Cultural Learners or Game Theorists?

Baimel, Adam*
Department of Psychology, Health and Professional Development
Oxford Brookes University
*abaimel@brookes.ac.uk

Juda, Myriam
Department of Psychology
Simon Fraser University

Birch, Susan
Department of Psychology
University of British Columbia

Henrich, Joseph*
Department of Human Evolutionary Biology
Harvard University
*henrich@fas.harvard.edu

June 5, 2020

Abstract

Theoretical efforts have sought to identify the key selection pressures that drove the expansion of brains in the human lineage and shaped the emergence of the central features of human cognition. Focusing on two theories, each capable of accounting for the rapid brain expansion in our genus, we present the results of an experiment designed to favor either the Machiavellian Intelligence or the Cumulative Cultural Brain Hypothesis. Children (aged 3-7 years) observed an interaction that provided them with behavioral information that could either be used to outmaneuver a partner in subsequent interactions or for cultural learning. The results show that, even after four rounds of repeated interaction and sometimes lower payoffs, children continued to rely on copying the observed behavior instead of harnessing the same information to strategically extract payoffs (stickers) from their partners. Analyses reveal that superior mentalizing abilities are associated with more targeted cultural learning—the selective copying of fewer irrelevant actions. By contrast, superior cognitive abilities are associated with increased imitation of irrelevant actions. Neither mentalizing capacities nor general cognitive abilities predicts children's ability to strategically use social information maximize payoffs. These results provide developmental evidence favoring the Cumulative Cultural Brain Hypothesis over the Machiavellian Intelligence Hypothesis.

1 Introduction

Brain size in the human lineage expanded roughly 3-fold in less than three million years (Bailey and Geary, 2009; Schoenemann, 2006). Accompanying this rapid expansion were energetically costly modifications to the female pelvis, permitting the birthing of big-headed infants, and a reduction in the length of gestation, making human births relatively premature from the perspective of other primates (Boyd and Silk, 2012). Because of the magnitude and fitness costs of these rapid genetic changes, current theorizing focuses on identifying the 'autocatalytic' or 'runaway' evolutionary processes responsible. Focusing on the two primary hypotheses capable of producing the requisite auto-catalytic evolutionary dynamics, the Machiavellian Intelligence Hypothesis (McNally et al., 2012) and the Cumulative Cultural Brain Hypothesis (Muthukrishna et al., 2017), we tested the psychological implications of these two theories among 3- to 7-year-old children. Our results support the Cultural Brain Hypothesis (Whiten and van Schaik, 2007; Muthukrishna and Henrich, 2016; Henrich, 2015; Laland, 2017; Boyd, 2018; Gavrilets and Vose, 2006), which proposes that cultural evolution generated an ever-increasing body of adaptive learned information that created selection pressures for bigger brains that were better equipped for cultural transmission. Meanwhile, we find little support for the Machiavellian Intelligence Hypothesis, which when applied to humans proposes that brain expansion in our species, and its associated psychological capacities, were driven by an arms race in strategic social reasoning (Byrne and Whiten, 1988).[1]

2 Theoretical Framework

To explain the origins of big human brains, along with their sophisticated cognitive abilities and associated life history characteristics, a pair of distinct theoretical approaches have emerged, both of which fall under the broader rubric of the Social Brain Hypothesis (Humphrey, 1976). Both approaches aim to explain the variation in the brain sizes and cognitive abilities among primates specifically, as well as having implications for other mammals and birds, and crucially both can be extended to supply the auto-catalytic dynamics necessary to explain the rapid expansion of brains and cognitive abilities in the human lineage. The leading version of the Social Brain Hypothesis, which we label the Machiavellian Intelligence Hypothesis (for clarity), proposes that big brains and sophisticated cognitive abilities result from the selection pressures for strategic thinking applied to managing relationships in larger, or more intensely social, groups (Dunbar, 1998). Bigger brained individuals, in this view, are better able to track and strategically deploy information, including about third-parties, regarding the strategies or choices of others. These psychological abilities allow them to better trick, manipulate, and deceive others, as well as to sustain longer-term alliances. In this view, the complexity of social life is driven by some external pressure, like predation (Dunbar and Shultz, 2007). For the human case, the required runaway dynamics arise from an ever-escalating social competition in strategic thinking in which selection favors competitors who can reason one step farther than others–as in the backward or forward induction required of agents in standard game theory (Trivers, 1971). One potential reason why this runaway social competition occurred in the human lineages, but not in other species, may be due to intergroup competition, where members of the same species became potentially dangerous predators (Bailey and Geary, 2009). Central to this reasoning is the ability to represent others' mental states (their beliefs and preferences) — mentalizing — and the use of this faculty to exploit or manipulate conspecifics.

By contrast, the Cultural Brain Hypothesis proposes that combinations of individual and social learning generate a pool of adaptive non-genetic information, which may take many forms including foraging skills, food preferences, tool-using techniques, communicative signals, ally preferences, or socially-strategic tactics (Muthukrishna and Henrich, 2016; Whiten and van Schaik, 2007; Gavrilets and Vose, 2006; Reader et al., 2011; Henrich, 2015). The emergence of this pool of adaptive information creates selection pressures favoring larger brains that are better able to acquire, store, organize and re-transmit this body of

[1] In this context, the Cumulative Cultural Brain Hypothesis converges with the Embodied Capital Hypothesis (?). We focus on the former because those working under this rubric have explicitly analyzed the role of cumulative cultural evolution. Note that while Gavrilets and Vose (2006) is presented under the rubric of the Machiavellian Intelligence Hypothesis, we feel the actual dynamic processes embedded in the model actually capture a version of the Cumulative Cultural Brain Hypothesis.

fitness-enhancing information. Applied to the human lineage, this becomes the Cumulative Cultural Brain Hypothesis (Muthukrishna et al., 2017), as human ancestors crossed a theoretical threshold in which the body of adaptive know-how and preferences began to accumulate over generations. This further increased the selection pressure for larger brains that were better able to acquire, store, organize and re-transmit this information. The better at cultural learning human ancestors became, the more rapidly cultural evolution could accumulate large pools of adaptive know-how, and the greater the selection pressures became on genes for building larger brains that were better able to copy the motivations, strategies, and beliefs from the minds of others in their social milieu. Here, mentalizing evolved in order to improve cultural learning, to better extract adaptive information, in the form of knowledge, motivations, beliefs, intentions and strategies, from the minds of others. By this account, better mentalizers should be more selective in their learning, and more attuned to copying, for example, intentional as compared to accidental or incidental behaviors. Indeed, while non-selective imitation ('overimitation') can be a good cultural learning strategy that ensures that all key behaviours are copied (e.g., Chudek et al., 2016) - a learner is even more adaptive if they can selectively infer the necessary from the irrelevant. Superior mentalizing may equip learners with the capacity to be more selective, in part as it helps learners figure out what is necessary and what is not (e.g., Brosseau-Liard et al., 2015).

While often making many similar predictions about brain size, sociality and breeding patterns, these two hypotheses make very different predictions about the nature of human psychology, and about how humans ought to deal with novel social dilemmas. And, although both approaches do emphasize the importance of our species' mentalizing or 'mind-reading' abilities, they propose that these mentalizing abilities will be put into primary service in quite different ways. Specifically, the Machiavellian Intelligence Hypothesis holds that humans readily develop the ability and motivation to out-smart others, by out-mentalizing them, especially in competition for desirable resources. By contrast, the Cultural Brain Hypothesis proposes that mentalizing abilities will first develop for, and be deployed most commonly, in the service of cultural learning, not primarily for strategically out-witting others. The empirical question is: do children *initially* put their mental abilities to work in learning from others or to exploiting their opponents for personal gain?

We allowed these hypotheses to face-off in a simple experimental design administered to 280 children (51% female) ranging in age from 3 to 7 years old in Vancouver, Canada. Our approach was two pronged. First, we observed children's decisions in a resource distribution game in which they could win stickers — a valued resource. Importantly, the particulars of the game and its conditions were designed such that children's decisions could reflect either the outcomes of imitative cultural learning or strategic social reasoning. Second, we assessed children's capacities for mentalizing in three ways: using (1) a classic false-belief task (N = 276); Wimmer and Perner, 1983), (2) a storybook instrument (N = 100; Blijd-Hoogewys et al., 2008), and (3) parental reports (N = 150; Tahiroglu et al., 2014). In a subset of the sample (N = 118), we also measured children's general cognitive abilities (McGrew and Woodcock, 2001). The Cumulative Cultural Brain Hypothesis (CCBH) predicts that mentalizing should be associated with the cultural acquisition of intentional actions, preferences or strategies while the Machiavellian Intelligence Hypothesis (MIH) predicts that mentalizing will be associated with behavior that maximizes payoffs by taking advantage of social information in a zero-sum interaction. Our measure of general cognitive abilities provides a valuable control.

