Reading Minds to Win: The Power of Cognitive and Affective Skills in Children's Strategic Play *

Antonio Alfonso
U Loyola

Pablo Brañas Garza
U Loyola

Isabelle Brocas
USC, IAST and CEPR

Juan D. Carrillo
USC, IAST and CEPR

Juan B. González
USC

María José Vázquez
Fundación Etea

January 2025

Abstract

Do children use private information to their own advantage? Is this ability related to emotional intelligence? To answer these questions, we conduct a large lab-in-the-field experiment with 1662 participants from 8 to 18 years old who play a game with two-sided private information. We show that participants of all ages understand the fundamental relationship between action and private information. The ability to select payoff-enhancing strategies steadily increases with age but the capacity to recognize subtle variations in incentives triggered by changes in game structure remains elusive even for individuals at their peak cognitive capacity. Remarkably, participants of all ages who have heightened emotional intelligence exhibit a greater tendency to anticipate the behavior of others, best respond to them and, consequently, achieve higher payoffs. The paper thus reveals a strong, robust connection between affective and cognitive-theory-of-mind in young populations. It also highlights the importance of empathic skills for decision making.

Keywords: developmental decision-making; private information; rationality; theory-of-mind.

JEL Classification: C91, C93, D82.

* We thank the members of the Los Angeles Behavioral Economics Laboratory (LABEL) and the Loyola Behavioral Lab for their insights and to Chris Crabbe for his outstanding programming skills. Special thanks to Pablo Montero, Mónica Vasco and Javier González for excellent research assistance. We are grateful to the staff of Colegios Salesianos –in particular Andrés Corbacho, Manuel Redondo and Mónica Solís– for their support in running the experiment in their schools. The study was conducted with the University of Southern California IRB approval UP-12-00528. We acknowledge the financial support of the National Science Foundation grant SES-2315770, the Spanish Ministry projects PID2021-126892NB-100 and PDC2022-133230-I00, and the European Project Horizon GAP-10-1095175.

1 Introduction

Scenarios involving private information are challenging. Laboratory and field experiments indicate that not only students, but also seasoned investors and experienced entrepreneurs fall prey to the winner's curse in common value auctions (Thaler, 1988; Kagel and Levin, 2002). Similar findings apply to bargaining (Grosskopf et al., 2007; Bereby-Meyer and Grosskopf, 2008) and asset trading (Carrillo and Palfrey, 2011). Although progress has been made in understanding the causes of deviations (Charness and Levin, 2009; Brocas et al., 2014) our knowledge is still incomplete.

In this paper, we address two issues. First, we study the developmental trajectory in the ability to optimally incorporate private information in our decisions. We consider a broad age spectrum and a large number of individuals (1662 participants ranging from 8 to 18 years old), and determine which fundamental aspects about the strategic use of private information are innate, learned with age, or remain undeveloped. By doing so, we seek to gain deeper insights into the challenges that individuals face regarding information asymmetries, even after reaching intellectual maturity. Second, we study the joint evolution of strategic thinking ability (cognitive theory-of-mind or c-ToM) and emotional intelligence (affective theory-of-mind or a-ToM). The goal is to determine if individuals with a heightened ability to recognize and empathize with the emotions of others also excel in deciphering their intentions, thereby achieving superior outcomes.

The game we study is a variation of the 'compromise game' introduced by Carrillo and Palfrey (2009) (from now on, [CP]), which can be illustrated with the following stylized example. Two parties engaged in a dispute privately know the strength of their case. They can settle the dispute and obtain intermediate payoffs. Or, they can go to litigation, where both strengths are revealed. The party with the stronger case wins and receives a high payoff, while the party with the weaker case loses and receives a low payoff. We consider four treatments that share the basic structure but differ in the way outcomes and payoffs are computed.[1]

This game has two highly desirable features compared to most games with two-sided private information. On the one hand, it is extremely simple to explain, even to an 8 years-old child. On the other hand, analyzing strategic thinking in this game is not limited to

[1] In section 2, we discuss several applications that align with the fundamental characteristics of each treatment and establish the equilibrium for each case. Naturally, "litigation" is not a wording presented to our young subjects.

understanding whether participants play at equilibrium or best respond to the empirical behavior of others. In fact, choices in this game can be classified into three nested degrees of rationality. This allows us to determine which ability is mastered at each age, thus providing a more nuanced view of the development in the ability to think strategically.

The most basic degree of rationality, called *choice monotonicity*, is computed at the aggregate level. It posits that the likelihood of litigation in a group increases with the strength of the case. The next degree, called *strategic consistency*, represents a refinement at the individual level of the first degree. It stipulates that individuals should employ "cutpoint" strategies: settle if the strength is below a certain value and litigate if the strength is above it. These two characteristics–choice monotonicity and strategic consistency–hold across all treatments and highlight the relationship between information and action in the decision-making process. The third and most sophisticated degree of rationality, called *environmental variability*, prescribes that cutpoints vary across treatments. This variation arises because the equilibrium and best response cutpoints are influenced by treatment-specific incentives of players, both exogenous (how payoffs are determined) and endogenous (how a rival is expected to employ their private information in that particular treatment). Achieving this degree of rationality is cognitively demanding because it necessitates an optimal response to subtle shifts in incentives and beliefs, and it can only be examined through our innovative experimental design with multiple treatments.

The experiment provides interesting insights into age-related evolution in behavior. We show that participants exhibit choice monotonicity from an early age: subjects in all grades align their choices with the principle that higher values warrant choosing the litigation option. The study also finds that strategic consistency is prominent even among the youngest participants (one-half of participants aged 8 to 10 employ cutpoint strategies). This tendency grows steadily with age, even though participants never reach full compliance (three-quarters of participants aged 16 to 18 employ cutpoint strategies). Finally, the study reveals that environmental variability is missing across all ages. This absence persists even when examining only the subset of the oldest participants who employ cutpoint strategies. In other words, the nuanced differences in incentives stemming from subtle alterations in the game structure are not recognized by even the most sophisticated players in the study. In sum, the experiment sheds light on the developmental shift in the ability to engage in strategic thinking. While there appears to be an inherent human capacity to grasp the fundamental importance of conditioning choices on private information, the skill of a step shift in behavior based on information is acquired during childhood and adolescence. The additional and more complex step of conditioning on subtle external factors is a cognitive challenge that even young adults find it difficult to master.

We then perform an emotional intelligence test and uncover a strong and robust correlation between cognitive and affective theory-of-mind. Controlling for age, individuals who are best at reading the emotions of others also perform best in our game: they react more to their information, they use cutpoint strategies more frequently, they are more likely to best respond to the strategies of others and they obtain significantly higher payoffs. This finding suggests that the capacity to physically perceive and interpret emotions is closely intertwined with the ability to mentally decipher intentions. Connecting the two dimensions of theory of mind is an important finding per se for behavioral sciences, as these aspects are frequently perceived as distinct abilities. Our finding also implies that investigations into strategic behavior should not exclusively focus on associations with cognitive capabilities but should also explore the emotional processes that steer our actions.

Our paper is related to two strands of the literature. It contributes to the growing research on decision-making in children and adolescents. Prior game theoretic experiments have predominantly focused on games with complete information. Exceptions include Murnighan and Saxon (1998) who study an ultimatum game where the receiver privately knows the size of the pie and Sher et al. (2014) who analyze a hide-and-seek game where the hider knows the location of the prize and sends a signal to the seeker. In contrast, our study aims to investigate how age-related changes manifest in more complex games where asymmetric information is two-sided and individuals must uncover the relationship between the rival's action and their information.

Our paper also contributes to the growing body of research examining the relationship between fundamental abilities and performance in strategic contexts. Prior studies have primarily emphasized cognitive abilities, such as fluid intelligence (Brañas-Garza et al., 2012; Proto et al., 2019, 2022), and, to a lesser extent, cognitive theory of mind (Fe et al., 2022), highlighting positive associations between these constructs. Similarly, research has established links between cognitive theory of mind and abstract reasoning in children within educational settings (Alan and Turkum, 2024). We extend this literature by shifting the focus to the relationship between strategic abilities and *affective* components. Emerging evidence suggests that affective processes—how individuals experience, regulate, and respond to emotional stimuli—play a significant role in shaping behavior

(Naqvi et al., 2006; Lerner et al., 2015). Moreover, studies indicate that affective and cognitive processes are supported by overlapping neural mechanisms, underscoring their interconnectedness (Corradi-Dell'Acqua et al., 2014). Despite this, affective components have largely been examined in the context of prosocial behavior or the development of socio-emotional skills (Alan et al., 2021). In contrast, our study directly tests the hypothesis that affective mechanisms, alongside cognitive abilities, are critical for shaping decision-making outcomes in strategic environments.

The paper is organized as follows. Section 2 presents the game and equilibrium. Section 3 describes the experiment and discusses the hypotheses. In section 4, we study the aggregate choices of our population. Section 5 reports an analysis of individual strategies, comparing them across ages, treatments and decisions. In section 6, we discuss the relationship between affective and cognitive theory-of-mind. Section 7 discusses cursed equilibrium in the context of our game. Section 8 gathers concluding remarks.

2 The game

2.1 Treatments

We study several variations of [CP], a two-player game with two-sided private information. Players 1 and 2 are indexed by i and j. Player i has private value x_i drawn from distribution $F_i(x_i \,|\, x_j)$ with support $[\underline{x}, \overline{x}]$. To prevent ties, we also impose $x_i \neq x_j$. Player i chooses action $a_i \in \{p, s\}$, where p stands for "pass" and s stands for "see" in the terminology of the experiment. Final utilities depend on the action and information of both players $u_i(a_i, a_j; x_i, x_j)$. We consider a benchmark treatment and three variants.

Benchmark.

T1. Players move simultaneously. If both 'pass' (p), the outcome is 'compromise' (e.g., settlement), and each player obtains a medium payoff. If at least one player chooses 'see' (s), the outcome is 'competition' (e.g., litigation): the player with highest value obtains a high payoff and the player with lowest value obtains a low payoff. Formally:

$$u_i(p,p) = m \quad \text{and} \quad u_i(a_i,a_j) = \begin{cases} h & \text{if} \quad x_i > x_j \\ l & \text{if} \quad x_i < x_j \end{cases} \quad \text{whenever } (a_i,a_j) \neq (p,p)$$

with $l < m < h$ and $\underline{x} < m < \overline{x}$.

Variants.

T2. Identical to **T1** except that if the outcome is 'competition', then the player i with highest value obtains x_i instead of h. Formally:

$$u_i(p,p) = m \quad \text{and} \quad u_i(a_i,a_j) = \begin{cases} x_i & \text{if} \quad x_i > x_j \\ l & \text{if} \quad x_i < x_j \end{cases} \text{ whenever } (a_i,a_j) \neq (p,p)$$

T3. Identical to **T1** except that the outcome is 'compromise' unless both players choose s. Formally:

$$u_i(a_i,a_j) = m \text{ whenever } (a_i,a_j) \neq (s,s) \quad \text{and} \quad u_i(s,s) = \begin{cases} h & \text{if} \quad x_i > x_j \\ l & \text{if} \quad x_i < x_j \end{cases}$$

T4. Identical to **T1** except that players move sequentially.

T1 and **T4** are the two variants studied in [CP]. They capture in a stylized way a large range of situations, including war, litigation and market competition: nations, individuals and firms can coexist peacefully, reach a settlement and share the market (both choose p) or, instead, go to war, litigate and race to innovate (at least one chooses s). The likelihood of obtaining the larger reward through competition hinges on factors that are private information, such as the nations' military capacity, the strength of the legal cases, and the quality of firms. Furthermore, the compromise option requires mutual consent, while the competitive option can be decided unilaterally, since it takes only one player to initiate a conflict, a legal dispute, or an innovation race.

However, and as explored in the new variant **T2**, utilities may also depend on private values. A solid litigation case may increase both the probability of winning and the damages awarded to the plaintiff. A top quality research department enhances the likelihood of innovation and, if successful, also boosts profits. Possessing exceptional talent increases both the probability of securing a promotion and the salary in the new role, etc.

Finally, and as modeled in the second new variant **T3**, compromise is sometimes the default outcome which is overruled only if both players choose s. For example, in an electoral campaign, a public debate is often used to reveal the relative quality of the candidates, and therefore their chances of winning the election. However, a debate occurs only if both contenders accept to participate. Analogously, in financial settings, a trade occurs only if both parties agree.

2.2 Equilibrium

Despite the simplicity of the game, the equilibrium solution is non-trivial. Proposition 1 summarizes the properties of the equilibrium in each variant **Tk** (k $\in \{1, 2, 3, 4\}$).

Proposition 1. *The optimal strategy in Treatment* **Tk** *always involves a cutpoint:* $x^*_{\mathbf{Tk}} \in [\underline{x}, \overline{x}]$ *such that* $a^*_i(x_i) = p$ *if* $x_i < x^*_{\mathbf{Tk}}$ *and* $a^*_i(x_i) = s$ *if* $x_i \geq x^*_{\mathbf{Tk}}$.

The equilibrium cutpoint of both players varies across treatments: (i) $x^*_{\mathbf{T1}} = x^*_{\mathbf{T4}} = \underline{x}$, *(ii)* $x^*_{\mathbf{T2}} = m$, *and (iii)* $x^*_{\mathbf{T3}} = \overline{x}$.

Proof. See Appendix B1. □

In all treatments, it is always optimal to best respond with a cutpoint strategy. The reason is that a participant's expected payoff increases more steeply with their private value if they choose s than if they choose p.

However, subtle differences across treatments have dramatic consequences for equilibrium behavior. Indeed, depending on the rival's choice, the action of a player is sometimes inconsequential. Therefore, players must condition on the choice of the rival to compute their optimal cutpoint. Because their action is relevant in different cases, the optimal strategy depends on treatment specific features. More precisely, the equilibrium cutpoint unravels downwards to $\underline{x}$ in **T1** and **T4**, the treatments studied in [CP]. The equilibrium cutpoint unravels downwards but only to m in **T2** and it unravels upwards to $\overline{x}$ in **T3**. Studying these different treatments allows us to determine whether individuals react to elements of the environment.

Notice that, just like in the guessing game (Nagel, 1995), the equilibrium involves playing a weakly dominated strategy (it is therefore a best response only to an individual who also plays at equilibrium). Hence, we do not expect even the most sophisticated participants to employ the equilibrium strategy. On the other hand, and as we will develop below, we do expect sophisticated players to adapt their behavior to the treatment.

3 The experiment

3.1 Design

The paper studies the choices of children and teenagers in discrete versions of the four treatments described in section 2.1. Working with young participants presents some methodological challenges. We follow the guidelines developed in Brocas and Carrillo (2020).

Game. For simplicity, we modify the game outlined in section 2.1. Participants are matched in pairs, and each pair is informed that there is a deck of 10 cards, numbered from 1 to 10. The computer assigns one card to each player, so that types are correlated and ties are not possible. Participants only see their own card and must choose between "pass" or "see." This approach offers two main advantages compared to random and independent draws (as, for example, in [CP]): it eliminates the possibility of ties and it is straightforward to explain, suppressing the need to delve into distributional considerations.[2]

We employ the strategy method to elicit behavior. It is well-known that the strategy method may be more challenging to understand and it may elicit different responses compared to the direct-response method (Brandts and Charness, 2011). In our case, we employ an elicitation strategy similar to the multiple price list method often adopted in the literature to elicit risk preferences, time preferences and willingness to pay. More precisely, we tell our participants: "You are going to get a card. What do you want to do if you get a 1? If you get a 2? ..." and so on. Questions are ordered from lowest to highest value, and we do not impose any restriction on their strategy. This method allows us to observe the strategy employed by individuals –for example, whether it involves using a cutpoint (zero or one switch) or not (two or more switches)– instead of having to infer the strategy from the choice. It is especially important if, as in our case, (i) we believe that a non-negligible fraction of participants may not use a cutpoint strategy, (ii) we want to learn what these other strategies are, and (iii) we want to examine age-related changes in the fraction of individuals who use cutpoint strategies.[3]

We study the four treatments described in section 2.1 using a between-subject design. We consider integer private values $x_i \in \{1, 2, ..., 10\}$, and point payoffs $h = 10$, $m = 5$ and $l = 1$. Adapting Proposition 1 to our discrete case with no ties, the equilibrium cutpoint in each treatment is: $x^*_{\mathbf{T1}} = x^*_{\mathbf{T4}} = 2$, $x^*_{\mathbf{T2}} = 6$ and $x^*_{\mathbf{T3}} = 10$.[4,5]

One strength of the design is that the presentation, instructions, cognitive require-

[2]As discussed in Proposition 1, the equilibrium is unchanged if a participant misunderstands the instructions and believes that the two cards are drawn with replacement.

