

“But He’s My Brother”: How Family Obligation Impacts Moral Judgments

Junho Lee (leejunho@ucla.edu)

Keith J. Holyoak (holyoak@psych.ucla.edu)

Department of Psychology, University of California, Los Angeles
Los Angeles, CA 90095 USA

Abstract

We created practical moral dilemmas for which participants role-played witnessing a transgression by a target person. The identity of the transgressor was manipulated to be either a stranger or the participant’s brother. Participants made factual and unethicality judgments regarding the incident and reported their willingness to report the transgressor to the police. When the factual situation was ambiguous, participants interpreted the facts in favor of the target person when that target was their brother. This family favoritism in turn led to partial moral judgments and decisions, while creating overall coherence. When it was made clear that their brother actually committed the transgression, partiality in unethicality judgment was reduced but partiality in the decision to report persisted, even though overall coherence was thereby reduced. Using path analyses, we show how strong moral constraints such as family obligation can shift moral reasoning processes.

Keywords: morality; judgment; decision making; family obligation; motivated reasoning; path analysis

Introduction

Moral reasoning is often viewed as an individual’s assessment of his or her responsibilities toward strangers or near-strangers within some broadly defined group (e.g., fellow citizens). Particularly when moral judgments are considered within the ethical framework of utilitarianism (e.g., Singer, 1979), the value of each affected person is held to be independent of the unique perspective of the individual decision maker. Under this view, moral judgments are *agent-neutral* (Nagel, 1986). The great majority of studies in moral psychology (notably, those focusing on sacrificial dilemmas based on variations of the trolley problem; Foot, 1978) involve scenarios in which a hypothetical decision is made concerning the fates of anonymous strangers (e.g., Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Cikara, Farnsworth, Harris, & Fiske, 2010; Uhlmann, Pizzaro, Tannenbaum, & Ditto, 2009). Any hint of a personal relationship between the reasoner and those affected by the moral judgment is avoided in the interests of supposed objectivity, consistent with experimental controls commonly employed in psychological research.

This focus in moral psychology on agent-neutral judgments has contributed to the neglect of moral issues that arise in everyday life. Few of us will ever have to choose whether to redirect a runaway trolley so as to kill one stranger in order to save five others; many of us will have to choose whether to support our own children or donate all our income to charity. Of course, for most people the latter

decision is not likely to be a difficult one—our own children need to be fed first. The rare parent who exhibits “pathological generosity” consistent with extreme utilitarianism likely suffers from brain damage (Ferreira-Garcia, Fontenelle, Moll, & de Oliveira-Souza, 2014; see also Kahane et al., 2012).

In general, introducing a personal relationship between the moral agent and the affected parties emphasizes *agent-relative* responsibilities (Nagel, 1989). These include family obligations, patriotism, or in-group loyalty. Agent-relative decisions inevitably evoke the deontological concepts of right and duties (Holyoak & Powell, 2016). Rights and duties cannot be universal given the potential conflicts between agents upholding the interests of different parties. For example, the duties of soldiers in two opposing nations at war generally cannot be reconciled, since their individual duties to be loyal to their respective countries conflict with each other. Yet even mortal enemies may recognize and respect each other’s duties—a deserter from one side may be despised as a traitor by the other.

As Nagel (1979) pointed out, agent-relative moral decisions are inevitably personal and situated. One’s moral obligation to provide care and protection for a child is not equal between the case where that child is one’s own daughter and the case where that child is anonymous and unrelated. Family obligations are essentially agent-relative, in that the ethicality of a decision depends on who is making the decision, or on the relationship between the moral agent and the affected parties.

