

Belief Updating in Moral Dilemmas

Zachary Horne, Derek Powell & Joseph Spino


**Review of Philosophy and
Psychology**

ISSN 1878-5158

Volume 4

Number 4

Rev.Phil.Psych. (2013) 4:705-714
DOI 10.1007/s13164-013-0159-y



Your article is protected by copyright and all rights are held exclusively by Springer Science +Business Media Dordrecht. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Belief Updating in Moral Dilemmas

Zachary Horne · Derek Powell · Joseph Spino

Published online: 22 September 2013
© Springer Science+Business Media Dordrecht 2013

Abstract Moral psychologists have shown that people's past moral experiences can affect their subsequent moral decisions. One prominent finding in this line of research is that when people make a judgment about the Trolley dilemma after considering the Footbridge dilemma, they are significantly less likely to decide it is acceptable to redirect a train to save five people. Additionally, this ordering effect is asymmetrical, as making a judgment about the Trolley dilemma has little to no effect on people's judgments about the Footbridge dilemma. We argue that this asymmetry is the result of a difference in how each dilemma affects people's beliefs about the importance of saving lives. In two experiments, we show that considering the Footbridge dilemma disconfirms these beliefs, while considering the Trolley dilemma does not significantly affect them. Consistent with predictions of sequential learning models, our findings offer a clear and parsimonious account of the asymmetry in the ordering effect.

A growing body of research in moral psychology has shown that people's past moral experiences can affect their subsequent moral decisions (Liao et al. 2011; Lombrozo 2009; Sinnott-Armstrong 2008; Wiegmann et al. 2008; Wiegmann and Okan 2012). One prominent finding in this line of research is that when people make a judgment about the Trolley dilemma after considering the Footbridge dilemma, they are significantly less likely to decide it is acceptable to redirect a train in order to save five people.¹ Additionally, this ordering effect is asymmetrical. Whereas making a judgment about

¹In the Trolley dilemma, participants are told they are at the wheel of a runaway train that will cause the deaths of five workmen if it proceeds on its present course. The only way to avoid the deaths of these workmen is to hit a switch, redirecting the train to a side track where it will kill a single workman instead. This dilemma is often contrasted with another famous moral dilemma, the Footbridge dilemma, which asks participants to consider a situation in which they are on a footbridge, in between a runaway trolley and five workmen who will be killed if nothing is done. Participants are then told that the only way to save the lives of the five workmen is to push a stranger off the bridge and onto the tracks below where his large body will stop the trolley.

Z. Horne (✉) · J. Spino
Department of Philosophy, University of Illinois - Urbana-Champaign, Urbana, USA
e-mail: horne2@illinois.edu

D. Powell
Department of Psychology, University of California - Los Angeles, Los Angeles, USA

the Footbridge dilemma affects people's judgments about the Trolley dilemma, making a judgment about the Trolley dilemma has little to no effect on their judgments about the Footbridge dilemma (Schwitzgebel and Cushman 2012). We call this the *Footbridge-Trolley ordering effect*.

Some have argued that the Footbridge-Trolley ordering effect provides critical insight into the mental processes that underlie people's moral judgments (Schwitzgebel and Cushman 2012), while others have interpreted this finding as being substantial evidence against the use of intuitions in ethical theorizing (Liao et al. 2011; Sinnott-Armstrong 2008; Tobia et al. 2012). Despite extensive empirical research and philosophical discussion of this effect, there is no consensus explanation as to why it occurs.

Researchers have proposed a number of explanations for ordering effects on moral judgments. Some researchers suggest that these ordering effects are produced when participants compare similar moral vignettes (Liao et al. 2011; Sinnott-Armstrong 2008; Wiegmann et al. 2008). It seems likely that a process of comparison plays a role in producing ordering effects. However, this suggestion does not offer an explanation of why the Footbridge-Trolley ordering effect is asymmetrical. Given that the dilemmas are so similar, it is unclear why comparing the Footbridge dilemma to the Trolley dilemma affects people's judgments about the Trolley dilemma but not vice versa.

Schwitzgebel and Cushman (2012) have proposed that this asymmetry is explained by a difference in the mental processes that are recruited when people make utilitarian or deontological moral judgments. These researchers assume a dual-process account of moral reasoning wherein different mental processes produce different types of moral judgments; emotional processing leads to deontological moral judgment whereas reasoning processes lead to utilitarian moral judgments. Schwitzgebel and Cushman argue that the emotional processes typically associated with deontological moral judgment are not sensitive to new evidence, such as a previously considered moral dilemma. In contrast, the reasoning processes typically associated with utilitarian moral judgment are responsive to new evidence. This would explain why people's judgments about the Trolley dilemma (utilitarian) are influenced by their previous judgment about the Footbridge dilemma (deontological), but that the converse does not hold.

