The evolution of choice and learning in the two-person beauty contest game from kindergarten to adulthood *

Isabelle Brocas
University of Southern California and CEPR

Juan D. Carrillo
University of Southern California and CEPR

December 2019

Abstract

We develop a graphical, non-analytical version of the two-person beauty contest game to study the developmental trajectory of *instinctive behavior* and *learning* from kindergarten to adulthood. These are captured by observing behavior when the game is played in two consecutive trials. We find that equilibrium behavior in the first trial increases significantly between 5 and 10 years of age (from 17.9% to 61.4%) and stabilizes afterwards. Children of all ages learn to play the equilibrium, especially when they observe an equilibrium choice by the rival. Our younger children are the weakest learners mainly because they are less frequently paired with rivals who play at equilibrium. Finally, the choice process data suggests that participants who play at equilibrium in the second trial are also performing fewer steps before reaching a decision, indicating that they are less hesitant about their strategy.

Keywords: developmental game theory, laboratory experiment, steps of dominance, two-person beauty contest.

JEL Classification: C72, C92.

*We thank the members of the Los Angeles Behavioral Economics Laboratory (LABEL). We are grateful to the staff of the Lycée International de Los Angeles (LILA) –in particular Emmanuelle Acker, Nordine Bouriche, Mathieu Mondange and Anneli Harvey– for their help and support running the experiment in their school. The study was conducted with the University of Southern California IRB approval UP-12-00528. We acknowledge the financial support of the National Science Foundation grant SES-1851915.
Address for correspondence: Isabelle Brocas, Department of Economics, University of Southern California, 3620 S. Vermont Ave., Los Angeles, CA 90089, USA, <brocas@usc.edu>.

1 Introduction

Strategic thinking is a cornerstone of multi-person decision-making. Researchers in experimental economics have proposed a variety of games to investigate whether subjects are able to anticipate the actions of others and best respond to them. A prominent example in this literature is the "guessing game" or "beauty contest" (Nagel, 1995; Stahl, 1996; Ho et al., 1998), where n players (from three to thousands) simultaneously announce a number and the person closest to p (< 1) times the average wins. An attractive feature of this game is that differences in depths of reasoning translate into predictable differences in choices.[1] Camerer et al. (2004) provide a behavioral theory that serves as a natural template. However, one challenge of this game is the difficulty to empirically distinguish between subjects with limited ability to reason strategically and subjects who believe that others have such limited ability.

To address this challenge, Costa-Gomes and Crawford (2006), Grosskopf and Nagel (2008) and Chou et al. (2009) study a two-player version of the game. In this variant, the lowest number weakly dominates any other, making the optimal, payoff-maximizing behavior reasonably easy to determine. Perhaps surprisingly, these studies still highlight significant departures from the theoretical predictions. At the same time, and as one would expect, suboptimal choices are less prevalent when subjects are experienced and when instructions are simplified.

In the current study, we adopt a developmental perspective to identify the reasons why strategic play is absent in dominance solvable games. We argue that observing the trajectory through which we acquire (or perhaps, do not acquire) such ability should help us understand and predict behavioral shortcomings. To this purpose, we ask school-age participants from 5 to 18 years old to play a version of the two-person beauty contest. We perform the exercise twice in order to unveil the capacity of participants to learn from their mistakes. Two types of feedback are available in the game: they may learn from suffering a low payoff or from observing a more successful strategy by the rival. Also, a major challenge in developmental studies is to design a protocol that is accessible and clear for children as young as 5 years old and, at the same time, non-trivial for adults. We address this issue by proposing a novel, graphical, non-analytical interface.

The experimental economics literature has recently witnessed an increased interest in

[1]See Coricelli and Nagel (2009) for a study of the neural correlates of depths of reasoning.

developmental aspects of economic behavior in children and teens. However, as noted in the surveys by Sutter et al. (2019) and List et al. (2018), the majority of the literature focuses on individual decision making paradigms. Notable exceptions include Sher et al. (2014), Chen et al. (2016), Czermak et al. (2016), Brocas and Carrillo (2018a,b) and Fe and Gill (2018). However, these studies focus on games that are quite challenging to play and/or they restrict to a limited age range. By proposing a significantly simpler design and observing behavior in a wide age range, we hope to unveil any cognitive milestone in the acquisition of the ability to select dominant strategies. We are also interested in determining which kind of feedback (if any) facilitates learning.

Our study delivers three main insights. First, we show that the game is neither impossible to solve for our youngest participants nor trivial for the adult population. Equilibrium behavior increases significantly between Kindergartners and 5^{th} graders and stabilizes afterwards. It means that, in this simple game, 10 year-old children are as good decision makers as college undergraduates. The result mirrors our findings in other, more complex dominance-solvable games (Brocas and Carrillo, 2018b) and suggests that the contribution of age to equilibrium play vanishes early in life. This pattern is also reminiscent of the observed development of rational choice in individual decision-making contexts. Indeed, Harbaugh et al. (2001) and Brocas et al. (2019) have shown that 10 years old children are significantly more rational than their younger peers and they are as rational as adults. Second, equilibrium behavior increases significantly between the first and the second trial. The absolute improvement is similar across all age groups. However, among participants who miss the first trial, the proportion who learns to play the equilibrium in the second trial increases significantly between K and 5^{th} grade and remains constant afterwards. We show that observing a rival playing at equilibrium is the most powerful learning channel in our game. Since younger children face less sophisticated peers, they are not exposed as often to these learning opportunities, thereby improving their behavior less frequently. Finally, our data reveals an intriguing link between choices and processes. In the second trial, equilibrium players are significantly more likely to play fast (in our case, make their choice in one step) than slow (go through intermediate steps before submitting their final choice) whereas non-equilibrium players are equally likely to play fast or slow. This suggests a negative relationship between learning and hesitation.

The article is organized as follows. We describe the experiment in section 2. Section 3 reports the results organized in different ways. Section 3.1 focuses on developmental changes and reports results by age group. Section 3.2 addresses learning and reports findings by trial. Section 3.3 focuses on decision processes associated with choices and section 3.4 presents a regression analysis that investigates the robustness of our findings. Final remarks are relegated to section 4.

2 Design and procedures

The paper studies a twice-repeated, two-person beauty contest game with children and adolescents. This population presents significant methodological challenges, and we follow the guidelines proposed by Brocas and Carrillo (2019) to address them.[2]

Participants. We recruited 304 participants from grades K to 10^{th} at the Lycée International de Los Angeles (LILA), a French-English bilingual private school in Los Angeles. We ran 28 sessions, each with 8 to 12 participants, in a classroom at the school using touch-screen PC tablets. The task was programmed with the open source software Multistage.[3] For each session, we mixed male and female participants from the same grade although, for logistic reasons, we sometimes included participants of two consecutive grades.[4] For comparison, we also recruited 34 USC students (adults - A) and we ran 3 sessions at the Los Angeles Behavioral Economics Laboratory (LABEL) in the department of Economics at the University of Southern California, using *identical* procedures. Table 1 reports the distribution of participants by grade.

