COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 45 (2021) e13043

ISSN: 1551-6709 online
DOI: 10.1111/cogs.13043

Constrained Choice: Children's and Adults' Attribution of Choice to a Humanoid Robot

Teresa Flanagan,[a] Joshua Rottman,[b] Lauren H. Howard[b]

[a]*Department of Psychology, Cornell University*
[b]*Department of Psychology, Franklin & Marshall College*
jrottman@fandm.edu, lauren.howard@fandm.edu

Received 10 December 2019; received in revised form 16 August 2021; accepted 18 August 2021

Abstract

Young children, like adults, understand that human agents can flexibly choose different actions in different contexts, and they evaluate these agents based on such choices. However, little is known about children's tendencies to attribute the capacity to choose to robots, despite increased contact with robotic agents. In this paper, we compare 5- to 7-year-old children's and adults' attributions of free choice to a robot and to a human child by using a series of tasks measuring agency attribution, action prediction, and choice attribution. In morally neutral scenarios, children ascribed similar levels of free choice to the robot and the human, while adults were more likely to ascribe free choice to the human. For morally relevant scenarios, however, both age groups considered the robot's actions to be more constrained than the human's actions. These findings demonstrate that children and adults hold a nuanced understanding of free choice that is sensitive to both the agent type and constraints within a given scenario.

Keywords: Freedom of choice; Robots; Action prediction; Agency; Developmental psychology

1. Introduction

The ability to choose freely is at the foundation of being human, explaining why we view institutions that restrict free choice, such as prison and slavery, as dehumanizing. Even at a young age, children understand the difference between a choice and a constraint (Kushnir, 2018) and children negatively evaluate people who limit others' ability to choose (Zhao, Zhao, Gweon, & Kushnir, 2020). Is this capacity for choice attributed only to humans or can it also be ascribed to other agents, like robots? Even though children are increasingly beginning to develop alongside these new, interactive technologies (e.g., Amazon Alexa, Cozmo, and

Correspondence should be sent to Teresa Flanagan, Department of Psychology, Cornell University, Ithaca, NY 14850, USA. E-mail: tmf87@cornell.edu

Roomba), and readily engage with such agents as social partners (Kahn et al., 2012a; Martin et al., 2020; Meltzoff, Brooks, Shon, & Rao, 2010), little is known about children's tendencies to attribute robots with the capacity to choose.

From early in life, children understand choice as the ability to consider alternative means for achieving a goal, and to flexibly switch between these alternatives when faced with new constraints. For example, 10-month-old infants expect human agents to use different actions to obtain an object depending on whether physical constraints are present or absent (Brandone & Wellman, 2009), and toddlers use this understanding to differentially respond to agents who could have acted one way, but chose another (Behne, Carpenter, Call, & Tomasello, 2005; Dunfield & Kuhlmeier, 2010; Hamlin, Wynn, & Bloom, 2008). By the preschool years, children are not only able to use causal reasoning across different domains (physical, social, and psychological) to explain why an action did or did not occur (Schult & Wellman, 1997), but are also able to verbally generate alternative options when the main goal of an agent is constrained (Sobel, 2004). Thus, early in life, children understand that human agents can choose to act in certain ways, and that these actions may be constrained by internal or external barriers. This ability to entertain alternative actions suggests that children understand that agents can "choose to do otherwise," a hallmark for a mature understanding of free choice (see Kushnir, 2018).

Children's understanding of choice, however, is not monolithic, as children seem to struggle with understanding how alternatives can be applied in certain circumstances. For example, 4- to 5-year-old children seem to believe that it is not possible to act against desires even without physical constraints (e.g., wanting to eat a tasty cracker but choosing not to; Kushnir, Gopnik, Chernyak, Seiver, & Wellman, 2015), and often choose to act in accordance with their desires at the expense of reaching a salient goal (Yang & Frye, 2018). Relatedly, 3- to 5-year-old children are likely to say that a choice is more moral if it is consistent with an agent's desires (e.g., cleaning up toys because they wanted to) versus conflicting with an agent's desires (e.g., cleaning up toys even if they wanted to go play outside), a pattern that is reversed in older children and adults (Starmans & Bloom, 2016), perhaps because young children do not appreciate that virtuous people frequently need to choose to act in ways that are inconsistent with initial inclinations.

Even research with adults demonstrates conflicting theories of choice in light of constraints. While seemingly counterintuitive, adults appear to believe that it is possible for a person to have choice in a determined universe (i.e., a universe in which every event that happened *had to* happen). For example, Nichols and Knobe (2007) presented a determined universe to adults, such that everything that happened in that universe was completely caused by whatever happened before. Even in this determined universe, adults still held a man morally responsible for killing his family. This reflects the philosophical theory of *compatibilism*, which argues that even if the world at large was discovered to be fixed and determined, it is still possible for humans to have some degree of causal responsibility within that system (Fischer, 2006). Importantly, adults' attributions of choice are dependent on the action and the context of the constraints (Monroe, Dillon, & Malle, 2014; Nichols, 2011). For example, in one scenario, a man named Bill is forced to kill his friend, Joe, by orders of a group of hijackers. Despite the highly constrained context, adults still held Bill morally responsible for killing his friend if it they had information that Bill had previously wanted to kill Joe (Woolfolk, Doris, & Darley, 2006). As such, there are certain scenarios, particularly those relating to internal desires or moral decisions, that appear to muddle both children's and adults' understanding of human choice.[1]

Research in both adults and children clearly demonstrates a nuanced interpretation of choice for *human* targets. However, much less is known about ascriptions of choice to other agents, even though our social world involves numerous kinds of nonhuman agents (e.g., animals, plants, and robots). Of particular interest to this study is how individuals interpret the actions of nonhuman robots. Some work has been done exploring this question in adults, finding that adults are ambivalent concerning robots' free will capacities (Kahn et al., 2012b; Monroe et al., 2014; Weisman et al., 2017). However, this is even more interesting to explore in children because, unlike most adults, they have grown up in a world where robots are significantly more commonplace. Research shows that children can interact with robots as if they are social partners (Kahn et al., 2012a; Martin et al., 2020; Meltzoff et al., 2010). Even though children engage with robots in a sophisticated way, it is unclear if children are willing to believe that robots have free choice. Previous research has demonstrated that children think robots have abilities that are related to free choice (e.g., doing things on purpose, choosing to move, thinking for themselves, and knowing good from bad), but children do not think that robots have physiological experiences like humans (e.g., feeling pain, feeling scared, and feeling hungry; Brink, Gray, & Wellman, 2019). Though some work has shown that children assign more freedom of choice to human agents than to inanimate objects (e.g., a toy ball; Nichols, 2004), we are unaware of any research that has explicitly examined children's attributions of free choice to humanoid robots. Investigating the development of free choice attribution to robots would, therefore, be an initial step toward demonstrating that free choice may not be perceived as human specific. This research has the further potential to reveal the types of features of an agent that are critical for free choice attribution.

Both adults and children ascribe a mixture of animate and inanimate characteristics to humanoid robots, suggesting that robots are perceived as a functionally separate category from either humans or mere objects (Kahn et al., 2012a; Kahn et al., 2012b; Severson & Carlson, 2010). For example, children tend to think that robots hold a certain level of intelligence and some sensory abilities (e.g., can think, can see, and can be tickled), but not emotions or biological capabilities (e.g., can feel happiness, needs sleep, and can grow), though these ascriptions vary with both participant age and robot type (e.g., Bernstein & Crowley, 2008; Jipson & Gelman, 2007; Saylor, Somanader, Levin, & Kawamura, 2010). Furthermore, children often require prior information or experience with robots before they consider them as agentive beings; 18-month-old infants only follow the gaze of a robot they previously saw acting contingently with an adult (Meltzoff et al., 2010), 4- to 7-year-old children are more likely to assume a robot has intelligence if they have more exposure to robots (Bernstein & Crowley, 2008), and 5- to 7-year-old children are more likely to attribute emotional and physical characteristics to a robot that was previously framed as autonomous (Chernyak & Gary, 2016). Thus, robots straddle the animate and inanimate worlds, making them particularly interesting as a test case for children's ascriptions of free choice.