3 Methods

3.1 Participants

In the greater Vancouver area (Canada), 280 children (136 Males; 144 Females; 2 with sex unreported) aged from 2.91 to 6.93 years (M = 4.48, SD = 0.94) were recruited to participate in this study from 22 daycare centres (N = 201), a local science museum (N = 55), and the child subject pool at the University of British Columbia (N = 24). The family income of participating children ranged from 20,000 CAD to 220,000 CAD at our different sampling sites around the city (Median = 100,000 CAD). The median family income in the greater Vancouver area in 2015 was around 72,000 CAD (Statistics Canada, 2017). Most of the children had one ($N = 234$) or two siblings ($N = 27$).[2]

3.2 Materials and Procedures

Participants completed a battery of assessments: (1) the sticker bargaining game, (2) a false-belief test (Wimmer and Perner, 1983), (3) the Theory of Mind Storybooks (Blijd-Hoogewys et al., 2008), and (4) a test of general cognitive abilities (McGrew and Woodcock, 2001). Participating children recruited from daycares completed the assessments in a round-robin style with different research assistants making one to three visits per daycare centre. Parents or guardians of participating children provided demographic information and filled out an observation instrument on their child's mentalizing capacities (Tahiroglu et al., 2014). The parent/guardian questionnaire was completed either as the child participated in the other tasks (at the Science Museum and in-lab) or was sent home with participating children at daycare centers and collected at a later time. As some of the assessments (e.g., the ToM Storybooks and the cognitive ability tests) required lengthy and/or returning sessions with the children, we do not have complete data for all participating children. This is primarily due to children's absence on returning visits to the daycare centres and parents/guardians not returning the take-home questionnaire.

3.2.1 The Sticker Bargaining Game

The sticker game involved two active players, a proposer and a responder. The proposer had to decide how to allocate four stickers between two baskets; the responder then had to pick which basket they wanted, which left the proposer with the remaining basket. When responders are assumed to prefer more stickers to fewer stickers, game theory predicts that sticker-maximizing proposers will make a 2-2 division between the baskets. Procedurally, children first watched a live demonstration of the game in which two adult models interacted for three rounds, ostensibly as an instructional aid. A third adult experimenter laid out the the stickers in front of the proposer at the beginning of each round. In all demonstrations, the proposer initiated each round by performing five actions: announcing that they had four stickers (”I have four stickers”), counting them out-loud (”One, two, three, four”), tapping on each sticker twice with a finger, shuffling them around into a different order, and realigning them into a straight line. The proposer then allocated the stickers to the baskets, and asked the responder, ”Which do I get to keep?”—prompting the responder's decision. Importantly, the demonstration varied (1) the proposer's allocations, (2) the responder's preferences and (3) whether or not the participants could actually observe any of the allocations (the sticker payoffs).

Table 1: Sticker game decision matrix and by condition predictions

Conditions	Proposer's actions	Responder's actions	Cultural learners' divisions	Machiavellian divisions
CONTROL	Unseen	Unseen	Baseline	Baseline
EVEN	Allocates 2-2	Picks 2 sticker cup	Even (2-2)	Baseline
NICE	Allocates 3-1	Picks 1 sticker cup	Uneven (3-1)	Uneven split (4-0, 3-1)
SELFISH	Allocates 3-1	Picks 3 sticker cup	Uneven (3-1)	Even split (2-2)

Children were randomly assigned to one of four conditions, labeled CONTROL, EVEN, SELFISH, NICE (Table 1), and played for four rounds in the role of proposer against the *same person they had just observed in the responder role in demonstration*:

1. CONTROL condition: The proposer in the live demonstration was given the four stickers, they then performed the five actions described above. After stating that they were to put the stickers in the baskets, and *before* asking which they got to keep, the experimenter in the demonstration placed an

[2] This study was approved by the University of British Columbia Behavioural Research Ethics Board. Written informed consent was obtained by the parents of participating children in addition to children's verbal assent to participate, and children were given the option to withdraw at any point during the study.

occluding box over the baskets that had one side cut out such that the adult proposer and the responder could see the baskets and the placement of the stickers but the participant could not. The responder selected one of the baskets from within the box, so participants also didn't get to see how many stickers either player retrieved. The box was removed and then replaced prior to the proposer's allocations in the following rounds. This treatment provides a baseline for how children will allocate stickers at test in the absence of information about the proposer's or the responder's preferences. In the other conditions, children had full view of all decisions and outcomes.

2. EVEN condition: The proposer split the stickers evenly, leaving the responder with no choice but to return two stickers and keep two for themselves. Since this 2-2 split provides no additional information as to the responder's preferences, Machiavellians who adjust their strategies in light of the responders' behaviours should match the baseline condition. Cultural learners, however, should tend to copy the model and split the stickers evenly.

3. NICE condition: The proposer distributed the stickers unevenly with three stickers in one basket and one in the other (the order of which was counterbalanced across participants). The responder was then 'nice' and always picked the basket with only one sticker. Here, both good cultural learners and Machiavellians should allocate unevenly, with cultural learners copying the model and Machiavellians adjusting to best exploit their opponent.

4. SELFISH condition: The proposer allocated the stickers unevenly but now the responder was 'selfish' and always took the basket with three stickers. Here, cultural learners should copy the uneven allocation tendencies of the model, while good Machiavellians should recognize the sticker-maximizing tendencies of their opponent and pick an even 2-2 allocation. Table 1 summarizes these treatments and predictions.

After the demonstrations, children were told that it was their turn to play and were placed into the role of the proposer, playing against the same responder that they had just seen in the demonstration. When the child first took the place of the proposer, they were asked whether they liked stickers (in general). Four participants said they did not. At the outset of each round of the testing phase, children were again asked whether or not they liked the specific stickers that had been laid out in front of them by the experimenter. To these queries, 17 children indicated that they did not like the *specific* stickers on that round[3]. Participants played the game for four rounds with the responder playing the same strategy that the child saw in the demonstration. In the CONTROL condition, however, the responder's behaviors at test were dependent on the child's allocations but were pre-determined and counterbalanced across the sample. If the child distributed evenly, the responder always chose either the left or right basket. If allocations were uneven, the responder was randomly assigned apriori to be either nice ($N = 12$) or selfish ($N = 14$). This was also the case in response to uneven allocations in the EVEN condition (Nice, $N = 10$; Selfish, $N = 4$). After each round, the stickers that the children obtained in the game were placed in a small plastic bag for them to take home.

3.2.2 Measures of Mentalizing

We measured children's mentalizing abilities in three ways:

1. The "Sally Anne Task" (Wimmer and Perner, 1983): in this task children were presented with a live demonstration of a false belief test using hand puppets in a 'change of location paradigm'. The test involved two characters, "Sally" who had a basket and "Anne" who had a box. The test began with Sally placing a toy in her basket. Sally then left the scene to "go play outside". While Sally was away and could not see what takes place, Anne took the toy out of Sally's basket to put it into her box. Sally then returned and the child participant is asked three questions: "where is the toy now?", one memory question "where was the toy at the beginning?" and the focal belief question: "where will

[3] Analyses reveal no robust relationship between children's report of liking of the stickers and their choices, so we did not exclude children based on their sticker preferences

Sally look for her toy?". Children are said to pass the test when they reply that Sally will look for the toy inside her own basket (1 = Pass; 0 = Fail)—that is where she had left it (and *not* where the child knows it to currently be). To insure that participants understood where Sally had actually placed the toy, and where it was in reality after it was moved by Anne, the experimenter corrected the participant if they had responded to either of these two questions incorrectly before asking the focal false-belief item. Incorrect responses on the false belief item were not corrected.

2. Theory of Mind Storybooks (Blijd-Hoogewys et al., 2008): This instrument consists of six storybooks portraying a protagonist, Sam, who experiences various emotions, desires and thoughts in a series of brief stories about this character, his friends and his family. The storybooks, which were read aloud by an adult experimenter, consist of 34 tasks, with assessments of five components of mentalizing: (1) emotion recognition, (2) distinguishing between physical and mental entities, (3) understanding that seeing leads to knowing, (4) prediction of behaviors and emotions from desires, and (5) prediction of behaviors and emotions from beliefs. An overall "Theory of Mind" score is indexed by the sum-total of coded responses, ranging from 0-110 on the basis of a continuous scoring system. The task takes 40-50 minutes to complete (including a brief break). As an instrument of various aspects of mentalizing, these storybooks have been shown to have robust internal consistency, test-retest reliability, inter-rater reliability, construct validity and convergent validity.

