Received: 9 January 2019 | Revised: 24 September 2019 | Accepted: 14 January 2020
DOI: 10.1111/desc.12951

SHORT REPORT

Developmental Science WILEY

Children's understanding of habitual behaviour

Micah B. Goldwater[1] | **Samuel J. Gershman**[2] | **Caroline Moul**[1] | **Charles Ludowici**[1] | **Amy Burton**[1] | **Brittany Killer**[1] | **Rebecca-Lee Kuhnert**[3] | **Kate Ridgway**[1]

[1]University of Sydney, Camperdown, NSW, Australia
[2]Harvard University, Cambridge, MA, USA
[3]Black Dog Institute, Randwick, NSW, Australia

Correspondence
Micah B. Goldwater, School of Psychology, Brennan MacCallum A18, University of Sydney, NSW 2006, Australia.
Email: Micah.goldwater@sydney.edu.au

Abstract
Research into the development of Theory of Mind (ToM) has shown how children from a very early age infer other people's goals. However, human behaviour is sometimes driven not by plans to achieve goals, but by habits, which are formed over long periods of reinforcement. Habitual and goal-directed behaviours are often aligned with one another but can diverge when the optimal behavioural policy changes without being directly reinforced (thus specifically hobbling the habitual learning strategy). Unlike the flexibility of goal-directed behaviour, rigid habits can cause agents to persist in behaviour that is no longer adaptive. In the current study, all children predict agents will tend to behave consistently with their goals, but between the ages of 5 and 10, children showed an increasing understanding of how habits can cause agents to persistently take suboptimal actions. These findings stand out from the typical way the development of social reasoning is examined, which instead focuses on children's increasing appreciation of how others' beliefs or expectations affect how they will act in service of their goals. The current findings show that children also learn that under certain circumstances, people's actions are suboptimal despite potentially 'knowing better.'

KEYWORDS
goal-directed action, habits, Theory of Mind

1 | INTRODUCTION

As children develop, what changes in how they understand the behaviours of other people? Research into the development of children's 'Theory of Mind' (ToM i.e. their naïve beliefs about how other minds work and govern behaviour) has primarily focused on whether children predict others' actions considering their beliefs and goals or whether they 'egocentrically' base their predictions on their own mental states (e.g. Wellman, Cross, & Watson, 2001). In both cases, whether the goals are those of the child or of the other person, the assumption is that children predict actions will be in the service of goals—that they are part of rational plans (e.g. Baker, Saxe, & Tanenbaum, 2009; Gergely & Csibra, 2003). However, many behaviours are in fact not based on plans; human actions are routinely rooted in reflexive habits, not reflective plans (e.g. Daw, Niv, & Dayan, 2005; Dickinson & Balleine, 1994).

Decades of work on behavioural control suggest that planning and habitual behaviour originate from different cognitive and neural systems (see Dolan & Dayan, 2013 for review). The habitual system generally follows Thorndike's (1927) *law of effect*: rewarded behaviours will be repeated. More precisely, given particular states of an environment, actions that have led to rewards in those states will be repeated. That is, habits are strengthened over time when the same responses, given the same stimuli, lead to consistent positive outcomes. This is a generally adaptive strategy, and therefore habits tend to be aligned with an agent's plans. However, a habit can become maladaptive when a goal changes or when a change in the environment mandates a different set of behaviours, with no opportunity to learn from trial and error. In these circumstances, habits and plans will predict different behaviours. This leads to a basic question: in what circumstances do humans act flexibly in the service of goals or stick to their habits? Research has shown people are more likely to act habitually the more consistently behaviours have been repeated (e.g. Dickinson, Balleine, Watt, Gonzalez, & Boakes, 1995), when cognitive resources are depleted due to stress (Otto, Raio, Chiang, Phelps, & Daw, 2013) or working-memory load (Otto et al., 2013), and when the decision to act is made very quickly (among others causes, Kool & Botvinik, 2014). Importantly for the current research, younger children are more likely to act habitually than are older children, who show more goal-directed flexibility (Decker, Otto, Daw, & Hartley, 2016; Munakata, Snyder, & Chatham, 2012). This discussion of habits and goals prompts the question: do people typically take this central psychological distinction into account when they predict others' behaviour?

Gershman, Gerstenberg, Baker, and Cushman (2016) were the first to research whether adults' ToM reflects a distinction between habits and goals, and whether this folk psychological distinction is aligned with the psychological research reviewed above. Their experiment had the following basic structure: subjects learned about an agent's goals and action history, and then made predictions about the agent's future behaviours. The critical predictions were made when something about the context changed such that the previous behaviours were no longer the optimal way to achieve a goal. For example, in one task, subjects predicted which direction an individual would turn a doorknob at work. To open the door at work, the knob needs to be turned clockwise, but can get stuck when turned counter-clockwise. However, this individual's doors at home are opened when turned counter-clockwise. What will this person do at work? The adult subjects' predictions were generally aligned with the habit/plan distinction: They predicted the worker to be more likely to turn the knob in the same direction as his knobs at home when the worker had a longer history with the home knobs, when the worker had less time to make a decision, and when the worker was under cognitive load.

It is worth emphasizing that this work only examines one key aspect of plans—that they are conceived of in the service of achieving a goal, and support the flexible adjustment of behaviour in order to do so. Typically, when we use 'plan' this further implies that there are multiple steps in order to achieve it. The work of Gershman et al. (2016) did not focus on that aspect of plans, and so neither do we here, as explained below.

Gershman et al. (2016) work suggested that, unlike previous theories of ToM that focus on the understanding of rational plans, reasoners are sensitive to the conditions under which people deviate from flexible planning. The current research investigates the developmental origin of the habit/plan distinction in people's reasoning about others' behaviours. Previous research in children's ToM has focused on inferences about beliefs and goals (e.g. Gergely & Csibra, 2003), but were not designed to probe whether children use the habit/plan distinction to make these inferences. One hypothesis is that because children's ability to override habits develops slowly (e.g. Munakata et al., 2012), perhaps they will extend knowledge of their own behaviour to others, and predict an agent would repeat behaviours that have been frequently repeated in the past. An alternative hypothesis is that dual-system theory of mind is more complex than a single-system (pure planning) theory, and therefore children require more experience to develop this theory, much like in other domains of cognition where theory complexity grows with age (e.g. Carey, 2011). This hypothesis predicts that younger children will show less understanding of habits than older children.

