Young children build consensus in networks with local information *

Isabelle Brocas
University of Southern California, IAST and CEPR

Juan D. Carrillo
University of Southern California, IAST and CEPR

Ulysses Rios
University of Southern California

December 2024

Abstract

This study investigates the foundational mechanisms of consensus-building in young children aged 5 to 8, focusing on their capacity to coordinate decisions in network environments with limited information. The findings reveal a sharp developmental progression, with older children reaching near-perfect consensus rates. Effective coordination stems from two main factors. First, imperfect imitation: children follow the majority behavior of the peers with high probability but not with certainty. Second, patience heterogeneity: children vary their timing, with some acting quickly to provide direction and others delaying their choice to lock consensus. Comparing children's choices with simple algorithms, we observe that while algorithms broadly capture behaviors, children's intrinsic flexibility and responsiveness to social cues allows them to outperform simple computational models, particularly in the more complex conditions. These findings suggest that coordination success stems from traits that support complementary roles and adaptable behaviors. Furthermore, this work contributes to understanding developmental milestones in decision-making, providing a foundation for future investigations into how children navigate complex social environments.

Keywords: developmental decision-making, coordination, network experiments.

JEL Classification: C90, C92, D85.

*We thank the members of the Los Angeles Behavioral Economics Laboratory (LABEL), especially Daniel Hermosillo, Prianca Nadkarni and Andrea Victoria, for their insights and help running the experiment and to Sergio Sapiña from playstudies.com for his outstanding programming skills. We are grateful to the staff of the Lycée International de Los Angeles (LILA) –in particular David Padilla– for their support. The study was conducted with the University of Southern California IRB approval UP-12-00528. We gratefully acknowledge the financial support of the National Science Foundation (NSF) grant SES-2315770 and the French National Research Agency (ANR) grant ANR-17-EURE-0010.

The ability to coordinate actions toward a shared goal is essential for success across economic and social activities. In the workplace, effective coordination boosts team efficiency, speeding project completion and improving outcomes. Large-scale scientific efforts rely on coordination among researchers to align methods, share data, and achieve breakthroughs. In healthcare, coordinated efforts among professionals enhance patient care, reducing errors and enhancing results. Environmental resource management also depends on stakeholder coordination to sustainably manage resources and tackle challenges like climate change and biodiversity loss. Understanding *how* and *when* people develop these coordination abilities is therefore critical to fostering effective collaboration.

Experimental economics research indicates that mutually advantageous coordination is difficult to achieve in single-play games (Dal Bó et al., 2021), yet tends to emerge when two actors engage repeatedly (McKelvey and Palfrey, 2001). However, in many real-world situations, people lack the opportunity to build expertise through repeated interactions. More critically, these interactions are often complex, as individuals typically operate within large social networks where they directly interact with only a subset of nearby contacts (Jackson, 2008; Goyal, 2012; Jackson et al., 2022). This networked structure adds layers of complexity to coordination, as individuals must navigate indirect influences and local information rather than relying on global knowledge or repetitive encounters.

In parallel, research in computer science has demonstrated that individuals in complex networks can still achieve coordination when allowed some form of communication before finalizing their choices. The seminal work by Kearns and colleagues (Kearns et al. (2006); Judd et al. (2010); Kearns (2012), etc. collectively referred to as [$\mathcal{K}$]) shows that individuals with only local information can effectively solve global problems, such as graph coloring and consensus-building, through an iterative "choose-observe-adjust" process. This tâtonnement-like approach resembles indirect communication: players choose a color, observe their neighbors' choices, and adjust their own choice accordingly. This method enables local interactions to support global coordination, highlighting the potential of structured observation and adaptation in complex social networks.

In this paper, we investigate the inherent nature of coordination abilities and identify specific features of the decision process that tend to facilitate or hinder coordination in simple networks. To achieve this, we examine a streamlined version of the consensus problem in [$\mathcal{K}$]. Our experiment involves a six-person complete graph network, where each participant has either two or four color options and can observe the choices of either two or three neighbors. Participants have 30 seconds to adjust their choices, and the task is deemed successful if all six individuals simultaneously select the same color. This setup allows us to explore how limited choice options and neighbor visibility impact coordination in a controlled environment.

Unlike the problems in [$\mathcal{K}$], this task is straightforward for adults. However, we focus on a *very young* population—ages 5 to 8—whose cognitive skills, such as impulse control and attention, are still developing. Our setup may present challenges for our participants while also allowing us to identify which elements of coordination are innate or acquired very early. We anticipate that our young participants may show stubbornness or inattention, making non-convergence likely, as a single deviation can disrupt group success. We also expect children to rely on very simple imitation algorithms, observing and following peers' choices to guide their own.

As a preview, the experiment yields two major findings. First, there is a marked discontinuity in coordination success between kindergarteners (ages 5-6) and first or second graders (ages 6-8): kindergarteners converge only 36% to 72% of the time, while older children almost always reach convergence, at rates of 81% to 100%. This suggests that while coordination ability is not entirely innate, it develops very early, with a noticeable shift around age 6. Second, simple imitation of neighbors' choices is insufficient for solving the task. Instead, the near-perfect convergence observed in first and second graders relies on two nuanced aspects of their decision-making process: (small) *choice frictions* and (large) *patience heterogeneity.* Small choice frictions refer to the advantage of imitating the actions of the majority of the neighbors with high probability, but not with absolute certainty. Large patience heterogeneity refers to the benefit of some players acting quickly to set direction while others wait to observe and reinforce consensus. To support our analysis, we develop several algorithms of increasing sophistication and compare simulated behavior to children's actual decisions. While these algorithms broadly capture the patterns in children's choices, they do not perform as well as the oldest students, highlighting the flexibility and responsiveness of children's heuristic strategies to social cues even at a young age.

1 Results

1.1 Research design

We conducted the experiment with 150 students from LILA, a K-12 private school in Los Angeles, comprising 54 kindergartners (33f, 21m, aged 5-6), 48 first graders (26f, 22m, aged 6-7) and 48 second graders (24f, 24m, aged 7-8). We refer to these groups as **K**, **1** and **2**. In the experiment, groups of six children from the same grade participated in the "consensus task" described in Figure 1, where they chose between either 2 or 4 colors (C2 or C4) while observing the decisions of either 2 or 3 neighbors (N2 or N3). We employed a 2×2 within-subject design. Participants played 4 rounds in each condition (C2N2, C2N3, C4N2, C4N3) in counterbalanced blocks of two rounds with the same five partners, totalling 16 rounds. We anticipated consensus building to be more difficult with more color choices and fewer peers to observe.