[3]Strategies could be estimated from choices using MLE methods. However, it would require large amounts of data (hence, many choices) and an ex-ante knowledge of the set of possible strategies.

[4]With integer values of x and no ties, $x^*_{\mathbf{T1}} = x^*_{\mathbf{T4}} = 1$, $x^*_{\mathbf{T2}} = 5$ and $a_{\mathbf{T3}}(x) = p$ for all x are also equilibria of the game, although they do not survive the Trembling Hand refinement. These details are important for theory but have limited consequences for our experimental analysis.

[5]We deliberately chose values such that $2m < h + l$. It implies that colluding on playing 'pass' all the time, decreases rather than increases the expected payoff of players.

ments, and logical inferences are virtually indistinguishable across **T1**, **T2** and **T3**. Therefore, it is difficult to argue that differences in behavior among these treatments stem from differences in understanding the rules. Conversely, **T4** is cognitively more demanding since we elicit two strategies –as the first and second mover, hereafter referred to as **T4F** and **T4S**– and the latter is contingent on the first mover selecting a specific action (pass). Therefore, it is prudent to approach comparisons between **T1**, **T4F** and **T4S** with caution.

Finally, we elicit player's strategies twice, called the first and second decisions. Initially, we request them to report a strategy to be employed in six consecutive matches against six distinct opponents. We utilize the random draws (x_i, x_j) and present for each of the six games the cards drawn by both players, the choices given their strategies, and the payoffs given the rules of the treatment.[6] We instruct participants to review and comprehend the table, after which they select a new strategy to be used in six new matches against six new opponents. Figure 1 presents screenshots of the strategy elicitation (left) and the outcomes of the first six games (right) for **T1** (with the text translated from the Spanish version). In Appendix A1, we provide the full set of instructions.

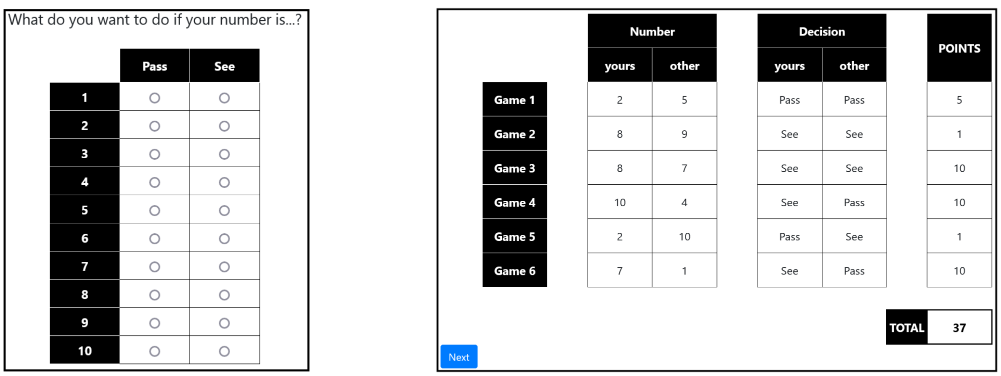

Figure 1: Screenshots. Strategy elicitation (left) and summary of outcomes (right) in **T1**

Theory of mind. We conduct a novel, child-friendly version of an affective Theory-of-mind (a-ToM) test. In our a-ToM task, participants observe images of a person's eyes and are tasked with selecting the most appropriate adjective from four options to describe the conveyed emotion. We then examine the correlation between performance in the affective and cognitive tasks. Further details about the task are presented in section 6.

[6]The strategies of other players are not shown (only their actions given the values drawn). This way, we prevent participants from blindly mimicking someone else's strategy.

Timing. The experiment consists of four tasks programmed in 'oTree' (Chen et al., 2016) and implemented on PC tablets. Due to the challenge of accessing such a large population of children, we combined two projects. We started with a third-party dictator game, the results of which are analyzed in (Brañas-Garza et al., 2024). After a break, we administered the game under study here, followed by the affective theory-of-mind task. We finished with a Big Five personality questionnaire, which was designed to complement the dictator game. All tasks were incentivized, except for the personality questionnaire. While the two games exhibit sufficient differences to mitigate concerns about cross-contamination, we took precautionary measures to ensure experimental integrity. We always conducted the dictator game first followed by the game under investigation, and employed random and anonymous subject re-matching between the two. Most importantly, we did not announce the results of the dictator game until the end of the experiment.

Population and procedures A strength of this study lies in its extensive age range and large sample size, particularly noteworthy for a lab-in-the-field experiment. We recruited 1662 school-age students from low- to middle-income families, covering the Spanish equivalent of grades 3 to 12 in the US educational system. We conducted this study across four schools situated in four distinct cities in the south of Spain. All schools are part of the Salesianos, a private network of catholic schools.[7] Given that the oldest participants are young adults close to their prime cognitive ability (18 years old), we decided not to include a an undergraduate control group, who would only be slightly older (18 to 22 years old) and possess a distinct background. We counterbalanced the assignment of participants to treatments. Table 1 summarizes the participants by grade (age) and treatment.

Grade	3	4	5	6	7	8	9	10	11	12	Total
Age	8-9	9-10	10-11	11-12	12-13	13-14	14-15	15-16	16-17	17-18	
T1	47	34	45	24	55	30	50	25	23	43	376
T2	43	33	25	44	27	49	63	41	52	20	397
T3	32	43	49	48	59	52	23	46	29	57	438
T4	35	47	43	44	59	57	39	48	53	26	451
Total	157	157	162	160	200	188	175	160	157	146	1662

Table 1: Number of participants by grade (age) and treatment

[7] The Salesianos is a missionary school system with presence in 132 countries. It is heavily subsidized by the Spanish government and provides free education. Schools are typically located in low to middle income neighborhoods, and they are particularly attractive to families with low to middle SES. This network has 59 schools in Spain alone serving almost 40,000 students (see `https://www.salesianos.es/escuelas`).

We visited each school and conducted the experiment one class at a time, using a portable lab. Participants were organized into subgroups of 8 to 12 individuals, each led by one experimenter. The experiment was integrated into the school's academic activities, and our Institutional Review Board stipulated an opt-out procedure. Importantly, no children opted out, thereby eliminating any potential selection bias. Only children who were absent on the day of the experiment did not participate. The overall participation rate stood at 92%. For visual clarity and to enhance statistical power, we grouped two consecutive grades together, which we refer to as "grade groups" (e.g., 3-4, 5-6, etc.). For the regression analysis, we use instead the age (in months) of the participant.

Payments and duration. Finding a single medium of payment that can provide comparable incentives to participants of varying ages can be challenging. As outlined in Brocas and Carrillo (2020), an appropriate approach may involve distinct payment mediums tailored to different grades. However, for this particular experiment, we capitalized on the presence of integrated cafeterias and the unique opportunity to use food as an accepted form of payment (an option forbidden in most schools). We introduced our own currency that participants could use in the cafeteria to purchase food, drinks, and snacks of their choice. The cafeteria offers affordable options (e.g., a large steak sandwich for 1.80€), is highly popular, and widely accessible during recess and lunchtime. Our currency had no expiration date, although the majority of participants spent their endowments within a week. The tasks examined in this paper had a duration of approximately 25 minutes, and the entire experiment never exceeded the length of one standard school period (50 minutes). Participants earned 0.05€ for each point acquired in this game. The average payoff was 3.23€ in the game studied here, 1.16€ in the a-ToM task and 4.00€ in the dictator game, for a total of 8.39€.

3.2 Why this experiment?

This study leverages three novel features to examine behavior in games with two-sided private information. Here, we address how each feature provides unique insights.

Multiple treatments. [CP] demonstrates that the behavior of adults is consistent with cutpoint strategies. While it rarely adheres to equilibrium, a significant fraction of strategies are close to best responding to the rival's behavior. Since [CP] focuses on two treatments that predict the same equilibrium (**T1** and **T4**), it is not possible to determine whether participants realize that they must condition only on the cases where their action is relevant in order to determine their optimal strategy. To investigate this phenomenon, it is crucial to compare across treatments that exhibit different unraveling properties and therefore different equilibria. By introducing **T2** and **T3**, we can evaluate if participants are responsive to the unique incentives of each treatment and adjust their strategies in the direction predicted both by theory and by empirical payoff maximization.

Developmental approach. Games of strategy represent intricate cognitive tasks that rely on a combination of logical and social skills. These foundational abilities undergo gradual and distinct processes during childhood and adolescence. Therefore, the developmental path of behavior in strategic games is influenced by the timing of these skill developments. Among games of strategy, incomplete information is arguably the characteristic most closely related to cognitive theory-of-mind (c-ToM), defined as the ability (i) to recognize that other people have distinct beliefs and (ii) to infer and predict behaviors based on those beliefs. Taking a developmental perspective in a game characterized by hierarchical degrees of rationality allows us to discern which facets of c-ToM are mastered at different ages, and which abilities never fully mature.

Affective mechanisms. Having established the importance of c-ToM in games with two-sided private information, we ponder if this skill is linked to affective theory of mind (a-ToM), the ability to recognize and empathize with the emotions of others. By correlating performance between our game and an emotional intelligence task, we ascertain whether and to what extent affective faculties also contribute to performance in complex strategic environments.

3.3 Hypotheses

A notable feature of our game is that behavior can be categorized into hierarchical degrees of rationality. At its fundamental level lies choice monotonicity: in all treatments, the likelihood of playing "see" in a grade group should increase as the private value rises. In an intermediate level is strategic consistency: individuals should employ cutpoint strategies, again in all treatments. Finally, the highest level of sophistication prescribes environmental variability: the optimal cutpoint should be contingent on the game's structure. Given the large range of ages in our sample, we have the opportunity to identify the age at which participants' behavior aligns with each degree of rationality. We put forth several hypotheses concerning the age-related shifts in behavior.

Hypothesis 1. *'Choice monotonicity' is satisfied in the entire population: the aggregate probability of choosing 'see' is increasing in x in all grade groups.*

Hypothesis 2. *'Strategic consistency' improves with age. Also, the fraction of participants who use a cutpoint strategy is small (but significantly above zero) among the youngest and high (but significantly below one) among the oldest.*

Hypothesis 3. *'Environmental variability' is satisfied by all participants who adhere to strategic consistency. Formally, the empirical distributions of cutpoints in the simultaneous treatments can be ranked according to theory: $F(x^*_{\mathbf{T3}})$ F.O.S.D. $F(x^*_{\mathbf{T2}})$ F.O.S.D. $F(x^*_{\mathbf{T1}})$.*

According to Hypothesis 1, we anticipate that all participants, including the youngest ones, will grasp the fundamental principles of the game and choose "see" more frequently as their value x_i increases. If this hypothesis holds true, it suggests that the foundational yet non-trivial principles of the game, linking actions to information, can be understood at an early age. Hypothesis 2 implies that young children may struggle to recognize the optimality of a cutpoint strategy. As they develop, the prevalence of this crucial aspect of optimal decision-making is expected to increase, although this improvement might be gradual and incomplete. In essence, we expect this property to be neither insurmountable for our younger participants nor trivial for our older ones. In accordance with Hypothesis 3, our more cognitively developed participants (those employing cutpoint strategies) are likely to grasp the intricacies of the treatment assigned to them. Consequently, we should observe a shift in cutpoints across treatments as they adapt to each scenario.

In addition to examining the developmental aspects of decision-making, our experiment also explores the interplay between cognitive and affective theory of mind.

Hypothesis 4. *Participants with higher affective theory of mind understand the strategic aspects of the game better: they use cutpoint strategies more often and reach higher payoffs.*

Hypothesis 4 posits that the ability to read others' emotions is connected to the ability to infer their intentions, a pivotal aspect of strategic decision-making. A more detailed motivation for this hypothesis is provided in section 6.

4 Aggregate analysis

In our initial data analysis, we report for each treatment and each grade group the average probability of seeking the competition outcome (choose s) as a function of the private value x ($\in \{1, ..., 10\}$). We denote this probability $S(x)$. Results are presented in Figure 2.

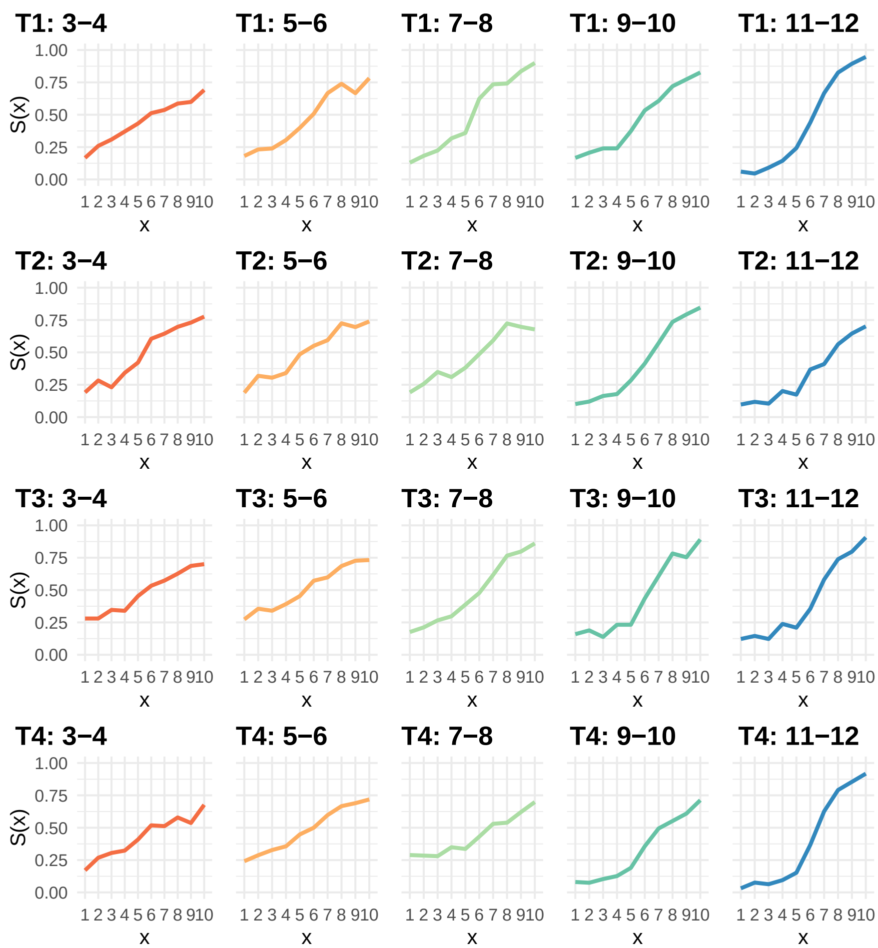

Figure 2: Average probability of choosing to see ($S(x)$) by treatment and grade group

Strongly supporting Hypothesis 1, we observe a monotonic increase in the probability of choosing s as the private value increases. This trend holds true across all treatments and grade groups. Specifically, when $x = 1$, the probability ($S(1)$) tends to hover around or fall below 0.25, while at $x = 10$, it typically reaches or exceeds 0.75. The increase with

x is often steep, although this pattern does exhibit some variation across grade groups, as we will elaborate on shortly. Overall, despite the young age of some participants and the complexity of contingent reasoning, a majority of them appear to understand the fundamental relationship between own information and own action.

To study in more detail the sensitivity of s to private value across different grade groups, we perform Probit regressions. Our unit of observation is the participant's action (see = 1). We conduct the regressions separately for each treatment. In our regressions, the independent variables are the private value (x), the participant's age in months (Age), an interaction term between the two included in some regressions ($x \times Age$), a dummy indicating gender ($Male$ = 1) and a dummy representing the second decision ($2ndDec$ = 1). For **T4**, we also include a dummy to account for being the second mover ($T4S$ = 1). Unless otherwise noted, all regressions incorporate school fixed effects and standard errors are clustered at the individual level. Results are presented in Table 2.