In the current study, we examined whether 5- to 7-year-old children and adults ascribe the freedom to choose to either a human child or a robot agent. In doing so, we explored folk beliefs about how agent type impacts choices and constraints, bridging research on the development of thinking about robots with research on the development of choice attribution. The children's age range was selected because children at this age can attribute choice to other people (Kushnir et al., 2015) and have a sophisticated understanding of robots (Bernstein & Crowley, 2008). An adult sample was collected to compare beliefs about free choice to beliefs about compatibilism (i.e., the idea that people can have some degree of free will even if the world was discovered to be fixed), and to make comparisons between the age groups. Judgments of free choice were measured by participants' predictions of the agent's action (i.e., which of two games would the agent play?) and attribution of choice to the agent (i.e., did the agent choose to or have to play the game?) in three different contexts, manipulated within-subjects.

We introduced the human child and robot agent as being extremely likely to play a specific game (i.e., the "default game") when no external constraints were present (i.e., the child or robot either "only plays" or "was programmed to play" the game, respectively). The human child and robot agent were then presented in three scenarios (No Constraint, Moral Constraint, and Physical Constraint) in which the agent could play either the default game or an alternative game. In the No Constraint scenario, there were no external limitations on the agent to play the typical, default game. In the Moral Constraint and Physical Constraint scenarios, there were external limitations on the agent to play the default game: playing the default game would make a person cry or the default game was broken, respectively. We included these two scenarios because previous research has demonstrated that young children expect humans to act in accordance with social limitations (e.g., you should not make your friend cry even if you want to) and physical limitations (e.g., you cannot run through a wall even if you want to; Chernyak, Kushnir, Sullivan, & Wang, 2013). Including a Moral Constraint scenario and a Physical Constraint scenario in our study allowed us to investigate whether children's and adults' attribution of choice to a robot is domain-specific or domain-general.

For each of these scenarios, we explored whether children and adults would predict the agent to play the default game or instead would predict the agent to respond to situational constraints and play an alternative game. Children and adults were then asked whether the agent "chose to" play the predicted game or "had to" play the predicted game. In both samples, we also measured judgments of the agent's physical, mental, moral, and social capabilities to explore the relationship between free choice judgments and other agentive characteristics. In the adult sample only, we examined responses to an additional series of scenarios known to measure compatibilist thinking (Nichols & Knobe, 2007), a common way of assessing mature beliefs about choice and constraints.

Based on the previous work exploring children's and adults' trait attributions to robots and humans (Chernyak & Gary, 2016; Kahn et al., 2012a; Kahn et al., 2012b), along with previous work on children's and adults' differential reactions to context and constraint (Kushnir et al., 2015; Nichols, 2011), we hypothesized that children and adults would attribute more free choice to the human agent than to the robot agent. Overall, we hypothesized that participants would be more likely to predict that the human child would act against his default action in response to the constraints, but participants would be less likely to predict the robot to act against its programming in response to the constraints. We expected this because prior work has demonstrated that children and adults do not think robots have the same amount of mental and social capabilities as humans (Kahn et al., 2012a; Kahn et al., 2012b). We additionally predicted that participants would be more likely to say that the human agent "chose to" act, while the robot agent "had to" act. We expected this because prior work has demonstrated that children view physical rules (e.g., you cannot choose to fly) as more constraining than internal preferences (e.g., you can choose to not eat your favorite food; Kushnir et al., 2015), and a robot's programming is likely to be more aligned with physical rules, while a human's repeated behavior is likely to be more aligned with internal preferences. All of this would demonstrate that children and adults are able to attribute some free choice to a robot agent, albeit not as much as to a human agent. Ascribing robots with the ability to choose in response to constraints, therefore, would be a step toward demonstrating that children's free choice beliefs can extend to nonhuman entities.

2. Methods

2.1. Participants

Two separate samples were collected, the *Child Sample* and the *Adult Sample*. For the *Child Sample*, there was an a priori decision to stop at 32 children due to an experimenter time constraint and to match sample sizes from prior work (Kushnir et al., 2015). The *Child Sample* consisted of 5- to 7-year-old children ($M_{age} = 5.72$, $SD_{age} = 0.68$, 15 females). Twenty-six children were identified by their parents as White, one was Black/African American, one was Asian, one was Hispanic/Latino, and three were Mixed Race. Seventeen of the children's caregivers had a post-graduate degree, eight had a bachelor's degree, four had an associate degree, one had some college education, one had some high school education, and one caregiver did not report. Children were recruited from a participant database and tested in a laboratory in a small city in the northeastern United States in the Spring of 2018. The majority of participants (84%) had exposure to one or more robotic technology devices (e.g., toys, assistant devices, home appliance devices, and educational toys). One additional child participated but was excluded due to a developmental disability.

For the *Adult Sample*, we collected online data from 60 participants with the goal that at least half of the participants would have usable data points, thus matching the *Child Sample's* size. Twenty participants were excluded for failing a validation check that involved recalling key information from the videos ($N = 17$) or failing a validation check that involved recalling key information about the determined universe scenarios ($N = 3$). Thus, the final sample consisted of 40 U.S. adults ($M_{age} = 38.6$, $SD_{age} = 11.39$; 17 females), analogous in size to the child sample. Twenty-eight adults self-reported as White, three were Black/African American, four were Asian, three were Hispanic/Latino, one was Mixed Race, and one was not reported. Three participants had a post-graduate degree, 21 had a bachelor's degree, 13 had some college education, 2 had a high school diploma, and 1 did not report. Adults were

Fig 1. Screenshots of video stimuli used in the robot condition.
Note. Participants watched a short introduction video and then proceeded to view three constraint scenario videos. Videos in the human condition were identical, with the exception of a human child taking the place of the robot.

recruited via Amazon Mechanical Turk in the Summer of 2020. The majority of participants (95%) had exposure to one or more robotic technology devices.

2.2. Materials and procedure

Participants were randomly assigned to one of two conditions (robot or human agent; the *Child Sample* and the *Adult Sample* each had 16 participants in the robot condition). In the robot condition, participants were asked to watch and respond to the actions of a robot figure named Robovie. The robot was a black and white humanoid robot toy, approximately 35 cm tall (Fig. 1), which could carry out a number of actions. In the human condition, participants watched and responded to the actions of a human child. The actor was a boy approximately the same age as the child participants, named Billy. All stimuli were prerecorded and presented via video on a Dell laptop so that the actions could be matched across conditions. The stimuli, including the prerecorded videos, are available at our repository on the Open Science Framework (https://osf.io/jmdvg/).

Regardless of condition, all participants proceeded through the same paradigm to assess their ascriptions of free choice. After watching scenes that included either the robot or human child, participants were asked to predict the agent's next action (default or alternative game decision) and were asked to attribute choice to the agent (would they "choose to" or "have to" engage in the predicted action, adapted from Kushnir et al., 2015). Via these questions, we were able to investigate whether or not participants believed that an agent could act against its default action and respond to constraints in a way that indicated free choice. After the videos, participants were asked a series of questions related to their beliefs about the agent's physical, mental, moral, and social capacities (adapted from Chernyak & Gary, 2016). In the *Child Sample*, the experimenter advanced the videos and asked the questions. The session was video recorded, and the experimenter live-recorded the child's responses on a piece of paper.[2] In the *Adult Sample*, the survey instructed the participants when to advance the videos. The adults read the questions on the survey and wrote their responses in the text boxes provided to them.

2.3. Video paradigm

2.3.1. Introduction phase

During the introduction phase of the video, participants watched a short clip (60 seconds long) that introduced the agent (robot or human) and showed the agent performing simple actions. The purpose of this introduction was to demonstrate that the agent was autonomous, intentional, and had some basic intelligence, as these capacities are necessary for children to attribute agency (Chernyak & Gary, 2016; Meltzoff et al., 2010). The video consisted of a narrator first describing the agent ("This is Robovie, Robovie is a robot" or "This is Billy, Billy is a kid") paired with a still picture of the agent in a children's room (Fig. 1). Then, the agent performed two simple actions: dancing and throwing a bucket. The next video clip presented the agent's actions as constrained. In the robot condition, participants were shown a video of adults wearing lab coats tinkering on the robot's back. During the video, the narrator said that "Robovie was programmed to know a lot about science and to play science games." In the human condition, participants were shown a video of adults wearing lab coats patting Billy on the back. During the video, the narrator said that "Billy's parents are scientists, so Billy knows a lot about science and plays science games."

