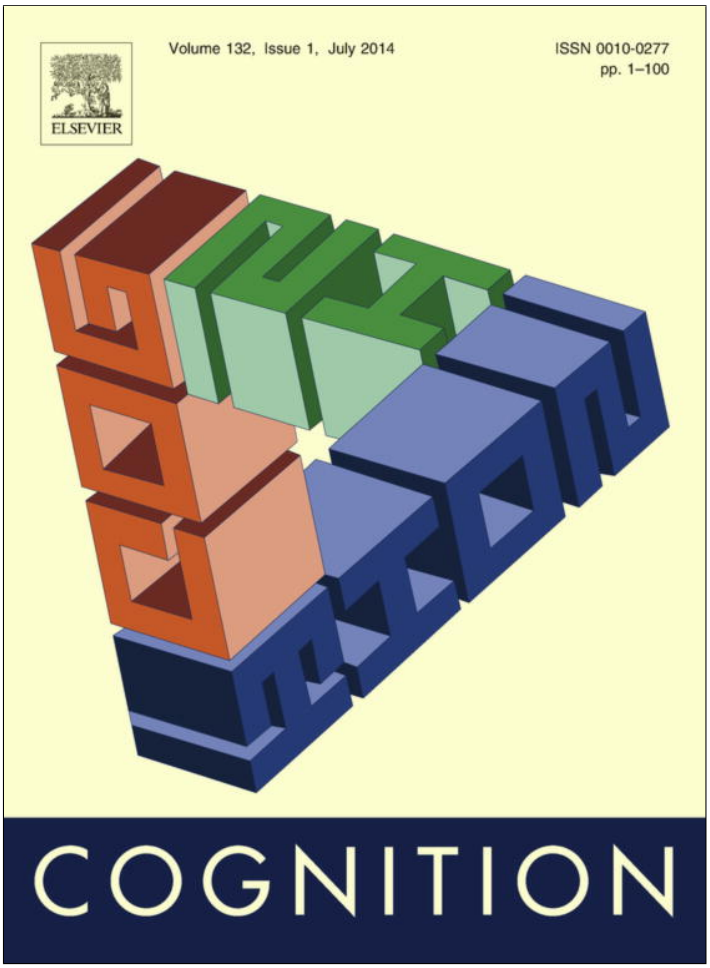

Cognition 132 (2014) 22–29

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

The moral pop-out effect: Enhanced perceptual awareness of morally relevant stimuli

Ana P. Gantman, Jay J. Van Bavel *

Psychology Department, New York University, 6 Washington Place, New York, NY 10003, USA

ARTICLE INFO

Article history:
Received 13 August 2013
Revised 20 February 2014
Accepted 23 February 2014
Available online 16 April 2014

Keywords:
Morality
Awareness
Vision
Perception
Ambiguity

ABSTRACT

People perceive religious and moral iconography in ambiguous objects, ranging from grilled cheese to bird feces. In the current research, we examined whether moral concerns can shape awareness of perceptually ambiguous stimuli. In three experiments, we presented masked moral and non-moral words around the threshold for conscious awareness as part of a lexical decision task. Participants correctly identified moral words more frequently than non-moral words—a phenomenon we term the *moral pop-out effect*. The moral pop-out effect was only evident when stimuli were presented at durations that made them perceptually ambiguous, but not when the stimuli were presented too quickly to perceive or slowly enough to easily perceive. The *moral pop-out effect* was not moderated by exposure to harm and cannot be explained by differences in arousal, valence, or extremity. Although most models of moral psychology assume the initial perception of moral stimuli, our research suggests that moral beliefs and values may shape perceptual awareness.

1. Introduction

In 2004, a woman from Florida sold a decade old, partially burnt, grilled cheese sandwich on eBay for $28,000 (Associated Press, 2004). The bidders clamored to pay over 14,000 times the value of the toast because an image of the Virgin Mary was perceived to be staring out from its charred center. Perceiving religious and moral iconography in natural phenomena, ranging from grilled cheese to bird feces, is surprisingly common (see http://jesusiseverywhere.net). In the current research, we examined whether moral concerns can shape the perception of ambiguous stimuli.

The vast majority of theories in moral psychology presume the perception of moral stimuli or "eliciting situations" (e.g., Haidt, 2001). In much of this research, participants are presented with vivid dilemmas and asked to render their moral judgment. Although moral perception is generally considered a necessary, pre-requisite for judgment and decision-making, there is good reason to believe that personal beliefs, moral identities, or moral motives may influence the basic awareness and interpretation of moral stimuli prior to action (see Aquino & Reed, 2002; Narvaez, Lapsley, Hagele, & Lasky, 2006). If so, these motives may literally lead people to see evidence of their moral values and beliefs in grilled cheese sandwiches or other perceptually ambiguous stimuli.

Research suggests that people have enhanced accessibility of highly valued or goal-relevant stimuli (Förster, Liberman, & Friedman, 2007), which may enhance perceptual awareness (Anderson, 2005; Anderson & Phelps, 2001; Bruner & Goodman, 1947; Vuilleumier, 2005). For example, food-related words are easier to recognize when one is hungry than when one is satiated (Radel & Clément-Guillotin, 2012; see also Balcetis, Dunning, & Granot, 2012). Given that morality satisfies multiple core motives, including the need for control (Kay, Gaucher, McGregor, & Nash, 2010), justice (Lerner & Miller, 1978), and to belong to and maintain social groups (Haidt & Graham, 2009), we hypothesized that perceptually ambiguous, moral stimuli

* Corresponding author. Tel.: +1 212 992 9627; fax: +1 212 995 4966.
E-mail address: jay.vanbavel@nyu.edu (J.J. Van Bavel).

http://dx.doi.org/10.1016/j.cognition.2014.02.007
0010-0277/

would be more likely to reach perceptual awareness than non-moral stimuli.

