References

1. Gantman, A.P. and Van Bavel, J.J. (2015) Moral perception. *Trends Cogn. Sci.* 19, 631–633
2. Sherman, G.D. *et al.* (2012) The faintest speck of dirt: disgust enhances the detection of impurity. *Psychol. Sci.* 23, 1506–1514
3. Anderson, E. *et al.* (2011) The visual impact of gossip. *Science* 332, 1446–1448
4. Callan, M.J. *et al.* (2013) Eye movements to audiovisual scenes reveal expectations of a just world. *J. Exp. Psychol. Gen.* 142, 34–40
5. Baumert, A. *et al.* (2010) Justice sensitivity and the processing of justice-related information. *Eur. J. Pers.* 25, 386–397
6. Gantman, A.P. and Van Bavel, J.J. (2014) The moral pop-out effect: enhanced perceptual awareness of morally relevant stimuli. *Cognition* 132, 22–29
7. Meyer, D.E. and Schvaneveldt, R.W. (1971) Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *J. Exp. Psychol.* 90, 227–234
8. Collins, A.M. and Loftus, E.F. (1975) A spreading-activation theory of semantic processing. *Psychol. Rev.* 82, 407–428
9. Firestone, C. and Scholl, B.J. (2015) Enhanced visual awareness for morality and pajamas? Perception vs. memory in 'top-down' effects. *Cognition* 136, 409–416

Letter

See for Yourself: Perception Is Attuned to Morality

Ana P. Gantman[1] and Jay J. Van Bavel[1,*]

Perception appears preferentially attuned to moral content [1]. Despite the centrality of both morality and perception in cognitive science, little work has attempted to bridge these fields. Research on moral perception has the potential to inform our understanding of morality and perception, and may have important consequences for policy.

Firestone and Scholl recently argued that moral perception does not exist [2]. They claimed that moral perception is not about morality, that evidence of moral perception reflects visual processing 'in a trivial and unexciting sense', or does not pertain to perception [2]. We list the evidence for their claims below and invite you, the reader, to be the judge.

Isn't that Morality?

The stimuli from previous research that Firestone and Scholl claim 'do not invoke morality', include hitting a small child, making a racist comment, and defecating on a crowded street [3]. Extensive evidence suggests that actions such as these shape judgments of moral character [4], but perhaps we have to agree to disagree with Firestone and Scholl.

Isn't that Interesting?

Firestone and Scholl assert that it is 'unexciting' to know that people who learn about a character's bad actions subsequently attend to depictions of bad outcomes because they expect justice [5]. They asserted 'it can hardly be a new or bold claim that people look at what they expect'. While we agree that the relation between moral concerns and attention is mediated by basic cognitive processes, such as expectations, we disagree that this is 'trivial and unexciting'.

The goal of cognitive science is building process-oriented models. Instead of dismissing this research, we humbly submit that more work should focus on the processes underlying morality. This approach seems especially important in domains such as conflict resolution and legal decision-making [6,7].

Are These Phenomena Analogous?

We also disagree with Firestone and Scholl's [2] interpretation of the moral pop-out effect. We found that people correctly detect moral words (e.g., kill) more frequently than non-moral words [e.g., die; matched for length and frequency (http://corpus.byu.edu/coca/)], but only when the words were presented near the threshold for awareness (~40–60 ms; [8]). Moreover, the moral pop-out effect remained after adjusting for ratings of word valence, emotionality, and intensity. We suggested that moral words more readily reached perceptual awareness compared with non-moral words.

Firestone and Scholl [9] recently successfully reproduced the moral pop-out effect, and allegedly similar fashion and transportation pop-out effects. They argued that moral pop-out can be fully explained by semantic priming because 'relatedness is the key factor in such effects, and thus that memory, not perception, improves detection of morally related words' ([12] p. 43). Their claim hinges on similarities between morality and fashion and/or transportation pop-out effects. However, they did not randomly assign participants to detect moral versus fashion and/or transportation words; neither did they obtain sufficient power to test their claim that these other semantic categories show 'entirely analogous' effects to morality ([9] p. 411). As such, any comparisons they made between moral versus fashion and/or transportation effects seem speculative.

