The development of rationality in games with hidden information [*]

Antonio Alfonso
U Loyola

Pablo Brañas Garza
U Loyola

Isabelle Brocas
USC and CEPR

Juan D. Carrillo
USC and CEPR

Juan B. González
USC

María José Vázquez
Fundación Etea

October 2023

Abstract

We present the first lab-in-the-field experiment that studies the strategic behavior of privately informed children and adolescents. We recruited 1662 participants from 8 to 18 years old to play a game of two-sided asymmetric information. We show that participants of all ages understand the fundamental relationship between action and private information (first level of rationality or 'choice monotonicity'). Older participants are more likely to select strategies that match basic features of the optimal strategy compared to their younger peers (second level of rationality or 'strategic consistency'). However, in none of our grades the individuals react to variations in incentives triggered by changes in game structure (third level of rationality or 'environmental variability'). Remarkably, participants with heightened emotional intelligence exhibit a greater tendency to play strategically, best respond to others and, consequently, achieve higher payoffs. It reveals a strong, robust connection between affective theory-of-mind and cognitive theory-of-mind.

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 project PID2021-126892NB-100, and the European Project Horizon GAP-10-1095175.

1 Introduction

Scenarios involving private information are notoriously challenging to grasp. Thorough examinations of both controlled 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 results apply in other asymmetric information contexts, such as the takeover game (Grosskopf et al., 2007; Bereby-Meyer and Grosskopf, 2008), the betting game (Sonsino et al., 2002), and asset trading (Carrillo and Palfrey, 2011).

Researchers have employed various experimental methodologies to understand the causes of deviations from optimal behavior. In the takeover game, Charness and Levin (2009) transforms the game into an individual decision making problem, simplifying it to just two states while Martínez-Marquina et al. (2019) converts the probabilistic framework into a deterministic one. For the case of the betting game, Brocas et al. (2014) utilizes non-choice measures (mousetracking) to determine whether deviations are the result of choice errors or cognitive constraints.

In this paper, we present a private information game that allows us to evaluate the behavior of participants across three distinct levels of rationality and introduce two novelties to examine the intricacies of asymmetric information. First, we provide the first attempt in the literature to trace the *developmental trajectory* of behavior in a game with private information. Our study encompasses a broad age spectrum and involves a large number of individuals (1662 participants ranging from 8 to 18 years old). The objective is to examine how age-related differences impact our comprehension of the different levels of rationality, distinguishing between those that are inherent, learned, or remain undeveloped. By doing so, we seek to gain deeper insights into the challenges that individuals, even those who have reached their cognitive maturity, continue to face in these games. Second, we relate performance in our highly cognitive game with *the ability to recognize and empathize with emotions*, which we evaluate using a novel affective theory of mind task. The goal is to determine if individuals with a heightened ability to perceive others' emotions 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]). It has an intuitive setup that can be illustrated using the following stylized example. Two parties are engaged in a dispute and they each know privately the strength of their case. They have the option to either settle the dispute or go to litigation. In the case of a settlement, both parties receive intermediate payoffs. If they proceed to litigation, the strengths of their respective cases are revealed. The party with the stronger case wins, receiving a high payoff, while the party with the weaker case loses, receiving a low payoff. We consider four treatments that share the same basic structure but differ in the way outcomes and payoffs are computed. In the benchmark case, settlement requires mutual agreement and the high, intermediate and low payoffs are fixed. We then manipulate the winner's payoff, conditions for settlement, and sequence of moves. The logic is the same in all treatments. Participants must always adopt a 'threshold strategy': seek litigation (action 'see' in our experiment) when the value of their case exceeds a certain threshold, and seek a settlement (action 'pass' in our experiment) otherwise. Importantly, both the theoretical equilibrium threshold and the empirical best response threshold vary across treatments, depending on the rules of the game. In section 2, we present various applications that align with the fundamental characteristics of the different treatments and establish the equilibrium for each case.

In our empirical analysis, we classify choices into *three hierarchical levels of rationality*. The first and most fundamental level, called *choice monotonicity*, posits that individuals are more inclined to choose 'see' when their private value (strength of their case) is higher. The second level, called *strategic consistency*, represents a refinement of the first level. It stipulates that individuals employ threshold strategies: if they choose 'see' for a specific value, they never switch to 'pass' when the value is higher. Importantly, a threshold strategy is optimal for any strategy of the rival, indicating that this level does not require to understand how others act on their private information. 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 level, called *environmental variability*, outlines that the observed distribution of cutpoints utilized by participants will vary between treatments. This variation arises because the equilibrium cutpoint is influenced by treatment-specific incentives of players, including their comprehension of how a rival will employ their private information in that particular treatment. Achieving this level of rationality is cognitively demanding because it necessitates an optimal response to subtle shifts in incentives, and it can only be examined through our innovative experimental design featuring multiple treatments. By discerning which level of rationality is mastered at each age, we can identify the cognitive challenges associated with each aspect of the decision-making process.

The experiment provides interesting insights into age-related shifts in rationality. First, we show that participants exhibit choice monotonicity from an early age: participants in all grades align their choices with the principle that higher values warrant selecting the 'see' option. Second, the study finds that strategic consistency is prominent even among the youngest participants (one-half of participants aged 8 to 10 employ threshold strategies). This tendency towards adopting threshold strategies grows steadily with age, although participants do not reach full compliance (three-quarters of participants aged 16 to 18 employ threshold strategies). Third, the study reveals that environmental variability is absent across all ages. This absence persists even when examining only the subset of the oldest participants who employ threshold 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. Finally, we estimate a cursed equilibrium model (Eyster and Rabin, 2005) and show that this behavioral theory captures some important elements of the behavior of our participants (partial unraveling of equilibrium cutpoint) but not others (lack of environmental variability).

Equally noteworthy, the experiment uncovers 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 threshold strategies more frequently, they are more likely to best respond to the strategies of others and, consequently, they obtain significantly higher payoffs. Overall, this finding suggests that the capacity to physically perceive and interpret emotions is closely intertwined with the ability to mentally decipher intentions—a pivotal aspect of optimal decision-making in strategic games. 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. First, it contributes to the growing research on decision-making in children and adolescents. To the best of our knowledge, prior game theoretic experiments have concentrated exclusively on games with complete information (see Sutter et al. (2019) and List et al. (2021) for detailed surveys). In contrast, our study aims to investigate how age-related changes manifest in games where individuals must base their choices on their private information.[1] Second, our paper contributes to the existing body of research that examines the positive association between cognitive ability, often assessed through measures of fluid intelligence, and performance in strategic contexts (Brañas-Garza et al., 2012; Proto et al., 2019, 2022; Fe et al., 2022). Consistent with the findings in this literature, our older participants, who typically possess more advanced cognitive abilities, exhibit superior performance in certain facets of the game, notably in terms of strategic consistency. However, our paper deviates from the conventional approach by contending that, after controlling for age, emotional intelligence, as measured through an affective theory of mind task, also plays a very significant role in performance. This perspective underscores the importance of affective mechanisms alongside cognitive abilities in shaping decision-making outcomes in strategic settings.

The paper is organized as follows. Section 2 presents the game and solves for the equilibrium. Section 3 describes the experiment in detail 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, across treatments and across decisions. Section 6 investigates potential reasons for the observed deviations, and provides a test for a behavioral theory. In section 7, we discuss the relationship between affective theory of mind and performance in the game, which we associate to cognitive theory of mind. Section 8 gathers concluding remarks.

2 Theory

2.1 Game and variants

We consider an extended version of [CP], a two-player game with two-sided private information. Two players, 1 and 2, are indexed by i and j. Player i has private value x_i drawn from distribution $F_i(x_i \,|\, x_j)$. $F_i(\cdot)$ and $F_j(\cdot)$ can be different but for simplicity have identical support $[\underline{x}, \overline{x}]$. To prevent ties, we also impose $x_i \neq x_j$. Player i chooses an action from a binary space $a_i \in \{p, s\}$, 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.

[1] Brocas and Carrillo (2022b) study the problem of adverse selection in Charness and Levin (2009)'s adaptation of the takeover game. The game, however, has no strategic considerations since it is cleverly transformed into an individual decision making problem. Most importantly, the decision-maker has no private information, so it is not possible to analyze strategies as mappings between information and action.

Benchmark.

T1. Players move simultaneously. If both choose ‘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:

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

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.

Many situations arise in which individuals holding conflicting objectives and possessing private information find themselves in a position to either compete for the high payoff or compromise on the medium one. [CP] –who study a version of **T1** and **T4**– discuss the examples of 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). In these examples, 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 (initiating a conflict, a legal dispute, or an innovation race requires only one party).

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

Finally, and as modeled in **T3**, compromise serves sometimes as the default outcome and it is overruled only if both players choose action 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.

As we will see below, these subtle differences between **T1**, **T2**, and **T3** have dramatic consequences for equilibrium behavior. Studying different treatments allows us to determine whether individuals react to elements of the environment.

2.2 Equilibrium

Despite the simplicity of the game, the equilibrium solution is non-trivial. Proposition 1 summarizes the properties of the Bayesian Nash Equilibrium (**T1**, **T2**, **T3**) and Perfect Bayesian Equilibrium (**T4**) in each variant of the game.

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 is: (i)* $x^*_{\mathbf{T1}} = x^*_{\mathbf{T4}} = \underline{x}$, *(ii)* $x^*_{\mathbf{T2}} = m$, *and (iii)* $x^*_{\mathbf{T3}} = \overline{x}$.