	T1		**T2**		**T3**		**T4**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
x	0.240***	-0.050	0.213***	0.080	0.212***	-0.053	0.181***	-0.126***
	(0.013)	(0.056)	(0.012)	(0.053)	(0.011)	(0.050)	(0.009)	(0.041)
Age	0.004	-0.010***	-0.004***	-0.009***	-0.001*	-0.011***	-0.003***	-0.014***
	(0.0009)	(0.002)	(0.0008)	(0.002)	(0.0008)	(0.002)	(0.0008)	(0.002)
$x \times Age$	—	0.002***	—	0.0008**	—	0.002***	—	0.002***
		(0.0004)		(0.0003)		(0.0003)		(0.0003)
Male	0.081	0.085*	0.033	0.033	-0.042	-0.043	-0.007	0.003
	(0.050)	(0.050)	(0.051)	(0.051)	(0.046)	(0.047)	(0.044)	(0.045)
2ndDec.	0.082***	0.087**	0.075***	0.076***	0.086***	0.090***	0.042**	0.044**
	(0.030)	(0.031)	(0.026)	(0.026)	(0.028)	(0.029)	(0.021)	(0.022)
T4S	—	—	—	—	—	—	0.044**	0.046**
							(0.021)	(0.022)
Obs.	7,520	7,520	7,940	7,940	8,760	8,760	18,040	18,040
Pseudo R^2	0.185	0.197	0.157	0.159	0.152	0.163	0.121	0.134

Significance levels: *** = 0.01, ** = 0.05, * = 0.1

Table 2: Probit regression of decision to choose s by treatment

The regressions results validate the patterns observed in Figure 2. From the regressions without interaction terms (columns (1), (3), (5), (7)), we notice that as the private value x increases, the likelihood of playing 'see' also increases across all treatments. On the other hand, age by itself has no significant effect. Adding the interaction term (columns (2), (4), (6), (8)) reveals that the absence of an age effect is due to the fact that older participants are less likely to choose s for the lowest values of x but they are more reactive to an increase in that parameter (adding a quadratic age term did not improve the fit). The increased sensitivity to changes in x is highly statistically significant across all treatments (albeit slightly smaller in magnitude in **T2**). In other words, older players demonstrate a better grasp of the empirical advantages of choosing s for high values of x and p for low values of x. Additionally, the likelihood of selecting 'see' is marginally higher among males in **T1**, among participants making their second decision in all treatments, and among second movers in the sequential treatment. Finally, Table 2 reveals an increased inclination to choose 'see' in the second decision (*2ndDec*).

In Appendix B2, we investigate the changes between first and second decision. The analysis confirms only a modest increase in s. Importantly, changes are not the result of participants learning how to best respond or play at equilibrium, as they often contradict both predictions. Indeed, the increase in s is most noticeable in **T3** (where theory predicts the opposite) and for low values of x (where it is empirically unlikely to be advantageous). Given the absence of systematic differences, we will combine data from both decisions for the remainder of the paper, unless otherwise specified. In Appendix B3, we further study the sensitivity of actions to private values and confirm that older players change their behavior more drastically than their younger peers.

The findings of this section are summarized as follows.

Result 1. *Hypothesis 1 is supported by the data. Choice monotonicity ($S(x)$ increasing in x) is satisfied across all grade groups. At the same time, the sensitivity to private value is more pronounced among older individuals.*

5 Individual behavior

5.1 Individual strategies

We next examine the strategies adopted by participants at different ages. We categorize strategies into different classes and calculate the proportion of participants who adhere to each. To mitigate this risk of overlooking certain strategies, we contemplate a very large set of options, even including some unconventional strategies that may be suboptimal and/or unexpected. Table 3 provides a description of all the strategies considered.

This set of strategies encompasses a wide range of options spanning from strategies that satisfy strategic consistency (aP, aS, C_0) or almost satisfy it (C_1) to fully suboptimal

strategy	description
aP	Always pass ($a_i(x) = p$ for all x)
aS	Always see ($a_i(x) = s$ for all x)
C_0	(Interior) cutpoint ($a_i(1) = p$, $a_i(10) = s$ and exactly one switch)
C_1	(Interior) cutpoint with one "mistake" (it becomes C_0 if we reverse one choice)
RC_0	Reverse (interior) cutpoint ($a_i(1) = s$, $a_i(10) = p$ and exactly one switch)
RC_1	Reverse (interior) cutpoint with one "mistake"
INT_0	Three intervals ($a_i(1) = a_i(10)$ and exactly two switches)
INT_1	Three intervals with one "mistake"
ALT_0	Perfect alternation p and s ($a_i(x) \neq a_i(x+1)$)
ALT_1	Almost perfect alternation p and s (at least four separate groups of consecutive s)
PAT	Patterns that are recognizable but different from ALT_0 and ALT_1*
Other	Behavior not covered by any of the above strategies

* Patterns: (i) 2-2 alternation (e.g., $(p, p, s, s, p, p, s, s, p, p)$); (ii) 2-1 alternation (e.g., $(p, s, s, p, s, s, p, s, s, p)$); (iii) symmetric strategies below $x = 5$ and above $x = 5$ (e.g., $(p, s, s, p, p, p, p, s, s, p)$); and (iv) strategies that are repeated for x below 5 and x above 5 (e.g., $(p, p, s, s, s, p, p, s, s, s)$).

Table 3: Individual strategies

ones (RC_0, RC_1).[8] It also includes strategies that fall in between (INT_0, INT_1), as well as some unconventional or "pleasant" patterns (ALT_0, ALT_1, PAT) that cannot be readily reconciled with any reasonable theory.

We categorized behavior into these strategies, prioritizing classification in the order presented in Table 3 (for example, if a strategy is consistent with both C_1 and INT_0, we classify it as C_1). We found that only 1.9% of observations match RC_0 and RC_1, 2.0% match INT_0 and INT_1, and 1.5% match PAT, so we decided not to retain these strategies. For visual ease, we grouped ALT_0 and ALT_1 under a single category, called ALT. This leaves us with a total of six strategies: aP, aS, C_0, C_1, ALT and Other (for the remaining participants). Figure 3 reports the empirical frequency of each strategy separated by treatment, grade group, and decision.

The distribution of strategies is similar between **T1**, **T2** and **T3**, as shown in Figure 3(A), although a paired χ^2-test of differences of proportion finds a difference between **T2** and **T3**, $p < 0.01$. The sequential treatment is different than all the simultaneous ones (χ^2-test, $p < 0.01$), mainly due to a higher frequency of aP at the expense of C_0. By contrast, differences between **T4F** and **T4S** are not significant (χ^2-test, $p = 0.51$). More

[8] In our analysis, we take a cautious approach and refer to aP and aS as "boundary cutpoint strategies."

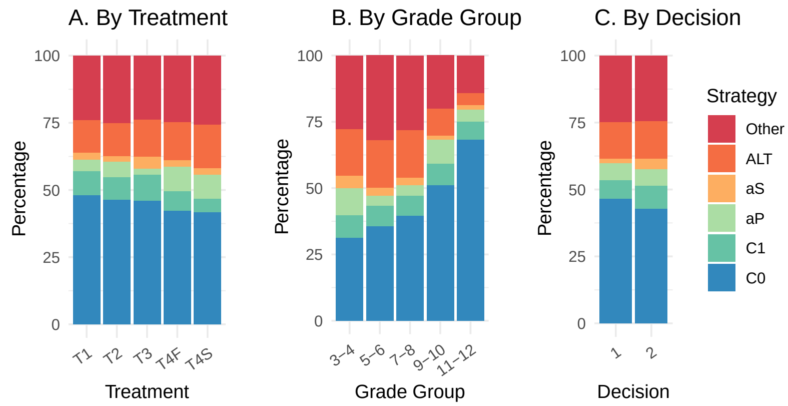

Figure 3: Individual strategies by treatment (A), grade group (B) and decision (C)

generally, let us call $C_+ = aP + aS + C_0$ the set of cutpoint strategies that perfectly align with an interior (C_0) or a boundary (aP + aS) cutpoint. More than half the observations in all treatments are classified as C_+. At the same time, there is also a significant fraction of alternation (12% to 16%) and unclassified (24% to 26%) strategies.

By contrast, strategies change very significantly across grade groups, as evident in Figure 3(B) (paired χ^2-tests, $p < 0.02$). In strong support of Hypothesis 2, strategic consistency increases steadily and very significantly with age. For example, C_0 changes from 31% in 3-4 to 68% in 11-12. At the same time, the level of strategic consistency is higher than initially anticipated. We hypothesized few cutpoint strategies among our youngest participants, and yet 46% of children in grades 3-4 align with C_+.[9] Finally, differences between first and second decision are very minor, albeit statistically significant (χ^2-test, $p < 0.01$), see Figure 3(C). We observe a slight increase in mistakes among interior cutpoint strategies (from C_0 to C_1) and also a small increase in aS.

5.2 Cutpoint strategies across ages and decisions

To further investigate the decision to use cutpoint strategies across ages, we perform a simple Probit regression for each treatment. In this regression, the dependent variable is whether the choice aligns with a perfect boundary or interior cutpoint strategy (C_+). The

[9] To put it in perspective, educated adults in much simpler settings (e.g., elicitation of risk preferences using multiple price lists) exhibit multiple switches 15%-20% of the time.

results are presented in Table 4.[10]

	T1 (1)	**T2** (2)	**T3** (3)	**T4** (4)
Age	0.005***	0.003	0.007***	0.010***
	(0.002)	(0.002)	(0.002)	(0.002)
Male	0.267**	0.214*	0.442***	0.064
	(0.116)	(0.115)	(0.108)	(0.102)
2ndDec.	-0.013	-0.241***	0.127**	-0.054
	(0.066)	(0.057)	(0.062)	(0.052)
T4S	—	—	—	-0.025
				(0.035)
Obs.	752	794	876	1,804
Pseudo R^2	0.043	0.030	0.056	0.089

Significance levels: *** = 0.01, ** = 0.05, * = 0.1

Table 4: Probit regression of adherence to cutpoint strategy C_+ by treatment

In support of Hypothesis 2 and in line with previous findings, we observe a rise in the likelihood of adopting a cutpoint strategy with age, although the effect is not statistically significant in **T2**. We notice that males are more likely to use cutpoint strategies. This effect of gender will be explored further in section 6.3. We also observe a difference between the first and second decisions. However, the sign changes across treatments, further supporting our observation that there are no consistent learning patterns.

Finally, in Appendix B4, we perform a Probit regression similar to the one in Table 2, where we divide the sample into individuals who use perfect cutpoint strategies and those who do not. We show that the sensitivity to information of individuals who use cutpoint strategies is high and relatively constant across age whereas that of individuals who do not employ cutpoint strategies is low and increasing with age. We also conduct a basic analysis of changes in strategies across decisions and show that strategic consistency is equally likely to increase than to decrease between the first and second decision.

We summarize the main finding of sections 5.1 and 5.2 as follows.

Result 2. *Hypothesis 2 is supported by the data. Strategic consistency increases monotonically with age. Nevertheless, the fraction of individuals who use cutpoint strategies is higher for all ages than initially anticipated.*

[10] We obtain very similar results both if we add C_1 to the set of cutpoint strategies, as well as if we only consider C_0 the interior cutpoint strategies.

5.3 Cutpoint strategies across treatments

We next focus on choices that align perfectly with interior or boundary cutpoint strategies (C_+), and hypothesize that participants adhering to strategic consistency will conform to environmental variability. In other words, the ordering of cutpoints across treatments predicted by theory will, in a stochastic sense, be preserved in our sample. To test this hypothesis, we construct the same function as in Figure 2 but only with players who employ cutpoint strategies C_+. We denote this function as $S_c(x)$. This time, $S_c(x+1) \geq S_c(x)$ *by construction*. Figure 4 presents the functions $S_c(x)$ for the simultaneous (top) and sequential (bottom) treatments, separated by grade group and combining all (All).

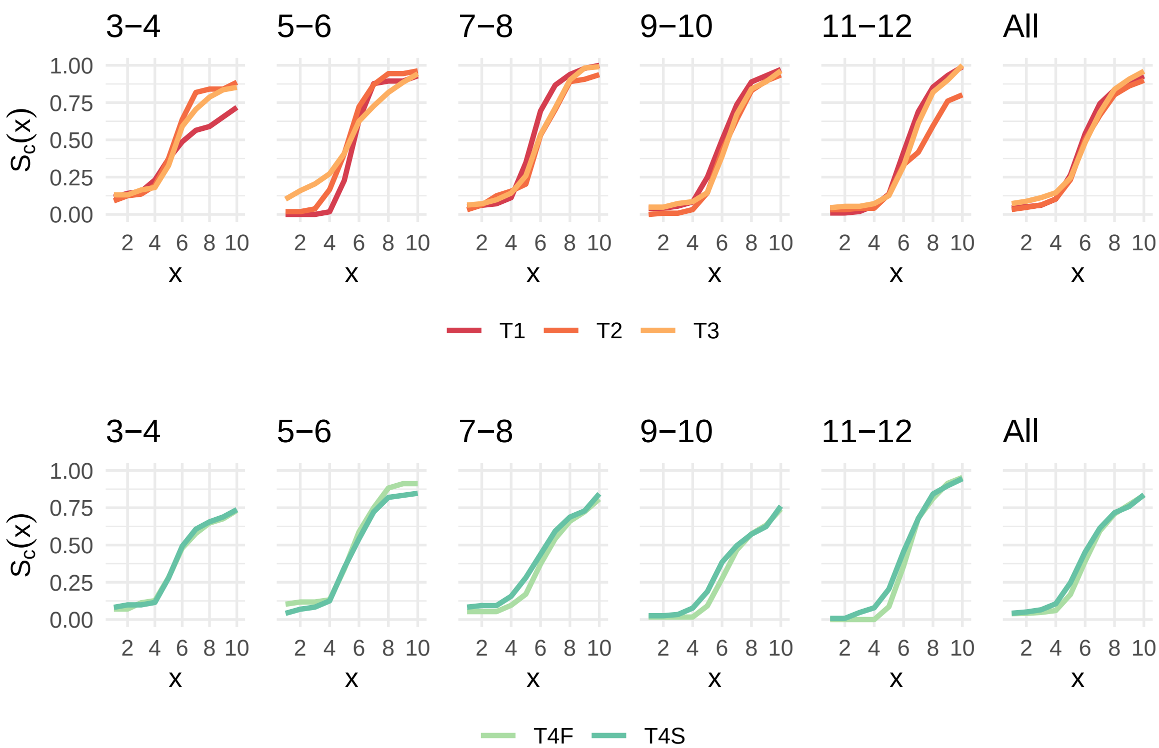

Figure 4: Comparison of $S_c(x)$ across treatments for each grade group

There is large heterogeneity in cutpoints. Only a few cutpoints are at the boundaries, with the majority (60% to 72%) falling between $x^* = 5$ and $x^* = 8$. Most importantly, we do not observe the differences we expected between **T1**, **T2** and **T3**. A Kolmogorov-

Smirnov (KS) test of differences between distributions reveals that the two supposedly most extreme treatments, **T1** and **T3**, are statistically different for grade group 3-4 (p = 0.04), with the sign in the opposite direction of the theoretical prediction. In grade group 11-12, 'see' is more frequent in **T1** than in **T2** as predicted by theory (KS-test, $p < 0.01$) but it is also more frequent in **T3** than in **T2**, contrary to the theory (KS-test, $p < 0.01$). All other differences are not significant.[11] Overall, Hypothesis 3 is strongly rejected: distributional differences in cutpoints between **T1**, **T2** and **T3** are small and inconsistent. The importance of conditioning on the rival's choice of p in **T1** and s in **T3**—the critical aspect to the downward and upward unraveling dynamics—is apparently not grasped by any group, including our oldest and arguably most sophisticated participants.