Each round began with a 30 second countdown, during which participants started without an initial color. Within this timeframe, they could select and change their color choice as often as desired. A round was considered a “success” if all players converged on the same color within the 30 seconds; otherwise, it was marked as a “failure” if the timer expired without convergence. At the end of each round, participants received feedback on the outcome (success or failure) and the distribution of participants across each color choice.

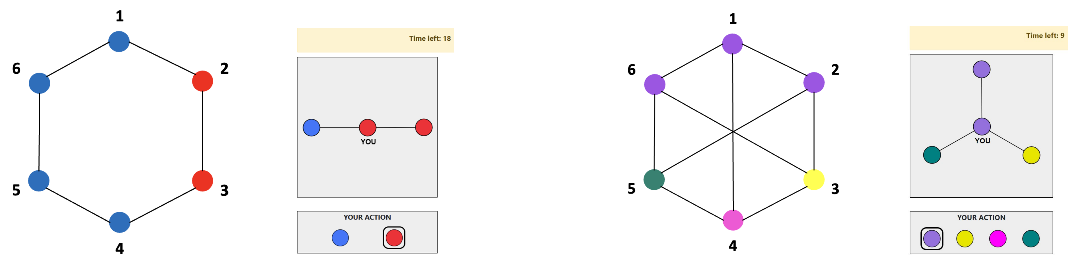

(a) **(b)** **(c)** **(d)**

Figure 1: Consensus Task. Panels (a) and (c) show the full network structure–visible only to the experimenter–for C2N2 and C4N3, respectively. Each node corresponds to a player’s current color choice, with lines representing the direct connections (observed neighbors) and the number next to the node capturing the player’s ID. Panels (b) and (d) show screenshots of the graphical user interface from the perspective of player ID #2. The player is marked as “you” in the center, with neighbors arranged equidistantly to prevent framing effects. Players can change their color choice simply by tapping on an option under “your action”, with a minimum interval of one second between changes. As soon as all players end up on the same color, the timer stops, marking the round as a success.

Participants were incentivized with points, which they could exchange for toys of their choice in our toy shop at the end of the session. They were explicitly instructed that earning more points meant receiving more toys. Each session lasted around 30 minutes, with all participants earning between 4 and 8 toys. The full set of instructions, read aloud to participants, are available in SI1.

1.2 Network outcomes

The primary question is to assess the network’s performance. Given the uncertain difficulty level in reaching consensus, we introduce a comparison benchmark using a *simple* algorithm, referred to as **AL**: “Each player initially chooses a color at random. Next, a player is randomly selected and adopts the color of the majority within their neighborhood, counting their own choice when the number of neighbors is even (N2) and excluding it when the number of neighbors is odd (N3). Afterward, another player is randomly chosen from the remaining five and follows the same rule. This process continues for a maximum of 30 moves after the initial selection."[1] While other algorithms could be considered, this serves as a simple yet reasonable reference benchmark.

Likelihood of convergence

In Figure 2, we begin by calculating the probability of convergence within the allotted 30 seconds and comparing it to our algorithm's performance.

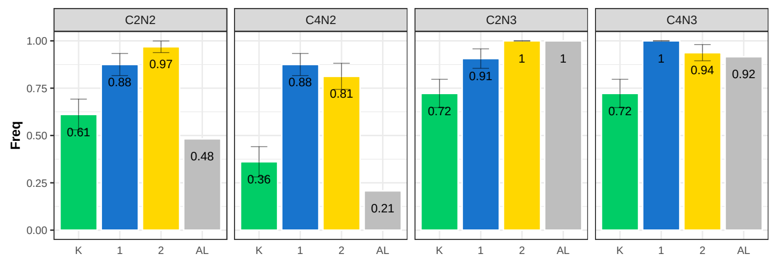

Figure 2: Percentage of network convergence for grades **K**, **1** and **2**, and comparison with 10,000 simulations of our behavioral algorithm, **AL**, computed separately for each condition (C2N2, C4N2, C2N3, C4N3).

The results demonstrate a very significant age-related improvement within our window of observation. First and second graders exhibit an impressive ability to reach consensus across all conditions, with network success rates ranging from 26 to 32 out of 32. There are no statistically significant differences between **1** and **2** (test of differences in proportions, $p > 0.30$). By contrast, kindergartners face greater difficulty, with consensus achieved in only 13 to 26 out of 36 rounds, with all proportions significantly different from the older grades ($p < 0.05$) except for **K** v. **1** in C2N3. The algorithm performs (surprisingly) poorly with two neighbors (even below **K**'s performance) but excels with three neighbors, underscoring the role of increased connections in facilitating consensus. More generally, **K**'s behavior is consistent with the hypothesis that consensus is hardest with more colors and fewer neighbors (C4N2). These findings are further supported by a Probit regression analysis of network convergence probability (see SI2 for details).

Speed of convergence

Next, we examine the speed of convergence as an additional performance metric. When examining the empirical number of choices needed to reach consensus (conditional on suc-

[1]Note that the only reason the algorithm is different in N2 and N3 is to prevent ties.

cessful convergence), there are no statistically significant differences across grades, except between **2** v. **K** and **2** v. **1** in condition C4N2 (pairwise t-test of mean comparison with pooled SD, Holm adjustment, 5% significance level). To enhance statistical power, Figure 3 presents the distribution of choices until consensus, pooled across all grades but separated by condition. This distribution is also compared to the algorithm's performance.

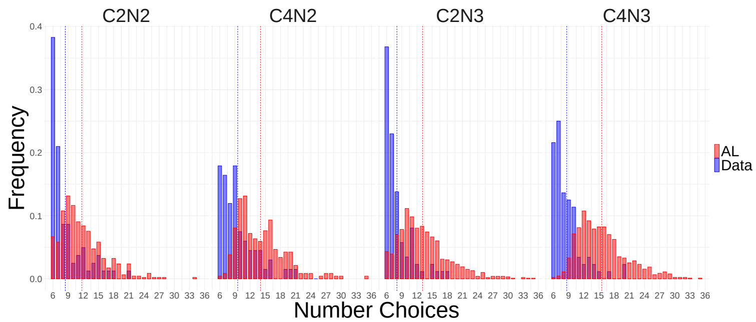

Figure 3: Distribution of the total number of choices in the network, conditional on successful convergence, for our combined population (all grades grouped) and for the algorithm, calculated separately for each condition (C2N2, C4N2, C2N3, C4N3) [vertical lines indicate the average values in the graph].

