Children's strategic theory of mind

Itai Sher[a], Melissa Koenig[b,1], and Aldo Rustichini[a]

[a]Department of Economics and [b]Institute for Child Development, University of Minnesota, Minneapolis, MN 55455

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved July 30, 2014 (received for review February 25, 2014)

Human strategic interaction requires reasoning about other people's behavior and mental states, combined with an understanding of their incentives. However, the ontogenic development of strategic reasoning is not well understood: At what age do we show a capacity for sophisticated play in social interactions? Several lines of inquiry suggest an important role for recursive thinking (RT) and theory of mind (ToM), but these capacities leave out the strategic element. We posit a strategic theory of mind (SToM) integrating ToM and RT with reasoning about incentives of all players. We investigated SToM in 3- to 9-y-old children and adults in two games that represent prevalent aspects of social interaction. Children anticipate deceptive and competitive moves from the other player and play both games in a strategically sophisticated manner by 7 y of age. One game has a pure strategy Nash equilibrium: In this game, children achieve equilibrium play by the age of 7 y on the first move. In the other game, with a single mixed-strategy equilibrium, children's behavior moved toward the equilibrium with experience. These two results also correspond to two ways in which children's behavior resembles adult behavior in the same games. In both games, children's behavior becomes more strategically sophisticated with age on the first move. Beyond the age of 7 y, children begin to think about strategic interaction not myopically, but in a farsighted way, possibly with a view to cooperating and capitalizing on mutual gains in long-run relationships.

practical reasoning | child development | game theory

Strategic environments determine outcomes as a function of the decision of many players. Behavior in these environments is directed by a capacity we term strategic theory of mind (SToM), the capacity to infer other people's mental processes and predict their behavior on the basis of knowledge of their incentives and assumption of their rationality. SToM requires two more primitive capacities. The first is "ordinary" theory of mind (ToM), which is a person's ability to "impute mental states to himself and others" (1). The second is recursive thinking (RT), which is the ability to use the output of one step of a reasoning process as input to a following step. In addition, SToM requires that agents reason about the incentives of all involved.

To take an example of the combination of ToM, RT, and reasoning about incentives that will be relevant in the context of this study, suppose that Ann and Bob play a game with rules known to both players, where the rules imply that Ann has an incentive to lie to Bob when Bob believes her. Ann can conclude that she should lie to Bob. Similarly, Bob can conclude that Ann has an incentive to lie to him, and hence he will not believe her. However, in the same manner, Ann concludes that Bob will perform exactly the same reasoning, so that Bob will be skeptical in view of Ann's incentive to lie. Ann will use this output of her reasoning about Bob's incentives as input for her next step of reasoning, and conclude that it is best for her to be truthful instead. A further application of the pattern will suggest that Bob will realize that Ann will perform this reasoning as well, and so on. In each step, the information about the incentives of others is used to predict their actions, given their beliefs and incentives.

Understanding SToM relies on a more subtle distinction within ToM, between the epistemic capacity to understand what others will believe and the practical capacity to understand what others will decide to do in light of their beliefs. In strategic reasoning, a child may need to apply these two types of understanding in sequence. For example, in one of our experiments, if a child attributes to her adult opponent the mistrustful belief that the child will lie (using the epistemic capacity), the child might conclude that the opponent will do the opposite of what the child suggests (using the practical capacity). SToM requires not only that a child be able to answer specific epistemic or practical questions related to ToM but that, without being asked, the child be able to call on answers to such questions recursively to decide what she should do, in light of the incentives applying to all.

The ability to use first-order ToM reasoning flexibly develops at around 3–4 y of age (2–4). The ability to perform RT emerges at a later age [about 7–8 y of age (5)], although the ability to understand recursive notions, such as the successor function and the numbering system, appears earlier (6, 7). Thus, if SToM indeed results from integrating the two capacities, one might expect SToM to emerge at a later age, perhaps substantially later if integrating the two is more complex than using the two in parallel. Although various accounts have been proposed for how first-order ToM develops, second-order developments have received less attention and are less understood (4, 8). Using two strategic games played by children aged 3–9 y, and showcasing different kinds of incentives, we test the hypothesis that SToM results from an integration of ToM and RT, and identify the age at which SToM emerges in children.

Perner and Wimmer (9) show that children demonstrate competence in higher order thinking relating to ToM at the age of 6–7 y. Unlike us, Perner and Wimmer (9) neither study higher order thinking in games nor offer an account of what explains (or supports) the development of higher order thinking. Moreover, such ToM research isolates conceptual developments, giving little attention to how children integrate their conceptual understanding with their practical decisions about how to act in light of considerations about how others will likely act. In philosophical terms, much psychological research targets developments in children's theoretical reasoning but neglects the interaction of theoretical and practical reasoning. A typical first- or second-order false belief task highlights questions about belief (where will X expect the ice cream truck to be?) while making information about goals and actions transparent conditional on those beliefs (where

Significance

Human interaction requires reasoning not only about other people's observed behavior and mental states but also about their incentives and goals. The development of children's strategic thinking is not well understood, leaving open critical questions about early human capacity for strategic interaction. We investigated strategic reasoning in 3- to 9-y-old children and adults in two strategic games that represent prevalent aspects of social interaction: incentives to mislead and competition. We find that despite strategic differences in the two games, by the age of 7 y, children's behavior is similar to that of adults. Our findings also show an early sophisticated ability to think strategically about others in both static and repeated interactions.

