



Can coherence-based interventions change dogged moral beliefs about meat-eating?[☆]

Zachary Horne^{a, *}, Joshua Rottman^b, Caroline Lawrence^c

^a Department of Psychology, University of Edinburgh, UK

^b Department of Psychology, Franklin & Marshall College, USA

^c Yale Law School, Yale University, USA

ARTICLE INFO

Keywords:

Coherence-based reasoning
Morality
Belief change
Animal ethics

ABSTRACT

What causes people to change their beliefs about right and wrong? Coherence-based interventions can change people's moral beliefs about abstract moral principles (Holyoak & Powell, 2016), but it is unclear whether these interventions would be similarly effective for everyday moral beliefs that can impact routine behavior. In the present research, we examined whether coherence-based "memes" highlighting the moral similarities of pigs and dogs can shift moral beliefs about consuming meat. Across three preregistered experiments ($N = 2281$), we found that self-reported beliefs about the permissibility of eating some animals can be subtly shifted by brief coherence-based interventions which highlight morally relevant capacities (e.g., intelligence, emotional capacities) of an animal that is frequently eaten in Western society (pig) and an animal that is typically considered forbidden to eat by Westerners (dog). We discuss the implications of these findings for psychological and ethical theory.

1. Introduction

People do not always have clear reasons for their moral commitments (e.g. Haidt, 2001; Kelly, 2011), but they are nonetheless steadfast in maintaining them (Mullen & Skitka, 2006; Skitka, 2010). Because many moral commitments may be unprincipled and culturally driven, they can appear to be immune to reasoned argument (Lord, Ross, & Lepper, 1979; Miske, Schweitzer, & Horne, 2019; Schwitzgebel & Cushman, 2015; but see May, 2018). Still, people can change their minds—even about morality—thus giving up core commitments. Voters become disenchanted with their leaders and leave their political parties, omnivores become vegetarians, and atheists become religious. All of these changes can entail taking on an entirely new set of moral commitments. What then causes people to change their moral beliefs (e.g. Kohlberg & Kramer, 1969), and what mechanisms underlie those changes?

In a recent line of research, psychologists and philosophers have examined how moral beliefs change when they are indirectly countered (e.g., Holyoak & Powell, 2016; Horne, Powell, & Hummel, 2015; Horne, Powell, & Spino, 2014; May, 2018). In philosophy and the law, theories, explanations, and even individual beliefs are evaluated by how they

cohere with other beliefs (e.g., Dancy, 1984). Coherence has been argued to be particularly important in the domain of ethics because empirical observations are unlikely to tell us if, for example, utilitarianism is a defensible ethical theory, or whether it is morally wrong to eat meat (e.g., Campbell & Kumar, 2012; Horne et al., 2015; May, 2018; Norcross, 2004). In line with the rhetorical strategies typical of philosophy and the law, there is emerging evidence that coherence-based arguments can change moral beliefs by way of *coherence-shifts* (e.g., Festinger, 1962; Holyoak & Powell, 2016; Horne et al., 2015, 2014). To consider one example, when people are presented with a situation (e.g., a moral dilemma) that elicits a case-based judgment inconsistent with a structurally similar moral principle they otherwise accept (e.g., the belief that morally right actions are those that maximize impersonal happiness), a tension arises due to an internal conflict among their beliefs about the dilemma and this general moral principle. This tension can induce belief revision because people have an implicit or explicit desire to maintain coherence in their network of beliefs (e.g., Festinger, 1962; Holyoak & Powell, 2016; Petty, Wheeler, & Tormala, 2003). These data suggest at least some moral beliefs can be indirectly countered by way of coherence-based interventions that manipulate auxiliary beliefs.

Although a promising tactic for persuasion, recent work examining

[☆] This paper has been recommended for acceptance by Dr. Rachel Barkan.

* Corresponding author.

E-mail address: zhorne@ed.ac.uk (Z. Horne).

the effects of coherence-based interventions on moral beliefs has focused on shifting beliefs about (comparatively) abstract moral principles (see Holyoak & Powell, 2016, for a review). This is notable because people may agree with unfamiliar abstract principles at first blush without being committed to the principle or assured in their agreement. Indeed, Horne et al. (2015) suggest that people do not realize that adopting a principle like maximize impersonal happiness (arguably) entails certain behaviors that people might find supererogatory (e.g., giving one-third of their money to charity) (Singer, 1972) or objectionable (e.g., removing the organs of a single healthy person to save five sick patients) (Norcross et al., 2008). According to Horne and colleagues, this fact explains why providing a single counterexample can induce belief revision. One possibility, then, is that coherence-based interventions are effective at shifting moral beliefs in these and similar studies because, although people may *initially* find an ethical principle plausible, they might not be particularly *confident* in that assessment.

This concern highlights an open question in the literature on coherence-based reasoning: to date, there is little evidence that coherence-based interventions affect concrete moral beliefs – those which are sufficiently specific that commitment to them is transparent, and thus when people endorse them, they are confident in that belief. For example, if a person thinks it is wrong to eat pork, it is clear they believe that they ought not eat pork. In contrast, if a person initially thinks they should maximize impersonal happiness, it is not clear what exactly they ought to do. Consider the fact that even ethicists who adopt this moral principle do not know exactly what it commits them to (Ridge & McKeever, 2020).

Furthermore, psychologists studying attitude change have historically focused on whether attitudes about issues like abortion (e.g., Huckfeldt & Sprague, 2000), capital punishment (e.g. Lord et al., 1979), or animal ethics can be changed (e.g. Loughnan, Bastian, & Haslam, 2014). Attitudes about these issues are substantially more concrete than sweeping ethical principles. And in this literature researchers have often found that shifting ethical attitudes is difficult, and can even lead to polarization (Lord et al., 1979). These points lead us to conclude that a strong test of coherence-based arguments typical of ethics and the law requires studying whether these interventions shift beliefs in cases in which (1) the implications of holding the belief are unambiguous and (2) people are confident in their beliefs. To this end, we investigated how coherence impacts perceptions of the ethics of eating meat.