To test for semantic priming, they predicted that 'moral words (e.g., crime) may be easier to detect when presented in the context of other moral words (e.g., guilty)–whereas random non-moral words (e.g., steel) are no easier to detect in the context of other random words (e.g., tired)' [2]. The authors predicted that fashion and/or transportation words were easier to detect when presented in the context of repeated fashion and/or transportation words ($M = 81.3\%$) compared with nonrepeated fashion and/or transportation words ($M = 76.0\%$), whereas random control words were no easier to detect in the context of other control words ($M = 74.8\%$) compared with nonrepeated control words ($M = 72.7\%$) [9]. Thus, fashion and/or transportation words do appear more related to one another than do control words.

Curiously, however, Firestone and Scholl did not report the analogous means for their morality study, despite the fact that it was central to their hypothesis (which we quoted above). We are keen to see these values in print.

It is trivially true that semantic memory is implicated in moral word detection as humans learn what stimuli are relevant to the moral domain via acculturation

[10,11]. However, it seems that the burden of proof is on Firestone and Scholl to design and fully report a study that demonstrates how fashion and/or transportation pop-out is 'entirely analogous' to moral pop-out.

Concluding Remarks

We leave it to you, the reader, to evaluate Firestone and Scholl's claims. We fail to see how hitting a small child does not pertain to morality, knowing that expectations of justice affect attention is unexciting, or formal claims about similarity can be made without randomly assigning people to conditions or even presenting similar data. Clearly, moral perception is a provocative topic. Perhaps the only thing we can agree on is that more research is needed.

Acknowledgments

The authors would like to thank Chaz Firestone and Brian Scholl for thoughtful comments about this work, as well as Leor Hackel, Peter Mende-Siedlecki, Greg Murphy, Mickey Inzlicht, Daryl Cameron, Wil Cunningham, Bob Rehder, Jon Winawer, Todd Gureckis, and all members of the NYU Social Perception and Evaluation Lab for helpful feedback on this research.

[1]Psychology Department, New York University, 6 Washington Pl, New York, NY 10003, USA

*Correspondence: jay.vanbavel@nyu.edu (J.J. Van Bavel).

http://dx.doi.org/10.1016/j.tics.2015.12.001

References

1. Gantman, A.P. and Van Bavel, J.J. (2015) Moral perception. *Trends Cog. Sci.* 19, 631–633
2. Firestone, C. and Scholl, B.J. (2015) 'Moral perception' reflects neither morality nor perception. *Trends Cog. Sci.* 19, 631–633
3. Anderson, E. *et al.* (2011) The visual impact of gossip. *Science* 332, 1446–1448
4. Uhlmann, E.L. *et al.* (2015) A person-centered approach to moral judgment. *Pers. Psychol. Sci.* 10, 72–81
5. Callan, M.J. *et al.* (2013) Eye movements to audiovisual scenes reveal expectations of a just world. *J. Exp. Psychol. Gen.* 142, 34–40
6. Ginges, J. *et al.* (2007) Sacred bounds on rational resolution of violent political conflict. *Proc. Natl. Acad. Sci. U.S.A.* 104, 7357–7360
7. Granot, Y. *et al.* (2014) Justice is not blind: visual attention exaggerates effects of group identification on legal punishment. *J. Exp. Psychol. Gen.* 143, 2196–2208
8. Gantman, A.P. and Van Bavel, J.J. (2014) The moral pop-out effect: Enhanced perceptual awareness of morally relevant stimuli. *Cognition* 132, 22–29
9. Firestone, C. and Scholl, B.J. (2015) Enhanced visual awareness for morality and pajamas? Perception vs. memory in 'top-down' effects. *Cognition* 136, 409–416
10. Haidt, J. *et al.* (1993) Affect, culture, and morality, or is it wrong to eat your dog? *J. Pers. Soc. Psychol.* 65, 613–628
11. Rozin, P. (1999) The process of moralization. *Psychol. Sci.* 10, 218–221
12. Firestone, C. and Scholl, B.J. (2015) Cognition does not affect perception: evaluating the evidence for 'top-down' effects. *Behav. Brain Sci.* Published online July 20, 2015. http://dx.doi.org/10.1017/S0140525X15000965

Forum

Neural Antecedents of Spontaneous Voluntary Movement: A New Perspective

Aaron Schurger,[1,2,3,*] Myrto Mylopoulos,[4,5] and David Rosenthal[6]