Research Highlights

- The majority of research on children's social reasoning concerns how children understand the relationship between a persons' goals, knowledge or beliefs, and their behaviour.
- We examined the development of children's appreciation of how people often act out of habit, even when it is in conflict with their goals.
- We sampled 250+ children, aged 5 to 10 who generally predicted that people would act consistently with their goals.
- However, the appreciation that sometimes people act habitually first appeared at 7, and continued to increase throughout middle childhood.

Related to work on how children infer others' goals, there is a distinct line of research that examines how children predict the behaviour of others based on inferring their traits or preferences. For example, Boseovski and Lee (Study 1; 2006) show that children as young as three predict an individual who they have seen be nice to others repeatedly will continue to do so. Ma and Xu (2011) show that children as young as 2 years of age infer preferences from another's consistent choices, and predict these individuals will continue to choose in alignment with their preferences. Perhaps then, children's early ability to predict consistent behaviour from inferred traits and preferences suggests that we will likewise observe the youngest children in our sample predict that people will persist in their frequently repeated behaviours.

In the current study, we adapted one of Gershman et al.'s tasks into child-appropriate vignettes and asked children to predict the behaviour of the agents in the vignettes. Each vignette described some behaviour of the agent, either performed twice (the Non-Habit condition), or repeated regularly for a long period of time (e.g. 'every day for a year'; the Habit condition, see Figure 1 for task design). Then, the vignettes introduced a context-shift that rendered the previous behaviour less efficient for achieving the agent's goal, and we asked the child whether the agent would change to this new optimal behaviour, or stick to what they had done in the past. If a child reported that the agent was less likely to change behaviour in the Habit condition, we posited that the child has a naïve theory of habits.

The children who participated in the research ranged from 5 to 10 years of age. We sought a wide age range because we did not have strong expectations about when understanding of habitual behaviour would emerge. We started at age 5 because children of this age show a robust ability to make explicit predictions that agents will

FIGURE 1 Task flow

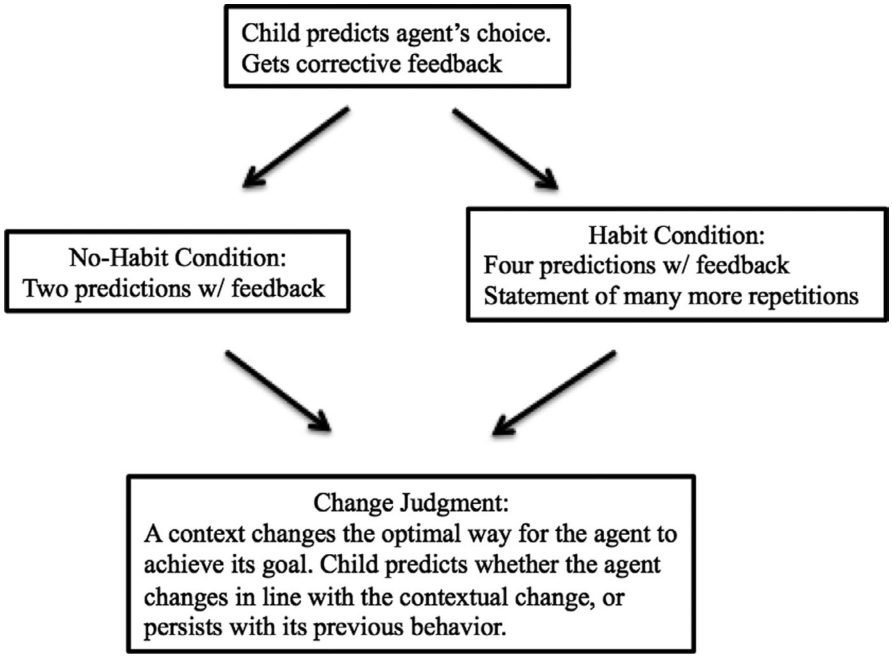

act in accordance with their beliefs and goals, and not egocentrically base their predictions on their own beliefs (Wellman et al., 2001). Likewise, children of this age readily predict consistent behaviours based on traits and preferences (see above).

2 | METHODS

2.1 | Subjects

Two hundred and sixty-two children (135 female) between the ages of 5 and 10 years participated in the study (30 five-year-olds; 49 six-year-olds; 46 seven-year-olds; 50 eight-year-olds; 49 nine-year-olds; 38 ten-year-olds). Seventeen children were removed from analysis after failing to answer at least three of five questions correctly querying their memory for the content of the experimental materials (see below), leaving 245. The majority of children either were accurate on four of five (49) or five of five (180); see Table 1.[1]

The children were recruited by students in an advanced undergraduate developmental psychology class at the University of Sydney in 2015 and 2016. Specifically, for course credit every student recruited and conducted the study with a single child. The students were fully trained in recruitment ethics and obtained consent from their child's legal guardian, in addition to giving their own consent for the data they collected to be used for purposes outside the class (such as inclusion in this manuscript). These student-researchers were trained in the methods of the study but were naïve to hypotheses until after data collection. Their primary writing assignment for the course was on the interpretation of the aggregated results.

This research was approved by the University of Sydney Human Research Ethics Committee, #2016/601, titled 'Children's understanding of goals and habits: How do children predict other people's behaviours?'