As for the sequential treatment, contrary to [CP], we do not find any significant differences between first and second movers in any grade group (KS-test, $p > 0.25$). However, [CP] argue that second movers choose s more often than first movers because they observe their rival's decision, make inferences about their information, and react accordingly. This reasoning is only possible under the direct response method. In the strategy method, both first and second movers engage in similar contingent reasoning, as they are asked to provide their strategies in advance. Consequently, it is reasonable that their decisions are also similar. Appendix B5 compares the choices of participants as first and second movers. It shows that changes are more frequent for younger participants and for intermediate private values. The results of this section are summarized as follows.

Result 3. *Hypothesis 3 is strongly rejected by the data. Environmental variability (which would manifest as distributional differences in cutpoints between* **T1***,* **T2** *and* **T3***) is not observed in our population, even in the subset of our oldest, strategically consistent players.*

5.4 Understanding deviations

5.4.1 Best response to empirical behavior

A potential explanation for the lack of environmental variability could be that the empirical best response to the behavior of rivals in the same grade group is the same across treatments. To study this possibility, we compute the cutpoint $\tilde{x}$ that maximizes the empirical expected payoff for each grade group and treatment (for **T4**, we determine the best

[11] If we include all participants (not only those satisfying strategic consistency), differences in $S(x)$ across treatments are still not significant (data omitted for brevity).

response of a first (second) mover to the empirical behavior of second (first) movers). We also determine the optimal cutpoint with pooled data (all). For comparison, we report the theoretical equilibrium (theory) computed in section 3. Results are compiled in Table 5.

Grade group	3-4	5-6	7-8	9-10	11-12	all	theory
T1	4	4	4	4	4	4	1-2
T2	7	7	7	7	7	7	5-6
T3	7	7	8	8	8	7	10
T4F	5	4	5	4	4	4	1-2
T4S	4	4	5	4	4	4	1-2

Table 5: Optimal empirical cutpoint $\tilde{x}$ by grade group and treatment

The empirical best response cutpoint $\tilde{x}$ is never at the boundary but it approximately preserves the same order across treatments as the equilibrium cutpoint x^*: $\tilde{x}_{\mathbf{T1}} < \tilde{x}_{\mathbf{T2}} \leq \tilde{x}_{\mathbf{T3}}$ and $\tilde{x}_{\mathbf{T1}} \simeq \tilde{x}_{\mathbf{T4F}} \simeq \tilde{x}_{\mathbf{T4S}}$, with the exception that the best response is similar in **T2** and **T3**. This is an important finding. It implies that the absence of significant differences in choices across treatments, especially between the polar cases **T1** and **T3**, cannot be attributed to individuals understanding the deviations of their less sophisticated peers and best responding to it. The result is also natural. Variations in treatment design (upward vs. downward unraveling) imply differences in best response strategies across treatments, even when the empirical behavior of rivals is similar.

5.4.2 Empirical loss

Another possible reason for the lack of environmental variability could be that deviations from the empirically optimal strategy result in losses that are insignificant. This is a critical issue because, if that were true, deviations could be attributed to an efficient conservation of cognitive effort. To address this issue, we calculate for each strategy type the revenue loss relative to selecting the empirically optimal cutpoint within their grade group, as determined in Table 5. Figure 5 illustrates this information for each treatment, with error bars representing 2 standard errors of the mean.

The largest losses are incurred when the individual chooses the boundary cutpoint opposite to the theoretical prediction: aP in **T1**, **T4F** and **T4S** and aS in **T3**. Both aP and aS are suboptimal when theory prescribes an intermediate cutpoint (**T2**). In general, individuals who do not use cutpoint strategies (ALT and Other) experience at least twice

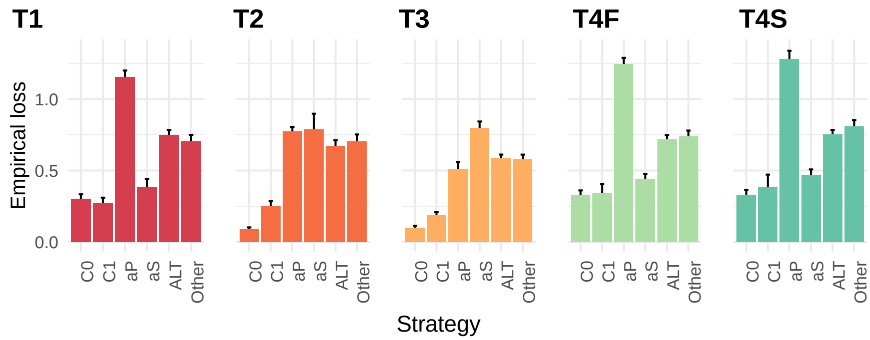

Figure 5: Empirical loss as a function of strategy

the revenue loss compared to those who utilize interior cutpoint strategies. The difference is similar across grades (data not reported for brevity). It is also substantial, especially considering that (i) the decision of a participant is often inconsequential due to the rival's behavior, and (ii) random behavior aligns with the optimal decision one-half of the time.

All in all, neglecting environmental variability results in significant losses and cannot be explained by a best response behavior to others: it is likely due to the cognitive difficulty of updating beliefs based on the rival's choice.

6 Cognitive and affective theory-of-mind

6.1 Measuring and evaluating theory-of-mind

"Theory of mind" encompasses the overarching idea of mentalizing about others. Cognitive theory of mind (c-ToM) centers on cognitive processes whereas affective theory of mind (a-ToM) involves empathizing with emotions, a concept closely tied emotional intelligence.

Traditionally, c-ToM tests have relied on false belief tasks (Wellman et al., 2001), which gauge an individual's ability to grasp that another person may hold beliefs that differ from reality.[12] Most children grasp the basic concept of false belief by age 4 or 5. To adapt the task to older participants, false belief tasks rely on complex narratives featuring numerous characters, objects, and locations that participants must track throughout the storyline (Kinderman et al., 1998; Valle et al., 2015). Consequently, these tasks necessitate

[12] In the simplest example, the participant is shown a person observing the placement of an object. When the person leaves the room, the object is displaced. The participant is then asked where they think the person believes the object is after coming back into the room (Wimmer and Perner, 1983).

participants to store and manipulate large amounts of information, placing a strain on working memory and concentration ability, two cognitive resources that supports but do not equate to reasoning ability.

We argue that a game of incomplete information is a better tool to measure c-ToM. As in false belief tasks, individuals must establish connections between actions of others and the information they possess. Furthermore, game theoretic paradigms align more closely with real-world scenarios and are better equipped to isolate deficiencies in mentalizing about others rather than deficits in comprehension or computation. Finally, our game contains multiple dimensions of cognitive ability that can be measured, such as sensitivity of s to private value, likelihood of choosing a cutpoint strategy, and payoffs obtained.

As for a-ToM, it is commonly evaluated using the "Read the Mind in the Eyes Task" (RMET, Baron-Cohen et al. (1997)), a task where participants view images of a person's eyes and are required to select the most appropriate adjective to describe the emotion conveyed. The RMET has been widely used in research on social cognition. The ability develops gradually and reaches maturity by adolescence (Moor et al., 2012).

The relationship between a-ToM and c-ToM is not established but there are compelling reasons to believe in their interrelation. First, individuals with autism exhibit poor performance in both tasks. Second, there is preliminary evidence suggesting a positive association between proficiency in a-ToM tasks and fluid intelligence, as gauged by IQ tests (Baker et al., 2014). Third, recent research has revealed that both c-ToM and a-ToM are mediated by a shared network of brain regions referred to as the default mode network. This network typically handles processes related to understanding the thoughts and feelings of others and links up with regions responsible for emotional processing and cognition (Smallwood et al., 2021). The coexistence of both a-ToM and c-ToM within a single network implies that these two facets of theory of mind may have evolved to address interconnected social and cognitive challenges.

In this section, we test the idea that affective processing plays a crucial role in strategic thinking (Hypothesis 4). Specifically, we investigate the potential relationship between a-ToM and c-ToM abilities by examining whether participants who excel in our game also demonstrate a heightened ability to decipher emotions.

6.2 Testing affective theory of mind

The original RMET was designed to study a-ToM in adults. The emotions sometimes involve complex nuances ("suspicious", "thoughtful") and are expressed in a complicated language ("flustered", "skeptical"). The existing adaptations for children and adolescents have been limited to reducing the number of trials (Baron-Cohen et al., 2001), which we find unsatisfactory for the reasons detailed below. We therefore decided to adapt and validate a novel, child-friendly version of RMET with 16 images of the eyes of the same 12 years old child expressing different emotions, which we called the "Developmental Read the Mind in the Eyes Task" (DeRMET).[13] For each image, we provided four answers, one of which best captured the emotion. Participants were incentivized, earning 0.10€ for every correct response. In Figure 6, we present one example (73% of the population provided the correct answer, namely "disgusted"). Detailed instructions, including the complete set of images, possible responses, correct answers, and percentage of participants who chose correctly each of them are available in Appendix A2.

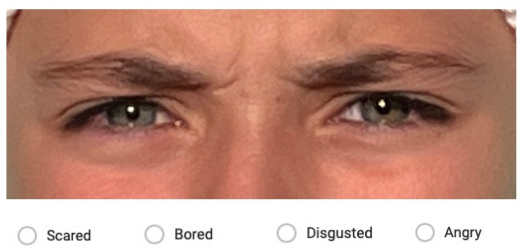

Figure 6: An example of an image and the four possible answers in the DeRMET

Compared to the original task, there are three advantages in using our version. By sourcing all images from a single individual, we make it more consistent. By using a child as a model, we make the images more relatable to young participants. By focusing on simpler emotions and using a straightforward language, we make the same task suitable for any age, from very young participants to adults. In Appendix B6, we provide aggregate statistics of performance in DeRMET. In the next section, we use the individual performance in this task as an input to explain behavior in the game.

[13]For its validation, we enlisted 43 teenagers and adults from different countries who were asked to identify emotions in 20 original images and to provide feedback on the difficulty of identifying these emotions. We then removed four emotions that participants found too difficult. The task and validation materials were registered with the Open Science Framework to ensure transparency and accessibility.

6.3 Affective Theory of Mind and optimal decision-making

To examine the link between a-ToM and c-ToM, we perform the same Probit regressions as in Tables 2 and 4 (likelihood of choosing s and C_+), but we now include performance in a-Tom as an additional independent variable. We also conduct OLS regressions to study the expected loss of an individual (relative to best responding to the empirical behavior of others in their grade group and treatment) as a function of the previous variables, including performance in a-ToM. The OLS regressions are performed both in the entire sample and in the subset of individuals who employ cutpoint strategies. We pool together all the information, with treatment dummy variables. The results are compiled in Table 6. Columns 1-2-3-4 report the regressions when we consider absolute performance in a-ToM (*abs-ToM*) while columns 5-6-7-8 report the same regressions when a-ToM is normalized to the behavior of other participants in their own grade (*norm-ToM*).

The results unambiguously indicate a strong, positive and robust link between a-ToM and different measures of strategic thinking. Individuals with high performance in a-ToM are more likely to pass for small values of x, more reactive to increases in x (columns (1) and (5)) and more likely to adopt cutpoint strategies (columns (2) and (6)) than participants with low performance in a-ToM. Effects are highly significant and complementary to the effect of age. Performance in a-ToM is also associated with lower losses unconditional (columns (3) and (7)) or conditional on using cutpoint strategies (columns (4) and (8)).[14] The results are very similar whether a-ToM performance is absolute (*abs-ToM*) or relative to other participants in their grade (*norm-ToM*).

As notice earlier, males adopt cutpoint strategies more often than females in their grade group and, consequently, achieve higher payoffs. This goes against our previous findings in a backward induction game, where females outperformed males (Brocas and Carrillo, 2021). However, participants in the two studies come from different schools with distinct demographics and cultures, which may account for the differences.

Result 4. *Hypothesis 4 is strongly supported by the data. There is a strong positive correlation between strategic thinking (reacting to value, adopting a cutpoint strategy, obtaining a high payoff) and affective theory of mind.*

[14] High and low a-ToM players who use cutpoint strategies set similar cutpoints on average (5.86 v. 5.84) but the former are more likely to concentrate them in intermediate values (standard deviation of 2.33 v. 2.65), which is typically a best response to the behavior of others (see section 5.4.1).

	Probit (1) $S(x)$	Probit (2) C_+	OLS (3) Loss	OLS (4) Loss | C_+	Probit (5) $S(x)$	Probit (6) C_+	OLS (7) Loss	OLS (8) Loss | C_+
Age	-0.010***	0.006***	0.0005*	0.002***	-0.011***	0.006***	0.0001	0.001***
	(0.001)	(0.0009)	(0.0003)	(0.0004)	(0.0010)	(0.0009)	(0.0003)	(0.0004)
x	-0.103***	—	—	—	-0.043*	—	—	—
	(0.029)				(0.024)			
$x \times$ *Age*	0.001***	—	—	—	0.002***	—	—	—
	(0.0002)				(0.0002)			
abs-ToM	-0.709***	0.528**	-0.241***	-0.361***	—	—	—	—
	(0.233)	(0.216)	(0.068)	(0.105)				
$x \times$ *abs-Tom*	0.125***	—	—	—	—	—	—	—
	(0.039)							
norm-ToM	—	—	—	—	-0.093***	0.101***	-0.027***	-0.033**
					(0.033)	(0.029)	(0.009)	(0.014)
$x \times$ *norm-Tom*	—	—	—	—	0.015***	—	—	—
					(0.006)			
Male	0.010	0.187***	-0.048***	-0.023	0.007	0.198***	-0.047***	-0.019
	(0.027)	(0.059)	(0.018)	(0.025)	(0.027)	(0.059)	(0.018)	(0.025)
2ndDec.	0.070***	-0.040	0.015**	0.012	0.070***	-0.040	0.015**	0.012
	(0.014)	(0.031)	(0.008)	(0.012)	(0.014)	(0.031)	(0.008)	(0.012)
T2	-0.072**	-0.012	-0.103***	-0.169***	-0.073**	-0.008	-0.104***	-0.167***
	(0.037)	(0.087)	(0.024)	(0.027)	(0.037)	(0.087)	(0.024)	(0.028)
T3	0.026	-0.054	-0.147***	-0.181 ***	0.025	-0.045	-0.149***	-0.182***
	(0.034)	(0.079)	(0.021)	(0.025)	(0.034)	(0.080)	(0.021)	(0.025)
T4F	-0.168***	-0.036	0.118***	0.164***	-0.168***	-0.031	0.117***	0.165***
	(0.036)	(0.080)	(0.024)	(0.034)	(0.036)	(0.080)	(0.024)	(0.035)
T4S	-0.109***	-0.071	0.152***	0.166***	-0.109***	-0.066	0.150***	0.166***
	(0.037)	(0.080)	(0.025)	(0.036)	(0.037)	(0.080)	(0.025)	(0.036)
Obs.	39,960	3,996	3,996	2,184	39,960	3,996	3,996	2,184
Pseudo R^2	0.156	0.045	0.117	0.189	0.155	0.047	0.114	0.181

Significance levels: $^{***} = 0.01$, $^{**} = 0.05$, $^{*} = 0.1$

Table 6: Regressions to evaluate the effect of the a-ToM task

7 A behavioral theory

Since departures from Nash equilibrium are significant and costly (section 5.4), it is natural to wonder if a behavioral theory could successfully explain the behavior of our participants. A first natural candidate would be a model of social preferences. However, with the parameters of our experiment, seeking the compromise outcome would require high levels of social preferences.[15] Also, even participants with these preferences should seek the competition outcome when assessing the overall earnings accrued throughout the experiment, rather than focusing on the payoff of each game separately (if participants choose always s, the average payoff is $(h + l)/2 > m$ for all players). Finally, an individual with so-

[15]For example, Fehr and Schmidt (1999)'s inequality aversion model would require $\beta > 5/9$ for participants to deviate from the behavior predicted by the standard theory.

cial concerns will always adopt a boundary cutpoint strategy (aP or aS), so these models cannot effectively predict the observed interior cutpoints.