3. The Children's Social Understanding Scale (CSUS-short form; Tahiroglu et al., 2014): parents of a subset of our sample also completed an 18-item parent-report questionnaire of their child's mentalizing capacities. The CSUS asks parents to reflect on their child's capacities for reasoning about mental states such as beliefs (e.g., "My child understands that telling lies can mislead other people"), knowledge (e.g., "My child uses words that express uncertainty"), perception (e.g., "My child thinks that you can still see an object even if you're looking in the opposite direction" (*reverse-coded)*), desires (e.g., "My child talks about what people like or want"), intentions (e.g., "My child talks about the difference between intentions and outcomes"), and emotions (e.g., My child talks about conflicting emotions"). The 18-item scale is reported to have good psychometric properties, and has been validated with regards to children aged 3 to 8 years of age.

3.2.3 General Cognitive Abilities

To assess children's general cognitive abilities, a subset of our sample completed the Brief Intellectual Ability test [BIA] (McGrew and Woodcock, 2001). The BIA was designed to assess cognitive abilities in children older than 2 years. An overall score is derived from the outcomes of three cognitive tests involving verbal comprehension, concept formation, and visual matching that assessed verbal skills, fluid reasoning, and processing speed. For our analyses, scores on the test were age-normalized using the scoring program provided by the test creators.

3.3 Sticker game response coding

Children's behaviors provided us with a rich set of data. We first coded children's sticker allocations and tracked their relative frequency across rounds and conditions to assess the extent of imitative cultural learning. Then, we recoded allocations such that they reflected payoff-maximizing-choices and estimated the contributions of mentalizing and general cognitive abilities to this strategic exploitation of the responder. Lastly, we counted if and how many of the proposer's seemingly-irrelevant behaviours—as seen in the demonstration (e.g., counting, tapping, shuffling)—the child reproduced on each test round to provide a measure of overimitation.

4 Results

We analyze our data in two steps. First, we consider how well the data fit the predictions arising from the Cumulative Cultural Brain Hypothesis by asking if, and how much, children tended to imitate the allocations and behaviours of their model/demonstrator. Second, we contrasted this analysis with how well children's behavior fit the predictions derived from the Machiavellian Intelligence Hypothesis.

4.1 Are children cultural learners in this zero-sum situation?

To assess the impact of our four treatments (t), we began by coding children's allocations into a binary variable, as either even splits (2/2, so $d_{i,j,s} = 1$) or uneven (i.e., 3 and 1 or 4 and 0, so $d_{i,j,s} = 0$). The variable i indexes the round, j indexes the individual, and s marks the sampling site. We modelled these decisions in a series of logistic regressions. To account for the non-independence of repeated responses across rounds and data collection in different sites, we adjusted all standard errors by clustering both within subjects and within sample sites (22 Daycares, Science Museum or in-lab). We estimate the regression equation (1) below. The coefficient on condition, C_t, captures the effects of our four treatments, using our control condition as the reference. The coefficient on round, R_i, reveals the average effect of personal experience per round of repeated play. β_t captures the effect of the interaction of treatment and round, which is crucial since we expect individual learning to have different effects in different treatments. The coefficient on children's ages, A_j, allows us to examine how children's inclination to offer even splits develops from ages 3 to 8 years in this population. M_j controls for the reported sex of our participants ($sex = 1$ is male).

$$logit[\Pr(d_{i,j,s} = 1)] = C_t condition + R_i round + A_j age + M_j sex + \beta_t * condition * round + constant \quad (1)$$

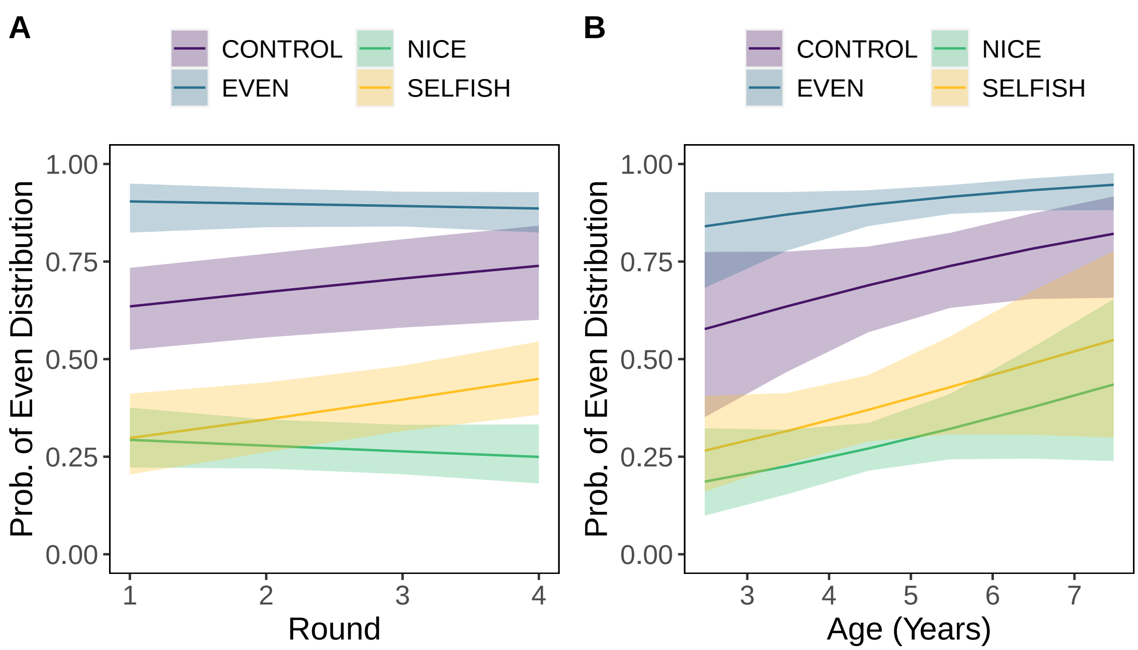

Figure 1: **Predicted probability of even distributions in each condition across the four rounds (left panel) and age (right panel)**. Predictions were generated from Model 4 in 2. The shaded regions show the 95% confidence intervals based on two-way clustering.

Table 2: Logistic regression models to predict uneven vs. even sticker allocations

	Sticker Allocations (0 = Uneven; 1 = Even)			
	Model 1 OR (95% CI)	Model 2 OR (95% CI)	Model 3 OR (95% CI)	Model 4 OR (95% CI)
Intercept	2.218***	1.968**	1.946***	1.735**
	(1.249, 3.940)	(1.124, 3.447)	(1.182, 3.202)	(1.092, 2.757)
Even Condition	3.691***	3.694***	3.882***	5.408***
	(2.006, 6.791)	(2.007, 6.801)	(2.129, 7.078)	(2.579, 11.337)
Nice Condition	0.171***	0.170***	0.168***	0.238***
	(0.090, 0.325)	(0.089, 0.325)	(0.089, 0.318)	(0.131, 0.433)
Selfish Condition	0.263***	0.262***	0.268***	0.244***
	(0.125, 0.555)	(0.123, 0.556)	(0.132, 0.547)	(0.115, 0.517)
Round (0 = Round 1)		1.085**	1.085**	1.176**
		(1.006, 1.171)	(1.004, 1.173)	(1.036, 1.335)
Age (Years, Centered)			1.273	1.275
			(0.940, 1.723)	(0.941, 1.727)
Sex (0 = Prop. of Males)			1.518*	1.523*
			(0.979, 2.355)	(0.980, 2.366)
Even Condition X Round				0.798*
				(0.610, 1.043)
Nice Condition X Round				0.790***
				(0.665, 0.938)
Selfish Condition X Round				1.058
				(0.828, 1.352)
Observations	1080			
Participants	273			
Sites	24			

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. The CONTROL condition (Intercept; controlling for other variables) is the reference category for condition effects. Round of the game was treated as a continuous variable, and thus condition by round interactions represent changes across the rounds in each condition. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

Figure 1 illustrates our key results and Table 2 provides greater detail. For each of our conditions, the left panel reveals the predicted probabilities of even allocations across the four rounds of play. The right panel shows the age trajectories in allocation strategies for each condition. The most striking result is the tendency of children to imitate the proposer they observed in the demonstration. Relative to the CONTROL condition, children who saw an even distribution were much more likely to distribute their stickers evenly (the blue line at the top of both plots). In Round 1, for example, the percentage of equal allocations increased from 62% in the CONTROL condition to 91% in the EVEN condition. By the last round, those who observed the model distribute the stickers evenly still remained 18 percentile points greater than the CONTROL condition. Similarly, when children saw a proposer divide the stickers unevenly in either the SELFISH or NICE conditions, they allocated their stickers much less evenly at test. In Round 1, the percentage of even allocations dropped to 28% in both the NICE and the SELFISH conditions. This is 34 percentile points below the frequency of equal splits observed in the CONTROL condition. Table 2 shows that, even holding participants' age, sex and round of play constant, those who observed an even split are 2.5 to 11.4 times more likely to also make an even allocation compared to the CONTROL (Model 4 in Table 2). When children see an UNEVEN split, they are 1.9 to 8.7 times less likely to offer an even split (Model 4).