1.1. Present research

In three experiments, we examined whether perceptually ambiguous moral stimuli would be more likely to reach perceptual awareness than matched non-moral stimuli—a phenomenon we termed the *moral pop-out effect*. We hypothesized that morally relevant stimuli presented close to the threshold of perceptual awareness—a point at which they are perceptually ambiguous—would be recognized more often than non-moral stimuli. In order to do this, we presented words and non-words very briefly in a lexical decision task, and varied whether the words pertained to morality or not.

In Experiment 1, participants completed the lexical decision task with moral and non-moral words presented for 40 ms to ensure the words were close to the threshold of perceptual awareness (Gelskov & Kouider, 2010). In Experiment 2, we manipulated the presentation time of the stimuli to examine the entire time course during which morally relevant words "pop-out". We reasoned that words presented too quickly would fall below perceptual awareness and words presented too slowly would be perceived accurately, regardless of content. If moral concerns influence the awareness of perceptually ambiguous stimuli, then we should only find evidence of the *moral pop-out effect* for stimuli presented close to the threshold for perceptual awareness. In Experiment 3, we replicated the *moral pop-out effect* and investigated whether it might be strengthened after exposure to harm, a determinant of moral construal (Gray & Schein, 2012).

1.2. Experiment 1: The moral pop-out effect

In Experiment 1 we examined whether moral words would reach perceptual awareness (i.e., "pop-out") more frequently than non-moral words. We adapted a typical lexical decision task in which participants see a string of letters and indicate whether or not they comprise a word. Previous research has shown that faces presented for short durations (17 and 33 ms) are correctly identified at chance levels, whereas faces presented for longer durations (50 ms or longer) are correctly identified more frequently until they level off at nearly 100% accuracy (Gelskov & Kouider, 2010). We presented stimuli for 40 ms (an estimated threshold for perceptual awareness), to examine whether moral words had a lower threshold for perceptual awareness than non-moral words.

2. Methods

2.1. Participants

Twenty undergraduate students at New York University participated for partial course credit. One participant was excluded because the computer program crashed.[1]

[1] It was determined *a priori* to run this experiment until the end of the semester. This applies to all subsequent experiments reported here.

2.2. Procedure

Participants were told that the experiment was about visual acuity. The concept of morality was never mentioned. Instructions for the lexical decision task were administered in DirectRT on a Dell Optiplex 760 with a 60 Hz refresh rate. Participants completed the study alone in a dimly lit room and sat approximately 16 in. from the monitor. Stimuli appeared in white letters on a black background, size 24 font in the center of the computer monitor. The experiment began with a brief tutorial with five trials of non-moral words and non-words (*apple*, *speilc*, *building*, *kroaf*, *parrot*) at decreasing stimulus durations (500, 300, 100, 80, and 60 ms) to allow participants to learn the task. On every trial, participants saw a fixation cross in the center of the screen for 100, 200 or 300 ms (randomized to prevent participants from feeling lulled by a repetitious rhythm). The fixation cross was followed by the stimulus word presented in the center of the screen for approximately 40 ms, and then a 200 ms backwards mask of ampersands that corresponded to the number of letters in the word (e.g., 'useful' was followed by '&&&&&&'). The screen was black until participants responded (see Fig. 1). There were 82 moral/non-moral words and 40 non-words presented in random order. All materials (including full moral and non-moral word lists) are available online at: https://osf.io/7fk9b/.

After the lexical decision task, participants completed a number of exploratory individual difference measures we thought might be associated with the *moral pop-out effect*. These were global belief in a just world (Lipkus, 1991), religiosity (Batson, 1976), the moral foundations questionnaire (Graham, Nosek, Haidt, Ravi, & Ditto, 2011), and revised disgust sensitivity (Olatunji et al., 2007). None of these individual difference measures were significantly correlated with the accurate recognition of moral vs. non-moral words in this or any subsequent experiment ($ps > .08$) and we do not discuss these measures further.[2]

Participants then completed a manipulation check intended to validate the distinction between moral and non-moral words. The experimenter explained to each participant that they were to rate whether the words were related to the domain of morality (and not whether the words were moral vs. immoral or whether they could imagine a moral situation involving the word). For example, "hero" and "devil" are both in the moral domain, but "pilot" should be considered non-moral. Participants then rated 82 randomly presented words (for a full word list, see https://osf.io/7fk9b/), 41 that we assumed were moral (e.g., *moral*, *virtue*, *steal*, *sin*, *should*) and 41 that were non-moral (e.g., *useful*, *virtual*, *steel*, *trick*, *could*) on a five-point scale (from 1 = "*not at all moral*" to 5 = "*very moral*"). Participants rated the moral words used in the lexical decision task as more morally relevant ($M = 3.84$, $SD = 0.50$) than non-moral words ($M = 2.03$, $SD = 0.49$), $t(18) = 16.36$, $p < .001$, $\eta^2 = .94$. Paired samples *t*-tests revealed no

[2] We did detect, however, a marginally significant interaction between word type and the moral foundation of harm ($p < .08$), such that for those participants who reported that harm was relevant to their moral judgments, the moral pop-out effect was accentuated, $B = 0.32$, $SE = .17$, $p = .07$, $z = 1.83$, despite the lack of main effect, $B = -0.40$, $SE = .26$, $p = .31$, $z = 1.56$.