<u>Proof</u>. **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 experimental implementation we will 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 threshold strategy.

Step 2. Assume player j uses a threshold strategy with cutpoint 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 $x_i^* = x_j^* = m$.[2] 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**. □

For *any* strategy chosen by the rival (threshold or not), it is always optimal to best respond with a threshold strategy, simply because the participant's expected payoff increases more steeply with their private value if they choose s than if they choose p.

However, players must realize that their action is sometimes inconsequential, and they must condition on the choice of the rival to compute the optimal cutpoint. Because their action is relevant in different scenarii –when the rival chooses p in **T1**, **T2** and **T4**, and when the rival chooses s in **T3**– the location of the optimal cutpoint depends on treatment specific features. Specifically, the equilibrium cutpoint unravels downwards to $\underline{x}$ in **T1** and **T4**, unravels downwards but only to m in **T2** and unravels upwards to $\overline{x}$ in **T3**.

3 Experiment

3.1 Design and procedures

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).

[2]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.

Game. To simplify the setting for young participants, we modify the game outlined in section 2.1 in the following manner. 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, which implies that types are correlated and ties are not possible. Participants can 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 is straightforward to explain, suppressing the need to delve into distributional considerations.

We employ the strategy method to elicit behavior. It is well-known that the strategy method may in some environments 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 value 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 threshold (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 threshold 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 each type of strategy.[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]

[3]Naturally, 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, then $a_i(1) = s$ is weakly dominated by $a_i(1) = p$ in **T1** and **T4**, $a_i(5) = s$ is weakly dominated by $a_i(5) = p$ in **T2**, and $a_i(10) = p$ is weakly dominated by $a_i(10) = s$ in **T3**. This means that $x^*_{\mathbf{T1}} = x^*_{\mathbf{T4}} = 1$, $x^*_{\mathbf{T2}} = 5$ 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.

One notable strength of the design is that the presentation, instructions, cognitive requirements, and logical inferences involved in this game 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 of the game. 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 during the game, called the first and second decisions. More precisely, initially, we request them to report a strategy to be employed in six consecutive matches against six distinct opponents. We then utilize the random draws (x_i, x_j) and present for each of the six games the cards drawn by both players, the choices of both players 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 are required to 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.

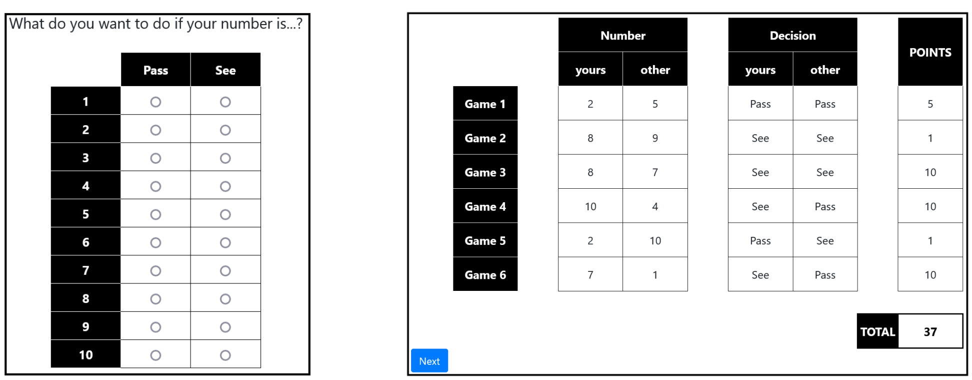

Figure 1: Screenshots. Strategy elicitation (left) and summary of outcomes (right) in **T1**

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

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

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 and the rationale behind its selection are presented in section 7.

Timing. The experiment consists of four tasks programmed in 'oTree' (Chen et al., 2016) and implemented on touchscreen PC tablets through a wireless closed network. Due to the challenge of accessing a very large population of children, we combined two projects. Initially, we conducted a third-party dictator game, the results of which are analyzed in a separate article (Brocas and Carrillo, 2022a). After a short 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, for obvious reasons. While the two games exhibit sufficient differences to mitigate concerns about cross-contamination, we took precautionary measures to ensure experimental integrity. Specifically, we consistently followed the same order, conducting 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. We concluded the session by administering a questionnaire in which we collected information about participants' age and gender.

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. Specifically, we recruited a diverse cohort of 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 can already be classified as young adults (18 years old), we made a deliberate departure from our prior research practices and opted not to include a control group of undergraduate participants, who would only be slightly older (ranging

[7]The Salesianos educational network, with a global presence in 132 countries, is a missionary school system. 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`).

from 18 to 22 years old) and possess a distinct background. 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

We visited each school and conducted the experiment one class at a time, utilizing a portable lab stationed in a dedicated classroom. Participants were organized into small subgroups comprising 8 to 12 individuals, each group led by one experimenter. The experiment was integrated into the school's academic activities, and our Institutional Review Board (IRB) stipulated an opt-out procedure. Importantly, no children opted out, thereby eliminating any potential selection bias. The only children who did not participate were those who were absent on the day of the experiment. The overall participation rate stood at 92%. For the sake of 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 often be challenging. As outlined in Brocas and Carrillo (2020), an appropriate approach typically involves employing distinct payment mediums tailored to different grades, such as toys for children and money for teenagers, with the goal of equalizing the value of rewards across individuals rather than providing identical rewards. 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€ and a sports drink for 1.10€), is highly popular, and widely accessible during recess and lunchtime. This 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 leverages three novel features to examine behavior in games of private information. Here, we address how each feature provides unique insights.

Multiple treatments. [CP] demonstrated that the behavior of adults from elite institutions at the peak of their cognitive capacities (Princeton undergraduates) is consistent with threshold strategies but does not adhere to equilibrium. It indicates that adults correctly reason about the game's features but do not necessarily condition their decisions on their rival's behavior. Since [CP] focus on treatments that predict the same equilibrium (**T1** and **T4**), it is challenging to fully ascertain the underlying causes of this result. For instance, participants may not realize that their threshold should be lower than that of their rival. Alternatively, they may not fully carry out the recursive thinking process.

To better understand this phenomenon, it is crucial to compare across treatments that predict different equilibria. By introducing **T2** and **T3**, we can evaluate whether participants are responsive to the specific incentives presented in each treatment and adjust their strategies in the direction predicted 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, games with incomplete information are arguably the most challenging ones. To our surprise, no study to date has examined age-related changes in when subjects hold private information. Taking a developmental perspective in a game characterized by hierarchical levels of rationality allows us to discern which facets of reasoning are mastered at different ages, and which abilities never fully mature. It also enables us to uncover the algorithms employed by individuals when devising strategies.

Affective mechanisms. Games with two-sided private information necessitate the formation of beliefs about others, a skill that falls under the purview of theory-of-mind. Cognitive theory of mind (c-ToM) is centered on understanding the cognitive processes and thoughts that underlie human behavior, while affective theory of mind (a-ToM) involves recognizing and empathizing with emotions –a component closely linked to emotional intelligence. Private information games clearly demand the cognitive ability to decipher the intentions of others. By investigating variations in affective mechanisms and their correlation with our game, our objective is to ascertain whether and to what extent affective abilities also contribute to performance in complex strategic environments.

3.3 Hypotheses

Behavior within our game can be categorized into hierarchical levels of rationality. At its fundamental level lies choice monotonicity: the likelihood of choosing to "see" in a grade group should increase as the private value rises. Moving to an intermediate level, we should observe strategic consistency: individuals should employ threshold strategies across all treatments, regardless of their beliefs about others' behavior. Finally, the highest level of sophistication prescribes environmental variability: the optimal cutpoint should be contingent on the game's structure. Correctly responding to the specific features of the environment requires conditioning on the choice of the rival, that is, understanding the relationship between their action and their information, the hallmark of asymmetric information games. Given the diversity of ages within our sample, we have the opportunity to identify the age at which participants' behavior aligns with each level of rationality. Consequently, we put forth several hypotheses concerning the age-related shifts in behavior.

Hypothesis 1. *'Choice monotonicity' is satisfied in all grade groups: 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 threshold 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. $F(x^*_{\mathbf{T1}})$.*

In line with Hypothesis 1, we anticipate that all participants, including the youngest ones, will grasp the fundamental principles of the game and opt for "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. Regarding Hypothesis 2, it implies that young children may struggle to recognize the optimality of a threshold strategy. Instead, many may adopt strategies involving multiple switches. As they develop, the prevalence of this crucial aspect of optimal decision-making is expected to increase, although this improvement will likely be gradual and incomplete. In essence, we expect this property to be neither insurmountable for our younger participants nor trivially understood by our older ones. In accordance with Hypothesis 3, our more cognitively developed participants (those employing threshold strategies) are likely to grasp the intricacies of the treatment assigned to them (e.g., whether "pass" or "see" necessitates agreement). Consequently, we should observe a shift in cutpoints across treatments as they adapt to the subtleties of 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 game better: they use threshold strategies more often and obtain higher payoffs.*

Hypothesis 4 posits that the ability to read others' emotions is intertwined with the ability to infer their intentions, a pivotal aspect of strategic decision-making. A more detailed motivation for this hypothesis is provided in section 7.

