

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Do Children Ascribe the Ability to Choose to Humanoid Robots?

Permalink

<https://escholarship.org/uc/item/7jr8x5rs>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Flanagan, Teresa

Rottman, Joshua

Howard, Lauren

Publication Date

2019

Peer reviewed

Do Children Ascribe the Ability to Choose to Humanoid Robots?

Teresa Flanagan (tmf87@cornell.edu)

Department of Human Development, Cornell University, 116 Reservoir Ave
Ithaca, NY 14853 USA

Joshua Rottman (jrottman@fandm.edu)

Department of Psychology, Franklin & Marshall College, P.O. Box 3003
Lancaster, PA 17604 USA

Lauren Howard (lauren.howard@fandm.edu)

Department of Psychology, Franklin & Marshall College, P.O. Box 3003
Lancaster, PA 17604 USA

Abstract

Investigating folk conceptions of choice and constraints have been problematic given that human actions are rarely considered constrained. In this paper, we utilize humanoid robots (more clearly influenced by determined programming) to empirically test children's developing concepts of choice and action. Using a series of agency attribution and choice prediction tasks, we examined whether children differentiate free will abilities between robots and humans. Results indicated that 5–7-year-old children similarly attributed the ability to choose to both a robot and human child. However, for moral scenarios, participants considered the robot's actions to be more constrained than the human. These findings demonstrate that children appear to hold a nuanced understanding of choice across agents and across context.

Keywords: choice attribution; human–robot interaction; free will; cognitive development

Introduction

Attributing another entity with the ability to make choices lies at the foundation of treating that individual as having moral responsibility and as being deserving of rights. As such it is critically important to understand when and why people attribute choice to others. Because it is difficult to imagine typical human actions being constrained or devoid of choice (Nichols, 2011), research on free will has been limited by participants' inabilities to conceive of a deterministic world (Sommers, 2010). In other words, since it is difficult for us to imagine our own actions as being constrained, it is similarly challenging to conceive of others' actions as being constrained. Investigating ideas of free will in non-human entities, such as robots, shows promise in terms of releasing participants from anecdotal notions of choice and constraint.

Modern technology has resulted in a growing presence of interactive robots (e.g., Siri, Alexa, Roombas), particularly in the environments of younger generations (Wei et al., 2011). Robots present an increasingly important category for which to investigate choice attribution, as they are known to be largely programmed by their designers. Previous research has indicated that adults are ambivalent about robots' capacities to make choices (Weisman, Dweck, & Markman, 2017). However, nothing is known about children's tendencies to

attribute robots with the freedom to choose, despite children now growing up in technologically rich environments.

From early in life, children understand the possibility of completing "alternative" actions, denoting a basic grasp of free will. For example, 10-month-old infants expect human agents to use different actions to obtain an object depending on whether there are physical constraints 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 et al., 2005; Dunfield & Kuhlmeier, 2010; Hamlin, Wynn, & Bloom, 2008). By the preschool years, children not only anticipate and react to alternative actions (e.g., Nichols, 2004), but are able to verbally generate alternative options when the main goal of an agent is constrained (Sobel, 2004). Thus, early in life, children show a relatively sophisticated understanding 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 will (see Kushnir, 2018). However, a reliance on experiments involving human agents as targets of judgment means that the boundaries of children's free will ascriptions have not been fully charted. Though some work has shown that children assign more freedom of choice to human agents than inanimate objects (Nichols, 2004), none have explicitly examined attribution of free will to humanoid robots.

As in adults (e.g., Kahn et al., 2012b), research has shown that children ascribe a mixture of animate and inanimate characteristics to humanoid robots, suggesting an ontological category that is functionally separate from either (Kahn et al., 2012a; Severson & Carlson, 2010). For example, children may assume that robots hold a certain level of intelligence and some sensory abilities (e.g., can think, can see, can be tickled), but not emotions or biological capabilities (e.g., can feel happiness, needs sleep, 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 agentic

beings. For example, 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). This work highlights the ways in which robots straddle the animate and inanimate worlds, making them particularly interesting as a test case for children's ascriptions of free will.

Importantly, it appears that children's understanding of free will, even for human agents, 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 et al., 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). As such, there are certain scenarios, particularly those relating to internal desires or moral decisions, that appear to muddle children's understanding of free will for human actors.

In the current study, we asked whether 5- to 7-year-old children's predictions of action and choice varied across target agent (human child or robot) and constraint scenario (No Constraint, Moral Constraint, Rational Constraint). Children in this age range undergo relevant changes in their free will beliefs and their perceptions of robots (Bernstein & Crowley, 2008; Kushnir et al., 2015). During testing, both the human and robot agents were introduced as being similarly likely to make a particular choice when no constraints were present (i.e., was either 'programmed to' or 'born to' play a certain game). Within each scenario, we explored whether children predicted that the agent would follow the typical, default object choice or would respond to the constraints and pick an alternate object.