7.1 α-cursed equilibrium: theory

We decided to study instead α-cursed equilibrium (α-CE). It is a behavioral theory based on an imperfect understanding of the relationship between information and action (Eyster and Rabin, 2005), that has been successful at explaining deviations in games of asymmetric information. Formally, in α-CE, individuals believe that the actions of opponents depend on their information with probability $(1-\alpha)$ and are independent of their information with probability α. Thus, $\alpha = 0$ corresponds to Bayesian Nash Equilibrium and $\alpha = 1$ is the polar case where individuals do not make any inferences from the rival's actions. Denoting $x^{\alpha}_{\mathbf{Tk}}$ the equilibrium cutpoint in treatment k under α-CE, the main properties of α-CE are (see Appendix B7 for the formal derivation):

(i) $x^{\alpha}_{\mathbf{T1}} = l$, $x^{\alpha}_{\mathbf{T2}} = m$ and $x^{\alpha}_{\mathbf{T3}} = h$ when $\alpha = 0$.
(ii) $x^{\alpha}_{\mathbf{T1}} = m$, $x^{\alpha}_{\mathbf{T2}} \in (m, h)$ and $x^{\alpha}_{\mathbf{T3}} = m$ when $\alpha = 1$.
(iii) $\frac{\partial x^{\alpha}_{\mathbf{T1}}}{\partial \alpha} \geq 0$, $\frac{\partial x^{\alpha}_{\mathbf{T2}}}{\partial \alpha} \geq 0$ and $\frac{\partial x^{\alpha}_{\mathbf{T3}}}{\partial \alpha} \leq 0$.
(iv) $x^{\alpha}_{\mathbf{T1}} = x^{\alpha}_{\mathbf{T4F}} = x^{\alpha}_{\mathbf{T4S}}$ for all α.

(i) holds by construction. (ii) states that fully cursed players set interior cutpoints, with identical choices in **T1** and **T3**. (iii) is arguably the most robust and interesting property, as it says that the amount of unraveling is inversely related to cursedness. (iv) implies that timing does not affect the equilibrium.

7.2 α-cursed equilibrium: estimation

We conduct a structural estimation of a two-parameter (α, λ) Behavioral Random Utility Model (BRUM), where α captures the cursedness of the population and λ reflects the mistakes in choices ($\lambda = 0$ corresponds to random choice and $\lambda \to \infty$ corresponds to best response to rivals).

Table 7 reports the estimated parameters α and λ of BRUM by treatment and grade group for individuals who use cutpoint strategies C_+ (in Appendix B8, we report the derivation of the structural model and a comparison between the the empirical c.d.f. and the estimated best fit of the BRUM.) For the goodness of fit (Fit), we compute the absolute value of the difference between the estimated probability of 'see' given BRUM and the empirical probability for every x, and then average over all values of x.

		grade group				
		3-4	5-6	7-8	9-10	11-12
T1	α	1.00	1.00	1.00	1.00	1.00
	λ	0.67	2.07	2.25	1.77	1.83
	Fit	0.13	0.10	0.04	0.10	0.14
T2	α	0.14	0.06	0.42	0.68	1.00
	λ	1.13	1.97	1.26	1.44	0.88
	Fit	0.01	0.00	0.00	0.02	0.08
T3	α	0.81	0.88	0.78	0.70	0.67
	λ	1.14	1.14	2.09	2.37	2.79
	Fit	0.03	0.00	0.00	0.02	0.02
T4F	α	1.00	1.00	1.00	1.00	1.00
	λ	0.81	1.41	0.93	0.79	1.69
	Fit	0.16	0.06	0.19	0.25	0.17
T4S	α	1.00	1.00	1.00	1.00	1.00
	λ	0.82	1.28	0.94	0.80	1.63
	Fit	0.15	0.10	0.14	0.22	0.13

Table 7: Estimation of α-CE for strategic consistent participants (C_+)

There are several interesting lessons to be learned from the estimations results in Table 7 and the graphical representations in Figure B6. Since the mean empirical cutpoint is above m in all treatments, even full cursedness is insufficient to capture the behavior of participants in **T1**, **T4F** and **T4S** (even with $\alpha = 1$, the estimated probabilities of s are still above the empirical counterparts for all x). The model is particularly unsuccessful in the sequential treatment, where 'pass' is very frequent and the behavior is asymmetric with respect to x. By contrast, the estimation provides an excellent fit in **T2** and **T3** (with the exception of 11-12 in **T2**), with interior levels of cursedness and higher responsiveness to utility differences as participants get older.

More generally, this behavioral theory exhibits both merits and limitations. The observed partial unraveling aligns with the behavioral predictions of cursed players, who underestimate (or ignore) the relationship between the rival's action and information. On the other hand, the large differences in α-estimates suggests that our participants struggle to incorporate treatment-specific incentives.

8 Conclusion

We have reported the results of the first controlled lab-in-the-field experiment designed to study the strategic behavior of children and adolescents in a game with two-sided private information, and the relationship between cognitive and affective theory of mind. To that end, we have recruited a large sample of 1662 participants with a wide age range, 8 to 18 years old. Our study led to several critical observations.

The results indicates a stepwise hierarchy of considerations individuals need to weigh when devising a strategy: (i) Should my private information influence my action? (ii) How does my payoff relate to my private information under each action? (iii) When is my action consequential and what does that mean for optimal conditioning? The answers to these questions serve as input in the calculation of a best response strategy. While all participants successfully navigate step (i), proficiency in step (ii) tends to improve with age. However, the fact that they do not adjust their strategies in response to observed behavior within a given treatment, and that strategies remain similar across treatments, suggests that the calculations involved in step (iii) may be too intricate. This implies that players often struggle to fully comprehend their rivals' incentives and the broader contextual details that inform strategic decision-making under asymmetric information.

We also observe a large heterogeneity in strategies. Some of the youngest participants select strategies closely resembling either equilibrium or best response, while a subset of the older participants lacks discernible, coherent strategies altogether. Importantly, individuals with a heightened sense of affective theory of mind select more often cutpoint strategies, especially those aligned with optimal responses to others' behavior. Consequently, they achieve higher payoffs, even after accounting for age differences. This result underscores the close relationship between affective and cognitive theory of mind. In terms of the hierarchy of considerations highlighted above, affective theory of mind seems to facilitate deliberation in step (ii) and, to a lesser extent, also in step (iii).

The positive correlation between a-ToM and c-ToM in our game offers insight into the interplay between the affective and cognitive facets of our capacity to understand others. While our results do not provide any causal evidence between affective processing and cognitive deliberations, they suggest that the ability to detect emotional states in others supports our representations of their reasoning, enabling us to incorporate this insight more efficiently into our decision-making. There is a substantial body of recent literature emphasizing the significant role of our emotions in influencing our decision-making (Lerner et al., 2015). Crucially, this perspective finds support in neuroscience research as well (Naqvi et al., 2006; Phelps et al., 2014). Our results suggest that our representations of the emotions of others also guide our decisions.

Shifting the examination of game theoretic paradigms beyond purely cognitive aspects may provide insights into why cognitive abilities alone are frequently insufficient for success in interactive settings. Our decision-making processes do not strictly adhere to the logical steps of the mathematical proofs we write to compute equilibria. Instead, we often utilize algorithms, largely unconsciously, and frequently rely on heuristics influenced by intuition. Emotions are important factors in these decision-making processes. They serve as key elements in translating our perception of external stimuli, such as facial expressions, and our interpretations of constructed emotional representations into actionable information. Individuals who excel at interpreting emotions have access to richer information about others. It logically follows that they utilize this information as an input for shaping their responses, which may explain the very robust connection we have observed between performance in a-ToM and final payoffs in our game. Future research should aim to dig deeper into the causal aspects of these mechanisms.

References

Sule Alan and Betul Turkum. Abstract reasoning, theory of mind and character development in the school. *Journal of Economic Behavior & Organization*, 221:307–326, 2024.

Sule Alan, Ceren Baysan, Mert Gumren, and Elif Kubilay. Building social cohesion in ethnically mixed schools: An intervention on perspective taking. *The Quarterly Journal of Economics*, 136(4):2147–2194, 2021.

Crystal A Baker, Eric Peterson, Steven Pulos, and Rena A Kirkland. Eyes and iq: A meta-analysis of the relationship between intelligence and "reading the mind in the eyes". *Intelligence*, 44:78–92, 2014.

Simon Baron-Cohen, Therese Jolliffe, Catherine Mortimore, and Mary Robertson. Another advanced test of theory of mind: Evidence from very high functioning adults with autism or asperger syndrome. *Journal of Child psychology and Psychiatry*, 38(7):813–822, 1997.

Simon Baron-Cohen, Sally Wheelwright, Amanda Spong, Victoria Scahill, John Lawson, et al. Are intuitive physics and intuitive psychology independent? a test with children with asperger syndrome. *Journal of developmental and learning disorders*, 5(1):47–78, 2001.

Yoella Bereby-Meyer and Brit Grosskopf. Overcoming the winner's curse: an adaptive learning perspective. *Journal of Behavioral Decision Making*, 21(1):15–27, 2008.

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

Pablo Brañas-Garza, Isabelle Brocas, and Juan D Carrillo. Developmental meritocracy. USC Working Paper, 2024.

Jordi Brandts and Gary Charness. The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14:375–398, 2011.

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, 2020.

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, Stephanie W Wang, and Colin F Camerer. Imperfect choice or imperfect attention? understanding strategic thinking in private information games. *Review of Economic Studies*, 81(3):944–970, 2014.

Juan D Carrillo and Thomas R Palfrey. The compromise game: two-sided adverse selection in the laboratory. *American Economic Journal: Microeconomics*, 1(1):151–181, 2009.

Juan D Carrillo and Thomas R Palfrey. No trade. *Games and Economic Behavior*, 71(1): 66–87, 2011.

Gary Charness and Dan Levin. The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics*, 1(1):207–36, 2009.

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.

Corrado Corradi-Dell'Acqua, Christoph Hofstetter, and Patrik Vuilleumier. Cognitive and affective theory of mind share the same local patterns of activity in posterior temporal but not medial prefrontal cortex. *Social cognitive and affective neuroscience*, 9(8):1175–1184, 2014.

Erik Eyster and Matthew Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

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

Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868, 1999.

Brit Grosskopf, Yoella Bereby-Meyer, and Max Bazerman. On the robustness of the winner's curse phenomenon. *Theory and Decision*, 63:389–418, 2007.

John H Kagel and Dan Levin. *Common value auctions and the winner's curse*. Princeton University Press, 2002.

Peter Kinderman, Robin Dunbar, and Richard P Bentall. Theory-of-mind deficits and causal attributions. *British journal of Psychology*, 89(2):191–204, 1998.

Jennifer S Lerner, Ye Li, Piercarlo Valdesolo, and Karim S Kassam. Emotion and decision making. *Annual review of psychology*, 66(1):799–823, 2015.

Bregtje Gunther Moor, Zdeňa A Op de Macks, Berna Güroğlu, Serge ARB Rombouts, Maurits W Van der Molen, and Eveline A Crone. Neurodevelopmental changes of reading the mind in the eyes. *Social cognitive and affective neuroscience*, 7(1):44, 2012.

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

Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American economic review*, 85(5):1313–1326, 1995.

Nasir Naqvi, Baba Shiv, and Antoine Bechara. The role of emotion in decision making: A cognitive neuroscience perspective. *Current directions in psychological science*, 15(5): 260–264, 2006.

Elizabeth A Phelps, Karolina M Lempert, and Peter Sokol-Hessner. Emotion and decision making: multiple modulatory neural circuits. *Annual review of neuroscience*, 37:263–287, 2014.

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.

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.

Jonathan Smallwood, Boris C Bernhardt, Robert Leech, Danilo Bzdok, Elizabeth Jefferies, and Daniel S Margulies. The default mode network in cognition: a topographical perspective. *Nature reviews neuroscience*, 22(8):503–513, 2021.

Richard H Thaler. Anomalies: The winner's curse. *Journal of economic perspectives*, 2 (1):191–202, 1988.

Annalisa Valle, Davide Massaro, Ilaria Castelli, and Antonella Marchetti. Theory of mind development in adolescence and early adulthood: The growing complexity of recursive thinking ability. *Europe's journal of psychology*, 11(1):112, 2015.

Henry M Wellman, David Cross, and Julanne Watson. Meta-analysis of theory-of-mind development: The truth about false belief. *Child development*, 72(3):655–684, 2001.

Heinz Wimmer and Josef Perner. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13 (1):103–128, 1983.

Appendix A. Details of design and procedures

A1 Instructions of the game

[This is a translation of the Spanish instructions for **T1** (other treatments are analogous)]

Hi, my name is Juan. Today, we are going to play a game. In this game, you will earn points. At the end of the experiment you will exchange each point for one cent that you will be able to spend in the cafeteria. You will get a lot of points so you will get enough money to buy something you like. Just remember: the more points you get, the more money you will have.

The computer has 10 cards numbered 1 to 10. It will match you with another person in the room. You will not know who that person is and it is not the objective of the game to learn that. The computer will give one card to each of you. You will see the number in your card but not the number in the other person's card. (Note. See Figure A1 for the slides projected.)

[**SLIDE 1, SLIDE 2, 'CLICK' TO SHUFFLE, 'CLICK' TO PICK A CARD, SLIDE 3**]

The rules of the game are very easy. You will both have to say "pass" or "see".
- If you both say 'pass', [**CLICK**] you each win **5** points, no matter what card you have.
- If at least one of you says 'see', [**CLICK**] then whoever has the highest card wins **10** points and the other wins only **1** point.

It is important to understand that the only way for both of you to win 5 points is if you both say 'pass'. So, if you say 'pass' but the other says 'see' <u>or</u> if the other says 'pass' but you say 'see' <u>or</u> if you both say 'see', in all these cases, whoever has the highest card wins 10 points and the other wins 1 point. Let's see if you have understood.

[**SLIDE 4**]

Suppose you got a 6 and you decided to say 'pass'. This is just an example. How many points will you get if:

- [**SLIDE 5**] The other has a 7 and says 'pass'. [**CLICK**]. Answer: 5 points
- [**SLIDE 6**] The other has a 7 and says 'see'. [**CLICK**]. Answer: 1 point
- [**SLIDE 7**] The other has a 4 and says 'see'. [**CLICK**]. Answer: 10 points

Ok. Now that we know the rules, let me explain you how we are going to play. You are going to see in your computer a screen like this:

[**SLIDE 8**]

The computer is going to ask you: what do you want to do if you get a 1 [**CLICK**]? if you get a 2 [**CLICK**]? if you get a 3 [**CLICK**]? And so on. You have to respond 'pass' or 'see' for each card. It is very important that you say what you want to do for all cards, from 1 to 10. The computer is going to ask the same questions to the other person at the same time.

When you have both responded to all the questions, the computer will give you one card each and will do whatever you have decided. For example, if the computer gives you a 6 and for that number you said 'pass' [**CLICK**], the computer will choose 'pass' for you. Is that clear?

In order to earn more points (and therefore make more money!), you are going to play this game six times, each time with a different person and a random number between 1 and 10. Besides, whatever you win in the six games, I am going to give you FIVE times that amount of money. That is, I will multiply all the points by 5.

At the end, you will see for each of the six games [**SLIDE 9**]:

- Your number
- The number of the other person
- What you told the computer to do if you got that number (‘pass’ or ‘see’)
- What the other person told the computer to do if they got that number (‘pass’ or ‘see’)
- And finally, the points you got given the rules of the game

Do you want to play?

[**AT THE END OF THE GAME**]

You can now see in your screen the choices and points you got in the six games. Do you want to play again? Once again, I am going to give you 5 times what you get. But before I am going to ask you again what you want to do if you get each number. Do you remember this screen?

[**SLIDE 10**]

You can answer the same as before or, you can say something different if you think you can win more points. It is totally up to you.

[**AT THE END OF THE GAME**]

You can see once again in your screen the choices and points you got in the six games.