While Hypotheses 1 through 4 address the most fundamental questions in this study, additional issues of interest warrant exploration. First, we expect a positive though limited amount of learning. This expectation aligns with findings from previous asymmetric information games where feedback is limited and counterfactuals are subtle (Selten et al., 2005; Grosskopf et al., 2007). As in those studies, the structure of our game does *not* readily facilitate extensive learning.[8]

Second, in line with the findings in [CP], we expect that individuals who employ threshold strategies in the sequential treatment will tend to play closer to equilibrium after the rival's choice is determined (second mover) than when they need to anticipate what is coming next (first mover). Formally, this can be expressed as: $F(x^*_{\mathbf{T4F}})$ F.O.S.D.

[8] To illustrate this point, suppose for example that i chooses $a_i(3) = p$ in **T1** (against the theoretical prediction). If j chooses s, there is nothing i could have changed. If j chooses p and $x_j > 3$, then i is glad of their choice. Only if j chooses p and $x_j < 3$ will i learn their 'mistake'.

$F(x^*_{\mathbf{T4S}})$ in **T4**. The rationale behind this expectation is straightforward: it is typically easier to base a strategy on an actual observed outcome than on a hypothetical one.

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 $S(x)$. Remember that we use the strategy method and we elicit strategies twice during the game. Therefore, we have two measures $a_i(x) \in \{p, s\}$ for each individual at each value x. Consequently, we have a total of 20 observations per individual. The results are presented in Figure 2.

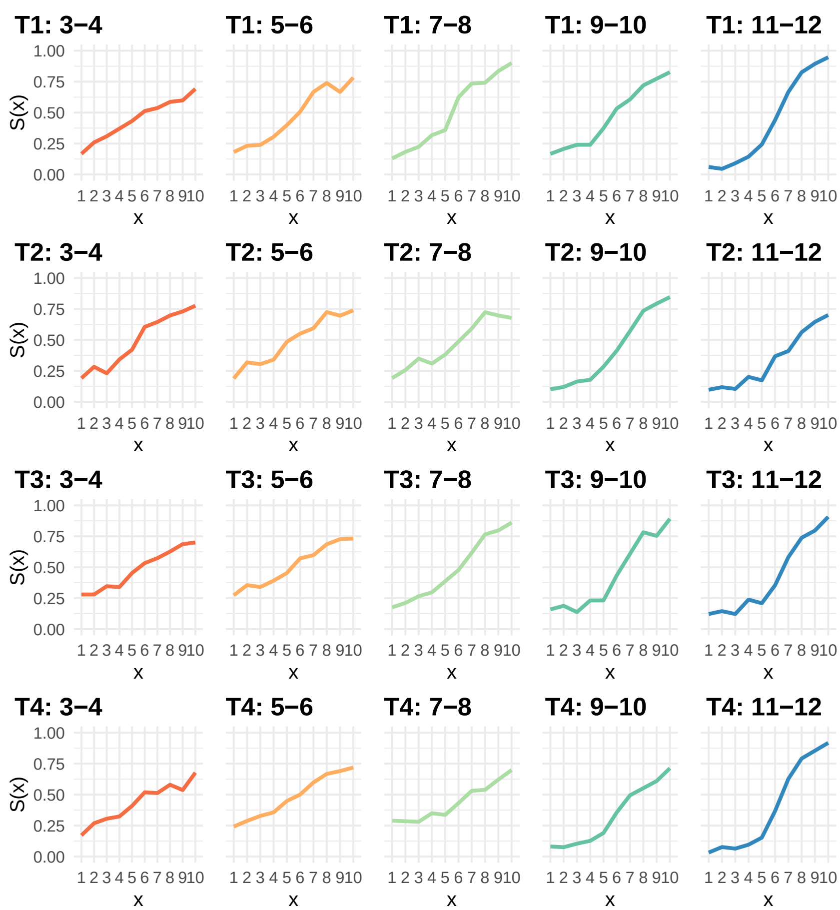

Figure 2: Average probability of choosing to see ($S(x)$) by treatment and grade group

Strongly supporting Hypothesis 1, we observe a monotonic increase in the probability of choosing s as the 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 choosing to '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.

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 3 (error bars represent the 95% confidence interval).

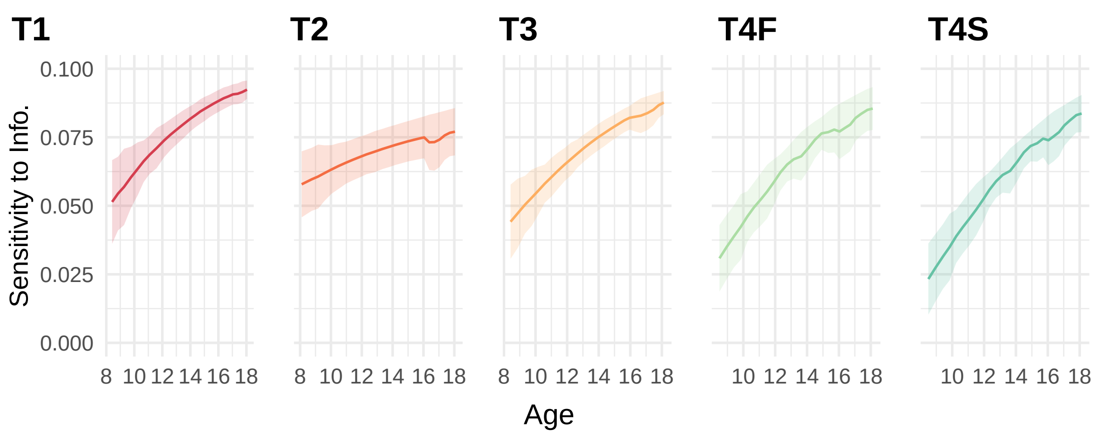

Figure 3: Sensitivity of action to private value by treatment and age

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.

Finally, Table 2 reveals an increased inclination to choose 'see' in the second decision (*2ndDec*). In Appendix B1, we investigate these changes in behavior in more detail. The analysis confirms only a modest tendency to increase the choice of s. Furthermore, 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 significant changes, we will combine data from both decisions for the remainder of the paper, unless otherwise specified.

The findings of this section are summarized as follows.

Result 1. *Hypothesis 1 is supported by the data. Choice monotonicity ($S(x)$ is 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 now turn to examine the strategies adopted by participants at different ages. We categorize these strategies into different classes and calculate the proportion of participants who adhere to each category. A potential caveat of this method is the possibility of overlooking certain strategies. To mitigate this risk, we contemplate a large set of options, even including some unconventional strategies that may be suboptimal and/or unexpected. Table 3 provides a detailed 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 ones (RC_0, RC_1).[9] It also includes strategies that fall in between (INT_0, INT_1), as well as

[9] One could argue that aP and aS are suboptimal strategies since $a_i(1) = s$ and $a_i(10) = p$ are weakly dominated actions. On the other hand, and as discussed in footnote 4, they are part of the equilibrium set. We take a cautious approach and refer to aP and aS as "boundary threshold strategies."

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

some unconventional patterns (ALT_0, ALT_1, PAT) that cannot be readily reconciled with any reasonable theoretical framework.

We categorized the behavior of our participants 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. We therefore opted for not retaining these strategies. For visual ease, we also decided to group ALT_0 and ALT_1 under a single category, denoted as ALT. This leaves us with a total of six strategies: aP, aS, C_0, C_1, ALT and Other (for participants not classified under any of the other five categories). Figure 4 reports the empirical frequency of each strategy categorized by treatment, grade group, and decision.

The distribution of strategies appears very similar between treatments **T1**, **T2** and **T3**, as shown in Figure 4(A) (although a paired χ^2-test of differences of proportion finds a statistically significant difference between **T2** and **T3**, $p < 0.01$). In the sequential treatment, we observe a higher frequency of the 'always pass' strategy, often at the expense of interior cutpoint strategies. The differences between **T4** and the other treatments are all statistically significant (χ^2-test, $p < 0.01$) whereas differences between **T4F** and **T4S**

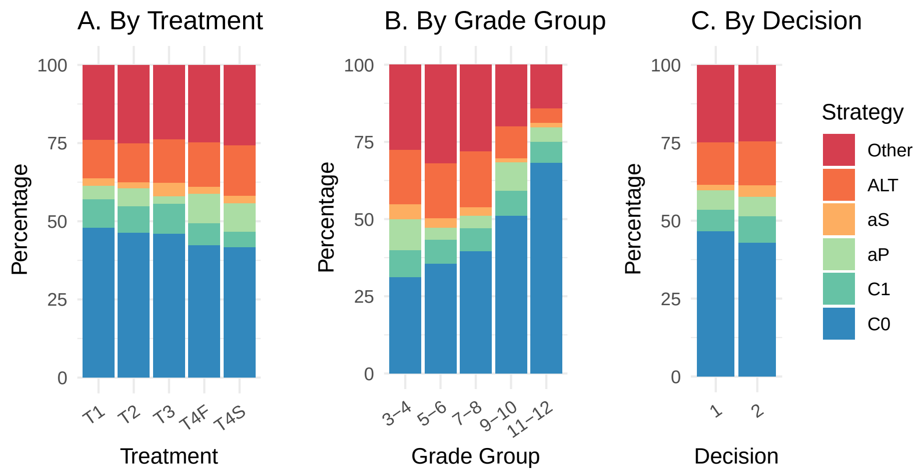

Figure 4: Individual strategies by treatment (A), grade group (B) and decision (C)

are not (χ^2-test, p = 0.51). More generally, let us call $C_+ = aP + aS + C_0$ the set of threshold strategies that perfectly align with an interior cutpoint (C_0) or a boundary cutpoint (aP + aS). We observe that more than half the observations in all treatments are classified as C_+. The proportion increases to around 60% when we allow one mistake (add C_1). Additionally, there is also a significant fraction of alternation (12% to 16%) and unclassified (24% to 26%) strategies.