Author contributions: I.S., M.K., and A.R. designed research; I.S., M.K., and A.R. performed the research; A.R. analyzed data; and I.S., M.K., and A.R. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: mkoenig@umn.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1403283111/-/DCSupplemental.

will X look if she expects the ice cream truck to be at Y?). In contrast, in the games used in the current experiment, instead of merely attributing beliefs, participants need to anticipate one another's actions and recognize that those predictions turn on what the other person believes I will do, which, in turn, depends on what the other person believes I believe the other person will do. It is common for children to have information about people's goals, and thus many of their interactions with others, such as siblings, friends, and parents, set the stage for not only theoretical but interactive practical reasoning.

Some papers have studied children playing games with a focus on fairness and trustworthiness (distinct from truthfulness) in ultimatum and trust games (10–12); this literature builds on earlier experimental economics research with children (13–15). Sally and Hill (16) examined the effects of autism on children's performance in games, with a focus on fairness. Other related papers on children's strategic sophistication in games have been written by Perner (17) and Shultz and Cloghesy (18). Perner (17) studies two-by-two games in matrix form, with each containing a dominant strategy for at least one player. In contrast, neither of our games has a dominant strategy, so Perner (17) cannot differentiate, as we do, between games with different levels of strategic complexity.

In the developmental literature on deception, young children are often seen to have difficulty in deceiving others (19, 20) and in interpreting overtly deceptive points from others (21, 22). In the literature on children's ability to mistrust others, children are typically presented with putatively unreliable speakers (of inaccurate, ignorant, and antisocial varieties) (23–25). Notably, in all of this research, children are not asked to consider others' incentives. Thus, we examine how games that contain incentives to be dishonest lead children to mistrust, and to be deceptive themselves. Furthermore, one potential limit of the studies mentioned above is that children are typically asked to perform in only one communicative role: that of sender or receiver of communicative signals.

In our study, the child played the same game both as sender and as receiver. This feature of our design allowed us to examine whether children's strategic reasoning emerges in a similar fashion across these two roles. An additional novelty of our design is that children played two games, each against an experimenter. The two games were chosen because they have completely different game theoretical properties, so they provide an independent test of our hypothesis because ToM and RT operate differently in the two games.

In the sender–receiver game, the child played one of two roles: She was either the sender or the receiver. An experimenter occupied the other role. The sender knows the location of a piece of candy placed in one of two boxes, and the receiver does not. The sender points to one of the boxes, not necessarily the box containing the candy. The receiver then selects a box. If the receiver selects the box with the candy, the receiver keeps the candy; otherwise, the sender gets the candy. As in our introductory example, the sender has an incentive to deceive the receiver if the receiver believes the deception.

In the stickers game, the child and experimenter simultaneously select a number of stickers between one and five. Whoever selects strictly fewer stickers gets to keep her stickers, whereas the other player receives nothing. If both players select the same number of stickers, neither keeps any stickers. Each player has an incentive to undercut the other player by as little as possible: For example, a player who knows that her opponent will select three stickers should select two stickers.

Both the sender–receiver game and the stickers game were played repeatedly, allowing observation of both children's first moves and the subsequent evolution of play. Observation of first moves reveals the outcome of the child's a priori reasoning process, allowing us to draw inferences on the role of SToM therein, and the evolution of this role with age. Observation of subsequent moves reveals the incremental outcomes of the child's learning processes, allowing us to draw inferences about how SToM interacts with learning from experience.

We ran a parallel experiment on adults, which is described in *SI Appendix*, to allow for a comparison between children and adult behavior. *SI Appendix* contains a detailed description of our procedure for children, as well as additional results on children's behavior in later rounds in the two games. In this setting, the specific questions we address are the following. First, at what age do children show evidence of SToM? Second, is the age of acquisition of SToM constant across simple games of different kind, or does it depend on specific properties of the game? Finally, how is the acquisition of SToM modulated by other cognitive capacities, such as working memory?

Strategic Analysis of the Games

There are two important strategic dimensions, widely discussed in the game theoretical literature, along which the two games differ: scope for cooperation and action selection [with the latter corresponding to availability of (iteratively) dominated actions and existence of pure strategy equilibrium].

The sender–receiver game is zero-sum, meaning that the interests of the players are exactly opposed. The zero-sum nature and the symmetry of the game imply that if the players are rational, each will win half of the time on average. Thus, there is no scope for cooperation in the sender–receiver game.

