### 4.4 Reaction to empirical risk in the long games

As noted earlier, the incentives to stop at any point in time are modulated by several motives, including reward seeking effects, reasoning abilities and beliefs about empirical risk. As stages progress, a participant should realize that not stopping has both an increased risk (likelihood that the rival stops next) and a higher opportunity cost (difference between the current payoff of stopping and the 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 and reveal differences across ages, 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, $h_t = \frac{p_t}{\sum_{i=t}^{T} p_i} \in [0, 1]$. Strategic decision-making predicts an increasing hazard rate ($h_{t+1} > h_t$). Figure 5 reports the data by age group. Since the number of observations decreases significantly as we move to later stages and since stopping after round 6 is rare, we present hazard rates only for the relevant range of stages 1 through 6. We also pool together Rounds 1, 3, 4 and 5 to increase statistical power, even though we are aware that behavior changes over rounds, and observations across rounds are collected from the same individuals.

As typical in experiments with adults, individuals in our control population (**C4**) display an increasing hazard rate, indicating that they are each time more likely to stop

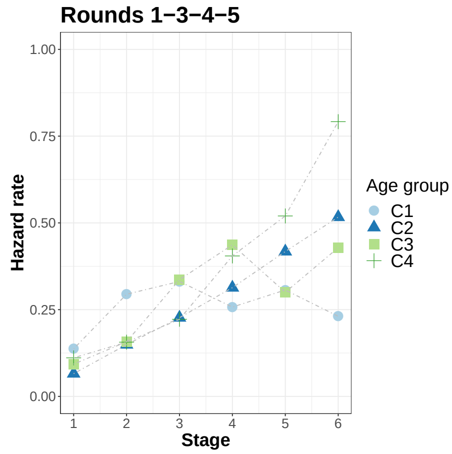


**Figure 5:** Stopping hazard rate in long games by age group

conditional on reaching a given stage. In this group, the estimated slope of the time series is steep (0.124) and very significantly different from zero (p = 0.008). Hazard rates also increase steadily in **C2** and **C3**, although the slope is statistically significant in the former (0.089, p = 0.009) but not in the latter (0.067, p = 0.132). By contrast, our youngest participants (**C1**) exhibit an almost constant hazard rate (with an estimated slope of 0.004, p = 0.999). Within a round, groups **C2**, **C3** and **C4** act as if the short term risk of stopping is higher as more stages have passed. Participants in groups **C1** however play in a memoryless fashion. This behavior is difficult to reconcile with a theory where rivals are perceived as strategic, and expected costs and benefits of passing at each stage are traded-off accordingly. The result reinforces the findings in section 3, where we argued that some of our youngest participants do not stop because they do not anticipate the strategic incentives of others and rather focus on their potential increasing payoffs.

### 4.5 Factors of behavior in the long games

We have noted that behavior converges with repeated play. Table 7 illustrates that age is not a predictor of behavior in the long games R3 to R5. 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 abilities may. We retain the participants who played first in R1 and R2 and whose behavior was categorized as *Logic*, *ToM* and *Rewards* depending on their initial choice in both games (we omitted the 2 participants classified as *Illogical*). 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 *Rewards* 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 |
| *Logic* | -1.389* | -1.895** | -1.219 |
| | (0.635) | (0.616) | (0.712) |
| *ToM* | -0.117 | -0.451 | -0.828* |
| | (0.388) | (0.379) | (0.393) |
| *Male* | -0.146 | -0.682 | -0.700 |
| | (0.375) | (0.355) | (0.383) |
| *const.* | 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, both those who behaved logically and those who applied ToM reasoning continue to stop earlier in all games. These differences are significant in R3 and R4 for participants classified as *Logic* and in R5 for participants classified as *ToM*. This indicates that tendencies observed early in the experiment are persistent.

### 4.6 Summary

The main lesson of this section is the overall decrease in the stopping stage in all age groups over the course of the experiment. Behavior converges to a large extent. Given that behavior in the first round was 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 increasing incentives to stop within round. Furthermore, despite the general tendency of all participants to learn to stop earlier, initial tendencies (in R1 and R2) to act purely logically, to respond to rewards or to apply ToM abilities are still present in R3, R4 and R5.