By contrast, there is significant variation in strategies across grade groups, as evident in Figure 4(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. Age effects disappear in elementary and middle school if we consider boundary cutpoints and allow for one mistake (55% in 3-4, 51% in 5-6, and 55% in 7-8). This occurs primarily because our youngest participants tend to opt for the 'always pass' strategy. However, discrepancies reemerge in the behavior of older players, with 69% in 9-10 and 82% in 11-12 adhering to strategic consistency. At the same time, the level of strategic consistency is higher than initially anticipated. We hypothesized few threshold strategies among our youngest participants. Instead, 46% of children in grades 3-4 *perfectly* align with either an interior or a boundary cutpoint, despite the technical challenges involved in reporting a strategy.[10] Finally, differences between first and second

[10]We believe this number is high. To put it in perspective, notice that educated adults in much simpler

decision are very minor, albeit statistically significant (χ^2-test, $p < 0.01$), see Figure 4(C). We observe a slight increase in mistakes among interior cutpoint strategies (from C_0 to C_1) and a tendency to 'see always' (aS).

5.2 Threshold strategies across ages and decisions

We further investigate the decision to use threshold strategies across ages. More specifically, 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 threshold strategy (C_+). The results are presented in Table 4.[11]

	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 threshold 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 threshold strategy with increasing age. However, the effect is not statistically significant in **T2**. This finding is surprising, particularly considering that choosing s in **T2** is a weakly dominated strategy when $x \leq 5$. Notice that males are more likely to use threshold strategies. This effect of gender will be explored further in section 7.3. 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 distinct and consistent learning patterns.

settings (e.g., elicitation of time or risk preferences using multiple price lists) engage in inconsistent behavior that takes the form of multiple switches 15%-20% of the time.

[11] We obtain very similar results both if we add C_1 to the set of threshold strategies as well as if we remove aS and aP (that is, we only consider C_0, the interior threshold strategies).

Finally, in Appendix B2, we perform a Probit regression similar to the one in Table 2, except that we divide the sample into individuals who use perfect threshold strategies and those who do not. We then perform the same exercise as in Figure 3 and show that the sensitivity to information of individuals who use threshold strategies is high and relatively constant across age. By contrast, the sensitivity to information of individuals who do not employ threshold strategies is low, although it increases 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. At the same time, the fraction of individuals who use threshold strategies is higher for all ages than initially anticipated.*

5.3 Threshold strategies across treatments

We next focus on choices that align perfectly with interior or boundary threshold strategies (C_+), and hypothesized that participants adhering to strategic consistency would conform to environmental variability. In other words, the ordering of cutpoints across treatments predicted by theory should be preserved in our sample, in a stochastic sense. To test this hypothesis, we construct the same function as in Figure 2 but only with players who employ threshold strategies C_+. We denote this function as $S_c(x)$. This time, $S_c(x+1) \geq S_c(x)$ *by construction*. In Figure 5, we present the functions $S_c(x)$ for the simultaneous treatments **T1**, **T2**, **T3** (top row) and the sequential treatments **T4F**, **T4S** (bottom row). The functions are categorized by grade group and are also aggregated for all individuals (All).

There is significant heterogeneity in cutpoints. As conjectured in section 3.3 (and in contrast to theoretical predictions for all treatments except **T2**), only a few cutpoints are at the boundaries, while the majority (60% to 72% depending on the treatment) fall 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 significantly different only for grade group 3-4 (p = 0.04). Moreover, the sign (higher choice of s in **T3**) points in the opposite direction of what we expected.

In the oldest grade group (11-12), we observe some differences, although they do not

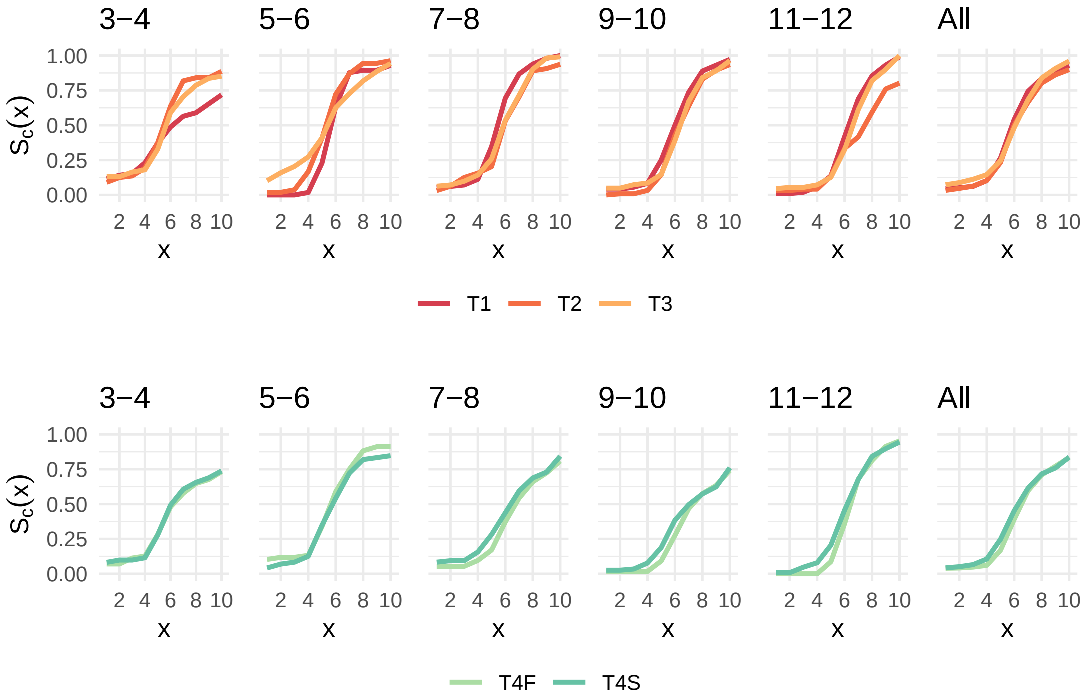

Figure 5: Comparison of $S_c(x)$ across treatments for each grade group

consistently align with our expectations. For instance, the frequency of choosing 'see' is indeed higher in **T1** compared to **T2** as predicted (KS-test, $p < 0.01$). However, it is also more frequent in **T3** compared to **T2**, contrary to our prediction (KS-test, $p < 0.01$). All other differences are not significant.[12] Overall, Hypothesis 3 is strongly rejected: distributional differences in cutpoints across treatments **T1**, **T2** and **T3** are small and inconsistent. Specifically, it appears that the significance of conditioning on the rival's choice of p in **T1** and s in **T3**—which is a critical aspect of the downward and upward unraveling dynamics in these treatments—is not grasped by any group of players, including our oldest and arguably most sophisticated participants. We conjecture that this difference would likely be overlooked by educated young adults as well.

Regarding the sequential treatment, we do not find any significant differences between

[12]Notice that we restricted attention to individuals satisfying strategic consistency (our most rational players) to increase the chance of observing environmental variability. If we include all participants, differences in $S(x)$ across treatments are still not significant (data omitted for brevity).

the first and second movers in any grade group (KS-test, $p > 0.25$). This contradicts both our initial hypothesis and the results presented in [CP]. In hindsight, we believe this result is unsurprising. [CP] argue that second movers opt for "see" more frequently than first movers because they can observe their rival's decision, make inferences about their information, and react accordingly. This reasoning, however, 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, which may explain the lack of significant differences between them.

Appendix B3 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.*

6 Understanding deviations

6.1 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 others within the same grade group is the same across treatments. Indeed, we know from Eq. (1) that optimal behavior requires a threshold strategy. However, if participants deviate from equilibrium and some resort to highly suboptimal rules (e.g., alternation between p and s), it is possible that empirically optimal cutpoints in each treatment will differ from their theoretical counterparts.

To study this issue, we compute the cutpoint $\tilde{x}$ that maximizes the empirical expected payoff for each grade group and each treatment.[13] We also determine the optimal cutpoint by pooling the data from all grades (all). Finally, and for comparison, we report the theoretical equilibrium (theory) computed in section 3. Results are compiled in Table 5.

As expected, the empirical best response cutpoint $\tilde{x}$ is never located at the boundary. However, $\tilde{x}$ approximately preserves the same order across treatments as x^*, the equilib-

[13] For **T4**, we determine the best response of a first (second) mover to the empirical behavior of second (first) movers.

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

rium cutpoint: $\tilde{x}_{\mathbf{T1}} < \tilde{x}_{\mathbf{T2}} \leq \tilde{x}_{\mathbf{T3}}$ 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**.[14] 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. Theory predicts full downward unraveling (**T1**, **T4**), partial downward unraveling (**T2**), and full upward unraveling (**T3**). The same qualitative forces hold when the empirical behavior of the population is similar across treatments. In other words, variations in treatment design lead to differences in the optimal response strategy, even when the empirical behavior is almost the same across all treatments.

