

# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

### Title

Underdetermination and Obligation Rules: Adult and Children's use of Closure Principles in Moral Learning

### Permalink

<https://escholarship.org/uc/item/3179162t>

### Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

### Authors

Stegall, Jessa

Mikhail, John

Nichols, Shaun

et al.

### Publication Date

2023

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Underdetermination and Obligation Rules: Adult and Children's use of Closure Principles in Moral Learning

**Jessa Stegall** (jessa.stegall@duke.edu)

Department of Psychology and Neuroscience, 417 Chapel Drive  
Durham, NC 27708 USA

**John Mikhail** (john.mikhail@law.georgetown.edu)

Georgetown University Law Center, 600 New Jersey Ave., NW  
Washington, D.C. 20001

**Shaun Nichols** (sbn44@cornell.edu)

Department of Philosophy, Sage School of Philosophy, 218 Goldwin Smith Hall  
Ithaca, NY 14853

**Tamar Kushnir** (tamar.kushnir@duke.edu)

Department of Psychology and Neuroscience, 417 Chapel Drive  
Durham, NC 27708 USA

## Abstract

Moral learners face an underdetermination problem – the rules they are taught cannot account for all novel cases. One way learners solve this problem is through closure principles, through which a learner assumes that anything that isn't explicitly forbidden is permitted, and vice versa. The current work aims to explore whether closure rules are used when reasoning about obligations. Building on previous work, we ask 1) whether adults and children use obligation rules in a similar manner to rules about permissibility and impermissibility, and 2) how early in life these inference abilities emerge. Across two studies, we explore inferences about obligation from both adults ( $N = 120$ ,  $M_{\text{age}} = 33.73$  years) and children ( $N = 103$ ,  $M_{\text{age}} = 5.52$  years). We found that while both adults and children rationally learn closure principles consistent with deontic logic, children and adults make opposite inferences about novel cases when provided with obligation rules.

**Keywords:** moral learning; cognitive development; deontic logic; obligation

## Introduction

Much of moral reasoning within philosophy relies on normatively specified rules of deontic logic. The deontic logic system has three main concepts: permissibility, obligation, and impermissibility.<sup>1</sup> These three concepts provide a baseline of reason to most moral, legal, and religious systems. They also bear logical relationships to one another, allowing information about one to give valuable information about another (Mikhail 2007).

From the time we are young, we learn about moral rules by hearing that certain things are permissible and that others are prohibited. For example, if a parent tells a child “You can't

play in the living room”, the child will learn that playing in the living room is forbidden. If the parent says the opposite, the child will learn that playing in the living room is permissible. But teaching a small set of rules does not tell the learner how explicit rules extend to novel cases, nor does it inform learners of what they are obliged to do. This is referred to as the *underdetermination problem*, where the information given to a learner will always be insufficient for determining what the learner should believe or do in a novel, unspecified situation.

One answer to the underdetermination problem is to specify *closure principles* (Mikhail, 2011; Nichols, 2021). Theories of deontic logic discuss two such principles (Mikhail, 2011). The Natural Liberty Principle (NLP) states that any action that isn't explicitly forbidden is permitted. To continue our example, if a parent teaches their child a rule such as “You can't play in the living room”, a child applying something like NLP (even implicitly) may infer that they may play anywhere else in the house. The Residual Prohibition Principle (RPP) states that any action that isn't explicitly permitted is forbidden. Under RPP, a child taught that “You can play in the living room”, may infer that they may *only* play in the living room and nowhere else in the house.

Permissibility is not the only important aspect of intuitive moral reasoning. We often do things out of a sense of obligation, even when we would prefer to do something else. Obligation tends to be upheld in agreement-like social interactions and impacts our everyday moral decisions (Tomasello, 2020). Like rules of permissibility and impermissibility, obligation may have an underlying logical structure that helps moral learners determine what they are and are not obliged to do, as well as helping to determine

<sup>1</sup> There are four main concepts within deontic logic, but we are choosing to leave out “omissible rules”. The reason for this is because the three discussed deontic concepts encapsulate

information gained from omissible rules and would therefore be redundant to include in our studies.

what further inferences can be learned from explicit obligation rules.