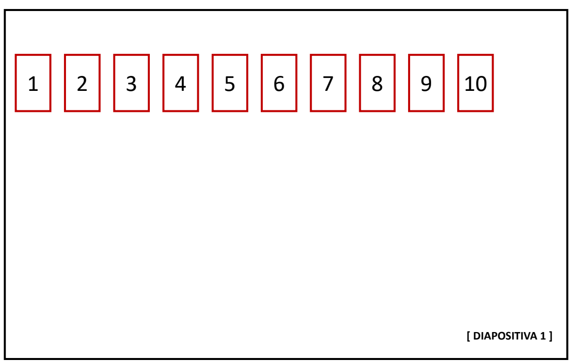

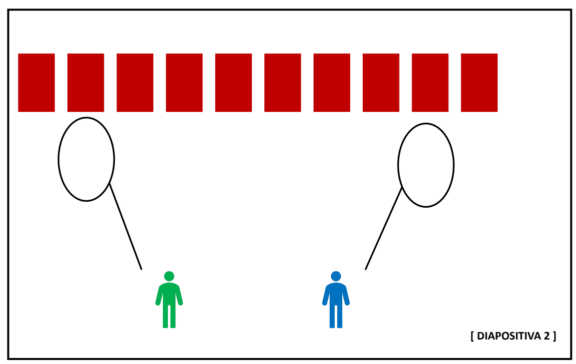

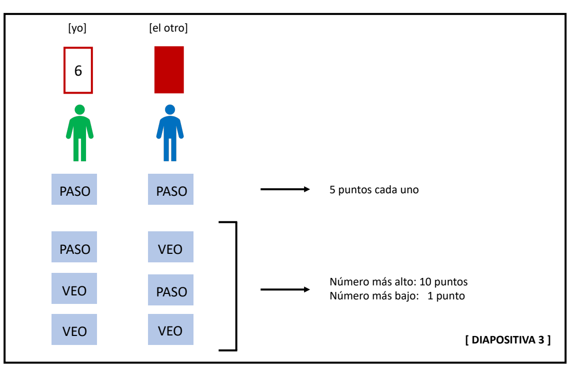

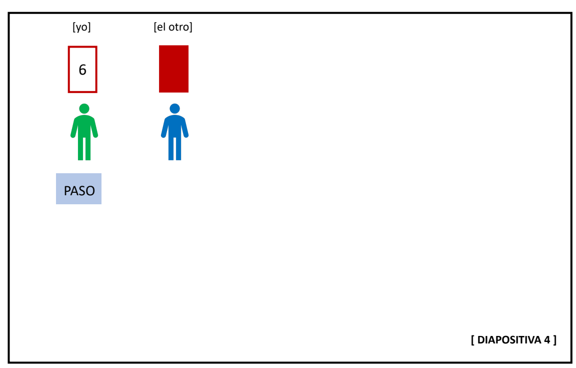

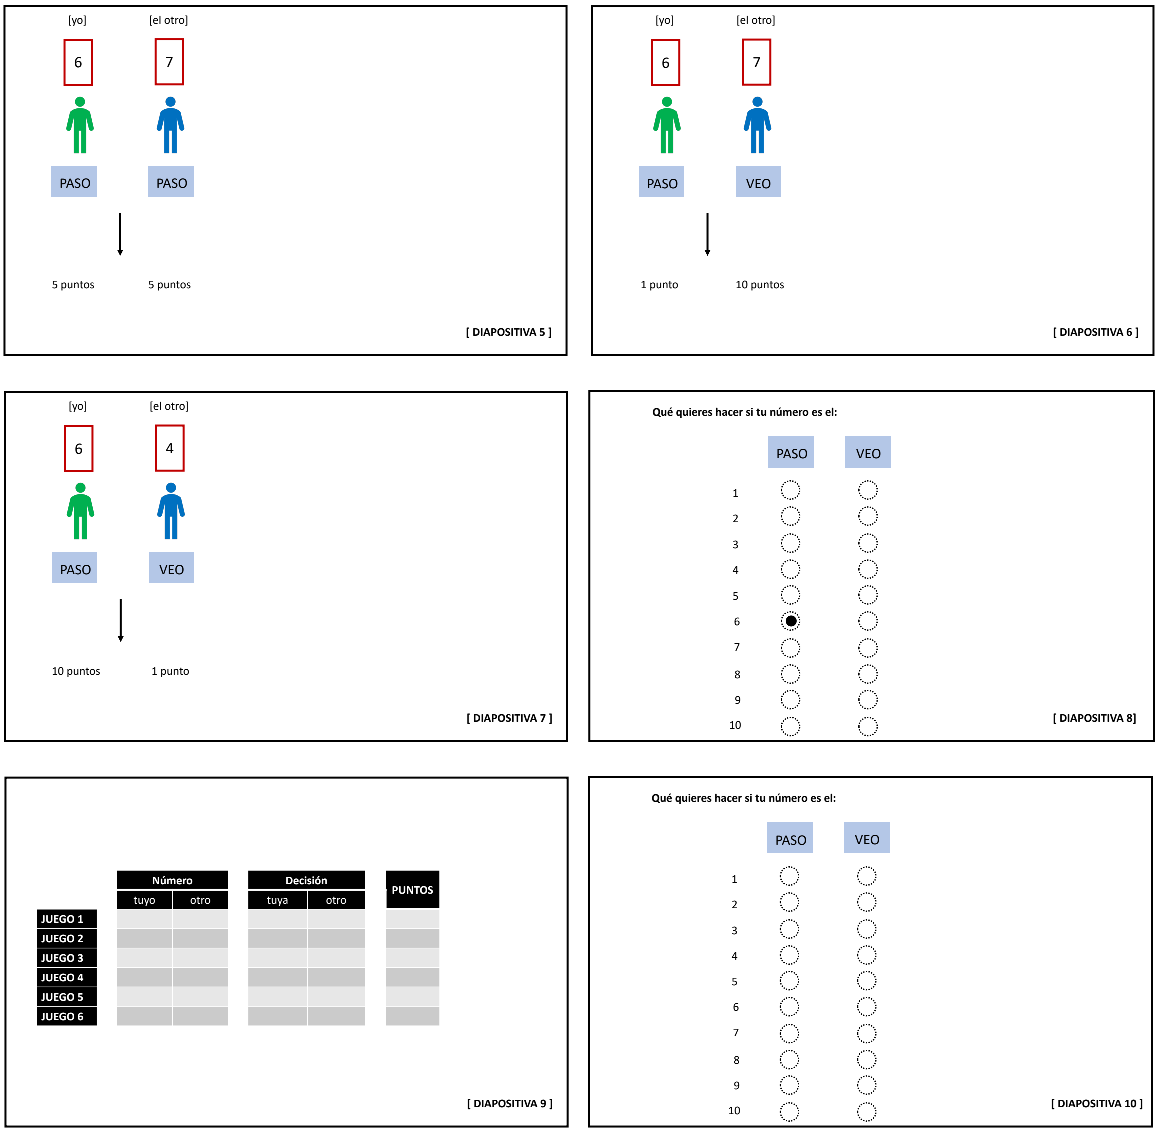

Figure A1: Slides (in Spanish) projected on screen for instructions of the game

A2 Developmental Read the Mind in the Eyes (DeRMET)

Instructions.

Below, you will see a series of images with a boy's eyes expressing a certain emotion. Look closely and choose, from the four options, the one that best represents the person's emotion. When you are sure of your answer, press NEXT, and the next image will appear. Once you press NEXT, you will not be able to change your answer. For each correct answer, you will receive 10 points, which is equivalent to 10 cents. The more correct answers, the more money you will earn.

Note to the instructor.

Each image is presented in a different page with four emotions to choose from (see the list below). The participant must select exactly one. After a choice is made, the next image appears on the screen and the participant cannot go back. There is no feedback between images.

List of images.

1. 
2. 
3. 
4. 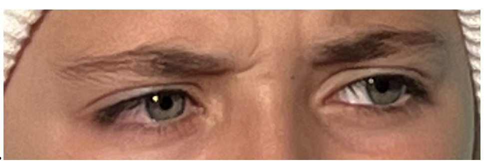
5. 
6. 
7. 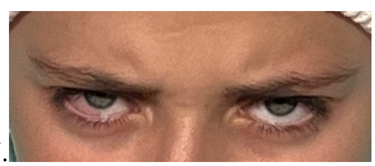
8. 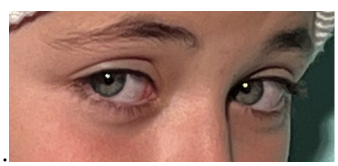
9. 
10. 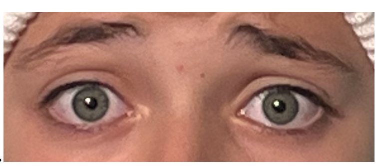
11. 
12. 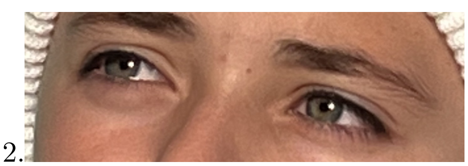

13.  14. 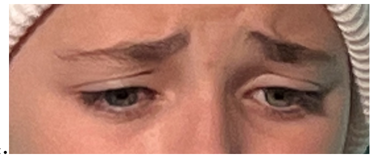 15. 

16. 

List of emotions for each image.

Image #		Emotions			% correct†
1	1. Bored	2. Confused*	3. Hopeful	4. Guilty	57.6
2	1. Sneaky	2. Angry	3. Guilty	4. Tired*	80.0
3	1. Guilty	2. Tired	3. Confused	4. Sad*	62.3
4	1. Admiring	2. Disapproving*	3. Shocked	4. Bored	48.0
5	1. Surprised	2. Admiring	3. Happy*	4. Sneaky	75.1
6	1. Confused	2. Ashamed*	3. Scared	4. Tired	73.4
7	1. Angry*	2. Scared	3. Sad	4. Worried	97.8
8	1. Excited	2. Confused	3. Thinking	4. Flirting*	71.7
9	1. Excited*	2. Happy	3. Hopeful	4. Flirting	76.1
10	1. Disgusted	2. Sad	3. Angry	4. Scared*	89.8
11	1. Scared	2. Bored	3. Disgusted*	4. Angry	73.4
12	1. Surprised	2. Happy	3. Thinking*	4. Disapproving	89.3
13	1. Hopeful	2. Sneaky*	3. Thinking	4. Excited	92.3
14	1. Guilty	2. Bored	3. Disgusted	4. Worried*	61.0
15	1. Happy	2. Surprised	3. Hopeful*	4. Sneaky	49.0
16	1. Surprised*	2. Happy	3. Flirting	4. Admiring	67.7

*: correct answer.

† It reports for each image the aggregate empirical probability of finding the correct answer

Note. What we call the "correct" answer is the emotion the actor was instructed to express in a subtle manner. For all images, it also corresponds to the choice of the majority of players.

Appendix B. Proofs and additional analyses

B1 Proof of Proposition 1

The proof is a slight variation of Proposition 1 in [CP].

Proof of **T1**. Step 1. Suppose for technical convenience that $F(\cdot)$ is continuous and differentiable with $f(x_i) > 0$ for all x_i (although in the experiment we use integer values of x_i). By construction, player i's strategy is relevant only if $a_j = p$. Conditional on $a_j = p$, player i's payoff is:

$$u_i(p,p) = m \quad \text{and} \quad u_i(s,p) = \Pr(x_j < x_i \,|\, a_j = p)h + \Pr(x_j > x_i \,|\, a_j = p)l \tag{1}$$

From (1), it is immediate that $\frac{\partial u_i(p,p)}{\partial x_i} = 0$ and $\frac{\partial u_i(s,p)}{\partial x_i} \geq 0$, which implies the existence of a unique (but not necessarily interior) value $x_i^* \in [\underline{x}, \overline{x}]$ such that $u_i(p,p) \gtreqless u_i(s,p)$ iff $x_i \lesseqgtr x_i^*$. This proves that the equilibrium must involve a cutpoint strategy.

Step 2. Assume player j uses a cutpoint strategy x_j^*. Using (1), then for a given x_i, player i's expected payoff under p and s and given $a_j = p$ are:

$$u_i(p,p) = m \quad \text{and} \quad u_i(s,p) = \Pr(x_j < x_i \,|\, x_j < x_j^*)h + \Pr(x_j > x_i \,|\, x_j < x_j^*)l \tag{2}$$

From (2) and given $h > m$, it is immediate that $\lim_{x_i \to (x_j^*)^-} u_i(s,p) > \lim_{x_i \to (x_j^*)^-} u_i(p,p) = m$. This means that $x_i^* < x_j^*$. Using a symmetric reasoning, $x_j^* < x_i^*$, and therefore the unique Bayesian Nash Equilibrium (BNE) unravels to $x_i^* = x_j^* = \underline{x}$.

The proof in **T2** is identical when $x_i \geq m$. When $x_i < m$, $a_i(x_i) = s$ is weakly dominated by $a_i(x_i) = p$, so that $x_i^* = x_j^* = m$.[16] The proof for **T3** is symmetric to **T1**, with each player willing to choose a cutpoint higher (instead of lower) than the rival. The proof for the Perfect Bayesian Equilibrium in **T4** is identical to the BNE in **T1**. □

B2 Learning

In the analysis shown in Table 2 we observed a trend where the inclination to seek the competition outcome (i.e., choosing 'see') increases between the first and second decision (*2ndDec*). To investigate these changes in behavior in more detail, we combine data from all grades and calculate $S_2(x) - S_1(x)$ for each treatment. This variable captures the change in the overall likelihood of choosing s as a function of x between the second and first decision. The results are displayed in Figure B1, where the error bars represent ± 2 standard errors of the mean. A positive value represents an increase in s, and a negative value represents a decrease in s (as in Figure 2, we cluster two consecutive grades together for visual ease and to increase statistical power).

If participants were to play closer to equilibrium in their second decision, we would observe $S_2(x) > S_1(x)$ for all x in **T1** and **T4** and for $x \geq 6$ in **T2** (positive values in Figure B1). We would also observe $S_2(x) < S_1(x)$ for all x in **T3** and for $x \leq 5$ in **T2** (negative values in Figure B1). Conversely, if we expected that participants would learn how to best respond to their peers (the

[16] Strictly speaking, $x_i^* = x_j^* = \underline{x}$ is also a BNE in **T2**: if $x_j^* = \underline{x}$, player i's action is irrelevant and viceversa. However, this knife-edge equilibrium is in weakly dominated strategies and it is not Trembling Hand. We will therefore focus on the more natural one.

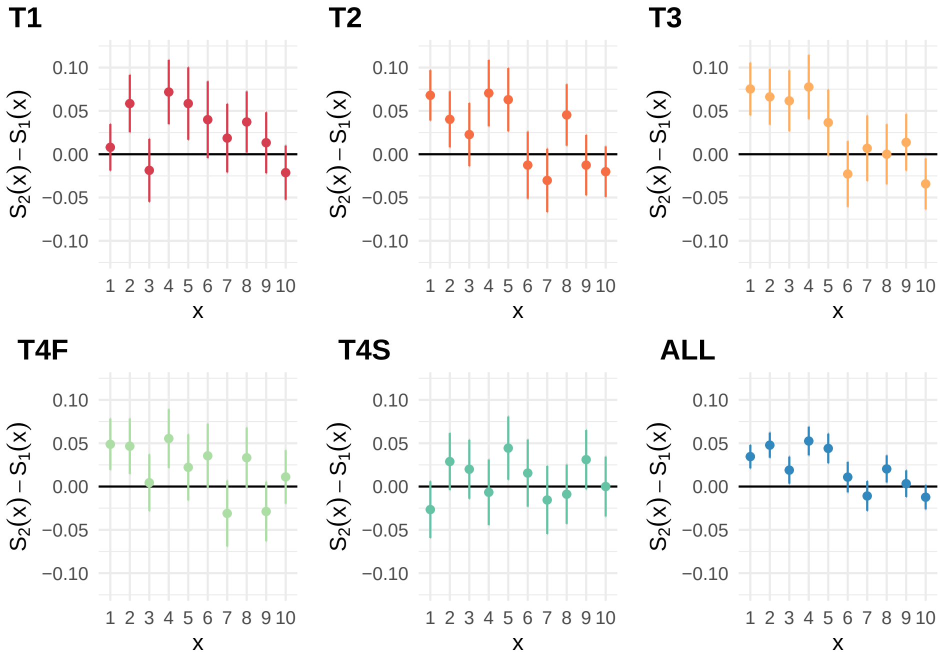

Figure B1: $S_2(x) - S_1(x)$ by treatment

vast majority of whom do not play at equilibrium), we should observe a decrease in the likelihood of playing s for 'low' x and an increase for 'high' x across all treatments. Also, the exact cutoff points would vary depending on the treatment and grade groups.