6.2 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 various choices lead to similar outcomes, deviations may be attributed to an efficient conservation of cognitive effort. To assess the consequences of our participants' decisions, we calculate the expected empirical revenue loss (relative to selecting the empirically optimal cutpoint within their grade group, as determined in Table 5) averaging across individuals who employ the same strategy type. Figure 6 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 boundaries aP and aS are suboptimal when theory prescribes an intermediate cutpoint

[14] The computation assumes risk neutrality. If individuals are risk averse, optimal cutpoints will be slightly affected but the order of best response across treatments will still be preserved.

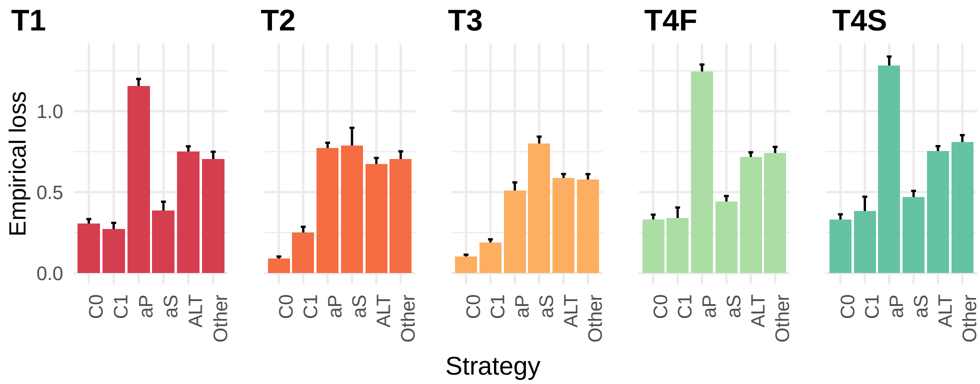

Figure 6: Empirical loss as a function of strategy

(**T2**). More generally, individuals who do not use threshold strategies (ALT and Other) experience twice the revenue loss compared to those who utilize interior threshold strategies (C_0 and C_1), even when we aggregate all interior cutpoints. The difference is similar across treatments (Figure 6) and across grades (data not reported for brevity). The difference is also substantial, especially when considering that (i) the decision of a participant is often inconsequential due to the behavior of their rival, and (ii) even if the participant made random choices, it would align with the optimal decision one-half of the time.

6.3 A behavioral theory

Since departures from Nash equilibrium are significant (section 6.1) and costly (section 6.2), 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, other regarding concerns cannot explain the behavior observed in our game. Indeed, 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, but importantly, an individual with social concerns will always adopt a boundary threshold strategy

[15] For example, consider the inequality aversion model by Fehr and Schmidt (1999), where $u_i(y_i, y_j) = y_i - \alpha \max\{y_j - y_i, 0\} - \beta \max\{y_i - y_j, 0\}$ (with $0 \leq \beta \leq \alpha$). This model would require $10 - 9\beta < 5 \Rightarrow \beta > 5/9$ for participants to deviate from the behavior predicted by the standard theory.

(either aP or aS). That is, in the absence of additional parameters or assumptions, these models cannot effectively predict the interior cutpoints widely observed in practice.

6.3.1 α-cursed equilibrium: theory

We decided to study instead α-cursed equilibrium (α-CE), a behavioral theory based on an imperfect understanding of the relationship between information and action (Eyster and Rabin, 2005). As researchers have shown, this prominent theory partly explains the deviations of individuals in several games of asymmetric information, including trade, adverse selection and common value auctions. 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.

Denote $x^{\alpha}_{\mathbf{Tk}}$ the equilibrium threshold in treatment k under α-CE (i.e., the equivalent of $x^{*}_{\mathbf{Tk}}$ for our behavioral theory). The main properties of α-CE are (see Appendix B4 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.

Property (i) holds by construction. According to property (ii), fully cursed players set interior cutpoints in all treatments. Given the parameters of our model, they make identical choices in **T1** and **T3**. Property (iii) is arguably the most robust and interesting one. It states that the amount of unraveling in all treatments is inversely related to the level of cursedness of individuals. Finally, property (iv) says that the sequence of moves does not affect the equilibrium, although this is a consequence of the way α-CE is defined.

6.3.2 α-cursed equilibrium: estimation

Given the potential of α-CE to represent our participants' choices, we conduct a structural estimation of a Behavioral Random Utility Model (BRUM). More precisely, we assume that participants are α-cursed (they underestimate the relationship between the rivals' action and their information) and commit mistakes in choices (the probability of choosing an action over another is monotonically increasing in the utility difference rather than a step function). Formally, we estimate a two-parameter (α, λ) model, where α captures the mean cutpoint of the population and λ reflects the responsiveness to the utility differential. Here, $\lambda = 0$ corresponds to random choice and $\lambda \to \infty$ corresponds to best response.

Table 6 reports the estimated parameters α and λ of BRUM by treatment and grade group for individuals who use threshold strategies C_+ (the derivation of the structural model can be found in Appendix B5). 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. Figure 7 reports the empirical c.d.f. ($S_c(x)$, as in Figure 5) and the estimated best fit of the BRUM.

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

There are several interesting lessons to be learned from these estimations. The theoretical properties (i) to (iv) together with the fact that the mean empirical threshold is above m in all treatments, implies that 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

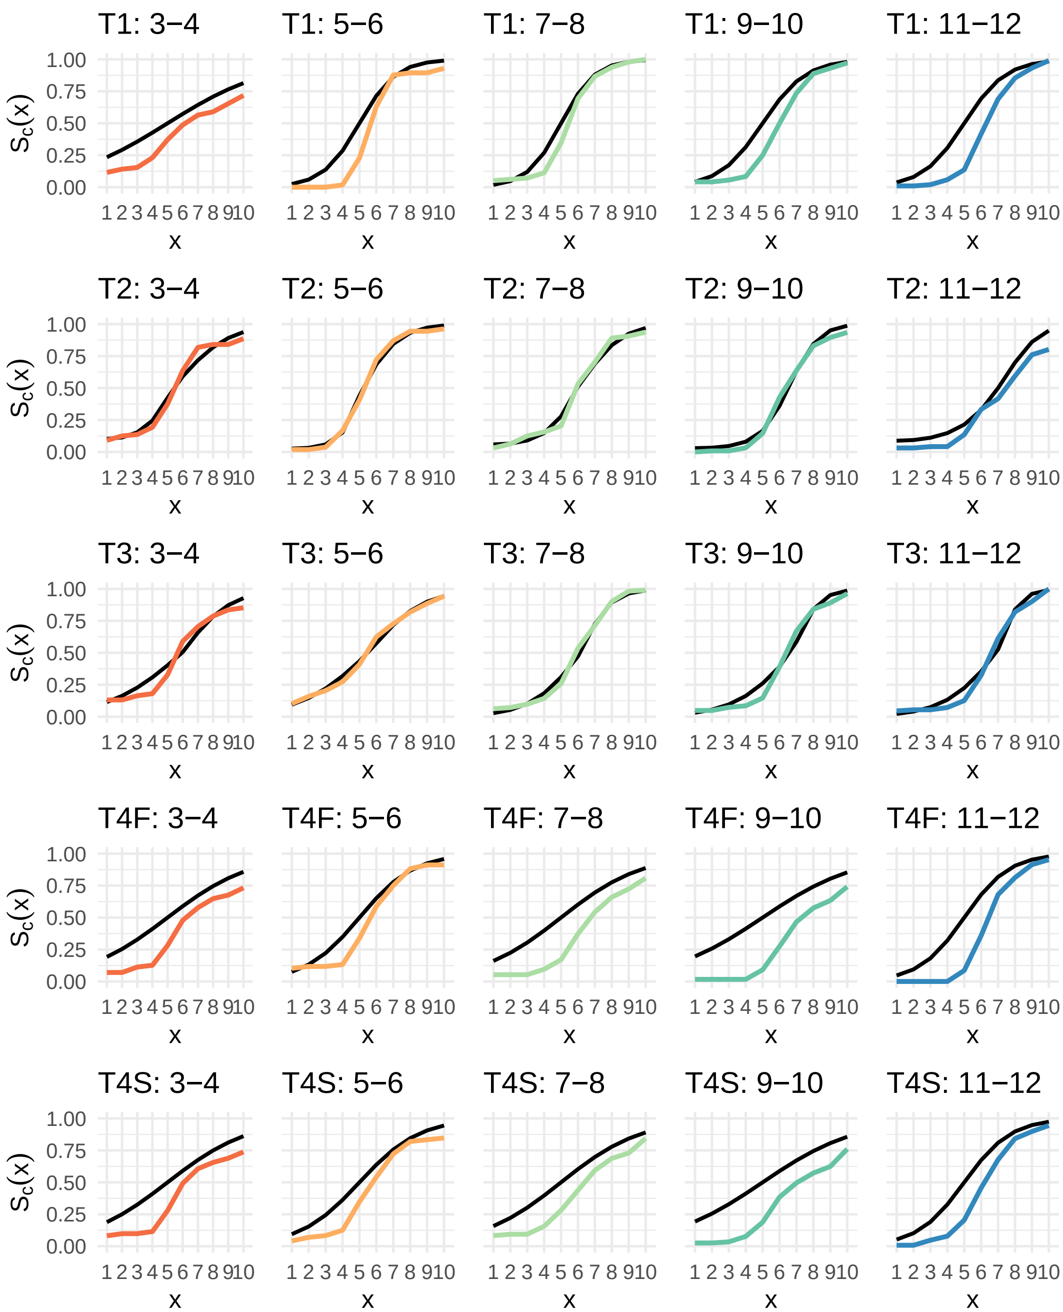

Figure 7: Best fit (α, λ) of BRUM for strategic consistent participants (C_+)

in the sequential treatment, where not only 'pass' is very frequent but the behavior is also 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 (especially in **T3**).