Unlike the impact of cultural learning illustrated above, individual learning plays little role over four rounds of repeated play (Figure 1, left panel). In three of our conditions, all except the EVEN condition, children alter their allocations in ways that increase their payoffs—see the coefficients in Model 4 (Table 2) for Round (the control) and the interactions of each Condition and Round. However, these effects are small and not always estimated with precision. The EVEN condition (interacted with Round) appears slightly anomalous but this results from the fact that nearly all children in this condition made even allocations in the first round. Overall, the impact of cultural learning from the demonstrator dominates individual experience, even in the last round.

Across conditions, as shown in the right panel of Figure 1, children distributed increasingly evenly with age, although the impact of imitation still shows through across the entire age span. Models 3 and 4 in Table 2 suggest that for each additional year of life, the odds of making an even allocation go up by about 27% (1.27). The confidence intervals on these estimates, however, include 1, stretching down to 0.94, so we can't be too confident about this result. Such a rising preference for equality, assuming it exists, means that children in the NICE condition increasingly missed out on potential payoffs as they aged—they became less payoff-maximizing.

To verify these results, we conducted a supplemental study with 39 additional participants that sought to (1) replicate our main finding for the EVEN and SELFISH conditions and (2) probe children's understanding of the task. The results, detailed in Section S 3, replicate the relevant findings just discussed and reveal how children understood the rules of the the game.

4.1.1 Do mentalizing abilities improve cultural learning?

To further test predictions from the CCBH, we analysed the relationship between selective imitation in the sticker game and our three measures of mentalizing, controlling for general cognitive ability (BIA). If mentalizing is 'for' sharpening cultural learning to be more selective, then we'd expect that better mentalizers would copy fewer of the demonstrators' irrelevant actions (e.g. tapping, counting, shuffling, etc.). Recall that before distributing the stickers in each round, the experimenter consistently performed five actions that were not connected to the actual sticker allocations. At test, we tallied how many of these behaviors children imitated, and modelled the total counts in each round in a series of Poisson regressions (Table 3), again using robust standard errors adjusted by clustering on both subjects and sampling site.

Across models, we observed a clear association between children's cognitive abilities, mentalizing capacities and the extent of their overimitation. Figure 2 illustrates that greater mentalizing capacities as indexed by (A) passing the false-belief test (Table S1), (B) higher scores on the ToM Storybooks (S2), or (C) greater parent-reported capacities for reasoning about mental states (CSUS; Table S3) was associated with decreased overimitation (see Table 3. The effects of mentalizing are large, as illustrated in Figure 2, though the point estimate for the coefficient on the CSUS is poorly estimated. In contrast to mentalizing, greater cognitive abilities as measured by the BIA are associated with *more* overimitation. Indeed, the data hint

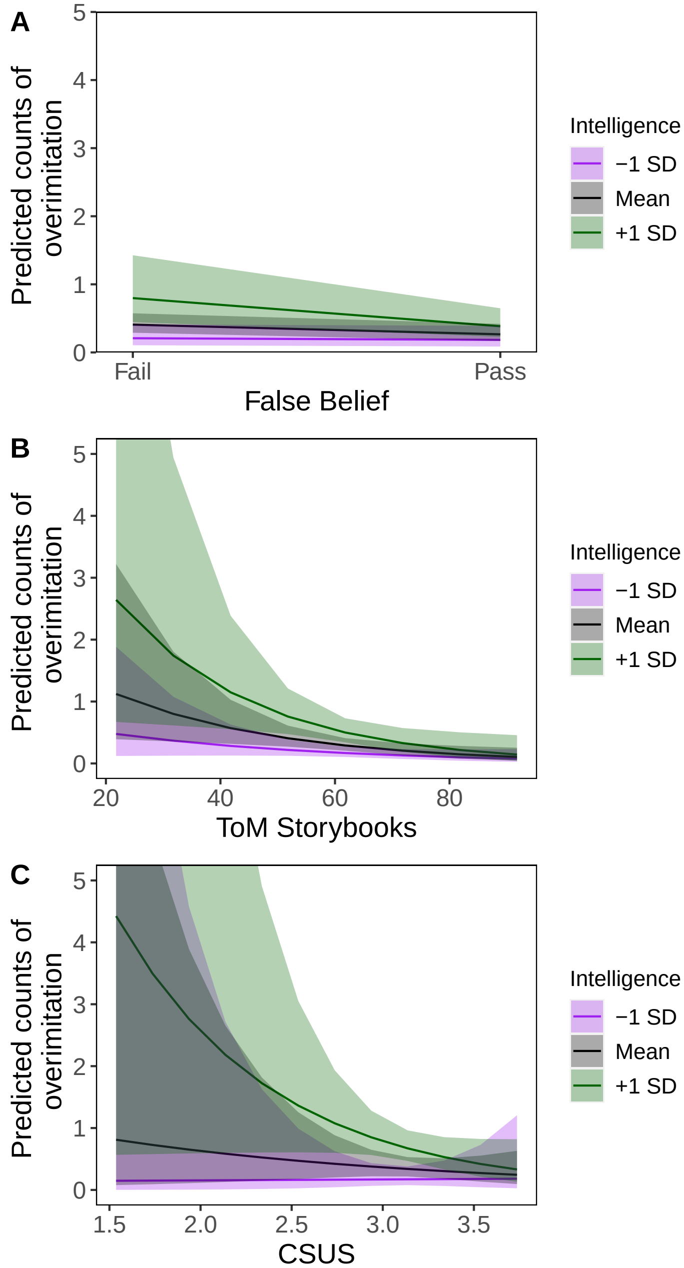

Figure 2: Predicted overimitation counts by mentalizing and general cognitive scores. Shaded regions are 95% confidence intervals. Predictions were generated from models presented in Table 3.

that the stronger the cognitive performance of children on the BIA, the greater the impact of mentalizing on overimitation. Detailed analyses indicate that no one of the three subscales on the BIA is driving the observed relationship with overimitation (Table S4).

Table 3: Poisson regression models to predict counts of overimitation from mentalizing and cognitive ability

	Overimitation (Counts of irrelevant actions by round)		
	False Belief	ToM Storybooks	CSUS
	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	0.408***	0.299***	0.345***
	(0.300, 0.556)	(0.211, 0.423)	(0.224, 0.530)
Mentalizing	0.669**	0.975*	0.487
	(0.464, 0.967)	(0.951, 1.000)	(0.136, 1.739)
Cog. Ability (Centered)	1.048**	1.045***	1.044***
	(1.008, 1.089)	(1.023, 1.068)	(1.030, 1.058)
Mentalizing X Cog. Ability	0.979	0.999	0.963
	(0.941, 1.019)	(0.998, 1.001)	(0.870, 1.067)
Observations	463	300	272
Participants	116	75	68
Sites	18	17	17

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. False Belief was coded as pass/fail (1/0). ToM Storybook and CSUS scores were centered. Models with additional controls are presented in the supplemental: False Belief (Table S1), ToM Storybooks (Table S2) and CSUS (Table S3). For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

4.2 Are children good Machiavellians in this bargaining context?

To test the focal predictions of the Machiavellian Intelligence Hypothesis, we estimated the contributions of mentalizing and cognitive abilities on children's capacities to exploit the responder in order to maximize their own sticker payoffs. The Pay-off Maximizing Choice—the PMC—varied by condition such that 'even' allocations were pay-off maximizing in the EVEN, and SELFISH conditions while 'uneven' distributions were pay-off maximizing in the NICE condition. In the CONTROL condition, participants were blind to the responder's strategy in the demonstration, and thus 'even' allocations were coded as pay-off maximizing until the participant distributed stickers unevenly, and thus revealing the responder's selfish or nice strategy (apriori counterbalanced). If selfish, then 'even' distributions on the *following* round were coded as pay-off maximizing. If the responder was nice, then 'uneven' distributions on the *following* rounds was coded as pay-off maximizing. Children's allocations, indexed as being either payoff-maximizing (PMC = 1) or not (PMC = 0) were modelled in a series of logistic regressions with standard errors adjusted by clustering on subjects and sampling sites[4].