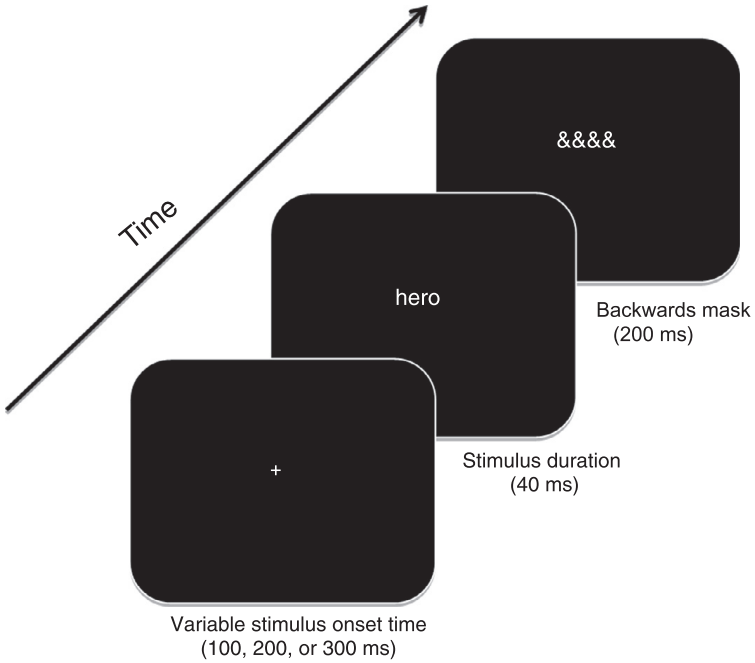

Fig. 1. Schematic of lexical decision task (Experiment 1). Participants saw a fixation cross, followed by either a moral word, non-moral word, or non-word displayed for approximately 40 ms. A backwards mask was presented for 200 ms. The screen remained black until "*w*" or "*o*" was pressed to indicate whether the string of letters appeared as a word or non-word, respectively.

discernible differences in word length ($p = .46$) or frequency in the English language ($p = .55$. Frequency data was taken from the 2012 edition of the Corpus of Contemporary American English (COCA; Davies, 2008).

3. Results and discussion

3.1. Analytic strategy for lexical decision-task

Given the categorical dependent measure and within-subjects design, we used generalized estimating equations (GEE) to estimate our regression parameters instead of ordinary least-squares regression (Zeger & Liang, 1986). This allowed us to take learning effects and other forms of interdependence among participants' ratings into account. Because our stimuli were presented in random order, an exchangeable correlation matrix was specified for all models (Ballinger, 2004). For analyses using GEE models, we report unstandardized regression coefficients (B), standard errors (SE) and Wald Z's.

3.2. Moral pop-out effect

In order to test whether moral words were more accurately recognized than non-moral words, we regressed categorization accuracy (word, non-word) against word type (non-moral = −1, moral = 1). As predicted, moral words were categorized more accurately as words ($M = 72\%$, $SE = 2\%$) than non-moral words ($M = 66\%$, $SE = 2\%$), $B = -.32$, $SE = .11$, $p < .01$, $z = 2.96$[3] (see Fig. 2). This *moral pop-out effect* was evidenced despite comparing two sets of words that were matched on frequency and length. These findings are consistent with the hypothesis that moral stimuli are privileged in perceptual awareness (see Table 1).

3.3. Experiment 2: The moral pop-out effect under conditions of perceptual ambiguity

To test whether the *moral pop-out effect* was only evident for perceptually ambiguous stimuli, we systematically varied the duration of stimulus presentation in Experiment 2. We hypothesized that we would replicate the results of Experiment 1, but only when stimuli were presented close to the threshold of perceptual awareness—approximately equidistant from durations associated with chance accuracy (50%) and perfect accuracy (100%); at durations that allow for approximately 75% accuracy. We hypothesized that at very fast stimulus durations, participants would be unable to discern whether the string of letters comprise a word, and at slower stimulus durations participants would be able to see all the stimuli easily and respond accurately.

4. Methods

4.1. Participants

Thirty-eight undergraduate students at New York University participated for course credit.

4.2. Materials

Participants completed the lexical decision task as described in Experiment 1, with the addition of variation in stimulus durations, such that participants saw words on the screen anywhere from 20 to 100 ms at approximately

[3] Because the stimulus duration was short, we did not predict a difference in reaction time responding to moral vs. non-moral words. Indeed, there is no effect of word type on reaction time when we regress log-transformed reaction time on word type ($p = .90$, $z = .14$).

Table 1
Summary statistics for experiments 1–3 and meta-analysis.

Experiment	Sample threshold of perceptual ambiguity	Mean	SE
Experiment 1	40 ms		
Moral words		.72	.02
Non-moral words		.66	.02
Experiment 2	40–50 ms		
Moral words		.75	.02
Non-moral words		.68	.02
Experiment 3	50–60 ms		
Moral words		.67	.02
Non-moral words		.61	.02
Meta-analysis	40–60 ms		
Moral words		.69	.01
Non-moral words		.65	.01

Note: thresholds of perceptual ambiguity were determined by accuracy falling between chance (50%) and complete accuracy (100%), hovering around 75%. Means and standard errors were calculated across participants.