Population	LILA										USC
Grade	**K**	**1**$^{\text{st}}$	**2**$^{\text{nd}}$	**3**$^{\text{rd}}$	**4**$^{\text{th}}$	**5**$^{\text{th}}$	**6**$^{\text{th}}$	**7**$^{\text{th}}$	**8**$^{\text{th}}$	**10**$^{\text{th}}$	**A**
Participants	42	32	24	27	30	27	20	40	34	28	34

Table 1: Number of participants by grade

Notice that with some exceptions (e.g., Cobo-Reyes et al. (2019)) studies with children usually do not recruit an adult population. We believe it is important to include an adult control group to establish a behavioral benchmark, even if the comparison should be taken

[2]In short, the principles are: (i) simplify the procedures given the participants' limited attention; (ii) offer age-appropriate incentives; (iii) present the task in a simple and attractive way; (iv) describe carefully the children population(s), and (v) include a benchmark adult comparison group whenever possible.

[3]Documentation and instructions for downloading the software can be found at the website `http://multistage.ssel.caltech.edu`.

[4]Students from 9^{th}, 11^{th} and 12^{th} grade did not participate in the study because they were taking or preparing for national french or US exams during this period.

with a grain of salt (Brocas and Carrillo, 2019). In our case, the majority of students at LILA are from caucasian families of upper-middle socioeconomic status. After graduating, they typically attend well-ranked colleges in Europe and North America (including USC). Despite some differences (nationality, family background, size of peer group, etc.), the two populations match 'reasonably' well. Notice also that our population is homogenous, making it possible to perform meaningful age comparisons. As shown in Brocas and Carrillo (2018b), behavior in children is highly dependent on a variety of economic and demographic characteristics (see also Charness et al. (2019)). Pooling participants from different schools is likely to introduce confounds that hide any developmental trajectory. We avoid these confounds by recruiting children from the same school, who follow the same school curriculum and come from similar social and economic backgrounds. We also collect information about gender and number of siblings to capture potential remaining sources of individual heterogeneity.

Tasks. The experiment consisted of two tasks always performed in the same order. We started with the "Army" game studied in this paper. After a short break, we ran a mixed strategy game called "Farmers and Pirates". We report and analyze the results of the second game in a different article (Brocas and Carrillo, 2020).

Army Game. Participants completed *two* consecutive trials of the Army game. The game is a graphical variant of the 2-person beauty contest that has been extensively studied in the literature (Costa-Gomes and Crawford, 2006; Grosskopf and Nagel, 2008). In our version, two players (1 and 2) simultaneously select one cell from a two-dimensional 8×8 rectangular grid. A player obtains points by choosing a cell to the left of the cell chosen by the other player. If both players choose cells in the same column, then the points go to player 1. A player also obtains points by choosing a cell below the cell chosen by the other player. If both players choose cells in the same row, then the points go to player 2.

As mentioned earlier, it is of paramount importance to provide a simple, graphical interface and a story accessible and appealing to children as young as 5 years-old. We therefore developed a battle role-playing story. Participants are matched in groups of two and assigned a role, either as "tree leader" or "volcano leader." Each leader starts with a possession (trees or volcanoes) and an army (green dot or red dot), as depicted in the screenshots of Figure 1. Armies are worth 20 points and possessions are worth 6 points.

The objective of a leader is to protect his possession and seize his rival's. This is

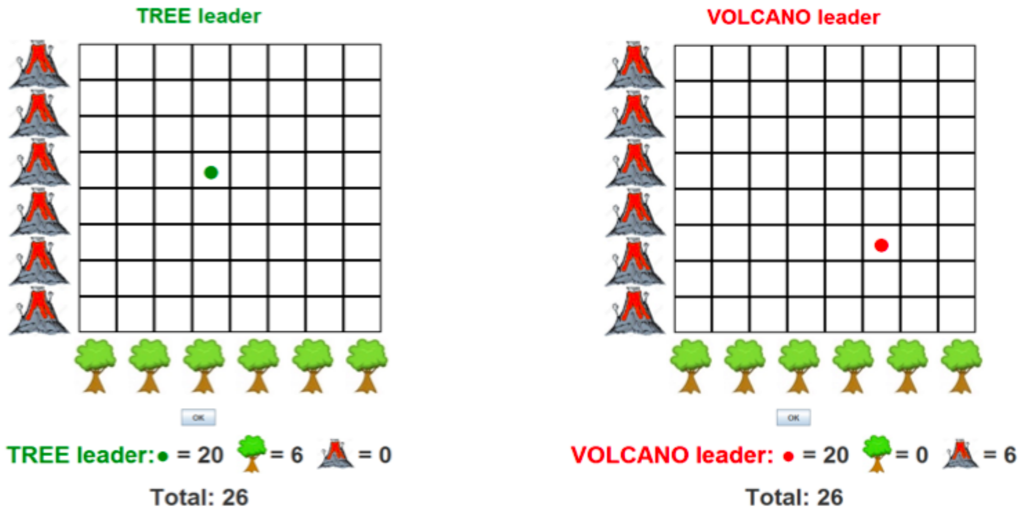

Figure 1: Army game. Computer screenshots of tree leader (left) and volcano leader (right) at the beginning of a trial.

achieved by strategically locating the army. More precisely, the leader who locates his army closest to the trees (lower row) seizes the trees, and the leader who locates his army closest to the volcanoes (leftward column) seizes the volcanoes. If both armies are positioned at the same distance from the trees, the tree leader keeps the trees. If both armies are positioned at the same distance of the volcanoes, the volcano leader keeps the volcanoes.[5]

To implement their choices, participants simultaneously drag their army with their finger over the touchscreen PC tablet, from the initial position to one of the 64 cells in the grid, without observing their rival's choice, so that armies can be placed in the same cell. Participants can "pause" at any cell and move their army again. A choice is final only after the subject presses the OK button. Once participants have selected their final locations, we reveal them to each other and display the payoffs. We instruct the participants to observe and interpret the results and randomly form new groups. Participants then play a second trial of the same game with their new partner. A copy of the instructions is included in Appendix A.

The battle metaphor for the two-person beauty contest game was first introduced by

[5]Allocating points for the army (which can never be lost) implies a positive reward for every participant and a low payoff-variance. This ensures a pleasant experience for all children.