Although most people strongly believe that it is permissible to eat meat, like many other moral norms they adopt, beliefs about the moral permissibility of eating meat are often contradictory and seemingly inconsistent. While only 5% of the United States population is vegetarian (Newport, 2012), the majority think it is impermissible to eat certain kinds of meat (for example, dogs, cats, and horses). Certain meats are tabooed in every society, but the details of these taboos vary across cultures (Fessler, Arguello, Mekdara, & Macias, 2003). Ethicists have also argued that the coherence of people's beliefs about eating meat is normatively suspect (Norcross, 2004; Singer, 1972). For example, pigs are more intelligent than cats, but Westerners often think that it is wrong to eat cats and acceptable to eat pigs. Even within Western culture, there is considerable variability. For instance, the French eat horses, but this seems impermissible to most Americans (Harris, 1998). Within the United States, there is also substantial diversity, for instance with regard to beliefs about the ethics of eating in vitro meat (Wilks & Phillips, 2017). These observations suggest that, rather than stemming from reasoned or principled considerations, beliefs about the permissibility of eating meat are predominantly rooted in

biological, cultural, historical, and ecological contingencies that have produced an (often) disjointed set of moral norms (Harris, 1998).¹ Whether these contingencies are in fact indefensible, the *prima facie* incoherence of people's moral norms provides a testbed for examining the efficacy of coherence-based interventions to reveal the psychological mechanisms underlying moral belief change.

To summarize, people have a diverse set of moral beliefs about meat consumption and the norms they adopt are often arbitrary and internally inconsistent (again see Norcross, 2004, for a review). This set of observations led us to examine whether coherence-based interventions can elicit belief change about the moral permissibility of eating some animals (i.e., pigs) but not others (i.e., dogs). Coherence-based models of moral reasoning (e.g., Campbell & Kumar, 2012; Holyoak & Powell, 2016; Horne et al., 2015; May, 2018) predict that the desire to maintain a coherent set of moral beliefs would lead people to revise their beliefs about eating animals for which there is no taboo in some corners of Western society (e.g., pigs). Specifically, these accounts predict that increasing the perceived moral similarity between dogs and pigs will cause people to judge that eating pigs (and perhaps even other similar animals) is morally wrong.² This finding would be consistent with prior work demonstrating that coherence-based arguments can influence agreement with some abstract moral principles (Holyoak & Powell, 2016; Horne et al., 2015), and, more broadly, that rule-based coherence theories of moral reasoning can explain an important aspect of everyday moral reasoning (e.g., Campbell & Kumar, 2012; Holyoak & Powell, 2016).

1.1. The present experiments

In the present experiments, we sought to provide a strong test of the hypothesis that coherence-based interventions can shift moral beliefs. To this end, we created minimal, naturalistic interventions aimed at highlighting the similarities between dogs and pigs. The logic of this design was that if a very brief, naturalistic intervention can shift moral beliefs that routinely affect people's everyday behavior—even if only slightly—then this would provide compelling evidence for the coherence theories detailed above. In this way, the experiments we present follow a logic of discovery recommended by Popper in which we consider a bold prediction made by coherence theories and determine if even in these cases we see evidence which corroborate the theories in question (Popper, 2005). The implications of following this approach are twofold: First, we did not expect the magnitude of the minimal, naturalistic interventions we used to be large. Second, even if the effects we observed were comparatively small, this would not indicate that stronger coherence interventions would not yield larger effects.

With this experimental approach in mind, in three preregistered experiments, we presented participants with brief memes resembling those frequently seen online and on social media. These memes were intended to highlight the similarities between two animals for which Western society has different moral norms (i.e., dogs and pigs), while simultaneously maintaining both tight experimental control and high ecological validity. In Experiment 1, we presented participants with three memes comparing pigs with dogs, highlighting either morally relevant similarities or less morally valenced forms of anatomical or physiological (i.e., somatic) similarity. To maintain control in both

¹ To clarify, the aim of this work is to test an *empirical hypothesis* about how people respond to common coherence-based interventions. This should be separated from what people ought to do—the normative question of whether it is morally permissible to eat meat is well beyond the scope of this paper (for a recent review see Gruen, 2017).

² Strictly speaking, on the coherence-accounts outlined above, it is possible that increasing the perceived similarity between dogs and pigs would make people more likely to think it's permissible to eat dogs, a point we discuss below.

conditions, these verbal comparisons were paired with graphic images clearly intended to persuade participants of the horrific nature of the practices that make it possible to eat pigs or dogs, and thus of the immorality of these practices. In Experiment 2, we presented participants with the same memes but included additional measures aimed at measuring participants' commitments to eating meat and the similarity they perceived between dogs and several other animals (described in detail below). In Experiment 3, we also compared pigs and dogs, but removed mention of how these animals are differentially consumed. We also changed the pictures so both memes displayed cute images of the animals, and measured different ways in which dogs and pigs are considered to be similar. All three experiments sought to address the question of whether brief coherence-based interventions highlighting moral similarities could affect moral beliefs likely to impact everyday behavior.

We note that the present approach differs from other methods for investigating attitude change about animal ethics, which have typically solely focused on emotional dispositions. While the moralization of specific kinds of meat is associated with feelings of disgust (Rozin, Markwith, & Stoess, 1997), disgust may be a consequence rather than a cause of changing these moral convictions (Fessler et al., 2003), suggesting that emotion alone may not be sufficient for shifting moral beliefs toward meat (but see Feinberg, Kovacheff, Teper, & Inbar, 2019). Nevertheless, because moral change may occur most readily when both hot and cold cognition are involved (Nichols, 2004), we supplemented our arguments with both emotionally evocative photographs (Experiments 1 and 2) and pleasant photographs (Experiment 3) (see Slovic, 2007).

All studies were preregistered on the Open Science Framework, and we report all measures, manipulations, data exclusions, and sample size determinations. Exploratory analyses are located in Appendices A and B.