Models that include only the mentalizing and cognitive ability variables are presented in Table 4 (Tables S5-S7 include additional controls). This analysis yields two results. First, we find no consistent or reliable effect of mentalizing on pay-off maximizing decisions. Two measures suggest that greater mentalizing is associated with less payoff-maximizing (the opposite of the MIH's prediction) and one measure suggests a tiny positive effect of mentalizing on payoff-maximizing choices; but, all estimates reveal great uncertainty with large confidence intervals that stretch across 1. Second, a child's BIA scores reveals a small positive association with payoff-maximization, though this too is poorly estimated. These results provide no support

[4]The results that follow were robust to alternative codings of allocations in the CONTROL condition. In additional models, we treated all uneven allocations in the CONTROL condition as *not* pay-off maximizing, and in others treated the first uneven allocation (if the responder was 'nice' as pay-off maximizing despite the child likely 'lucking' into the higher payoff) - neither of which made any substantial changes to the estimates presented in Table 4

Table 4: Logistic regression models to predict payoff maximizing choices from mentalizing and cognitive ability

	Sticker allocations (1 = Payoff maximizing choice)		
	False Belief	ToM Storybooks	Parental Report [CSUS]
	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	2.486***	2.472***	1.891***
	(1.368, 4.517)	(1.583, 3.860)	(1.229, 2.909)
Mentalizing	0.845	1.015	0.646
	(0.422, 1.690)	(0.989, 1.042)	(0.159, 2.618)
Cog. Ability (Centered)	1.005	1.007	1.006
	(0.977, 1.033)	(0.969, 1.047)	(0.968, 1.046)
Observations	463	299	271
Participants	116	75	68
Sites	18	17	17

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. False Belief was coded as pass/fail (1/0). ToM Storybook and CSUS scores were centered. Models with additional controls are presented in the supplemental: False Belief (Table S5), ToM Storybooks (Table S6) and CSUS (Table S7). For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

for the idea that either a child's mentalizing skills or cognitive abilities are deployed to anticipate the predictable actions of one's partners in order to select the payoff maximizing behavior.

5 Discussion

The speed and magnitude of the evolutionary changes in our species—in our brains, life histories and behavior—over the last two million years would seem to demand an auto-catalytic, runaway evolutionary process. The two leading evolutionary hypotheses capable of generating the kinds of feedback loops required are two variants of the social brain hypothesis (Humphrey, 1976): the Machiavellian Intelligence Hypothesis (McNally et al., 2012) and the Cumulative Cultural Brain Hypothesis (Muthukrishna et al., 2017). The former suggests that the driving selection pressures in human evolution arose from an arms race in strategic thinking, with a focus on deception, manipulation and exploitation, created by living in larger and/or more social groups. By contrast, the latter hypothesis argues for a synergy between genes and culture in which cultural evolution generates an ever expanding body of adaptive cultural information that, in turn, favors brains that are better at acquiring, storing and organizing that information. Such markedly distinct evolutionary pressures, if one of them did indeed drive much of human brain expansion, should be readily detectable in modern human psychology.

To test these two leading hypotheses about the driving selection pressures on human brains and psychology—the Machiavellian Intelligence Hypothesis and the Cumulative Cultural Brain Hypothesis—we designed a simple bargaining experiment in which children had the opportunity to use social information in one of two ways, either strategically to exploit an opponent for payoff advantage or for cultural learning to adapt to a novel circumstance. To incorporate individual experience, we also permitted them to engage in individual learning by playing the game over four rounds with the same opponent. Our main results show that the choices of children are strongly shaped by cultural learning while showing little strategic use of the readily available social information.

Supplementing our main analysis, we also collected individual-level measures of children's mentalizing skills and their general cognitive abilities. We focused on mentalizing because both the MIH and the CCBH have pointed to mentalizing as a key capacity in humans that was likely under auto-catalytic selection. Crucially, while the MIH predicts that mentalizing skills will be deployed in the service of Machiavellian efforts to strategically outwhit opponents by anticipating their actions, the CCBH predicts that our greater mentalizing abilities evolved primarily in the service of improving cultural learning. Straightforwardly, all of our measures of mentalizing predict superior cultural learning (more selective imitation, less overimitation) but are not reliably associated with using the available social information to select payoff maximizing options. This further supports the CCBH over the MIH.

In designing this experiment, we strove to 'tilt' the situation in favor of Machiavellian thinking by (1) using a zero-sum bargaining game, (2) allowing learning over repeated interactions with valuable rewards and (3) partnering participants with strangers. In our judgment, the only factor that might tilt towards cultural learning was the fact that the child's opponent was an adult–but this is because evidence suggests that children preferentially copy those who are older and more experienced (VanderBorght and Jaswal, 2009). Future work should explore how the choice of opponents influences the child's decisions. Of course, children and adults are capable of both cultural learning and strategic thinking to some degree, so our enterprise was to look at children's initial and strongest deployment of these abilities and not for a presence vs. absence assessment.

In closing, we note that our study has limitations. First, there may be other evolutionary hypotheses that we've not considered that could deliver this pattern of results. Second, while much cross-cultural evidence supports the centrality of cultural learning for children, it remains an important concern that we've sampled but a single population. Having refined our protocol and obtained interesting results, we hope to collect similar data in diverse populations. If we are truly seeing a robust product of deep evolutionary forces, we should find qualitatively similar results elsewhere. Finally, here we focused on several measures of mentalizing and used the BIA as a control; future work should collect and explore a larger battery of cognitive measures across a more diverse range of contexts.

Acknowledgements

We thank the participating children, their parents, and the daycare centres we visited in the greater Vancouver area. We thank Andy Baron and the Living Lab at Science World. We thank the team of research assistants who worked on this project: Isabel Avelino, Alison Beaumont, Jabeene Bhimji, Kelsey Chan, Kelsey Chapman, Phoebe Cheng, Kiran Chatha, William Cash Hateley, Jessica Hsu, Rosemary Jacobsen, Emi Levy, Alison Lowe, Donya Mahdavi, Lily Mak, Daniel Maldoff, Tierney Mielne, Ariella Nudell, Michelle Packham, Jamie Seo Hee Lee, John Tran, Oscar Velasquez, Nicole Yam, Xuanni Yang, Victor Yin, Ruo Bing Zhang, Roy Zenan Zhou, and Mona Zhu. JH thanks CIFAR and the Social Science and Humanities Research Council in Canada.

Author Contributions

MJ and JH conceived and designed the study. AB and MJ conducted data collection. AB and MJ completed the statistical analyses. AB, MJ, SB and JH wrote the paper. AB and MJ share first co-authorship.

Conflicts of Interest

All authors declare no conflicts of interest.

Research Transparency and Reproducibility

All data and analysis scripts are publicly available on the Open Science Framework: https://osf.io/dgyrj/ (DOI: 10.17605/OSF.IO/DGYRJ)

References

Bailey, D.H., Geary, D.C., 2009. Hominid Brain Evolution. Human Nature 20, 67–79. doi:10.1007/s12110-008-9054-0.

Blijd-Hoogewys, E.M.A., van Geert, P.L.C., Serra, M., Minderaa, R.B., 2008. Measuring Theory of Mind in Children. Psychometric Properties of the ToM Storybooks. Journal of Autism and Developmental Disorders 38, 1907–1930. doi:10.1007/s10803-008-0585-3.

Boyd, R., 2018. A Different Kind of Animal. Princeton University Press, NJ, USA.

Boyd, R., Silk, J.B., 2012. How Humans Evolved. W. W. Norton & Co., New York.

Brosseau-Liard, P., Penney, D., Poulin-Dubois, D., 2015. Theory of mind selectively predicts preschoolers' knowledge-based selective word learning. British Journal of Developmental Psychology , n/a–n/adoi:10.1111/bjdp.12107.

Byrne, R.W., Whiten, A., 1988. Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans. Oxford University Press, Oxford.

Chudek, M., Baron, A.S., Birch, S., 2016. Unselective Overimitators: The Evolutionary Implications of Children's Indiscriminate Copying of Successful and Prestigious Models. Child Development 87, 782–794. doi:10.1111/cdev.12529.

Dunbar, R.I.M., 1998. The social brain hypothesis. Evolutionary Anthropology: Issues, News, and Reviews 6, 178–190. doi:10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8.

Dunbar, R.I.M., Shultz, S., 2007. Understanding primate brain evolution. Philosophical Transactions of the Royal Society B-Biological Sciences 362, 649–658.

Gavrilets, S., Vose, A., 2006. The dynamics of Machiavellian intelligence. Proceedings of the National Academy of Sciences 103, 16823–16828.

Henrich, J., 2015. The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter. Princeton University Press, Princeton.

Humphrey, N., 1976. The social function of intellect, in: Bateson, P.P.G., Hinde, R.A. (Eds.), Growing Points in Ethology, pp. 303–317.

Laland, K.N., 2017. Darwin's Unfinished Symphony: How Culture Made the Human Mind. Princeton University Press, Princeton, NJ.

McGrew, K.S., Woodcock, R.W., 2001. Woodcock-Johnson III Technical Manual. Riverside Publishing, Itasca, IL.