More generally, this behavioral theory exhibits both merits and limitations when tested against our dataset. On the one hand, the observed partial unraveling aligns with the behavioral predictions of cursed players, who underestimate (or even fully ignore) the relationship between the rival's action and information. On the other hand, the large differences in α-estimates for treatments where unraveling goes in opposite directions suggests that our participants struggle to incorporate treatment-specific incentives, a deviation that α-CE fails to capture.

7 Theory of Mind

7.1 Cognitive theory of mind v. affective theory of mind

"Theory of mind" encompasses the overarching idea of mentalizing about others. Cognitive theory of mind (c-ToM) centers on understanding the cognitive processes driving human behavior whereas affective theory of mind (a-ToM) involves empathizing with emotions. This latter component is closely tied to the concept of emotional intelligence, which is increasingly recognized as a pivotal asset within the corporate realm.[16]

Traditionally, the assessment of c-ToM has 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 actual reality.[17] Typically developing children as young as four or five years old grasp the concept of false belief. When more intricate tasks involving higher-order beliefs are introduced, older children and even adults tend to exhibit a bias towards their own information. However, to introduce challenges for 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; Meinhardt

[16] Emotional intelligence refers to the ability to perceive and manage emotions effectively and is often studied in professional contexts. a-ToM refers more specifically to the ability to attribute emotions. It involves recognizing that individuals have their own emotions and perspectives, and can be leveraged to predict the behavior of others.

[17] 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).

et al., 2011; Valle et al., 2015). Consequently, these tasks necessitate participants to store and manipulate a substantial amount of information, placing a strain on working memory, a purely cognitive resource that supports but does not equate to reasoning ability.

We argue that a game of incomplete information is a better tool to measure c-ToM. On the one hand, it demands the same cognitive skills as false belief tasks. In both scenarios, individuals must establish connections between the actions of others and the information they possess. On the other hand, 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 and computation. Finally, our specific game contains several dimensions of cognitive ability that can be empirically measured, such as sensitivity of s to private value, likelihood of choosing a threshold 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)), which assesses an individual's capacity to comprehend and interpret emotional and mental states. In this task, 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. It also proves particularly pertinent in the investigation of conditions such as autism spectrum disorders, where theory of mind deficits may be prevalent. 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 poorer performance in both tasks. Second, there is preliminary evidence suggesting a positive association between proficiency in a-ToM tasks and general 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). It is noteworthy that c-ToM tasks engage a subnetwork connected to cognitive functions, whereas a-ToM tasks involve a subnetwork linked to emotional processing (Atique et al., 2011; Healey and Grossman, 2018). 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 complex strategic situations (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.

7.2 Testing affective theory of mind

The original RMET was designed to study a-ToM in adults. The emotions sometimes involve complex nuances (e.g. "suspicious", "thoughtful") and are expressed in a complicated language (e.g. "flustered", "skeptical"). We therefore decided to adapt and validate a novel, child-adapted 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).[18] For each image, we provided four possible answers, one of which best captured the depicted emotion. Participants were incentivized, earning 0.10€ for every correct response. In Figure 8, we present one example as seen by our participants (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.

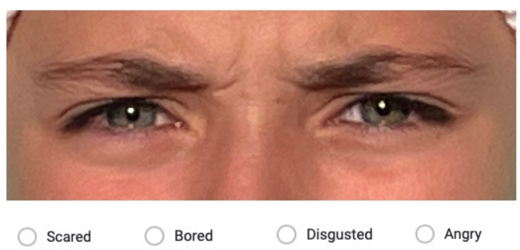

Figure 8: An example of an image and the four possible answers in the DeRMET

Compared to the original task, there are three advantages in using our version with our population. By sourcing all images from a single individual, we make it more consistent. By using a child as a model to depict emotions, we make the images more relatable to children and teenagers. By focusing on simpler emotions and using a straightforward language, we make the same task suitable for any age, from very young participants to

[18] 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 difficult to identify. The task and validation materials were registered with the Open Science Framework to ensure transparency and accessibility.

adults. Figure 9 summarizes the distribution of correct answers in DeRMET by grade group, with the group average at the top.

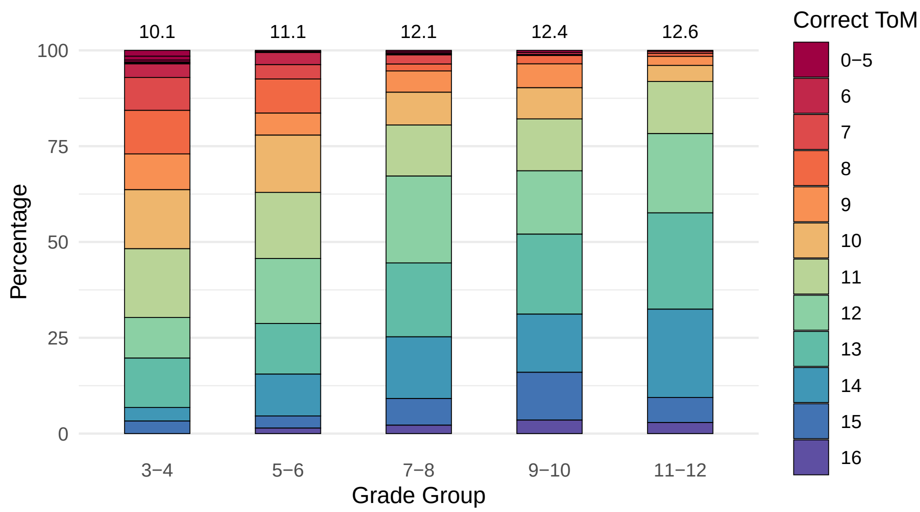

Figure 9: Distribution of correct answers in the DeRMET (a-ToM) task by grade group

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.[19] This means that the task is well-calibrated: neither trivial nor impossibly challenging for any grade group. 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" consistently aligns with the choice of the majority of participants.[20] In the next section, we use the performance in DeRMET as an input to explain behavior in the game.

7.3 Affective Theory of Mind and optimal decision-making

To examine the possibility of a link between a-ToM and strategic thinking, we perform the same Probit regressions as in Tables 2 and 4 (likelihood to choose s and likelihood to choose

[19] By the properties of the binomial distribution with $p = 0.25$ and $n = 16$, 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.

[20] 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.

a threshold strategy C_+), but we now include the performance in the a-Tom task as an additional independent variable. We also conduct OLS regressions to study the expected loss of an individual (relative to playing best response to the empirical behavior of others in their grade group and treatment) as a function of the previous variables, including performance in a-ToM. These regressions are performed both in the entire sample and in the subset of individuals who employ threshold strategies. For conciseness and given that behavior in the previous regressions is similar across treatments, we pool together all the information, with treatment dummy variables. The results are compiled in Table 7. 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*).

	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 7: Regressions to evaluate the effect of the a-ToM task

The results unambiguously indicate a strong, positive and robust link between a-ToM and different measures of strategic thinking in our game. Individuals with high a-ToM performance are more likely to pass for small values of x and they are also more reactive to increases in x compared to participants with low a-ToM performance (columns (1) and (5)). Participants displaying high a-ToM performance are also more likely to adopt threshold strategies (columns (2) and (6)). Both effects are highly significant and independent of the (still significant) effect of age. Performance in a-ToM is also associated with lower losses relative to best response behavior within grade groups and treatments (columns (3) and (7)). This is expected since they employ threshold strategies more often. Finally, even among individuals who use threshold strategies, payoffs are higher for those with high a-ToM performance than for those with low a-ToM performance (columns (4) and (8)).[21] The results are very similar whether we consider absolute performance in the a-ToM task (*abs-ToM*) or performance relative to other participants in their own grade (*norm-ToM*), which is intuitive since the regressions also account for age.[22]

Table 7 also reinforces the gender differences highlighted earlier: males are significantly more likely to adopt threshold strategies and, consequently, achieve higher payoffs than females within their grade group. This goes against our previous findings in a backward induction game, where females outperformed males (Brocas and Carrillo, 2021). However, it is worth noting that participants in that study came from a different school with distinct demographics and cultural influences, which may explain 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 threshold strategy, obtaining a high payoff) and affective theory of mind.*

8 Conclusion

We have reported the results of the first controlled lab-in-the-field experiment designed to study the strategic behavior of privately informed children and adolescents and the relationship between cognitive and affective theory of mind. To that end, we have recruited a

[21] High and low a-ToM players who use threshold strategies set similar cutpoints on average (5.86 v. 5.84) but the former are significantly 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 6.1).

[22] Notice the positive effect of age in the OLS regressions of losses. However, one must realize that these are not absolute losses, but losses relative to the highest achievable *in that grade group*.

large sample of 1662 participants ranging in age from 8 to 18 years old to play a computerized game of asymmetric information. Our study led to several critical observations.

