1 Rational mechanisms can produce immoral evaluations

2 Derek Powell

3 Department of Psychology

4 Stanford University

5 Zachary Horne

6 School of Social and Behavioral Sciences

7 Arizona State University

Abstract

People’s expectations play an important role in their reactions to events. There is often disappointment when events fail to meet expectations and a thrill to having one’s expectations exceeded. Here, we propose that expectations influence moral evaluations through rational information-theoretic principles: less expected events do more to inform us about the state of the world than do more expected events thereby altering our evaluations of morally significant events. An implication of this proposal is that people may have inappropriately muted responses to highly common but nonetheless tragic events. In two preregistered experiments, we found that people’s judgments of morally-significant events were affected by the likelihood of that event. People were more upset about events that were unexpected (e.g., a robbery at a clothing store) than events that were more expected (e.g., a robbery at a convenience store). We argue that this bias has pernicious moral consequences, including leading to reduced concern for victims in most need of help.

Keywords: judgment; decision-making; morality; open science

Rational mechanisms can produce immoral evaluations

Introduction

When we learn of tragic events—a friend’s loss of a loved-one, a violent killing, or a terrorist attack—we are mournful, concerned, angry, and driven to action. But our responses to these events can differ depending on the context in which they occur (e.g., Trope & Liberman, 2010). In particular, our expectations seem to play an important role in our reactions to events. This seems true of the ordinary, the alarming, and even the catastrophic. When a friend’s loved-one passes, we often ask, “was it sudden?” A murder in a sleepy town prompts a very different reaction than a similar crime in a major city. And even the most serious crimes, like acts of terrorism, seem to evoke more outrage, horror, and concern when they occur in peaceful rather than unstable regions. For instance, the 2015 Paris attacks were felt around the world. Yet, most of those mourning had little to say 15 hours earlier, when another tragic attack killed at least 43 people in Beirut (Graham, 2015). Likewise, stories of the Virginia Tech massacre, which killed 32 people, overshadowed coverage of attacks killing nearly 200 people in Iraq—one of the bloodiest days since the 2003 invasion (Greenslade, 2007). In each of these situations, the surprise elicited by an event appears to have affected people’s reactions to it. In some contexts, harm is tragic. In others, “these things happen.”

In non-moral domains, several controlled studies have demonstrated how people’s expectations affect their reactions to outcomes. For instance, Mellers and colleagues (1997) found expectations influenced affective reactions during a gambling task: Participants were often more excited by an small unexpected gain (e.g., winning \$30 with 10% probability) than they were by a larger but more expected gain (e.g., winning \$60 with 90% probability). Beyond the lab, several studies have demonstrated that people’s evaluations of events, along with their subsequent emotional responses, are affected by their expectations about how those events were supposed to unfold. For example, Villano and colleagues (2020) observed that students’ emotional responses to their grades were better predicted by

their prediction errors about their grades than was the grade itself (also see Golub, Gilbert, & Wilson, 2009). In a similar vein, *Norm theory* (Kahneman & Miller, 1986) emphasizes the role of counterfactual comparisons (thought to be determined by expectations of how an event was supposed to unfold) in evaluations. On this account, the salience of different counterfactuals can explain, for example, why Olympic Bronze medalists are often happier than Silver medalists (Medvec, Madey, & Gilovich, 1995; Medvec & Savitsky, 1997) – Bronze medalists tend to compare their place against failing to medal and thereby feel victorious, whereas Silver medalists tend to compare themselves to Gold, thereby feeling defeated. McGraw, Mellers, and Tetlock (2005) have also argued that Olympic Silver medalists’ disappointment relative to Bronze medalists is actually a function of comparison against expectations, rather than resulting from a direct comparison against the others on the podium (e.g., Bell, 1985; Gul, 1991; Loomes & Sugden, 1986; Mellers et al., 1997).

Intuitively, and in many cases empirically, our expectations shape our reactions to parties, sporting events, grades, and gambles. However, as an ethical principle, no serious normative account would justify drawing a distinction in the utility of moral outcomes based on how expected they are to occur (e.g., Alexander & Moore, 2016; Hursthouse & Pettigrove, 2016; Sinnott-Armstrong, 2019). Tragedies are tragic whether expected or not, and failure to recognize this could lead to tacit endorsement of the pain and suffering of others wherever pain and suffering are the status quo. Consequently, if expectations in fact influence our reactions to morally-significant world events as anecdote and some naturalistic data suggest (e.g., Olivola & Sagara, 2009), then this may not only pose both serious social problems, but also may provide further evidence for the ways in which highly-intuitive judgments (e.g., that unexpected morally-significant events *seem* more upsetting) may subtly mislead us about morality.