TABLE 1 Proportion of Change Judgments wherein the agent was predicted to repeat its previous actions as function of memory question accuracy

	3 out of 5 correct	4 out of 5 correct	5 out of 5 correct
Habit	0.34	0.41	0.40
Non-Habit	0.38	0.35	0.31
n	16	49	180
Mean age (years)	7.44	7.22	7.86

2.2 | Materials

The experimental materials consisted of four short vignettes modelled after Gershman and colleagues' materials. Each vignette set up a choice for the main character (the agent) in the story. For example:

> *Sally sometimes has trouble sleeping, but reading helps her get to sleep. One day, her mum bought a lamp to put next to her bed so she could read if she woke up in the middle of the night. The light is on the left side of her bed. The next night, Sally woke up in the middle of the night, and thought she would like to turn on the light and read. Which direction did she reach for the lamp, to the left side of the bed or to the right side of the bed?*

These vignettes were accompanied by pictures (see Appendix B) and the experimenter pointed to the relevant objects and characters, and they were used to help elicit child responses. Children guessed the answer to the first question in each story, and the experimenter would either confirm or correct their answer (e.g. in this case, the answer was the left side). The stories continued after the initial situation, setting up the same choice, for example:

> *The next night, Sally woke up in the middle of the night, and thought she would like to turn on the light and read. Which direction did she reach for the lamp this time, to the left side of the bed or to the right side of the bed?*

To establish a consistent action history, the agent repeated the same action, and the experimenter continued to give corrective feedback based on the children's guesses. For two of the four stories, the action was performed a total of twice (the No-Habit condition). In the other two stories (counter-balanced across subjects), the child guessed two more times (for a total of four; the Habit condition). We did not expect children to make accurate predictions on this first question, so the feedback was important to help establish that the agents would behave consistently. It is possible children would not see any intrinsic differences in value between the choices, and so this feedback shows behavioural consistency of the agents regardless of the children's own perceptions. Children were guessing on this first trial (53% accuracy in both conditions), but quickly learned from the feedback (70% accuracy on the 2nd prediction in both conditions; 85% accurate on the 3rd predictions in the Habit condition; 87% accurate on the 4th prediction in the Habit condition).

After the 4th prediction in the Habit condition, the child then heard about how the agent continued this action, for example:

> *Sally found that reading helped her to fall back to sleep and so continued to do this each time she woke up. She woke up in the middle of the night very frequently. Many times a week, every week, for a whole year! Every time, Sally turned the light on to read. Every time Sally reached to the left side of the bed to switch on her light*

Both additional guesses and this verbal description were used to make a salient distinction between the No-Habit and Habit conditions. Our working hypothesis was that the additional guesses would help make the distinction more salient than a described distinction alone, but only had an additional two guesses to avoid the procedure from seeming overly repetitive and lose the children's attention (and ingrain the children themselves in a habit of predicting the same response).

The stories then set up the primary question of the task, 'the Change-Judgment,' by introducing a change in context or new option for the agent and asked the children if the agent would change to this new option, or stay with the same choice, for example:

'Sally grew out of her bed, and needed a bigger one. But the new bed wasn't going to arrive for another night, so Sally was going to sleep in the guest bedroom where her grandma would sleep when she visited. Sally woke up in the middle of the night and wanted to read. Unlike in her own room, the lamp was on the right side of the bed in this room. Which direction did she reach for the lamp, to the left side of the bed or to the right side of the bed?' Children were told what the agent did and a further elaboration as feedback. For two of the vignettes, the agents stayed the course, and for two of the vignettes they changed their behaviour (one each for the Habit and Non-Habit condition, respectively).

Last, there were memory questions about key pieces of information that would have to be accurately remembered for the response to the change judgment to be meaningful. The number varied across stories because they varied in how many pieces of information were seen as critical. In the lamp story, there were two (there were a total of five memory questions for the four vignettes, see Appendix A):

> *On which side of the bed is the lamp in the guest bedroom that Sally is sleeping in this one night?*
>
> *On which side of the bed is the lamp in Sally's bedroom where she usually sleeps?*

The other three vignettes followed the same form, see Appendix A for details.[2]

2.3 | Procedure

After asking the child if he/she was willing to listen to stories and answer some questions about them, the experimenter led the child through each story, asking questions and recording responses. The experimenter held up the relevant pictures at the relevant time of the story and pointed to them (again, see Appendix B). The experiment lasted for approximately 20 minutes.

3 | RESULTS

Before analysing the primary Change-Judgment data, we confirmed that children of all ages were correctly answering which choice the agent would make by the 4th trial of the Habit condition (before the context changed which choice was optimal). Children of all ages learned from the corrective feedback that the agent would behave consistently (see Table 2 for the whole pattern). Children responded accurately on 87% of the 4th trials, and this did not differ across ages. This ensured that any age differences in how children predicted the agent would behave after the context change cannot be attributed to age differences in how children predicted the agent would behave before the context change.

Overall, children predicted the agents would change their behaviour aligned with the contextual change. That is, they predicted the agent would persist with their previous behaviour only 36% of the time. Importantly, however, this proportion changes with condition and age, see Table 3 and below.

TABLE 2 Proportion of accurate predictions for the agents' behaviour before the context change, as a response to corrective feedback

	Prediction 1	Prediction 2	Prediction 3	Prediction 4
All ages				
Habit	0.53	0.68	0.85	0.87
Non-Habit	0.52	0.72		
5- and 6-year-olds				
Habit	0.57	0.71	0.86	0.87
Non-Habit	0.51	0.70		
7- and 8-year-olds				
Habit	0.52	0.65	0.82	0.84
Non-Habit	0.54	0.70		
9- and 10-year-olds				
Habit	0.51	0.68	0.86	0.92
Non-Habit	0.51	0.77		

TABLE 3 Proportion of Change Judgments wherein the agent was predicted to repeat its previous actions

	5- and 6-year-olds	7- and 8-year-olds	9- and 10-year-olds
Habit	0.28	0.44	0.46
Non-Habit	0.31	0.30	0.36

3.1 | Model selection

There were two vignettes in the Habit condition, and two vignettes in the Non-Habit condition, and so each individual child predicted whether the agent repeated the same action in the Change-Judgment either 0, 1 or 2 times in each condition. To analyse these data, we used Christensen's (2015) Ordinal package for R (R Core Team, 2017). The package fits a model of the form:

$$\mathrm{logit}\,(P\,(Y_i \le j)) = \mathrm{logit}\,[\theta_j - \beta_{age}\,(\mathrm{ChildAgeCentered}) - \beta_{habit}\,(\mathrm{Habit}) - u\,(\mathrm{Subject}_i)]$$

In this model, an 'observation' is the pair of change-judgments for each condition. This model states that the logit of the probability that the *i*th observation falls in the *j*th category of the outcome variable or below is a function of:

1. θ_j: The intercept for response category *j*.
2. β_{age} and β_{habit}: Effects for age and habit. These are equal for the different levels of *j*. In this parameterization of the model, $\beta > 0$ means that the odds of responses at the higher end of the scale increase with this variable.
3. u(Subject$_i$): A random effect for the *i*th participants. u(Subject$_i$) ~ $N(0, \sigma^2_1)$. The standard deviation is estimated by the software.