The relationship between obligation and permission is underdetermined in deontic logic. Rules about obligation have logical implications on permissibility – if you must play in the living room, that logically implies that it is also permissible for you to play in the living room (“ought” implies “may”). However, this implication does not hold in the opposite direction. If it is permissible for you to play in the living room, it does not follow that you have an obligation to do so.

The inability to determine if permission implies obligation extends to novel actions – if you are told what is permissible, you may still not know whether anything else you might want to do is obligatory. Likewise, if you are given explicit obligation rules, you may not know whether a novel action is permissible. The process by which learners infer what they may, must not, and must do in novel, unspecified cases is the central question of the current work.

In a recent study, Nichols and Gaus (2018) showed that U.S. adults can in principle infer NLP and RPP rationally from evidence and a set of domain-general assumptions about pedagogical sampling (the teacher of the rules is knowledgeable and maximally informative, e.g., Bonawitz et al., 2011; Shafto et al., 2014). In one study, participants were presented with a fictional rule system (the “rules of mice”) that had only two rules: either two actions were permitted, or both were forbidden, depending on condition. Results showed that when trained on two permission rules (i.e., mice are permitted to be in the red barn and mice are permitted to be in the yellow barn), participants reliably inferred that a novel action (i.e., entering the green barn) was not permissible, which is consistent with the Residual Prohibition Principle. When participants were trained on two prohibition rules (i.e., mice are not permitted to be in the red barn and mice are not permitted to be in the yellow barn), participants reliably inferred that a novel action (i.e., entering the green barn) was permissible, consistent with the Natural Liberty Principle.

Recent work has shown that other moral generalizations are equally learnable. For example – adults can infer the scope of rules, knowing to whom they do and do not apply beyond the examples they are given. Partington et al., (2023) showed that adults can infer the scope of a norm (whether it applies broadly or narrowly) from a handful of instances of impermissibility and can do so rationally.

Prior work also leaves open the question of whether young learners can use rational inference mechanisms to infer closure rules. Some support for this comes from Partington et al. (2023), who found that by age 7 children infer the scope of rules rationally from statistical evidence in the same way adults do. However, no study to date has looked at the use of closure rules throughout development, nor at the underdetermined relationship between all three main deontic logic concepts (permissibility, impermissibility, obligation).

The current studies address the following three questions about how obligation may influence inferences about novel

actions: 1) Do we infer anything about what is obligatory from evidence of what is permitted or forbidden? 2) What, if anything, can we infer from evidence of obligation rules? Finally, 3) Do these intuitions change across development or remain the same? Further, we discuss why inferences about obligation may differ from other deontic logic rules, as well as why this might lead to opposite intuitions about what is permissible or prohibited for different people.

## Experiment 1: Adults

Using the general methods of Nichols & Gaus (2018) we explored adults’ inferences about obligation from two directions – first if a novel action is inferred to be obligatory from *permission* or *prohibition* rules, then what can be inferred from explicitly given *obligation* rules. Experiment 1A examined whether adult participants use permitted and forbidden rules to make inferences about whether a novel action is obligatory. If obligations about novel actions are inferred the same way as their permissibility or impermissibility, then we will see an inference pattern emerge that is consistent with either NLP or RPP.

In Experiment 1B, we provided adults with obligation rules to see if obligations are used to make inferences about the permissible, impermissible, or obligatory status of novel actions.

### Experiment 1A

#### Method

**Participants.** For experiment 1A, eighty 18- to 66-year-old adults ( $M_{age} = 33.8$  years,  $SD = 11.51$ ) participated in the study using the online recruiting system Prolific. The current sample consisted of 51% women, 41% men, 4% non-binary, and 4% identified as other. One additional person was tested but not included in the dataset, as they chose to opt-out before completing the task.

**Procedure.** In a between-subjects design, participants were randomly assigned to one of two conditions: *Permission Rules* ( $N = 40$ ) or *Prohibition Rules* ( $N = 40$ ). In each condition, participants read the same opening vignette: “There is a Farm with lots of mice, and all the mice are supposed to follow The Rules of Mice. The rules are taught to them by one of the older mice. The Farm has four barns: Red, Blue, Yellow, and Green. The Rules of Mice Book has only two rules.” Participants in the Permission Rules condition read the following two rules about permissibility:

- 1) Mice are allowed to visit the Red Barn.
- 2) Mice are allowed to visit the Yellow Barn.