Participants of all ages understand that private information is meaningful for the selection of actions and they select different actions according to the strength of their information, a property we called 'choice monotonicity'. However, older participants are more likely to select strategies that match qualitative features of the optimal strategy, a property we termed 'strategic consistency'. Contrary to choice monotonicity, which is present at an early age, strategic consistency develops over time. Yet, it is not always mastered by the oldest participants. The findings are consistent with documented developmental trajectories in games of complete information, showing that very young children grasp fundamental strategic concepts (Brocas and Carrillo, 2018) but learn to perform closer to predictions as they develop, although they may not reach a level of mastery (Brocas and Carrillo, 2021, 2022b).

The presence of private information introduces a series of challenges that are not overcome as we develop. Players fail to respond effectively to the strategies of others and reach the equilibrium. The clearest evidence is that participants do not adapt their play to specific changes in incentives due to changes in the structure of the game. They are oblivious to what we termed 'environmental variability'. This result is particularly intriguing among individuals who have reached or are close to their peak cognitive capacity. Overall, participants do not accurately anticipate how others utilize their private information to shape their own strategies. While partial (rather than full) unraveling in the choice of threshold can be partly explained by "cursed" reasoning, the lack of adjustment to the treatment condition cannot.

These collective observations indicate 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 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 threshold 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 some extent also in step (iii).

The positive correlation between a-ToM and performance 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 and representing 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 delve deeper into the causal aspects of these mechanisms.

References

Bijoy Atique, Michael Erb, Alireza Gharabaghi, Wolfgang Grodd, and Silke Anders. Task-specific activity and connectivity within the mentalizing network during emotion and intention mentalizing. *Neuroimage*, 55(4):1899–1911, 2011.

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.

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.

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. Iterative dominance in young children: Experimental evidence in simple two-person games. *Journal of Economic Behavior & Organization*, 2018.

Isabelle Brocas and Juan D Carrillo. Introduction to special issue "understanding cognition and decision-making by children." studying decision-making in children: challenges and opportunities. *Journal of Economic Behavior & Organization*, 179:777–783, 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 and Juan D Carrillo. Developmental meritocracy. USC Working Paper, 2022a.

Isabelle Brocas and Juan D Carrillo. Adverse selection and contingent reasoning in preadolescents and teenagers. *Games and Economic Behavior*, 133:331–351, 2022b.

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.

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.

Meghan L Healey and Murray Grossman. Cognitive and affective perspective-taking: evidence for shared and dissociable anatomical substrates. *Frontiers in neurology*, 9: 491, 2018.

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*, 66:799–823, 2015.

John A List, Ragan Petrie, and Anya Samek. How experiments with children inform economics. Technical report, National Bureau of Economic Research, 2021.

Alejandro Martínez-Marquina, Muriel Niederle, and Emanuel Vespa. Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10):3437–74, 2019.

Jörg Meinhardt, Beate Sodian, Claudia Thoermer, Katrin Döhnel, and Monika Sommer. True-and false-belief reasoning in children and adults: An event-related potential study of theory of mind. *Developmental Cognitive Neuroscience*, 1(1):67–76, 2011.

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.

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.

Reinhard Selten, Klaus Abbink, and Ricarda Cox. Learning direction theory and the winner's curse. *Experimental Economics*, 8(1):5–20, 2005.

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.

Doron Sonsino, Ido Erev, and Sharon Gilat. On rationality, learning and zero-sum betting–an experimental study of the no-betting conjecture. Working Paper, 2002.

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

Richard 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. Supplementary material: 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 10 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.