Participants reach consensus significantly faster than the algorithm predicts (Kolmogorov-Smirnov test, $p < 0.001$ for all four conditions). Indeed, 72% to 84% of participants converge within 10 choices, compared to only 9% to 49% in the algorithm. Convergence is rare after 30 choices (or even earlier) validating the experimental time limit and the theoretical choice cap in the algorithm.

Overall, **AL** accurately captures the likelihood of convergence for **K** in N2 and for **1** and **2** in N3, though it fails to capture the speed of convergence in any condition.

1.3 Individual behavior

Next, we analyze individual behavior to gain insight into the strategies our participants employ. Our methodology is as follows: first, we identify the empirical features of participants' choices, and then we incorporate these features into **AL** to determine if the modified algorithm better aligns with the observed behaviors.

Dynamic imitation

For each participant, we calculate how frequently they follow the strategy prescribed by the algorithm. Specifically, we observe the network from the perspective of individual i at a given moment and check if their choice aligns with **AL**. We then examine the network after one neighbor has changed their action, repeating this analysis iteratively until no further changes occur. From this, we compute p_i, the fraction of i's choices consistent with the algorithm and designate $1-p_i$ as the level of non-compliance with the algorithm, or "choice frictions". This procedure is repeated for all participants.[2] Figure 4 presents a histogram of the empirical distribution of p_i across **K**, **1** and **2**.

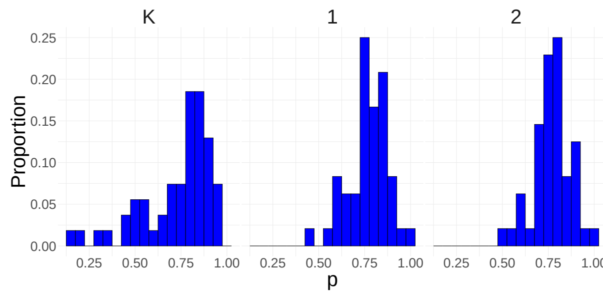

Figure 4: Empirical fraction of p_i, representing the choices consistent with **AL**, presented by grade.

The average empirical fraction of choices consistent with **AL** by grade is $\hat{p}_{\mathbf{K}} = 0.73$, $\hat{p}_{\mathbf{1}} = 0.77$ and $\hat{p}_{\mathbf{2}} = 0.77$, yielding a population average of $\hat{p} = 0.75$. Friction disparities are higher in **K**, supporting the idea that confusion or inattention ($p_i < 0.50$)—though rare—occurs primarily among the youngest participants. Since the deviation by just one individual can disrupt convergence, this might explain their lower success rate in reaching consensus. In general, compliance with **AL** is high ($p_i \in [0.7, 0.9]$ for two-thirds of participants), yet perfect or almost perfect compliance remains uncommon.

To assess whether these frictions are beneficial or detrimental to performance, we extend our algorithm to incorporate the possibility of deviations. Figure 5 displays the convergence of **AL** as we vary p ($\in [1/2, 1]$), which we call $\mathbf{AL_p}$ (where, naturally, $\mathbf{AL_1} \equiv$ **AL**). In this extended model, each individual follows the majority with probability p ($\in [1/2, 1]$) and deviates with probability $1-p$. Deviations involve either keeping their current choice when the strategy requires switching or randomly changing to one of the minority colors when the strategy requires maintaining their choice.

From a theoretical perspective, *some* departures from full compliance with **AL** are highly beneficial in N2. They can be either beneficial or detrimental in N3. The reason is straightforward: while deviations generally delay convergence, they also prevent scenarios where two subgroups lock into different colors with no incentives to switch—a configuration more common in N2. In such cases, deviations help overcome impasses, facilitating consensus that might otherwise be unreachable.

The level of frictions that maximizes convergence across all conditions is $p^* = 0.86$, higher but close to the observed empirical level ($\hat{p} = 0.75$). Notably, convergence under

[2] We consider all configurations where at least one neighbor has already made a choice.

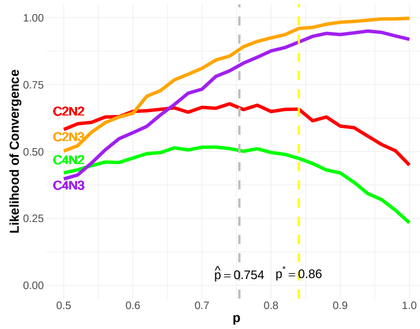

Figure 5: $\mathbf{AL_p}$: convergence of the algorithm by condition as a function of choice frictions (we run 3,000 simulations for $p \in [0.5, 1]$ in 0.02 intervals). We also report the empirical average ($\hat{p}$) and the optimal average level of choice frictions (p^*).

the empirical friction level ($\hat{p}$) surpasses convergence under full compliance ($p = 1$) when averaged across all conditions, suggesting that moderate frictions can enhance overall performance. At the same time, choice frictions notably decrease the speed of convergence (see SI3 for details).

Initial choice

Another key feature of our algorithm is that all six participants make their initial choices randomly. However, several indicators suggest this assumption does not hold empirically. Firstly, the probability of achieving convergence in exactly six moves (one per participant) is 30% in C2 and 19% in C4—much higher than the expected 3.1% and 0.03% if initial choices were entirely random. Additionally, conditional on achieving convergence, the likelihood that the first color selected by the first player to choose aligns with the final consensus color is 73% in C2 and 67% in C4 (compared to 50% and 25% respectively if initial choices were random). Finally, for rounds converging in more than six moves, there is a 43% likelihood in C2 and 37% in C4 that consensus is reached immediately after the first move of one of the players. The first two observations suggest that imitation occurs even within the first choices of the players (hence, not random). The third observation suggests a "sniping" behavior, where at least one participant delays their choice, waiting for an opportune moment to make the decisive, consensus-reaching move.

To evaluate the degree of sequentiality in participants' initial choices, we construct a ranking measure. For each participant, we determine the order of their initial choice within their network for each round (e.g., first one to move, second one to move, etc.) and calculate the average ranking across all 16 rounds. Figure 6 shows the empirical distribution of these average rankings across **K**, **1** and **2**.