After, in both conditions, participants were shown a photo of the agent in a room with a board game. In the photo, the agent was specifically looking and waving their arms at the science game. The narrator repeated that the agent "plays science games every day. [The agent] has a lot of games to play but [the agent] only plays science games. [The agent] plays science games every day." Previous research has shown that children attribute consistent (nonrandom) actions as denoting not just "programing" but desires (Kushnir, Xu, & Wellman, 2010). Therefore, the introduction video portrayed that playing the science game, given no other constraints, was the default action for both agents. (For the full script, see Supplementary Material, Table S1.)

2.3.2. Action prediction

After the introduction video, participants watched three further video segments that presented the three constraint scenarios (No Constraint, Moral Constraint, and Physical Constraint). These segments described the objects in the room (a science game and a history game), presented the relevant constraints on the agent's ability to play these games, and asked the participant to predict which game the agent would play within each of these scenarios.

Table 1
Free choice questionnaire for each constraint scenario

Category	Question
Action Prediction Question:	What game will Robovie/Billy play?
Choice Question:	Did Robovie/Billy *choose to* play the science/history game or did it/he *have to* play the science/history game?
Explanation:	Why?
Additional Moral Constraint Questionnaire	
Responsibility Check Question:	Did Robovie/Billy make Sally happy or sad by playing the science/history game?
Moral Responsibility Question:	Should we give Robovie/Billy a sticker for making Sally happy?/ Should we put the Robovie/Billy in time-out for making Sally sad?
Explanation:	Why?/Why not?

Based on the introduction video segment, the science game should be the *default* game if no other constraints are present. The history game is the *alternative* game decision.

In the first video (No Constraint scenario), participants were asked to pick one of the two games (default or alternative) that the agent would play without any limitations present. Since participants had previously been told that the agent only plays the default game, we hypothesized that participants would predict both agents to play the default game. The second video (Moral Constraint scenario) was identical to the first, but with the limitation that playing the default game would result in hurting someone else's feelings and making that person cry. In this video, it would be wrong for the agent to play the default game, thus requiring the agent to play the alternative game if they wanted to stay within moral bounds. The third video (Physical Constraint scenario) asked the participants to predict the agent's game decision when the default game was broken. In this video, it would be irrational to play the broken (unplayable) game, requiring the agent to play the alternative game. For the Moral Constraint and Physical Constraint scenarios, we hypothesized that participants would be more likely to say the agent would play the default game in the robot condition than in the human condition, indicating differential attributions of the capacity for choice.

2.3.3. Choice attribution

After each scenario, participants were asked two follow-up questions to explore choice attribution, adapted from Kushnir et al. (2015). Specifically, participants were asked whether the agent "chose to" or "had to" play the default or alternative game (whichever was chosen in the Action Prediction question). For each constraint scenario, we hypothesized that participants would ascribe more choice to the human agent than to the robot agent.

Next, participants were asked to explain their answer (Table 1). Participants' verbal explanations were transcribed offline and coded into one of five possible explanation categories, adapted from Kushnir et al. (2015): Internal (reference to desires or psychological states that could influence action), External (reference to physical, social, or biological conditions that could influence action), Autonomy (reference to general ability to choose autonomously),

Table 2
Explanation coding scheme

Explanation Category	Examples
External: reference to physical, social, or physiological conditions that constrain or influence the action	“Because the other one is broken”
Internal: reference to desires or internal states that constrain or influence the action	“Because he likes science”
Autonomy: reference to general ability to choose autonomously	“It’s his choice”
Determined: reference to lack of ability to choose autonomously based on the introduction video	“Because it’s programmed to play science”
I don’t know/Other/No response	

Determined (reference to lack of ability to choose autonomously), and Other (Table 2). Two coders independently scored all explanation categories for each explanation ($\kappa = 0.77$, $p < .001$) and discrepancies were resolved through discussion.

Furthermore, in the Moral Constraint scenario, we also explored whether participants differentially held the agent morally responsible for the agent’s moral action. After the choice attribution question in the Moral Constraint scenario, participants were asked a Responsibility Check question to ensure the participant understood what would happen next. Participants passed the Responsibility Check question if the participant’s response aligned with their game prediction (e.g., saying the agent made the person sad for playing the default game or saying the agent made the person happy for playing the alternative game). Participants did not pass the Responsibility Check question if the participant’s response did not align with their game prediction (e.g., saying the agent made the person happy for playing the default game or saying the agent made the person sad for playing the alternative game). Participants were then asked whether the agent should be praised or blamed for the action, indicating moral responsibility. Since it is often required that an agent has free choice in order to be morally responsible (Nichols & Knobe, 2007), we predicted that participants would attribute moral responsibility to the child agent more than to the robot agent.

2.3.4. Agency attribution

After viewing the three scenarios, participants were asked about the agent’s physical, mental, moral, prosocial,[3] and friendship characteristics (see Supplementary Materials, Table S2), following Chernyak and Gary (2016). The purpose of asking these questions was to compare these attributions to the free choice attributions, and therefore possibly demonstrate certain characteristics that may be required for free choice attribution. This also served to test whether participants were judging this robot as an agent or as a mere object.

2.3.5. Determined universe

In the *Adult Sample*, participants were also asked about two different hypothetical universes (adapted from Nichols & Knobe, 2007). Universe A was causally determined, such that everything that happened was completely caused by whatever happened before. Universe

Table 3
Likelihood ratio chi-square tests statistics for free choice questions

		Main Effect of Agent Condition (Human vs. Robot)	Effect of Sample Group (Child vs. Adult)	Interaction Effect (Agent Condition * Sample Group)
Action Prediction	No Constraint	$\chi^2(1) = 1.42$	$\chi^2(1) = 1.87$	$\chi^2(1) = –3.03$
	Moral Constraint	**$\chi^2(1) = 20.83$**	**$\chi^2(1) = 10.71$**	$\chi^2(1) = –0.15$
	Physical Constraint	**$\chi^2(1) = 15.50$**	**$\chi^2(1) = 22.00$**	**$\chi^2(1) = 5.05$**
Choice Attribution	No Constraint	**$\chi^2(1) = 4.38$**	$\chi^2(1) = 0.87$	$\chi^2(1) = 2.63$
	Moral Constraint	$\chi^2(1) = 3.28$	**$\chi^2(1) = 8.17$**	$\chi^2(1) = 1.77$
	Physical Constraint	$\chi^2(1) = 3.46$	$\chi^2(1) = 3.56$	$\chi^2(1) = 0.37$

Note: Significant results are bolded.

B was almost determined, such that *almost* everything that happened was completely caused by whatever happened before. Participants were first asked which universe they thought was most similar to ours and were asked to give an explanation for their answer. Then, participants were told about a man in Universe A (the causally determined one) who killed his wife and children. Participants were asked if it was possible for that man to be morally responsible for killing his family (see Supplementary Materials for scenario and questionnaire). The purpose of this addition was to directly measure adults' compatibilist beliefs (i.e., belief that it is possible to have free will in a determined universe). Participants were coded as endorsing compatibilism if they believed that it was possible for the man in Universe A to be morally responsible.[4]

3. Results

3.1. Action prediction

Participants' action predictions were coded for each of the three constraint scenarios (No Constraint, Moral Constraint, and Physical Constraint). Within each constraint, participants received a score of 1 if they predicted the default game, and 0 if they predicted the alternative game. We ran separate general logistic models for each constraint scenario with participant's action prediction as the outcome variable and agent condition (Human vs. Robot) and sample group (Child vs. Adult) as between-subjects variables with a bias-reduction method[5] (Kosmidis, 2021). Test statistics are presented in Table 3. We then ran binomial tests to determine whether the percentage of the default game prediction differed from chance (50%). Percentages are presented in Fig. 2.