Several lines of research have suggested that people’s intuitions can lead them to draw moral distinctions where they ought not (e.g., Norcross, 2004), or attend to features of events that should have little bearing on the “moral facts.” For instance, (irrelevant)

changes in the syntax of sentences have been shown to affect philosophical judgments (Horne & Cimpian, 2018), and in the moral domain, intuitive judgments have been shown to be affected by morally irrelevant changes in the *framing* of moral scenarios (Nichols & Knobe, 2007; Rai & Holyoak, 2010; Tanner, Medin, & Iliev, 2008). To consider just one example, Rai and Holyoak (2010) observed that participants' willingness to sacrifice lives depends not only on how many lives are saved in a given situation, but on the number of lives at risk (also see Fetherstonhaugh, Slovic, Johnson, & Friedrich, 1997) – a finding predicted by the consumer-behavior literature highlighting lapses in intuitive judgments (e.g., Thaler, 1999). These studies—among several others—highlight the ways in which even highly-intuitive judgments can lead to questionable ethical principles if relied on too extensively (again see Norcross, 2004). Thus, while it may be intuitive, and even rational in other domains (Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010), that unexpected morally significant events seem “morally worse,” may not constitute a strong reason to rely on this intuition in developing an ethical theory.

In light of prior work demonstrating expectations influence our evaluation of non-moral events, and both the potential social and theoretical importance of expectancy effects in the moral domain, we now develop a theoretical account of how expectations influence evaluations. We will argue that expectations influence our evaluation of events not as a strange quirk or uniquely human bias, but instead as a fundamental component of how rational agents learn and reason.

A rational framework for how expectancy affects evaluations

Prior research has attempted to formalize how expectations can impact evaluations in specific domains. For instance, existing *disappointment theories* posit a disappointment function which decrements or increments utility evaluations according to some comparison of obtained and expected outcomes (e.g., Bell, 1985; Gul, 1991; Loomes & Sugden, 1986; Mellers et al., 1997). While promising insofar as they attempt to reconcile prior utility

theories with the substantial evidence that expectancy can impact perceived utility, these theories have a few limitations. First, these theories are greatly limited in scope, focused only on explaining choice behavior in artificial gambling scenarios – formally they are not stated in a sufficiently general way to articulate how they might extend beyond this domain. Second, these theories rest on a relatively arbitrary and *ad hoc* augmentation of expected utility theory, and thus their explanatory power and predictions are less clear in the moral domain. Third, although there is strong empirical support for the qualitative prediction that expectations influence evaluations, there is little quantitative support for any of the specific models proposed under these theories (see Abdellaoui & Bleichrodt, 2007; Abdellaoui, Bleichrodt, & Paraschiv, 2007). These considerations lead us to develop a more general account of how expectancy could influence evaluations, including moral evaluations.

Events, as we will conceive of them, are changes in the state of the world – transitions from one state to another. For this reason, we construe event evaluation as a comparison between the states of the world before and after the event. Similar comparison processes play crucial roles in many theories of decision-making (e.g. Bell, 1985; Gul, 1991; Loomes & Sugden, 1986; Mellers et al., 1997). For instance, in Prospect Theory (Kahneman & Tversky, 1979), the utility of a prospect is evaluated by comparison to a current reference point. Our proposal is similar: we argue that expectations can set the reference points against which people compare future outcomes, so that assigning value to an event consists in comparing the state of the world following the event to the state of the world just prior to the event. Further, we assume that people’s knowledge of the world is uncertain, so we represent the states of the world probabilistically. This means that a person’s model for the state of the world is simply their expectations about what has happened, what is happening, and what will happen in the future. A direct implication of this proposal is that the evaluation of an event is linked to what is learned about the world as a result of that event’s occurrence. That is, we evaluate events positively or negatively

to the extent that they inform our understanding of the world in positive or negative ways.

One advantage of this account is it allows us to understand how expectations influence people’s evaluations using a well-established theoretical framework. A fundamental insight of Information Theory is that the information carried by an event is a function of its prior probability; low probability events carry more information than high probability events (Shannon, 1948). We propose that people learn more about the state of the world when their expectations are violated by shocking world events as compared to when they are affirmed by less surprising events. For example, compared to a bombing in Lebanon, if a bombing occurs in Paris, a less expected location, we learn that the world is more dangerous than we had previously believed. Although this suggestion is novel in the moral domain, these points are widely accepted elsewhere. For instance, expectations have long been recognized as fundamental to models of human and animal learning (Rescorla & Wagner, 1972). However, as we will discuss below, the ethical implications of expectancy effects are more problematic – while adjusting our evaluations according to the information they carry may help us efficiently navigate the world, they may also blunt our responses to equally tragic events whose utilities are identical.