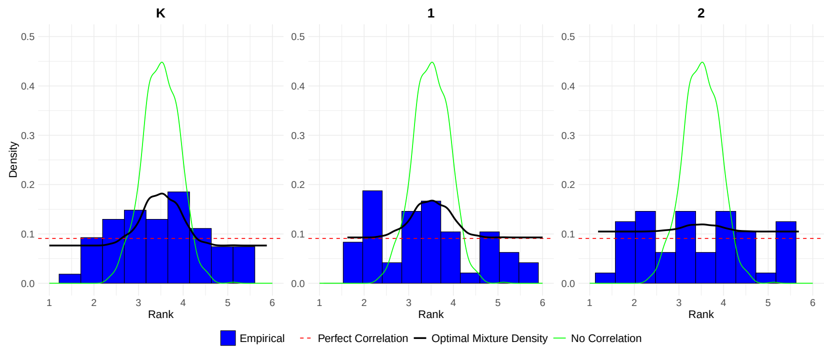

Figure 6: Distribution of average rankings in the population (blue histogram), presented by grade. For comparison, we include the two polar cases, where the 16 rankings are random across rounds (solid green line) and perfectly correlated across rounds (dashed red line).

The distribution of initial choice rankings reveals substantial heterogeneity across individuals (blue histogram): some participants systematically make their first decision rapidly, while others systematically wait and observe their peers. This pattern aligns more closely with a consistent ranking across rounds (dashed red line) than with a random ranking across rounds (solid green line). The dispersion is also higher in older children. This, in turn, reinforces the idea that $\mathbf{AL}_{\hat{\mathbf{p}}}$'s assumption of random initial choice does not accurately capture the behavior of our participants.

To evaluate the level of patience heterogeneity in our population, we estimate the average degree of sequentiality in participants' initial choices. It involves creating a weighted combination of simulated densities with perfect rank correlation and no rank correlation. Using maximum likelihood estimation techniques, we find that the mix that best matches the data's histograms in each grade (thick black solid line in Figure 6) are $\hat{q}_{\mathbf{K}} = 0.24$, $\hat{q}_{\mathbf{1}} = 0.17$ and $\hat{q}_{\mathbf{2}} = 0.03$, for a population average of $\hat{q} = 0.19$. The result confirms a higher degree of sequentiality in **2** than in **1** or **K**.

Does choice sequentiality improve performance? To explore this, we further extend our algorithm by assuming that each player's initial choice is simultaneous with the preceding player with probability q and sequential with probability $1 - q$. Here, $q = 1$ represents minimal patience heterogeneity, while $q = 0$ represents maximal patience heterogeneity.[3]

[3] Specifically, one player is randomly chosen to make the first initial action, which they select randomly. A second player is then chosen at random. If this second player is not a neighbor of the first, they make a random choice. However, if the first player is a neighbor, then with probability q, the second player's choice is "simultaneous" (i.e., made without observing the first player's action, thus remaining random). With probability $1 - q$, the choice is "sequential" (i.e., made after observing the first player's action) and follows the prescriptions of **AL**. This process repeats until all six players have made their initial choice, at which point the algorithm proceeds as before.

Figure 7 depicts $\mathbf{AL}_{\hat{\mathbf{p}},\mathbf{q}}$, the theoretical likelihood convergence as we vary patience heterogeneity q ($\in [0,1]$) and given the empirically observed level of choice frictions $\hat{p}$ (naturally, $\mathbf{AL}_{\hat{\mathbf{p}},\mathbf{1}} \equiv \mathbf{AL}_{\hat{\mathbf{p}}}$).

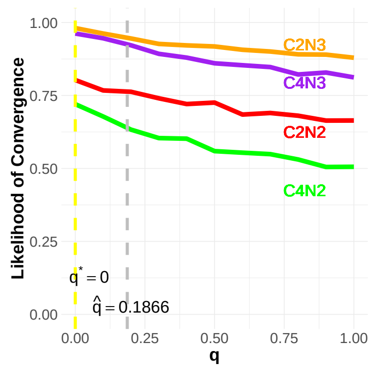

Figure 7: $\mathbf{AL}_{\hat{\mathbf{p}},\mathbf{q}}$: convergence of algorithm given $\hat{p}$ choice frictions as a function of q. We also report the empirical average ($\hat{q}$) and the optimal average patience heterogeneity (q^*).

Convergence is monotonically decreasing in q. Sequentiality in initial choices significantly enhances imitation and, consequently, convergence, as it allows participants to follow the initial choices of others. This approach, however, depends on a high level of coordination regarding who acts and who waits, given the absence of a pre-specified order of moves. The effect is roughly linear and slightly more pronounced in N2, where achieving consensus is inherently more challenging. Furthermore, sequentiality impacts the distribution of choices until convergence even more than the likelihood of convergence itself. Indeed, the updated algorithm $\mathbf{AL}_{\hat{\mathbf{p}},\hat{\mathbf{q}}}$ aligns much more closely with the observed speed of convergence among participants than the initial **AL** (see SI3 for details).

To summarize our findings, Table 1 compares the likelihood of convergence across our three participant groups and the three algorithmic models.

	empirical			algorithm *		
	K	**1**	**2**	**AL**	$\mathbf{AL}_{\hat{\mathbf{p}}}$	$\mathbf{AL}_{\hat{\mathbf{p}},\hat{\mathbf{q}}}$
% convergence						
C2N2	0.61	0.88	0.97	0.48	0.67	0.77
C4N2	0.36	0.88	0.81	0.21	0.50	0.64
C2N3	0.72	0.91	1.00	1.00	0.89	0.95
C4N3	0.72	1.00	0.94	0.92	0.82	0.92

* $\hat{\mathbf{p}} = 0.75$ and $\hat{\mathbf{q}} = 0.19$

Table 1: Summary comparison of performance between participants and algorithms

The algorithm incorporating small choice frictions and high patience heterogeneity, $\mathbf{AL}_{\hat{\mathbf{p}},\hat{\mathbf{q}}}$, underestimates convergence in N2 but otherwise captures reasonably well the probability of consensus-building in grades **1** and **2**.

We also perform an OLS regression and show that individuals who move earlier (with ranks closest to 1 in Figure 6) are more likely to follow the prescriptions of the algorithm. Conversely, impulsive participants (who repeatedly click on the same action even though it has no practical effect in the game) are less likely to adhere to it (see SI4 for details).