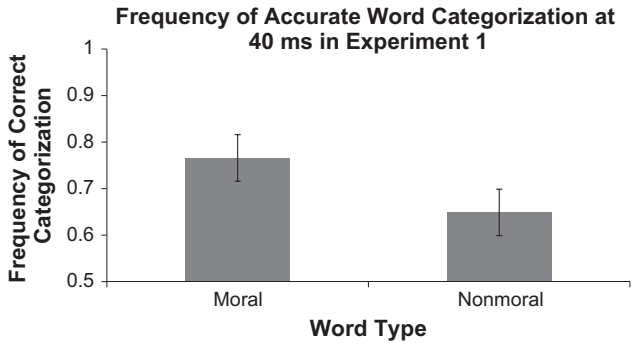

Fig. 2. Moral words were correctly categorized as words more frequently than non-moral words during a lexical decision task (Experiment 1). Bars represent standard errors.

10 ms intervals. This allowed us to identify the stimulus durations closest to the threshold of perceptual awareness in our sample (i.e., 75% accuracy). Due to randomization, the total number of letter strings presented at a given stimulus duration varied for each person but led to no systematic differences.

4.3. Procedure

Participants were told that the experiment was about following current events and language skills. Prior to the lexical decision task, participants read an excerpt from a news article surveying the Flea Markets in New York City (Parker, 2011) to maintain the cover story. The excerpt contained no words from the lexical decision task nor any mention of morality. After being told that they would answer questions about the article later, participants completed the lexical decision task and a number of individual difference measures, including the global belief in a just world (Lipkus, 1991), religiosity (Batson, 1976), and revised disgust sensitivity (Olatunji et al., 2007).

5. Results and discussion

5.1. Overall visibility curve

As expected, participants' performance on the lexical decision task improved as the presentation times increased. At 20 ms, participants correctly identified letter strings as words 27% of the time, well below chance (27% moral, 26% non-moral). At longer presentation times, performance increased significantly, with participants correctly identifying 65% (69% moral, 62% non-moral), 78% (81% moral, 76% non-moral), 83% (84% moral, 81% non-moral), 90% (89% moral, 91% non-moral), 93% (94% moral, 92% non-moral), 96% (97% moral, 94% non-moral), and 98% (100% moral, 97% non-moral) of words successfully at 40, 50, 60, 70, 80, 90 and 100 ms, respectively (see Table 1). We ran a logistic regression to fit the log odds of word categorization (word or not word) on stimulus duration treated as a continuous variable. We found a significant effect of stimulus duration on accuracy: $B = .06$, $SE = .003$, Wald $X^2 = 548.31$, $p < .001$, $r = .49$. Fig. 3 shows an overall increase in visibility curve as stimulus durations increase, consistent with previous research (Gelskov & Kouider, 2010).

5.2. Moral pop-out effect

Our primary hypothesis was that moral stimuli would "pop-out" when they were presented close to the threshold of perceptual awareness—approximately equidistant from durations associated with chance accuracy (50%) and perfect accuracy (100%). Therefore, we grouped stimulus durations accordingly such that 30 ms and below (42% accuracy) were considered "fast", 40 and 50 ms together (72% accuracy) were considered "moderate", and durations of 60 ms or higher (89% accuracy) were considered "slow." Moderate stimuli were close to the threshold of perceptual awareness.

In order to detect the *moral pop-out effect*, we separately analyzed each group of stimulus durations by performing the same GEE analysis as in Experiment 1. Replicating the previous experiment, when the stimuli were presented close to the threshold for perceptual awareness (i.e., when accuracy was approximately 75%), we observed evidence of the *moral pop-out effect*. We found a marginally significant effect of word type on accuracy: $B = .07$, $SE = .09$, $p = .06$, $z = .83$ (see Table 1). As seen in Fig. 3, the *moral pop-out effect* is largest at stimulus durations close to the threshold for perceptual awareness (40–50 ms). This is

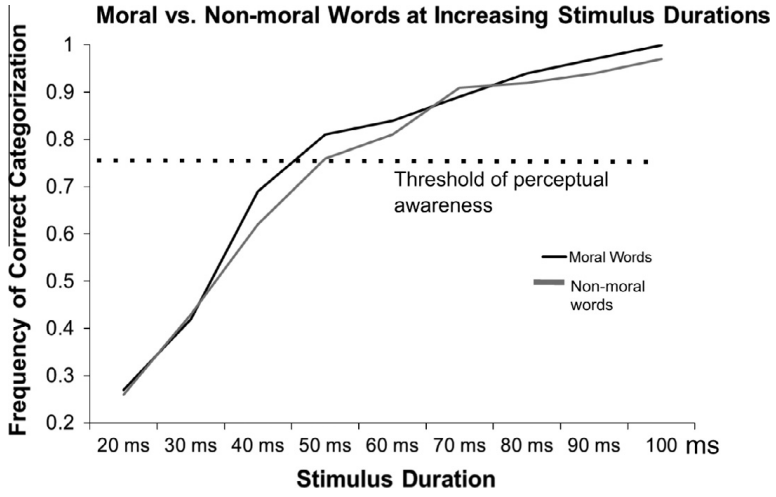

Fig. 3. Words become increasingly likely to be correctly categorized as words as stimulus presentation times increase (Experiment 2). Moral words were correctly identified more frequently than non-moral words when the stimulus was presented for a moderate length of time (40–50 ms). Overall means are displayed for ease of interpretation despite interdependence (ms = milliseconds).

consistent with our hypothesis that moral concerns shape awareness of perceptually ambiguous stimuli.