Fifty years ago, Kornhuber and Deecke first reported their discovery of the *Bereitschaftspotential* [1], or cortical **'readiness potential' (RP)** (see Glossary), a slow build-up of scalp electrical potential preceding the onset of subjectively spontaneous voluntary movements (SVMs). The RP was interpreted as 'the electro-physiological sign of planning, preparation, and initiation of volitional acts' [2], implicitly presumed to reflect the consequence of a **decision** process in the brain. Then, in the early 1980s, Benjamin Libet found that the onset of the RP precedes subjective estimates of the time of the conscious 'urge' to move by 300 ms or more [3] – a result that has since been confirmed at the single-neuron level [4]. This counterintuitive discovery, which we call 'Libet's paradox', led to the view that the conscious decision emerges well after the action has already been initiated unconsciously in the brain, as reflected in the apparent build-up of the RP. While controversy over Libet's findings has churned vigorously for many years, crucial assumptions about the nature of the RP itself have gone unquestioned.

Now a series of new developments has begun to unravel what we thought we knew about the brain activity preceding SVMs.

Glossary

Bounded integration: also known as integration to bound or evidence accumulation, the term refers to a computational model of decision making wherein sensory evidence and internal noise (both in the form of neuronal activity) are integrated over time by one or more decision neurons until a fixed threshold-level firing rate is reached, at which point the animal issues a motor response. In the case of spontaneous self-initiated movement there is no sensory evidence, so the process is dominated by internal noise.
False-positive rate: how often we decide that a movement will occur when in fact it does not.
Neural decision to move: a neural event or state that commits some part of the body to an imminent movement. It is not necessary for this state to be conscious for it to qualify as a 'decision'.
Readiness potential (RP): originally dubbed the *Bereitschaftspotential* by Kornhuber and Deecke [1], the name refers to a slow build-up of scalp electrical potential, measured using (EEG) or electrocorticography (ECoG), preceding the onset of spontaneous self-initiated movements. When measured using magnetoencephalography (MEG) the build-up is referred to as a 'readiness field' and when measured as a change in firing rate in single neurons it is referred to as a 'readiness discharge'.
Time-unlocked forecasting: movement-locked data give us, at each time point $t - \tau$, the probability of the signal at time $t - \tau$ given a movement at $t0$: $p(S_{t-\tau}|M_{t0})$. Time-unlocked forecasting tells us the probability of a movement at time $t + \tau$ in the future given the signal now and in the recent past: $p(M_{t+\tau}|S_t, S_{t-1}, \ldots, S_{t-n})$.
True-positive rate: how often we decide that a movement will occur when in fact it does.

The main new revelation is that the apparent build-up of this activity, up until about 200 ms pre-movement, may reflect the ebb and flow of background neuronal noise, rather than the outcome of a specific neural event corresponding to a 'decision' to initiate movement. In particular, two independent studies, one using electroencephalography (EEG) recordings in humans [5] and the other using single-unit recordings in rats [6], have converged in showing that **bounded-integration** processes, which involve the accumulation of noisy evidence until a decision threshold is reached, offer a coherent and plausible explanation for the apparent pre-movement build-up of neuronal activity. Bounded integration or 'evidence

Look Closer: A Point-by-Point Reply Regarding 'Moral Perception'

(Or: Do We Really Think That Hitting Small Children Is Okay?)

Chaz Firestone & Brian J. Scholl

A primary focus in our laboratory is the study of how perception connects up with the rest of the mind, and in particular how visual processing itself traffics in factors that we typically associate with higher-level cognition — e.g. physical causality (e.g. Scholl & Nakayama, 2002), animacy and intentionality (e.g. Gao et al., 2010), and causal history (e.g. Chen & Scholl, in press). As a result, we were deeply interested in the notion of 'moral perception' — the novel suggestion that "perception is preferentially attuned to moral content" — as introduced in a recent *TICS* forum piece by Gantman and Van Bavel (2015). We were then delighted to engage in a brief dialogue about this notion, as realized in our recent commentary (Firestone & Scholl, 2016) and in their subsequent response (Gantman & Van Bavel, 2016). Reacting to this dialogue, some colleagues asked us what we thought of their response, and so the purpose of this document is just to offer a few brief informal rejoinders.