3.1.1. No constraint

We did not find a significant main effect of agent condition or sample group and we did not find a significant interaction effect between the two (Table 3). Participants overwhelmingly

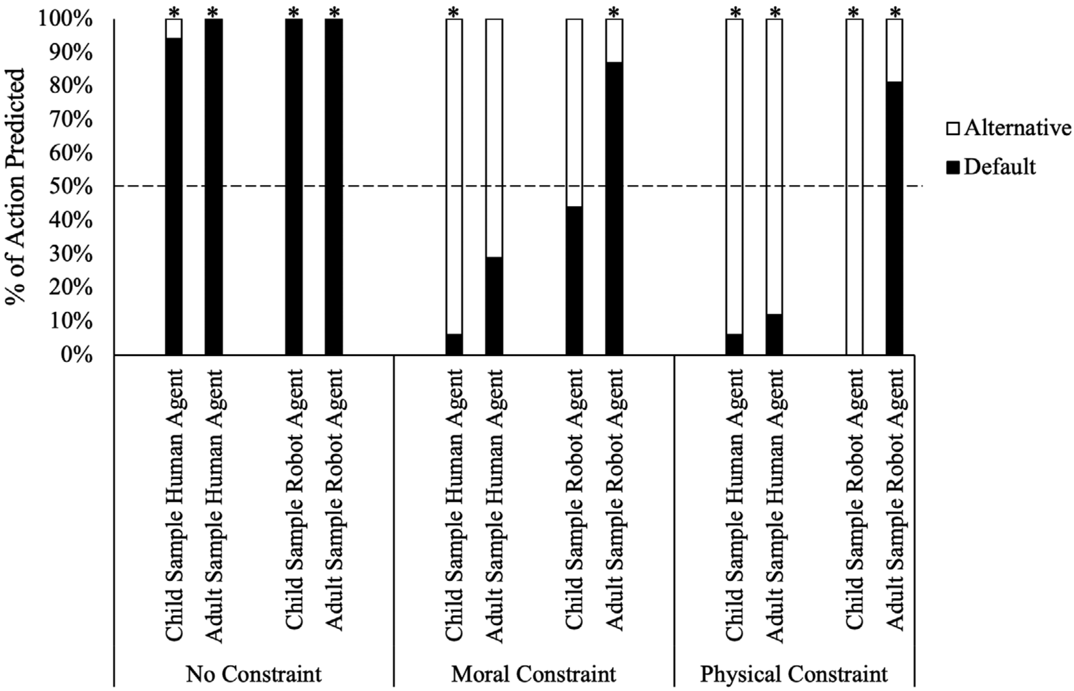

Fig 2. Percentage of the child and adult participants' action predictions, separated by agent and constraint type. *Note*. Asterisks signify response patterns that are significantly different than chance (50%).

tended to predict that both the human child and the robot would play the default game when no external constraints were present, binomial $ps < .001$ (Fig. 2).

3.1.2. Moral constraint

We found a main effect of agent condition, $\chi^2(1) = 20.83$, $p < .001$, and a main effect of sample group, $\chi^2(1) = 10.71$, $p = .001$. We did not find a significant interaction between the two (Table 3). Despite an absence of an interaction, an inspection of patterns yielded by the two main effects suggested that it would be most informative to analyze conditional effects of agent type within each age group. Thus, in an exploratory analysis investigating the main effects further, we first analyzed the conditional effects of the agent condition for each sample group and then analyzed the conditional effects of the sample group for each agent condition. For the *Child Sample*, the odds of predicting the default game were 8.13 times higher for the robot condition (44%) than the human condition (6%), $p = .039$. Children tended to predict that the human agent would not play the default game, binomial $p < .001$, but were divided for the robot agent, binomial $p = .804$. For the *Adult Sample*, the odds of predicting the default game were 13.53 times higher for the robot condition (87%) than the human condition (29%), $p = .002$. Adults tended to predict that the robot agent would play the default game, binomial $p = .004$, but were divided for the human agent, binomial $p = .064$. For the human condition, there was no significant difference between sample groups, $OR = .226$, $p = .132$. For the robot condition, however, the odds of predicting the robot to play the default game were 7.35 times higher for adults than children, $p = .021$.

3.1.3. Physical constraint

We found a main effect of agent condition, $\chi^2(1) = 15.51$, $p < .001$, and a main effect of sample group, $\chi^2(1) = 22.00$, $p< .001$. We also found a significant interaction between the agent condition and sample group, $\chi^2(1) = 5.05$, $p = .025$ (Table 3). To investigate the main effects further, we first analyzed the conditional effects of the agent condition for each sample group and then analyzed the conditional effects of the sample group for each agent condition. For the *Child Sample*, there was no significant difference in the odds of predicting the default game for the robot condition (0%) compared to the human condition (6%), $OR = 3.19$, $p = .500$. Children overwhelmingly tended to predict that both agents would not play the default game, binomial $ps < .001$. For the *Adult Sample*, however, the odds of predicting the default game were 23.70 times higher for the robot condition (81%) than the human condition (12%), $p < .001$. Adults overwhelmingly tended to predict that the human agent would *not* play the default game, binomial $p < .001$, but predicted that the robot agent would play the default game, binomial $p = .021$. For the human condition, there was no significant difference between sample groups, $OR = 0.59$, $p = .624$. For the robot condition, however, the odds of predicting the default game were 127.23 times higher for adults than children, $p = .003$. To investigate the interaction effect further, we then compared the conditional effects of the agent condition between sample groups. We found a significant difference in the magnitude of the agent conditional effect between sample groups, $OR = 75.7$, $p = .025$.

3.2. Choice attribution

In each scenario, participants were asked whether the agent "chose to" or "had to" play the predicted game, regardless of which game type was predicted (default or alternative). Choice was indicated if the participant said the agent "chose to" play the game, receiving a score of 1, while constraint was indicated if the participant said the agent "had to" play the game, receiving a score of 0. We ran separate general logistic models for each constraint scenario with participant's choice attribution as the outcome variable and agent condition (Human vs. Robot) and sample group (Child vs. Adult) as between-subjects variables with a bias-reduction method (Kosmidis, 2021). Test statistics are presented in Table 3. We then ran binomial tests to determine whether the percentage of the "choose to" response differed from chance (50%). Percentages are presented in Fig. 3.

3.2.1. No constraint

We found a main effect of agent condition, $\chi^2(1) = 4.38$, $p = .037$. We did not find a main effect of sample group or a significant interaction effect between the two (Table 3). However, an inspection of patterns across conditions suggested a particularly high pattern of choice attribution to the human agent within the adult sample. Thus, we conducted an exploratory analysis investigating the conditional effects of the agent condition for each sample group. For the *Child Sample*, there was no significant difference in the odds of saying the agent chose to

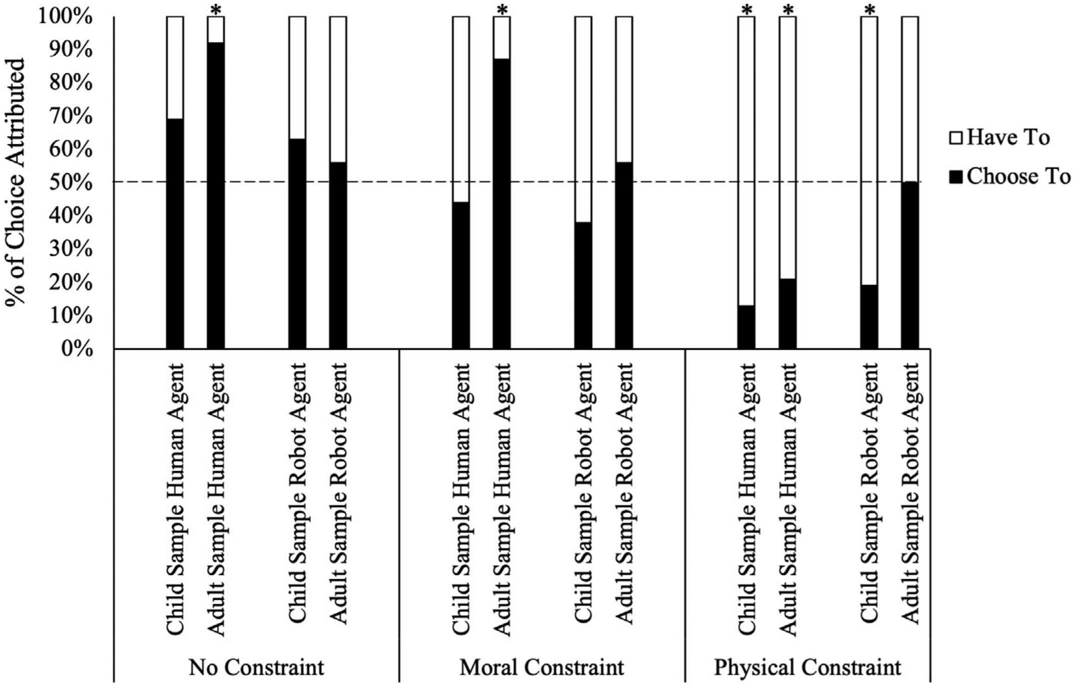

Fig 3. Percentage of child and adult participants' attributions of choice, split by agent and constraint type.
Note. Asterisks signify choice response patterns that are significantly different than chance (50%).

play the game in the human condition (69%) compared to the robot condition (63%), $OR = 1.29$, $p = .728$. Children's responses on whether the agent "chose to" play the game did not differ from chance for the human or robot agent, binomial $ps \geq .210$. For the *Adult Sample*, however, the odds of saying the agent chose to play the game were 7.11 times higher for the human condition (92%) than the robot condition (56%), $p = .021$. Adults overwhelmingly said that the human agent "chose to" play the game, binomial $p < .001$, but were divided for the robot agent, binomial $p = .804$.