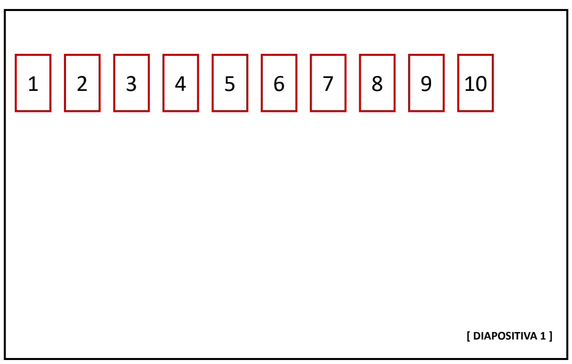

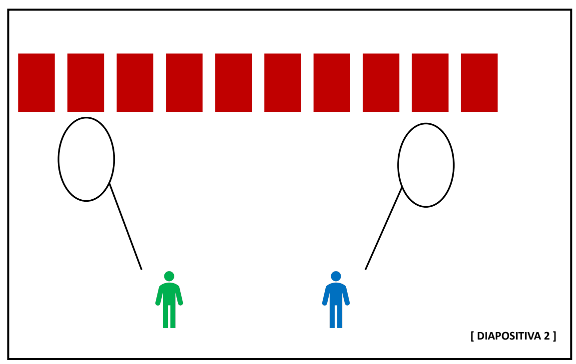

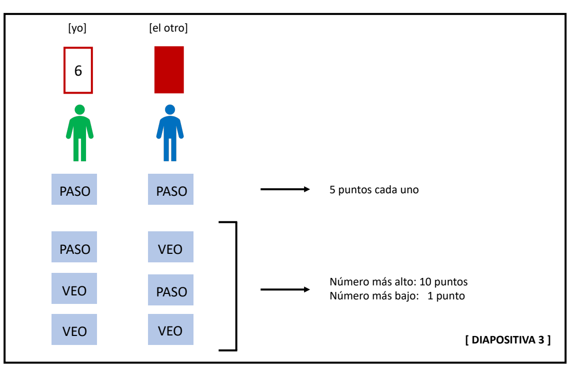

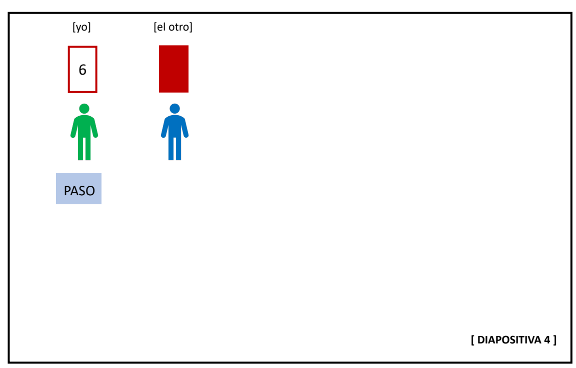

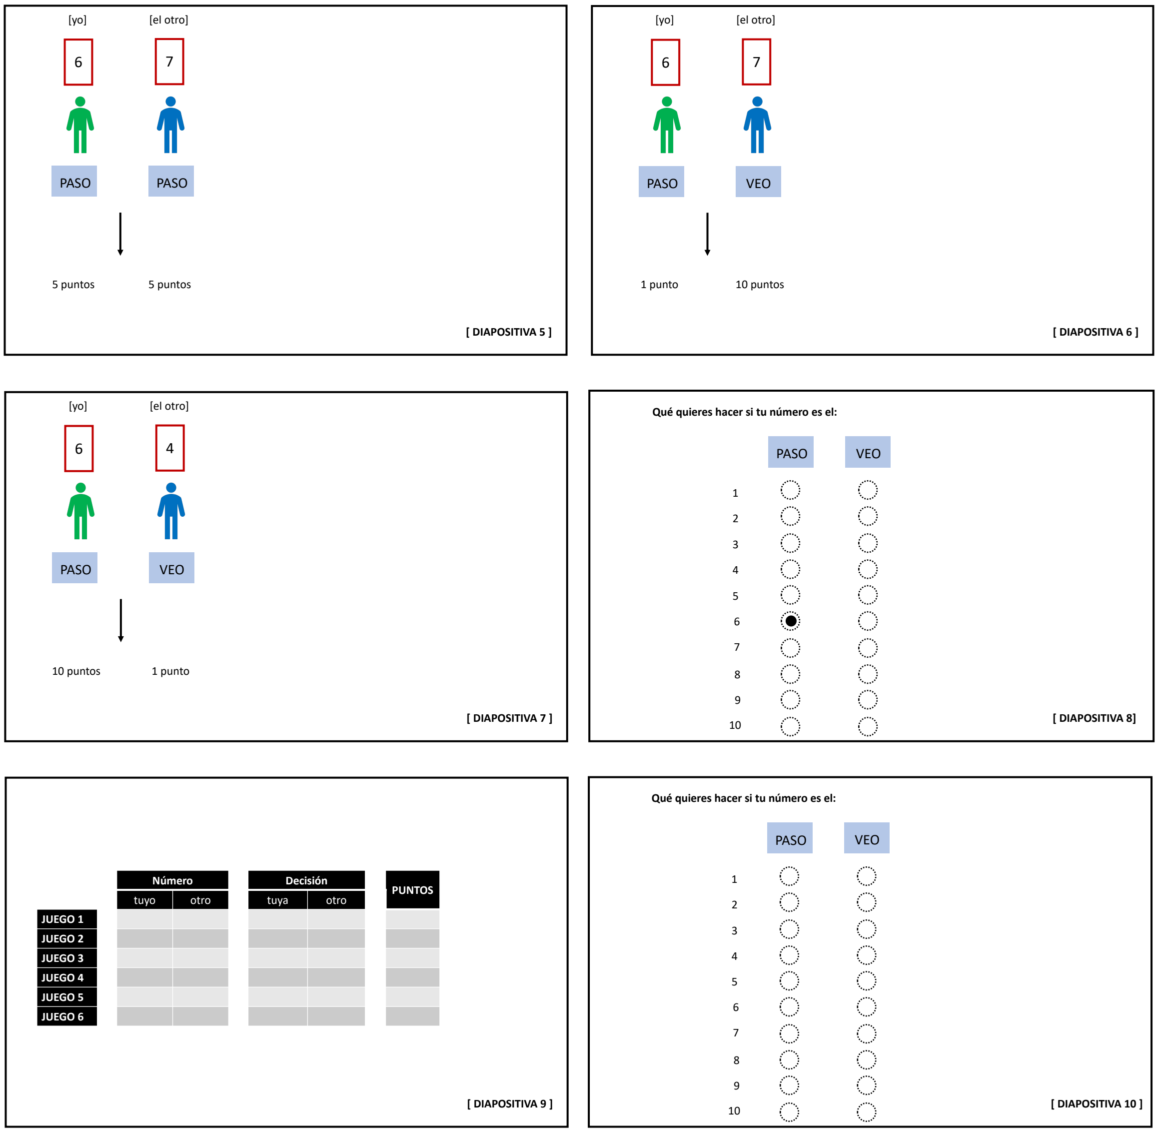

Figure 10: Slides (in Spanish) projected on screen for instructions of the game

A2. Developmental Read the Mind in the Eyes Test (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. 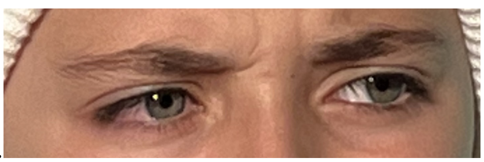
5. 
6. 
7. 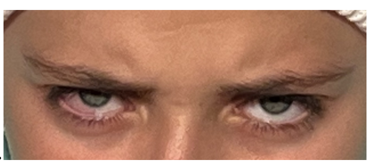
8. 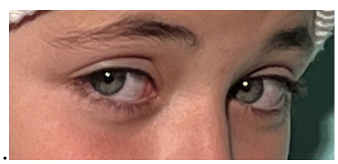
9. 
10. 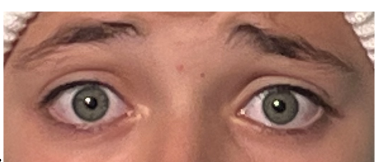
11. 
12. 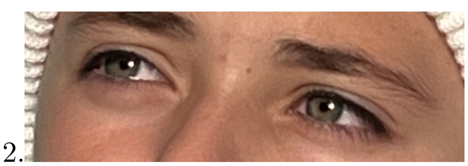


13.

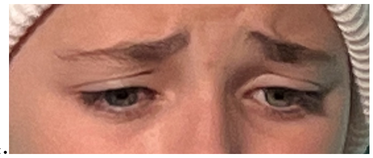
14.


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

Appendix B. Supplementary material: additional analyses

B1. 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 11, 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).

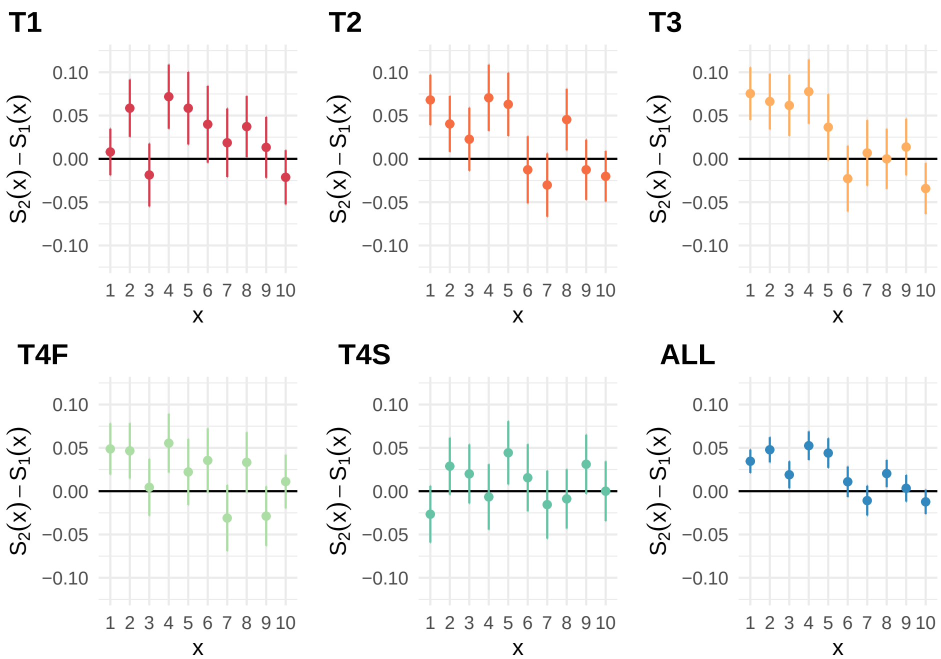

Figure 11: $S_2(x) - S_1(x)$ by treatment

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 11). 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 11). Conversely, if we expected that participants would learn how to best respond to their peers (the 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 11, 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.

B2. Behavior of individuals who use and do not use threshold 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 threshold strategies ($\mathrm{C}_+ = \mathrm{C}_0 + \mathrm{aS} + \mathrm{aP}$) and individuals who do not use perfect threshold strategies ($\mathrm{notC}_+ = \mathrm{C}_1 + \mathrm{ALT} + \mathrm{Other}$). Using the predictions obtained from these Probit models, we perform the same exercise as in Figure 3 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 12 for all players together (All, left), players who use threshold strategies (C_+, center), and players who do not use threshold strategies (notC_+, right).

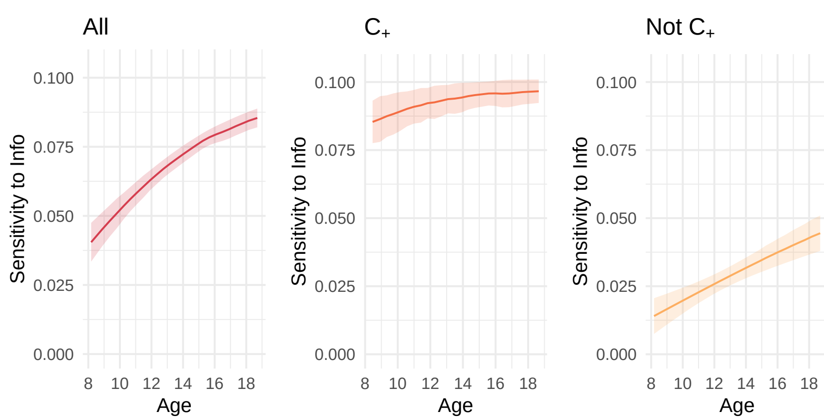

Figure 12: Sensitivity of action to private information by type of strategy

As already indicated in section 4, Figure 12 (left) reaffirms that older participants react more to private value than their younger peers. According to Figure 12 (center), the subgroup of individuals who use threshold 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 threshold 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 12 (right). It's noteworthy that our youngest players who employ threshold strategies display a sensitivity to information comparable to our oldest "average" player. Conversely, our oldest participants who do not employ threshold strategies exhibit a sensitivity to information similar to our youngest "average" player.

In summary, participants who employ threshold strategies exhibit consistent reactions to information regardless of their age. The primary differences across grade groups are twofold: (i) the proportion of individuals who adhere to strategic consistency and (ii) the behavior of those who do not adhere to it.

We also conduct a basic analysis of changes in strategies across decisions. In Table 8, we present a 2×2 matrix. Each cell represents the percentage of observations that correspond to a threshold (C_+) v. no threshold (notC_+) strategy in the first (rows) and second (columns) decision.

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

Table 8: Percentage of strategies in first (rows) and second (columns) decision

Once again, Table 8 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.

B3. 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 13 reports that information by grade group (left) and private value (right).

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 13(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 13(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.

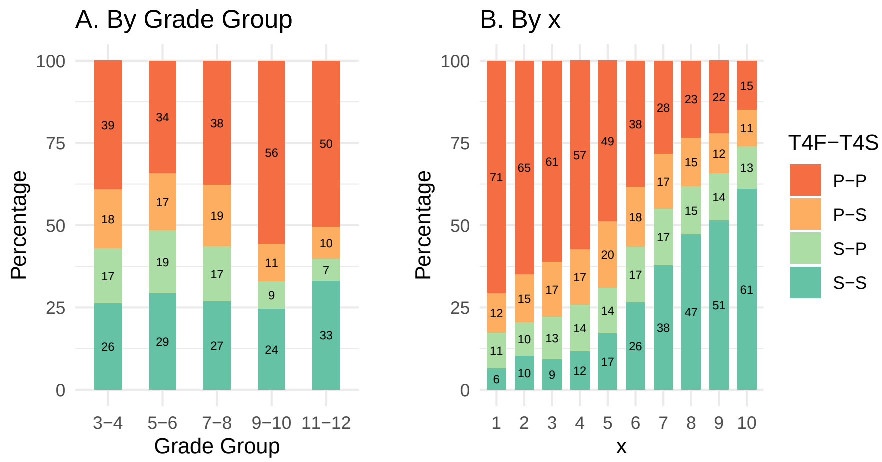

Figure 13: Change in individual choice between **T4F** and **T4S**

B4. 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:

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

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:

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

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

B5. 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:

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

(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.