Finally, we explore the prevalence and underlying factors of sniping behavior, and show that younger players are less likely to converge in C4 but more likely to do so through sniping (see SI5 for details). Then, we show that convergence increases when the first player to move does so faster, and when the individuals in the network are more efficient at (endogenously) adopting different roles, leader v. sniper (see SI6 for details).

2 Discussion

This experiment uncovers a developmental progression in children's ability to coordinate actions in groups. Kindergarteners (ages 5-6) show lower consensus-building proficiency compared to first and second graders (ages 6-8), indicating that cognitive mechanisms crucial to coordination—such as impulse control, sustained attention, and processing social information—sharpen notably between these ages. This likely stems from improvements in executive functions, working memory, and abstract thinking (Diamond, 2013). The task also demonstrate a fine sensitivity to developmental nuances within this narrow age range, as our previous research in strategic behavior among young children has shown (Brocas and Carrillo, 2020, 2021).

An interesting empirical aspect is the observed "patience heterogeneity"—variability in speed of the initial decision. Groups with both fast and slow decision-makers reach consensus more effectively, with quick responses providing direction and slower responses locking consensus. This balance of speeds may enhance decision-making in real-world group scenarios, where diversity in approach allows adaptation to changing information. Patience heterogeneity aligns with evolutionary theory, which holds that diversity in traits, like decision-making speed, can benefit group survival by fostering adaptability (Simons, 2011). In social contexts, variation prevents premature convergence on suboptimal choices and supports resilience. This diversity reflects the idea from evolutionary game theory that variation in behavior benefits groups in cooperative settings (McNamara et al., 2004).

The study also shows that pure imitation does not guarantee coordination success. Children who adjust choices based on social cues, which we label as "choice frictions", coordinate better and faster, indicating that they use more complex mechanisms than previously assumed. This aligns with theories of social cognition, which argue that children are not passive imitators but active participants in social learning (Tomasello, 2009). They employ adaptive strategies based on peers' behaviors, reflecting early-developed skills like shared intentionality and theory of mind (Wellman et al., 2001), and problem-solving skills that are fine-tuned in social contexts (Gopnik and Wellman, 2012).

Network complexity, in the form of fewer neighbors or more options, makes consensus significantly more challenging. This finding aligns with research suggesting that network structure and visibility critically influence collective decision-making. For instance, Kearns et al. (2006) demonstrates that in networks with limited connectivity, individuals struggle to reach consensus, and increasing visibility of others' choices or reducing options can simplify coordination by clarifying social cues. At the same time, our participants exhibit an impressive collective ability to *reduce complexity*, even when full coordination is not achieved. In fact, when we consider networks with 4 choices where consensus is not reached, participants manage to coordinate on just 2 choices in 83% of cases and five out of six participants successfully coordinate at some stage in 35% of cases. Non-convergence is often due to a single stubborn player who blocks progress towards a complete alignment.

Simple behavioral algorithms predict children's choices in broad strokes but miss critical nuances, particularly regarding decision speed and friction, underscoring the flexibility of children's heuristic strategies shaped by social cues. This aligns with theories of heuristic decision-making, where children make good enough choices suited to their cognitive limits and environmental constraints (Simon, 1955; Todd and Gigerenzer, 2000; Kahneman, 2011). Ecological rationality further asserts that these strategies are adapted to specific contexts, allowing humans, including young children, to outperform rigid models in real-life settings by relying on simple, context-sensitive strategies (Gopnik, 2012).

More broadly, our findings suggest that integrating structured group tasks that promote coordination and peer adaptation into early education may enhance the cooperative skills, social cognition, and iterative problem-solving ability of children. Practical applications include designing collaborative team-building activities, early interventions for impulsive children, and educational tools that align with developmental stages in social cognition. Future research could build on this by exploring diverse socio-economic groups and cultural differences or incorporating more complex tasks, such as fairness and resource allocation, to deepen our understanding of developmental coordination skills.

3 Methods

The study was conducted with the University of Southern California IRB approval UP-12-00528. We distributed consent forms to parents through the school administration with an opt-out option. No student (or parent) asked to not participate. In the end, only absent students and those who could not form a group of six within their grade did not take part in this experiment.

We conducted 10 sessions at LILA with 12 or 18 participants each (2 or 3 networks). We brought a portable lab to each classroom, gave PC tablets to every student, and connected the tablets to each other and to the portable server in a closed circuit with a wireless router. The experiment was programmed in oTree (Chen et al., 2016). Since we needed groups of six players, we sometimes mixed students from two different classes but always from the same grade. When multiples of six were not possible, the extra student(s) were randomly occupied on a different task. We used cardboard separations to preserve anonymity and made sure that students seated next to each other always belonged to different networks.

The timing of the experiment was as follows. First, in order to elicit their interest, we showed a sample of the toys the students would be playing for. These include 20 to 25 pre-screened, currently fashionable small toys such as bouncy balls, pop-up bracelets, scented pens, slime, emoji keychains, etc. We then read the instructions aloud with the support of a powerpoint presentation (as described in SI1). After, students played one practice round, where they could raise their hand and privately ask clarification questions. We then conducted the 16 rounds of the "network" task with 2 and 4 colors (C2 and C4) and with 2 and 3 neighbors (N2 and N3), in the following counterbalanced blocks of two: C2N3 C2N3 C2N2 C2N2 C4N3 C4N3 C4N2 C4N2 C2N2 C2N2 C2N3 C2N3 C4N2 C4N2 C4N3 C4N3, with a brief stretching pause halfway through the experiment (after the eighth counting round). Finally, they learned the number of toys they had won. We accompanied them to another classroom where we had setup the "toy shop" and they selected their favorite items.

For payments, participants obtained 2 points for each network convergence and 1 point for non-convergence. The conversion rate was one toy for each 10 points, rounded to the next ten (for example, 12 points = 2 toys). The procedure entailed incentives for performance but, at the same time, relatively small variance, which ensured that every participant was satisfied with their final outcome. Among our participants, 84% obtained between 6 and 8 toys.

Finally, we employed 16 different colors so that no four consecutive rounds were played with the same colors. This ensured that choices were unlikely to be driven by a "favorite" color. It also prevented the possibility that participants who converged in a round decided all to start the next round on the same color.

References

Isabelle Brocas and Juan D Carrillo. Iterative dominance in young children: Experimental evidence in simple two-person games. *Journal of Economic Behavior & Organization*, 179:623–637, 2020.