2. Analytic strategy

Throughout, we perform Bayesian estimation using the R package brms (Bürkner, 2017). Bayesian statistics are becoming common in psychological science for several reasons. First, Bayesian statistics enable researchers to draw inferences about the parameters of interest in light of the data they have collected, which is the typical aim of scientific research (Kruschke, 2014). Second, Bayesian statistics allow researchers to make easily interpretable, intuitive statements about the probability that an estimated parameter falls in a certain interval given the observed data (Morey, Hoekstra, Rouder, Lee, & Wagenmakers, 2016). Throughout, we report Bayesian 95% posterior uncertainty intervals for all parameters of interest. Third, Bayesian posterior uncertainty intervals are useful because they allow inferences about the relative credibility of the values within them: Values close to their center are more credible than values close to their limits. As a result, whether a posterior uncertainty interval for a parameter estimate just crosses vs. just avoids crossing 0 is immaterial. In either case, 0 may be a low-credibility value of the parameter because it is in the tail of its posterior (Kruschke, 2014). For readers who are unfamiliar with Bayesian statistics, we provide an introduction to key concepts, as well as a detailed explanation of our priors, in Appendix D.

From an estimation perspective, Bayesian statistics also provide a principled way to regularize parameter estimates (namely, by changing the prior distribution). Regularizing priors reduce the likelihood of model overfitting (McElreath, 2016). Throughout, we set regularizing priors on the parameters in our models, which we detail in the Supplemental Online Material (SOM). However, less conservative (non-regularizing) prior choices, or parameter estimates using Maximum Likelihood Estimation, often yielded similar parameter estimates – see the Open Science Framework page for Experiment 3 (<https://osf.io/wgczf/>).

We are unaware of any widely-accepted methods for calculating statistical power in mixed-effects Bayesian cumulative regression

models, which we use throughout the paper, beyond using prior predictive checking and data simulation. The motivation for our prior selections and corresponding prior predictive checks can be found in Appendix D. Still, to aid the reader in interpreting our data and the statistical power of our studies, we also report frequentist sensitivity analyses throughout even though they provide only a very coarse-grained approximation of the precision our sample sizes are able to provide.

3. Experiment 1

3.1. Method

3.1.1. Participants

Participants were 386 Amazon Mechanical Turk workers ($M_{age} = 35$, 65% female) who were compensated \$0.40 for their participation, or approximately \$7.00 an hour. An additional 14 participants were tested but were excluded according to our preregistered analysis plan. We required that participants be located in the United States to accept the study HIT. Our sample size and exclusion criteria were decided in accordance with our preregistration: <https://osf.io/uufuru/>.

A sensitivity power analysis for an independent samples *t*-test indicated we could detect an effect as small as a Cohen's *d* of 0.28 with 80% power, approximately a log odds ratio of 0.51 using the equation $\log \text{odds ratio} = \frac{d\pi}{\sqrt{3}}$ (Sánchez-Meca, Marín-Martínez, & Chacón-Moscoso, 2003).

We note that our exclusions diverged from the preregistration in one respect: Rather than excluding current or former vegetarians, vegans, or those who keep kosher, we directly modeled these factors in a series of exploratory analyses (see the Exploratory Analysis Models section of the Open Science Framework).

3.1.2. Materials and procedure

We presented participants with one of two sets of memes in a between-subjects design. Interventions were structured as memes in order to maintain high ecological validity; memes like those shown in Fig. 1 are typical of, for example, activist posts on many social media sites (e.g., <https://twitter.com/peta>, Twitter, 2021). These memes displayed pigs and dogs in nets, on spits, or in slaughterhouses, along with statements comparing them along either morally relevant or somatic dimensions.

The pictures in both conditions were identical, but we manipulated the type of the similarity that was highlighted. For instance, in the Somatic Similarity condition, participants read that pigs have better smell than sight, just like dogs, whereas in the Moral Similarity condition, they read that pigs bond with their farmers, just like dogs with their owners. The ordering of the pictures and text was randomized and counter-balanced. Example stimuli are shown below in Fig. 1 and the text in each meme is shown below in Table 1. All other materials for Experiment 1 can be found in the SOM.

The stimuli in these conditions were equated on a number of dimensions to eliminate deflationary explanations of experimental effects. For example, the inclusion of identical, emotionally evocative photographs in each condition prevented emotion induction from being a likely explanation of any condition differences. Most importantly, there were no differences across conditions in how demand characteristics could shape participants' responses. Even though one aim of the study was not masked—to change participants' beliefs—the presence of the same photographs across both conditions ensured participants would think *this hypothesis* was what we sought to test in both conditions, when we were instead examining the differential impact of comparing moral and somatic similarity. However, one limitation of this design is that the intercepts (more accurately, thresholds – see Appendix D) in our statistical models do not represent a truly neutral reference group and thus the model thresholds need to be interpreted with considerable caution.



Fig. 1. Sample stimuli from Experiments 1 and 2. Each participant saw three memes from either the Somatic Similarity (left panel) or Moral Similarity (right panel) condition in a between-subjects design.

Table 1

Statements highlighting the similarities between pigs and dogs in three memes in the Somatic Similarity and Moral Similarity conditions in Experiments 1 and 2.

Somatic Similarity	Moral Similarity
Pigs have sweat glands, just like dogs...	Pigs have 200 million more neurons than dogs...
Pigs have better smell than sight, just like dogs...	Pigs experience fear and sorrow, just like dogs...
Pigs have canine teeth, just like dogs...	Pigs bond with farmers, just like dogs...

After viewing each set of three memes, participants advanced to the next screen where they answered four questions about the morality of eating meat in general (General Meat Consumption Scale) and four questions about the morality of eating specific animals: chickens, pigs, cows, and dogs (Specific Meat Consumption Scale). An example of an item in the General Meat Consumption Scale is "Eating meat is morally acceptable" and an example of an item in the Specific Meat Consumption Scale is "Eating chicken is morally wrong." The order in which these scales were presented was counterbalanced, and the order of the items within these scales was randomized. The complete set of scale items can be found in Table 3 of the SOM.