Participants in the Prohibited Rules condition read two rules about impermissibility:

- 1) Mice must not visit the Red Barn.
- 2) Mice must not visit the Yellow Barn.

Participants in both conditions were then asked the same three test questions (concerning *permission*, *obligation*, and *prohibition*) in a fixed order using a 1–6-point Likert scale. All the questions concerned another novel action (i.e.,

visiting the green barn), which was not part of the original set. For the first question, participants were asked to respond to a permission question by indicating “the extent to which you think it is permissible for Marky to visit the Green Barn”, where 1 was “not permissible” and 6 was “permissible”. Participants were then asked to respond to an obligation question by indicating “the extent to which you think it is obligatory for Marky to visit the Green Barn” where 1 was “not obligatory” and 6 was “obligatory”. Finally, participants were asked to respond to a prohibition question by indicating “the extent to which you think it is forbidden for Marky to visit the Green Barn”, where 1 was “not forbidden” and 6 was “forbidden”. Test question phrasing matched Nichols and Gaus (2018).

## Results

**Permissibility and Impermissibility Questions.** Our first set of analyses replicate the permissibility and impermissibility results from Nichols & Gaus (2018). As expected, a t-test showed that participants in the Permission Rules condition rated novel actions less permissible than participants in the Prohibition Rules condition (Permission Rules:  $M = 1.93/6$ ,  $SD = 1.57$ , Prohibition Rules:  $M = 5.74/6$ ,  $SD = 1.08$ ;  $t(58) = 13.89$ ,  $p < .001$ ). Using the midpoint of 3.5 as the hypothetical mean, one-sample t-tests confirmed that on average, participants in the Permission Rules condition rated novel actions as less permissible than would be expected by chance, while participants in the Prohibition Rules condition rated a novel action as more permissible (Permission Rules:  $t(39) = 6.19$ ,  $p < .001$ , Prohibition Rules:  $t(39) = 18.53$ ,  $p < .001$ ).

Answers for the impermissibility question followed the opposite pattern. A t-test showed that participants in the Permission Rules condition rated a novel action as more prohibited than those in the Prohibition Rules condition (Permission Rules:  $M = 5.02/6$ ,  $SD = 1.52$ , Prohibition Rules:  $M = 1.52/6$ ,  $SD = 1.48$ ;  $t(76) = 11.10$ ,  $p < .001$ , see Figure 1). In addition, participants in the Permission Rules condition on average rated a novel action as more prohibited than those in the Prohibition Rules condition (Permission Rules:  $t(39) = 6.18$ ,  $p < .001$ , Prohibition Rules:  $t(39) = 9.81$ ,  $p < .001$ ).

**Obligation Test Question.** Our focal question concerned inferences about obligation from permission and prohibition rules. A t-test showed that participants in the Permission Rules condition rated a novel action as less obligatory than those in the Prohibition Rules condition (Permission Rules:  $M = 1.34/6$ ,  $SD = 0.85$ , Prohibition Rules:  $M = 2.36/6$ ,  $SD = 1.72$ ;  $t(57) = 3.50$ ,  $p < .001$ , see Figure 1). Though participants’ inferences across conditions were different, they were also low endorsements of the obligation to perform the new action overall: one sample t-tests against the midpoint of the scale (3.5/6) showed that the average ratings were significantly lower than the midline in *both* the Permission and Prohibition conditions (Permission Rules:  $t(39) = 15.74$ ,  $p < .001$ ; Prohibition Rules:  $t(39) = 3.91$ ,  $p < .001$ ).




Figure 1. Adult participant ratings of a novel action (“visiting the green barn”) as permissible, obligatory, or impermissible in each condition in Experiment 1A.