In contrast, the stickers game allows scope for cooperation when repeated. If played once, rational play predicts that each player selects just one sticker (see the discussion of iterated dominance below). In that case, neither player wins anything. If players play repeatedly, however, as they do in our experiment, they can benefit from cooperation: If on each round, one player selected five stickers and the other selected four, and the players alternated between these two roles, each player would then alternate between winning zero and four stickers, winning an average of two stickers per round.

Putting cooperation aside, the game theoretical concept of iterated dominance provides a unique prediction when the game is played once by fully rational players: Both should select just one sticker. This reasoning starts by arguing that no player should ever select five stickers because, regardless of the other's play, selecting five leads to zero winnings for the player who does so. If both players recognize this point and neither selects five stickers, a similar argument shows that neither should select four stickers. Iterating, one concludes that both should select one sticker. The choice of one sticker by each player is also a symmetrical Nash equilibrium. There are two asymmetrical Nash equilibria: In each, one player selects one sticker and the other player selects two stickers. One-shot play makes it difficult to coordinate on one of the two asymmetrical equilibria, especially because each player would prefer the equilibrium where she/he is the one choosing one sticker. Thus, we do not expect to observe these asymmetrical equilibria. In sum, both iterated dominance and symmetrical Nash equilibrium predict one sticker.

In contrast, no simple prediction is possible in the sender–receiver game: The rule prescribing that the sender should tell the truth could be exploited by a receiver choosing the box the sender indicates, and the rule prescribing that the sender should lie could be exploited by a receiver choosing the opposite box. Iterated dominance does not eliminate any possibilities in the sender–receiver game; thus, complete perfectly rational application of ToM and RT leads to no clear prescription in the sender–receiver game. In addition, the sender–receiver game does not contain a pure strategy Nash equilibrium. Rather, Nash equilibrium prescribes not a single action but a probability distribution over actions: Rational players should randomize equally among actions to make it impossible for the opponent to exploit any systematic tendency to choose one action over the other.

In our analysis, we are primarily interested in the first round of play, which isolates children's ability to reason about the game before play. The optimal rational reasoning processes described above used ToM and RT extensively. Thus, we do not expect all children to behave according to that prescription. The youngest children are likely to play naively; as they age, we expect children to apply recursive thinking at successively higher levels. We assume that at each stage in their reasoning, children best respond to their beliefs about opponents' play, meaning that children choose the action that is best in light of their beliefs.

In the sender–receiver game, this analysis implies that children, as they age, will move along the following series of steps. Each substep below labeled (*a*) involves a use of ToM in which the child comes to understand what the other player will believe and do, and each substep labeled (*b*) involves a decision of the child about what to do in light of her beliefs about the other player:

i) Children play naively, reporting truthfully as sender and trusting as receiver.

ii) (*a*) Children realize that the other player will play naively as in step *i*.
(*b*) Children best respond to belief *ii* (*a*), lying as sender and trusting as receiver.

iii) (*a*) Children realize that the other player will reason and play as in step *ii*.
(*b*) Children best respond to belief *iii* (*a*), lying as sender and mistrusting as receiver.

It is possible that older children will iterate the above reasoning process beyond step *iii*, or even that they will eventually realize that there is no end to the process and also realize that their opponent will come to this conclusion. This reasoning may cause children to converge to randomization as in the fully rational benchmark to prevent their opponent from exploiting any systematic behavior that they would otherwise exhibit. In summary, the implications of the assumption that as they age, children will move along the sender–receiver reasoning sequence are as follows: *I*) Lying as the sender increases with age. This implication follows because lying occurs in steps *ii* and *iii*, but not in step *i*. *II*) Mistrusting as the receiver increases with age. This implication follows because mistrusting occurs at step *iii*, but not at step *i* or *ii*. *III*) A more subtle implication is that children start lying before they start mistrusting. This implication follows because lying starts at step *ii*, whereas mistrusting starts only at step *iii*. A moral, rather than strategic, theory of behavior, assuming that younger children are inhibited from selecting immoral actions, would make a prediction opposite from that in *III*, namely, that children mistrust at a younger age than that at which they lie. The reason is that one can only gain from lying if the opponent is trusting (i.e., moral), whereas one can gain from mistrusting only if the opponent is lying (i.e., immoral). In the latter case, the (perceived) immorality of the opponent may release one from one's own moral obligation, whereas in the former case, lying exploits the moral behavior of one's opponent.

For the stickers game, we conjecture that children, as they age, follow a reasoning sequence parallel to the one we conjectured for the sender–receiver game:

i) Children play naively, selecting five stickers, the maximum possible number.

ii) (*a*) Children realize that the other player will play naively as in step *i*.
(*b*) Children best respond to belief *ii* (*a*), selecting four stickers.

iii) (*a*) Children realize that the other player will reason and play as in step *ii*.
(*b*) Children best respond to belief *ii* (*b*), selecting three stickers.

iv) Children either continue iterating the above until they reach one sticker or "jump" to one sticker, realizing where the process is heading.