We thought nearly every one of G&VB's points missed the mark, and that the arguments for 'moral perception' continue to be based on misinterpretations, overinterpretations, and flawed evidence. Here's why:[1]

Our 1st Argument: Much of the trumpeted work doesn't reflect morality.
The search for evidence of 'moral perception' ended up focusing on phenomena that may not be about morality in the first place. In one case, the original study had already explicitly ruled out morality empirically (Sherman et al., 2012, Experiment 3), though G&VB did not mention this. In another case (Anderson et al., 2011), the effect worked with immoral actions (e.g. hitting a child) but also (a) worked with actions that were socially negative but not necessarily *immoral* (e.g. defecating in public, or being in a car accident while borrowing your friend's car), and (b) *failed* to work with morally good actions (e.g. helping the elderly) — suggesting that morality *per se* may not be driving this effect.

G&VB
"We fail to see how hitting a small child does not pertain to morality".

Response: G&VB are missing the point.
Of course hitting a small child is immoral. (Pro tip: If ever you think someone is arguing that hitting small children is okay, consider whether you might be misinterpreting them!) Our point was that *in addition to* those clearly immoral

[1] G&VB made some related arguments when commenting on our forthcoming *BBS* target article, "*Cognition Does Not Influence Perception: Evaluating the Evidence for 'Top-Down' Effects*". As a result, a few of our responses here are cribbed from the reply to our *BBS* commentators. Both the target article and the response are available here:
http://perception.research.yale.edu/preprints/Firestone-Scholl-BBS-Full.pdf

cases, "such effects *also* occurred with mere norm violations that do not invoke morality (e.g. indecent public behavior)" (Firestone & Scholl, 2016, emphasis added); so, for this and other reasons, we don't know whether morality is really responsible. (Indeed, Anderson et al. themselves think this implicates "negative social affect" rather than morality.) In short, G&VB focus too much on hitting small children, and not enough on public defecation.

Additional Response: What happened to Sherman et al.?
In the initial forum article, the very first piece of evidence that was trumpeted as indicating 'moral perception' was the work by Sherman et al. — which G&VB interpreted as reflecting "moral emotions". We pointed out that these authors had explicitly ruled this out — discovering that their results were driven by physical disgust, but *not* moral disgust (something that the *TICS* forum piece failed to mention). G&VB did not respond to this in their reply; perhaps they agree on this point?

Our 2nd Argument: Much of the trumpeted work doesn't reflect perception.
The discovery that people attend to just outcomes in moral scenarios (Callan et al., 2013) tells us something about what we expect in moral scenarios, but it doesn't tell us anything new about how *attention itself* works. We already knew that attention is drawn to what people expect.

G&VB
"Isn't that interesting?"

Response: Sure, it's interesting, but it still doesn't support 'moral perception'.
These findings are indeed interesting for certain purposes (which is why our *TICS* comment explicitly described them as "interesting"!). But they're interesting in revealing *what our moral expectations are*, not in revealing new ways that attention works.

G&VB
"The goal of cognitive science is building process-oriented models".

Response: Exactly!
This is exactly what motivated our studies. The initial allegations of 'moral popout' simply reported the effect and gave it a shiny new name — without any consideration of just *how* such effects could instead be produced by the underlying mechanics of familiar cognitive processes (such as memory). Our entire point is that when you consider such "process-oriented models", you realize that this exciting new phenomenon is really just semantic priming:

Our 3rd Argument: 'Moral popout' doesn't exist; it's just semantic priming.
Moral words (e.g. "crime", "punishment") are easier to see than random non-moral words (e.g., "steel", "ownership") simply because moral words are semantically related (whereas random non-moral words are not), and related words prime each other. As a result, words from *any* arbitrary category —

including 'silly' categories such as fashion — will "pop out" when compared to unrelated words (Firestone & Scholl, 2015). There is no evidence that morality plays a role.

G&VB
"They argued that moral pop-out can be fully explained by semantic priming because 'relatedness is the key factor in such effects'".

Response: No, we argued this on the basis of our *experiments*.
The relatedness confound alone doesn't mean that moral popout must be fully explained by semantic priming; it merely means that semantic priming *could* be responsible, such that Gantman and Van Bavel's (2014) original conclusions wouldn't follow. (Note that G&VB never mentioned the 'relatedness' confound in their initial studies.) Nevertheless, we actively tested this alternative, and semantic relatedness *in fact* produced analogous "popout" effects. Our experimental results — not just the confound itself — are what suggest that 'moral popout' is really just semantic priming.

G&VB
Moral and non-moral words were "matched for length and frequency".