Chou et al. (2009) to study the effect of instructions on behavior. A two-dimensional grid implementation of the game was first proposed by Chen et al. (2018) to track the eye movements of subjects prior to their final choice and, through that method, infer their reasoning process. Our protocol combines and extends both procedures. We remove any numerical, verbal or analytical literacy requirement to play the game, thereby making it understandable and attractive to very young children. At the same time, we hope the game is non-trivial for educated adults. We believe that our presentation facilitates understanding relative to previous designs. However, the main goal is not to analyze the role of instructions but to make it accessible to all ages, thereby allowing a comparison of behavior over a long age span.

Inspired by Chen et al. (2018), we recorded not only the final location, but also the trajectory employed by participants before making a choice, that is, the intermediate resting locations (if any). The Army game lasted 15 to 20 minutes, including instructions. The entire experiment ("Army" followed by "Farmers and Pirates") lasted 60 to 75 minutes.

Payoffs. Subjects accumulated points. Following Brocas and Carrillo (2019), we implemented different payments depending on the population. LILA participants in grades 6^{th} to 10^{th} and USC students had points converted into money paid at the end of the experiment with an amazon e-giftcard. We set a 33% higher conversion rate at USC (and a $5 show up fee) to account for differences in marginal value of money and opportunity cost of time. Payments for both games ("Army" and "Farmers and Pirates") averaged $11.3 at LILA and $21.2 at USC. For LILA participants in grades K to 5^{th}, we set up a shop with 20 to 25 pre-screened, age-appropriate toys and stationery (bracelets, erasers, figurines, die-cast cars, trading cards, apps, earbuds, gel pens, etc.). Before the experiment, children were taken to the shop and showed the toys they were playing for. They were told the point prices of each toy. For the youngest subjects, we explicitly clarified that more points would result in more toys. At the end of the experiment, subjects learned their point earnings and were accompanied to the shop to exchange points for toys. We made sure that every child earned enough points to obtain at least three toys.[6]

Taxonomy of strategies. In a Nash equilibrium, the tree leader positions his army in one

[6] The procedure emphasizes the importance of accumulating points while making the experience enjoyable. At this age, a toy is also a significantly more attractive reward than money. We spent on average $5 per child in toys. Most children are familiar with this method of accumulating points that are subsequently exchanged for rewards, since it is commonly employed in arcade rooms and fairs.

of the bottom row cells and the volcano leader positions his army in one of the leftmost column cells. Of these cells, all but the bottom left corner cell are weakly dominated. From now on and with a slight abuse of language, we call *Equilibrium* (**E**) the strategy that consists in choosing the bottom left cell. It is the unique, non-dominated Nash equilibrium. We call *Defensive* (**D**) the strategy of a participant who focuses on protecting his possessions, namely, the remaining bottom row cells for the tree leader and leftmost column cells for the volcano leader. It corresponds to the weakly dominated Nash equilibria of the game. We call *Offensive* (**O**), the strategy of a participant who focuses on trying to seize the possessions of the rival, namely, the remaining leftmost column cells for the tree leader and the remaining bottom row cells for the volcano leader. This strategy is not a Nash equilibrium and does not guarantee a possession. Finally, we call *Irrational* (**I**), any other strategy. Notice that a participant who chooses **E** demonstrates an understanding of both dimensions of the problem, a participant who chooses **D** or **O** reveals that he has focused and understood one dimension but not the other, and a participant who chooses **I** has understood no dimension of the problem. Also, the large grid (8×8) ensures that strategies **E**, **D** and **O** are unlikely to be played by pure chance.

3 Data analysis

The software did not record the decisions of 10 participants: 4 from K and 2 from 3$^{\text{rd}}$ grade in the first trial and 3 from K and 1 from 3$^{\text{rd}}$ grade in the second trial. We report the results obtained for the remaining 328 participants. To increase the statistical power, we group the school kids in 5 age groups: K-1, 2-3, 4-5, 6-7, 8-10. The USC adult population (A) constitutes our control group. For the statistical analysis, we perform two-sided tests unless otherwise noted.

3.1 Aggregate behavior

Following the definitions in the previous section, Figure 2 depicts the proportion of observations that fall in each of the four strategies: Equilibrium (**E**), Defensive (**D**), Offensive (**O**) and Irrational (**I**). We separate the analysis by trial and age group.

Despite the significant developmental differences, the optimal strategy is neither impossible to find nor trivial to determine for any given age group. Indeed, averaging over both trials, the proportion of equilibrium play is 0.26 for K-1 and 0.74 for adults. Of the

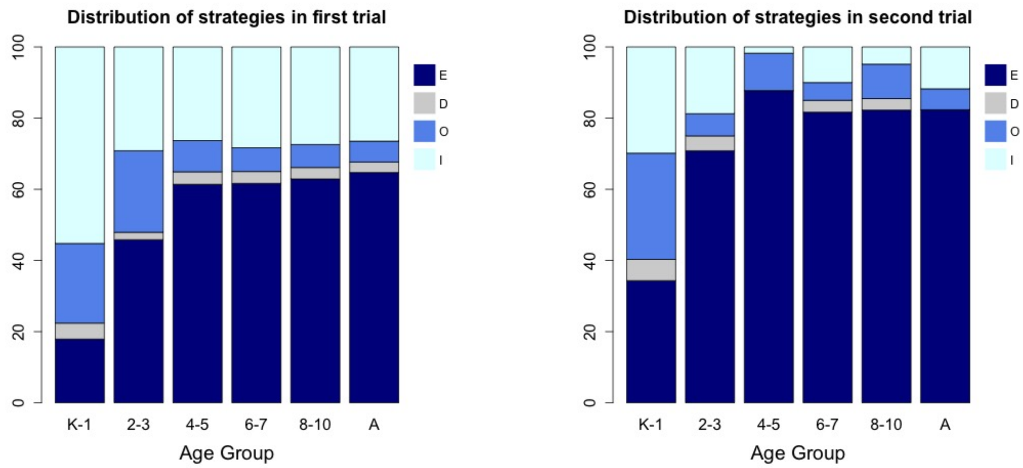

Figure 2: Distributions of strategies in first (left) and second (right) trial.

four strategies, only the weakly dominated Nash equilibrium **D** is rarely used. Among the participants who do not play the equilibrium, missing both dimensions of the problem (**I**) is more frequent than missing one (**O** + **D**).