This sequence assumes that children begin with five stickers. Given the rules of the stickers game, choosing five stickers can never lead the child to win any stickers, regardless of the choice of the other player. In other words, the choice of five stickers is weakly dominated by any choice of strictly fewer stickers. This reasoning about the choice of five stickers does not require ToM but only requires knowledge of the rules of the game. Although this inference is likely to be too difficult for the youngest children, it is worth noting that the qualitative features of the stickers reasoning sequence would not change if we assumed that naive play at step *i* consists of choosing four, rather than five, stickers.

The implications of the assumption that, as they age, children will move along the stickers reasoning sequence are as follows: (*I′*) The number of stickers declines with age, and (*II′*) the number of stickers eventually converges to one.

Experimental Results

In all tables, the *P* value is indicated as follows: **P* < 0.1; ***P* < 0.05; ****P* < 0.01.

Total sample size was 69, and mean age was 5.99 y (SE = 0.19); the minimum age was 3.3 y, and the maximum age was 8.83 y. The age distribution of children was approximately uniform (*SI Appendix*, Fig. S1). Thirty-six children (52.1%) were male.

As a proxy for intelligence, we measured working memory by administering an *n*-backward task (26). We find that performance substantially improves with age. *SI Appendix*, Fig. S2 suggests that approximately one additional item is remembered with 2 additional years of age, until 7 y of age. We gave children a maximum number of five items to recall; thus, the flat performance for older children is, at least in part, due to this fact. *SI Appendix*, Table S1 estimates the gain per year as 0.45. There is no significant sex effect in *n*-backward task performance.

Stickers Game. Fig. 1 shows the change in the numbers of stickers played in first move as age progresses. The number, as predicted, steadily declines with age. Most of the youngest children (aged less than 4 y) play the maximum number of five stickers. The estimated local average [locally weighted scatterplot smoothing (LOWESS)] reaches the minimum of one sticker by the age of 6.5 y. Fig. 1 shows that in the age interval of 6.5–8 y, all children in our sample choose only one sticker. Table 1 shows an estimated average reduction of the number of stickers played of ~0.4 stickers per year of age. There is no significant sex effect or an interaction between sex and age. There is a large and significant effect of intelligence, or working memory, as measured by the score in the *n*-backward task: A higher score is associated with a smaller number of stickers played. The score in the regression reported in Table 1 is normalized between 0 and 1; thus, the total effect from maximum to minimum score is a reduction of 2.4 stickers over five stickers.

The choice of a single sticker by both players is the unique iterated dominance solution in the game. Our data are consistent with the hypothesis that children follow the iterated reasoning

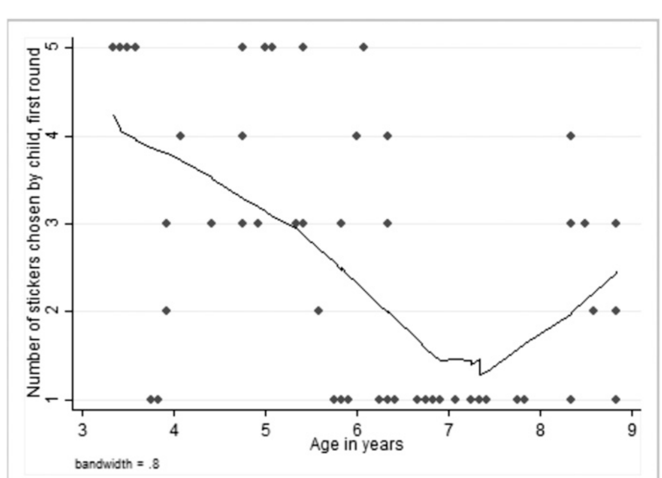

Fig. 1. First move in the stickers game. The number of stickers chosen by the child in the first move of the game and age of the child are shown. The continuous line is the locally weighted scatterplot smoothing (LOWESS).

Table 1. Number of stickers in the first move of the stickers game: Ordinary least squares

	1 B/SE	2 B/SE	3 B/SE	4 B/SE
Age	−0.510***	−0.403***	−2.465***	−1.597**
	(0.098)	(0.145)	(0.692)	(0.705)
Age squared			0.174***	0.128**
			(0.057)	(0.056)
Male		1.858	2.254*	1.807
		(1.215)	(1.150)	(1.139)
Male × age		−0.227	−0.306	−0.272
		(0.195)	(0.186)	(0.180)
n-backward score				−0.603***
				(0.171)
Constant	5.584***	4.714***	10.390***	8.942***
	(0.610)	(0.875)	(2.042)	(2.027)
N	67	7	67	65

The *n*-backward score is normalized in the unit interval. Age is in years. SEs are provided in parentheses. $*P < 0.1$; $**P < 0.05$; $***P < 0.01$.

process described above. Thus, an understanding of equilibrium behavior, or behavior implied by iterated dominance, grows with age and is complete by the age of 7 y. However, after the age of 8 y, there is a clear trend for children to play a larger number of stickers, with an average of ∼2.5 stickers at an age older than 8 y. A plausible explanation of this behavior is that children understand that the equilibrium behavior in the single-shot game leads to a zero payoff for both; thus, the cost of deviating from the equilibrium is small. In addition, the possibility of cooperation of some form between the two players (e.g., by randomly playing a number of stickers larger than one) makes the choice of a number of stickers larger than one preferable. Such cooperation could be profitable to both players, given that play is repeated over several rounds.