## Discussion

Responses to the obligation question showed that participants made higher obligatoriness ratings for a novel action when given two prohibition rules rather than two permission rules. Importantly, both averages fell significantly below the midpoint, implying that overall, participants were not inferring that a novel action is obligatory based on a previously given set of permission or prohibition rules. This generally supports the idea that participants may be using a different inferential process when it comes to inferring an obligation – participants do not infer obligations in new cases from either rule emphasizing what is permitted or rules emphasizing what is forbidden. In both the Permission and Prohibition Rules conditions, participants ranked a novel action as not obligatory when compared to chance. These results leave an open question as to why there are differences between the Permission and Prohibition conditions on the obligation test question, and we discuss potential limitations and future directions in the discussion. In addition, these results replicate Nichols & Gaus (2018) and are consistent with the idea that participants infer closure rules rationally based on the available evidence. When given rules permitting certain actions, participants inferred other actions are prohibited, consistent with the Residual Prohibition Principle. Conversely, given rules that certain actions are forbidden, participants infer that novel actions are permissible, consistent with the Natural Liberty Principle.

In Experiment 1A, we examined whether participants inferred that a novel action was obligatory from permission or prohibition rules. In study 1B, we provided participants with obligation rules to see if these rules are used to make inferences about the permissible, impermissible, or obligatory status of a novel action. We were interested to see if obligation rules led participants to make inferences similar to either permission or prohibition rules - or if they created a new inference pattern entirely.

## Experiment 1B

In order to explore how adult participants use obligation rules to make inferences about novel actions, we used the same vignette from 1A to explore a novel *Obligation Rules* condition. Participants were given two obligation rules and were asked the permissible, impermissible, and obligatory character of a novel action based on those rules.

## Method

**Participants.** Forty 18- to 67-year-old adults ( $M_{age} = 33.58$ ,  $SD = 13.63$ ) participated through the online recruitment system Prolific. The current sample consisted of 58% women, 35% men, 3% non-binary, and 2% identified as other.

**Procedure.** Participants were told an identical vignette to Experiment 1A, but were given two obligation rules, rather than two permission or prohibition rules. Participants in this Obligation condition encountered the following two rules:

- 1) Mice must visit the Red Barn.
- 2) Mice must visit the Yellow Barn.

The rest of the methods were identical to Experiment 1A. After given the two obligation rules, participants were asked about the permissible, impermissible, and obligatory status of a novel action in a fixed order on a 1-6 Likert scale.

## Results

Figure 2 shows the results for all three test questions. A one-way ANOVA showed a significant main effect of question ( $F(2, 117) = 20.84$ ,  $p < .001$ ). Follow-up t-tests showed that ratings for the obligation question ( $M = 1.32/6$ ,  $SD = 1.08$ ) were significantly lower than both the permissibility and impermissibility questions (permissibility question:  $M = 4.02/6$ ,  $SD = 2.27$ ;  $t(39) = 6.93$ ,  $p < .001$ ; impermissibility question:  $M = 2.68/6$ ,  $SD = 2.13$ ;  $t(39) = 3.60$ ,  $p < .001$ ). One sample t-tests showed that obligation ratings were significantly lower than the midpoint of the scale ( $t(39) = 14.14$ ,  $p < .001$ ). Participants were slightly more likely to judge a novel action as permissible when compared to the midline ( $t(39) = 1.37$ ,  $p = ns$ ), and significantly more likely to say novel actions were not prohibited ( $t(39) = 2.51$ ,  $p = .016$ ).

## Discussion

In Experiment 1B, we sought to determine if, after learning that two actions are obligatory, participants would make inferences about the deontic status of novel actions. Our data suggests adult learners infer that actions which are not obligatory are permissible; participants in our sample believed that agents were permitted to – and even more strongly, and not forbidden from – engaging in novel actions so long as their obligations were met. Our data suggests that adult moral learners infer Natural Liberty from explicitly stated obligations similarly to how they do for explicitly stated permission.




Figure 2. Adult participant ratings of a novel action (“visiting the green barn”) as permissible, obligatory, or impermissible the Obligation condition in Experiment 1B.

## Experiment 2: Children

In Experiment 2, we sought to both replicate the findings in Experiments 1A and 1B with children and explore the developmental trajectory of inference abilities in deontic logic. We designed a child-friendly version of the previous experiments to determine if young children infer the closure rules of NLP and RPP in a similar way as adults. We simplified the adult vignette and adjusted the language of both the condition rules and the test questions so that it matched ordinary ways in which children hear about permission and obligation in daily life. For that reason, children heard “okay”, “not okay”, and “have to”, in place of “permissible”, “impermissible”, and “obligatory”, respectively (Göckeritz et al., 2014).