As noted, Schwitzgebel and Cushman's account presumes that the Trolley and Footbridge dilemmas recruit fundamentally different mental processes. Cast in this light, the asymmetry in the ordering effect may be construed as additional evidence that different mental processes underlie deontological and utilitarian moral judgment. However, asymmetrical ordering effects have been widely reported in research on judgments of non-moral problems, and have been explained without appeal to differences in mental processes (Hogarth and Einhorn 1992). Moreover, some researchers have found evidence that moral judgments recruit domain general cognitive processes such as decision-making heuristics (Rai and Holyoak 2010), and mechanisms associated with causal reasoning (Waldmann et al. 2012). In contrast to Schwitzgebel and Cushman's proposal, these researchers argue that the field should seek domain-general explanations of moral behavior before positing more specialized mental processes that appeal specifically to the details of the dilemmas under consideration.

In the present paper we sought to examine whether domain general models of learning and judgment can provide insight into the Footbridge-Trolley ordering effect. Sequential learning models have been developed across a variety of domains, describing how people and (non-human animals) form stimulus-response

associations (Rescorla and Wagner 1972), how people learn causal relationships (Shanks 1985; Van Hamme and Wasserman 1994; Danks et al. 2003), and how they update their beliefs (Hogarth and Einhorn 1992). Sequential learning models describe how the state of some mental representation is affected by new experiences. In many models, the possible states of these mental representations are bounded their minimum and maximum states are often expressed as probabilities of 0 and 1 (Danks et al. 2003).

In order to constrain the possible states of mental representations within sequential learning models, the degree to which some experience affects a given mental representation is contingent on the current state of that representation. Specifically, the weight assigned to new experiences on a given mental representation is inversely proportional to the difference from the representation's current state to whichever bound it is approaching. So when a belief or association is strong, experiences supporting the representation are less influential than experiences that undermine it. And just the opposite is true when the representation is weak. This dynamic can lead to asymmetrical ordering effects.

We propose that sequential learning models can explain the Footbridge-Trolley ordering effect. In the next section, we consider Hogarth and Einhorn's (1992) Belief-Adjustment model, a sequential learning model that was explicitly developed to explain ordering effects. However, as discussed above, the general principles we summarized are at work in a number of models that describe a wide range of human (and nonhuman animal) behavior (Danks et al. 2003; Hogarth and Einhorn 1992; Rescorla and Wagner 1972; Shanks 1985; Van Hamme and Wasserman 1994).

1 A Belief-Adjustment Explanation

Hogarth and Einhorn's Belief-Adjustment model (1992) specifies how the state of some belief is affected by new evidence. The following equation describes the model:

$$S_k = S_{k-1} + w_k[s(x_k) - R].$$

S_k is the state of some belief after k pieces of evidence and S_{k-1} is the state of that belief just before the k^{th} piece of evidence. The variable $s(x_k)$ is the individual's subjective evaluation of the k^{th} piece of evidence. The variable R is a reference point against which the k^{th} piece of evidence is evaluated and w_k is the weight assigned to the difference between the individual's subjective evaluation and their reference point. Each of these terms takes on a value between 0 and 1.

To help frame the Footbridge-Trolley ordering effect in the context of prior research on ordering effects, we must consider what belief, if any, is adjusted when people consider these moral dilemmas. For naive participants, S_{k-1} cannot be about either the Trolley or the Footbridge dilemmas in particular, since participants have never considered either dilemma prior to participating in the experiment. Rather, S_{k-1} is a general moral belief that is recruited when one considers either of these moral situations. For example, a candidate belief might be "In life or death situations, one should always take whatever means necessary to save the most lives." In this case, participants' judgments about a dilemma constitute their evaluation of a new piece of evidence $s(x_k)$ that is relevant to their more general belief S_{k-1} . Hogarth and Einhorn's model states how these

evaluations affect S_{k-1} , resulting in an updated belief S_k . One standing issue is why the Footbridge and Trolley dilemmas affect some belief S_{k-1} to different degrees.