## 5 Conclusion

Behavior in strategic games results from the combination of cognitive abilities, ToM and intrinsic motivations. In the centipede game, each option to pass is a gamble that may lead to higher rewards and a participant might be driven by these prospects.[14] However, equilibrium theory prescribes stopping at each stage, independently of one's risk preferences, reward evaluation and game length. Still, numerous experiments show that adults do delay stopping in long games. There are at least two potential motives. On the one hand, a participant may not be able to perform a logical reasoning, in which case they do not represent others as performing such reasoning either. They evaluate the decision based on their motivations. On the other hand, a participant may be able to backward induct and correctly eliminate the dominated 'passing' options but knows that others empirically do not stop. Their decision is to best respond to the empirical risk and take a measured chance. Our study sheds light on these motives and on how reasoning develops with age.

We find that logical abilities develop gradually, leading to a decrease in stopping stage with age. Interestingly, while young participants are least likely to act logically, those who do, tend to over-estimate the ability of their peers to behave similarly. As a consequence, they apply the logic blindly. With age, participants learn to correctly assess 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 in the long game is in line with the majority of the literature on the centipede that documents strong deviations from backward inductions even after experience. It is consistent with Levitt et al. (2011) who show not only that chess Grandmasters do not behave differently from standard subjects, but that there is no correlation between early stopping in the centipede game and ability to backward induct in more challenging, dominance-solvable, zero-sum games.

The intuitive way to approach the game for participants with limited cognition is as a series of gambles leading to potential increasing rewards. The behavioral differences we observe in the short game between participants who stop immediately and those who do not indicates that reward seeking is a strong motive of behavior among those who do not apply logic. A majority of young participants reason in that fashion. The inclination of young children to take risks has been widely documented in the literature (e.g., 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 disproportionally more attention to rewards. In our case, it makes them not only willing to take more risks than older participants, but also willing to stop at each stage with the same conditional probability independently of the number of stages that have passed. Overall, passing in young children results from lack of cognition and ToM, while it follows a calculated empirical risk in older participants.

We have deliberately referred to motivations and reward seeking behavior rather than to risk attitude. While the observed behavior of each participant may be captured by a

[14]By contrast, social preferences are unlikely to drive behavior since, in our formulation, both players cannot strictly improve by passing multiple times.

typical risk preference in response to the empirical behavior of rivals, we believe that the underlying mechanisms that lead to choice not only are complex but they also change with age. While many young participants display a very late stopping behavior only consistent with salience effects and limited cognition, the tendency to delay stopping in teenagers more likely reflects intrinsic age-related impulsivity or immature risk avoidance mechanisms. Also, while the decrease in late stopping with age is consistent with documented decreases in risk taking attitudes during development (Sutter et al., 2019), other predictions from that literature are not supported by our data. For instance, many studies show that higher cognitive skills and more developed executive functions are associated with a preference for risk taking in children and adolescents (Andreoni et al., 2020). Here, more cognitively able children take fewer risks. It suggests that strategic risk in the centipede game is qualitatively different from preference-based risk in individual decision-making. Differences in behavior seem to stem from differences in abilities to reason about potential consequences and to downplay salient features or impulses during reward processing. 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 again not supported by our data. It may be due to the fact that the task does not explicitly describes probabilities. Instead, probabilities need to be inferred subjectively with the help of ToM abilities. We also conjecture that behavior in our strategic task relates to gender-related differences in competitiveness (Niederle and Vesterlund, 2011). Interpretations should also be taken with caution because our study is not designed to assess gender differences. Still, 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.

Middle- and high-schoolers understand the unraveling logic of the short game at least as well as participants in the four-stage version of the original game (McKelvey and Palfrey, 1992). And yet, there is heterogeneity within age. Some very young participants understand it equally well (21% of elementary children) while some educated adults (28%) do not. Interestingly, here we notice that cognitively equipped young children are not able to *not apply* their skill, while their older peers are. It underscores the importance of ToM in the context of the long centipede game: the best response to what others do is not to play the Nash equilibrium but to best respond to non-equilibrium beliefs. This is particularly important in this game, since once the first player passes in stage 1, backward induction ceases to be a relevant concept. The fact that some young children play the Nash equilibrium in R1 and R2 indicates that ToM develops after the logical backward induction skills required in this game.