Sender–Receiver Game. In the sender–receiver game, children were randomly assigned to play either the sender role or the receiver role first. Fig. 2 reports the children's choices by age in years. To make the choices in sender and receiver roles easily comparable, we label the choice of deception (in the role of sender) and mistrust (in the role of receiver) as 1, and the opposite choice as 0. This variable can be interpreted as a deception/mistrust index. In Fig. 2 (*Upper*), children's first moves in whichever role (sender or receiver) the child played first are shown, and in Fig. 2 (*Lower*), the first move in the second role is shown; the first moves as sender (Fig. 2, *Left*) and as receiver (Fig. 2, *Right*) are separately shown.

In Fig. 2 (*Upper*, first move in first role), there is a clear upward trend, going from fully truthful behavior for senders and fully trusting behavior for receivers among the youngest children to deception and mistrust among the oldest. There is a slight trend downward for older ages in the receiver role, moving toward the equilibrium choice of 50%. Unlike the stickers game, in the sender–receiver game, iteration of inferences does not lead to an equilibrium; thus, behavior is unlikely to settle down. As children play, learning and observation of the other's behavior substantially modifies a child's own behavior. Fig. 2 shows that when children play the game in their second role, having already observed the other player when playing in the other role, their behavior is muted and less extreme. In particular, younger children deceive/mistrust more often and older children tell the truth/trust more often. In Fig. 2 (*Lower*, first move in second role), the trend is flatter and not clearly monotonic.

Table 2 presents the linear probability model for the first move in the first role, so coefficients indicate the effect size (logit regressions are reported in *SI Appendix*, Table S2). The regressions confirm a significant change of an increase of ∼13–15% in the fraction of deceiving and mistrusting behavior for every year of age. There is no significant effect of sex or intelligence. As age increases, children appear to take the first steps in the strategic reasoning process (see *Strategic Analysis of the Games* above) according to which one should deceive against a trusting opponent and one should mistrust against a deceptive opponent. There is a significant additional effect from the quadratic term for age, indicating that for older children, the overall effect of age is reduced [as is also clear in Fig. 2 (*Upper Right*) for children playing as receiver].

To estimate the effect of experience, we compare first-round choices when children played the game for the first and second times in opposite roles. The second time variable is equal to 1 in the second time observation, and is 0 for the first. Results for the random effect linear model are presented in Table 3 (the logit regression is shown in *SI Appendix*, Table S3). Age interacts with experience: Moving from the first to the second role leads to a

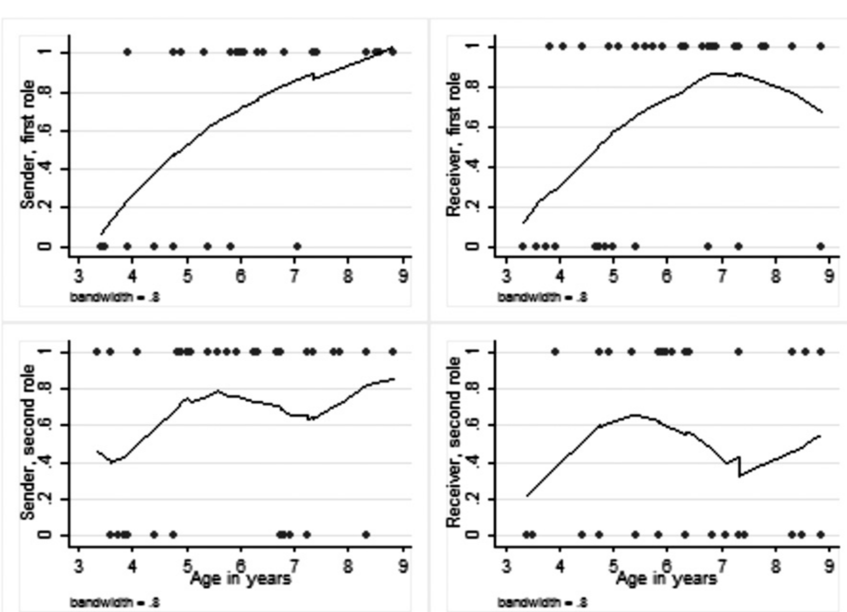

Fig. 2. First move in the sender–receiver game. The variable on the vertical axis is equal to 1 if the child deceives in the role of sender and mistrusts in the role of receiver; it is 0 otherwise. (*Upper*) First move in the rounds where the child played the game the first time and age (in years) of child. (Lower) First move in the rounds where the child played the game the second time. (*Left*) Move as sender. (*Right*) Move as receiver. The continuous line is the LOWESS.