From Figure B1, it is apparent that the majority of changes between the first and second decision are modest (below 0.08) and often lack statistical significance. Looking at **T1**, it initially seems that participants are learning to align their choices with equilibrium predictions since $S_2(x)$ tends to be higher than $S_1(x)$, especially for $x \in \{2, 4, 5\}$. However, this observation can be misleading. A more robust version of this pattern is observed in **T3**, where theory predicts precisely the opposite behavior. Even more unexpectedly, we observe an increase in choosing option s in **T2** when $x \leq 5$, despite it being a dominated and empirically suboptimal strategy. In **T4F**, we observe slightly more competition for low x and mixed results for high x. Overall, the general tendency (bottom right graph) shows an increase in the choice of option s, but primarily for low values of x, precisely when it is empirically unlikely to be advantageous.

In summary, our initial expectation of observing small but positive improvements between the first and second decisions was not met. The slight tendency to increase the choice of option s is unlikely to result from participants learning how to best respond or play at equilibrium, as it often contradicts both predictions. There are several possible explanations for this lack of improvement. Firstly, feedback in this game is not conducive to effective learning, as it often necessitates making inferences based on counterfactual scenarios. Secondly, feedback may have been presented in an excessively dense manner (Figure 1 (right)). Lastly, two decisions may have been insufficient to capture a learning trend. Given the absence of meaningful changes, we will combine data from both decisions for the remainder of the paper, unless otherwise specified.

B3 Sensitivity of action to private information

The variations in the sensitivity of actions to private values across different grade groups can best be appreciated by plotting the average marginal effect of information by age. We use the predictions derived from the fitted Probit model to assess the partial derivative of the probability of each individual of selecting option s as a function of the private value. We then calculate the average of these derivatives within each grade group. A rising trend in the average marginal change in the probability of choosing option s indicates that older participants are more responsive to variations in private values. We present this information for each treatment separately in Figure B2 (error bars represent the 95% confidence interval).

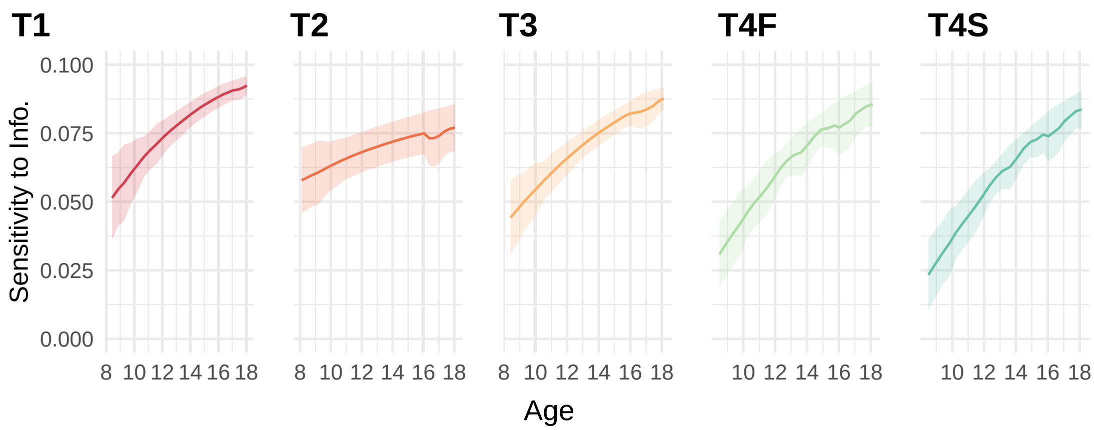

Figure B2: Sensitivity of action to private value by treatment and age

The responsiveness of actions to changes in private value consistently grows with age, though the strength of this effect varies among treatments. It is more pronounced in **T4** and milder in **T2**. This observation confirms that older players change their behavior more drastically in response to changes in information compared to their younger peers.

B4 Behavior of individuals who use and do not use cutpoint strategies

We carry out a Probit regression similar to the one in Table 2. However, we divide the sample into two groups: individuals who use perfect cutpoint strategies ($C_+ = C_0 + aS + aP$) and individuals who do not use perfect cutpoint strategies ($notC_+ = C_1 + ALT + Other$). Using the predictions obtained from these Probit models, we perform the same exercise as in Figure B2 to examine the sensitivity of action to private information across ages. For simplicity, we pool all treatments together. The results of this exercise are presented in Figure B3 for all players together (All, left), players who use cutpoint strategies (C_+, center), and players who do not use cutpoint strategies ($notC_+$, right).

As already indicated in section 4, Figure B3 (left) reaffirms that older participants react more to private value than their younger peers. According to Figure B3 (center), the subgroup of

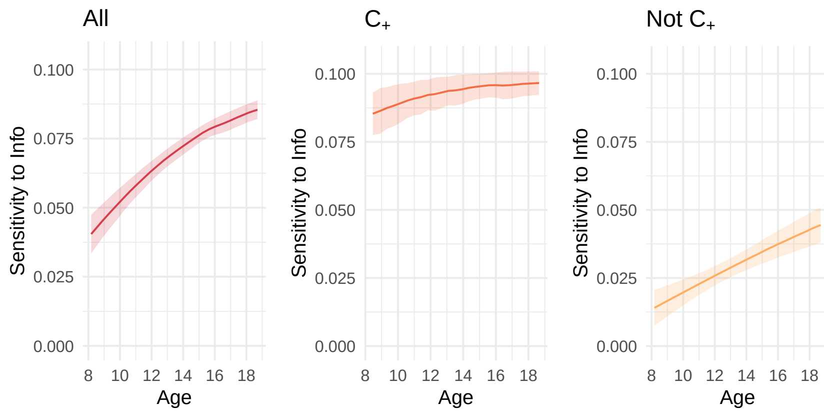

Figure B3: Sensitivity of action to private information by type of strategy

individuals who use cutpoint strategies display the highest sensitivity to information. Importantly, this sensitivity remains relatively consistent across different grade groups, with only a slight increase as they grow older. In contrast, individuals who do not employ cutpoint strategies exhibit low sensitivity to information, primarily because of their frequent switches between the 'pass' and 'see' options. However, with increasing age, these switches become less frequent, resulting in higher sensitivity to information, as depicted in Figure B3 (right). It's noteworthy that our youngest players who employ cutpoint strategies display a sensitivity to information comparable to our oldest "average" player. Conversely, our oldest participants who do not employ cutpoint strategies exhibit a sensitivity to information similar to our youngest "average" player.

In summary, participants in all ages who employ cutpoint strategies exhibit consistent reactions to information. Primary differences across grade groups are twofold: (i) proportion of individuals who adhere to strategic consistency and (ii) behavior of those who do not adhere to it.

We also conduct a basic analysis of changes in strategies across decisions. In Table B1, we present a 2×2 matrix. Each cell represents the percentage of observations that correspond to a cutpoint (C_+) v. no cutpoint (notC_+) strategy in the first (rows) and second (columns) decision.

	C_+	notC_+
C_+	43.2	11.4
notC_+	9.7	35.7

Table B1: Percentage of strategies in first (rows) and second (columns) decision

Once again, Table B1 provides evidence that systematic changes in behavior are largely absent. Indeed, for the vast majority of cases (78.9%), we observe the same type of strategy in both decisions. Furthermore, the likelihood of learning and unlearning are very similar: among participants who choose notC_+ in the first decision, 21.4% choose C_+ in the second, and among participants who choose C_+ in the first decision, 20.9% choose notC_+ in the second. This means that when there is a change, it is equally likely to reflect an increase or a decrease in strategic consistency. A similar pattern emerges when considering the finer partition of all six strategies.

B5 Choices of participants as first and second movers

Given our strategy elicitation method, participants in the sequential treatment make choices for all x, both as first and second movers. We can then perform a comparison of individual choices ('pass' v. 'see') in **T4F** and **T4S**. Figure B4 reports that information by grade group (left) and private value (right).

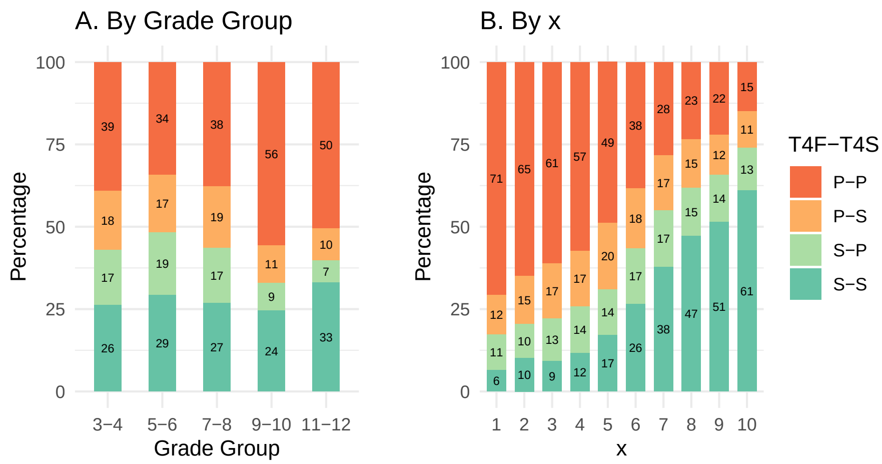

Figure B4: Change in individual choice between **T4F** and **T4S**

Although many individuals choose the same action as both first and second movers, there is still a significant fraction who change between the two roles. As one would expect, changes are more frequent for intermediate values of information, where the optimal decision is less clear (34-35% for $x \in \{5, 6, 7\}$) than for extreme ones (23-26% for $x \in \{1, 2, 9, 10\}$), as shown in Figure B4(B). Additionally, our younger participants are almost twice as likely to change their actions compared to their older counterparts (35-36% in 3-4, 5-6, 7-8 v. 17-20% in 9-10, 11-12, Figure B4(A)). On the other hand, changes are equally frequent in both directions at all ages and for all values, reinforcing the idea that players treat their decision as first and second mover very similarly.

B6 Descriptive statistics of performance in DeRMET

Figure B5 summarizes the distribution of correct answers in DeRMET by grade group, with the group average at the top.

As with most tasks that combine sensory perception with the ability to process social cues, a-ToM is not age-invariant. Nevertheless, the average improvement is relatively small despite the large age range (from 10.1 to 12.6, t-test $p < 0.001$). Most importantly, very few participants obtain a perfect score (16) or a score likely to occur under random choice.[17] This means that the

[17] A random player would obtain a score of 5 or less with probability 0.81, which is very far from what we observed in any grade group.

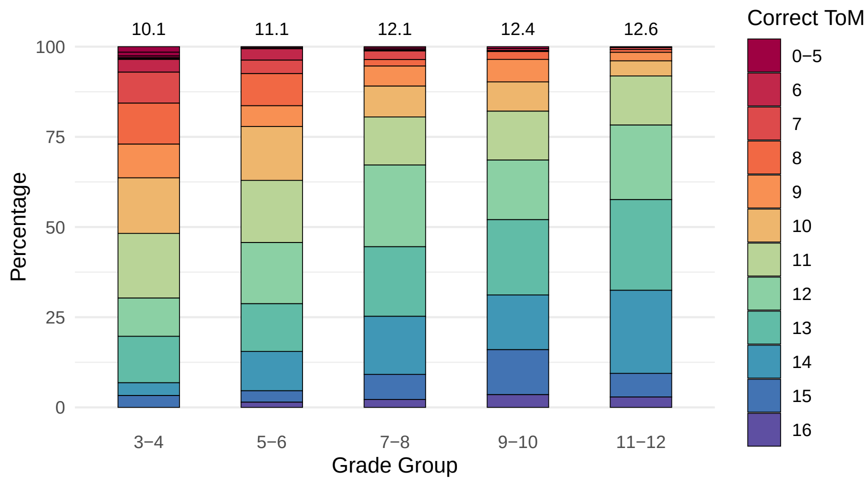

Figure B5: Distribution of correct answers in the DeRMET (a-ToM) task by grade group

task is well-calibrated: neither trivial nor impossibly challenging for any individual. Finally, the percentage of correct answers for each image varies between 48.0% and 97.8%, which means that the answer we designated as "correct" aligns with the choice of the majority of participants.[18]

B7 Theoretical derivation of the α-cursed equilibrium

We can compute for each treatment, the optimal decision of α-cursed players. Following the same steps as in (2) and assuming that α-cursed player j chooses the cutpoint x^α, the expected utility of α-cursed player i with value x_i in **T1** when they play s and given $a_j = p$ is:

$$\begin{aligned} u^\alpha_{i,\mathbf{T1}}(s,p;x_i) &= \alpha\Big[\Pr(x_j < x_i)h + \Pr(x_j > x_i)l\Big] \\ &\quad +(1-\alpha)\Big[\Pr(x_j < x_i \,|\, x_j < x^\alpha)h + \Pr(x_j > x_i \,|\, x_j < x^\alpha)l\Big] \\ &= \alpha\Big[F(x_i)h + (1-F(x_i))l\Big] + (1-\alpha)\Big[\min\{1, \tfrac{F(x_i)}{F(x^\alpha)}\}h + \max\{0, 1-\tfrac{F(x_i)}{F(x^\alpha)}\}l\Big] \end{aligned}$$

The equilibrium cutpoint $x^\alpha_{\mathbf{T1}}$ is given by $u^\alpha_{i,\mathbf{T1}}(p,p;x^\alpha_{\mathbf{T1}}) = u^\alpha_{i,\mathbf{T1}}(s,p;x^\alpha_{\mathbf{T1}})$, that is:

$$m = \alpha\Big[F(x^\alpha_{\mathbf{T1}})h + (1-F(x^\alpha_{\mathbf{T1}}))l\Big] + (1-\alpha)h \;\Leftrightarrow\; F(x^\alpha_{\mathbf{T1}}) = \begin{cases} 1 - \frac{h-m}{\alpha(h-l)} & \text{if} \quad \alpha \geq \frac{h-m}{h-l} \\ 0 & \text{if} \quad \alpha < \frac{h-m}{h-l} \end{cases} \tag{3}$$

A similar analysis in **T2** yields:

$$u^\alpha_{i,\mathbf{T2}}(s,p;x_i) = \alpha\Big[F(x_i)x_i + (1-F(x_i))l\Big] + (1-\alpha)\Big[\min\{1, \tfrac{F(x_i)}{F(x^\alpha)}\}x_i + \max\{0, 1-\tfrac{F(x_i)}{F(x^\alpha)}\}l\Big]$$

[18] This is important for our research. Indeed, if it were not the case, it would bring into question what constitutes a correct answer: the one we expect or the one that individuals choose.

so the equilibrium cutpoint $x^{\alpha}_{\mathbf{T2}}$ is implicitly defined as the solution of:

$$m = \alpha\Big[F(x^{\alpha}_{\mathbf{T2}})x^{\alpha}_{\mathbf{T2}} + (1 - F(x^{\alpha}_{\mathbf{T2}}))l\Big] + (1-\alpha)x^{\alpha}_{\mathbf{T2}} \;\Leftrightarrow\; \alpha(1 - F(x^{\alpha}_{\mathbf{T2}})) = \frac{x^{\alpha}_{\mathbf{T2}} - m}{x^{\alpha}_{\mathbf{T2}} - l} \tag{4}$$

Using implicit differentiation, it can be easily shown that:

$$\frac{\partial x^{\alpha}_{\mathbf{T2}}}{\partial \alpha} > 0, \quad x^{\alpha}_{\mathbf{T2}} = m \text{ when } \alpha = 0, \quad x^{\alpha}_{\mathbf{T2}} < h \text{ when } \alpha = 1$$