To our knowledge, this is the first experimental study where earnings monotonically decrease with age. Cognitive sophistication allows each player to anticipate empirical risk more accurately and to best respond to it. As players age and acquire these skills, their behavior comes closer to equilibrium behavior resulting in a decrease in payoffs. Our study also demonstrates that while reasoning abilities are important to make logical iterations and formulate best responses, ToM is critical to evaluate empirical risks. Players who reason logically but fail ToM obtain lower payoffs than players who do not display any of these skills. Overall, young children obtain large earnings because they are lured by the high rewards in late stages. As logical abilities and ToM develop, participants take fewer chances, resulting in earlier stopping and therefore lower earnings.

Experience heavily disciplines behavior, and by the fifth round participants in all grades end up making very similar choices on aggregate. The fact that young children rapidly adapt their decisions indicates that they use feedback as an input in their choices across rounds, and apply inductive logic to what they observe. It would be interesting to refine those ideas with a larger pool of participants and more rounds, in particular to see whether children learn to respond to empirical risk, to backward induct or both. It would also be interesting to link individual learning to the feedback received. For instance, we expect children to learn differently if the first rival they meet stops early or takes some chances.

As noted earlier, the centipede game is largely dependent on beliefs. For that reason, it is an ideal game to mix participants who differ in cognitive abilities but also in abilities to read other players' minds. This kind of exercise is instructive because it tells us about the capacity of people to integrate relevant information about potential rivals and to fine tune their behavior to this information. Examples include Palacios-Huerta and Volij (2009) who mixed undergraduate students and professional chess players in the centipede paradigm itself and Proto et al. (forthcoming) who mixed players with different intelligence levels in the repeated prisoner's dilemma. An interesting alley for future research would be to mix participants from different grades. On the one hand, it would allow us to measure the capacity of young children to act in a more sophisticated manner in the presence of older children. On the other hand, it would shed light on the ability and willingness of older children to take advantage of less cognitively able rivals.

# References


James Andreoni, Amalia Di Girolamo, John A List, Claire Mackevicius, and Anya Samek. Risk preferences of children and adolescents in relation to gender, cognitive skills, soft skills, and executive functions. *Journal of economic behavior & organization*, 179:729–742, 2020.

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.

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. Developmental meritocracy. USC Working Paper, 2020b.

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

Isabelle Brocas, Juan D Carrillo, T Dalton Combs, and Niree Kodaverdian. The development of consistent decision-making across economic domains. *Games and Economic Behavior*, 116:217–240, 2019.

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.

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

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.

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

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.

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.

Glenn W Harrison and John A List. Field experiments. *Journal of Economic literature*, 42(4):1009–1055, 2004.

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.

Rosa Li. Flexing dual-systems models: How variable cognitive control in children informs our understanding of risk-taking across development. *Developmental cognitive neuroscience*, 27:91–98, 2017.

John A List. Non est disputandum de generalizability? a glimpse into the external validity trial. Technical report, National Bureau of Economic Research, 2020.

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 strategic choices in the repeated prisoners' dilemma. *Review of Economic Studies*, forthcoming.