A formal treatment of this theory establishes the potential for its generality across domains. Let the state of the world prior to some event be represented by a random variable, X . We’ll assume X is discrete, and that every realization x represents some potential specific state of the world in all relevant aspects under consideration. An agent has some utility function u , that applies over the different realizations of X . Thus the expected utility of the present state of the world is:

$$E_X[u(X)] = \sum_x u(x)p(X = x)$$

When we learn an event has occurred, we can represent this as an observation of a second, binary random variable, S . We can think of S as a “signalling event” that informs us about about X (following the formalism in Arrow, 1996). That is, we assume there is a

dependency between S and X , so that observing S informs us about the state of the world, X . For instance, reading a news report of a terror attack in Brussels informs us about the state of the world, telling us about what realizations of X are no longer possible (e.g., any world wherein it was a peaceful day in Brussels) and which are more or less likely. Conditioned on S , the expected utility of the world is:

$$E_{X|S}[u(X)] = \sum_x u(x)p(X = x|s)$$

As we suggested above, our construal of event evaluation is as a comparison between states of the world following and prior to an event. Consequently, we define the value assigned to the signalling event, $V(S)$, as:

$$V(S) = E_{X|S}[u(X)] - E_X[u(X)]$$

From this it can be shown that the value assigned to event S is proportional to its prior probability (see Appendix). Specifically, it is the difference in the expected utility under different values of S , weighted by the inverse of the prior probability of S .

$$V(s = 1) = (1 - p(s = 1)) \sum_x u(x)p(X = x|s = 1) - u(x)p(X = x|s = 0)$$

This conclusion does not depend on the nature of the event in question nor on the utility function over events, but instead holds generally wherever we conceive of evaluation as a holistic comparison between uncertain states of the world. Consequently, this account predicts expectations will equally affect reactions to gambles with clearly specified values and probabilities (Mellers et al., 1997), as well as real-world events with ill-defined values, where reasoners' expectations are defined by their existing beliefs and world knowledge.

Given the fundamental role that expectations play in event evaluation, and the serious moral biases that may result, we sought to test the account described above by

examining whether evaluations of morally harmful events are affected by expectations about those events. We conducted two preregistered studies examining how even rationally-grounded cognitive mechanisms, in this case the perceived likelihood of an event, could impact moral evaluations.

The Present Experiments

We conducted two preregistered studies to test our formal account of how expectations could impact moral evaluations <https://osf.io/86rsw/>. Each study consisted of multiple trials in which participants were presented brief descriptions of two events. Their task was to decide which of the two events was “more upsetting.” In Study 1, they made a two-alternative forced-choice, and in Study 2 they made their judgments using a five-point rating scale. Both studies consisted of 16 experimental trials, five *equivalent* filler trials, and five *non-equivalent* filler trials, which we detail below. This research was approved by the Arizona State University Institutional Review Board (IRB: STUDY00007785).

The design of these studies was chosen with special consideration to two theoretical points. First, the information-theoretic account we have described makes a solely directional prediction – all else being equal, more surprising events will elicit stronger reactions than less surprising events. Asking participants to compare two similar events allowed us to test this directional prediction. Second, we are predicting the informativeness of events, in part determined by one’s prior expectations of the event, would shape their evaluations. Thus, reasoners should not have to be induced to hold expectations experimentally (as in experiments by Mellers et al., 1997), but instead their own prior beliefs should be sufficient to shape their expectations and affect their evaluations. For this reason, we manipulated the prior probability of events by describing two otherwise identical events occurring in different contexts. We explicitly chose contexts that differed in ways that could not be accounted for by construal-level explanations (Trope & Liberman, 2010) or explanations based on biases towards one’s in-group (Brewer, 1999). This allowed

us to test the role of expectations based on world knowledge and provided a more naturalistic set of items.

In experimental trials, the two events described identical harms, but one event was more expected and one less expected, as determined in a norming study we summarize below. For example, participants considered the following stimulus:

A 30 year old man in California dies in an earthquake [Expected]

A 30 year old man in Oklahoma dies in an earthquake [Unexpected]

In each of these events, the harm to the victim is the same (i.e., death) but one event is more probable than the other, given the different likelihoods of earthquakes occurring in California versus Oklahoma.

It is worth noting that in the context of gambles, it is possible to manipulate probabilities specifically and directly (Mellers et al., 1997). However, in the context of more realistic moral harms, any attempt to directly manipulate probabilities risks also changing participants' perceptions of the degree of harm. For instance, suppose we describe two cancer patients, one with a 5% chance of survival and one with a 50% chance of survival. Even if they both ultimately meet the same fate, participants might fairly assume the patient with the lower chance of survival had a more serious illness, was more debilitated, suffered more severely, and so forth. In contrast, changing the *context* of the two events explains away these likelihood differences, better equating the events in other respects. Thus, manipulating expectations through otherwise irrelevant contextual details helps ensure participants understand the harm the victims suffer is identical.

As we noted, both studies also included equivalent filler trials. For instance, participants judged which of these two events was more upsetting: “A family in Florida is evicted from their home after they run into financial troubles.” vs. “A family in Louisiana is evicted from their home after they run into financial troubles.” In these trials, the two events differed in contextual details that we did not expect would affect participants'

judgments (confirmed in a norming study). These filler trials were meant to prevent participants from becoming explicitly aware of the structure of the experimental trials.

Finally, both studies also included non-equivalent filler trials. In these trials, two events differed substantially in the degree of harm suffered by a victim. We expected that participants would see one event as considerably more upsetting than the other. For example, participants evaluated the following two outcomes, “A parked car is damaged in a parking lot in Oregon” vs. “A car is totaled on the highway in Washington.” These trials were included to allow participants a chance to use the endpoints of the response scale and to reduce any task demands that might drive them to make artificially fine-grained distinctions when generating their responses. All experimental materials for these studies are available in the Supplemental Materials at <https://osf.io/86rsw/>.