5.3. Experiment 3: Harm and the moral pop-out effect

In Experiment 3, we sought to examine the effect of harm on moral perception. Although most theories of moral psychology are silent with regard to factors that change the perceptual awareness of moral stimuli, recent work has proposed that the primary determinant of whether a situation is considered moral or not is whether that event contained an agent harming another being (Gray & Schein, 2012). Accordingly, harm may prime people to recognize moral words, yielding a stronger *moral pop-out* effect. On the other hand, the *moral pop-out effect* may be insensitive to non-motivational manipulations because it arises due to the chronic accessibility of potentially satiating moral stimuli.[4] In order to help test these competing accounts, participants who came into the lab were randomly assigned to read one of two vignettes (harm vs. no harm) before completing the lexical decision task. In the *no harm condition*, someone drives to a party, becomes drunk, and chooses not to drive home. In the *harm condition*, the story is the same except someone chooses to drive home, hits another car, killing someone in it, and drives away.

6. Methods

6.1. Participants and design

Seventy-seven undergraduates from New York University participated for course credit over two semesters.[5]

6.2. Materials

Participants read one of two vignettes prior to completing the same lexical decision task from the previous experiments. In both conditions, participants saw a photo of Steve, a white college-aged male. Beneath Steve's photo was a description of him driving to a party hosted by friends. Steve ends up enjoying the party and becomes intoxicated. In the *no harm* condition, a friend offered to drive Steve home and he accepted. In the *harm* condition, Steve decided to drive home despite being drunk and collided with another car, killing the person in the passenger's seat. Steve panicked and drove away.

6.3. Procedure

Participants were told that the experiment was about following current events and language skills. Prior to the lexical decision task, participants read one of two vignettes concerning drunk driving (with or without harm) and then completed the lexical decision task. Similar to Experiment 2, we randomly presented words for 40, 50, 60, and 70 ms to identify perceptually ambiguous stimuli. Again, this allowed us to identify the stimulus durations closest to the threshold of perceptual awareness in our sample (i.e., 75% accuracy).

7. Results

7.1. Overall visibility curve

As expected, participants' performance on the lexical decision task improved as the presentation times increased. At 40 ms, participants correctly categorized letter strings as words 51% of the time—near chance (52% moral, 51% non-moral). At 50 ms, participants correctly categorized 65% of the letter strings (67% moral, 62% non-moral), while at 60 ms accuracy jumped to 74% (75% moral, 73% non-moral). Finally, when letter strings were presented

[4] A motivational manipulation should activate a relevant need or motive (e.g., indicate an injustice) and then either successfully or unsuccessfully satiate that need (e.g., restore justice or not).

[5] Though we had previously decided to determine our sample size by the end of the semester, reviewers suggested that we rerun the study for an additional semester to increase statistical power.

for 70 ms, participants responded accurately 83% (86% moral, 80% non-moral) of the time (see Table 1). First, we ran a logistic regression on the log odds of word categorization (word or non-word) against stimulus duration (treated as a continuous factor). We found a significant effect of stimulus duration on categorization accuracy such that the longer a word was presented on the screen, the greater the likelihood that it was correctly categorized as a word, $B = .06$, $SE = .002$, Wald $X^2 = 759.38$, $p < .001$, $r = .38$.

7.2. Moral pop-out effect

Our primary hypothesis was that moral stimuli would "pop-out" when they were presented close to the threshold of perceptual awareness—approximately equidistant from durations associated with chance accuracy (50%) and perfect accuracy (100%). Based on accuracy scores in this sample, stimulus durations of 40 ms were considered "fast", durations of 50–60 ms were considered "moderate", and durations of 70 ms were considered "slow."[6]

In order to detect the *moral pop-out effect*, we separately analyzed each group of stimulus durations by performing the same GEE analysis as in Experiments 1 and 2. Replicating the previous experiments, when the stimuli were presented close to the threshold for perceptual awareness (here, 50–60 ms) we observed evidence of the *moral pop-out effect* (see Table 1). We found a significant effect of word type on accuracy: $B = .15$, $SE = .05$, $p = .005$, $z = 2.77$. In other words, the *moral pop-out effect* occurred at stimulus durations that were close to the threshold for perceptual awareness. As in the first two experiments, no *moral pop-out effect* was evident when stimuli were presented too briefly (40 ms; $p = .12$), such that the letter strings could not be detected, or when they were presented for so long that participants nearly always accurately categorized the letter strings (70 ms; $p = .11$).

We were not, however, able to detect a main effect of the drunk driving vignette at any of the stimulus durations (all $ps > .46$), or an interaction between the condition and whether the words were moral or non-moral (all $ps > .36$). The 95% confidence interval for the interaction between word type (moral vs. non-moral) and vignette (lower bound = −.19, upper bound = .19) includes zero and suggests that the effect of harm, if it exists, may be in either direction. While this is consistent with our hypothesis that moral concerns shape awareness of perceptually ambiguous stimuli, and the *moral pop-out* effect was not moderated by exposure to harm in this sample, limited conclusions should be drawn from this null effect. It is still possible that the *moral pop-out effect* may be increased via priming of morally relevant constructs if a different prime is used. Alternatively, a motivational manipulation may be needed. We highlight, instead, that the *moral pop-out effect* remains – when stimuli are presented close to the threshold for perceptual awareness, moral words are accurately categorized more frequently than non-moral words.