Response: No: they were confounded in a self-serving way.
The moral words used by Gantman and Van Bavel (2014) were 7% shorter and 25% more frequent than the non-moral words were — two self-serving biases that have never been pointed out. (We're not sure where the line should be drawn between "matched" and "confounded in an unacknowledged and self-serving way" — but surely it's at least an order of magnitude smaller than 25%.) Our experiments (Firestone & Scholl, 2015) carefully controlled for these factors and are thus the only interpretable published data on "moral popout".

G&VB
"The words were presented near the threshold for awareness".

Response: We can't be sure, since the reported duration was impossible.
Nobody knows how long the words were presented onscreen in G&VB's (2014) experiments, since they reported a critical presentation duration of 40ms on a 60Hz monitor. This is physically impossible: such a monitor can only present words for multiples of 1/60 of a second — e.g. 17ms, 33ms, 50ms, etc. (For this reason, our replication used a 50ms duration.)

G&VB
"They [F&S] did not randomly assign participants to detect moral versus fashion and/or transportation words."

Response: Look closer: Just how is this challenge supposed to work?
It is true that subjects were not "randomly assigned" to which experiment they participated in: first we discovered a silly 'fashion popout' effect, then we

generalized the effect to another arbitrary category (transportation), and then (in response to comments from reviewers, including G&VB) we replicated 'moral popout' too. So what is the challenge here? G&VB don't elaborate, so let's look closer. If we were claiming *differences* between the experiments, then non-random assignment could be problematic in a straightforward way (e.g., perhaps one group was more tired or stressed out). But we suggested that there is *no* evidence of any relevant differences among these popout effects, and our explanation is that the same underlying process — semantic priming — drives all the effects (with no evidence that morality plays any role). Can a lack of random assignment explain this apparent equivalence differently? Such an explanation would have to assume that the 'true' effect with morality in our experiments was much larger than for the other categories (due to the morality-specific boost) but that these particular subjects (i.e. Yale community members tested in one month rather than another) somehow deflated this previously undiscovered 'super-popout' down to … exactly the same magnitude (of 4%) previously reported by Gantman and Van Bavel (2014)! Random assignment is a red herring here; it cannot save the day for 'moral popout'.

G&VB
"Curiously, however, Firestone and Scholl did not report the analogous means for their morality study…"

Response: Look closer: The means support our account.
Let's explore this in detail. Our claim is that words from a given category (whether fashion, transportation, or morality) prime each other and enhance lexical decision performance. The critical evidence for this is simply that words from such categories show a "popout" effect: the mere fact that fashion and transportation pop out too is why we think there's no evidence that morality *per se* plays any role. But, in addition to this, we can also look for evidence of semantic priming at the trial-by-trial level, to see whether category words seen on one trial enhance performance for category words on the next trial (even though our experiments were not designed for this analysis). To test this — say, for fashion — we first need to ask whether fashion words preceded by fashion words (e.g. "blouse" preceded by "dress") were recognized more accurately than fashion words preceded by non-fashion words (e.g. "blouse" preceded by "limit"). In fact, they were, by 4.9%. However, this number is meaningless on its own: maybe fashion words are exciting or arousing (after all, they included the word "bikini"!) and would increase performance for *any* subsequent word (in which case the true priming effect could be nil); or, maybe fashion words are distracting and would, all else equal, *impair* performance on subsequent words (in which case the true repetition priming effect could be even larger than 4.9%). So, for this to be interpretable at all, we also need to compare performance on *non-fashion* words preceded by fashion words (e.g., "limit" preceded by "blouse") to *non-fashion* words preceded by non-fashion words (e.g., "limit" preceded by "diesel"). And indeed, fashion words in this case confer a 3.9% *disadvantage* for non-fashion-word recognition accuracy. So, fashion words (rather than non-fashion words) boost recognition of fashion words (rather than non-fashion

words) by [(4.9%) – (-3.9%)] = **8.8%**. For transportation, that value is [(5.7%) – (-0.2%)] = **5.9%**. And for morality, that value is [(1.1%) – (-3.4%)] = 4.5%. And we also reported a valenced-matched version of this morality analysis, which produces [(2.5%) – (-2.7%)] = **5.2%**. These values did not differ significantly — though, as might not be surprising, these three values of 8.8% (fashion), 5.9% (transportation), and 5.2% (morality) are rank-ordered just as are the magnitudes of the 'popout' effects themselves: 8.6% (fashion), 4.3% (transportation), and 3.9% (morality). These are the means G&VB asked for, and they plainly support our account.[2] In sum, looking closely and carefully at the data supports the claim that there is no qualitative difference between 'moral popout' and 'fashion popout', and that both can be entirely explained by semantic priming.