## Method

**Participants.** 4- to 6-year-old children ( $N = 103$ ,  $M_{age} = 5.52$ ,  $SD = 0.84$ ; 40% girls, data collection ongoing) were tested either online via Zoom, or in-person in local museums and lab spaces. Seven additional children were tested but were excluded (4 failed memory checks, 2 parental interference, 1 failed colorblindness check). A preregistered *a priori* power analysis indicated that we need 120 children in our final sample to achieve sufficient power.

**Procedure.** In a between-subjects design, children were randomly assigned to one of three conditions: *Permission*, *Prohibition*, or *Obligation*. All children saw an identical story about a mouse on a farm through PowerPoint either online via Zoom screensharing, or in person on a 13” MacBook. After a brief warm-up, children were told that the mice who lived on the farm had rules, and that the rules were taught to them by the older mice. Children in the *Permission* condition heard two rules about permissibility (i.e., “It’s okay for mice to go into the red barn”, and “It’s okay for mice to go into the yellow barn”); participants in the *Prohibition* condition heard two rules about impermissibility (i.e., “It’s not okay for mice

to go into the red barn”, and “It’s not okay for mice to go into the yellow barn”); and participants in the *Obligation* condition heard two rules about obligation (i.e., “Mice have to go into the red barn” and “Mice have to go into the yellow barn”).

Participants were then introduced to “Marky”, who was described as another mouse who lived on the farm and knew the rules of the farm. Finally, in all three conditions, participants heard the same two test questions in a counterbalanced order. Participants heard the *permissibility question*: “Is it okay for Marky to go into the green barn?” and the *obligation question*: “Does Marky have to go into the green barn?”.

## Results

**Permissibility Question.** Figure 3 shows the proportion of children who indicated that a novel action was permissible and/or required, across age group and by condition. For children’s inferences about the permissibility of novel actions, we constructed a Generalized Linear Model assuming a binary logistic response on whether condition predicted children’s answer to the permissibility question. We considered condition (Permission, Prohibition, and Obligation) and age as the independent variables in this analysis and built a factorial model. Our preliminary analyses suggested there was no main effect of condition, Wald  $\chi^2(2, N = 103) = 5.35, p = .067$ , and no main effect of age, Wald  $\chi^2(1, N = 103) = 0.02, p = .883$ . However, there was a significant interaction between condition and age, Wald  $\chi^2(2, N = 103) = 10.99, p = .004$ . Like adults, older children in the Prohibition condition were more likely to say a novel action was permissible than younger children, consistent with the NLP ( $OR = 7.1, p = .004$ ).

We then compared children’s responses to chance to see whether there was consistency in their permissibility judgments within each condition. Children in the Permission condition are marginally more likely than chance to say that novel action is impermissible (23/34 cases, Binomial test,  $p = .057$ ). Children in the Prohibition condition are significantly more likely than chance to say a novel action is permissible (25/35 cases, Binomial test,  $p = .012$ ). Children in the Obligation condition were significantly more likely than chance to say that a novel action was impermissible (23/33 cases, Binomial test,  $p = .035$ ).

**Obligation Question.** We ran a Generalized Linear Model assuming a binary logistic response on whether condition predicted children’s answer to the obligation question. We found no main effect of condition, Wald  $\chi^2(2, N = 103) = 1.91, p = .385$ , and no main effect of age, Wald  $\chi^2(1, N = 103) = 2.21, p = .137$ . Further, there was no interaction between condition and age, Wald  $\chi^2(2, N = 103) = 1.85, p = .396$ . Across the board, children, like adults, indicated that a novel action (i.e., entering the green barn) was not obligatory based on explicit rules about similar but distinct actions (i.e., entering the red and yellow barns) regardless of the type of rule (permission, prohibition, or obligation) or age.




Figure 3. Children’s ratings of a novel action (e.g., “visiting the green barn”) as (im)permissible or obligatory in each condition in Experiment 2.

We then compared children’s responses to chance to see whether there was consistency in their obligation judgments within each condition. Children in the Permission condition were significantly less likely than chance to say that novel action was obligatory (Consistent with RPP, 9/34 cases, Binomial test,  $p = .009$ ). Similarly, children in the Obligation condition were significantly less likely than chance to say that a novel action was obligatory (7/33 cases, Binomial test,  $p = .001$ ). In contrast to these, about half of the children in the Prohibition condition said that the novel action was obligatory (13/35 cases, Binomial test,  $p = ns$ ).