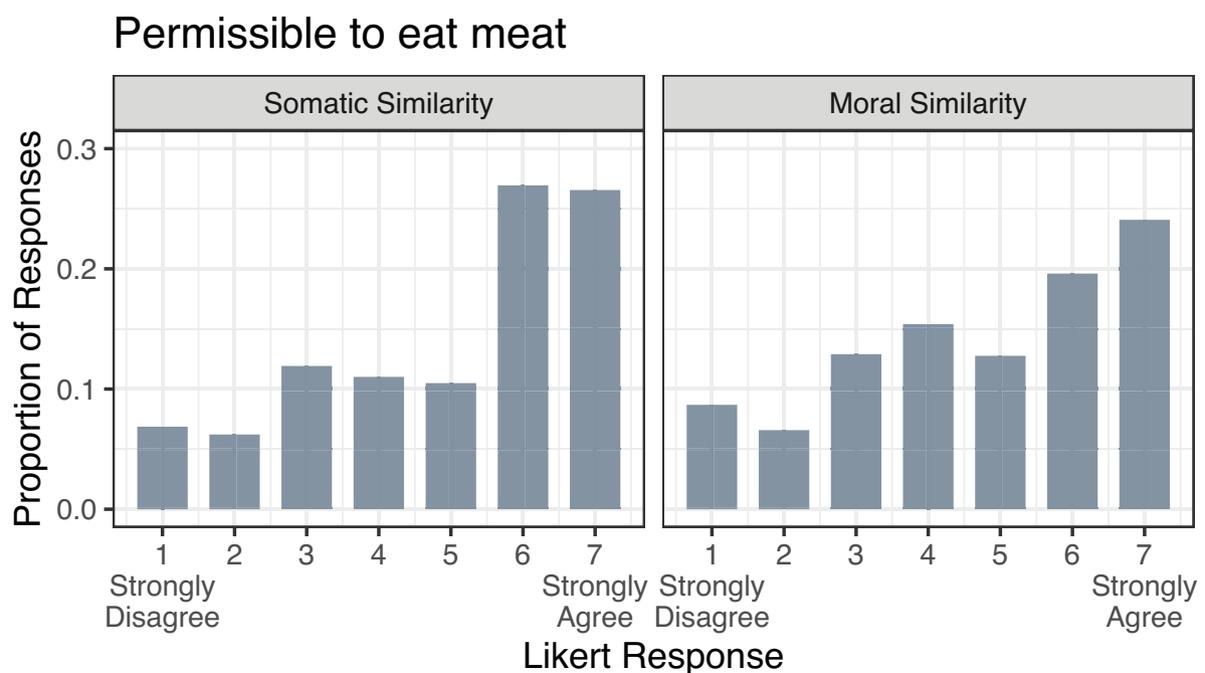


Fig. 2. Beliefs about how wrong it is to eat meat in general in Experiment 1. Higher Likert scale choices indicate more favorable attitudes toward eating meat.