Table 2. Deception and mistrust index in sender-receiver game: Ordinary least squares

	1 B/SE	2 B/SE	3 B/SE	4 B/SE
Age	0.141***	0.127***	0.648***	0.659**
	(0.032)	(0.046)	(0.228)	(0.260)
Age squared			−0.044**	−0.045**
			(0.019)	(0.021)
Male		−0.502	−0.643*	−0.716*
		(0.389)	(0.381)	(0.424)
n-backward score				−0.114
				(0.307)
Male × age		0.044	0.070	0.080
		(0.062)	(0.061)	(0.066)
Constant	−0.180	0.019	−1.417**	−1.362*
	(0.199)	(0.276)	(0.671)	(0.742)
N	67	67	67	64

The dependent variable is equal to 1 if the child deceives in the role of sender and mistrusts in the role of receiver; it is 0 otherwise. Only observations in which subjects played the game for the first time are considered. The *n*-backward score is normalized in the unit interval. Age is in years. SEs are provided in parentheses. $^*P < 0.1$; $^{**}P < 0.05$; $^{***}P < 0.01$.

greater movement toward deceive/mistrust among younger children than among older children. The younger children move away from their starting point of playing trust/truth most of the time, and, similarly, the older children move away from their starting point of deceive/mistrust. The overall effect among both younger and older children is that a movement from the first role to the second role causes children to be more deceptive/mistrusting.

Comparison with Adults. We compare children's behavior with that of adults (mean age = 21.4 y) in the same (although differently presented) games. In the adult version of the stickers game, almost all adults chose the number 1 or 2 (which, for the winner, was converted to dollars rather than stickers): In the first move, the mean choice was 1.81 (SE = 0.11). *SI Appendix*, Fig. S6 reports the number chosen in the first move for children younger than 7.5 y of age; then, for the others; and, finally, for the adults. Although no adult subject chose the number 5, a substantial number chose the number 2. Like adults, most children over 6.5 y of age chose either one or two stickers. In contrast, children younger than 6.5 y of age chose a larger number of stickers (Fig. 1). As we discuss in next section (*Discussion*), there is a difference between children younger than 8 y of age and children older than 8 y of age, but both of these groups are more similar to adults than children younger than 6.5 y of age.

In the adult version of the sender–receiver game, adults play the deceiving/mistrustful first move 43% of the time (n = 44), whereas children chose the deceiving/mistrustful first move 67% of the time (n = 67). For comparison, consider the LOWESS for children reported in Fig. 2. For children younger than 5.5 y of age, deceit/mistrust was chosen 38% of the time (n = 26); for children between 5.5 and 7.5 y of age, deceit/mistrust was chosen 85% of the time (n = 29); and for children older than 7.5 y of age, deceit/mistrust was chosen 80% of the time (n = 12). The youngest children chose the deceiving/mistrustful moves more rarely than adults. As age increases, the faction choosing the deceiving/mistrustful move grows until the older children choose this move more often than adults. Thus, the first-round behavior of the oldest and youngest children in the first role is more predictable than that of adults. In contrast, in the second, more experienced role, children's first-round behavior becomes more unpredictable and adult-like.

Discussion

Children appear to acquire SToM by the age of 6–7 y, as evidenced by a tendency to select sophisticated actions on the first move. In our data, this tendency holds in both games despite significant game theoretical differences. The stickers game has an incentive for cooperation, because taking turns in selecting a high number of stickers or randomizing in every round would lead to a higher payoff than the equilibrium of the one-shot game. In contrast, the sender–receiver game is a zero-sum game and provides no scope for cooperation even in repeated interaction. The stickers game can be solved by iterated dominance, and the sender–receiver game cannot; the unique mixed equilibrium of the sender–receiver game cannot be reached through iterated dominance. Our data are consistent with the view that as they age, children apply continually higher levels of recursive thinking, following the reasoning sequences presented in *Strategic Analysis of the Games* above, and conforming to implications (*I*–*II*) in the sender–receiver game and implications (*I′*–*II′*) in the stickers game. Such first-move sophistication shows that children apply strategic reasoning when given only information about incentives, before any feedback.

In the sender–receiver game, the process does not converge to a pure strategy Nash equilibrium because none exists; the mixed strategy equilibrium is much harder to understand and estimate precisely. Two findings speak to mixed strategy play in children. First, Table 3 shows that younger children move from playing truth/trust to sometimes playing deceive/mistrust and the older children move from playing deceive/mistrust to sometimes playing truth/trust. This pattern is consistent with the view that in the first role, children play a pure strategy, but with more experience in the second role, they move toward mixed strategies. Second, over multiple rounds, *SI Appendix*, Fig. S5 shows that play becomes more concentrated around equal probability randomization as children age. It is, however, difficult to separate randomization from alternation in response to the experimenter's last move. Senders are responsive to the experimenter's last move, whereas receivers are minimally so. Note that senders tend to choose the move that is not a best response to the experimenter's last move. For senders, responsiveness to the last move becomes more muted with age, suggesting an increase in randomization (all of the above results are shown in *SI Appendix*, Tables S7–S10).