As in other sequential learning models, in the Belief-Adjustment model the effect of evidence is dependent on the current state of the belief, S . In the Belief-Adjustment model, when evidence is disconfirmatory the weight is

$$w_k = \alpha S_{k-1},$$

whereas if evidence is confirmatory, the weight is

$$w_k = \beta(1 - S_{k-1}).$$

Thus, the weight of disconfirmatory evidence is proportional to the strength of the belief (S_{k-1}), whereas the weight of confirmatory evidence is proportional to the value of $1 - S_{k-1}$. The parameters α and β ($0 < \alpha, \beta < 1$) represent an individual's sensitivity to disconfirmatory and confirmatory evidence, respectively. For now, we will assume they are both equal to 1.

For the sake of argument, suppose that participants' belief, S_{k-1} , represents a proposition like "In life or death situations, one should always take whatever means necessary to save the most lives." If this is the case, then it is reasonable to think people would endorse this belief strongly. That is, if we assign it a value in the model, it would be greater than .5. Under these conditions, the Belief-Adjustment model predicts that disconfirmatory evidence (the Footbridge dilemma) will have a stronger effect on this belief than equivalent confirmatory evidence (the Trolley dilemma). As a toy example, suppose S_{k-1} takes on a value of .75. The weight of disconfirmatory evidence is $1 \times .75 = .75$, whereas the weight of confirmatory evidence is $1 \times (1 - .75) = .25$. Suppose further that the Trolley and Footbridge dilemmas provide equivalent evidence for and against the belief, such that the contrast $[s(x_k) - R]$ is equal to .5 in both cases. Assuming these values, considering the Trolley dilemma will increase S_{k-1} by .125 ($S_k = S_{k-1} + .125$), whereas considering the Footbridge dilemma will decrease S_{k-1} by .375 ($S_k = S_{k-1} - .375$). This is merely an illustration, but whenever we assume some value for S_{k-1} greater than .5, the weight of disconfirmatory evidence will be greater than the weight of confirmatory evidence (all else being equal). Consequently, if people strongly endorse the beliefs that underlie their judgments about the Footbridge and Trolley dilemmas, then the asymmetry in the Footbridge-Trolley ordering effect could be a direct consequence of how people weigh evidence.²

In the next section of the paper, we provide support for the claim that making judgments about the Footbridge and Trolley dilemmas affects people's moral beliefs to different degrees, that these beliefs are strongly held, and consequently, that the Belief-Adjustment model (1992) can explain the asymmetry in the Footbridge-Trolley ordering effect.

² Another possible explanation is that some beliefs influence participants' judgments about the Footbridge dilemma but are not recruited when making a judgment about the Trolley dilemma. For instance, making a judgment about the Trolley dilemma might recruit a set of beliefs S_1 , S_2 , and S_3 , whereas making a judgment about the Footbridge dilemma recruits shared beliefs S_1 , S_2 , and S_3 , along with additional beliefs S_4 , S_5 , and S_6 . This means it is possible that consideration of the Footbridge dilemma affects all of the beliefs recruited when making judgments about the Trolley dilemma, but that consideration of the Trolley dilemma only affects some subset of the beliefs recruited when making judgments about the Footbridge dilemma. If this were right, then people's judgments about the Trolley dilemma would have a weaker effect on people's judgments about the Footbridge dilemma than vice versa.

2 Experimental Design

We sought to investigate how making judgments about the Footbridge and Trolley dilemmas affects participants' moral beliefs and the degree to which they endorse these moral beliefs. We conducted two experiments wherein participants made judgments either about the Footbridge dilemma, the Trolley dilemma, or a non-moral control dilemma (following a between-subjects design) and then rated the extent to which they agree with a statement of a moral belief. People's responses to moral dilemmas are often discussed in terms of "utilitarianism" and "deontology" (e.g., Greene et al. 2001; Prinz 2007; Wiegmann and Okan 2012). Accordingly, in Experiment 1 participants rated a statement generally representing a utilitarian viewpoint: "In life or death situations, you should always take whatever means necessary to save the most lives." In Experiment 2, participants rated a statement that represented a deontological viewpoint: "You should never kill another person." Participants' agreement ratings in control conditions were used to measure their endorsement of the utilitarian and deontological beliefs stated above. By comparing the agreement ratings of participants who made moral judgments about the Footbridge and Trolley dilemmas with ratings of participants' in the control condition, we were able to measure the degree to which considering these dilemmas confirm or disconfirm the beliefs that are recruited when making judgments about the Trolley and Footbridge dilemmas.