The analysis is similar in **T3** except that we need to condition on $a_j = s$:

$$\begin{aligned} u^{\alpha}_{i,\mathbf{T3}}(s, s; x_i) &= \alpha\Big[\Pr(x_j < x_i)h + \Pr(x_j > x_i)l\Big] \\ &\quad +(1-\alpha)\Big[\Pr(x_j < x_i \,|\, x_j > x^{\alpha})h + \Pr(x_j > x_i \,|\, x_j > x^{\alpha})l\Big] \\ &= \alpha\Big[F(x_i)h + (1 - F(x_i))l\Big] + (1-\alpha)\Big[\max\{\tfrac{F(x_i)-F(x^{\alpha})}{1-F(x^{\alpha})}, 0\}h + \min\{1, \tfrac{1-F(x_i)}{1-F(x^{\alpha})}\}l\Big] \end{aligned}$$

The equilibrium cutpoint $x^{\alpha}_{\mathbf{T3}}$ is given by $u^{\alpha}_{i,\mathbf{T3}}(p, s; x^{\alpha}_{\mathbf{T3}}) = u^{\alpha}_{i,\mathbf{T3}}(s, s; x^{\alpha}_{\mathbf{T3}})$, that is:

$$m = \alpha\Big[F(x^{\alpha}_{\mathbf{T3}})h + (1 - F(x^{\alpha}_{\mathbf{T3}}))l\Big] + (1-\alpha)l \;\Leftrightarrow\; F(x^{\alpha}_{\mathbf{T3}}) = \begin{cases} \frac{m-l}{\alpha(h-l)} & \text{if} \quad \alpha \geq \frac{m-l}{h-l} \\ 1 & \text{if} \quad \alpha < \frac{m-l}{h-l} \end{cases} \tag{5}$$

Finally, since a cursed player does not link action to information, the second mover does not revise their beliefs after observing the choice of the first mover. As a result, $x^{\alpha}_{\mathbf{T4F}} = x^{\alpha}_{\mathbf{T4S}} = x^{\alpha}_{\mathbf{T1}}$. For uniform distributions with $\underline{x} = l$ and $\overline{x} = h$, we simply set $F(x) = \frac{x-l}{h-l}$ in (3)-(4)-(5) to obtain:

$$\mathbf{T1}: \begin{cases} x^{\alpha}_{\mathbf{T1}} = h - \frac{h-m}{\alpha} & \text{if} \quad \alpha \geq \frac{h-m}{h-l} \\ x^{\alpha}_{\mathbf{T1}} = l & \text{if} \quad \alpha < \frac{h-m}{h-l} \end{cases}, \quad \mathbf{T2}: \alpha\, \frac{h - x^{\alpha}_{\mathbf{T2}}}{h-l} = \frac{x^{\alpha}_{\mathbf{T2}} - m}{x^{\alpha}_{\mathbf{T2}} - l}$$

$$\mathbf{T3}: \begin{cases} x^{\alpha}_{\mathbf{T3}} = l + \frac{m-l}{\alpha} & \text{if} \quad \alpha \geq \frac{m-l}{h-l} \\ x^{\alpha}_{\mathbf{T3}} = h & \text{if} \quad \alpha < \frac{m-l}{h-l} \end{cases}$$

B8 Structural estimation of the α-cursed equilibrium

In each trial, subject i chooses between s and p. In treatment **Tk**, the expected payoff of action $a_i = \{s, p\}$ given a value x_i is:

$$U^{\alpha}_{\mathbf{Tk}}(a_i; x_i) = \Pr(a_j = p)u^{\alpha}_{i,\mathbf{Tk}}(a_i, p; x_i) + \Pr(a_j = s)u^{\alpha}_{i,\mathbf{Tk}}(a_i, s; x_i)$$

but the perceived payoff is:

$$V^{\alpha}_{\mathbf{Tk}}(a_i; x_i) = U^{\alpha}_{\mathbf{Tk}}(a_i; x_i) + \epsilon_{i,a_i}$$

The probability that i chooses option s over p in treatment **Tk** given x_i is therefore:

$$\begin{aligned} P^{s}_{i,\mathbf{Tk}}(x_i) &= \Pr[U^{\alpha}_{\mathbf{Tk}}(s; x_i) + \epsilon_{i,s} > U^{\alpha}_{\mathbf{Tk}}(p; x_i) + \epsilon_{i,p}] \\ &= \Pr[\epsilon_{i,p} - \epsilon_{i,s} < U^{\alpha}_{\mathbf{Tk}}(s; x_i) - U^{\alpha}_{\mathbf{Tk}}(p; x_i)] \end{aligned}$$

Let us assume that error terms are independent and follow an extreme value distribution with cumulative density function:

$$F_i(\epsilon_{i,a_i}) = \exp(-e^{-\lambda\epsilon_{i,a_i}}) \quad \text{with} \quad \lambda > 0 \quad \text{for all} \quad a_i = \{s, p\}$$

For any λ, the probability that subject i with value x_i chooses option s over p in **Tk** is the logistic function:

$$P^s_{i,\mathbf{Tk}}(x_i) = \frac{1}{1 + e^{-\lambda\left(U^\alpha_{\mathbf{Tk}}(s;x_i) - U^\alpha_{\mathbf{Tk}}(p;x_i)\right)}}$$

- In **T1**, $u^\alpha_{i,\mathbf{T1}}(s, s; x_i) = u^\alpha_{i,\mathbf{T1}}(p, s; x_i)$ and $u^\alpha_{i,\mathbf{T1}}(p, p; x_i) = m$. Therefore:

$$P^s_{i,\mathbf{T1}}(x_i) = \frac{1}{1 + e^{-\lambda\left(\Pr(a_j=p)\left(u^\alpha_{i,\mathbf{T1}}(s,p;x_i) - u^\alpha_{i,\mathbf{T1}}(p,p;x_i)\right)\right)}} = \frac{1}{1 + e^{-\lambda\left(F(x^\alpha_{\mathbf{T1}})(u^\alpha_{i,\mathbf{T1}}(s,p;x_i) - m)\right)}}$$

With our parameters, if $\alpha \leq \frac{5}{9}$, $x^\alpha_{\mathbf{T1}} = 1$. If $\alpha > \frac{5}{9}$, $x^\alpha_{\mathbf{T1}} = 10 - \frac{5}{\alpha}$ and there are two cases:
(i) If $x_i \leq 10 - \frac{5}{\alpha}$, then $F(x^\alpha_{\mathbf{T1}})(u^\alpha_{i,\mathbf{T1}}(s, p; x_i) - m) = \frac{4}{9}\left[x_i - 10 + \frac{5}{\alpha}\right]$
(ii) If $x_i > 10 - \frac{5}{\alpha}$, then $F(x^\alpha_{\mathbf{T1}})(u^\alpha_{i,\mathbf{T1}}(s, p; x_i) - m) = \frac{9\alpha-5}{9}\left[x_i - 10 + \frac{5}{\alpha}\right]$

- In **T2**, $u^\alpha_{i,\mathbf{T2}}(s, s; x_i) = u^\alpha_{i,\mathbf{T2}}(p, s; x_i)$ and $u^\alpha_{i,\mathbf{T2}}(p, p; x_i) = m$. Therefore:

$$P^s_{i,\mathbf{T2}}(x_i) = \frac{1}{1 + e^{-\lambda\left(\Pr(a_j=p)\left(u^\alpha_{i,\mathbf{T2}}(s,p;x_i) - u^\alpha_{i,\mathbf{T2}}(p,p;x_i)\right)\right)}} = \frac{1}{1 + e^{-\lambda\left(F(x^\alpha_{\mathbf{T2}})(u^\alpha_{i,\mathbf{T2}}(s,p;x_2) - m)\right)}}$$

We know that $F(x^\alpha_{\mathbf{T2}}) = 1 - \frac{x^\alpha_{\mathbf{T2}} - m}{\alpha(x^\alpha_{\mathbf{T2}} - l)}$ which, given the parameters of the experiment, can be rewritten as: $\frac{x^\alpha_{\mathbf{T2}} - 1}{9} = 1 - \frac{x^\alpha_{\mathbf{T2}} - 5}{\alpha(x^\alpha_{\mathbf{T2}} - 1)}$. Solving this equation, we get:

$$x^\alpha_{\mathbf{T2}} = \frac{11\alpha - 9 + \sqrt{(11\alpha - 9)^2 - 4\alpha(10\alpha - 45)}}{2\alpha} \quad \forall\alpha > 0$$

(i) If $x_i \leq x^\alpha_{\mathbf{T2}}$, $F(x^\alpha_{\mathbf{T2}})(u^\alpha_{i,\mathbf{T2}}(s, p; x_i) - m) = \frac{x^\alpha_{\mathbf{T2}} - 1}{9}\left[\alpha\left(\frac{x_i - 1}{9}x_i + \frac{10 - x_i}{9}\right) + (1 - \alpha)\left(\frac{x_i - 1}{x^\alpha_{\mathbf{T2}} - 1}x_i + \frac{x^\alpha_{\mathbf{T2}} - x_i}{x^\alpha_{\mathbf{T2}} - 1}\right) - 5\right]$

(ii) If $x_i > x^\alpha_{\mathbf{T2}}$, $F(x^\alpha_{\mathbf{T2}})(u^\alpha_{i,\mathbf{T2}}(s, p; x_i) - m) = \frac{x^\alpha_{\mathbf{T2}} - 1}{9}\left[x_i - \alpha\left(\frac{(x_i - 1)(10 - x_i)}{9}\right) - 5\right]$

- In **T3**, $u^\alpha_{i,\mathbf{T3}}(s, p; x_i) = u^\alpha_{i,\mathbf{T3}}(p, s; x_i) = u^\alpha_{i,\mathbf{T3}}(p, p; x_i) = m$. Therefore:

$$P^s_{i,\mathbf{T3}}(x_i) = \frac{1}{1 + e^{-\lambda\left(\Pr(a_j=s)\left(u^\alpha_{i,\mathbf{T3}}(s,s;x_i) - u^\alpha_{i,\mathbf{T3}}(p,s;x_i)\right)\right)}} = \frac{1}{1 + e^{-\lambda\left((1 - F(x^\alpha_{\mathbf{T3}}))(u^\alpha_{i,\mathbf{T3}}(s,s;x_i) - m)\right)}}$$

With our parameters, if $\alpha \leq \frac{4}{9}$, $x^\alpha_{\mathbf{T3}} = 10$. If $\alpha > \frac{4}{9}$, $x^\alpha_{\mathbf{T3}} = 1 + \frac{4}{\alpha}$ and there are two cases:
(i) If $x_i > 1 + \frac{4}{\alpha}$, then $(1 - F(x^\alpha_{\mathbf{T3}}))(u^\alpha_{i,\mathbf{T3}}(s, s; x_i) - m) = \frac{5}{9}\left[x_i - 1 - \frac{4}{\alpha}\right]$
(ii) If $x_i < 1 + \frac{4}{\alpha}$, then $(1 - F(x^\alpha_{\mathbf{T3}}))(u^\alpha_{i,\mathbf{T3}}(s, s; x_i) - m) = \frac{9\alpha-4}{9}\left[x_i - 1 - \frac{4}{\alpha}\right]$

- In **T4**, $P^s_{i,\mathbf{T4F}}(x_i) = P^s_{i,\mathbf{T4S}}(x_i) = P^s_{i,\mathbf{T1}}(x_i)$.

For the estimation, we look at each grade group and treatment separately. For each individual and for each x_i, λ and α, we compute $x^\alpha_{\mathbf{Tk}}$ and $P^s_{i,\mathbf{Tk}}(x_i)$. Then we compute the likelihood of the entire data given (α, λ), and look for the values of α and λ that maximize the likelihood of the observed frequencies of actions.

Finally, Figure B6 reports the empirical c.d.f. ($S_c(x)$, as in Figure 4) and the estimated best fit of the BRUM.

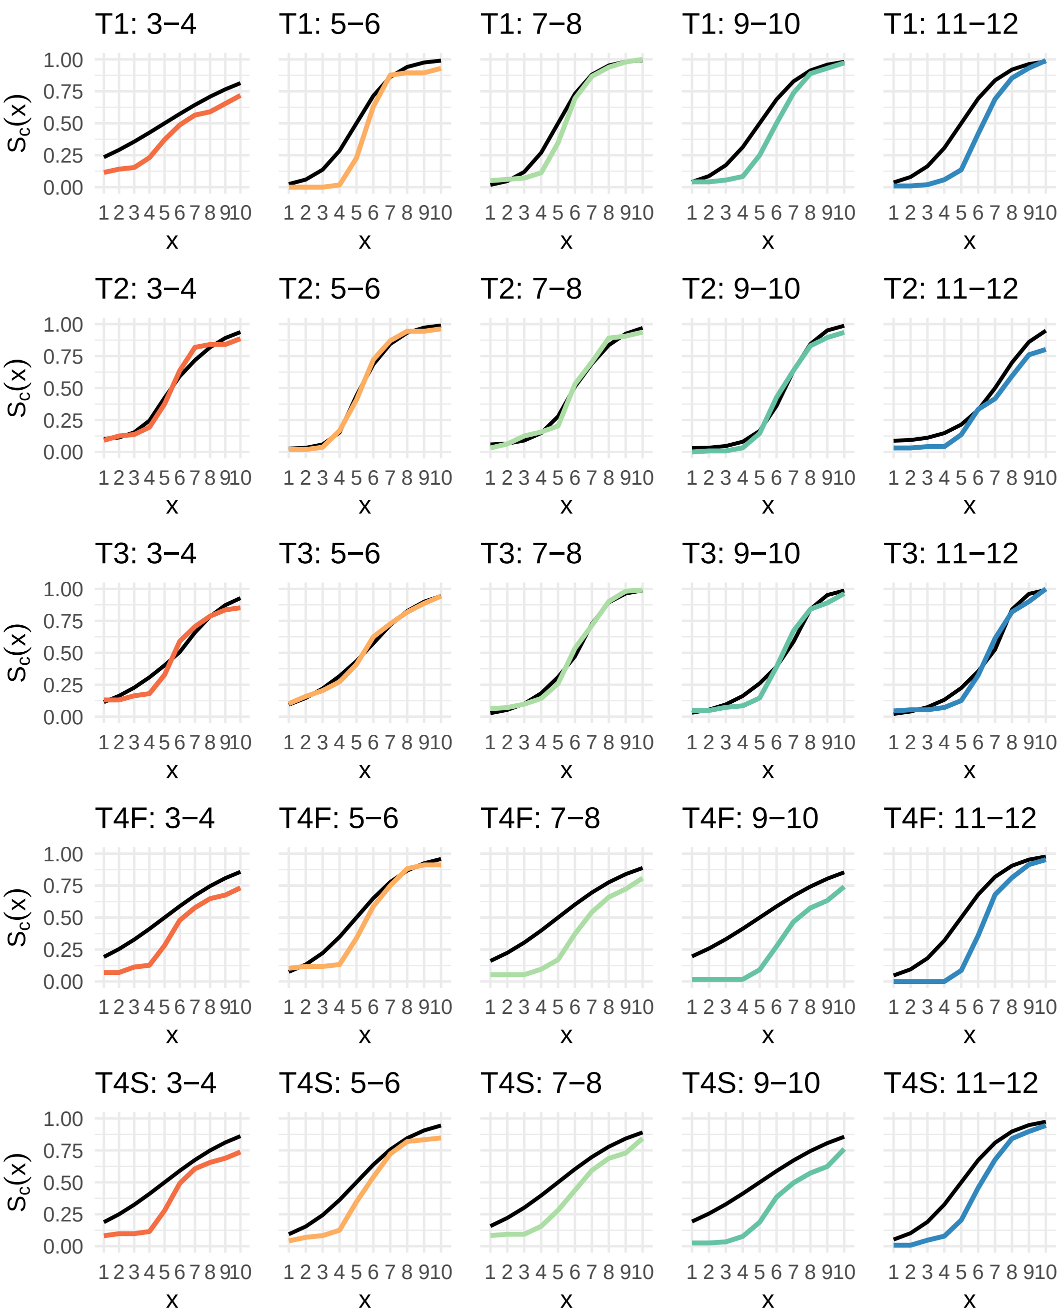

Figure B6: Best fit (α, λ) of BRUM for strategic consistent participants (C_+)