[6] Including 70 ms in the "moderate" group only makes the effect larger, $B = .21$, $SE = .07$, $p = .004$, $z = 2.85$.

7.3. Meta-analysis combining Experiments 1–3

In Experiments 1–3, we found evidence for the *moral pop-out effect*. When letter-strings are presented around the threshold for perceptual awareness, moral words are more frequently categorized as words than non-moral words. While all three experiments show the same general effect, each experiment has relatively few participants, leading us to conduct a meta-analysis across the three studies in order to examine the strength of the *moral pop-out effect*. Given that the studies use different stimulus durations, we have divided them into three groups, 30 ms and shorter are regarded as fast, 40–60 ms are regarded as moderate, and 70 ms and faster are regarded as slow, effectively pooling the data at moderate durations, as the shorter and longer presentation times are almost entirely from Experiment 2 (though 70 ms was also included in Experiment 3). Overall, participants correctly identified words at moderate stimulus durations 67% of the time (69% moral, 65% non-moral). Consistent with our experiments, moral words are accurately categorized as words more frequently than non-moral words, at moderate stimulus durations, $B = .20$, $SE = .05$, $z = 4.41$, $p < .0001$. The 95% confidence interval (lower bound $B = .11$, upper bound $B = .29$) indicates that the effect is robust. We also used meta-analysis to examine potential moderation of the *moral pop-out effect* by belief in a just world, religiosity, and disgust sensitivity, and found no effects, all $ps > .2$.[7]

Using meta-analysis, we were able to assess potential confounds relating to our moral and non-moral word lists. In order to ensure that the *moral pop-out effect* is not due to differences in elicited arousal or extremity, we asked a new sample of undergraduates to rate each of the words in terms of how "emotionally arousing" they found them (from "not at all emotionally arousing" to "extremely emotionally arousing" on a 7-point Likert scale from 1 to 7) and how positive or negative they found them (from "extremely negative" to "extremely positive" on a 7-point Likert scale from −3 to 3). Participants indeed found moral words to be more emotionally arousing, ($M = 3.98$, $SD = 1.05$) than non-moral words ($M = 3.50$, $SD = .94$), $t(18) = 2.94$, $p = .009$. Moral words were perceived as more negative ($M = -.52$, $SD = .27$) than non-moral words ($M = .05$, $SD = .20$), $t(20) = 10.37$, $p < .0001$. We also calculated the absolute value for the valence ratings to determine extremity or use of the anchors of the scale. Moral words were rated more extremely ($M = 1.57$, $SD = .48$) than non-moral words ($M = .44$, $SD = .10$), $t(20) = 11.11$, $p < .0001$.[8]

In order to ensure that these differences do not account for the *moral pop-out effect*, they were included in the meta-analysis along with the moral vs. non-moral dimension. While valence is a significant predictor of accurately categorizing a letter string as a word, $B = .04$, $SE = .02$, $z = 2.59$, $p = .009$, including this factor in the analysis did not eradicate the effect moral vs. non-moral words,

[7] A main effect for disgust sensitivity was found, such that those higher in disgust sensitivity show greater accuracy of word recognition regardless of word-type, $B = .01$, $SE = .005$, $z = 2.20$, $p = .02$.

[8] Differences in degrees of freedom arose from two participants failing to complete the "emotional arousal" part of the survey.

$B = .18$, $SE = .05$, $z = 3.77$, $p < .0001$. Neither of the tested predictors significantly predicted accurate word categorization at slow or fast stimulus durations. Self-reported emotional arousal elicited by the stimulus words did not significantly predict accurate categorization of the letter strings, $B = -.03$, $SE = .03$, $z = .96$, $p = .34$ at moderate stimulus durations.[9] Finally, while extremity is also a significant predictor of accurately categorizing a letter string as a word, $B = -.15$, $SE = .04$, $z = 3.55$, $p < .0001$, including this factor in the analysis did not eliminate the effect of moral vs. non-moral words, $B = .13$, $SE = .05$, $z = 2.50$, $p = .01$ at moderate stimulus durations. Taken together, moral (vs. non-moral) words are more likely to be accurately categorized as words when presented ambiguously – here, around the threshold for visual perception. This effect cannot be explained by differences in moral vs. non-moral words in arousal, valence, or extremity even though moral words appear greater on all three of dimensions. While the meta-analysis indicates that this is a reliable effect, so far we have only tested undergraduates, and conclusions cannot be extended to all individuals.

8. General discussion

Across three experiments, we found evidence for the *moral pop-out effect*. People exposed to letter strings presented close to the threshold for perceptual awareness were able to correctly identify moral words more frequently than non-moral words. However, this *moral pop-out effect* was only evident at stimulus durations close to the threshold of perceptual awareness (i.e., around 75% accuracy). Finally, the *moral pop-out effect* was not moderated by exposure to harm and cannot be explained by differences in valence, extremity, or arousal. Our findings suggest that perceptually ambiguous moral stimuli may require fewer processing prerequisites than non-moral stimuli in order to reach conscious awareness—possibly because moral stimuli satisfy core motives.