Based on previous work exploring children's trait attributions to robots and humans (Chernyak & Gary, 2016; Kahn et al., 2012), along with children's differential reactions to context and constraint (Kushnir et al., 2015; Nichols, 2011), we hypothesized the following: Without constraints, participants would predict the default action for both the robot and the human agent, and each would be significantly above chance in this choice. With rational constraints (the default action being impossible to completely fulfill), participants would predict the default action significantly more for the robot than the human agent. In the robot condition, participants would predict the default action above or at chance, and in the human condition, participants would

predict the default action significantly below chance. With moral constraints (the default action causing harm), participants would predict the default action significantly more for the robot than the human agent. In the robot condition, participants would predict the default action above or at chance, and in the human condition, participants would predict the default action significantly below chance.

Method

Participants

The final sample consisted of 32 children, aged 5–7 years old ($M_{\text{age}} = 5.72$, $SD_{\text{age}} = 0.68$, 15 females, 26 White), who were recruited from a participant database and tested in a laboratory in a small city in the northeastern region of the United States. One additional child participated but was excluded due to a developmental disability.

Materials & Procedure

Participants were randomly assigned to one of two conditions (robot or human). In the robot condition, children were asked to watch and respond to the actions of a robot figure named Robovie. The robot was a black and white humanoid toy, approximately 35 cm tall. It was preprogrammed to complete 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 participants, named Billy. All stimuli were pre-recorded and presented via video on a Dell laptop so that the agent's actions and perceived agency could be matched across conditions.

Regardless of condition, all participants proceeded through the same paradigm to assess their understanding of free will across ontological kinds. Children watched the video of either the robot or human, during and after which all participants were asked to predict the agent's actions (default or alternative object choice) and asked to attribute choice to the agent (did they "choose to" do the action or not, adapted from Kushnir et al., 2015). Answers to these questions indicated whether children believed the agent could act against its default choice and respond to constraints in a way that indicated free will.

Video Paradigm

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 this has been found 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."/ "This is Billy. Billy is a kid") paired with a still picture of the agent in a children's room (see Figure 1). Then the agent performed two simple actions: dancing and

throwing a bucket. The next video clip presented the agent’s actions as determined, stating that the agent was either programmed (robot condition, “Robovie is programmed to know a lot about science and can play science games”) or born (human, “Billy’s parents are scientists, so Billy knows a lot about science and plays science games”) to play a certain type of game (science games). Following this presentation, it was reiterated that the agent *only* plays the science game, even if there are other games present. This conveyed to participants that the science game, given no other constraints, was the default choice for both agents. This also indicated that a choice to play an alternative game (e.g. a history game), given constraints, would require the agent to “override” their entrenched pattern of playing the default games.

Action Prediction After the introduction video, participants watched three further video segments that presented each of the constraint scenarios (No Constraint, Moral Constraint, Rational 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. As explained in 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 choice.

In the first video (No Constraint scenario), participants were asked to pick which of two games (default or alternative) the agent would play without any limitation. Since participants had previously been told that agent only plays the default game, we hypothesized that this question would elicit a default response without the need for inferring choice or free will. The second video (Moral Constraint scenario) was identical to the first, but with the limitation that the agent will be playing with another person and playing the default game would result in hurting that person’s feelings. In this video, it would be wrong for the agent to play the default game, thus requiring them to play the alternative game if they wanted to

stay within moral bounds. The third video (Rational Constraint scenario) asked the child to predict the game the agent would play if the default game was broken. In this video, it would be irrational to play the broken (unplayable) game, requiring them to play the alternative game in order to act rationally.

Choice Attribution After each video, the experimenter 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/alternative game, along with an open-ended prompt asking them why.

Coding

Children’s action predictions were coded for each of the video constraints. In the No Constraint scenario, picking the alternative game (rather than the default game) clearly indicated choice, as it went against the default pattern of behavior. However, as there was no obvious reason to select the alternative game, the alternative game was not expected in either the human or the robot condition. In the Moral Constraint scenario, picking the alternative game indicated a choice that was driven by a consideration of others’ feelings, whereas picking the default game indicated a disregard for others’ feelings. Thus, we expected participants would predict the agent to play the alternative game in the human condition, but not in the robot condition. In the Rational Constraint scenario, picking the alternative game indicated a choice that was driven by a rational consideration of which game was possible to play, whereas picking the alternative game indicated a disregard for rational considerations. Therefore, in this scenario we expected participants would predict the agent to play the alternative game in the human condition, but not in the robot condition.