As expected, choices depend critically on age and trial. We observe a very significant increase in **E** between the first and second trial in all age groups (between 16% and 26%), mostly at the expense of **I** (McNemar test of differences, p-value between 0.001 and 0.015). The improvement is similar (18%) but only marginally significant for our control population ($p = 0.077$), in part because of the smaller number of observations. Although behavior is different across trials, the changes with age are similar in both trials: an increase in equilibrium choice from K-1 to 4-5 which stabilizes afterwards. The difference in equilibrium behavior between K-1 and the other groups is statistically significant in each trial ($p < 0.01$). By contrast, some of the differences between 2-3 and the older groups are not statistically significant again due to the small number of observations within each trial (one-sided tests, p-values between 0.05 and 0.08 for the first trial and p-values between 0.03 and 0.17 for the second).[7]

[7] If we group all participants from 4$^{\text{th}}$ grade to adults together, the difference between 2-3 and the older participants is significant (one-sided test, $p = 0.026$ for the first trial and $p = 0.033$ for the second).

3.2 Learning

To determine whether participants learn to play the equilibrium, it is instructive to look at the evolution of behavior at the individual level. Figure 3 depicts the proportion of participants who play the equilibrium strategy in no trial (NN), only the first trial (EN), only the second trial (NE) and in both trials (EE).

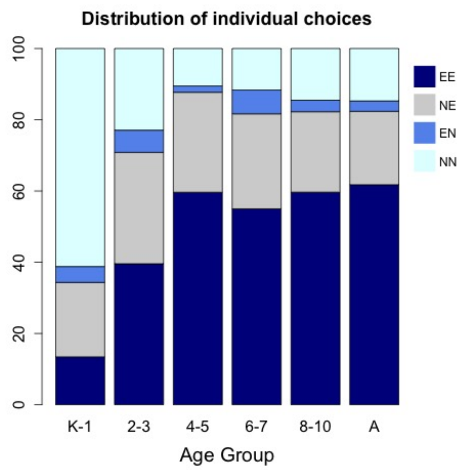

Figure 3: Equilibrium choices of individuals across trials by age group.

Perhaps not surprisingly and consistent with a theory of learning, among the participants who played differently in the two trials (NE and EN), the vast majority (80.0% to 94.1% depending on the age group) moved from not playing at equilibrium to playing at equilibrium.

Naturally, learning depends on feedback. There are two interrelated but somewhat different types of feedback a participant may pay attention to: the outcome in the first trial (win, lose or keep the initial points) and what the partner played in the first trial. We will refer to the first as *payoff feedback* and to the second as *strategy feedback*.

Payoff feedback. We consider the population of participants who do not play the equilibrium strategy in the first trial. For each age group, we compute the likelihood that they lose the trial (obtain no possession) and not lose it (obtain one or both possessions, so that they are not worse-off than initially). The results are reported in Figure 4 (left). We then compute the probability that they play the equilibrium strategy in the second trial as a function of the payoff feedback received. The results are reported in Figure 4 (right).

Since there are only 3 to 6 participants who do not lose after not playing at equilibrium in the first trial in each age group 4-5 and above, we pool them in a new category "Older".

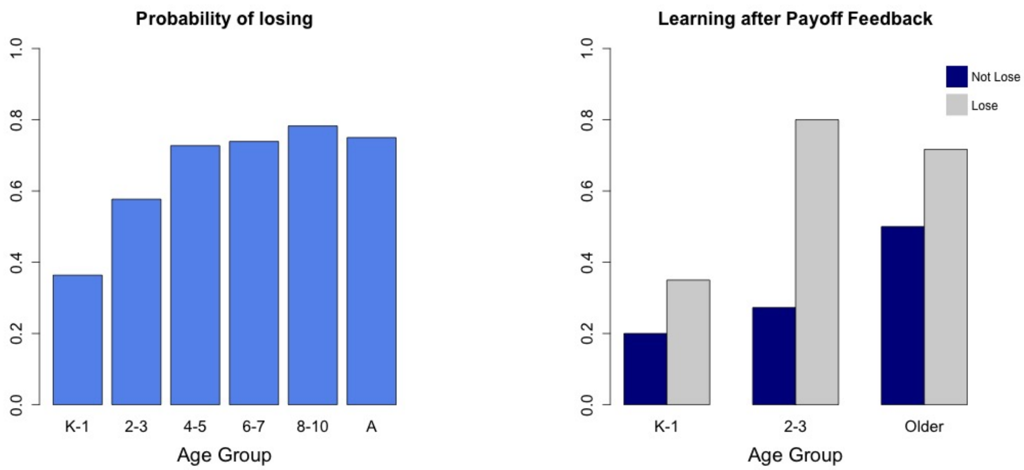

Figure 4: Payoff feedback.
Left: probability of losing after a non-equilibrium choice. Right: probability of learning to play equilibrium in the second trial conditional on losing and not losing in the first trial.

We know from Figures 2 and 3 that the absolute improvement between first and second trial is similar across age groups. Given the behavioral differences in the first trial, it means that among the subjects who miss the first trial, the percentage of participants who learn to play **E** steadily increases from K-1 to 4-5.

There are different reasons for the change in these learning rates. Participants in K-1 are "punished" less frequently for non-equilibrium choices than participants in 4-5 and older school age groups (Figure 4 - left, $p < 0.01$). They also learn less independently of the feedback received (Figure 4 - right, $p < 0.01$ for learning conditional on losing and $p = 0.044$ for learning conditional on not losing). By contrast, participants in 2-3 are punished only slightly less frequently for non-equilibrium choices than participants in 4-5 and older (differences not significant). The main difference is that 2-3 is the only age-group where learning depends significantly on losing vs. not losing the first trial (0.80 vs. 0.27, $p = 0.022$). In other words, in this age group, a negative payoff feedback is key to trigger learning.

Strategy feedback. We hypothesize that observing the strategy of the rival may be revealing as well. We therefore consider again the participants who did not play at equilibrium in the first trial and report in Figure 5 (left) whether they observe an equilibrium strategy played by the rival or not. We then compute the probability that these subjects learn to play the equilibrium as a function of the rival's strategy in the first trial. The results are presented in Figure 5 (right). Again, given the low number of non-equilibrium observations, we pool age groups 4-5 and above in the category "Older".