Bloom (2011) observed that those about whom we make moral judgments and decisions are often friends, colleagues, family, or in-group members who share some personal attributes with us (e.g., students from the same school or people of the same ethnicity or nation; Cikara et al., 2010; Haidt, 2007; Haidt & Graham, 2007). Consequently, our moral judgments may often be partial and biased, especially when the safety, well-being, or interest of a close group conflicts with those of a more distant group. People’s general tendency to favor in-group members, and more specifically family and kin, has been documented in experimental studies (Cikara et al., 2010; Burnstein, Crandall, & Kitayama, 1994; Haidt, 2007) and evolutionary analyses (Shackelford & Hansen, 2015; Hamilton, 1964).

In this paper, we take an approach that differs from previous studies that have examined family favoritism in the context of artificial sacrificial dilemmas, such as the trolley problem. Instead, we examine how perceived family obligation—a ubiquitous source of agent-relative issues—impacts moral decisions in more realistic situations.

Our central aim is to understand the mechanisms by which family obligation may impact moral judgments. One general possibility is that family relationships may enter into an interconnected network of beliefs and attitudes that collectively shift so as to maximize *coherence*. For example, people faced with a legal decision involving contradictory and ambiguous factors will shift their views on all the relevant factors to maximize coherence with their eventual decision (Holyoak & Simon, 1999). Coherence shifts, which can be modeled in terms of constraint satisfaction within a belief network, have been shown to affect decisions ranging from consumer choice (Simon, Krawczyk & Holyoak, 2004; Russo, Carlson, Meloy, & Yong, 2008) to judgments of legal and moral culpability (Simon, Stenstrom, & Read, 2015; Simon, Snow, & Read, 2004). In the case of an apparent transgression, family favoritism may alter assessments of uncertain aspects of the situation so as to make a family member appear to be less blameworthy than a stranger would be (an instance of motivated reasoning; Kunda, 1990). Such coherence shifts would be consistent with what Holyoak and Powell (2016) termed *deontological coherence*, whereby a network of beliefs and values is altered so as to minimize conflict between rival moral values.

Coherence-driven decision making takes advantage of factual and moral/legal ambiguities, which make it possible to shift beliefs without blatantly contradicting points of certainty. Sometimes, however, key facts are incontrovertible—the culprit may be caught red-handed. In such situations, the moral agent may be placed in a true dilemma, being forced to choose whether to act in accord with the dictates of the law and society, or to honor a perceived obligation to protect a family member. In the present study, we examined the influence of family favoritism in situations that either did or did not suggest factual ambiguity.

Experiment

In our experiment, participants read a scenario in which they witnessed a hypothetical transgression. They were asked factual and moral questions including how willing they were to report the transgressor to the police. We manipulated the identity of the purported transgressor as either participant's brother or a stranger. This design pits two moral goals against one another: Family obligation (to protect a family member, even from a social punishment) and civic duty (to report a transgression).

Method

Participants The Amazon Mechanical Turk system was used to recruit 341 participants (189 females, $M_{\text{age}} = 35.9$, $SD_{\text{age}} = 11.6$) residing in the United States. Ethnicity was self-reported as 75.4% European/European American, 10.6% African American/Black, 6.5% Hispanic or Spanish origin, 6.2% Asian/Asian American, with 1.5% classifying themselves as “others”. Participants received \$1.40 as

compensation for completing the study, which took a median of 6.5 minutes.

Design, Materials and Procedures Each participant read and made decisions about a single scenario, in which a target person is implicated in a purported crime. The scenarios used one of two basic cover stories (traffic violation or street battery). The rest of the design was a 2 (identity of transgressor: stranger/brother) \times 2 (severity of violation: misdemeanor/felony) \times 2 (situational ambiguity: ambiguous/unambiguous) factorial. About 20 participants completed each of the 16 conditions (including the variation in cover story). All factors were manipulated between-subjects.

After providing consent, participants were instructed to imagine themselves as the person in the given scenario and answer the questions after careful consideration. The participants role-played being a witness to a possible transgression. For example, one of the ambiguous scenarios involving a street battery incident (misdemeanor) was as follows:

One evening, you are walking home after a long workday. A block ahead, you see a man wearing a red baseball cap, who seems to be arguing with another man. Soon, the two men disappear into an alley. As you walk up to where they were, you see the other man lying on the ground in an alley, covering his face and groaning, though his injury doesn't seem to be serious. The injured man is drunk, so he may not be reliable or truthful about what happened.