3.2.2. *Moral constraint*

We found a main effect of sample group, $\chi^2(1) = 8.17$, $p = .004$. We did not find a significant main effect of agent condition or an interaction effect between the two (Table 3). Again, however, an inspection of patterns indicating that adults were particularly likely to attribute choice to the human agent led us to conduct an exploratory analysis investigating the conditional effects of the sample group for each agent condition. For the human condition, the odds of saying the human chose to play the game were 7.75 times higher for adults than children, $p = .008$. Children's responses on whether the human "chose to" play the game did not differ from chance, binomial $p = .804$, while adults overwhelmingly said that the human "chose to" play the game, binomial $p < .001$. For the robot condition, however, there was no significant difference in the odds of saying the robot chose to play the game for adults compared to children, $OR = .489$, $p = .320$. Children and adults were divided on whether the robot "chose to" play the game, binomial $ps \geq .454$.

For the Moral Responsibility question, analyses excluded participants who failed the Responsibility Check question (*Child Sample*: five participants in the human condition and one participant in the robot condition; *Adult Sample*: one participant in the human condition). We then ran a general logistic model with attribution of moral responsibility as the outcome variable and agent condition (Human vs. Robot) and sample group (Child vs. Adult) as between-subjects variables.

We found a main effect of agent condition, $\chi^2(1) = 16.78$, $p < .001$, and a main effect of sample group, $\chi^2(1) = 21.61$, $p < .001$. We did not find an interaction effect, $\chi^2(1) = 0.73$, $p = 1.00$. Again, patterns in the data led us to conduct an exploratory analysis investigating each of the conditional effects resulting from the 2×2 design. For the *Child Sample*, there was no significant difference in the odds of saying the agent should be held responsible in the human condition (100%) compared to the robot condition (80%), $OR = 6.44$, $p = .254$. Children tended to say the agent should be held responsible for both the human and robot agents, binomial $ps \leq .035$. For the *Adult Sample*, however, the odds of saying the agent should be held responsible were 12.8 times higher for the human condition (70%) than the robot condition (13%), $p = .002$. Adults were divided on whether the human agent should be held responsible, binomial $p = .093$, but overwhelmingly said that the robot agent should not be responsible, binomial $p = .004$. For the human condition, there was no significant difference between sample groups, $OR = 10.5$, $p = .136$. For the robot condition, however, the odds of saying the robot should be held responsible were 20.7 times higher for children than adults, $p = .001$.

3.2.3. Physical constraint

We did not find a significant main effect of agent condition or sample group and we did not find a significant interaction effect between the two (Table 3). Participants overwhelmingly tended to say that the agents "had to" play the game, binomial $ps \leq .021$, except for adults in the robot condition, where responses were divided, binomial $p = 1.00$.

3.3. Explanations of choice attributions

In addition to asking participants which game the agent would select under the three constraints, we also asked why they predicted that decision (see Table 2). Across constraint scenarios, a majority of child participants gave External explanations for both the human agent (62.5%) and the robot agent (54.2%). The majority of adult participants gave External explanations for the human agent (60.13%) and Deterministic explanations for the robot agent (52.1%). See Fig. 4 for distributions of explanations within constraint scenarios.

3.4. Attributions of agency

Participants' agency attributions (i.e., their attributions of physical, mental, moral, and friendship characteristics) were entered into a repeated measures general logistic model with attribution of agency (yes/no) as the outcome variable and item type (Physical vs. Mental vs.

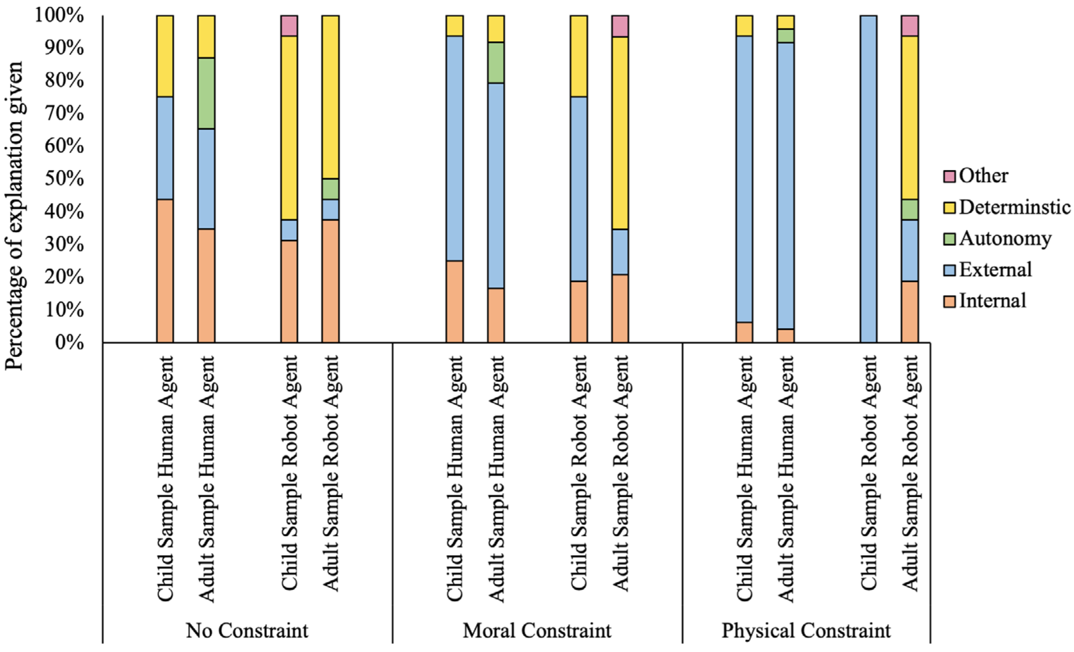

Fig 4. Percentage of child and adult participants' explanation type, separated by agent condition and constraint scenario.

Moral vs. Friend) as a within-groups factor, agent condition (Human vs. Robot) and sample group (Child vs. Adult) as between-groups factors, and ID and item type as random factors. Percentages are presented in Fig. 5.

We found a main effect of agent condition, $\chi^2(1) = 26.54$, $p < .001$, item type, $\chi^2(3) = 26.98$, $p < .001$, and sample group, $\chi^2(1) = 9.36$, $p = .002$, and a significant two-way interaction between agent condition and item type, $\chi^2(3) = 26.04$, $p < .001$. Controlling for sample group, the odds of attributing agentic capabilities were significantly increased in the human condition as compared to the robot condition for the mental capabilities (e.g., getting upset, being intelligent, and having feelings), $OR = 7.21$, $p < .001$, physical capabilities (e.g., feel being tickled and get hurt), $OR = 877.77$, $p < .001$, moral treatment (e.g., should not be hit and should not be yelled at), $OR = 7.21$, $p = .035$, and for saying the agent could be their friend, $OR = 8.37$, $p = .031$. In the robot condition, controlling for sample group, the odds of attributing agentic capabilities were significantly higher for saying the robot deserves moral treatment and could be their friend than for mental capabilities (Moral/Mental: $OR = 25.45$, $p < .001$; Friend/Mental: $OR = 9.11$, $p = .001$) and physical capabilities (Moral/Physical: $OR = 75.55$, $p < .001$; Friend/Physical: $OR = 27.02$, $p < .001$). There were no significant differences between questions in the human condition, controlling for sample group, $ps \geq .27$. Controlling for agent condition and item type, child participants had a significant increase in the odds of attributing agentic capabilities as compared to adult participants, $OR = 6.4$, $p = .002$.