McNally, L., Brown, S.P., Jackson, A.L., 2012. Cooperation and the evolution of intelligence. Proceedings of the Royal Society B: Biological Sciences 279, 3027–3034. doi:10.1098/rspb.2012.0206.

Muthukrishna, M., Doebeli, M., Chudek, M., Henrich, J., 2017. The Cultural Brain Hypothesis: How culture drives brain expansion, underlies sociality, and alters life history. bioRxiv , 209007doi:10.1101/209007.

Muthukrishna, M., Henrich, J., 2016. Innovation in the collective brain. Phil. Trans. R. Soc. B 371, 20150192. doi:10.1098/rstb.2015.0192.

Reader, S.M., Hager, Y., Laland, K.N., 2011. The evolution of primate general and cultural intelligence. Philosophical Transactions of the Royal Society B-Biological Sciences 366, 1017–1027.

Schoenemann, T., 2006. From Monkey Brain to Human Brain: A Fyssen Foundation Symposium. American Journal of Human Biology 18, 722–724. doi:10.1002/ajhb.20556.

Statistics Canada, 2017. Focus on Geography Series, 2016 Census - Census metropolitan area of Vancouver. https://www12.statcan.gc.ca/census-recensement/2016/as-sa/fogs-spg/Facts-CMA-Eng.cfm?TOPIC=6&LANG=Eng&GK=CMA&GC=933.

Tahiroglu, D., Moses, L.J., Carlson, S.M., Mahy, C.E.V., Olofson, E.L., Sabbagh, M.A., 2014. The Children's Social Understanding Scale: Construction and validation of a parent-report measure for assessing individual differences in children's theories of mind. Developmental Psychology 50, 2485–2497. doi:10.1037/a0037914.

Trivers, R.L., 1971. The evolution of reciprocal altruism. The Quarterly Review of Biology 46, 35–57.

VanderBorght, M., Jaswal, V.K., 2009. Who knows best? Preschoolers sometimes prefer child informants over adult informants. Infant and Child Development 18, 61–71. doi:10.1002/icd.591.

Whiten, A., van Schaik, C.P., 2007. The evolution of animal 'cultures' and social intelligence. Philosophical Transactions of the Royal Society B-Biological Sciences 363, 603–620.

Wimmer, H., Perner, J., 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. Cognition 13, 103–128. doi:10.1016/0010-0277(83)90004-5.

Cultural Learners or Game Theorists?

Supplemental Materials

Baimel, Adam*
Department of Psychology, Health and Professional Development
Oxford Brookes University
*abaimel@brookes.ac.uk

Juda, Myriam
Department of Psychology
Simon Fraser University

Birch, Susan
Department of Psychology
University of British Columbia

Henrich, Joseph*
Department of Human Evolutionary Biology
Harvard University
*henrich@fas.harvard.edu

June 5, 2020

S 1 The contributions of mentalizing and cognitive ability to overimitation

S 1.1 False Belief

Table S1: Poisson regression models to predict overimitation from false belief and cognitive ability

	Overimitation				
	Model 1	Model 2	Model 3	Model 4	Model 5
	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	0.438***	0.429***	0.408***	0.494***	0.499***
	(0.271, 0.706)	(0.307, 0.600)	(0.300, 0.556)	(0.372, 0.656)	(0.362, 0.687)
False Belief (1 = Pass)	0.708*	0.583***	0.669**	0.669**	0.652*
	(0.489, 1.026)	(0.389, 0.873)	(0.464, 0.967)	(0.463, 0.966)	(0.416, 1.020)
Cog. Ability (Centered)		1.037**	1.048**	1.048**	1.048**
		(1.006, 1.069)	(1.008, 1.089)	(1.008, 1.089)	(1.010, 1.089)
False Belief X Cog. Ability			0.979	0.979	0.979
			(0.941, 1.019)	(0.941, 1.019)	(0.942, 1.018)
Round (0 = Round 1)				0.874**	0.874**
				(0.774, 0.986)	(0.774, 0.986)
Age (Yrs, Centered)					1.016
					(0.610, 1.690)
Sex (Prop. of Males)					1.223
					(0.727, 2.060)
463 observations					
116 participants					
18 sites					

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 1.2 ToM Storybooks

Table S2: Poisson regression models to predict overimitation from ToM storybooks and cognitive ability

	Overimitation				
	Model 1	Model 2	Model 3	Model 4	Model 5
	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	0.332***	0.285***	0.299***	0.379***	0.375***
	(0.234, 0.469)	(0.198, 0.409)	(0.211, 0.423)	(0.256, 0.562)	(0.254, 0.555)
ToM Storybooks (Centered)	0.989	0.974**	0.975*	0.975*	0.967**
	(0.966, 1.012)	(0.950, 0.998)	(0.951, 1.000)	(0.951, 1.000)	(0.942, 0.992)
Cog. Ability (Centered)		1.044***	1.045***	1.045***	1.042***
		(1.017, 1.071)	(1.023, 1.068)	(1.023, 1.068)	(1.022, 1.064)
ToM Storybooks X Cog. Ability			0.999	0.999	0.999
			(0.998, 1.001)	(0.998, 1.001)	(0.998, 1.001)
Round (0 = Round 1)				0.842**	0.842**
				(0.738, 0.961)	(0.738, 0.961)
Age (Years, Centered)					1.396
					(0.858, 2.273)
Sex (Prop. of Males)					1.008
					(0.547, 1.859)
300 observations					
75 participants					
17 sites					

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 1.3 Parental Report - Children's Social Understanding Scale [CSUS]

Table S3: Poisson regression models to predict overimitation from CSUS and cognitive ability

	Overimitation				
	Model 1	Model 2	Model 3	Model 4	Model 5
	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	0.427***	0.333***	0.345***	0.374***	0.368***
	(0.257, 0.711)	(0.215, 0.515)	(0.224, 0.530)	(0.231, 0.604)	(0.223, 0.609)
CSUS (Centered)	0.625	0.389**	0.487	0.487	0.580
	(0.308, 1.266)	(0.163, 0.925)	(0.136, 1.739)	(0.136, 1.743)	(0.137, 2.454)
Cog. Ability (Centered)		1.049***	1.044***	1.044***	1.046***
		(1.021, 1.079)	(1.017, 1.072)	(1.017, 1.072)	(1.017, 1.076)
CSUS X Cog. Ability			0.963	0.963	0.959
			(0.870, 1.067)	(0.870, 1.067)	(0.861, 1.070)
Round (0 = Round 1)				0.947	0.947
				(0.872, 1.029)	(0.872, 1.029)
Age (Years, Centered)					0.849
					(0.495, 1.454)
Sex (Prop. of Males)					1.314
					(0.790, 2.185)
272 observations					
68 participants					
17 sites					

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 1.4 Cognitive Ability Subscale Analysis

We modeled whether the relationships between cognitive ability and overimitation presented in the main text (see Table 3) could be further qualified by examining the associations between the three subscales of the Brief Intellectual Ability [BIA] test and amount of overimitation. The BIA score is made up from the equally-weighted results of three individual tests - a test of (1) concept formation, (2) verbal comprehension, and (3) visual matching. The concept formation test asks participants to identify rules that define patterns in sequences of geometric figures. The verbal comprehension test asks participants to name pictured objects, identify synonyms and antonyms of said word. The visual matching test has participants identify (e.g., point to) as many of matching pairs of numbers in a row of six numbers as quickly as they can in a three-minute time period. The subcsale analyses reveal that the already small effect of cognitive ability on overimitation may be driven mostly by verbal and visual matching scores rather than concept formation (see Table S4 for details).