In plain terms, the model tested how age and experimental condition affected judgments about action repetition despite a contextual change indicating the action was no longer the optimal way to achieve the agent's goal.

To show that this model was the appropriate way to analyse the data, we compare the fit to other potential models. It is ideal to test the maximal mixed effects structure because maximal mixed-effects models offer the best generalizability by accounting for all sources of dependence in a design (Barr, Levy, Scheepers, & Tily, 2013). Because observations are clustered by subject in our design, the maximal mixed effects structure includes random intercepts and random effects for Habit (condition) by subject. Each observation is the pair of change-judgments for each condition (respectively), and so we end up with two observations and two random effects (slope and intercept) for each subject, which isn't enough to estimate the model (but is a result of the nature and pragmatics of collecting data with this population). We adjusted to this by dropping the random effect of the Habit condition, which enabled model estimation. We then tested models that were all combinations of the possible fixed effects, interactions and random slopes. The best-fitting model is the one described above, see Table 4 for the complete list.

3.2 | Statistical analysis with the best-fitting model

As can be seen in Table 5, the child's age, the condition and an interaction between age and condition were all significant predictors in our selected model. For children of mean age, the odds of predicting the agent will repeat the actions *j* or less times in the Non-Habit stories is 1.59 times that of the Habit stories. The odds of predicting that the agent will repeat the previous action *j* or less times increase by 1.32 per year for children in the Habit condition but just 1.04 in the Non-Habit condition (which is not significantly different than 1, meaning no change).

This interaction effect is important because it shows that the sensitivity to the distinction between the Habit and Non-Habit condition increases with age. The logit model predictions were calculated using this equation from Agresti (1996, p. 183)

$$P(Yi \le j) = \frac{e^{\theta j - \beta age(ChildAgeCentered) - \beta habit(Habit) - \beta interaction(Habit^{*}Age) - \mu(Subject i)}}{1 + e^{\theta j - \beta age(ChildAgeCentered) - \beta habit(Habit) - \beta interaction(Habit^{*}Age) - \mu(Subject i)}}.$$

TABLE 4 Comparing model fits with Akaike information criterion (AIC). Smaller AIC values indicates better fits

Model	AIC
Main effect of habit, no random intercepts	981.29248
Main effect of age, no random intercepts	980.780919
Main effects of habit and age, no random intercepts	976.332264
Main effects and interaction of habit, no random intercepts	974.474868
Main effect of habit with random intercepts	980.091297
Main effect of age with random intercepts	980.779583
Main effects of habit and age with random intercepts	975.848885
Main effects and interaction of habit and age with random intercepts	973.845472

TABLE 5 Ordinal regression coefficients for the Change-Judgment

	Estimate	Std. Error	z value	Pr(>|z|)
Habit/Non-Habit	−0.459	0.179	−2.571	0.010*
Child (age-centred) 0.270	0.086	3.151	0.002**	
Habit/Non-Habit X age	−0.230	0.116	−1.993	0.046*

*$p < .05$,
**$p < .01$.

We visualize how the model predictions fit the data across four figures. Figure 2 shows the model predictions for children, as a function of their age, indicate whether the agent would persist with its past behaviour 0, 1 or 2 (out of 2) times, respectively. Then to show how the model fits the actual data, Figures 3-5 show children's indications of 0, 1 or 2 times, respectively. We split these data fits across three figures to avoid visual cluttering.

3.3 | Potential item-differences

We next consider some potential differences between our four items. An anonymous reviewer suggested that the vignettes about the walking the dog, and switching off the lamp concerned behavioural or motor routines, while the vignettes about sports and playing with dolls were more akin to preferences than 'just' motor responses. Table 6 shows that both pairs of items show the same key effect—children predicted the agent would persist more frequently in the Habit condition than the Non-Habit condition. Further when broken down by age, we see the pattern is quite similar to the pattern in Table 3 which includes all four items together. The 5- and 6-year-olds showed no increased frequency of 'persist' judgments in the Habit condition for either pair of vignettes, while the 7–10 year-old children showed more frequent persist judgments in the Habit condition for both pairs of vignettes.

We note the one apparent deviation from this overall pattern is the elevated rate that the 9- and 10-year-olds predicted the agent would persist in the Non-Habit condition (0.46 of the judgments) for the 'behaviour' items. So we examined whether there was something non-random about how these items for this age group deviated from the overall pattern with two analyses. The first compared the relative rates of persist judgments for the Habit and Non-Habit conditions elicited by the 'behaviour' vignettes for the 9- and 10-year-olds to the relative rates elicited by the 'preference' vignettes for the same age group. This analysis revealed no reliable differences, Yates χ^2 (1) = 1.75, p = .19. The second analysis again examined the relative rates of persist judgments for the Habit and Non-Habit conditions elicited by the 'behaviour' vignettes for the 9- and 10-year-olds, but now compared them to the same items for the 7- and 8-year-olds. Again, there was no evidence of a reliable difference, Yates χ^2 (1) = 0.32, p = .57. In sum, we find no evidence that there are meaningful differences between these pairs of items, and are confident that our primary findings are general to the entire set of vignettes.

4 | DISCUSSION

This experiment was the first to examine whether children have a folk theory of habitual behaviour: that is, whether they think an agent may be more likely to persist in a previous behaviour despite a contextual change rendering a different action optimal. Similar to the adult subjects in Gershman et al. (2016), on the whole the children were sensitive to the action history, and were less likely to predict the agent would change to the new optimal behaviour when the agent had persisted in the previous action more consistently over a longer period of time. While there was a main effect of the Habit condition, crucially there was also an interaction with age. The pattern displayed in Figures 2-5 suggests that children younger than 7 years did not show sensitivity to the distinction between conditions, and the sensitivity (according to the model), continued to increase through age 10. Crucially, children of all ages learned from feedback to help establish the agent's action history.