Children’s responses to the choice question were coded for whether or not they responded that they agent “had to” or “chose to” play a certain game. The ability to choose was indicated by a response that the agent “chose to” play the game, regardless of which game the agent chose.

Results

Action Prediction Results

Across all three constraint scenarios, action predictions were explored by running an omnibus binomial test to determine whether the percentage of game prediction (default and alternative) differed from chance (50%). The percentage of the default game predicted was marginally lower than chance in the human condition ($p = .059$) and at chance in the robot condition ($p = .665$). A Mann-Whitney test indicated that the default game prediction did not differ by agent condition ($U = 1032, p = .301, r = .182$). Within each of the three constraint scenarios, action predictions were explored by running binomial tests to determine whether the percentage of game prediction (default and alternative) differed from chance (50%). Further, Mann-Whitney U tests were run to test for



Figure 1: Screenshots of video stimuli used in the robot condition. 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 child in place of the robot.

differences between agents (human and robot) in the predicted percentage of default game prediction in each of the constraint scenarios. Frequencies are presented in Figure 2.

Action Prediction: No Constraint We ran binomial tests on participants' prediction of the game that would be played (default or alternative), for both the human and robot agents. Participants overwhelmingly tended to predict that both the human and the robot would play the default game ($ps < .001$). This demonstrates that participants predicted the agent to act in accordance with the ways that it had always acted in the past (i.e., playing the default game that it was programmed or born to play), indicating that participants understood that there was a strong likelihood for both the human and the robot to select the default game and confirming that the participants understood the introduction video. A Mann-Whitney test indicated that the default game prediction did not differ by agent condition ($U = 120, p = .317, r = .18$).

Action Prediction: Moral Constraint A binomial test indicated that the percentage of the default game predicted in the human condition was significantly lower than chance ($p < .001$). Thus, participants believed the human would go against his desires in order to act morally. In contrast, the percentage of the default game predicted in the robot condition was not significantly different from chance ($p = .804$). This demonstrates that participants were unsure if a robot would go against its programming in order to act morally. A Mann-Whitney test indicated that the prediction of the default game was significantly higher for the robot condition than the human condition ($U = 80, p < .05, r = .42$).




Figure 2: Percentage of the participants' action prediction across conditions for each constraint. Asterisks signify predictions that are significantly different than chance.

Action Prediction: Rational Constraint Binomial tests indicated that the percentage of the default game predicted in both the human condition and the robot condition was significantly lower than chance ($ps < .001$). This demonstrates that participants believed that both agents would go against their programming or desires in order to act rationally. A Mann-Whitney test indicated that action prediction did not differ by condition ($U = 120, p = .317, r = .18$).

Choice Attribution Results

In each scenario, participants were asked whether the agent "chose to" or "had to" play the predicted game, regardless of game type (default or alternative). Across all three constraint scenarios, choice attributions were explored by running an omnibus binomial test to determine whether the percentage of choice attribution ("choose to" and "have to") differed from chance (50%). The percentage of "choose to" responses did not differ from chance in either the human condition or the robot condition (human: $p = .312$; robot: $p = .193$). A Mann-Whitney test indicated that the choice attribution did not differ by agent condition ($U = 1128, p = .836, r = .037$). Within each of the three constraint scenarios, choice attributions were explored by running binomial tests to determine whether the percentage of choice attribution ("choose to" and "have to") differed from chance (50%). Further, Mann-Whitney U tests were run to test for differences between agents (human and robot) in the predicted percentage of default game prediction in each of the constraint scenarios. Frequencies are presented in Figure 3.




Figure 3: Percentage of participants' attribution of choice across conditions for each constraint. Asterisks signify attributions that are significantly different than chance.

Choice Attribution: No Constraint Binomial tests indicated that the percentage of “choose to” responses in both the human condition and the robot condition did not differ from chance (human: $p = .210$; robot: $p = .454$). A Mann-Whitney test indicated that choice attribution did not differ by condition ($U = 120, p = .714, r = .065$).

Choice Attribution: Moral Constraint Binomial tests indicated that the percentage of “choose to” responses in both the human condition and the robot condition did not differ from chance (human: $p = .804$; robot: $p = .454$). Mann-Whitney test indicated that choice attribution did not differ by condition ($U = 120, p = .723, r = .063$).