A comparison of child and adult behavior reveals important similarities. First, in the adult analog of stickers, adults tend to select a small number (1 or 2), and children move from a large number to a small number with age. Second, adults do not exclusively choose 1, just as the oldest children (aged 8 y and older) do not exclusively choose 1. Third, in the sender–receiver game, adults do not overwhelmingly choose the sophisticated or naive action but rather mix their play, both on the first and later moves. Children do not appear to mix on their first move in their first role but tend to move toward mixing with experience as detailed above.

Table 3. Choice in first round, for the first and second times a child played the game: Random effects linear model, robust SEs

	1 B/SE	2 B/SE	3 B/SE	4 B/SE
Second time	−0.075	−0.075	0.609**	0.609**
	(0.075)	(0.075)	(0.282)	(0.283)
Age		0.084***	0.141***	0.591***
		(0.025)	(0.031)	(0.177)
Age × second time			−0.113**	−0.113**
			(0.047)	(0.047)
Age squared				−0.037**
				(0.014)
Constant	0.672***	0.162	−0.180	−1.451***
	(0.058)	(0.177)	(0.201)	(0.519)
N	134	134	134	134

The variable second time is the indicator of the trial in which the subject played the game for the second time. Age is in years. SEs are provided in parentheses. $^*P < 0.1$; $^{**}P < 0.05$; $^{***}P < 0.01$.

One developmental explanation for the gradually increasing sophistication is an increasing capacity to perform more steps of RT. Support for this explanation comes from the finding of Perner and Wimmer (9) that children demonstrate competence in higher order thinking at the age of 6–7 y, which is precisely when children demonstrate strategic sophistication in our experiments. Further support derives from the finding that strategic play in the stickers game was correlated with children's working memory performance. An alternative explanation is development in either the epistemic capacity to attribute beliefs to others in response to incentives or the practical capacity to transition from attributed beliefs to attributed actions. Indirect support for the failure of the practical capacity comes from experiments showing that children may recognize faulty beliefs in others early, although failing to use this knowledge to anticipate the other's surprise (27), actions such as false statements (28), or emotional reactions (29) correctly. Future experiments should discriminate the contributions of children's developing recursive reasoning, epistemic understanding, practical reasoning, and more general factors (e.g., behavioral control) (30, 31).

The analysis of stickers presents a particularly interesting finding. Between the ages of 6.5 and 8 y, all children select one sticker, the smallest possible number, in the first round. This choice is the most sophisticated move for a player who considers each round in isolation. Younger children progressively reduced the number of stickers with age. However, after the age of 8 y, children begin to select a larger number of stickers. This outcome suggests a possible cooperative motive if children view the game as a repeated interaction and recognize that both players can benefit if each gets a chance to win in different trials. Call this interpretation the collusive interpretation. A different feigning interpretation is that older children select a higher number of stickers to fool the opponent into thinking that they will do so again, allowing the child to undercut the opponent later. Both collusion and feigning posit that the oldest children become farsighted, viewing the game as a multiple-round affair, but collusion is cooperative and feigning is deceptive. The distinction between one-shot and repeated play is important in the game theoretical literature. Future research should focus on children's developing capacity to perceive interactions more broadly as repeated rather than one-shot.

There are several additional promising directions for future research. In experiments with two children playing against each other, we could examine whether the movement toward mixed strategies in later rounds of the sender–receiver game replicates. In a two-child version of stickers, the oldest children might actually achieve the higher payoff collusive outcomes that they seem to signal in the early rounds of our experiments. Another relevant experiment would have children play only one round (against an experimenter) and explicitly tell them they will play only one round. If a farsighted interpretation is correct, the oldest children should choose only one sticker in this version. This potential outcome, however, would not discriminate between the collusive and feigning interpretations, which are both farsighted. To separate these explanations, one could run another multiple-round experiment and ask children to explain their choices, with an interest in whether they offer "collusive" or "feigning" justifications.

This paper has explored an important understudied aspect of child development: children's ability to reason strategically, predicting others' behavior on the basis of knowledge of incentives, and adjusting to it. We focused on competition and incentives to mislead, two prevalent aspects of the environment to which children must adapt. Children demonstrate strategic sophistication at a surprisingly young age, and even appear to be able to think about interaction in a farsighted manner, considering the ramifications of the current game on the game to follow.

ACKNOWLEDGMENTS. M.K. is supported by National Science Foundation (NSF) Grant DLS-656007. A.R. is supported by NSF Grant SES-1357877.