## Discussion, Experiment 2

Our data suggests that children are both similar to and different from adults in their inferences about closure rules. Like adults, explicitly teaching that two actions are permitted or forbidden led 4- to 6-year-old children in our sample to make inferences about the permissibility and impermissibility of novel actions consistent with deontic logic. Children as young as 4 in the Permission condition inferred Residual Prohibition. With age, adult-like inferences about Natural Liberty emerged in the Prohibition condition. Like adults, children in our sample did not infer obligation in any of the three conditions. All together, these results are the first to show that young moral learners can make rational inferences about deontic rules, and that these inferences emerge between ages 4 and 6.

However, we also found at least one important difference between young learners and their adult counterparts. When deciding if a novel action was permissible or not given obligation rules, children, unlike adults, inferred that a novel action was not permissible. This suggests children might be cautious about performing anything other than explicitly stated obligations, a finding we discuss further below.

## General Discussion

The present studies aimed to investigate the way both adult and child moral learners solve the underdetermination

problem. After being explicitly taught two moral rules about what is permissible, forbidden, or obligatory, what if anything do learners infer about novel cases?

Our findings replicate past work showing that adults solve the underdetermination problem through a process of rational learning in a manner consistent with deontic logic, and we extend prior work by showing that these inferences emerge between ages 4 and 6. But also, our findings point to interesting developmental differences when it comes to what learners infer about novel actions from explicitly taught obligations. While adult learners endorse Natural Liberty for novel actions, young learners lead towards a more conservative Residual Prohibition for novel actions. This finding is consistent with several findings on early normativity: first, that young children tend to enforce strict norms (Tomasello, 2019) and that young children tend to conflate *ought* with *can* (and *oughtn't with can't*, e.g., Shtulman & Phillips, 2018; Chernyak et al., 2013, 2019). For example, children as young as 3 years will infer a norm even when one is not explicitly expressed and this will influence their behavior and norm-enforcing (Schmidt et al., 2016). Either of these reasons might lead children to infer Residual Prohibition from learning about two actions which are obligatory.

It is also possible that the developmental differences we found here are specific to Western cultural contexts, where all else equal adults are more inclined towards individualism and valuing autonomy and choice (Hui & Villareal, 1989; Kim & Sherman, 2007). It may be the case that U.S. adults are working from this individualistic perspective to infer NLP from obligation rules – if one's explicit obligations are fulfilled, then one is free to engage in any novel action (and therefore all are permissible and not prohibited). In cultures with “tight” systems of norms (Gelfand et al., 2011) that are strictly adhered to, adults might be more inclined to infer that explicit obligations restrict future action (Atari et al., 2022).

Across studies, we also found that neither adults nor children inferred obligations from being explicitly taught any permission, prohibition, or obligation rules. This is also consistent with deontic logic – learning about obligations licenses inferences about permission, but not vice versa. It may be the case that obligations must always be explicitly taught, although this seems unlikely from previous work finding that young children judge that others are obliged to help people even in the case where the helpee is a novel person (Marshall et al., 2022). Questions therefore remain about what leads learners to infer that a novel action is obligatory when it isn't explicitly taught. One answer is that we generalize obligations broadly (such as when we infer broad, inclusive norms from a few cases e.g., Partington et al., 2023). Another answer is that we infer novel obligations to maintain a self-consistent moral identity (as a “good” or “moral” person). Indeed, prior work has shown that identity can be a source of moral motivation (Hardy & Carlo, 2005). Exploring these alternative pathways to learning what counts as obligatory, and how these changes with age, are interesting questions for future research.

Taken together, these findings suggest both adults and children rationally learn closure principles consistent with deontic logic, and that there are developmental changes in moral learning specific to our concept of obligation.

## Acknowledgments

We thank Emma Xiong, Athena Wells, and Bri'Yon Watts for help with data collection, and Danielle Schugars for stimuli creation. We also thank the staff of the Durham Museum of Life and Science for their partnership and ongoing support of our research. Finally, we thank the children and families who participated in this project.