Choice Attribution: Rational Constraint Binomial tests indicated that the percentage of “have to” responses in both the human condition and the robot condition was significantly higher than chance ($ps < .05$). Mann-Whitney test indicated that choice attribution did not differ by condition ($U = 120, p = .632, r = .111$). These findings indicate that the participants believed acting in accordance with physical limitations was a constraint on choice regardless of the agent acting, such that free will was significantly limited by this rational constraint.

Discussion

This research investigated children’s predictions and choice attributions about the actions of a robot or a human child. Results indicated that, overall, children tended to judge a humanoid robot as having a similar amount of freedom of choice as a human child. For example, without any constraints, or with a rational constraint (such as the default game being broken), participants predicted similar actions for both the human and the robot. However, when the robot had an opportunity to change its actions in order to avoid making another child sad, children judged the robot as less likely to act in this way as compared to the human.

Action Prediction

In the No Constraint scenario, participants significantly predicted the agent to play the default game, in both the human and robot condition. This demonstrates that children’s basic understanding of the introduction video, where it was made clear that the agent only plays the science game. In the Rational Constraint scenario, results demonstrated that participants believed that both the robot and the human could act against its programming in order to act rationally (i.e., to play the alternative game). This may be an indication of choice attribution, as children believe a robot can be responsive to reasons and is not entirely constrained by its programming (Fischer, 2006). In the Moral Constraint, however, participants believed the human would be responsive to reasons and act morally (i.e., play the alternative game), but they were at chance in the robot condition. This demonstrates that in moral situations, children were unsure if the robot could go against its programming. This could be explained in a number of ways.

Children may have thought that the robot did not “care” about moral reasons or they may have thought that the robot did not have the capacity to recognize moral reasons. The latter explanation could be due to the fact that children did not interact with the robot in this study, making the robot’s social capacities were ambiguous. Future research could include various interaction components between the participant and the robot, which might unveil the types of social capacities that are required for a robot to appear responsive to moral reasons.

Choice Attribution

Overall, our results indicate that children’s attribution of choice is not unitary across situations. Similar to previous work (Kushnir et al., 2015), it appears that children are sensitive to both the agent type and the context that an agent is presented in when attributing free will. Furthermore, these results suggest that robot programming is not *always* a constraint on freedom of choice for young children. Most importantly, these results demonstrate that children are sensitive to constraints, such that some constraints (e.g. physical impossibility) are more restrictive to an agent’s choice than other constraints (e.g. desires, morals), and finding consistent for both the robot and the human agent.

In the No Constraint scenario, we were surprised to see that participants did not attribute choice above chance to the human agent. This may have been due to the presentation in the introduction video. For example, similar to the robot, we presented the human as having an entrenched disposition to play the default game. However, unlike the robot, who was programmed by scientists to choose this game, the human’s “programming” was that his parents were scientists and that he played science games every day. Previous research has shown that child participants attribute this type of consistent (non-random) choice as denoting not just ‘programming’, but desires (Kushnir, Xu, & Wellman, 2010). Since participants were at chance in attributing choice to the human agent, this could mean that children varied in believing whether or not having a strong desire is a constraint on actions.

In contrast to the human condition, we did not anticipate that participants would attribute so *much* choice to the humanoid robot in the No Constraint scenario. Here, the percentage of children that said the robot “chose to” select the default game was not significantly different from chance, suggesting that approximately half of our participants gave some semblance of free choice to the robot agent. This could be due to the varied exposure children have to robots, which research has shown is correlated with children’s propensity to attribute agency (Bernstein & Crowley, 2008). Alternatively, this could also be an indication of children’s general understanding of choice under minimal constraints. Specifically, children varied in the amount of choice they thought an agent had if the agent was constrained to perform a certain action based on how the agent was programmed or raised. Future research should investigate what underlies this individual difference; for example, do children who have

more interactions with robots also attribute them more choice in unconstrained scenarios?

In the Moral Constraint scenario, children were at chance in attributing choice in both the human condition and robot condition. This demonstrates that participants were almost equally likely to say that a moral action was a constraint or a choice. These results dovetail with previous research that has shown that children under 7-years-old are less likely to say that people have a choice in moral actions in comparison to older children (Chernyak et al., 2013). For the Rational Constraint, results also followed previous research, such that children assumed the agent “had to” make a certain choice when there were physical constraints (Kushnir et al., 2015), regardless of agent type. While these results support previous work, they also extend previous findings to other agents. Since results were similar for the human condition and the robot condition, this demonstrates that children extend their existing beliefs of choice and constraint to a fundamentally different type of agent. Future research should investigate other agents, such as plants or animals, to see if choice attribution in light of moral and rational constraints is a general or agent specific attribution.