Consistent with standard theories of ToM (e.g. Baker et al., 2009; Gergely & Csibra, 2003), children appear to generally predict that agents will make the optimal choice by changing their behaviours depending on the context. However, these standard theories do not incorporate a folk theory of habits, as children increasingly showed with age. Increasing sensitivity to the habit/plan distinction over the course of childhood is consistent with the observation that folk theories become increasingly complex with age in other cognitive domains (such as biology; Carey, 2011). Younger children have a very early assumption that behaviour is planning-based, and then throughout middle childhood, children increasingly consider a more complex theory of behaviour that distinguishes habits from plans.

We note that because the current work was the first in this domain, the tasks we used are certainly not the most sensitive *possible* tasks to show children's understanding of habits.[3] Often, it appears children do not show understanding of some distinction in early research on some developmental ability, but then building on this initial research, a more sensitive task is created to show that younger children do indeed have some understanding (e.g. Krettenauer, Malti, & Sokol, 2008 within the domain of moral judgment; Low

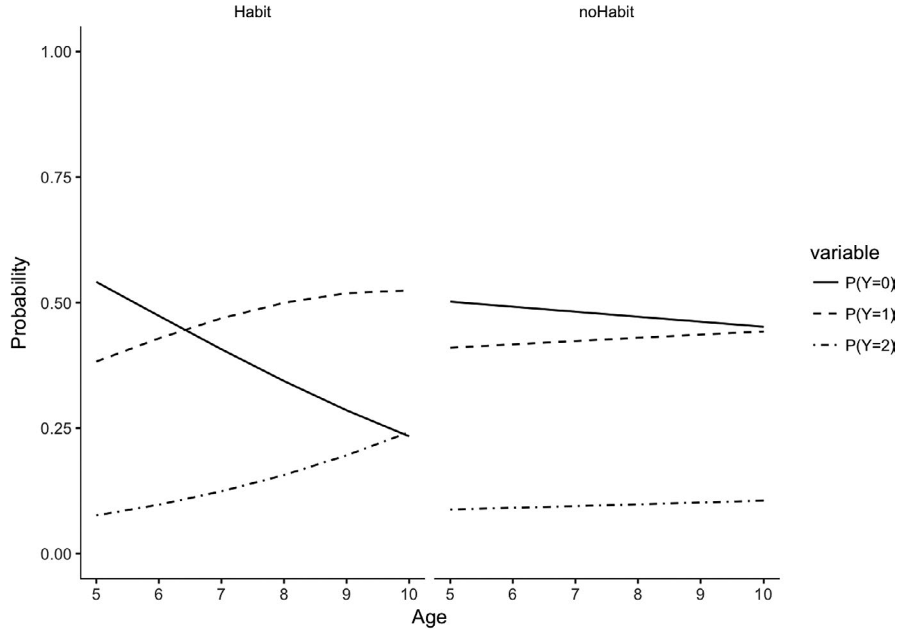

FIGURE 2 Model predictions for the probability of predicting the agent will repeat their previous actions 0, 1 or 2 times in the Habit and Non-Habit conditions across ages, respectively. In the Habit condition, the probability of predicting a repetition 0 times decreases with age, while the probability of predicting a repetition 1 or 2 times increases. The Non-Habit condition shows no such change with age

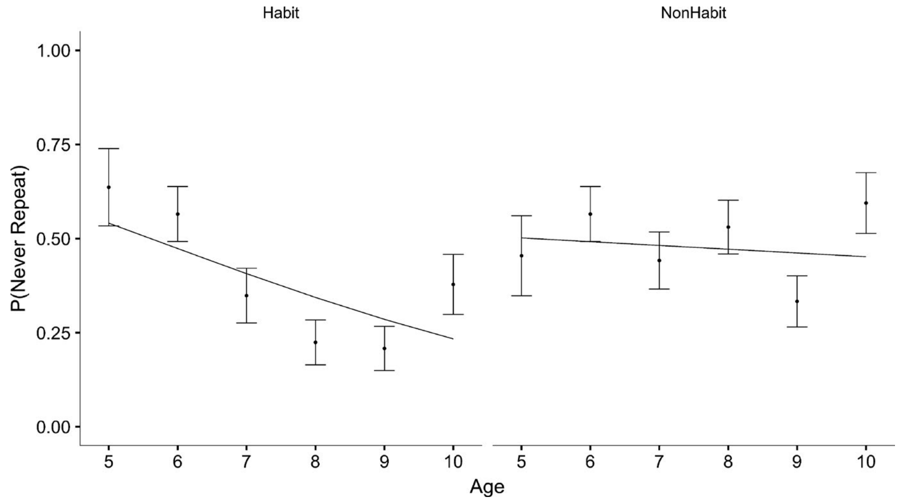

FIGURE 3 Model fits to data of predicting the agent would repeat past behaviours in zero of the two Change Judgments. The points represent the empirical proportions of these responses (with standard errors of the proportion). The Habit condition showed a decrease with age, while the Non-Habit condition showed no such change

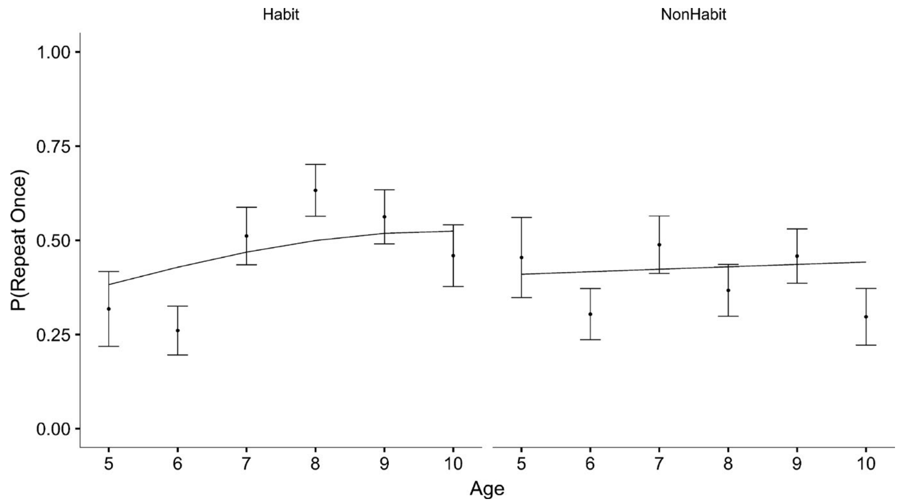

FIGURE 4 Model fits to data of predicting the agent would repeat past behaviours one of the two Change Judgments. The points represent the empirical proportions of these responses (with standard errors of the proportion). The Habit condition showed an increase with age, while the Non-Habit condition showed no such change