Isabelle Brocas and Juan D Carrillo. Young children use commodities as an indirect medium of exchange. *Games and Economic Behavior*, 125:48–61, 2021.

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.

Pedro Dal Bó, Guillaume R Fréchette, and Jeongbin Kim. The determinants of efficient behavior in coordination games. *Games and Economic Behavior*, 130:352–368, 2021.

Adele Diamond. Executive functions. *Annual review of psychology*, 64(1):135–168, 2013.

Alison Gopnik. Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102):1623–1627, 2012.

Alison Gopnik and Henry M Wellman. Reconstructing constructivism: causal models, bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, 138(6): 1085, 2012.

Sanjeev Goyal. *Connections: an introduction to the economics of networks*. Princeton University Press, 2012.

Matthew O. Jackson. *Social and economic networks*. Princeton University Press Princeton, 2008.

Matthew O Jackson, Suraj Malladi, and David McAdams. Learning through the grapevine and the impact of the breadth and depth of social networks. *Proceedings of the National Academy of Sciences*, 119(34):e2205549119, 2022.

Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences*, 107(34):14978–14982, 2010.

Daniel Kahneman. Thinking, fast and slow. *Farrar, Straus and Giroux*, 2011.

Michael Kearns. Experiments in social computation. *Communications of the ACM*, 55 (10):56–67, 2012.

Michael Kearns, Siddharth Suri, and Nick Montfort. An experimental study of the coloring problem on human subject networks. *science*, 313(5788):824–827, 2006.

Richard D McKelvey and Thomas R Palfrey. Playing in the dark: Information, learning, and coordination in repeated games. Mimeo, California Institute of Technology, 2001.

John M McNamara, Zoltan Barta, and Alasdair I Houston. Variation in behaviour promotes cooperation in the prisoner's dilemma game. *Nature*, 428(6984):745–748, 2004.

Herbert A Simon. A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118, 1955.

Andrew M Simons. Modes of response to environmental change and the elusive empirical evidence for bet hedging. *Proceedings of the Royal Society B: Biological Sciences*, 278 (1712):1601–1609, 2011.

Peter M Todd and Gerd Gigerenzer. Précis of simple heuristics that make us smart. *Behavioral and brain sciences*, 23(5):727–741, 2000.

Michael Tomasello. *The cultural origins of human cognition.* Harvard university press, 2009.

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.

Supplementary Information

SI1. Instructions

Hi everyone. My name is Juan. Today, we are going to play several games and you are going to win points. At the end, you are going to exchange the points for toys in our toy shop. You will all get many toys, but the better you do, the more points you earn, and the more points you earn the more toys you get.

NETWORK GAME

The computer is going to form several groups of six, but you will not know who is in your group and it is not the point to find out [**SLIDE 2**].

Each of you will be connected to some people in your group but not to others. When you are connected to someone, it means that you can see their choices. For example, here Bob sees the choices of John and James but not the other kids in his group. John sees the choices of Bob and Ann but not James or the others, and so on [**SLIDES 3 and 4**].

In some games, each of you will see the choices of two other people in your group. In some other games, you will see the choices of 3 other people [**SLIDES 5 and 6**]. This is an example of what the people in a group will see [**SLIDE 7**].

What you have to do in this game is very simple! You have to choose a color. That's it. If you all choose the same color at the same time, you all win two points! Otherwise, you will only win one point. So, it's not about choosing your favorite color but about choosing the same color as others. The good news is that you can change the color as many times as you want. There is only one rule: you cannot talk.

Let me give you an example. This is what you will see in your screen [**SLIDE 8**]. This is you (point), and these are the people around you (point). You are always at the center of your own group. In this case how many other people can you see? Two. Very good. Remember there are other people in your group. It's just that you can't see them. This is the screen when there are three other people in your group [**SLIDE 9**].

Let's go back to the case where you see two people. Now, this is where you choose a color [**SLIDE 10**]. In this case, everyone chooses between two colors, BLUE and RED, but in different rounds you will see different colors. As soon as you choose a color, it will appear here. And the people connected to you will be able to see it.

In this example, you chose red [**SLIDE 11**]. Here you see that one person chose blue and the other red [**SLIDE 12**]. Then, you can change your color as many times as you want. It is totally up to you. Remember the goal is to have everyone in your group of six choose the same color. This means not only the people you can see but also the others you do not see.

So, for example, imagine you are Bob. In this case, [**SLIDE 13**] everyone you see is choosing blue (including yourself), but you are still not winning two points because these guys are choosing red. While here [**SLIDE 14**], everyone in your group is choosing blue, so you all get two points. Does that make sense to everyone?

Now two more things. Every time you play, you will see a clock at the top of the screen [**SLIDE 15**]. This tells you the amount of time left to play. You start with 30 seconds and the clock runs backwards (29, 28, 27). Everyone can change their action as many times as they want. When everyone is in the same color, the clock stops, and you get two points. As long as everyone is not in the same color, the clock keeps running. If it hits 0 and not everyone is in the same color, you will only get one point.

You are going to play many rounds. Sometimes, you will see the choices of two people and sometimes you will see the choices of three people. Also, sometimes you will choose between two colors as in this example, and sometimes you will choose between four colors as in this example [**SLIDE 16**]. However, the rules are always the same: everyone has to be in the same color, it does not matter which one.

At the end of each round, you will know how many people chose each color. You will see a screen like this [**SLIDE 17**]. In this example, did you get two points? Why? How about this one? [**SLIDE 18**]. After you see if you won that round, you press the green button and move to the next round to win more points. Are you ready to play?

You are going to play many times, so you are going to have many chances to win points. Before we start, we are going to play a pretend round. This is only to make sure you understand the game. This round does not count for real so feel free to try different things. If you don't understand what's going on, raise your hand and we'll be happy to come and help you. Any questions?