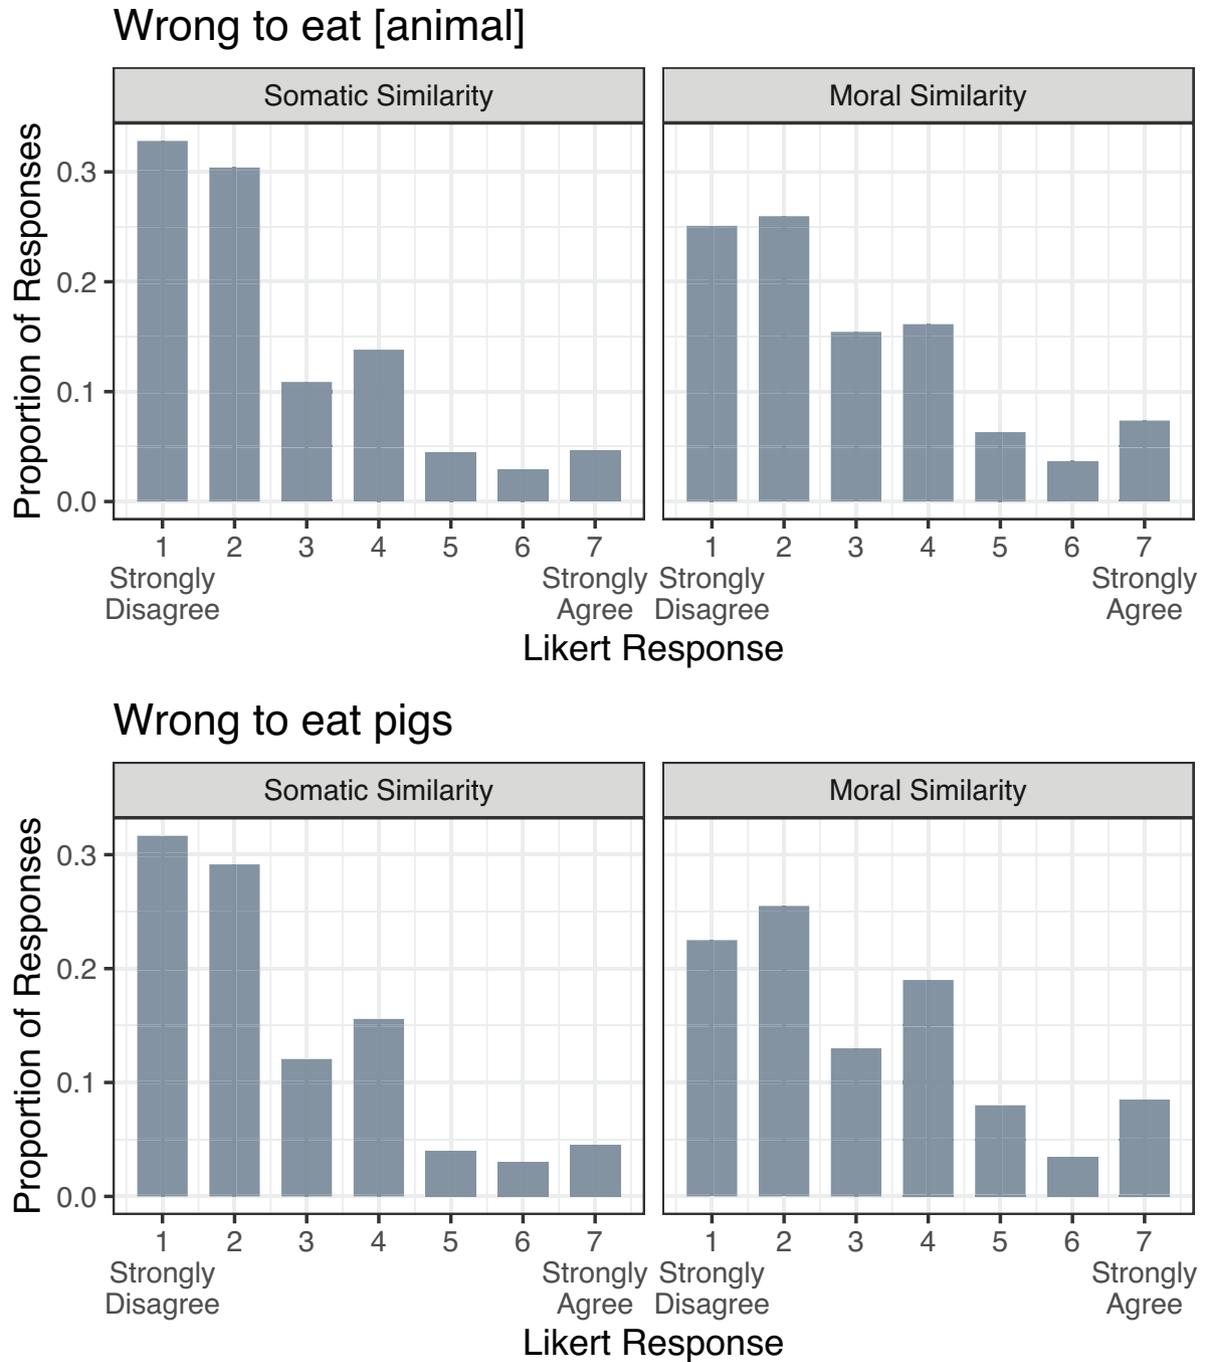


Fig. 3. Beliefs about the wrongfulness of eating meat from specific animals (top - i.e. cows, chicken, pigs) and pigs (bottom) in Experiment 1. Higher Likert scale choices indicate less favorable attitudes about eating specific animals.

3.2. Results

3.2.1. Preregistered analyses

We tested whether participants in the Moral Similarity condition were more likely than those in the Somatic Similarity condition to think that it was wrong to eat meat in general (i.e., the General Meat Consumption Scale), specific animals (i.e., the Specific Meat Consumption Scale), and pigs in particular—that is, the only animal compared to dogs in both conditions. To test these predictions, we fit two models to estimate the effect of highlighting moral similarity on beliefs about the ethics of eating meat.

Following the recommendations of Liddell and Kruschke (2018), Likert data were modeled with a cumulative probability distribution. The cumulative distribution is recommended for Likert scale data because it assumes that ordered nominal responses represent a continuous latent construct (in this case, strength of a particular moral belief). These models, rather than linear models, are particularly necessary when Likert responses are highly skewed, as we expected our data would be about eating meat.

We first tested how increasing the perception of moral similarity between dogs and pigs shifted beliefs about eating meat in general. To this end, we performed Bayesian cumulative mixed-effects modeling using the R package *brms* (Bürkner, 2017). Our first analysis regressed general meat consumption judgments on Condition (0 = Somatic Similarity, 1 = Moral Similarity), allowing for the effect of Condition to vary for each question in the General Meat Consumption scale (i.e., the maximal model given our design; see Barr, Levy, Scheepers, & Tily, 2013). Details about the priors, convergence behavior, and posterior density plots can be found in Appendix D and the SOM.

This analysis suggests the coherence-based intervention only weakly shifted beliefs in the wrongfulness of eating meat in general (see Fig. 2), $b_{\text{Condition}} = -0.45$, 95% Posterior Uncertainty Interval [-1.15, 0.23]. This is perhaps unsurprising because the intervention targeted attitudes about eating pigs specifically, rather than eating meat in general. Thus, it is still possible the intervention affected beliefs about eating specific animals people viewed as similar to dogs. We examined how the intervention changed beliefs about eating specific animals, predicting judgments about the ethics of eating specific animals (e.g., pigs) on condition, a predictor indicating the animal in question (i.e., Specific Animal), and included a random intercept for participant. We specify the priors and R-syntax for this model in the SOM.

Although the coherence-based intervention did not materially affect moral beliefs about eating meat in general, it had a stronger impact on beliefs about eating certain animals (see Fig. 3 – top panel), with the largest effect on beliefs about the ethics of eating pigs, $b = 1.23$, 95% Posterior Uncertainty Interval [0.12, 2.30]; see Fig. 3 – bottom panel. This analysis provides initial, if tentative, evidence that even brief coherence-based interventions highlighting the moral similarity between dogs and pigs can impact people's moral beliefs.

Although this model provides some evidence for the effect we predicted, the amount of uncertainty in the estimated effect of Condition is quite large, spanning a very small condition effect (e.g., Odds-Ratio = 1.12) to very large effects (e.g., Odds-Ratio = 9.97) (Richard, Bond, & Stokes-Zoota, 2003). The uncertainty in this parameter estimate suggests that the estimated effect is inflated (e.g., Gelman & Carlin, 2014) and thus should be interpreted with considerable caution. Consistent with this possibility, the posterior of the condition effect was substantially impacted by more informative, regularizing priors i.e. $N(0.00, 0.50)$, $b = 0.78$, 95% Posterior Uncertainty Interval [-0.09 to 1.63].³ These results motivated a direct replication and extension of Experiment 1.