Analytic Strategy and Exclusions. We analyzed our data by performing Bayesian estimation in the probabilistic programming language Stan (Carpenter et al., 2017). We set regularizing priors (McElreath, 2016) on the parameters in our models to reduce the likelihood of model overfitting, and assessed the sensitivity of our inferences by reanalyzing our data with less conservative priors. To ensure the generality of our findings, we tested participants’ reactions to a variety of items and treated both participants and items as random effects throughout our analyses. The use of random effects over items allows us to consider these items as a sample from a larger population of possible items, and to generalize our findings to that wider population, much as we generalize from our sample of participants to the larger population (Judd, Westfall, & Kenny, 2012).

In both studies, we excluded participants who failed to correctly answer attention-check questions or who indicated at the end of the study they had not taken their participation seriously. Exclusion criteria were decided a priori and in accordance with our study preregistrations.

Materials Norming

We created and normed 20 test items, each consisting of a pair of events that differed in their prior likelihood. The two events were manipulated so what one would seem more expected than the other. Fifty participants, 42 of whom passed questions checking their attention, were asked to choose which event seemed more likely using a 5-point scale (1 = “Event 1 seems a lot more likely”, 2 = “Event 1 seems a little more likely”, 3 = “Neither event seems more likely”, 4 = “Event 2 seems a little more likely”, 5 = “Event 2 seems much more likely”). As predicted, for all items we expected to be “more probable” outcome was judged as such, but the magnitude of this difference varied somewhat by the item in question. The plot below displays the mean likelihood rating by item, along with 95% confidence intervals.

Figure 1. Perceived likelihood of expected trials displaying means and 95% Confidence Intervals.

From among these 20 items we selected the top 16 most likely items for use in our studies.

The full text for all of these items can be found in the Supplementary Materials.

Study 1

Method

Participants. Meta-meta analyses of effect sizes typical of research in social psychology indicate the modal effect size is a Cohen's d of approximately 0.20 (Richard, Bond, & Stokes-Zoota, 2003). We sought to achieve 80% power to detect an effect of this magnitude which required recruiting 200 participants. We anticipated excluding as many as 20% of participants for failing questions checking their attention, so we recruited a total

of 255 participants from Amazon’s Mechanical Turk service. Of these 230 passed attention checks and were included in the final analyses (105 male, 125 female, $M_{Age} = 33$). All participants were paid \$0.60 for their participation.

Materials and procedure. Participants judged 16 experimental event-pairs and 10 equivalent filler event-pairs in random order. On each trial participants were presented with the event-pair stimulus and had to judge which outcome was more upsetting in a two-alternative forced choice task. The two events were labeled “Outcome 1” and “Outcome 2” and the order, and whether Outcome 1 or 2 was the unexpected event, was randomized and counterbalanced.

Results

We predicted that people would judge that events where unexpected harm occurred were more upsetting than events where expected harm occurred. To estimate this effect, we fit a Bayesian logistic random-effects model using the R package `brms` (Bürkner, 2017). Participants’ judgments were the dependent variable (1 = unexpected event is more upsetting; 0 = expected event is more upsetting) and the model included random intercepts for Trial and Subject, allowing us to generalize to the wider populations of both trials and subjects:

$$\text{Evaluation} \sim \text{Intercept} + (1|\text{Trial}) + (1|\text{Subject})$$

Bayesian analyses formulate model parameters as probability distributions wherein the joint probability distribution of the data, y , and a given parameter θ is computed via the prior probability of θ and the probability of y given θ .

$$p(y, \theta) = p(y|\theta) \times p(\theta)$$

We can derive this result from Bayes’ Theorem, which is used to calculate the posterior probability of θ given y :

$$p(\theta|y) \propto p(y|\theta) \times p(\theta)$$

To model the joint probability distribution of participants' responses, we specified priors over the possible effects each model parameter could have on the response variable:

$$\alpha \sim \mathcal{N}(0.00, 1.00)$$

$$\text{All } \sigma \sim \text{Half-}\mathcal{N}(1.00, 1.00)$$

Here, the α parameter in the regression model represents the log-odds of selecting the unexpected event as being more upsetting. Thus, by examining the population-level intercept, we can test whether participants were biased toward selecting the unexpected event ($\alpha > 0$), the expected event ($\alpha < 0$), or both equally ($\alpha = 0$) as more upsetting. All σ parameters specify group-level effects (e.g., the variability in intercepts across trials and subjects). These priors weakly regularize parameter estimates, but less conservative, uninformative priors yielded similar estimates.

Consistent with the predictions of our theoretical formulation, this model revealed people were considerably more likely to think that unexpected events were more upsetting than events that were expected, $\alpha = 1.40$, 95% Posterior Uncertainty Interval [1.07, 1.72]. Figure 2 shows responses broken-down by individual items. Participants' bias toward selecting the unexpected event as more upsetting was consistent across the 16 experimental items.