& Perner, 2012 in ToM; Fisher, 2002 within syntactic development). We do not make any claims about what children at any one age *absolutely* understands or does not. However, within a single task environment, age differences reveal that the robustness (i.e. the ease of detecting the child's sensitivity) of that understanding changes with development, and suggests the developmental

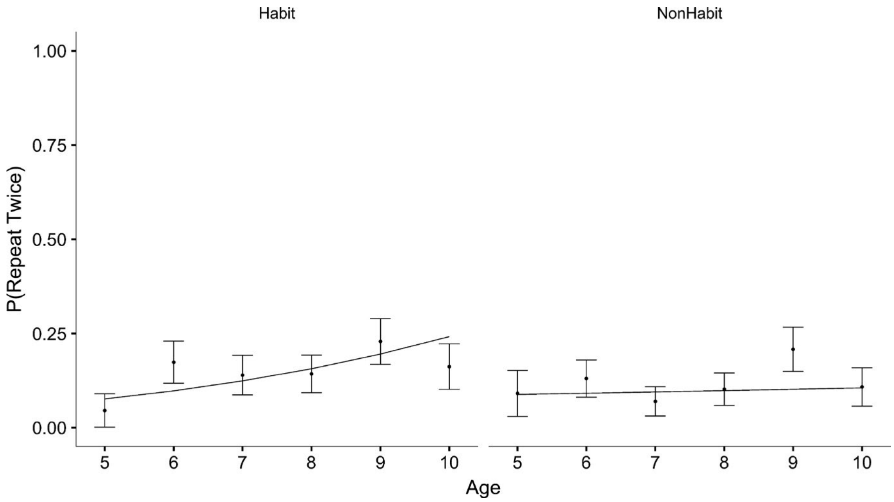

FIGURE 5 Model fits to data of predicting the agent would repeat past behaviours two of the two Change Judgments. The points represent the empirical proportions of these responses (with standard errors of the proportion). The Habit condition showed an increase with age, while the Non-Habit condition showed no such change

TABLE 6 Proportion of Change Judgments wherein the agent was predicted to repeat its previous actions, further broken down by items that concern behaviours versus preferences

	Behaviours		Preferences
All ages			
Habit	0.38		0.41
Non-Habit	0.33		0.31
	5- and 6-year-olds	**7- and 8-year-olds**	**9- and 10-year-olds**
Behaviours			
Habit	0.22	0.42	0.47
Non-Habit	0.21	0.30	0.46
Preferences			
Habit	0.32	0.43	0.45
Non-Habit	0.38	0.30	0.26

sequence. The sequence is still relevant even if younger children are later shown to have a more sophisticated understanding than the initial research could reveal. Thus, even if a task is developed to show that younger children are sensitive to the habit/plan distinction, this would not change the current interpretation that a more complex theory of behaviour emerges with development (without further evidence to re-explain the current pattern).

For the current task, this lack of nuance in the ToM of the younger children is quite different from the classic egocentric ToM mistakes preschool children often make, given how difficult children below 7 years may find it to override habitual behaviours in novel contexts (Decker et al., 2016; Munakata et al., 2012). That is, young children (at the least in this experimental context) do not extend what governs their own behaviour (in regards to the effects of action history) such that they may predict similar patterns in others. Whether this is rooted in a lack of understanding of their own behaviour is a topic for future research.

We further note that future research should aim to distinguish the differences between thinking about repeating behaviour due to habit from repeating behaviour due to traits (e.g. Boseovski & Lee, 2006) or preferences (e.g. Ma & Xu, 2011). In many ways, behavioural predictions may be quite similar for all three because each may be the cause of a suboptimal action (in relations to achieving some goal), but the traits, preferences and habits differ in the mental state attribution that is the cause of that suboptimal action. Perhaps distinguishing these in future work will require eliciting explanations from children for why they made their behavioural predictions.[4]

For interpreting the current results, we argued that if children did infer a trait or preference from the behavioural repetition, then even our youngest children should have been more likely to predict the agent would persist, as the above cited research showed 2- and 3-year-olds use behavioural frequency to infer a trait or preference, and then predict future consistency. However, we found just the opposite; children overall predicted the agents would change their behaviour. Further, the high rates of predicting behaviour change in the Habit condition for our youngest children held even for vignettes that were seemingly describing preferences, rather than just behaviour sequences. Though to be clear, we do not make any strong conclusions about different developmental patterns between understanding habits, traits and preferences because our study was not directly designed to distinguish amongst them.

We do note however, that Figures 3-5 appear to show that the peak of sensitivity to the habit plan distinction is at 8-year-olds, and then the 9–10-year-olds appear to regress somewhat. Currently, there is no evidence that this apparent regression is not just random fluctuation. First, the model fits the development change monotonically. Second, the numerical distinction between the middle age group and the oldest age group appears to be entirely in the 'behaviour' vignettes (see Table 6), and the above analysis showed no evidence that this age difference is non-random. Of course it is quite possible that the understanding of habits and plans continues to change throughout childhood past 8 years old, but to understand whether and how that is true requires future research.

This paper represents the first direct attempt to measure children's understanding of the distinction between habitual versus planned action, and it has provided evidence that this understanding increases from 5 to 10 years of age. The current findings prompt future research into whether we can teach children this distinction earlier, as many persistent behavioural problems (such as aggression and binge eating;

e.g. Hortsmann et al., 2015; Watson & de Wit, 2018) can be caused by maladaptive habits. How would children's understanding of habits affect their response to interventions aimed at changing these behaviours? A critical question for future research is to understand how a child's folk theory of habits interacts with the developing balance of power between plans and habits in the child's own behaviour.