8.1. Perceptual ambiguity

The current research suggests that moral stimuli only *"pop-out"* when they are presented at the threshold for perceptual awareness. In contrast, when stimuli are presented too quickly they fall below perceptual awareness and when they are presented too slowly they are perceived accurately, regardless of content. As a result, it is likely that the *moral pop-out effect* may be seen at different stimulus durations based on the sample. If the sample contains a majority of people with a lower visibility threshold, then the *moral pop out effect* may manifest at 40 ms. However, if the sample has a majority of people with a higher visibility threshold the *moral pop-out effect* may manifest at longer presentation times (50–60 ms). Indeed, people may have different thresholds for visibility (Radel & Clément-Guillotin, 2012). As a result, we cannot conclude from these studies that moral stimuli "pop out" at any particular time, but merely that moral stimuli are more likely to reach awareness under conditions of perceptual ambiguity. The same logic likely applies to the perception of religious iconography in everyday objects: only perceptually ambiguous objects—like the pattern of char on a grilled cheese sandwich—afford the perceptual system the opportunity to perceive religious or moral iconography.

8.2. Future directions and implications

Classic models of moral psychology have focused on the role of reasoning (e.g., Kohlberg, 1975) and more recent models have argued that moral intuitions are rapid, reflexive responses to certain stimuli, such as incest (e.g., Haidt, 2001). However, many of these models are silent with regard to the perceptual processes that precede the initial intuition. The current research—as well as frequent reports of religious and moral iconography appearing in every day objects—suggests that perceptual factors, such as ambiguity, may also play an important role in moral psychology. Future work should examine the downstream consequences of perception on moral judgment and decision-making, attempt to develop models of morality that incorporate the role of perception, and replicate these findings in variegated samples.

Future work should also examine whether or not perceptual awareness can be affected by various top-down influences. Personal beliefs may alter perceptual awareness, as religious individuals and those with paranormal beliefs are more likely to see faces in inanimate objects than individuals without these beliefs (Riekki, Lindeman, Aleneff, Halme, & Nuortimo, 2012). Moreover, although we did not find evidence that priming harm alters perceptual awareness, there is evidence that ostensibly automatic moral responses may be susceptible to top-down influences. For example, people who view a scene with a morally good vs. bad actor show gaze preferences for good vs. bad outcomes for those actors based on their just world beliefs and expectations, increasing the likelihood they will see a deserved outcome (Callan, Ferguson, & Bindemann, 2013). This work suggests that moral concerns may not only change the way we evaluate and act in a situation (Van Bavel, Packer, Haas, & Cunningham, 2012), but that this change may occur early enough in processing to tune perceptual and attentional processes.

Recent research has also shown that individuals vary in their sensitivity to justice based on the accessibility of these concerns, and that this changes their interpretation of ambiguous situations, both for those dispositionally high in this accessibility and for those individuals recently exposed to an unjust situation in the lab (Baumert & Schmitt, 2011). Similarly, individuals high in disgust sensitivity see finer gradations of gray, and after exposure to disgusting stimuli, are able to detect small deviations from white (Sherman, Haidt, & Clore, 2012). Though we found that neither belief in a just world, disgust sensitivity, religiosity nor recent exposure to a moral vignette moderated our effect, we do suspect that accessibility of moral concepts in the form of motivational concerns may increase the magnitude of the *moral pop-out effect*. For example, there is increasing evidence that current motivation shapes

[9] Interestingly, at slow stimulus durations (above 70 ms), arousal and extremity but not moral vs. non-moral words, predicts accurate categorization, $B = -.22$, $SE = .10$, $z = 2.13$, $p = .03$ and $B = -.31$, $SE = .13$, $z = 2.31$, $p = .02$, respectively.

perception of both threatening (Xiao & Van Bavel, 2012) and goal-relevant targets (Balcetis, Dunning, & Granot, 2012), such that they appear closer than irrelevant stimuli. If the *moral pop-out effect* is due to morality's role in fulfilling chronic core motives across individuals, such as belief in a just world, we expect that individuals might be more likely to see religious and moral iconography in their breakfast, if they just learned about an injustice in the morning paper.

9. Conclusion

The current research suggests that moral concerns shape our basic awareness of perceptually ambiguous stimuli. One implication is that these perceptions may reinforce the pre-existing beliefs that gave rise to them in the first place, creating a "virtuous cycle" between beliefs and perceptions. This potential moral confirmation bias may prove to be something of a double-edged sword. On the one hand, people with strong beliefs may find conviction, clarity and solace in these perceptual confirmations of their beliefs. On the other, if these signs depend on their perceived miraculous origin for their value, they may not be as powerful as they seem. Regardless, the enhanced perceptual awareness of moral stimuli may help shed some light on sightings of religious and moral iconography in everyday objects like grilled cheese sandwiches.

Acknowledgments

The authors would like to thank Lisa Kaggen, Jenny Ray, Sharareh Noorbaloochi, Shona Tritt and members of the *NYU Social Perception and Evaluation* Lab (@vanbavellab) for helpful comments on this manuscript, and Justin Lieberknecht, Shalini Sivathasan, and Nick Ungson for help with data collection. Responsibilities: AG and JVB designed studies, AG analyzed studies with input from JVB, AG and JVB wrote the manuscript.

References

Anderson, A. K. (2005). Affective influences on the attentional dynamics supporting awareness. *Journal of Experimental Psychology: General, 134*, 258–281.

Anderson, A. K., & Phelps, E. A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature, 411*, 305–309.