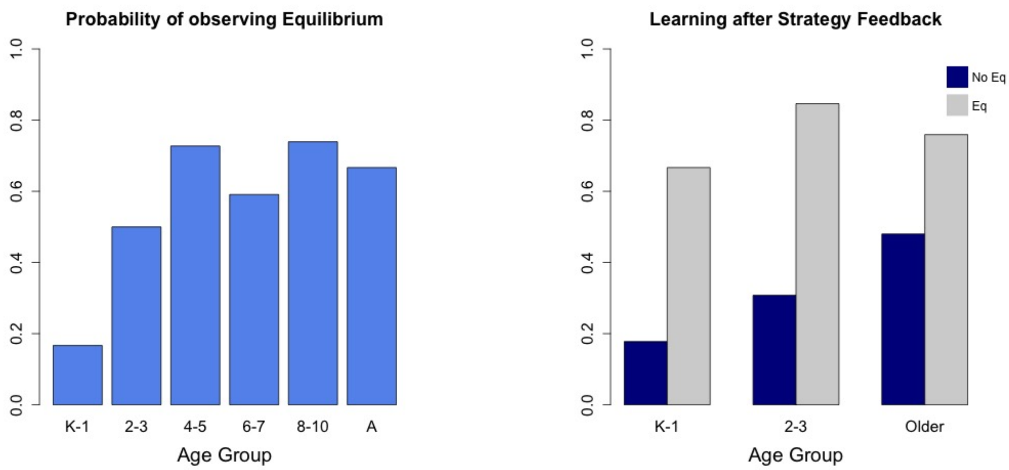

Figure 5: Strategy feedback.
Left: probability that a non-equilibrium player observes an equilibrium choice by the rival. Right: probability of learning to play equilibrium in the second trial conditional on rival playing and not playing the equilibrium strategy in the first trial.

Figure 5 (left) is similar to the equilibrium proportions in Figure 2 (left), except for a higher volatility, simply because the behavior of a participant is a priori independent of the rival's choice in that same trial. It is therefore not surprising that K-1 participants observe significantly less often an equilibrium strategy by the rival than participants in any other age group ($p < 0.01$).

According to Figure 5 (right), learning depends very much on whether the individual observes an equilibrium vs. a non-equilibrium choice by the rival. The difference is large in magnitude (27.9% to 53.8%) and statistically significant in all age groups ($p = 0.008$ for K-1, $p = 0.017$ for 2-3 and $p = 0.028$ for Older).

Overall, the results suggest that there are two potential channels to learn in the two-person beauty contest game: losing and witnessing equilibrium play. These two are evidently correlated: 47% of non-equilibrium players in the first trial both lose and observe an equilibrium choice by the rival. However, the second channel seems to offer a much more effective feedback. Finally, younger children update their behavior less than their older peers, in part because they face less sophisticated rivals, which limits their learning opportunities.

A remaining issue is to determine how participants who did not play at equilibrium in the first trial respond to feedback. We consider several behavioral strategies. *Learn* is the strategy consisting in learning to play **E**. *Naive* refers to a best response to the strategy of the first partner, that is, assuming that the second partner would play the same. A participant naively best responds by moving the army closer to equilibrium (that is, below and to the left of that cell), but not all the way to the bottom left cell. *Miss* refers to any strategy for which the participant does not move closer to equilibrium in the second trial.[8] Figure 6 summarizes the data.

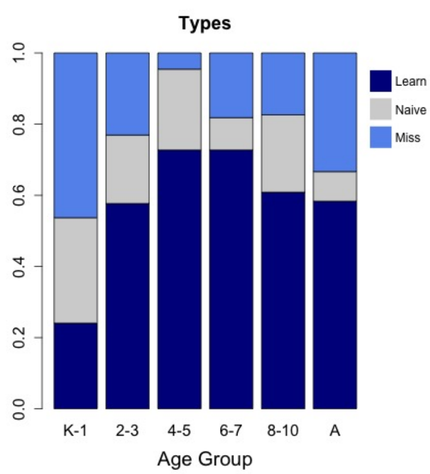

Figure 6: Distribution of types by age group among non-equilibrium players in first trial

The changes in the proportion of participants who learn at different ages has already been discussed, as it corresponds to $\frac{\text{NE}}{\text{NE+NN}}$ in Figure 3. Interestingly and as can be seen from Figure 6, moving closer to equilibrium is not more frequent than playing again a sub-optimal, non-discernible strategy (differences in proportions are not statistically

[8]We also considered *Imitate* for subjects who choose in the second trial the same strategy as the rival chose in the first. Since only three players adopted it, we did not retain it for the analysis.

significant). In other words, many of our non-equilibrium players in the first trial do not recognize the strategic features of the game and miss their second attempt as well.

3.3 Choice processes

By looking at the trajectory of the army, we can obtain relevant information regarding the participant's underlying *reasoning* process. For each individual, our software recorded the intermediate steps prior to settling in a final location. We hypothesize that participants who (formally or intuitively) find the dominant strategy are more likely to drag their army directly to the bottom left corner. By contrast, participants who are hesitant may use a 'tatonnement' process through which they first try a location, then think about the consequences of that choice, and maybe move their army to a different location, until they are satisfied. The objective of this section is to test for the relationship between choices and number of steps.

As it is well known, non-choice data is usually very noisy. We therefore decided to split our sample between participants who use one step to make a decision (*fast* movers) and those who use more than one step (*slow* movers). In other words, anything but a one-step choice is interpreted as providing some (imperfect) evidence of hesitation. Figure 7 depicts the proportion of fast (one step) players by age group and trial.

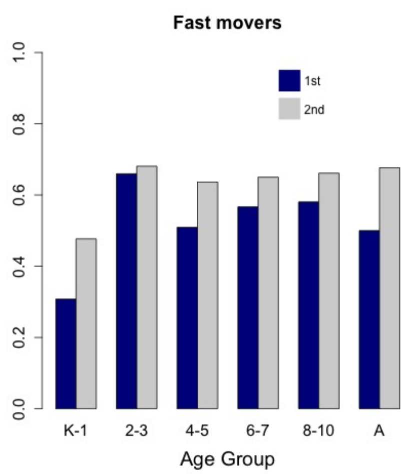

Figure 7: Proportion of fast players by age group in first and second trial.

The overall proportion of fast participants is lower in the first than in the second trial

(51.4% vs. 62.2%, Mc Nemar test, p = 0.003). When we look at each group separately, the difference across trials is marginally significant only for K-1 (p = 0.063). Also, the proportion of fast participants is significantly smaller in K-1 than in all other school age groups in the first trial (p-values between 0.001 and 0.039) and also marginally smaller in the second trial (p-values between 0.051 and 0.077). We found no significant differences between the other age groups.

We next study whether one-step choices correlate with equilibrium behavior. Figure 8 reports the number of fast and slow players in each trial as a function of their final choice, equilibrium vs. non-equilibrium.