We predicted that people would endorse both the utilitarian and deontological beliefs. Additionally, we predicted that the Footbridge dilemma would disconfirm the utilitarian belief more than the Trolley dilemma would confirm it. This result would explain the asymmetry in the Footbridge-Trolley ordering effect. However, it is also possible that the Trolley dilemma affects other relevant beliefs more strongly than the Footbridge dilemma, which could undermine this explanation. In Experiment 2 we examined the effects these dilemmas on the belief "You should never kill another person" to explore this possibility.

3 Experiment 1

3.1 Method

3.1.1 Participants

In Experiment 1, 265 participants (97 female, mean age of 30 years) were recruited online, through the Amazon Mechanical Turk (mTurk) work-distribution website. To be eligible for the mTurk posting, workers had to reside in the U.S. and have at least a 95 % approval rate. Participants were compensated \$.20 for their participation.

3.1.2 Procedure

This experiment was conducted online, administered using Qualtrics survey software. Participants were required to pass a CAPTCHA test at the beginning of the experiment to eliminate the possibility of automated computer responses.

After reading instructions, participants first read one dilemma and made a judgment about the appropriate action. Participants indicated their judgment about the moral dilemma using a 6-point Likert scale with the endpoints labeled as “Completely Inappropriate” and “Completely Appropriate.” Then, they rated the extent to which they agreed with a statement of a moral belief. In Experiment 1, this statement was: “In life or death situations, one should always take whatever means necessary to save the most lives”. Participants rated their agreement using a 7-point Likert scale with endpoints labeled as “Completely Disagree” and “Completely Agree.” Afterwards, participants answered a pair of simple reading comprehension questions about the dilemma they read, and also indicated whether they had ever seen or heard the dilemma before. Participants advanced through the experiment at their own pace, but timing controls ensured they stayed on each page long enough to completely read instructions, dilemmas, and questions.

3.2 Results and Discussion

Participants who reported having previously seen the dilemma (52) or who failed the reading comprehension checks (an additional 24) were excluded from analysis. This left 189 participants in the final analysis.

Consistent with our hypothesis, participants in the control condition strongly endorsed the utilitarian belief statement (mean rating=5.62; SD=1.2). Endorsement was also high in the Trolley condition (mean=5.26; SD=1.56) and there were no significant differences between people’s ratings in the Trolley and control conditions, $t(137)=1.657, p=.07$. In contrast, participants who read the Footbridge dilemma rated their agreement with the statement significantly lower (mean=4.10; SD=1.52), $t(137)=-5.79, p<.001$. Moreover, participants in the Footbridge condition gave significantly lower agreement ratings than participants in the Trolley condition, $t(98)=-4.278, p<.001$ (see Fig. 1). These results suggest that people construe the Footbridge




Fig. 1 Participants’ agreement ratings for the utilitarian belief statement across conditions

dilemma as evidence against the utilitarian belief statement, but do not see the Trolley dilemma as evidence in favor of it (in fact, the trend in the Trolley condition was exactly the opposite).

Participants' moral belief agreement ratings significantly correlated with their moral judgments in both the Footbridge condition ($r(48)=.326, p<.05$) and in the Trolley condition ($r(48)=.459, p<.01$), suggesting that this moral belief is importantly related to people's moral judgments about these dilemmas.

4 Experiment 2

4.1 Method

4.1.1 Participants

In Experiment 2, 202 participants (85 female, mean age=31 years) were recruited via mTurk using the same criteria and with the same compensation as in Experiment 1.

4.1.2 Procedure

The design and materials of Experiment 2 were identical to those of Experiment 1, with the exception of the moral belief statement. In Experiment 2, participants rated their agreement with the statement: "You should never kill another person".

4.2 Results and Discussion

As in Experiment 1, participants who reported having previously seen the dilemma (40) or who failed the reading comprehension checks (21) were excluded from analysis. This left 141 participants in the final analysis.

On average, participants in all three conditions endorsed the deontological belief statement. There were no significant differences between the control (mean=5.08; SD=1.56) and Footbridge conditions (mean=5.43; SD=1.47, $t(94)=0.972, p=.33$), nor between the control and Trolley conditions (mean=5.49; SD=1.63, $t(102)=1.168, p=.25$) (see Fig. 2). Participants' moral belief agreement ratings correlated only weakly (and non-significantly) with their moral judgments in both the Footbridge condition ($r(35)=-.212, p=.21$) and in the Trolley condition ($r(43)=-.287, p=.06$). Somewhat surprisingly, it appears that neither the Footbridge nor the Trolley dilemmas are taken as evidence for or against the deontological belief statement, and that this belief may not be as relevant for the formation of judgments about these dilemmas.