Figure 2. Probability with which participants chose the less expected event as more upsetting across trials in Study 1 sorted by the strength of the item’s expectancy bias. As predicted, participants exhibited a bias toward choosing the less expected event as more upsetting across all items. Error bars represent 95% Posterior Uncertainty Intervals calculated from the mixed-effects logistic regression model.

Study 2

In Study 1 we found that people’s judgments of events were biased by their expectations about those events. When choosing between two events, participants decided that unexpected events were more upsetting than expected events. In Study 2, we sought to test our hypothesis using a more conservative method: we provided participants with a more expressive response scale so that if they viewed the events under consideration as equally harmful, their responses could reflect this.

Method

Participants. Due to the nature of the task used in Study 2, we anticipated the effects in this study might be smaller than in Study 1. Consequently, we doubled our sample size to ensure adequate power to detect any effect of interest. A total of 504 participants were recruited from Amazon’s Mechanical Turk work distribution website (mTurk). Of these, 372 passed attention checks and were included in the final analyses (174 male, 198 female, $M_{Age} = 34$ years old). All participants were paid \$0.60 for their participation.

Materials and Procedure. The materials in Study 2 were the same as those in Study 1. On each trial of participants were presented with a pair of actions labeled “Outcome 1” and “Outcome 2” and were asked, “Which outcome seems more upsetting?” In Study 2 participants made their rating on a five-point scale (1 = Outcome 1 seems more upsetting, 2 = Outcome 1 seems a little more upsetting, 3 = neither seems more upsetting than the other, 4 = Outcome 2 seems a little more upsetting, 5 = Outcome 2 seems more upsetting). The assignment of an expected or unexpected outcome to Outcome 1 or 2 was counterbalanced and randomized.

Results

We predicted that participants would evaluate unexpected events as more upsetting than expected events. We estimated the magnitude of this effect by fitting a Bayesian cumulative (ordinal) (Liddell & Kruschke, 2018) random effects model with participants’ scale responses as the dependent variable (1 to 5) and random intercepts for Trial and Subject. This model produces four intercept coefficients, representing the cumulative log-odds of responses at each scale point or lower. For instance, the second coefficient represents the log-odds participants chose a 2 (“the expected outcome seems slightly more upsetting”) or lower on the scale. Similarly, the third intercept coefficient represents the log-odds participants chose a 3 or lower on the scale. We specify the priors on this model

below:

$$\alpha_1 \sim \mathcal{N}(-2.00, 0.50)$$

$$\alpha_2 \sim \mathcal{N}(-1.00, 1.00)$$

$$\alpha_3 \sim \mathcal{N}(1.00, 1.00)$$

$$\alpha_4 \sim \mathcal{N}(2.00, 1.00)$$

$$\text{All } \sigma \sim \text{Half-}\mathcal{N}(1.00, 1.00)$$

To test our prediction, we compared the third intercept coefficient (representing the log-odds of choosing 4 or higher) to the inverse of the second intercept coefficient (representing the log-odds *not* choosing a 3 or higher—i.e., the judgment that the expected event was a little more upsetting or a lot more upsetting), allowing us to test whether participants were more likely to choose the expected or unexpected event as being more upsetting in cases where they did not choose the neither option. This analysis indicated that people were more likely to think that events that were unexpected were more upsetting than events that were expected, $\alpha_{\text{difference}} = 1.40$, 95% Posterior Uncertainty Interval [0.42, 2.08]. We observed this difference consistently across 15 of 16 items (see Figure 3 below).

We note that the most common response was to view both events, which involve identical harms, as equivalently upsetting. However, across almost every item participants showed a bias towards choosing unexpected events as more upsetting even though the harms incurred in both cases are identical. Thus, Study 2 replicates the tendency to view unexpected moral events as more upsetting using a more conservative judgment task.

Figure 3. Participants' responses by item in Study 2. Upper: Heatmap displaying probability of choosing each response on the 1 to 5 scale across items. Lower: Marginal expected response across items, with 95% Posterior Uncertainty Intervals calculated from an ordinal regression model. In both plots, a response of three indicates indifference, responses below three indicate a bias to choose the more expected event as more upsetting, and responses above three indicates a bias to choose the less expected event as more upsetting. Trials sorted by response, and scale was limited to aid readability.

Discussion

A foundational principle of information theory is that events with lower prior probability carry more information than those with relatively higher prior probability

(Shannon, 1948). Inspired by this observation, we have presented a formal theory which posits that people evaluate events by comparing their mental model of the world—what they know about the world—before and after an event. Because these mental models incorporate uncertainty over future events, our account predicts people’s expectations about the future states of the world will influence their evaluations of events. This theory offers a rational account of prior findings demonstrating the influence of people’s expectations on their evaluations of non-moral events (e.g., Mellers et al., 1997; Villano et al., 2020). But, our account makes more far-ranging predictions than prior theoretical perspectives (Bell, 1985; Gul, 1991; Loomes & Sugden, 1986; Mellers et al., 1997), predicting that expectations will influence people’s evaluation of events in general, including their evaluations of *morally-significant* outcomes. These predictions were borne out in two preregistered studies: people reported feeling more upset about unexpected harmful events than expected harmful events, even when the harm a victim suffered was the same. For instance, people judged that a robbery at a clothing store (a relatively unexpected event) was more upsetting than a robbery at a convenience store (a more expected event).