[After the practice round] We are going to start the real game

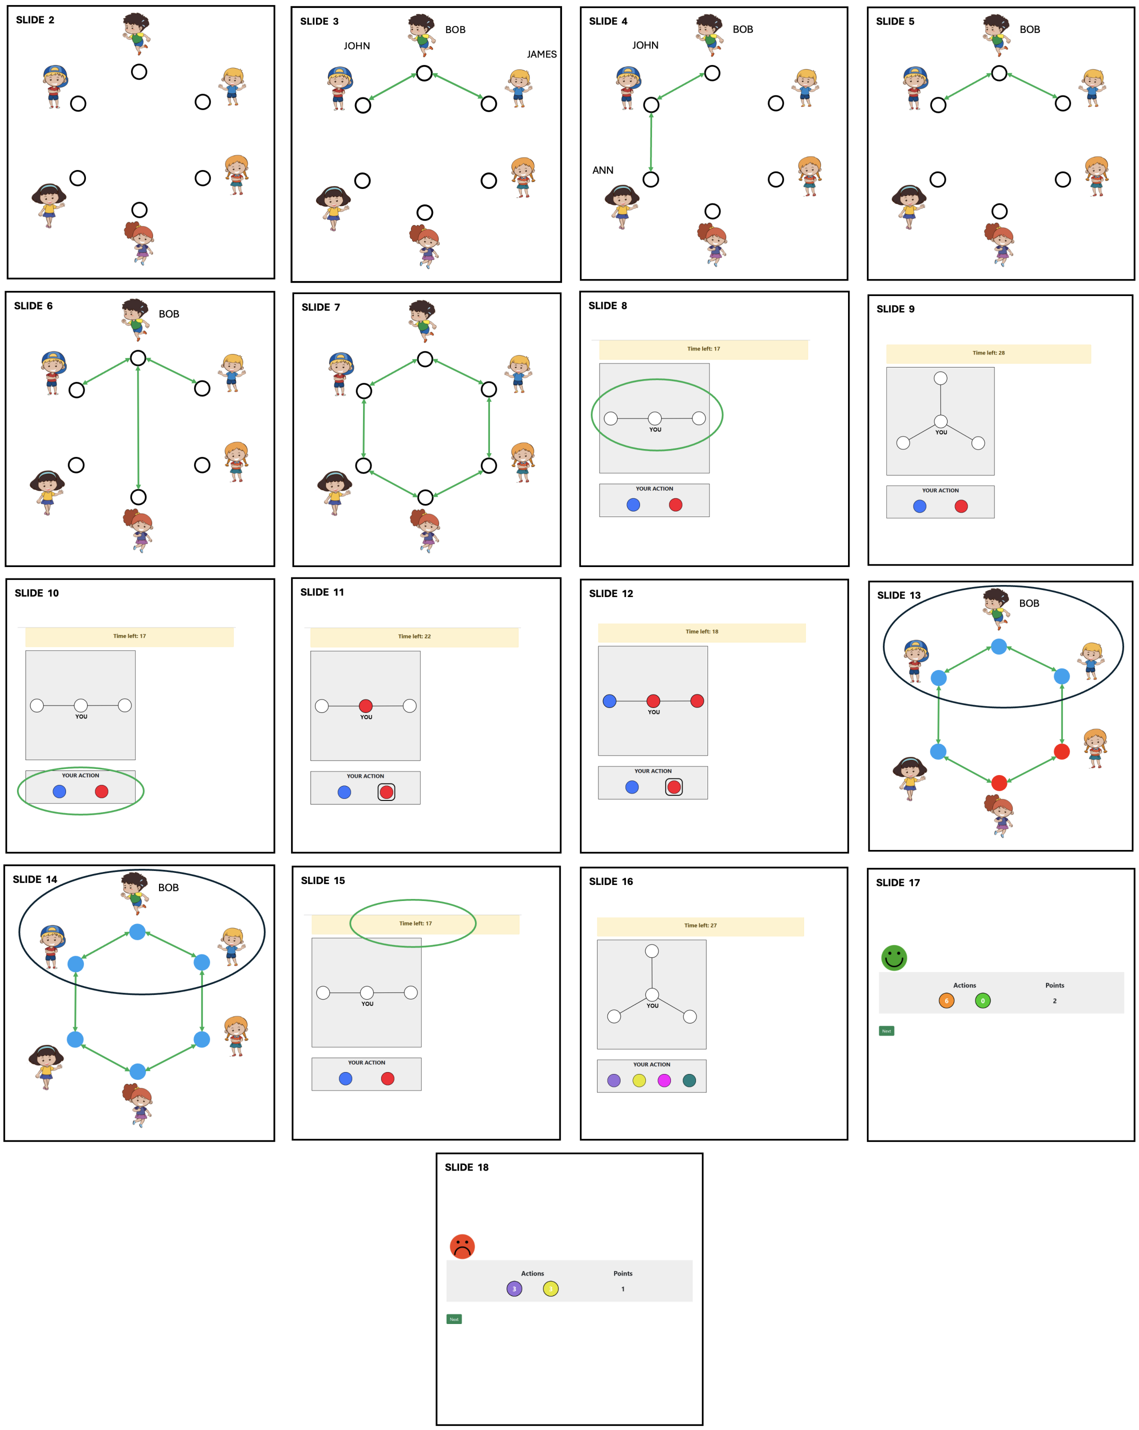

Figure SI1: Slides projected on screen for instructions

SI2. Probit regression of likelihood of network convergence

Table SI1 presents a Probit regression of network convergence probability with dummy variables for grade, number of males in the network, number of links, number of colors, and first v. second half of the experiment.

	Prob. Conv.
(Intercept)	1.840***
	(0.455)
Grade **K**	-1.190***
	(0.261)
Grade **2**	0.129
	(0.246)
# males	-0.040
	(0.126)
Second-Half	0.378°
	(0.203)
C4	-0.297*
	(0.148)
N2	-0.621***
	(0.150)
AIC	324.5
Num. obs.	400

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

Table SI1: Probit regression of probability of network convergence (standard errors are clustered at the session level). Convergence is significantly less frequent with more colors (C4), with fewer neighbors (N2), and in **K** (compared to **1** and **2**). Convergence is also marginally more frequent towards the end of the experiment. Gender composition of the network has no effect.

SI3. Choices conditional on convergence under frictions

Figure SI2 compares the empirical distribution of number of choices conditional on convergence with the distribution under 10,000 simulations of algorithms $\mathbf{AL}_{\hat{\mathbf{p}}}$ (left) and $\mathbf{AL}_{\hat{\mathbf{p}},\hat{\mathbf{q}}}$ (right), that is, the algorithm with frictions or with frictions and heterogenous patience.

Convergence is much slower under $\mathbf{AL}_{\hat{\mathbf{p}}}$ than under $\mathbf{AL}$ (see Figure 3) and therefore also much slower than that of our participants. The introduction of choice frictions decreases the number of instances where the network gets stuck in configurations where two subset of players coordinate in two different colors. However, these small "mistakes" increase the expected number of choices it takes to achieve consensus.