It is important to note that we don’t fully understand the way children were interpreting some of our events, particularly those relating to programming (“Robovie is programmed to know a lot about science”) or genetic inheritance (“Billy’s parents are scientists so he knows a lot about science”). However, previous research has shown that, starting at 5-years-old, children display knowledge of biological inheritance (Gimenez & Harris, 2002) and children understand a robot is programmed (Bernstein & Crowley, 2008). Furthermore, all participants were told multiple times that the agent (whether robot or human) only played the default game. Future research should explicitly measure children’s understanding of programming and inheritance in relation to their action prediction and choice attribution to an agent.

In sum, this research indicates that children are able to attribute choice to actions that are programmed or hard-wired. This suggests that “compatibilist” theories of free will – in which choice can be said to exist even in a fully determined universe (e.g., Fischer, 2006) – may be an intuitive aspect of folk psychology. This dovetails with previous studies that have advanced this claim, but which have based their conclusions on adults’ assessments of complicated thought experiments (e.g., Nahmias et al., 2005). The present research has shown this to be true in children and for an everyday case of pre-determinism. New advances in technology, which have introduced robots into children’s everyday environments, have not only improved quality of living but have also allowed for improvements in testing for folk attributions of choice to agents that straightforwardly exist in a constrained environment.

Acknowledgments

We would like to thank the Nissley Scholar Grant at Franklin & Marshall College for funding this project. We would like

to thank Julia McAleer for her work as a research assistant. Also, we thank the families from the Lancaster community for participating in this project.

References

- Behne, T., Carpenter, M., Call, J., & Tomasello, M. (2005). Unwilling versus unable: Infants’ understanding of intentional action. *Developmental Psychology, 41*, 328-337.
- Bernstein, D., & Crowley, K. (2008). Searching for signs of intelligent life: An investigation of young children’s beliefs about robot intelligence. *The Journal of the Learning Sciences, 17*(2), 225-247.
- 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.
- 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.
- 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.
- Dunfield, K. A., & Kuhlmeier, V. A. (2010). Intention-mediated selective helping in infancy. *Psychological Science, 21*, 523-527.
- Fischer, J. M. (2006). Responsiveness and moral responsibility. *My Way: Essays on Moral Responsibility*. New York, NY: Oxford University Press, 63-83.
- Gimenez, M., & Harris, P. (2002). Understanding constraints on inheritance: Evidence for biological thinking in early childhood. *British Journal of Developmental Psychology, 20*, 307-324.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2008). Social evaluation by preverbal infants. *Nature, 450*, 557-559.
- Jipson, J. L., & Gelman, S. A. (2007). Robots and rodents: Children’s inferences about living and nonliving kinds. *Child Development, 78*(6), 1675-1688.
- Kahn, P. H., Kanda, T., Ishiguro H., Freier, N. G., Severson, R. L., Gill, B. T., Ruckert, J. H., & Shen, S. (2012a). “Robovie, you’ll have to go into the closet now”: Children’s social and moral relationship with a humanoid robot. *Developmental Psychology, 48*(2), 303-314.
- Kahn, P. H., Severson, R. L., Kanda, T., Ishiguro, H., Gill, B. T., Ruckert, J. H., ... Freier, N. G. (2012b). Do people hold a humanoid robot morally accountable for the harm it causes? In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction - HRI '12* (p. 33). New York, New York, USA: ACM Press.
- Kushnir, T. (2018). The developmental and cultural psychology of free will. *Philosophy Compass, 13*(11), 1-17.
- 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.

- 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.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2005). Surveying freedom: Folk intuitions about free will and moral responsibility. *Philosophical Psychology*, 18(5), 561-584.
- Nichols, S. (2004). The folk psychology of free will: Fits and starts. *Mind & Language*, 19(5), 473-502.
- Nichols, S. (2011). Experimental philosophy and the problem of free will. *Science*, 331, 1401-1403.
- 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.
- 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.
- 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.
- Sommers, T. (2010). Experimental philosophy and free will. *Philosophy Compass*, 5(2), 199-212.
- 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.
- Wei, C. W., Hung, I. C., Lee, L., & Chen, N. S. (2011). A joyful classroom learning system with robot learning companion for children to learn mathematics multiplication. *Turkish Online Journal of Educational Technology*, 10(2), 11-23.
- Weisman, K., Dweck, C. S., & Markman, E. M. (2017). Rethinking people's conceptions of mental life. *Proceedings of the National Academy of Sciences*, 201704347.
- Yang, F., & Frye, D. (2018). When preferences are in the way: Children's predictions of goal-directed behaviors. *Developmental Psychology*, 54(6), 1051-1062.