You consider the possibility that the man with the red cap may have attacked the drunken man and then ran away, and should be reported to the police. However, you also consider the possibility that the drunken man may have been trying to pick a fight with the man in the red cap. The man with the red cap may have tried to defend himself, or perhaps hit the drunken man accidentally while trying to run away to avoid a fight.

The purported transgressor (man in the red cap in the above scenario; car driver in the traffic violation scenario) was the target person for moral judgment and decision. The transgression was either a misdemeanor, as above (purportedly punched a drunken man in the face and then ran away, or purportedly drove under the influence), or a felony (stabbed a drunken man and then ran away, or purported hit-and-run). In the unambiguous conditions, any doubt that the target person committed the offense was eliminated because the witness was said to have clearly witnessed the target person committing the transgression.

Following the description of the scenario, the target person was described as either a stranger or the participant's brother, depending on the condition. If participants did not have a brother, they were told to imagine they have a brother about their own age.

The participants in the ambiguous conditions were told that they saw a policeman a few blocks back, and were asked questions in the following fixed order: (1) “Do you

think your brother/stranger actually [committed a transgression?]" (different transgression inserted for each condition), using a 6-point scale (1: *Certainly not*, 6: *Certainly*). (2) "Given what you believe happened, how unethical was the behavior of your brother/stranger?", using a 5-point scale (1: *Not problematic*, 5: *Extremely unethical*). (3) "Given what you believe happened, would you report your brother/stranger to the policeman?", using a 6-point scale, same as (1). For brevity, these three questions and the corresponding mean scores will be termed *factual*, *unethicality*, and *report*. *Factual* judgment was not elicited from participants in the unambiguous conditions because the scenario stated directly that the target person committed the transgression. After the questions¹, the participants were asked to briefly explain the reason why they reported (or did not report) the target person to the police. Lastly, basic demographic questions were asked, including whether the participant had a sibling (87% reported that they did²).

Results

Mean Differences and Correlations As the basic pattern of results was similar for the two cover stories, all analyses were collapsed over that variable.

We conducted a 2 (identity of transgressor) \times 2 (severity) ANOVA for the *factual* score with the participants from the ambiguous conditions ($n = 171$). The identity of the target person had a significant effect on the *factual* score ($F(1, 167) = 23.65, p < .001, \eta_p^2 = .124$), while the main effect of severity ($F(1, 167) = 2.74, p = .10$) and its interaction with identity did not ($F(1, 167) = 0.15, p = .70$). As shown in Figure 1a, left, participants had a stronger belief that the target person committed a transgression when the target person was a stranger than when he was a brother.


The ANOVAs for *unethicality* and *report* scores included the additional variable of ambiguity. For *unethicality* (Figure 1, middle columns), this analysis yielded significant effects of identity ($F(1, 333) = 10.32, p = .001, \eta_p^2 = .03$), severity ($F(1, 333) = 43.64, p < .001, \eta_p^2 = .116$), and ambiguity ($F(1, 333) = 71.75, p < .001, \eta_p^2 = .177$). The interaction between identity and ambiguity approached significance ($F(1, 333) = 3.19, p = .075$); other interactions were not significant ($ps > .34$). Thus, the *unethicality* of the behavior was judged as significantly higher when the target person was a stranger, transgression was more severe, and the situation was unambiguous.

Based on previous findings (e.g., Holyoak & Simon, 1999; Simon et al., 2015), we predicted that ambiguity would enable participants to decrease judged *unethicality* when the target was their brother rather than a

¹ Additional questions were asked after the three main questions, but for brevity these will not be discussed.