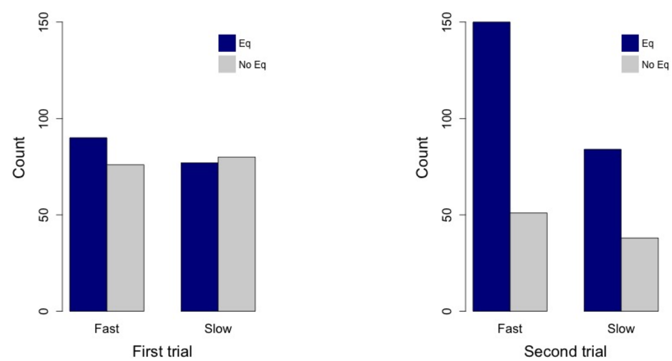

Figure 8: Fast vs. slow play by trial as a function of equilibrium behavior.

There is no relationship between playing fast or slow and choosing the equilibrium strategy in the first trial. The behavior, however, changes in the second trial. Indeed, among equilibrium players, fast play is significantly more frequent than slow play ($p < 0.001$) whereas no such difference is found among non-equilibrium players. Stated differently, fast play is relatively more frequent among subjects who play **E** than among those who do not play **E**. Overall, although our paper does not offer a formal theory to explain a relationship between processes and behavior, it provides suggestive evidence that individuals who learn to play the optimal strategy are also less hesitant.

3.4 Regression analysis

To complement the previous analysis, we run several Probit regressions. First, we consider only participants who do not play at equilibrium in the first trial. We construct a *Learning* dummy variable that takes value 1 if the participant plays at equilibrium in the second trial and 0 if he does not. We regress this variable on *age group* dummies (with 2-3 as the omitted category), a payoff feedback dummy (*Pay-feed*, that takes value 1 if the subject does not lose the trial and 0 if he does) and a strategy feedback dummy (*Strat-feed*, that takes value 1 if the partner plays at equilibrium in the first trial and 0 if he does not). We also control for two demographic variables: gender (Male = 1) and *Siblings* (a dummy that takes value 1 if the participant has one or more siblings). The results are presented in Table 2.

	(1)	**(2)**	**(3)**†	**(4)**†	**(5)**†	**(6)**†,§
K-1	-0.85**	-0.75*	-0.56^{o}	-0.56^{o}	-0.56^{o}	-0.70*
4-5	0.41	0.31	0.16	0.16	0.15	0.13
6-7	0.31	0.21	0.33	0.35	0.34	0.26
8-10	0.08	-0.07	-0.20	-0.20	-0.20	-0.23
A	0.02	-0.11	-0.19	-0.18	-0.18	—
Pay-feed	—	-0.70**	—	0.13	0.13	-0.11
Strat-feed	—	—	1.08***	1.18***	1.18***	0.91*
Male	—	—	—	—	0.09	0.33
Siblings	—	—	—	—	—	-0.21
Const.	0.19	0.50^{o}	-0.31	-0.41	-0.45	-0.15
# Obs.	161	161	159	159	159	136
AIC	211.0	203.2	188.9	190.7	192.5	165.0

^{o}, *, ** and ***: significant at the 10%, 5%, 1% and 0.1% level.

† differences in number of observations reflect missing data

§ adult population is omitted because we have no information on siblings

Table 2: Probit regression of learning to play equilibrium in the second trial

Confirming previous results, K-1 children learn to play the equilibrium at lower rates than their older peers, and there are no significant differences in learning across the other age groups (Column (1)). Losing the first trial or observing the rival play at equilibrium have individually a strong impact on learning and it accounts for part (though not all) of the effect of age (Columns (2) and (3)). Note, however, that the two types of feedback are strongly correlated (Pearson correlation coefficient PCC = -0.75, $p < 0.001$). When we combine both channels, only strategy feedback remains significant (Column (4)). This reinforces the previous finding that, in this game, witnessing the equilibrium strategy is a more powerful learning channel than losing the game. Our demographic variables, gender and siblings, have no significant effect on learning (Columns (5) and (6)).

We then run a second set of Probit regressions on the full sample to establish the determinants of equilibrium play in the second trial. We construct an *Equilibrium* dummy variable that takes value 1 if the participant plays at equilibrium in the second trial and 0 otherwise. We regress this variable on the equilibrium outcome in the first trial *NoEq1* (a dummy that takes value 1 if the subject does not play at equilibrium in the first trial and 0 if he does) the same payoff feedback and strategy feedback variables as before, as well as a choice process variable *Steps2* that captures the number of steps used in the second trial. Given the effect of age on learning described in Table 2, we control for age by simply including a dummy variable *Grade2*$^+$ that takes a value of 0 for participants in K-1 and a value of 1 for grades 2 and above. Finally, we also incorporate an interaction term between Strategy feedback and the outcome of the first trial. The results are presented in Table 3.

	(1)	**(2)**	**(3)**	**(4)**†	**(5)**†	**(6)**†
Steps2	-0.12*	-0.11*	-0.08	-0.08	-0.05	-0.04
NoEq1	—	-1.34***	-1.84***	-1.57***	-1.37***	-1.75***
Pay-feed	—	—	-0.84***	-0.48$^{\cdot}$	-0.40	0.13
Strat-feed	—	—	—	0.50*	0.34	-0.24
Grade2$^+$	—	—	—	—	0.74***	0.67**
NoEq1 × *Strat-feed*	—	—	—	—	—	1.32**
Const.	0.80***	1.58***	2.36***	1.73***	1.02*	0.85$^{\cdot}$
# Obs.	326	326	326	322	322	322
AIC	385.6	318.9	304.9	297.6	288.11	280.91

*, ** and ***: significant at the 5%, 1% and 0.1% level
† differences in number of observations reflect missing data

Table 3: Probit regression of equilibrium play in the second trial

As before, payoff feedback and strategy feedback are exogenously linked (PCC $= -0.34$, $p < 0.001$). Also, participants who play at equilibrium in the first trial cannot lose, resulting in a correlation between *NoEq1* and *Pay-feed* (PCC $= -0.65$, p-value < 0.001). Table 3 shows that behavior in the second trial is strongly predicted by choice in the first trial and strategy feedback. Performing fewer intermediate steps is an indicator of equilibrium play only in the absence of feedback variables, because the latter convey more accurate information. As in previous regressions, K-1 participants are significantly less likely to play the equilibrium strategy than the older individuals. Finally, the high significance of the interaction term in Column (6) confirms that quality feedback –especially, observing equilibrium behavior by the other player after a deviation– is an essential element for learning to strategize.