Table S4: Poisson regression models to predict overimitation by mentalizing and the subscales of the cognitive ability test

	Overimitation		
	False Belief	ToM Storybooks	CSUS
	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	0.448***	0.280***	0.342***
	(0.314, 0.640)	(0.197, 0.398)	(0.203, 0.577)
Mentalizing	0.578***	0.977*	0.378**
	(0.394, 0.847)	(0.953, 1.001)	(0.175, 0.819)
BIA - Concept Formation (Centered)	1.001	1.013	0.984
	(0.978, 1.025)	(0.995, 1.031)	(0.950, 1.019)
BIA - Verbal (Centered)	1.016**	1.015*	1.022***
	(1.002, 1.030)	(0.999, 1.031)	(1.008, 1.037)
BIA - Visual Matching (Centered)	1.015	1.027**	1.037**
	(0.979, 1.052)	(1.000, 1.055)	(1.008, 1.067)
Observations	447	292	256
Participants	112	73	64
Sites	18	17	16

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 2 The contributions of mentalizing and cognitive ability to payoff maximizing decisions

S 2.1 False Belief

Table S5: Logistic regression models to predict payoff maximizing decisions from false belief and cognitive ability

	Sticker allocations (1 = Payoff maximizing choice)			
	Model 1	Model 2	Model 3	Model 4
	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	2.486***	15.735***	12.393***	14.302***
	(1.368, 4.517)	(4.850, 51.053)	(3.253, 47.222)	(4.044, 50.577)
False Belief (1 = Pass)	0.845	0.729	0.726	0.592
	(0.422, 1.690)	(0.312, 1.703)	(0.308, 1.709)	(0.240, 1.461)
Cog. Ability (Centered)	1.005	1.003	1.003	1.000
	(0.977, 1.033)	(0.976, 1.031)	(0.976, 1.031)	(0.972, 1.030)
Even Condition		0.511	0.510	0.483
		(0.140, 1.860)	(0.140, 1.860)	(0.150, 1.558)
Nice Condition		0.216**	0.215**	0.210***
		(0.061, 0.766)	(0.061, 0.758)	(0.067, 0.662)
Selfish Condition		0.047***	0.046***	0.046***
		(0.015, 0.143)	(0.015, 0.138)	(0.018, 0.118)
Round (0 = Round 1)			1.187	1.189
			(0.964, 1.460)	(0.964, 1.467)
Age (Yrs. Centered)				1.495
				(0.831, 2.690)
Sex (0 = Prop. of Males)				0.910
				(0.452, 1.830)
Observations = 463				
Participants = 116				
Sites = 18				

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. The CONTROL condition (Intercept; controlling for other variables) is the reference category for condition effects. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 2.2 ToM Storybooks

Table S6: Logistic regression models to predict payoff maximizing decisions from ToM storybooks and cognitive ability

	Sticker allocations (1 = Payoff maximizing choice)			
	Model 1	Model 2	Model 3	Model 4
	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	2.472***	14.616***	11.122***	11.336***
	(1.583, 3.860)	(4.107, 52.011)	(2.401, 51.514)	(2.260, 56.852)
ToM Storybooks (Centered)	1.015	1.012	1.013	1.018
	(0.989, 1.042)	(0.984, 1.042)	(0.983, 1.043)	(0.982, 1.056)
Cog. Ability (Centered)	1.007	1.004	1.004	1.001
	(0.969, 1.047)	(0.964, 1.046)	(0.964, 1.046)	(0.953, 1.051)
Even Condition		0.560	0.559	0.564
		(0.113, 2.788)	(0.112, 2.792)	(0.113, 2.820)
Nice Condition		0.153**	0.151**	0.142**
		(0.030, 0.786)	(0.030, 0.768)	(0.025, 0.796)
Selfish Condition		0.051***	0.050***	0.050***
		(0.012, 0.210)	(0.012, 0.201)	(0.011, 0.223)
Round (0 = Round 1)			1.217	1.217
			(0.922, 1.606)	(0.920, 1.610)
Age (Yrs. Centered)				0.870
				(0.379, 1.995)
Sex (0 = Prop. of Males)				0.734
				(0.353, 1.527)
Observations = 299				
Participants = 75				
Sites = 17				

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. The CONTROL condition (Intercept; controlling for other variables) is the reference category for condition effects. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 2.3 Parental Report [CSUS]

Table S7: Logistic regression models to predict payoff maximizing decisions from CSUS and cognitive ability

	Sticker allocations (1 = Payoff maximizing choice)			
	Model 1	Model 2	Model 3	Model 4
	OR (95% CI)	OR (95% CI)	OR (95% CI)	OR (95% CI)
Intercept	1.891***	9.172***	6.859***	7.339***
	(1.229, 2.909)	(3.528, 23.845)	(2.070, 22.731)	(2.332, 23.096)
CSUS (Centered)	0.646	0.601	0.595	0.301
	(0.159, 2.618)	(0.124, 2.902)	(0.121, 2.925)	(0.053, 1.699)
Cog. Ability (Centered)	1.006	1.007	1.007	1.002
	(0.968, 1.046)	(0.972, 1.044)	(0.971, 1.045)	(0.965, 1.039)
Even Condition		0.462	0.459	0.403
		(0.146, 1.463)	(0.144, 1.462)	(0.117, 1.394)
Nice Condition		0.293*	0.290*	0.263*
		(0.076, 1.134)	(0.075, 1.120)	(0.066, 1.057)
Selfish Condition		0.057***	0.055***	0.057***
		(0.020, 0.163)	(0.020, 0.154)	(0.021, 0.157)
Round (0 = Round 1)			1.232	1.244
			(0.952, 1.593)	(0.953, 1.623)
Age (Yrs. Centered)				2.165*
				(0.926, 5.062)
Sex (0 = Prop. of Males)				0.705
				(0.244, 2.038)
Observations = 271				
Participants = 68				
Sites = 17				

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and use two-way clustering on both individuals and sites. 95% confidence intervals are reported below each coefficient in parentheses. The CONTROL condition (Intercept; controlling for other variables) is the reference category for condition effects. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 3 Supplemental Study: Exploring children's comprehension of the sticker game

In our main study, children only ever saw the demonstrator perform one particular allocation and this had a big impact on children's own allocations. An important question is what children inferred from this demonstration. Since our instructions implied that participants could allocate the stickers however they wanted, children most likely inferred that the model's action represented either a 'good strategy' in this interaction or the normatively correct standard in this situation. Either inference is consistent with view assumed in the main text. However, children may have inferred from the model's demonstration that the only permissible action was to allocate the stickers in precisely the same manner as the model. The view is subtly but importantly different from inferring something normative. As an analogy, young basketball players might watch an experienced player shoot using an underhand technique (e.g. NBA star Rick Barry). They might assume that you must shoot underhanded in basketball (or else it doesn't count and causes a 'turnover'); or, they might see this as the usual approach that people take in shooting, but that you can shoot overhand if you prefer (but others may think it is a bit odd). To examine this question, we conducted a small supplemental study in which children played the sticker game in an identical manner as in our main study. Following the game, children were asked a series of questions regarding the interaction to determine how children understood the 'rules' of the game.

S 3.1 Methods

Forty-four children were recruited from the Living Lab at The Telus World of Science Museum in Vancouver, Canada. Five of these participants were excluded from all analyses for three reasons (1) experimenter error (incorrect instructions were given to the child during the observation phase), (2) difficulties with answering the comprehension check questions in English or (3) having watched a sibling play the sticker game prior to participating. Our final sample of 39 contained 17 females and ranged in age from 3.58 to 6.93 years (M = 5.22, SD = 1.07).

Participants in this study were randomly assigned to one of two conditions (EVEN: N = 20; SELFISH: N = 19. The CONTROL condition from Study 1 was not replicated here, as there was no cause for concern regarding imitation effects as allocations were occluded from the participants' view. We included the SELFISH condition (but not the NICE condition) because if responses to follow-up questions in the SELFISH condition indicate that children understood that the stickers could be distributed differently than how they had observed, yet continued to imitate the unfavorable uneven distribution that resulted in reduced sticker payoffs, we could be more confident that these behaviors are the result of a propensity for imitation and not a lack of understanding or strict rule following. All participants played the game with the same two female experimenters who played the same role (proposer or responder) with each participant. Otherwise, the sticker game proceeded exactly as described in the main study. After the game, the experimenter who had played as proposer in the sticker game asked the participant six questions. These questions are described in tandem with the results below.

S 3.2 Results

In this section, we first show that we replicated the relevant results from the main text in this supplemental experiment and then explore how our participants understood the game using our interview protocol.

S 3.2.1 Replicating relevant results

As in the main study, children's allocations were strongly influenced by the allocation strategy they saw in the observation phase (see Figure S1 and Table S8 for model summary details). Note that the regression coefficients here, expressed in odds ratios, are relative to the SELFISH condition, (not a CONTROL condition as is presented in the main text), which is why they are so large. The confidence intervals are large because with 80 total observations in EVEN Condition, we have only 5 *uneven* observations. Nevertheless, the main results for these conditions in the main text are replicated here.

Table S8: Logistic regression models to predict uneven/even allocations in Study 2

	Sticker Allocations (0 = Uneven; 1 = Even)			
	Model 1	Model 2	Model 3	Model 4
	(1)	(2)	(3)	(4)
Intercept	1.303	0.910	0.930	0.878
	(0.614, 2.764)	(0.400, 2.071)	(0.383, 2.258)	(0.360, 2.139)
Even Condition	11.512***	11.828***	11.444**	14.943***
	(1.816, 72.963)	(1.810, 77.308)	(1.386, 94.501)	(2.012, 110.981)
Round (0 = Round 1)		1.274*	1.275*	1.328*
		(0.999, 1.625)	(0.998, 1.629)	(0.987, 1.785)
Age (Yrs. Centered)			1.121	1.121
			(0.402, 3.123)	(0.399, 3.147)
Sex (0 = Prop. of Males)			1.105	1.106
			(0.222, 5.514)	(0.219, 5.586)
Even Condition X Round				0.821
				(0.576, 1.168)
Observations = 156				
Participants = 39				

Notes: Coefficients are presented as odds ratios, so "1" indicates no effect. Standard errors and confidence intervals are robust and clustered on individuals. 95% confidence intervals are reported below each coefficient in parentheses. The SELFISH condition (Intercept; controlling for other variables) is the reference category for condition effects. Round of the game was treated as a continuous variable. Sex was centered on the percentage of males to ease interpretation of the other coefficients for the entire sample. For those interested in significance testing, ***, **, and * indicate p-values below 0.01, 0.05 and 0.1.