In demonstrating the influence of people’s expectations on their evaluations of morally consequential events, our studies add to the literature highlighting how moral intuitions are impacted by morally-irrelevant psychological factors. Prior research has found that people’s moral judgments are influenced by their “psychological distance” from the event (Trope & Liberman, 2010), how events are framed (e.g., Fetherstonhaugh et al., 1997; Rai & Holyoak, 2010), whether they are viewed as a norm (Alicke, Rose, & Bloom, 2011; Kahneman & Miller, 1986), their graphicness (Salerno, 2017), and how emotionally evocative an event is (Greene, 2017; Haidt, 2001; Horne & Powell, 2016). Our studies identified a novel factor impacting evaluations of morally-significant events: the likelihood of that event occurring.

Research suggesting that moral intuitions are influenced by morally-irrelevant factors has led many psychologists and philosophers to question the normative theses supported by

intuitive judgments on the grounds that these normative theses may be a by-product of flaws or quirks in our reasoning capacities (Greene, 2017; Haidt, 2001). On these accounts, moral intuitions are produced by “sub-optimal” heuristic cognitive mechanisms (see Haidt, 2001; Horne & Livengood, 2017), which often lead to (arguably) inconsistent ethical principles (Greene, 2017; Norcross, 2004). Perhaps the clearest challenges to the reliability of moral intuitions come from psychological experiments that show they’re are sensitive to purely psychological manipulations. For instance, Salerno and colleagues (2017) found that jurors’ intuitions about the guilt of a defendant are affected by color, but not black-and-white, crime scene photographs. Black-and-white and color photographs convey the same substantive facts about the case, meaning that the difference between these two conditions is purely psychological.

The present studies diverge from this prior work in that participant’s expectations were manipulated by substantive, but morally irrelevant, changes to the context of the events. However, the cognitive mechanism through which this morally-irrelevant factor impacts people’s intuitions is rationally-defensible – the probability of an event (arguably) corresponds to real features of the world, which people are appropriately attuned to in a wide variety of non-moral contexts (e.g., Joyce, 2019; Piantadosi, Tily, & Gibson, 2011; Rescorla & Wagner, 1972). Consequently, our findings highlight a way in which even rational cognitive processes (in this case, processes correctly attuned to an event’s prior probability) can give rise to indefensible moral intuitions.

These results raise questions for future research, in particular, about the “reliability” of moral intuitions (e.g., Sinnott-Armstrong, 2008). Several researchers have argued that moral intuitions are unreliable on the grounds that they are the products of non-normative cognitive processes (e.g., Carlsmith, 2006; Greene, 2017; Horne & Livengood, 2017; Sinnott-Armstrong, 2008). However, our findings illustrate cases where even rational cognitive mechanisms, at least as they are now commonly defined in research in cognitive science (Griffiths et al., 2010), may also produce flawed moral intuitions. Thus, it may be

that whether the psychological mechanism underlying a moral judgment (e.g., heuristic processes; deliberative reasoning) is generally regarded as “suboptimal” or “optimal” (Kahneman, 2011) does not itself imply that moral intuitions should be dismissed or trusted (see Horne & Livengood, 2017).

The present findings may also pose trouble for those who have argued that moral intuitions are generally reliable, and existing empirical work only impugns a small subset of our intuitions (for a discussion, see Liao, 2008). For instance, an ethicist might concede intuitions are only unreliable for “non-central” cases or (Weatherson, 2003). Our findings raise interesting questions about the viability of these arguments; by revealing the sensitivity of moral intuitions to expectations, our studies suggest that all moral intuitions may be at least somewhat contaminated, so to speak, by the influence of this morally-irrelevant factor. Our formal account assumes that human evaluations are always made with uncertain knowledge of the world, and will therefore partly reflect their expectations about the present and future. Thus, our findings could be interpreted as showing that all moral intuitions are, to some extent, contaminated. A question that remains then is where and when moral intuitions are sufficiently contaminated that they should be considered unreliable and that they should not guide our ethical theorizing. A full treatment of this issue is beyond the scope of this paper, but these findings highlight an important direction for future research.

To our knowledge, there are no existing ethical frameworks that differentiate between moral harms based on their prior probability alone (e.g., Alexander & Moore, 2016; Hursthouse & Pettigrove, 2016; Sinnott-Armstrong, 2019). On these grounds—among others—we have suggested people’s expectations should not impact their evaluations of morally-significant events. Still, it is clearly adaptive that a violation of our expectations would direct our attention to the events and situations about which our understanding is most in need of revision (Joyce, 2019; Rescorla & Wagner, 1972). There might even be instances where directing concern in this way is, if not morally defensible, at least

reasonable. For instance, an act of terrorism in a peaceful nation might be less expected because of our beliefs about the effectiveness of security forces, the stability of national relationships, and positive relationships among people groups in that nation. A surprising act of terrorism might therefore cast doubt on our deeper understanding of the geopolitical situation, leading us to suspect similar acts are more likely to occur in the future. In such contexts, it might be reasonable that this would cause us great moral concern.