DATA AVAILABILITY STATEMENT

Raw data are accessible here: https://osf.io/t3qgx/

ORCID

Micah B. Goldwater https://orcid.org/0000-0001-8052-9497

ENDNOTES

[1] Table 1 also shows the key DV predicting persistent behaviour in the "change judgments," and how that is affected by condition (see explanation below). We note that there were similar judgments rates between children who scored either 4 or 5 out of 5. Because there were so few children who scored 3 out of 5 correct (n = 16), it does not seem appropriate to conduct inferential statistical analysis on the relationship between memory accuracy and judgment.

[2] There were further elaborations of these vignettes, and then additional vignettes that were designed to answer distinct research questions about children's social and affective development that go beyond the scope of the current paper, and were only presented to a subset of the children analysed in the current data set.

[3] For example, Gerhman et al. (2016) shows that explicitly using the term "habit" increases the effect size of the Habit manipulation for adults. Perhaps this would help our youngest children. Another possibility is using a more "implicit" measure of predicting other's behaviour, such as with eye-movements, as discussed in Low & Perner, (2012).

[4] We did not record any explanations spontaneously offered by the children.

[5] Here, there is no intrinsic reason for a preference, we demonstrate the preference with the feedback to their prediction. The appendices show a schematic of the scripts to distinguish the vignettes from each other, but there are enriching elements via the questions and feedback.

REFERENCES

Agresti, A. (1996). *An introduction to categorical data analysis* (Vol. *135*). New York: Wiley.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329–349. https://doi.org/10.1016/j.cognition.2009.07.005

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Boseovski, J. J., & Lee, K. (2006). Children's use of frequency information for trait categorization and behavioural prediction. *Developmental Psychology*, *42*(3), 500–513.

Carey, S. (2011). Précis of the origin of concepts. *Behavioral and Brain Sciences*, *34*(03), 113–124. https://doi.org/10.1017/s0140525x10000919

Christensen, R. H. B. (2015). Ordinal—Regression models for ordinal data version 2015.6-28. Retrieved from https://CRAN.R-project.org/package=ordinal.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, *8*(12), 1704–1711.

Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychological Science*, *27*(6), 848–858. https://doi.org/10.1177/0956797616639301

Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. *Animal Learning and Behavior*, *22*(1), 1–18. https://doi.org/10.3758/BF03199951

Dickinson, A., Balleine, B., Watt, A., Gonzalez, F., & Boakes, R. A. (1995). Motivational control after extended instrumental training. *Animal Learning and Behavior*, *23*(2), 197–206. https://doi.org/10.3758/BF03199935

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*(2), 312–325. https://doi.org/10.1016/j.neuron.2013.09.007

Fisher, C. (2002). The role of abstract syntactic knowledge in language acquisition: A reply to. *Cognition*, *82*(3), 259–278. https://doi.org/10.1016/s0010-0277(01)00159-7

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naıve theory of rational action. *Trends in Cognitive Sciences*, *7*(7), 287–292.

Gershman, S. J., Gerstenberg, T., Baker, C. L., & Cushman, F. A. (2016). Plans, habits, and theory of mind. *PLoS ONE*, *11*(9), e0162246. https://doi.org/10.1371/journal.pone.0162246

Horstmann, A., Dietrich, A., Mathar, D., Pössel, M., Villringer, A., & Neumann, J. (2015). Slave to habit? Obesity is associated with decreased behavioural sensitivity to reward devaluation. *Appetite*, *87*, 175–183. https://doi.org/10.1016/j.appet.2014.12.212

Krettenauer, T., Malti, T., & Sokol, B. W. (2008). The development of moral emotion expectancies and the happy victimizer phenomenon: A critical review of theory and application. *International Journal of Developmental Science*, *2*(3), 221–235.

Kool, W., & Botvinick, M. (2014). A labor/leisure tradeoff in cognitive control. *Journal of Experimental Psychology: General*, *143*(1), 131. https://doi.org/10.1037/a0031048

Low, J., & Perner, J. (2012). Implicit and explicit theory of mind: State of the art. *British Journal of Developmental Psychology*, *30*(1), 1–13. https://doi.org/10.1111/j.2044-835X.2011.02074.x

Ma, L., & Xu, F. (2011). Young children's use of statistical sampling evidence to infer the subjectivity of preferences. *Cognition*, *120*(3), 403–411. https://doi.org/10.1016/j.cognition.2011.02.003

Munakata, Y., Snyder, H. R., & Chatham, C. H. (2012). Developing cognitive control: Three key transitions. *Current Directions in Psychological Science*, *21*(2), 71–77. https://doi.org/10.1177/0963721412436807

Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(52), 20941–20946. https://doi.org/10.1073/pnas.1312011110

R Core Team. (2017). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/.

Thorndike, E. L. (1927). The law of effect. *The American Journal of Psychology*, *39*(1/4), 212–222. https://doi.org/10.2307/1415413

Watson, P., & de Wit, S. (2018). Current limits of experimental research into habits and future directions. *Current Opinion in Behavioral Sciences*, *20*, 33–39. https://doi.org/10.1016/j.cobeha.2017.09.012

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, *72*(3), 655–684. https://doi.org/10.1111/1467-8624.00304

How to cite this article: Goldwater MB, Gershman SJ, Moul C, et al. Children's understanding of habitual behaviour. *Dev Sci*. 2020;23:e12951. https://doi.org/10.1111/desc.12951

APPENDIX A: Vignettes

STORY 1: WALKING A DOG TO SCHOOL

George's mum often brings the family dog to work. However, one day George's mum had to work late, but the dog needed to get back home before George's mum could bring him home. So George decided to volunteer to get the dog from his mum's work and walk him back home.

Here's a picture of his neighbourhood, where his mum's work is, and where his family lives. That's George's house. These are the three roads connecting work and home. As you can see, one road is closed down because workers are fixing it, and then there are two more roads. Which of the two roads do you think George will take to walk his dog back home?