5 General Discussion

We examined the roles that two moral beliefs, one representing a utilitarian perspective and the other representing a deontological perspective, play in producing the




Fig. 2 Participants' agreement ratings for the deontological belief statement across conditions

Footbridge-Trolley ordering effect. We hypothesized that three factors would provide insight into the asymmetry of the effect: how strongly people endorsed each belief, how their beliefs are related to their moral judgments about the Footbridge and Trolley dilemmas, and how considering these dilemmas affect people's endorsement of these beliefs.

Although we found that people endorsed both beliefs, only the utilitarian belief was significantly related to people's moral judgments in these two dilemmas. As we predicted, making a judgment about the Footbridge dilemma reduced participants' endorsements of the utilitarian belief, but making a judgment about the Trolley dilemma had no effect on their endorsements of this belief. Together, our results indicate that people recruit the utilitarian belief when making judgments about the Footbridge and Trolley dilemmas, and that this belief is more strongly affected by considering the Footbridge dilemma than it is affected by considering the Trolley dilemma. We propose that this pattern of belief recruitment and updating explains the asymmetry in the Footbridge-Trolley ordering effect.

These findings are consistent with the Belief-Adjustment model (Hogarth and Einhorn 1992). If we consider the utilitarian belief as S_{k-1} , then we should conclude that $w_{\text{Footbridge}}[s(x_{\text{Footbridge}}) - R]$ is not equal to $w_{\text{Trolley}}[s(x_{\text{Trolley}}) - R]$. There are two possible sources of this inequality: the subjective evaluations of the dilemmas ($s(x_k)$) and the weights (w_k) assigned to those evaluations. When a to-be-adjusted belief is strong, as is the case with the utilitarian belief at issue, the Belief-Adjustment model predicts that disconfirmatory evidence will be weighted strongly, whereas confirmatory evidence will be weighted weakly. This is precisely what we observed. Disconfirmatory evidence—the Footbridge dilemma—had a strong effect on their endorsement of the utilitarian belief, whereas confirmatory evidence—the Trolley dilemma—had no significant effect. Although it is also possible that the dilemmas are given different evaluations (after all, they are similar but not identical), the difference in weights is sufficient to explain our findings regarding the utilitarian belief, and consequently, the asymmetry in the Footbridge-Trolley ordering effect.

In Experiment 2, we did not find a significant correlation between people's endorsement of the deontological belief and their judgments about either the Footbridge or the Trolley dilemmas. Additionally, people's endorsements of the deontological belief were not affected by considering either of these dilemmas. Together these results are consistent, as it is unlikely that a belief would be strongly influenced by consideration of a dilemma if the belief were not recruited when forming judgments about that dilemma.

However, it is unclear why a relevant deontological belief would not be recruited when making judgments about the Footbridge and Trolley dilemmas. It is possible we failed to reject the null hypothesis when testing the correlation between people's endorsement of the deontological belief and their judgments about the Trolley dilemma due to insufficient power. Given that we did observe the expected trend in the data (and the obtained p -value was .06), collecting additional data might reveal that this modest correlation ($r(43) = -.287$) is statistically reliable.

Supposing this is true, it is somewhat surprising that considering the Trolley dilemma had no effect on our participants' endorsement of the deontological belief. At first blush, it seems that people's judgments about the Trolley dilemma provide evidence against the deontological belief at issue. And, as people endorse the deontological belief, sequential learning models like the Belief-Adjustment model (Hogarth and Einhorn 1992) predict that disconfirmatory evidence will be weighted strongly.

There are a few explanations for this finding that are consistent with the Belief-Adjustment model. First, it could be that people do not regard their judgments about the Trolley dilemma as strong evidence against the deontological belief that we probed. For instance, this could be because people tacitly allow for exceptions to some strict moral rules. This view, known as *Threshold Deontology*, states that deontological rules apply up until a point where the consequences of continuing to adhere to the rule become too terrible to allow (see Moore ch. 17, 1997; Moore 2012). For example, while torture is normally categorically prohibited, if torture were the *only* way to prevent a nuclear weapon from detonating in a major city, torturing would no longer be morally prohibited. So although participants endorsed the deontological belief, they may have tacitly held that this rule does not apply beyond some threshold. In this way, their deontological belief would not be undermined by considering the Trolley dilemma if the negative consequences of inaction exceeded the threshold of the deontological rule.