Yet, however well-grounded the underlying psychological processes, there is a clear moral danger whenever expectations influence our moral concerns. In particular, these effects pose a significant threat when considering people's reactions to humanitarian challenges brought on by poverty or regional instability. In these cases, the influence of expectations on moral evaluations threatens to impose vicious and morally pernicious cycles. For instance, people living in geo-politically unstable regions or in the developing world are often those who are most affected by terrorism, famine, and natural disasters, and are the very people in greatest need of assistance and concern from the world at-large. However, for precisely these reasons, it is often unsurprising when harm befalls them which makes us less likely to act. Thus, the influence of expectations on moral evaluations threatens to reduce observer's concern for these victims most in need, making people less likely to donate time or money to aid victims, or to take political action to prevent further harms. That all of these morally problematic actions may at least in part be the consequences of rational mechanisms highlights how even cognitive processes which are adaptive, efficient, and optimal can lead to indefensible ethical intuitions.

References

- Abdellaoui, M., & Bleichrodt, H. (2007). Eliciting gul's theory of disappointment aversion by the tradeoff method. *Journal of Economic Psychology*, 28(6), 631–645.
- Abdellaoui, M., Bleichrodt, H., & Paraschiv, C. (2007). Loss aversion under prospect theory: A parameter-free measurement. *Management Science*, 53(10), 1659–1674.
- Alexander, L., & Moore, M. (2016). *Deontological ethics*. Stanford Encyclopedia of Philosophy.
<https://stanford.library.sydney.edu.au/entries/ethics-deontological/>.
- Alicke, M. D., Rose, D., & Bloom, D. (2011). Causation, norm violation, and culpable control. *The Journal of Philosophy*, 108(12), 670–696.
- Arrow, K. J. (1996). The economics of information: An exposition. *Empirica*, 23(2), 119–128.
- Bell, D. (1985). Disappointment in decision making under uncertainty. *Operations Research*, 33(1), 1–27.
- Brewer, M. B. (1999). The psychology of prejudice: Ingroup love or outgroup hate? *Journal of Social Issues*, 55(3), 429–444.
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28.
- Carlsmith, K. M. (2006). The roles of retribution and utility in determining punishment. *Journal of Experimental Social Psychology*, 42(4), 437–451.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Fetherstonhaugh, D., Slovic, P., Johnson, S., & Friedrich, J. (1997). Insensitivity to the value of human life: A study of psychophysical numbing. *Journal of Risk and Uncertainty*, 14(3), 283–300.
- Golub, S. A., Gilbert, D. T., & Wilson, T. D. (2009). Anticipating one's troubles: The

costs and benefits of negative expectations. *Emotion*, 9(2), 277–281.

Graham, D. A. (2015, November 16). *The empathy gap between Paris and Beirut*.

<https://www.theatlantic.com/international/archive/2015/11/>

[paris-beirut-terrorism-empathy-gap/416121/](https://www.theatlantic.com/international/archive/2015/11/paris-beirut-terrorism-empathy-gap/416121/). (Accessed: 2018-06-14)

Greene, J. D. (2017). The rat-a-gorical imperative: Moral intuition and the limits of affective learning. *Cognition*, 167, 66–77.

Greenslade, R. (2007, April 19). *A hierarchy of death*. <https://www.theguardian.com/commentisfree/2007/apr/19/thirtytwodieinamericanuniv>. (Accessed: 2018-06-14)

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010).

Probabilistic models of cognition: Exploring representations and inductive biases.

Trends in Cognitive Sciences, 14(8), 357–364.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica*, 59(3), 667–686.

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108(4), 814–834.

Horne, Z., & Cimpian, A. (2018). Subtle syntactic cues affect intuitions about knowledge:

Methodological and theoretical implications for epistemology. In J. Knobe,

T. Lombrozo, & S. Nichols (Eds.), *Oxford studies in experimental philosophy, volume 2* (pp. 7–37). Oxford University Press.

Horne, Z., & Livengood, J. (2017). Ordering effects, updating effects, and the specter of global skepticism. *Synthese*, 194(4), 1189–1218.

Horne, Z., & Powell, D. (2016). How large is the role of emotion in judgments of moral dilemmas? *PloS ONE*, 11(7), Article e0154780.

Hursthouse, R., & Pettigrove, G. (2016). *Virtue ethics*. Stanford Encyclopedia of Philosophy.

<https://stanford.library.sydney.edu.au/entries/ethics-virtue/>.

Joyce, J. (2019). *Bayes' Theorem*. The Stanford Encyclopedia of Philosophy.

<https://plato.stanford.edu/archives/spr2019/entries/bayes-theorem/>.

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69.

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93(2), 136–153.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292.

Liao, S. M. (2008). A defense of intuitions. *Philosophical Studies*, 140(2), 247–262.

Liddell, T. M., & Kruschke, J. K. (2018). Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology*, 79, 328–348.