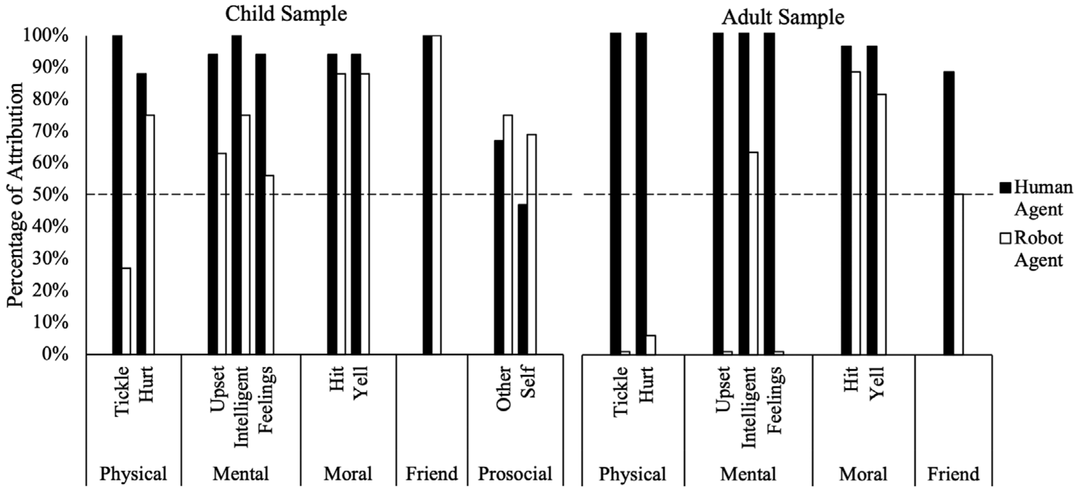

Fig 5. Percentage of attributions across item types, split by condition and sample groups.

The Prosocial item type questions were not included in the regression analysis because the *Adult Sample* did not answer these two questions. For the *Child Sample*, we ran separate general logistic models for each of the prosocial questions (Other: give the ball for another person or to the agent to play with; Self: keep sticker for themselves or give to the agent) with the agent condition (Human vs. Robot) as the between subjects variable. We did not find a significant difference between agent condition for either question (Other: $\chi^2(1) = 0.24$, $p = .627$; Self: $\chi^2(1) = 1.55$, $p = .213$).

A series of chi-square association analyses were also run to measure whether or not there was a relationship between measures of free choice (action prediction and choice attribution) and agency attributions (physical, mental, moral, prosocial, and friend) within each of the three constraint scenarios (No Constraint, Physical Constraint, and Moral Constraint) in the robot condition. In general, we did not find any significant relationships between free choice and agency attribution, except between saying the robot should be held responsible and saying the robot should not be hit, $\chi^2(1) = 4.286$, $p = .038$; for more information, see Supplementary Materials.

4. Discussion

From an early age, children understand that humans have the capacity to make choices (Kushnir, 2018) and, by the time they are adults, they have nuanced understandings about the ways in which those choices might be constrained (Nichols, 2011). Here, we utilized a novel approach to explore whether children and adults tend to extend the capacity to choose to humanoid robots. Our findings highlight both remarkable similarities and differences across these two samples. For example, both children's and adults' judgments of an agent's free choice were dependent both on the constraint and agent type. Without any constraints, children and adults predicted similar actions in both agent conditions. With a physical constraint, adults, but not children, predicted that the robot would play the broken, unplayable game. When the agent had an opportunity to change its actions in order to avoid making another person sad, both children and adults judged the robot as less likely to act in this way than the human, even though participants also believed the robot deserved moral treatment and could be their friend. Children's and adults' attributions of free choice, therefore, were dependent on the context of the scenario, incorporating both the agent's perceived ability and constraints on the action. Together, these results are a step toward demonstrating that free choice beliefs can extend to nonhuman entities, even during childhood.

The current study obtained these nuanced insights into choice attributions by utilizing a three-pronged approach to exploring children's and adults' understanding of free choice across agent types, asking questions about action prediction, choice attribution, and general abilities in order to provide a comprehensive set of distinct perspectives. Although versions of each of these individual strategies have previously been used to investigate children's attributions of choice, they have not previously been used in tandem. For example, some studies have investigated children's predictions of how entities would react when faced with a potential decision that differs from the norm (Chalik, Rivera, & Rhodes, 2014). If entities are thought to have the capacity for choosing an action that deviates from previous conventions, this can present compelling evidence that these agents are believed to have a capacity to choose freely. Other studies have measured children's explicit beliefs about how entities' actions were caused (Kushnir et al., 2015). If agents are described as "choosing to" engage in an alternative action, rather than "having to" engage the conventional action, this may suggest that participants have an explicit concept of the entities as being agents with the capacity to choose freely. Finally, some studies have investigated children's attributions of traits to entities, which are influenced by the entities' perceived level of autonomy (Chernyak & Gary, 2016). Each of these approaches provides a different perspective on how choice is construed, and their combination produced several intriguing patterns.

As anticipated, participants overwhelmingly predicted both the robot and the human to play the "default" game that the respective agent had always tended to play, given a context with no constraints that would lead them to act otherwise (No Constraint scenario). This served primarily as a manipulation check, assuring that participants clearly understood the introduction videos and the differential backstory for why each agent would play one game over another. However, although we expected participants to explicitly attribute more choice to the human child when picking the game, whereas the robot's programming might appear more constraining, we did not find support for this hypothesis with the child participants. Instead, children were divided as to whether they thought the agent "chose to" or "had to" to play a particular game for both agent types (robot and human child). Adults, on the other hand, attributed more choice to the human than they did for the robot.

Additionally, children's open-ended explanations tended to reference the robot's programming (56.3%) and the human's desires (43.8%), perhaps suggesting that young children view desires as being constraining in the same way as programming. Younger children are more likely to believe that you cannot choose to act against your desires (Kushnir et al., 2015; Yang & Frye, 2018) and as children get older, they become less likely to view desires as constraining (Chernyak et al., 2013), which can explain why the adult participants were more likely to ascribe choice to the human than the robot. This developmental pattern coincides with children's developing judgment of possibility (see Kushnir, 2018 for more discussion), so children's view of desires as limiting possibilities may be similar to programming. Alternatively, it may be that children did not view the robot's computational programming as being as constrained or determined as anticipated.

Though some level of action-switching was expected across both agent types when the default game was broken and thus unplayable (Physical Constraint scenario), we hypothesized that the robot might be considered more limited than the human in its ability to switch from the broken (default) game to the alternative game given the more constrained nature of its programming. Adults' responses supported this hypothesis; they were more likely to expect the robot, as compared to the human child, to play the broken, default game, suggesting that adults believed the robot to be relatively less responsive to external constraints. Contrary to this hypothesis, we found that children were overwhelmingly likely to say that both the human and the robot would switch to the playable, alternative game. Intriguingly, when participants expected agents to counteract their default action (which has been interpreted as an indicator that choice is operating), they generally reported that this action was not based on choice, but rather something they "had to" do. In other words, acting against a default tendency may not be sufficient for choice. Rather, true choice may only be attributed in cases when reasonable alternative possibilities are available.

The No Constraint and Physical Constraint scenarios, therefore, provided two complementary perspectives on children's tendencies to attribute free choice to humans and robots. By using multiple measures, we were able to obtain a nuanced picture of inclinations to attribute free choice, while showing across these two scenarios that children's tendencies were remarkably similar when applied to a human child and when applied to a humanoid robot.

Children and adults differentiated their predictions for how a robot or human would act in a case when playing the default game would cause emotional harm to a playmate (Moral Constraint scenario), supporting our hypothesis. Specifically, participants were more likely to believe that the human would be responsive to the moral constraints of the scenario, selecting the alternative game to avoid distressing another person, than the robot agent. This finding is particularly interesting in relation to the Physical Constraint scenario. Adults, on the one hand, did not think that the robot would play the alterative game in either scenario. This suggests that, in both scenarios, adults generally thought that the robot was not capable of acting against its programming, and thus incapable of having free choice.