² When participants actually had a sibling, *factual*, *unethicality* and *report* scores tended to be lower, although the extremely unequal numbers of participants in the two groups rendered statistical tests problematic. The central findings from ANOVAs and path analyses were consistent regardless of whether participants without siblings were included or excluded.

(a) Ambiguous condition



(b) Unambiguous condition




Figure 1: Means of *factual*, *unethicality*, and *report* scores for (a) ambiguous and (b) unambiguous conditions. *Factual* judgment was not asked in the unambiguous conditions. Error bars indicate standard error of the mean.

stranger, so as to cohere with their motivation to protect a family member. To test this prediction, two separate 2 (identity) \times 2 (severity) ANOVAs were conducted with *unethicality* scores for ambiguous and for unambiguous conditions. For the ambiguous conditions, identity ($F(1,167) = 10.18, p = .002, \eta_p^2 = .057$) and severity ($F(1,167) = 19.29, p < .001, \eta_p^2 = .100$) both yielded significant main effects, while the interaction between the two factors was not reliable ($F(1,167) = 0.32, p = .57$). For the unambiguous conditions, only severity had a significant main effect ($F(1,166) = 26.02, p < .001, \eta_p^2 = .135$); neither the main effect nor the interaction involving identity was reliable ($ps > .25$). Thus, situational ambiguity was indeed necessary in order for *unethicality* to be judged lower for the brother condition.

For *report* scores, a 3-way ANOVA revealed reliable main effects of identity ($F(1,333) = 80.1, p < .001, \eta_p^2 = .194$), severity ($F(1,333) = 31.88, p < .001, \eta_p^2 = .087$), and ambiguity ($F(1,333) = 25.43, p < .001, \eta_p^2 = .071$). None of the interactions were reliable ($ps > .19$). As shown in Figure 1, right columns, participants were more likely to report a transgression to a policeman when the target person was a stranger rather than a brother, whether or not the facts of the situation were ambiguous.

To seek converging evidence of how the influence of identity impacted judgments as a function of situational ambiguity, we examined the correlation between *unethicality* and *report* scores for the stranger and brother condition, computed separately for the ambiguous and unambiguous conditions. To control for the main effect of severity, both *unethicality* and *report* scores were standardized within each of the misdemeanor ($n = 170$) and felony ($n = 171$) conditions prior to the correlational analyses. Figure 2 presents scatterplots of the relationship between standardized *unethicality* scores and propensity to report the transgression. The lines show the best linear fits for the stranger and brother conditions, respectively. When

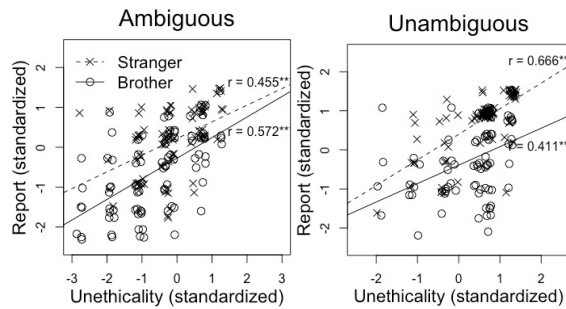


Figure 2: Scatterplots of standardized *unethicity* and *report* scores in ambiguous (left) and unambiguous (right) conditions, separated by stranger versus brother conditions. Jitter was added to the points for clarity of presentation.

the situation was ambiguous, correlation coefficients were not significantly different for the stranger ($r(82) = .46, p < .001$) versus brother condition ($r(85) = .57, p < .001$; $z_{\text{difference}} = 1.03, p = .30$). However, when the situation was unambiguous, the correlation coefficient for the stranger condition ($r(82) = .67, p < .001$) was significantly higher than that for the brother condition ($r(84) = .41, p < .001$; $z_{\text{difference}} = 2.35, p = .02$).