By contrast, sequentiality has a very large positive effect on speed of convergence. $\mathbf{AL}_{\hat{\mathbf{p}},\hat{\mathbf{q}}}$ matches the empirical behavior of our participants much better than the other algorithms. However, it is still not perfect. Indeed, the fraction of both 'immediate' convergence (6 choices) and

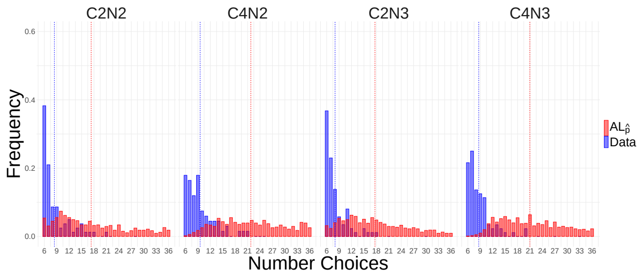

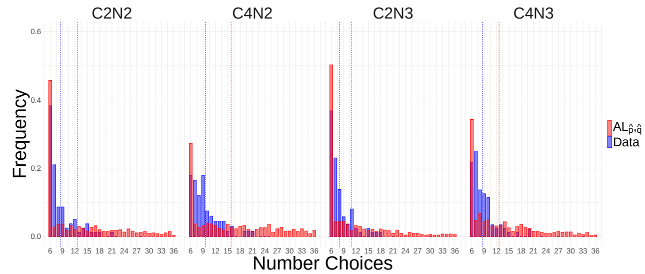

Figure SI2: Distribution of choices in the population and in the algorithms $\mathbf{AL}_{\hat{\mathbf{p}}}$ (left) and $\mathbf{AL}_{\hat{\mathbf{p}},\hat{\mathbf{q}}}$ (right) [averages are reported with vertical lines].

'delayed' convergence (12 choices or more) is larger than in the data while the fraction of 'intermediate' convergence (between 7 and 11 approximately) is smaller.

SI4. OLS regression of dynamic imitation

We investigate in Table SI2 the determinants of playing according to **AL**, p_i, as a function of the participant's age (in months), gender, and average rank of their first choice (q_i) which captures their leadership skills and ability to move outcomes in their direction. We also include extra clicks, which captures the impulsivity of the individual. Formally, it is the number of instances where the participant clicks in the same color twice or more. Those decisions have no effect on outcomes (they are not even observed by others) and only denote willingness to impose their choices.

	p_i
(Intercept)	0.914***
	(0.071)
Age	-0.001
	(0.001)
Male	-0.014
	(0.016)
Rank	-0.0174*
	(0.007)
ExtraClicks	-0.211***
	(0.017)
Adj. R^2	0.539
Num. obs.	148

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\circ}p < 0.1$

Table SI2: OLS regression of p_i.

Leaders (early movers) are more likely to follow the algorithm, thus showing a better understanding of the game. Conversely, impulsive participants are less likely to follow the algorithm, as they want to impose their choices on others. Interestingly, age has no effect on p_i, mostly because kindergartners—who are more likely to deviate from the algorithm—are also the more impulsive players.

SI5. Sniping behavior

We have previously noted some differences across grades in the sequentiality of initial choices (Figure 6). Nevertheless, conditional on convergence, the percentage of rounds where it is achieved in exactly 6 moves is similar across grades: 29% in **K**, 29% in **1** and 26% in **2**. Perhaps more interestingly, the percentage of rounds where convergence is reached in more than 6 moves but immediately after the first move of one participant–what we call sniping–is high and also similar across grades: 30% in **K**, 28% in **1** and 27% in **2**.

Table SI3 reports the proportion of sniped networks among networks that converged in more than 6 moves, separated by grade and number colors to choose from.

	C2	C4
K	0.37	0.52
1	0.43	0.38
2	0.48	0.27

Table SI3: Sniping by grade and complexity

Differences in sniping across grades is mostly related to the complexity of the situation. Indeed, conditional on convergence, children in **K** show a higher tendency than children in **2** to snipe in the more complex case with 4 colors (C4, test of differences, p = 0.053), even if they are less likely to reach overall convergence. This suggests that older children are better able to think and dynamically adapt to complex situations.

Finally, Figure SI3 illustrates for each of the 25 networks in the experiment, the number of rounds (out of 16) where convergence was achieved through sniping behavior. On average, sniping occurred in 22.8% of rounds but we notice very different behavior across networks, with a minimum of 0 and a maximum of 9 sniping in a given network.

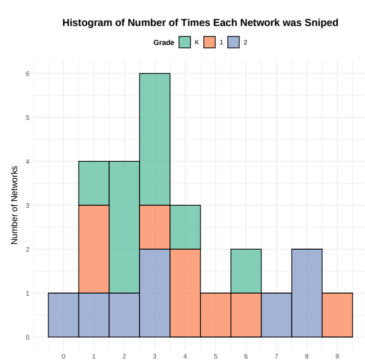

Figure SI3: Rounds of convergence through sniping behavior

SI6. Timing of the first choice

For each round, we calculate a measure called "Delay" which represents the time elapsed between the beginning of the game and the first move of the participant who moves first. We then include this measure in the original Probit regression model of probability of convergence (Table SI1), and present the results in Table SI4. As the regression shows, having one player in the network who moves fast (lower value of "Delay") helps convergence.

	Prob. Conv.
(Intercept)	2.790***
	(0.472)
Grade **K**	-1.213***
	(0.206)
Grade **2**	0.086
	(0.242)
# males	-0.060
	(0.080)
Second-Half	0.314°
	(0.16673)
Delay	-0.303**
	(0.106)
C4	-0.303°
	(0.164)
N2	-0.689***
	(0.17112)
AIC	318.0
Num. obs.	400

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

Table SI4: Probit regression of probability of network convergence, including 'Delay' as a predictor.

We also calculate the variance of the rank for each individual and then average these variances to obtain a measure of "spread" for each network. The spread differs significantly between networks in **K** and networks in **2** (1.77 v. 1.29, p = 0.041) and it is marginally correlated with the probability of convergence of the network (Pearson Correlation Coefficient PCC = -0.38, p = 0.063). This finding suggests that better self-sorting of players within the network which results in less ambiguity in roles (and thus, a smaller variance) enhances the likelihood of consensus, and is a skill that develops over time.