Loomes, G., & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies*, 53(2), 271–282.

McElreath, R. (2016). *Statistical rethinking: A Bayesian course with examples in R and Stan*. CRC Press.

McGraw, A. P., Mellers, B. A., & Tetlock, P. E. (2005). Expectations and emotions of olympic athletes. *Journal of Experimental Social Psychology*, 41(4), 438–446.

Medvec, V. H., Madey, S. F., & Gilovich, T. (1995). When less is more: counterfactual thinking and satisfaction among olympic medalists. *Journal of Personality and Social Psychology*, 69(4), 603–610.

Medvec, V. H., & Savitsky, K. (1997). When doing better means feeling worse: The effects of categorical cutoff points on counterfactual thinking and satisfaction. *Journal of Personality and Social Psychology*, 72(6), 1284–1296.

Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science*, 8(6), 423–429.

- Nichols, S., & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Nous*, 41(4), 663–685.
- Norcross, A. (2004). Puppies, pigs, and people: Eating meat and marginal cases. *Philosophical perspectives*, 18, 229–245.
- Olivola, C. Y., & Sagara, N. (2009). Distributions of observed death tolls govern sensitivity to human fatalities. *Proceedings of the National Academy of Sciences*, 106(52), 22151–22156.
- Piantadosi, S. T., Tily, H., & Gibson, E. (2011). Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Sciences*, 108(9), 3526–3529.
- Rai, T. S., & Holyoak, K. J. (2010). Moral principles or consumer preferences? alternative framings of the trolley problem. *Cognitive Science*, 34(2), 311–321.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical Conditioning II Current Research and Theory* (Vol. 21, pp. 64–99). Retrieved from <http://homepage.mac.com/sanagnos/rescorlawagner1972.pdf>
- Richard, F. D., Bond, C. F., & Stokes-Zoota, J. J. (2003). One Hundred Years of Social Psychology Quantitatively Described. *Review of General Psychology*, 7(4), 331–363.
- Salerno, J. M. (2017). Seeing red: Disgust reactions to gruesome photographs in color (but not in black and white) increase convictions. *Psychology, Public Policy, and Law*, 23(3), 336–350.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.
- Sinnott-Armstrong, W. (2008). *Framing moral intuitions in Moral psychology, vol 2: The cognitive science of morality: Intuition and diversity* (W. Sinnott-Armstrong, Ed.). MIT Press.
- Sinnott-Armstrong, W. (2019). *Consequentialism*. Stanford Encyclopedia of Philosophy.

<https://stanford.library.sydney.edu.au/entries/consequentialism/>.

Tanner, C., Medin, D. L., & Iliev, R. (2008). Influence of deontological versus consequentialist orientations on act choices and framing effects: When principles are more important than consequences. *European Journal of Social Psychology*, 38(5), 757–769.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, 12(3), 183–206.

Trope, Y., & Liberman, N. (2010, April). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–63.

Villano, W. J., Otto, A. R., Ezie, C., Gillis, R., & Heller, A. S. (2020). Temporal dynamics of real-world emotion are more strongly linked to prediction error than outcome. *Journal of Experimental Psychology: General*.

Weatherson, B. (2003). What good are counterexamples? *Philosophical Studies*, 115(1), 1–31.

Appendix

Formalizing how expectations influence evaluations

590 We defined the value assigned to the signalling event, $V(S)$, as:

$$V(S) = E_{X|S}[u(X)] - E_X[u(X)] \quad (1)$$

591 Which we write more specifically as:

$$V(s = 1) = \sum_x u(x)p(X = x|s = 1) - \sum_x u(x)p(X = x) \quad (2)$$

592 Our goal is to express this relation in terms of $p(S)$. First, we manipulate $E_X[u(X)]$ by
 593 applying the expectation of conditional expectation and some algebra.

$$\begin{aligned} \sum_x u(x)p(X = x) &= \sum_s u(x)p(X = x|s)p(s) \\ &= \sum_x u(x)p(X = x|s = 1)p(s = 1) - \\ &\quad \sum_x u(x)p(X = x|s = 0)p(s = 0) \\ &= p(s = 1) \sum_x u(x)p(X = x|s = 1) - \\ &\quad p(s = 0) \sum_x u(x)p(X = x|s = 0) \\ &= p(s = 1) \sum_x u(x)p(X = x|s = 1) - \\ &\quad (1 - p(s = 1)) \sum_x u(x)p(X = x|s = 0) \end{aligned}$$

594 Substituting this quantity back into Equation 2, we obtain:

$$\begin{aligned} V(s = 1) &= \sum_x u(x)p(X = x|s = 1) - \\ &\quad p(s = 1) \sum_x u(x)p(X = x|s = 1) - \\ &\quad (1 - p(s = 1)) \sum_x u(x)p(X = x|s = 0) \end{aligned}$$

595 Finally, we factor the first two sums and combine them with the last sum to yield:

$$V(s = 1) = (1 - p(s = 1)) \sum_x u(x)p(X = x|s = 1) - u(x)p(X = x|s = 0) \quad (3)$$