This pattern of correlations is consistent with the hypothesis that in the ambiguous condition, the same conduct is interpreted as less problematic when the target is a brother rather than stranger. Judged unethicity is then altered so as to cohere with the factual judgment; after this tacit “correction”, the relationship between *unethicity* and *report* scores is approximately the same regardless of target identity. In contrast, the unambiguous condition does not support a coherence shift in *unethicity*, so participants judge an action as about equally unethical regardless of identity. Faced, therefore, with the unresolved dilemma of duty to report versus duty to protect a family member, participants often favor the family member despite the perceived unethicity of his action, yielding a weaker correlation between *unethicity* and *report* for the brother condition.

Path Analysis To further analyze the relationships among the different judgments, separate path analyses for the ambiguous and unambiguous conditions were conducted using R package lavaan version 0.5-23 (Rosseel, 2012). Maximum likelihood estimation with robust standard errors and Satorra-Bentler scaled test statistics (Satorra & Bentler, 2010) were used for all analyses. During iterative testing of the models, suggestions from Lagrange multiplier tests that did not violate logic were considered when modifying the path models. The identity and severity variables were dummy coded (1: brother, 0: stranger; 1: felony, 0: misdemeanor, respectively).

For the ambiguous conditions, the model with the best fit ($\chi^2(2) = 3.21, p = .200$, RMSEA = .063, 90% CI of RMSEA = [.000, .183], CFI = .995, TLI = .972) had *factual* predicted by identity, *unethicity* predicted by *factual* and severity, and *report* predicted by all other variables: *unethicity*,

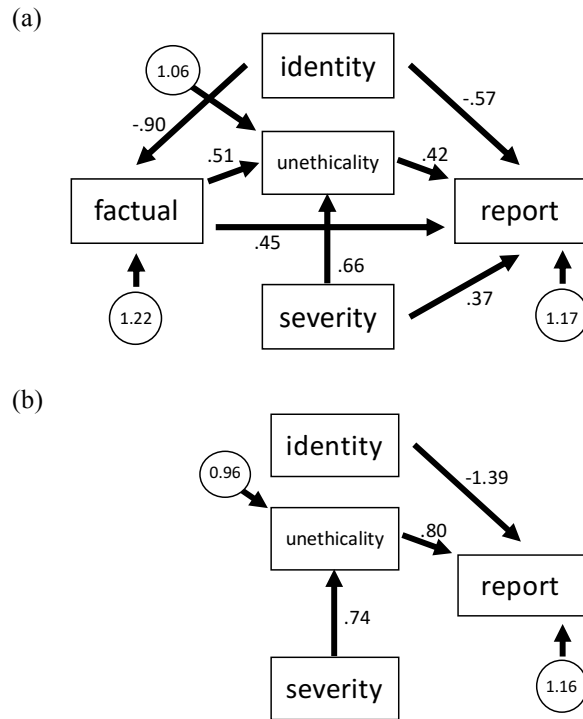


Figure 3: Path models that best fit the data from (a) ambiguous and (b) unambiguous conditions. Numbers in circles indicate standard errors after predicting the dependent (endogenous) variables.

factual, identity, and severity. Most paths were significant (path from severity to *report*: $p = .051$; all other $ps < .006$). This model (Figure 3a) implies that identity had a direct influence on *factual* and *report*, but not on *unethicity*. The same model was supported when separate path models were generated for each of the cover stories (traffic violation/street battery), with minor differences in fit indices.

We conducted mediation analyses using bootstrapping in PROCESS macro version 3.0 for SPSS (Preacher & Hayes, 2008) to test the significance of indirect effects of identity on *unethicity* and *report*. First, mediation between identity (predictor), *factual* (mediator) and *unethicity* (criterion) was tested, while severity also predicted *unethicity* (model 5 from the predefined models in PROCESS). Bootstrapping results ($n = 20,000$) indicated that identity did not have a significant direct effect on *unethicity* for either misdemeanor (-.220, 95% confidence interval from bootstrap = [-.698, .257]) or felony conditions (-.081, 95% CI = [-.541, .379]). In contrast, the indirect effect of identity on *unethicity* through *factual* was significant (-.439, 95% CI = [-.678, -.235]).