³ For more information about regularizing priors on β parameters, see the priors on β parameters section of Appendix D.

4. Experiment 2

4.1. Method

4.1.1. Participants

Participants in Experiment 2 were 745 Amazon Mechanical Turk workers ($M_{\text{age}} = 35$, 51% female), who were paid \$0.60 for participating in the study, or approximately \$6.00 an hour. An additional 158 were tested but excluded before analyzing the data for failing attention checks. We required that participants be located in the United States to accept the study HIT. The sample size and exclusion rules were determined prior to data collection according to our preregistration: <https://osf.io/zmr4s/>. The number of exclusions due to inattention in this study were higher than we anticipated, an issue we discuss in the results section below. Even with a smaller sample size than anticipated, a sensitivity power analysis for an independent samples *t*-test indicated we could detect an effect as small as a Cohen's *d* of 0.20 (log odds ratio = 0.36) with 80% power, the mean effect size in social psychology (Richard et al., 2003).

4.1.2. Materials and procedure

We used the same coherence-based memes in Experiment 2, and participants once again received either the Moral Similarity or Somatic Similarity coherence intervention in a between-subjects design. We made several other changes to our design, primarily aimed at decreasing sources of variability in our measurement. Specifically, we added several measures in order to better detect potential changes in moral beliefs about eating meat, including the revised Specific Meat Consumption Scale and a similarity judgment measure, which were administered at the end of the study. We describe each of these measures below.

4.1.3. Additional pretest and posttest measures

Prior to and after receiving the coherence-based intervention, participants rated their agreement on a seven question Meat Commitment Scale (Piazza et al., 2015) (see Table 5 in the SOM). The scale included items like "I don't want to eat meals without meat." After responding to items on this scale, participants rated their agreement with statements on a revised *General Meat Consumption Scale* (see Table 4 in the SOM), which again assessed beliefs about eating meat in general. For example, participants rated their agreement with the statement, "There is nothing morally wrong with eating meat." At the end of the study, participants completed a revised *Specific Meat Consumption Scale* (Table 4 in the SOM). The phrasing of the items on this scale were identical to Experiment 1, but the scale included additional questions about other animals, including lamb and fish – analyses on these additional questions were exploratory and are located in the Supplementary Online Materials. After completing these measures, participants were asked to judge how similar each animal on the Specific Meat Consumption Scale was to a dog. For example, participants were asked to rate how similar pigs are to dogs and fish are to dogs. This similarity measure was included for two reasons: First, we hypothesized that participants would be less likely to find it permissible to eat an animal they thought was similar to a dog. Second, we reasoned that, if the Moral Similarity intervention highlighted the more relevant similarity between dogs and pigs, then people should think that pigs are more similar to dogs relative to the Somatic Similarity condition. The details of these analyses are located in Appendix A.

4.2. Results

We tested whether participants in the Moral Similarity condition would think that it was more wrong to eat meat in general (i.e., the General Meat Consumption Scale) and that it was more wrong to eat specific animals (i.e., the Specific Meat Consumption Scale). However, in Experiment 2 we sought to reduce variance owing to pretest beliefs by

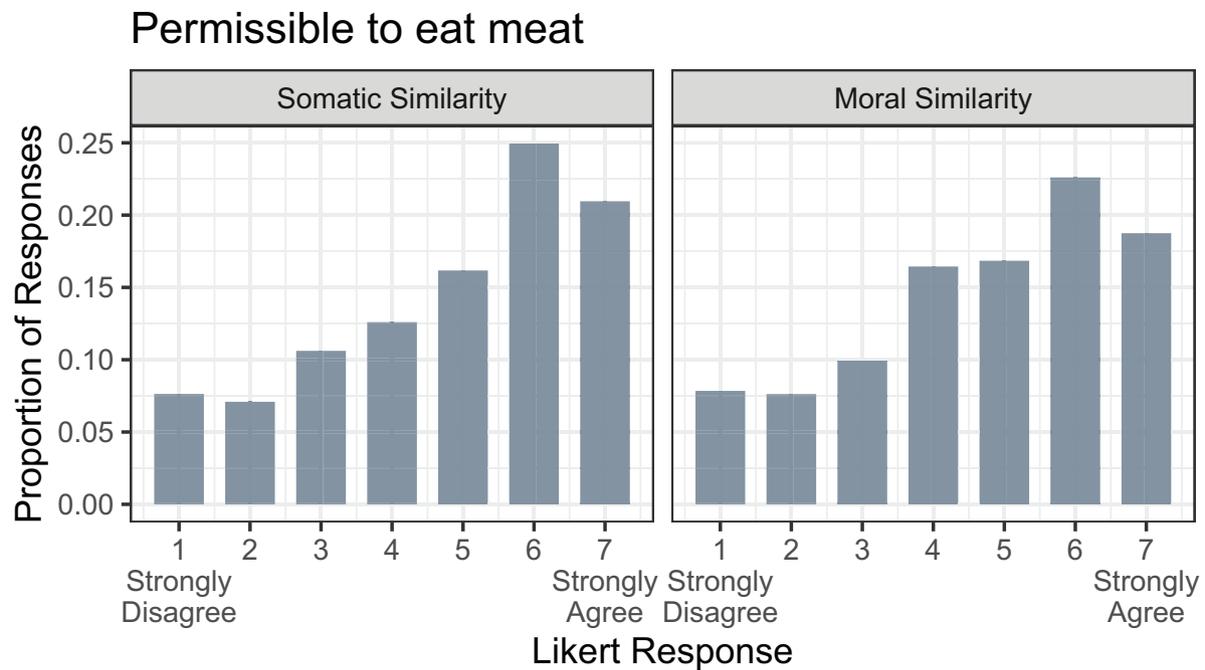


Fig. 4. Beliefs about the wrongfulness of eating meat in general in Experiment 2. Higher Likert scale choices indicate more favorable attitudes about eating meat.

including participants' pretest beliefs about their commitment to eating meat and their general beliefs about the morality of eating animals. General pretest beliefs were modeled as a monotonic effect when possible to reflect the cumulative nature of the response scale (Bürkner & Charpentier, 2018). The regression models and priors can be found in the supplement.