1. Premack D, Woodruff G (1978) Does the chimpanzee have a theory of mind? *Behav Brain Sci* 1(4):515–526.
2. Wellman HM, Cross D, Watson J (2001) Meta-analysis of theory-of-mind development: The truth about false belief. *Child Dev* 72(3):655–684.
3. Wellman HM, Liu D (2004) Scaling of theory-of-mind tasks. *Child Dev* 75(2):523–541.
4. Apperly IA, Butterfill SA (2009) Do humans have two systems to track beliefs and belief-like states? *Psychol Rev* 116(4):953–970.
5. Eliot J, Lovell K, Dayton CM, McGrady B (1979) A further investigation of children's understanding of recursive thinking. *J Exp Child Psychol* 28(1):149–157.
6. Wynn K (1992) Children's acquisition of the number words and the counting system. *Cogn Psychol* 24(2):220–251.
7. Hauser MD, Chomsky N, Fitch WT (2002) The faculty of language: What is it, who has it, and how did it evolve? *Science* 298(5598):1569–1579.
8. Miller SA (2009) Children's understanding of second-order mental states. *Psychol Bull* 135(5):749–773.
9. Perner J, Wimmer H (1985) John *thinks* that mary *thinks* that...: Attribution of second-order beliefs by 5-to 10-year-old children. *J Exp Child Psychol* 39(3):437–471.
10. Murninghan JK, Saxon MS (1998) Ultimatum bargaining by children and adults. *J Econ Psychol* 19(4):415–445.
11. Sutter M (2007) Outcomes versus intentions: On the nature of fair behavior and its development with age. *J Econ Psychol* 28(1):69–78.
12. Sutter M, Kocher MG (2007) Trust and trustworthiness across different age groups. *Games Econ Behav* 59(2):364–382.
13. Harbaugh WT, Krause K, Berry T (2001) Garp for kids: On the development of rational choice behavior. *Am Econ Rev* 91(5):1539–1545.
14. Harbaugh WT, Krause K, Vesterlund L (2001) Are adults better behaved than children? Age, experience, and the endowment effect. *Econ Lett* 70(2):175–181.
15. Harbaugh WT, Krause K, Vesterlund L (2002) Risk attitudes of children and adults: Choices over small and large probability gains and losses. *Exp Econ* 5(1):53–84.
16. Sally D, Hill E (2006) The development of interpersonal strategy: Autism, theory-of-mind, cooperation and fairness. *J Econ Psychol* 27(1):73–97.
17. Perner J (1979) Young children's preoccupation with their own payoffs in strategic analysis of 2 × 2 games. *Dev Psychol* 15(2):204–213.
18. Shultz TR, Cloghesy K (1981) Development of recursive awareness of intention. *Dev Psychol* 17(4):465–471.
19. Carlson SM, Moses LJ, Hix HR (1998) The role of inhibitory processes in young children's difficulties with deception and false belief. *Child Dev* 69(3):672–691.
20. Russell J, Mauthner N, Sharpe S, Tidswell T (1991) The windows task as a measure of strategic deception in preschoolers and autistic subjects. *Br J Dev Psychol* 9(2): 331–349.
21. Couillard NL, Woodward AL (1999) Children's comprehension of deceptive points. *Br J Dev Psychol* 17(4):515–521.
22. Jaswal VK, Croft AC, Setia AR, Cole CA (2010) Young children have a specific, highly robust bias to trust testimony. *Psychol Sci* 21(10):1541–1547.
23. Koenig MA, Woodward A (2010) Sensitivity of 24-month-olds to the prior inaccuracy of the source: Possible mechanisms. *Dev Psychol* 46(4):815–826.
24. Sabbagh MA, Shafman D (2009) How children block learning from ignorant speakers. *Cognition* 112(3):415–422.
25. Mascaro O, Sperber D (2009) The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition* 112(3):367–380.
26. Alloway TP, Alloway RG (2010) Investigating the predictive roles of working memory and IQ in academic attainment. *J Exp Child Psychol* 106(1):20–29.
27. Hadwin J, Perner J (1991) Pleased and surprised: Children's cognitive theory of emotion. *Br J Dev Psychol* 9(2):215–234.
28. Koenig MA, Harris PL (2005) Preschoolers mistrust ignorant and inaccurate speakers. *Child Dev* 76(6):1261–1277.
29. Harris PL, Johnson CN, Hutton D, Andrews G, Cooke T (1989) Young children's theory of mind and emotion. *Cogn Emot* 3(4):379–400.
30. Hughes C (1998) Finding your marbles: Does preschoolers' strategic behavior predict later understanding of mind? *Dev Psychol* 34(6):1326–1339.
31. Steinbeis N, Bernhardt BC, Singer T (2012) Impulse control and underlying functions of the left DLPFC mediate age-related and age-independent individual differences in strategic social behavior. *Neuron* 73(5):1040–1051.