Second, mediation between identity (predictor), *factual* (mediator), *unethicity* (mediator), and *report* (criterion) was tested. The residual *unethicity* and *report* scores obtained after regressing the two variables on severity were used in this model (model 6 in PROCESS). Bootstrapping results ($n = 20,000$) indicated that identity had a significant direct effect on *report* (-.577, 95% CI = [-.952, -.201]), and

its total indirect effect on *report* was also significant (-.646, 95% CI = [-.969, -.351]), although the indirect path from identity through *unethicality* to *report* was not significant (-.067, 95% CI = [-.228, .072]), consistent with the first mediation analysis.

For the unambiguous conditions, the best fitting model ($\chi^2(2) = 3.08, p = .214, RMSEA = .065, 90\% \text{ CI of RMSEA} = [.000, .200], CFI = .989, TLI = .973$) had *unethicality* predicted by severity and *report* predicted by identity and *unethicality* (Figure 3b). All paths were significant ($ps < .001$). Notably, identity had a significant direct effect on *report*. Again, the same model was supported when separate path models were generated for each of the cover stories (traffic violation/ street battery).

Overall, the path analyses favored similar models for the ambiguous and unambiguous conditions, with the salient difference that in the ambiguous condition, identity directly influenced the *factual* score (which in turn influenced both *unethicality* and *report*). In contrast, in the unambiguous condition the influence of identity on *report* was solely a direct one.

Discussion

The present study revealed that both factual and moral judgments are impacted by the personal relationship between the moral agent (i.e., a participant in our experiment) and a target person who appears to commit a transgression. Given situational ambiguity, participants judged the same behavior to be less likely to constitute an actual transgression (e.g., he may have only tried to defend himself) when the target was identified as their brother rather than a stranger. This effect can be interpreted as an instance of motivated reasoning (Kunda, 1990), in that participants' motivation to favor and protect their in-group member (brother) led them to judge the "facts of the case" in a way that favored that in-group member.

Both standard ANOVAs and path analyses showed that given situational ambiguity, participants rated an action as less unethical, and were less willing to report it to police, when the target person was their brother rather than a stranger. The path analysis indicated both a direct influence of target identity on willingness to report and also an indirect influence via the impact on the assessment of the situational facts. This pattern is consistent with a coherence shift (Holyoak & Powell, 2016; Holyoak & Simon, 1999): Ambiguity is exploited to interpret the facts in a way that favors a family member, thereby reducing judged unethicality and decreasing willingness to report.

When the situation was rendered unambiguous, such that the target person incontrovertibly committed the transgression, it was no longer possible to generate a coherent set of beliefs that would excuse the family member. The action was then judged equally unethical regardless of whether it was done by a brother or a stranger. However, when faced with a clear conflict between the agent-neutral duty to report crimes and the agent-relative duty to protect a family member, participants often elected

to fulfill the latter duty at the expense of the former. Consequential moral decisions such as reporting a family member to the police are not made easily because people typically have a strong sense of obligation to protect closely-related family members, even at a high cost (Burnstein et al., 1994; Hamilton, 1964). Comments provided by our participants suggested that some of them were also consciously aware of their use of familial (deontological) duty in making decisions (e.g., "I feel a great need to protect my brother"; "It is also my duty to protect my family members").

It would be useful in future studies to investigate the impact of in-group favoritism on moral reasoning process in different cultural contexts (Graham, Meindl, Beall, Johnson, & Zhang, 2016). Understanding cross- and within- societal differences in moral norms and reasoning is growing more crucial given the fierce clashes of cultures that we face today. For example, Chen, Brockner and Katz (1998) found evidence that people from individualistic (e.g., U.S.) versus collectivistic (e.g., China) cultures may have fundamentally different rationales for favoring in-group members (self-enhancement versus unconditional dedication to the group).