Both analyses revealed that the coherence-based intervention did not credibly effect participants' beliefs about the wrongfulness of eating meat, $b_{\text{General}} = -0.13$, 95% Posterior Uncertainty Interval [-0.53, 0.26] (see Fig. 4) and $b_{\text{Pig}} = 0.23$, 95% Posterior Uncertainty Interval [-0.28, 0.72] (see Fig. 5). These relationships were also assessed after interacting General Pretest and Pretest Meat Commitment beliefs with Condition, further confirming the Moral Similarity intervention did not exert a credible affect participants' beliefs about the ethics of eating meat relative to the Somatic Similarity condition (see Figs. 1–3 in the SOM). Still, we note the posteriors of these effects were compatible with the posteriors in Experiment 1. We address this point in Experiment 3.

As discussed, we also assessed people's beliefs using several other measures to better understand the impact and mechanisms underlying any effects of highlighting moral similarity. Specifically, we included a

scale examining people's commitment to eating meat and their perceptions of the similarity between dogs and other animals (e.g., pigs, chickens). Importantly, we observed that participants in the Moral Similarity condition did not judge dogs and pigs as more similar than in the Somatic Similarity condition (see Appendix A for further details); this may explain why we observed only a weak effect of the intervention on participants judgments about eating pigs.

Our preregistered analyses provided only weak evidence that an intervention highlighting moral similarity affected general beliefs about eating meat or specific beliefs about eating pigs. This may suggest that coherence-based interventions more generally are not an effective means of changing strongly held moral beliefs that impact everyday behavior. However, given the higher than expected number of exclusions in Experiment 2, and the fact that we did not observe that participants in the Moral Similarity condition rated dogs and pigs as more similar than they did in the Somatic Similarity condition, we conducted a third experiment designed to reduce the likelihood of participant exclusions due to inattention and diagnose the null similarity effect in Experiment 2.