Change-Judgment

George's mum had to stay late again. But now, all the construction was done, and that other road opened up. Now, which route will he choose? Will he stick with what he has done, or choose the new route?

STORY 2: PLAYING SPORTS

Kieran loves to play sports, and especially with his older brother Luke. Luke's favourite sport is basketball. But, Luke had to go on a trip, so he wasn't around to play with Kieran. When Kieran got home from school, he could either practice basketball on his own, or play soccer with his dad. Which do you think he did, did he play basketball on his own, or did he play soccer with his dad?

Change-Judgment

Then, Luke came home. He asked Kieran 'What sport do you want to play, basketball, or soccer?' Which do you think Kieran played?

Memory question

What is Luke's favourite sport?

STORY 3: TROUBLE SLEEPING

Sally sometimes has trouble sleeping, but reading helps her get to sleep. One day, her mum bought a lamp to put next to her bed so she could read if she woke up in the middle of the night. The light is on the left side of her bed. The next night, Sally woke up in the middle of the night, and thought she would like to turn on the light and read. Which direction did she reach for the lamp, to the left side of the bed or to the right side of the bed?

Change-Judgment

Sally grew out of her bed, and needed a bigger one. But the new bed wasn't going to arrive for another night, so Sally was going to sleep in the guest bedroom where her grandma would sleep when she visited. Sally woke up in the middle of the night and wanted to read. Unlike in her own room, the lamp was on the right side of the bed in this room. Which direction did she reach for the lamp, to the left side of the bed or to the right side of the bed?

Memory question

On which side of the bed is the lamp in the guest bedroom that Sally is sleeping in this one night?

On which side of the bed is the lamp in Sally's bedroom where she usually sleeps?

STORY 4: PLAYING WITH DOLLS

Ellie has two dolls, a Stacy doll and a Katie doll. She gets to choose which one to play with. On Monday after school, which doll did she play with?[5]

Change-Judgment

One day after school, Ellie goes to her friend Lilly's house. Lilly has a Katie doll, but also the brand new edition of the Stacy doll, which is way better than the Stacy doll used to be. Lilly lets Ellie have first choice of dolls to play with, which does she play with?

Memory question

Which doll did Lily have a new version of?

Which doll does Ellie usually prefer to play with?

APPENDIX B: Pictures accompanying the vignettes

Story 1: Walking a dog to school

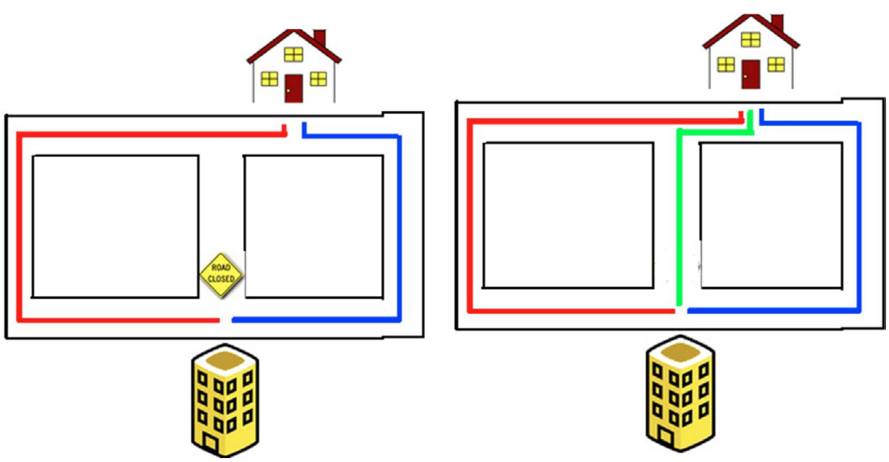

Story 2: Playing sports

Story 3: Trouble sleeping

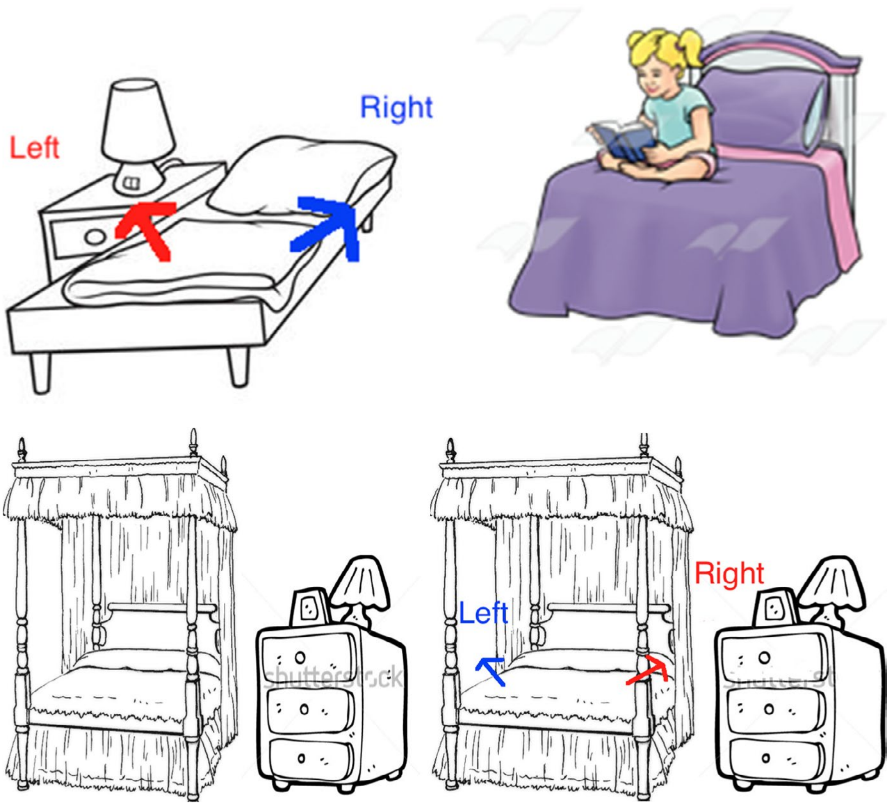

Story 4: Playing with dolls

Katie Doll

Stacy Doll