Second, as Hogarth and Einhorn acknowledge in their formulation of the Belief-Adjustment model, some beliefs are more sensitive to new evidence than others. For example, people may be largely insensitive to new evidence when updating certain moral beliefs because they have accumulated overwhelming evidence supporting those beliefs. Perhaps people have accumulated more evidence supporting the deontological belief than the utilitarian belief, making this belief more resistant to updating in light of new evidence. By choosing appropriate α and β parameter values in the weighting functions, the Belief-Adjustment model can be made to account for the effects of accumulated evidence. Subsequent investigations could examine whether people have accumulated different amounts of evidence for these two beliefs to justify these parameter values. The Belief-Adjustment model offers several explanations of our data, although determining whether these explanations are borne out requires further empirical research.

In sum, we examined how people's moral beliefs influence their judgments about the Footbridge and Trolley dilemmas and how making judgments about these dilemmas affects their moral beliefs. By using a mathematical model to explain the asymmetry in the Footbridge-Trolley ordering effect, we also demonstrated that sequential learning models can be used to make novel predictions about moral cognition.

Acknowledgments We would like to thank the two anonymous referees at Review of Philosophy and Psychology, Keith Holyoak, John Hummel, Jonathan Waskan, and Tage Rai.

References

- Danks, D., T.L. Griffiths, and J.B. Tenenbaum. 2003. Dynamical causal learning. In *Advances in neural information processing systems*, ed. S. Becker, S. Thrun, and K. Obermayer, 67–74. Cambridge, Mass: The MIT Press.
- Greene, J.D., R.B. Sommerville, L.E. Nystrom, J.M. Darley, and J.D. Cohen. 2001. An fMRI investigation of emotional engagement in moral judgment. *Science* 293(5537): 2105–2108.
- Hogarth, R.M., and H.J. Einhorn. 1992. Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology* 24(1): 1–55.
- Liao, S. M., A. Wiegmann, J. Alexander, and G. Vong. 2011. Putting the trolley in order: Experimental philosophy and the loop case. *Philosophical Psychology* :1–11.
- Lombrozo, T. 2009. The role of moral commitments in moral judgment. *Cognitive Science* 33(2): 273–286.
- Moore, M. 1997. *Placing blame: A general theory of the criminal law*. Oxford: Oxford University Press.
- Moore, M. 2012. Ethics in extremis: Targeted killings and the morality of targeted killings. In *Targeted killing: Law and morality in an asymmetrical world*, ed. C. Finkelstein, J. Ohlin, and A. Altman. Oxford: Oxford University Press.
- Prinz, J. 2007. *The emotional construction of morals*. Oxford University Press.
- Rai, T.S., and K.J. Holyoak. 2010. Moral principles or consumer preferences? Alternative framings of the Trolley problem. *Cognitive Science* 34(2): 311–321.
- Rescorla, R.A., and A.R. Wagner. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II*, ed. A.H. Black and W.F. Prokasy, 64–99. New York: Appleton.
- Schwitzgebel, E., and F. Cushman. 2012. Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*.
- Shanks, D. 1985. Forward and backward blocking in human contingency judgment. *The Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology* 37(1): 1–21.
- Sinnott-Armstrong, W. 2008. Framing moral intuitions. In *Moral psychology volume 2 the cognitive science of morality intuition and diversity*, ed. W. Sinnott-Armstrong, pp. 47–76. MIT Press.
- Tobia, K. et al. (2012). Moral Intuitions: Are Philosophers Experts?. *Philosophical Psychology* :1-10.
- Van Hamme, L.J., and E.A. Wasserman. 1994. Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation* 25: 127–151.
- Waldmann, M. R., J. Nagel, and A. Wiegmann. 2012. Moral Judgment. In *The Oxford handbook of thinking and reasoning*, eds. K. J. Holyoak, R. G. Morrison, 1–68.
- Wiegmann, A., and Y. Okan. 2012. Order effects in moral judgment: Searching for an Explanation. *Proceedings of the Thirty-Fourth Annual Conference of the Cognitive Science Society*. Sapporo, Japan.
- Wiegmann, A., J. Nagel, and S. Mangold. (2008). Order effects in moral judgment. *Philosophical Psychology* (March):2111–2116.