Children, on the other hand, thought that the robot would play the alternative game in the Physical Constraint scenario. This suggests that children thought that the robot was capable of acting against its programming to some extent. Therefore, there may be something unique about moral situations that make children less willing to believe that a robot would change its behavior than a human would. This may be either because children thought the robot did not care about moral reasons or because they thought the robot lacked the capacity to recognize moral reasons. Some support for the latter explanation can be found in prior research, which has demonstrated that children and adults view robots as having limited mental and emotional capabilities in comparison to humans (Kahn et al., 2012a; Kahn et al., 2012b). This was also echoed in our own results. On the Agency Attribution measure, adult and child participants were hesitant to attribute certain emotional capabilities to the robot, such as the capacity to become upset if someone was mean to it (Mental–Upset attribution) and the ability to experience feelings like happiness or sadness (Mental–Feelings). As emotional attunement is viewed as important for solving moral dilemmas (Danovitch & Keil, 2008), it is possible that children may have assumed a connection between the robot's lack of emotional capacity and its responsivity to emotional distress.

Adults were more likely to say that the human agent "chose" his action in the Moral Constraint scenario than children. This suggests either that some children viewed the game decision as being heavily constrained, while others did not or that children in general were not sure if the game decision was constrained. While we originally hypothesized that both child and adult participants would attribute more choice for the human agent, this finding is not surprising given prior work demonstrating that young children often believe moral actions to involve less free choice and more constraints than other types of nonmoral or nonsocial actions (Chernyak et al., 2013).

Even though children and adults did not think the robot had mental and physical capabilities, they were more likely to believe that the robot deserved to be spared moral harms, such as being hit or yelled at (Moral–Hit and Moral–Yell) and had the capacity to be their friend (Friend). One possibility is that children and adults assume that causing moral harm is bad (Schmidt, Rakoczy, & Tomasello, 2012) and this extends to nonliving agents, such as puppets (Vaish, Missana, & Tomasello, 2011). It is also possible that children and adults do in fact see humanoid robots as holding some moral standing, even in situations like those in the current study that involve noninteractive scenarios and limited prior information about the robot's agentive abilities (Sommer et al., 2019). If this is true, then increasing the robot's perceived agency via live social interactions with the participants might also increase the robot's perceived moral standing (see Kahn et al., 2012a; Meltzoff et al., 2010). Future research could explore this possibility by manipulating the type of interaction the participant has with the robot and by measuring the effects of this interaction on a host of agentic attributions (e.g., epistemic trust, moral capabilities, and free choice).

Differences in experience with robotic technologies may explain why children were more likely to view a humanoid robot as having free choice than adults. Children are increasingly developing alongside interactive technologies and so they may be more likely than adults to view such technologies as agentic beings (Severson & Carlson, 2010). On the other hand, children may be more naturally inclined to attribute free choice to any animate being, and this attribution may narrow as we gain more knowledge about specific agent types. For example, children who have increased experience with robots tend to attribute psychological characteristics to these agents, but still distinguish robots from living kinds (Bernstein & Crowley, 2008). The majority of child and adult participants in our study reported prior experience with robotic technologies, though adults may have had more knowledge about robots' mechanics, making them biased against viewing robots as agents (see Fiala, Arico, & Nichols, 2014). Therefore, it is possible that children may develop this bias against robot agents as they become more aware about robots' mechanisms or, as children continue to develop alongside such technologies, they may overlook this bias and continue to view robots as agents into adulthood. It is likely, therefore, that the type of experience with robots, may interact with age to create different beliefs about choice.

The limitations of the current study raise important questions for future investigations. First, we had a small number of participants who were mostly White, from an educated background, and who had experience with robotic technologies. Therefore, our findings may not be generalizable to the greater global population. With our small sample size, we were unable to detect strong relationships between free choice judgments and agency attributions or compatibilist beliefs. Second, our study did not investigate the role of individual differences, such as personality traits, cultural ideology, and religion, all of which have been shown to alter perceptions of free will in adults (Feltz & Cokely, 2009; Sarkissian et al., 2010; Wente et al., 2016). Thus, it stands to reason that these differences might also alter the way in which children and adults construe the agency of a humanoid robot. Future research with larger and more diverse samples could further explore how our perceptions of choice are tied to our lived experiences and surroundings, and thus can better attempt to capture the diverse human phenome (Barrett, 2020).

Finally, future research may investigate whether other conceptions of choice are manifested in children's and adults' attributions of robots' capacities. In our study, judgments of free choice were measured by action prediction and attribution of choice, but it is possible that our participants may have differed in their interpretation of "choosing to" do something. Therefore, different measures are needed in order to fully capture our judgments of free choice to robots. Other measures include attribution of intentionality (Reeder, Kumar, Hesson-McInnis, & Trafimow, 2002), action prediction with second-order desires (i.e., the desire to desire something; Frankfurt, 1971), and moral responsibility (Nichols & Knobe, 2007). We believe that such research would further provide indications that attributions of free choice are not fixed, but situational (see Woolfolk et al., 2006). To the extent that robots are consistently attributed with the capacity for choice, issues of moral standing and moral responsibility may increasingly come into play as technology continues to rapidly develop.

Notes

1 It is important to note that thus far and throughout the paper, the research presented on free choice has primarily been conducted in the United States. Free choice understanding varies by culture, both in adults (Sarkissian et al., 2010) and in children (Chernyak, Kang, & Kushnir, 2019), and therefore, the current literature should not be generalized to other populations.

2 If there was any confusion with the participant's response, two other researchers reviewed the video and agreed on an answer.

3 Adults were not asked the prosocial question (see Supplementary Materials, Table S2) because the questions involved giving physical items to the agent or keeping these items for themselves, which could not be directly replicated in an online measure.

4 Given that a series of Chi Square association analyses revealed no strong relationships between adults' compatibilist beliefs and their free choice attributions, we do not dis-

cuss this measure further in the main text. For more information, see the Supplementary Materials.

5 The bias-reduction method (Firth, 1993) places a penalty term on the standard maximum likelihood function used to generate a logistic model's parameter estimates and standard errors. The penalty term converges toward 0 as the sample size increases, making the method ideal for small sample bias.

Acknowledgments

We thank Dr. Robyn Waller and Dr. Tamar Kushnir for their feedback on this project, Erika Mudrak for her help with the statistical analyses, as well as Julia McAleer for her work as a research assistant. Also, we thank the families from the Lancaster community for participating in this project and Nathan Waller for his acting in the human child condition.

The research was supported by a Nissley Scholar Grant awarded by the Franklin and Marshall College Committee on Grants.

References

Barrett, H. C. (2020). Deciding what to observe: Thoughts for a post-WEIRD generation. *Evolution and Human Behavior*, *41*, 5, 445–453. https://doi.org/10.1016/j.evolhumbehav.2020.05.006

Behne, T., Carpenter, M., Call, J., & Tomasello, M. (2005). Unwilling Versus Unable: Infants' Understanding of Intentional Action. *Developmental Psychology*, *41*, (2), 328–337. http://doi.org/10.1037/0012-1649.41.2.328

Bernstein, D., & Crowley, K. (2008). Searching for signs of intelligent life: An investigation of young children's beliefs about robot intelligence. *Journal of the Learning Sciences*, *17*(2), 225–247. https://doi.org/10.1080/10508400801986116

Brandone, A. C., & Wellman, H. M. (2009). You can't always get what you want: Infants understand failed goal-directed actions. *Psychological Science*, *20*(1), 85–91. https://doi.org/10.1111/j.1467-9280.2008.02246.x

Brink, K. A., Gray, K., & Wellman, H. M. (2019). Creepiness creeps in: Uncanny valley feelings are acquired in childhood. *Child Development*, *90*(4), 1202–1214. https://doi.org/10.1111/cdev.12999

Chalik, L., Rivera, C., & Rhodes, M. (2014). Children's use of categories and mental states to predict social behavior. *Developmental Psychology*, *50*(10), 2360–2367. https://doi.org/10.1037/a0037729

Chernyak, N., & Gary, H. E. (2016). Children's cognitive and behavioral reactions to an autonomous versus controlled social robot dog. *Early Education and Development*, *27*(8), 1175–1189. https://doi.org/10.1080/10409289.2016.1158611

Chernyak, N., Kang, C., & Kushnir, T. (2019). The cultural roots of free will beliefs: How Singaporean and U.S. children judge and explain possibilities for action in interpersonal contexts. *Developmental Psychology*, *55*(4), 866–876. https://doi.org/10.1037/dev0000670

Chernyak, N., Kushnir, T., Sullivan, K. M., & Wang, Q. (2013). A comparison of American and Nepalese children's concepts of freedom of choice and social constraint. *Cognitive Science*, *37*(7), 1343–1355. https://doi.org/10.1111/cogs.12046