4 Concluding remarks

We have designed a strategic game that can be understood and played correctly by participants as young as 5 years old and, at the same time, can be moderately challenging for college undergraduates. We have found a significant increase in equilibrium behavior from K-1 to 4-5, and no improvement afterwards. Participants in all age categories learn to play closer to equilibrium by the second trial, but the improvement varies across ages. Our results suggest that feedback is an essential element of learning how to strategize. Observing the equilibrium choice of the partner guided a large fraction of our participants towards equilibrium play whereas only a few participants learned independently. A main difference between very young children and older ones is that the former are less often exposed to equilibrium strategies, and are therefore given fewer opportunities to learn. This explains part of the difference in choice improvements.

The observed trajectory is consistent with developmental theories that address logical thinking. Even though the game is simple, it requires the ability to form a hypothesis about what the partner may do and use this knowledge to select a best response. Evidence has shown that children can think ahead and have correct anticipations by 7 years of age (Eliot et al., 1979; Shultz and Cloghesy, 1981; Tecwyn et al., 2014). This explains why equilibrium play can be observed during elementary school. Children have also been shown to develop the ability to think logically about what they observe (inductive logic) between the ages of 8 and 12 (Feeney and Heit, 2007). This ability is likely to support learning from feedback, which is precisely observed to improve over that age range.

The surprising feature is to notice that behavior stabilizes at the end of elementary school. The ability to reason abstractly (hypothetical, counterfactual) is known to develop fully throughout middle school (Piaget, 1972; Rafetseder et al., 2013). We therefore expected some improvement in the rate of equilibrium play after 6^{th} grade: with age, it should become easier for a participant to form counterfactuals about the behavior of their partner and their own choices. Our result, also found in Brocas and Carrillo (2018b), suggests that even though the ability develops in tests of logic, it is not necessarily transferred to strategic environments.

We also observe heterogeneity in behavior within each age group: some participants always play at equilibrium and some participants never learn to play at equilibrium even after relevant feedback. This result is consistent with neo-piagetian theories of cognitive development that reject the concept of strict stages of development and emphasize individual differences (Morra et al., 2012).

The way children learn from feedback is a fascinating topic that requires further investigation. The idea that different kinds of feedback may have different effects on decision-making is reminiscent of the experimental literature on communication in adults. Schotter (2003), Schotter and Sopher (2003) and the subsequent literature shows that "naïve" advice has welfare improving effects in many environments compared to the mere observation of other's actions. In the context of our study, it may be interesting to consider a variant in which children are allowed to discuss their strategy after the first round. We could then introduce the children to a new game requiring the same logical features but a different presentation to distinguish between imitation and understanding. These issues are of primary importance to assess how logical concepts relevant to decision-making (such as inductive logic, deductive logic, counterfactual logic, backward induction, etc.) can be efficiently taught and transferred to day-to-day behavior.

References

Isabelle Brocas and Juan D Carrillo. Iterative dominance in young children: Experimental evidence in simple two-person games. *Journal of Economic Behavior & Organization*, 2018a.

Isabelle Brocas and Juan D Carrillo. Steps of reasoning in children and adolescents. USC Working Paper, 2018b.

Isabelle Brocas and Juan D Carrillo. Studying decision-making in children: challenges and opportunities. USC Working Paper, 2019.

Isabelle Brocas and Juan D Carrillo. The development of deceptive behavior in mixed strategy games. Work in Progress, 2020.

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

Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.

Gary Charness, John A List, Aldo Rustichini, Anya Samek, and Jeroen Van De Ven. Theory of mind among disadvantaged children: Evidence from a field experiment. *Journal of Economic Behavior & Organization*, 166:174–194, 2019.

Chun-Ting Chen, Chen-Ying Huang, and Joseph Tao-yi Wang. A window of cognition: Eyetracking the reasoning process in spatial beauty contest games. *Games and Economic Behavior*, 111:143–158, 2018.

Jingnan Chen, Daniel Houser, Natalia Montinari, and Marco Piovesan. Beware of popular kids bearing gifts: A framed field experiment. *Journal of Economic Behavior & Organization*, 132:104–120, 2016.

Eileen Chou, Margaret McConnell, Rosemarie Nagel, and Charles R Plott. The control of game form recognition in experiments: Understanding dominant strategy failures in a simple two person "guessing" game. *Experimental Economics*, 12(2):159–179, 2009.

Ramon Cobo-Reyes, Jose J Dominguez, Fernando García-Quero, Brit Grosskopf, Juan A Lacomba, Francisco Lagos, Tracy Xiao Liu, and Graeme Pearce. The development of social preferences. *Journal of Economic Behavior & Organization*, 2019.

Giorgio Coricelli and Rosemarie Nagel. Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 106(23): 9163–9168, 2009.

Miguel A Costa-Gomes and Vincent P Crawford. Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5):1737–1768, 2006.

Simon Czermak, Francesco Feri, Daniela Glätzle-Rützler, and Matthias Sutter. How strategic are children and adolescents? experimental evidence from normal-form games. *Journal of Economic Behavior & Organization*, 128:265–285, 2016.

John Eliot, Kenneth Lovell, C Mitchell Dayton, and BF McGrady. A further investigation of children's understanding of recursive thinking. *Journal of Experimental Child Psychology*, 28(1):149–157, 1979.

Eduardo Fe and David Gill. Cognitive skills and the development of strategic sophistication. *Available at SSRN 3285484*, 2018.

Aidan Feeney and Evan Heit. *Inductive reasoning: Experimental, developmental, and computational approaches*. Cambridge University Press, 2007.

Brit Grosskopf and Rosemarie Nagel. The two-person beauty contest. *Games and Economic Behavior*, 62(1):93–99, 2008.

William T Harbaugh, Kate Krause, and Timothy R Berry. GARP for Kids : On the Development of Rational Choice Behavior. *The American Economic Review*, 91(5): 1539–1545, 2001.

Teck-Hua Ho, Colin Camerer, and Keith Weigelt. Iterated dominance and iterated best response in experimental" p-beauty contests". *The American Economic Review*, 88(4): 947–969, 1998.

JA List, R Petrie, and A Samek. How experiments with children can inform economics. Working paper, 2018.

Sergio Morra, Camilla Gobbo, Zopito Marini, and Ronald Sheese. *Cognitive development: neo-Piagetian perspectives*. Psychology Press, 2012.

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

Jean Piaget. Intellectual evolution from adolescence to adulthood. *Human development*, 15(1):1–12, 1972.

Eva Rafetseder, Maria Schwitalla, and Josef Perner. Counterfactual reasoning: From childhood to adulthood. *Journal of experimental child psychology*, 114(3):389–404, 2013.

Andrew Schotter. Decision making with naive advice. *American Economic Review*, 93(2): 196–201, 2003.