Aquino, K., & Reed, A. II, (2002). The self-importance of moral identity. *Journal of Personality and Social Psychology, 83*, 1423–1440.

Associated Press (2004, November 23). 'Virgin Mary grilled cheese' sells for $28,000. *NBC News*. Retrieved from <http://www.msnbc.msn.com/id/6511148/ns/us_news weird_news/t/virgin-mary-grilled-cheese-sells/#.UN32MonjlEA>.

Balcetis, E., Dunning, D., & Granot, Y. (2012). Subjective value determines initial dominance in binocular rivalry. *Journal of Experimental Social Psychology, 48*, 122–129.

Ballinger, G. A. (2004). Using generalized estimating equations for longitudinal data analysis. *Organizational Research Methods, 7*, 127–150.

Batson, C. D. (1976). Religion as prosocial: Agent or double agent? *Journal for the Scientific Study of Religion, 15*, 29–45.

Baumert, A., & Schmitt, M. (2011). Justice-sensitive interpretations of ambiguous situations. *Australian Journal of Psychology, 61*, 6–12.

Bruner, J. S., & Goodman, C. C. (1947). Value and need as organizing factors in perception. *Journal of Abnormal and Social Psychology, 42*, 33–44.

Callan, M. J., Ferguson, H. J., & Bindemann, M. (2013). Eye movements to audiovisual scenes reveal expectations of a just world. *Journal of Experimental Psychology: General, 142*, 34–40.

Davies, M. (2008). *The Corpus of Contemporary American English*: *450 million words, 1990–present*. Available online at <http://corpus.byu.edu/coca/>.

Förster, J., Liberman, N., & Friedman, R. S. (2007). Seven principles of goal activation: A systematic approach to distinguishing goal priming from priming of non-goal constructs. *Personality and Social Psychology and Review, 11*, 211–233.

Gelskov, S. V., & Kouider, S. (2010). Psychophysical thresholds of face visibility during infancy. *Cognition, 114*, 285–292.

Graham, J., Nosek, B. A., Haidt, J. I., Ravi, K. S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology, 101*, 366–385.

Gray, K., & Schein, C. (2012). Two minds vs. two philosophies: Mind perception defines morality and dissolves the debate between deontology and utilitarianism. *Review of Philosophy and Psychology, 3*, 405–423.

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

Haidt, J., & Graham, J. (2009). Planet of the Durkheimians, where community, authority, and sacredness are foundations of morality. In J. Jost, A. C. Kay, & H. Thorisdottir (Eds.), *Social and psychological bases of ideology and system justification* (pp. 371–401). New York: Oxford University Press.

Kay, A. C., Gaucher, D., McGregor, I., & Nash, K. (2010). Religious belief as compensatory control. *Personality and Social Psychology Review, 14*, 37–48.

Kohlberg, L. (1975). The cognitive-developmental approach to moral education. *Phi Delta Kappan*, 670–677.

Lerner, M. J., & Miller, D. T. (1978). Just world research and the attribution process: Looking back and ahead. *Psychological Bulletin, 85*, 1030–1051.

Lipkus, I. (1991). The construction and preliminary validation of a global belief in a just world scale and the exploratory analysis of the multidimensional belief in a just world scale. *Personality and Individual Differences, 12*, 1171–1178.

Narvaez, D., Lapsley, D. K., Hagele, S., & Lasky, B. (2006). Moral chronicity and social information processing: Tests of a social cognitive approach to the moral personality. *Journal of Research in Personality, 40*, 966–985.

Olatunji, B. O., Tolin, D. F., Sawchuk, C. N., Williams, N., Abramowitz, J. S., Lohr, J. M., et al. (2007). The disgust scale: Item analysis, factor structure, and suggestions for refinement. *Psychological Assessment, 19*, 281–297.

Parker, A. (2011, May 13). The flea marketing of New York. *The New York Times*. Retrieved from <http://www.nytimes.com/2011/05/15/nyregion/the-flea-marketing-of-new-york city.html?pagewanted=all&_r=0>.

Radel, R., & Clément-Guillotin, C. (2012). Evidence of motivational indluences in early visual perception: Hunger modulates conscious access. *Psychological Science, 23*, 232–234.

Riekki, T., Lindeman, M., Aleneff, M., Halme, A., & Nuortimo, A. (2012). Paranormal and religious believers are more prone to illusory face perception than skeptics and nonbelievers. *Applied Cognitive Psychology*, Published online in Wiley Online Library (wileyonlinelibrary.com).

Sherman, G. D., Haidt, J., & Clore, G. L. (2012). The faintest speck of dirt: Disgust enhances the detection of impurity. *Psychological Science, 23*, 1506–1514.

Van Bavel, J. J., Packer, D. J., Haas, I. J., & Cunningham, W. A. (2012). The importance of moral construal: Moral versus non-moral construal elicits faster, more extreme, universal evaluations of the same actions. *PLoS ONE, 7*, e48693.

Vuilleumier, P. (2005). How brains beware: Neural mechanisms of emotional attention. *Trends in Cognitive Sciences, 9*, 585–594.

Xiao, Y. J., & Van Bavel, J. J. (2012). See your friends close, and your enemies closer: Social identity and identity threat shape the representation of physical distance. *Personality and Social Psychology Bulletin, 38*, 959–972.

Zeger, S. L., & Liang, K. Y. (1986). Longitudinal data analysis for discrete and continuous outcomes. *Biometrics, 42*, 121–130.