In short, our participants favored their brother when making consequential moral decisions. When possible (i.e., when the situation was ambiguous), they did so after coherence-based reasoning paved the way by creating a justification for leniency. But even when the situation was unambiguous, and their brother's guilt was clear, participants often refused to report him to the police. Family obligation is a powerful constraint on people's decisions. The reason some of our participants stated—"Because he is my brother"—is often sufficient justification for a moral decision.

References

- Bloom, P. (2011). Family, community, trolley problems, and the crisis in moral psychology. *The Yale Review*, 99(2), 26-43.
- Burnstein, E., Crandall, C., & Kitayama, S. (1994). Some neo-Darwinian decision rules for altruism: Weighing cues for inclusive fitness as a function of the biological importance of the decision. *Journal of Personality and Social Psychology*, 67(5), 773.
- Chen, Y. R., Brockner, J., & Katz, T. (1998). Toward an explanation of cultural differences in in-group favoritism: The role of individual versus collective primacy. *Journal of Personality and Social Psychology*, 75(6), 1490-1502.
- Cikara, M., Farnsworth, R. A., Harris, L. T., & Fiske, S. T. (2010). On the wrong side of the trolley track: Neural correlates of relative social valuation. *Social Cognitive and Affective Neuroscience*, 5(4), 404-413.
- Ferreira-Garcia, R., Fontenelle, L. F., Moll, J., & de Oliveira-Souza, R. (2014). Pathological generosity: An atypical impulse control disorder after a left subcortical stroke. *Neurocase*, 20(5), 496-500.
- Foot, P. (1978). The problem of abortion and the doctrine of double effect. In *Virtues and vices*. Oxford: Blackwell.

- Graham, J., Meindl, P., Beall, E., Johnson, K. M., & Zhang, L. (2016). Cultural differences in moral judgment and behavior, across and within societies. *Current Opinion in Psychology*, 8, 125-130.
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105-2108.
- Haidt, J. (2007). The new synthesis in moral psychology. *Science*, 316(5827), 998-1002.
- Haidt, J., & Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1), 98-116.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour. II. *Journal of Theoretical Biology*, 7(1), 17-52.
- Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, 128(1), 3-31.
- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin*, 142(11), 1179-1203.
- Kahane, G., Wiech, K., Shackel, N., Farias, M., Savulescu, J., & Tracey, I. (2012). The neural basis of intuitive and counterintuitive moral judgment. *Social Cognitive and Affective Neuroscience*, 7(4), 393-402.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480-498.
- Nagel, T. (1979). The fragmentation of value. In *Mortal questions* (pp. 128-141). Cambridge, UK: Cambridge University Press.
- Nagel, T. (1986). *The view from nowhere*. New York: Oxford University Press.
- Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The goal of consistency as a cause of information distortion. *Journal of Experimental Psychology: General*, 137, 456-470.
- Satorra, A., & Bentler, P. M. (2010). Ensuring positiveness of the scaled difference chi-square test statistic. *Psychometrika*, 75(2), 243-248.
- Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science*, 15, 331-336.
- Simon, D., Stenstrom, D. M., & Read, S. J. (2015). The coherence effect: Blending cold and hot cognitions. *Journal of Personality and Social Psychology*, 109(3), 369-394.
- Simon, D., Snow, C. J., & Read, S. J. (2004). The redux of cognitive consistency theories: Evidence judgments by constraint satisfaction. *Journal of Personality and Social Psychology*, 86(6), 814.
- Singer, P. (1979). *Practical ethics*. Cambridge, UK: Cambridge University Press.
- Uhlmann, E. L., Pizarro, D. A., Tannenbaum, D., & Ditto, P. H. (2009). The motivated use of moral principles. *Judgment and Decision Making*, 4, 476-491.