Danovitch, J. H., & Keil, F. C. (2008). Young Humeans: The role of emotions in children's evaluation of moral reasoning abilities. *Developmental Science*, *11*(1), 33–39. https://doi.org/10.1111/j.1467-7687.2007.00657.x

Dunfield, K. A., & Kuhlmeier, V. A. (2010). Intention-mediated selective helping in infancy. *Psychological Science*, *21*, 523–527. https://doi.org/10.1177/0956797610364119

Feltz, A., & Cokely, E. T. (2009). Do judgments about freedom and responsibility depend on who you are? Personality differences in intuitions about compatibilism and incompatibilism. *Consciousness and Cognition*, *18*(1), 342–350. https://doi.org/10.1016/j.concog.2008.08.001

Fiala, B., Arico, A., & Nichols, S. (2014). You, robot. In E. Machery & E. O'Neill, *Current controversies in experimental philosophy* (pp. 31–47). Retrieved from New York, NY: Routledge. https://philarchive.org/archive/FIAYRv1

Firth, D. (1993). Bias reduction of maximum likelihood estimates. *Biometrika*, *80*(1), 27–38. https://doi.org/10.1093/biomet/80.1.27

Fischer, J. M. (2006). Responsiveness and moral responsibility. F. Schoeman, *My way: Essays on moral responsibility* (pp. 63–83). New York: Oxford University Press.

Frankfurt, H. G. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, *68*, (1), 5–20. http://doi.org/10.2307/2024717

Hamlin, J. K., Wynn, K., & Bloom, P. (2008). Social evaluation by preverbal infants. *Nature*, *450*, 557–559. https://doi.org/10.1038/nature06288

Jipson, J. L., & Gelman, S. A. (2007). Robots and rodents: Children's inferences about living and nonliving kinds. *Child Development*, *78*(6), 1675–1688. https://doi.org/10.1111/j.1467-8624.2007.01095.x

Kahn, P. H., Kanda, T., Ishiguro, H., Freier, N. G., Severson, R. L., Gill, B. T., Ruckert, J. H., & Shen, S. (2012). "Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot. *Developmental Psychology*, *48*, (2), 303–314. http://doi.org/10.1037/a0027033

Kahn, P. H., Kanda, T., Ishiguro, H., Gill, B. T., Shen, S., Gary, H., Reichert, A., Freier, N. & Severson, R. (2012b). Do people hold a humanoid robot morally accountable for the harm it causes? In Proceedings of the 7th Annual ACM/IEEE International Conference on Human–Robot Interaction - HRI '12. New York: ACM Press. https://doi.org/10.1145/2157689.2157696

Kosmidis, I. (2021). Bias reduction in binomial-response generalized linear models. Retrieved from https://cran.r-project.org/web/packages/brglm/brglm.pdf

Kushnir, T. (2018). The developmental and cultural psychology of free will. *Philosophy Compass*, *13*(11), 1–17. https://doi.org/10.1111/phc3.12529

Kushnir, T., Gopnik, A., Chernyak, N., Seiver, E., & Wellman, H. M. (2015). Developing intuitions about free will between ages four and six. *Cognition*, *138*, 79–101. https://doi.org/10.1016/j.cognition.2015.01.003

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, *21*(8), 1134–1140. https://doi.org/10.1177/0956797610376652

Martin, D. U., Perry, C., MacIntyre, M. I., Varcoe, L., Pedell, S., & Kaufman, J. (2020). Investigating the nature of children's altruism using a social humanoid robot. *Computers in Human Behavior*, *104*. 106149. https://doi.org/10.1016/j.chb.2019.09.025

Meltzoff, A., Brooks, R., Shon, A., & Rao, R. (2010). "Social" robots are psychological agents for infants: A test of gaze following. *Neural Networks*, *23*(8–9), 966–972. https://doi.org/10.1016/j.neunet.2010.09.005

Monroe, A. E., Dillon, K. D., & Malle, B. F. (2014). Bringing free will down to Earth: People's psychological concept of free will and its role in moral judgment. *Consciousness and Cognition*, *27*(1), 100–108. https://doi.org/10.1016/j.concog.2014.04.011

Nichols, S. (2004). The folk psychology of free will: Fits and starts. *Mind & Language*, *19*(5), 473–502. https://doi.org/10.1111/j.0268-1064.2004.00269.x

Nichols, S. (2011). Experimental philosophy and the problem of free will. *Science*, *331*, 1401–1403. https://doi.org/10.1126/science.1192931

Nichols, S., & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Noûs*, *41*(4), 663–685. https://doi.org/10.1111/j.1468-0068.2007.00666.x

Reeder, G. D., Kumar, S., Hesson-McInnis, M. S., & Trafimow, D. (2002). Inferences about the morality of an aggressor: The role of perceived motive. *Journal of Personality and Social Psychology*, *83*(4), 789–803. https://doi.org/10.1037/0022-3514.83.4.789

Sarkissian, H., Chatterjee, A., De brigard, F., Knobe, J., Nichols, S., & Sirker, S. (2010). Is belief in free will a cultural universal? *Mind and Language*, *25*(3), 346–358. https://doi.org/10.1111/j.1468-0017.2010.01393.x

Saylor, M. M., Somanader, M., Levin, D. T., & Kawamura, K. (2010). How do young children deal with hybrids of living and non-living things: The case of humanoid robots. *British Journal of Developmental Psychology*, *28*(4), 835–851. https://doi.org/10.1348/026151009x481049

Schmidt, M. F. H., Rakoczy, H., & Tomasello, M. (2012). Young children enforce social norms selectively depending on the violator's group affiliation. *Cognition*, *124*(3), 325–333. https://doi.org/10.1016/j.cognition.2012.06.004

Schult, C. A., & Wellman, H. M. (1997). Explaining human movements and actions: Children's understanding of the limits of psychological explanation. *Cognition*, *62*(3), 291–324. https://doi.org/10.1016/S0010-0277(96)00786-X

Severson, R. L., & Carlson, S. M. (2010). Behaving as or behaving as if? Children's conceptions of personified robots and the emergence of a new ontological category. *Neural Networks*, *23*(8–9), 1099–1103. https://doi.org/10.1016/j.neunet.2010.08.014

Sobel, D. M. (2004). Exploring the coherence of young children's explanatory abilities: Evidence from generating counterfactuals. *British Journal of Developmental Psychology*, *22*, 37–58. https://doi.org/10.1348/026151004772901104

Sommer, K., Nielsen, M., Draheim, M., Redshaw, J., Vanman, E. J., & Wilks, M. (2019). Children's perceptions of the moral worth of live agents, robots, and inanimate objects. *Journal of Experimental Child Psychology*, *187*, 104656. https://doi.org/10.1016/j.jecp.2019.06.009

Starmans, C., & Bloom, P. (2016). When the spirit is willing, but the flesh is weak: Developmental differences in judgments about inner moral conflict. *Psychological Science*, *27*(11), 1498–1506. https://doi.org/10.1177/0956797616665813

Vaish, A., Missana, M., & Tomasello, M. (2011). Three-year-old children intervene in third-party moral transgressions. *British Journal of Developmental Psychology*, *29*(1), 124–130. https://doi.org/10.1348/026151010x532888

Weisman, K., Dweck, C. S., & Markman, E. M. (2017). Rethinking people's conceptions of mental life. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(43), 11374–11379. https://doi.org/10.1073/pnas.1704347114

Wente, A. O., Bridgers, S., Zhao, X., Seiver, E., Zhu, L., & Gopnik, A. (2016). How universal are free will beliefs? Cultural differences in Chinese and U.S. 4- and 6-year-olds. *Child Development*, *87*(3), 666–676. https://doi.org/10.1111/cdev.12528

Woolfolk, R. L., Doris, J. M., & Darley, J. M. (2006). Identification, situational constraint, and social cognition: Studies in the attribution of moral responsibility. *Cognition*, *100*(2), 283–301. https://doi.org/10.1016/j.cognition.2005.05.002

Yang, F., & Frye, D. (2018). When preferences are in the way: Children's predictions of goal-directed behaviors. *Developmental Psychology*, *54*(6), 1051–1062. https://doi.org/10.1037/dev0000490

Zhao, X., Zhao, X., Gweon, H. & Kushnir, T. (2020). Leaving a choice for others: Children's evaluations of considerate, socially-mindful actions. Child Development, 92(4), 1238–1253. https://doi.org/10.1111/cdev.13480

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supporting information