Andrew Schotter and Barry Sopher. Social learning and coordination conventions in intergenerational games: An experimental study. *Journal of political economy*, 111(3): 498–529, 2003.

Itai Sher, Melissa Koenig, and Aldo Rustichini. Children's strategic theory of mind. *Proceedings of the National Academy of Sciences*, 111(37):13307–13312, 2014.

Thomas R Shultz and Karen Cloghesy. Development of recursive awareness of intention. *Developmental Psychology*, 17(4):465, 1981.

Dale O Stahl. Boundedly rational rule learning in a guessing game. *Games and Economic Behavior*, 16(2):303–330, 1996.

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.

Emma C Tecwyn, Susannah KS Thorpe, and Jackie Chappell. Development of planning in 4-to 10-year-old children: Reducing inhibitory demands does not improve performance. *Journal of experimental child psychology*, 125:85–101, 2014.

Appendix A: Instructions

[Instructions for Kindergarten to 5$^{\text{th}}$ grade where points are exchanged for toys]

Hi, my name is [name] and these are my helpers [introduce helpers]. Today, we are going to play a few games with you. Do you want to play games? In all the games, you are going to win points. At the end, you can exchange the points for the toys you saw in the shop. Are you excited? Remember: if you win more points you will be able to get more toys.

Army Game

Our first game is called the "Army" game. In this game, you will play with someone in this room. The computer decides whom this other person is but you will not know it, and the point of the game is not to find out.

In the game, you are either the "Tree Leader" or the "Volcano Leader". The Tree Leader is GREEN (like the leaves) and the Volcano Leader is RED (like the lava). The computer decides who is the "Tree Leader" and who is the "Volcano Leader".

If you are the Tree Leader, your screen looks like this.

[SLIDE 1]

You own the Tree army, the GREEN dot on your screen [point]. The tree army is worth 20 points. You also own the Trees [point] that are worth 6 points. You do not own the volcanoes. So you own 26 points in total.

If you are the Volcano Leader, your screen looks like this.

[SLIDE 2]

You own the Volcano army, the RED dot on the screen [point]. The volcano army is also worth 20 points. You also own the Volcanoes [point] that are worth 6 points. You do not own the trees. So, just like the other person, you own 26 points in total.

These are both screens together.

[SLIDE 3]

If you are the Tree Leader, your job is to protect the trees and to win the volcanoes. And if you are the Volcano Leader, your job is to protect the volcanoes and to win the trees. You do so by moving your army on the grid. Moving the army is a secret, so you cannot see the army of the other person, and the other person cannot see your army [point to dots on screenshots].

How do you protect what you have?

If you are the Tree Leader, you protect the trees if your army is closer to the Trees than the other army. If the other army is closer to the trees than your army, then you lose the trees. If you are the Volcano Leader, you protect the volcanoes if your army is closer to the Volcanoes than the other army. If the other army is closer to the volcanoes than your army, then you lose the volcanoes.

How do you win what the other leader has?

If you are the Tree Leader, you win the volcanoes if your army is closer to the volcanoes than the other army. And if you are the Volcano Leader, you win the trees if your army is closer to the trees than the other army.

How many points do you get?

Anyone who ends up owning the trees gets the 6 points they are worth. Anyone who ends up owning the volcanoes gets the 6 points they are worth. If you are at the same distance then whoever had it in a first place keeps it. Also, you will never lose your army, which is worth 20 points. Are there any questions? OK, let's play then.

[Launch experiment]
[After choices are made, feedback page appears]

The screen in front of you is similar to the one on this screen.

[SLIDE 4]

It tells you the location of your army and the location of the other army. It also tells you whether you got the Trees, the Volcanoes or both, and the total number of points you earned.

[Go to each terminal and make sure every kid understands the payoffs]

Now, we will play the same game again. You will remain a Tree Leader or a Volcano Leader but you will play with a different person. Remeber the computer chooses the person but it does not tell you. Are you ready?

[After choices are made, feedback page appears]

Again, the screen in front of you tells you the location of your army and the location of the other army and the points you earned. Did you like this game?

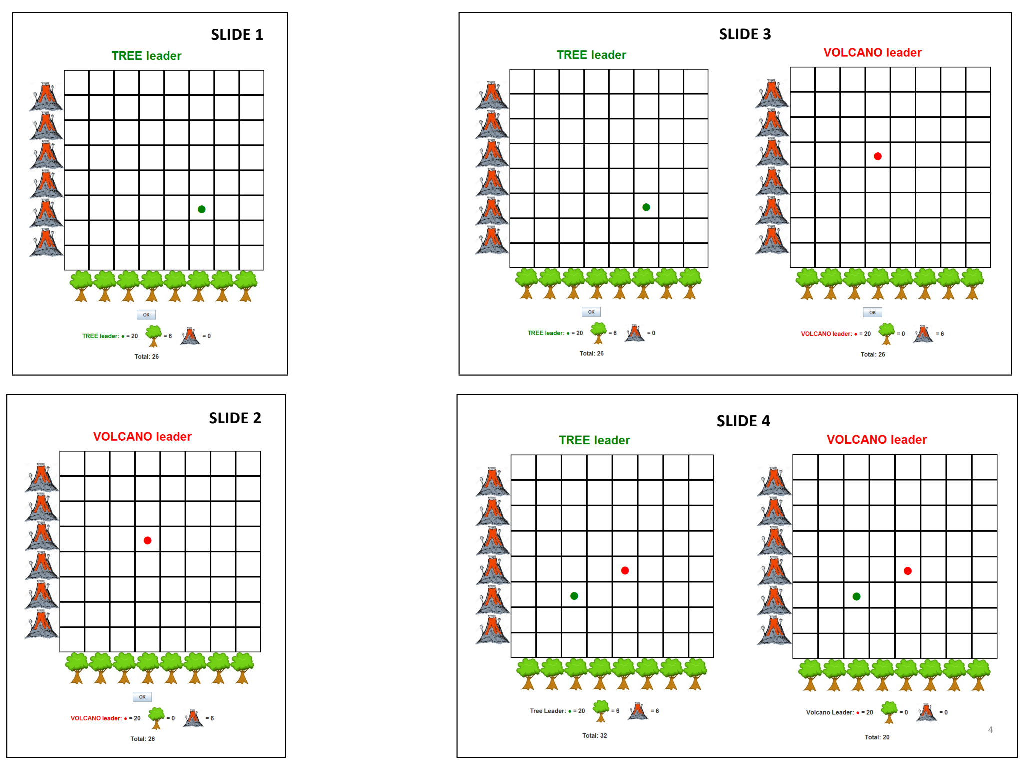

Figure 9: Slides projected during instructions