S 3.2.2 Participant's comprehension of the game

The post-game interviews of these participants unfolded as follows. First, at the completion of game, the experimenter exclaimed that the other research assistant had forgotten the rules of the game, and asked whether or not the child could teach her how to play the game. The child was then asked to indicate whether not the experimenter was allowed to distribute stickers in (1) an even manner (two in each basket), (2) uneven manner (three in one basket and one in the other), and (3) another uneven manner in which four stickers were placed in one basket and none in the other. Overall, across both conditions, roughly 60% of participants explicitly expressed the view that they could have done something different from the demonstrator and only 1 participant out of 39 said that an even distribution was not acceptable. This implies that participants didn't see deviations form the allocations they observed as rule violations.

However, children's inferences about the situation were not symmetrical across our two conditions. Crucially, participants in the SELFISH condition saw it as permissible to payoff maximize by making even offers; but, despite this recognition, they tended to copy the allocations of their demonstrator. Yet, in the EVEN condition, a small majority of participants (12 out of 20) thought that an uneven distribution would *not* be allowed. We cannot be sure whether children felt an uneven distribution was non-normative or an actual rule violation. This also means that 30% of participants thought that uneven distributions were permissible.

Participants were then asked if they remembered what the proposer in the observation phase did on her turns in the game and to indicate how many stickers she had put in each basket. Six children in the SELFISH condition and 7 children in the EVEN condition said they did not remember the allocations. Of those who did recall, 1 child out of 11 incorrectly stated the demonstrator's allocation in the EVEN condition

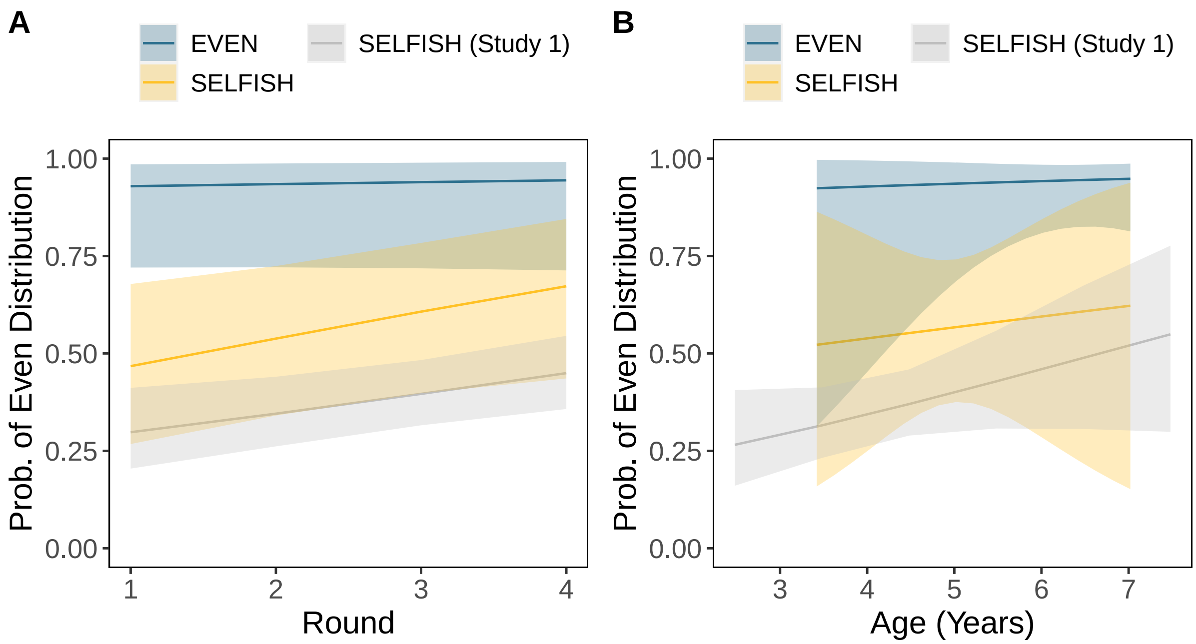

Figure S1: **Predicted probability of even distributions in the two conditions of Study 2 across the four rounds (Panel A) and age (Panel B)**. Predictions were generated from Model 4 in S8. The shaded regions show the 95% confidence intervals based on subject-level clustering. The grey lines reproduce predicted estimates from the SELFISH condition in Study 1 for comparison. Study 1 recruited children of a wider age-range than Study 2.

as did 2 out of 13 in the SELFISH condition. Then, children's memory of their own behaviors in the game was assessed in the same manner. Six children incorrectly remembered their own decisions: 2 in the EVEN condition and 4 in the SELFISH condition.

Following these memory checks, the experimenter recounted how the model distributed stickers in the observation phase and what the child did at test and then asked, "Could you have put the stickers in the baskets in any other way?". The results were almost identical with those above. Again, nearly two-third of participants explained that they could have deviated from the demonstrators' allocation. However, in the EVEN condition, 12 out of 20 children again thought that an uneven distribution would not be allowed. In the SELFISH condition, 2 children out of 19 thought only the demonstrators uneven allocation was allowed–that is, 17 children thought they could deviate from what they saw the demonstrator do.

Next, participants were asked, "Would you have been allowed to just take the stickers without even putting them into the baskets?" The answer to which is technically 'yes', however we wanted to see if children understood this situation to be a game with a certain set of boundary conditions. And unlike the other questions we asked, this question provided a response in which the expected modal answer would be 'no'. Indeed, only 8 participants (3 in the EVEN condition, and 5 in the SELFISH condition) said that they could have taken the stickers without first putting them in the baskets.

Lastly, we probed whether participants could explicitly reason about sticker distribution strategies by asking them, "While you were playing, if you thought [name of experimenter] was always going to choose the basket with the most stickers in it, how would you play the game in order to get the most stickers?" This was an open-ended question and responses were later coded for the presence/absence of mentioning an even distribution which is the strategic allocation given uncertainty regarding the responder's decisions in the EVEN condition, and knowing that the responder was SELFISH in the other condition. Many children provided no or irrelevant answers. Of those that did provide a relevant answer (11 in the EVEN condition and 12 in the SELFISH condition); 9 in the EVEN condition hinted at an explicit understanding that an even distribution was the best strategy, where as only 4 explicitly reported the same in the SELFISH condition.

S 3.2.3 Discussion

In this supplemental study, we sought to replicate certain key results from the main text and to probe children's explicit understanding of the rules of the sticker game. Despite the small sample size, the results from the main text replicate. On the question of children's inferences about normativity or permissibility of certain allocations in the game, we find a nuanced picture. Crucially, in the SELFISH condition where copying the model's allocations results in the participant getting fewer stickers, children overwhelmingly felt that they could deviate from the model's allocations, either by allocating 2/2 or 4/0 stickers. This means that the costly allocations of participants in the SELFISH treatment cannot be explained by confusion about the rules. This relieves an important methodological concern as it shows that our instructions themselves didn't lead children to automatically infer that they had to do whatever their demonstrator did.

However, we did find an interaction of the condition with our instructions. The impact of the demonstrators actions in the EVEN condition seemed to steer a small majority of participants toward the view that only the even allocation would be permitted. Here, the cultural transmitted information, perhaps because it dovetailed with some expectations that children brought into the lab with them about equal splits, caused some to infer that only an even split was permitted. Notably, older children were more likely to say that 2/2 was the only allocation permitted (Saying that alternative allocations were allowed was negatively correlated with age: r = -.64 for "3/1" allocations and r = -.76 for "4/0" allocations in the EVEN condition). Of course, some 40% of participants in the EVEN condition didn't make that inference. These data suggest how cultural learning shapes people's construction of the"rules of the game" and is likely relevant to understanding institutions. This finding underlines the centrality of cultural learning in children and certainly isn't the kind of mistaken inference that we'd expect under the Machiavellian Intelligence Hypothesis.

Note, although we find these results sufficient to relieve our concerns that our instructions may been misleading to children across the board, we haven't included this speculative discussion in the main text given the sample sizes and uncertainties involved.