## References

- Atari, M., Haidt, J., Graham, J., Koleva, S., Stevens, S. T., & Dehghani, M. (2022). Morality beyond the WEIRD: How the nomological network of morality varies across cultures. *PsyArXiv*. <https://doi.org/10.31234/osf.io/q6c9r>
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322-330. <https://doi.org/10.1016/j.cognition.2010.10.001>
- Chernyak, N., & Kushnir, T. (2013). Giving preschoolers choice increases sharing behavior. *Psychological Science*, 24(10), 1971-1979. <https://doi.org/10.1177/0956797613482335>
- Chernyak, N., Kang, C., & Kushnir, T. (2019). The cultural roots of free will beliefs: How Singaporean and US Children judge and explain possibilities for action in interpersonal contexts. *Developmental Psychology*, 55(4), 866. <https://doi.org/10.1037/dev0000670>
- Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., Duan, L., Almaliah, A., Ang, S., Arndt, J., Aycan, Z., Boehnke, K., Boski, P., Cabecinhas, R., Chan, D., Chhokar, J., D'Amato, A., Ferrer, M., Fischlmayr, I. C., Fischer, R., ... Yamaguchi, S. (2011). Differences between tight and loose cultures: a 33-nation study. *Science (New York, N.Y.)*, 332(6033), 1100-1104. <https://doi.org/10.1126/science.1197754>
- Hardy, S. A., & Carlo, G. (2005). Identity as a Source of Moral Motivation. *Human Development*, 48(4), 232-256. <https://doi.org/10.1159/000086859>
- Hui, C. H., & Villareal, M. J. (1989). Individualism-collectivism and psychological needs: Their relationships in two cultures. *Journal of Cross-Cultural Psychology*, 20(3), 310-323. <https://doi.org/10.1177/0022022189203005>
- Kim, H. S., & Sherman, D. K. (2007). "Express Yourself": Culture and the Effect of Self-Expression on Choice. *Journal of Personality and Social Psychology*, 92(1), 1-11. <https://doi.org/10.1037/0022-3514.92.1.1>
- Marshall, J., Gollwitzer, A., Mermin-Bunnell, K., Shinomiya, M., Retelsdorf, J., & Bloom, P. (2022). How development and culture shape intuitions about prosocial obligations. *Journal of Experimental Psychology*:



- General*, 151(8), 1866–1882.  
<https://doi.org/10.1037/xge0001136>
- Mikhail, J. (2007). Universal Moral Grammar: Theory, evidence, and the future. *Trends in Cognitive Sciences*, 11(8), 143–152. <https://doi.org/10.1016/j.tics.2006.12.007>
- Mikhail, J. (2011). *The elements of moral cognition: Rawls' linguistic analogy and the cognitive science of moral and legal judgment*. Cambridge, UK: Cambridge University Press.
- Nichols, S. and Gaus, J. (2018), Unspoken Rules: Resolving Underdetermination With Closure Principles. *Cognitive Science*, 42(8). <https://doi.org/10.1111/cogs.12674>
- Nichols, S. (2021). *Rational rules: Towards a theory of moral learning*. Oxford University Press.
- Partington, S., Nichols, S., & Kushnir, T. (2023). Rational learners and parochial norms. *Cognition*, 233. <https://doi.org/10.1016/j.cognition.2022.105366>
- Schmidt, M. F. H., Butler, L. P., Heinz, J., & Tomasello, M. (2016). Young Children See a Single Action and Infer a Social Norm: Promiscuous Normativity in 3-Year-Olds. *Psychological Science*, 27(10), 1360–1370. <https://doi.org/10.1177/0956797616661182>
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89. <https://doi.org/10.1016/j.cogpsych.2013.12.004>
- Shtulman, A., & Phillips, J. (2018). Differentiating “could” from “should”: Developmental changes in modal cognition. *Journal of Experimental Child Psychology*, 165, 161–182. <https://doi.org/10.1016/j.jecp.2017.05.012>
- Tomasello, M. (2019). *Becoming Human: A Theory of Ontogeny*. Harvard University Press.
- Tomasello M. (2020) The Moral Psychology of Obligation. *Behavioral and Brain Sciences* 43(56), 1–58. <https://doi.org/10.1017/S0140525X19001742>