G&VB
"…despite the fact that it was central to their hypothesis."

Response: This analysis was never "central" (though it does support us).
Our hypothesis is that semantic priming explains "popout" effects. To demonstrate this, we showed that "popout" effects emerge for arbitrary categories in the exact same way as in Gantman and Van Bavel (2014). Whether these effects also manifest at the trial-by-level is a completely separate matter that is not at all required by our hypothesis. Nevertheless, discovering such trial-level effects is *further evidence* for our hypothesis (as reported in detail above), even though our study was not designed to test for such effects. We have reported all of these means, which continue to suggest that "moral popout" is really just semantic priming.

G&VB
"[T]he burden of proof is on Firestone and Scholl to design and fully report a study that demonstrates how fashion and/or transportation pop-out is 'entirely analogous' to moral pop-out."

Response: This seems backward, and counter to general scientific practices.
G&VB are claiming that there is a revolutionary new phenomenon of "moral popout". We noted a stark and unrecognized confound in those studies — implicating an alternative factor (semantic priming) that has been well understood for decades — and then showed empirically that this confound alone produces similar effects. Surely this means that the burden of proof is now on G&VB to show that there is more to "moral popout" than semantic priming!

[2] You may notice that, relative to fashion and transportation, moral words preceded by moral words weren't recognized much more accurately than moral words preceded by non-moral words (either 1.1% or 2.5% [depending on whether valence is taken into account] compared to 4.9% for fashion, and 5.7% for transportation [though these differences are not statistically significant]). But as with fashion, *this value is entirely uninterpretable on its own*: perhaps reading morally relevant words like "murder" and "violence" is startling and impairs performance on the next trial in a way that counteracts the moral-moral repetition priming advantage. That's why we also need to look at the effect of moral words on subsequently presented *non-moral* words — and indeed we see that non-moral words are recognized less accurately when preceded by moral words than by non-moral words (-3.4% or -2.7%).

(This is the same reason why the burden of proof is generally held to be on proponents of alternative medicines to produce compelling evidence for their claims — and not on their critics to design and fully report studies that refute them.)

G&VB
"Perhaps the only thing we can agree on is that more research is needed."

Response: Amen!
This was the entire motivation for our studies. We think more — and more careful — research is needed before claiming that there is 'moral popout'; before telling millions of lay readers that "moral concerns can alter a person's perceptual experience" (http://nyti.ms/1efPqcZ); and before claiming that "moral perception" is a trend in cognitive science that has "important consequences for policy". To proclaim such lofty conclusions in the absence of this evidence seems premature, and even a form of uncareful overreach. In any case, we continue to find 'moral perception' to be a fascinating (if unlikely) possibility, and we hope that further research on the underlying mechanics of these effects will help the empirical basis for this idea to catch up to the bold pronouncements that are made in its name.

References

Anderson, E., Siegel, E. H., Bliss-Moreau, E., & Barrett, L. F. (2011). The visual impact of gossip. *Science*, *332*, 1446–1448.

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.

Chen, Y. -C., & Scholl, B. J. (in press). The perception of history: Seeing causal history in static shapes induces illusory motion perception. *Psychological Science.*

Firestone, C., & Scholl, B. J. (in press). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and Brain Sciences.*

Firestone, C., & Scholl, B. J. (2015). Enhanced visual awareness for morality and pajamas? Perception vs. memory in "top-down" effects. *Cognition*, *136*, 409–416.

Firestone, C., & Scholl, B. J. (2016). 'Moral perception' reflects neither morality nor perception. *Trends in Cognitive Sciences*, *20*, 75–76.

Gantman, A. P., & Van Bavel, J. J. (2014). The moral pop-out effect: Enhanced perceptual awareness of morally relevant stimuli. *Cognition*, *132*, 22–29.

Gantman, A. P., & Van Bavel, J. J. (2015). Moral perception. *Trends in Cognitive Sciences*, *19*, 631–633.

Gantman, A. P., & Van Bavel, J. J. (2016). See for yourself: Perception is attuned to morality. *Trends in Cognitive Sciences*, *20*, 76–77.

Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science*, *21*, 1845–1853.

Scholl, B. J., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science*, *13*, 493–498.

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.