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**

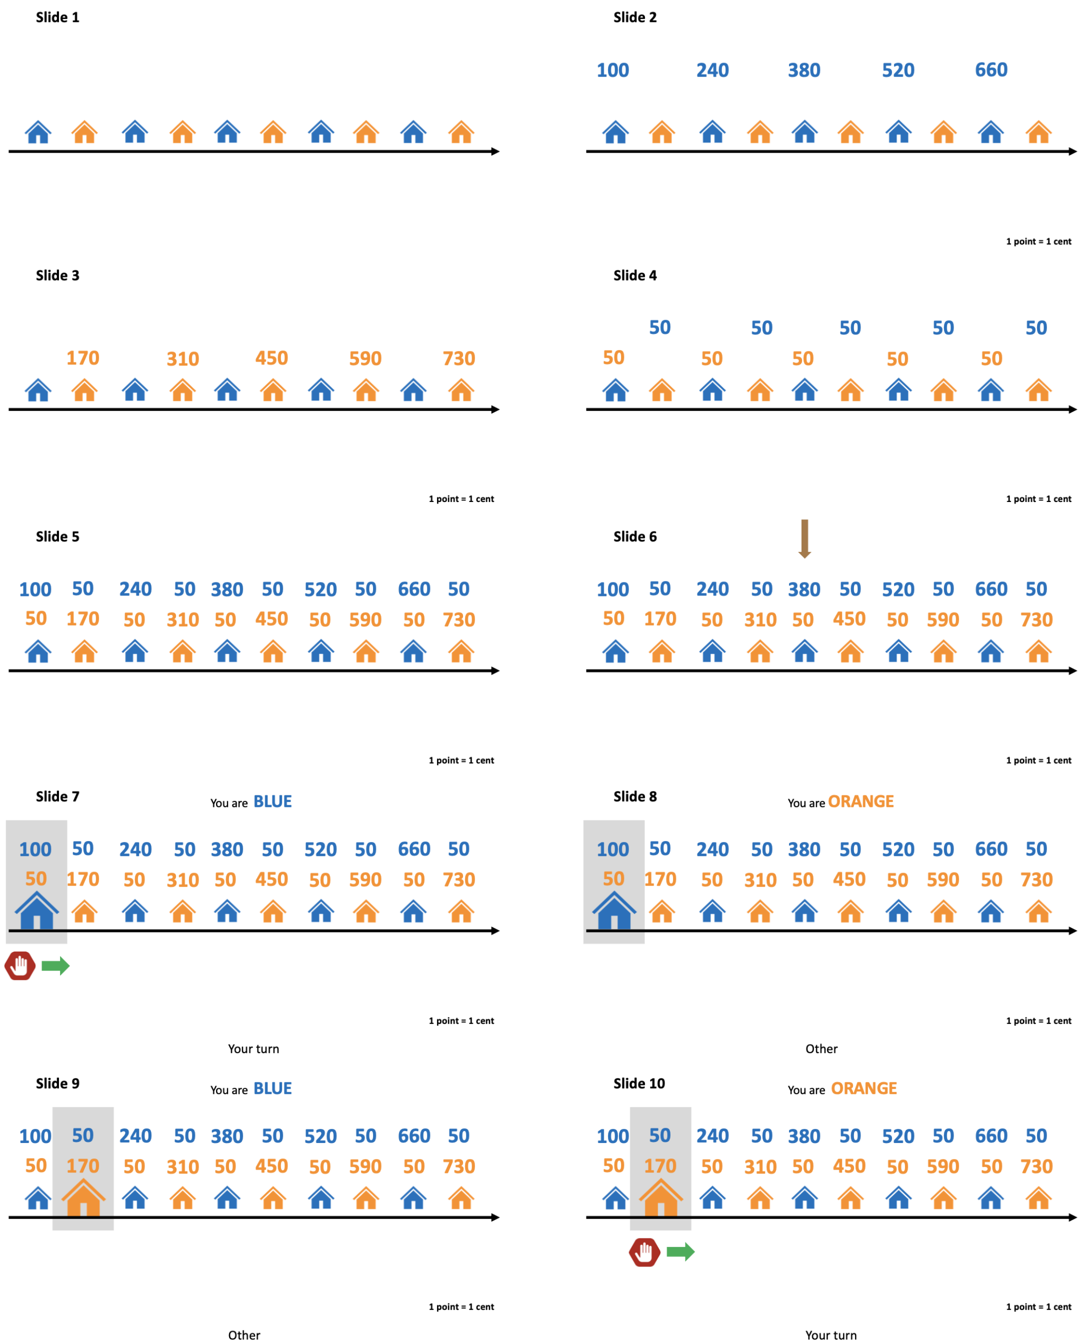


**Figure 6:** Slides projected on screen for instructions