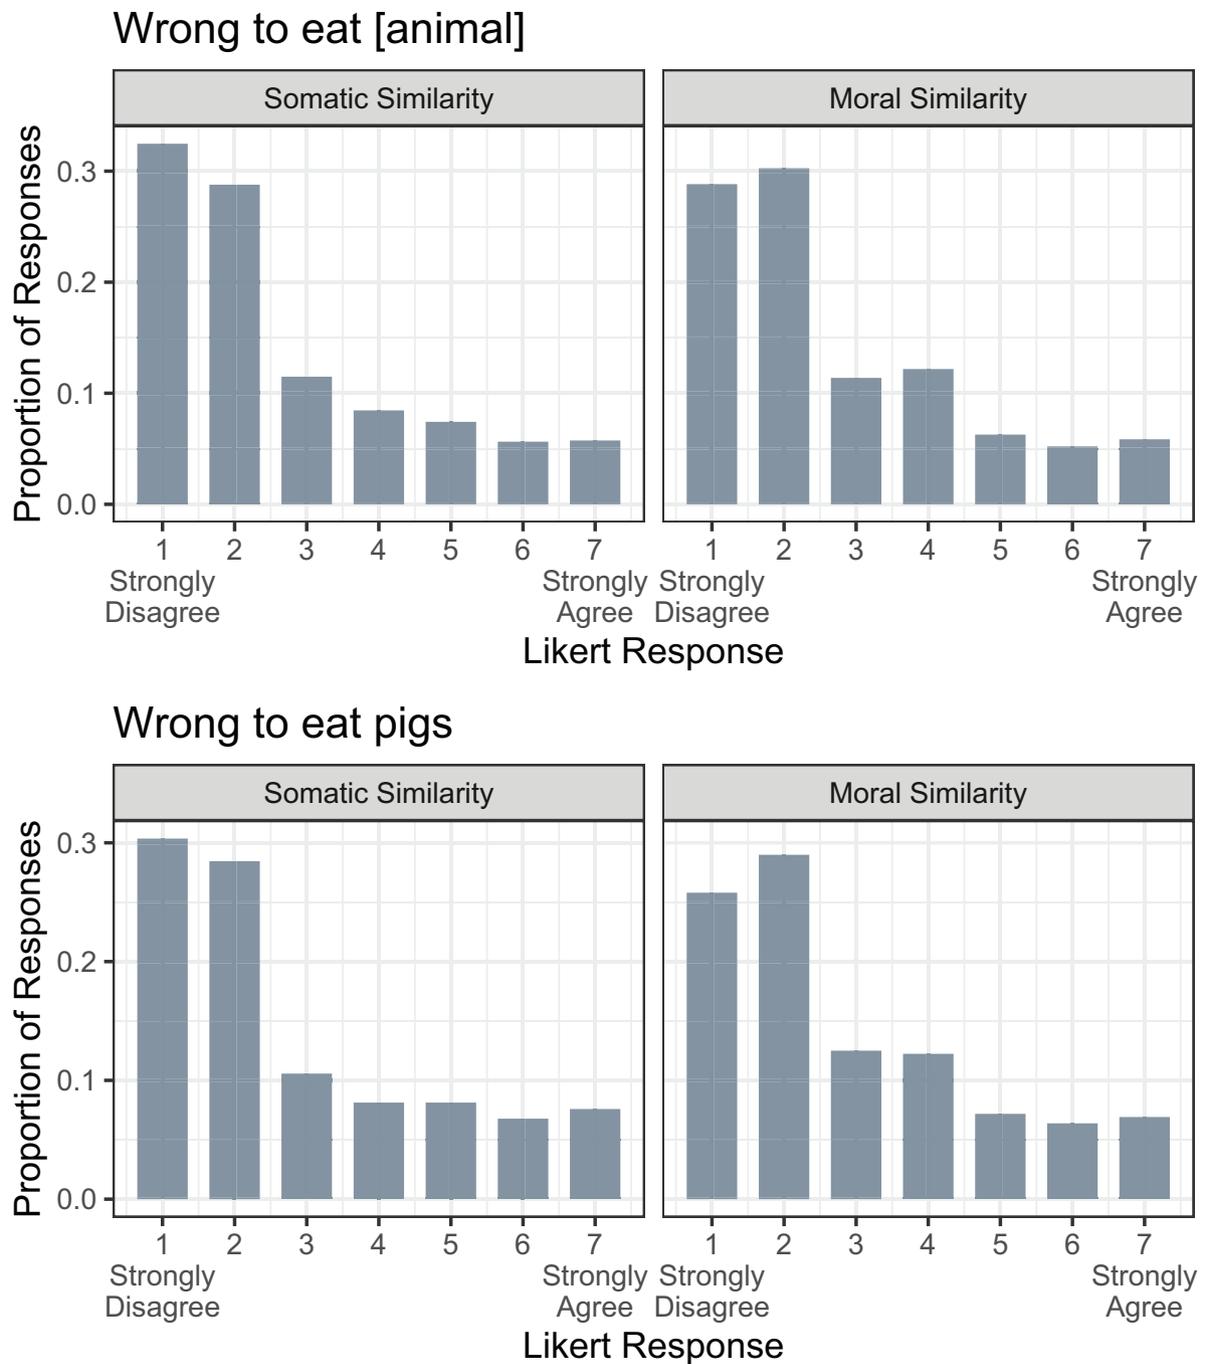


Fig. 5. Beliefs about the wrongfulness of eating meat from specific animals in Experiment 2. Higher Likert scale choices indicate less favorable attitudes about eating meat from specific animals.

5. Experiment 3

Together, Experiments 1 and 2 provided mixed evidence for the efficacy of coherence-based interventions which highlight the moral rather than somatic similarities on common moral beliefs. Experiment 3 sought to address the limitations of both of experiments to better understand how highlighting moral similarities impacts beliefs about animal ethics.

5.1. Method

5.1.1. Participants

Participants in Experiment 3 were 1150 Amazon Mechanical Turk workers ($M_{age} = 40$, 47% female) who were paid \$1.00 for participating in the study, or approximately \$8.00 an hour. An additional 159 were tested but excluded before analyzing the data for failing attention checks. We required that participants be located in the United States to accept the study HIT. The sample size and exclusion rules were determined prior to data collection in accordance with our preregistration: <https://osf.io/wgczf/>. A sensitivity power analysis for an independent samples *t*-test to compare judgments in the Moral Similarity condition

and control condition indicated we could detect an effect as small as a Cohen's d of 0.20 (log odds ratio = 0.36) with 80% power.

5.1.2. Materials and procedure

In Experiment 3, we sought to address key limitations of the second experiment. Specifically, because our similarity manipulation check did not provide evidence that participants perceived pigs and dogs as more similar to each other in the Moral Similarity condition than in the Somatic Similarity condition, we created a control condition that was even more sparse in the comparison that was made between dogs and pigs. This control condition was created to remain closely matched to the Similarity conditions in all relevant ways, and allowed us to determine whether the Somatic Similarity condition was subtly affecting participants' beliefs relative to a true baseline. To further address this limitation of Experiment 2, we increased the relevance and specificity of items in the similarity manipulation check. This allowed us to be certain that each intervention was in fact increasing the similarity that participants perceived between pigs and dogs (see Table 6 of the SOM). For example, participants were asked to rate how similar pigs were to dogs "in their capacity to suffer" and "their capacity to reason."

Second, we adjusted the memes to make them less offensive. Although typical of many common memes found online, the combination of disturbing photos used in Experiments 1 and 2, and the statement "you find it acceptable to eat pigs...but wrong to eat dogs" could have led to moral distancing (Barkan, Ayal, Gino, & Ariely, 2012) that diminished differences between conditions by causing participants to disengage from the task. Consequently, we (1) replaced the disturbing photographs of pigs and dogs with pictures that were more likely to elicit compassion and empathy, and (2) removed the acceptability statement from all memes (see Fig. 6 below and Tables 7–9 of the SOM).

Third, we shortened the number of questions participants were asked at pretest and posttest. At pretest, we combined items from the Meat Commitment Scale and the Revised General Meat Scale (e.g., the question I enjoy the taste of meat; see Table 10 of the SOM), and we only asked participants about eating dogs, pigs, and cows at posttest (see Table 11 of the SOM). Again, our aim was to improve participant attention to reduce the number of participants we excluded.

Finally, we added a short exploratory posttest measure to examine the possibility that highlighting the moral similarities of dogs and pigs impacts beliefs about the *treatment* of animals in general. We asked participants at posttest about the moral permissibility of animal testing (e.g., participants rated their agreement with the statement "Cosmetics research that uses animal testing cannot be justified and should be stopped"). The details of these exploratory analyses are located in Appendix B.

5.2. Results

5.2.1. Preregistered analyses

As in the prior two experiments, we tested whether participants in the Moral Similarity condition would think that it was more wrong to eat meat in general (i.e., the General Meat Consumption Scale) and to eat specific animals (i.e., the Specific Meat Consumption Scale). We examined these possibilities by fitting a series of Bayesian mixed-effects cumulative regressions (see the Experiment 3 priors specification section of the supplement).

First, these analyses revealed that participants in the Moral Similarity condition were slightly more likely than participants in the Control condition and the Somatic Similarity condition to judge it was wrong to eat meat in general, though this effect is compatible with effects in the opposite direction $b_{\text{Moral}} = 0.13$, 95% Posterior Uncertainty Interval [-0.05, 0.30] (see Fig. 7). In contrast, the Somatic Similarity appeared to have a truly null effect on participants' responses relative to a true control condition $b_{\text{Outward}} = -0.02$, 95% Posterior Uncertainty Interval [-0.19, 0.16].

Second, and most importantly, we found that participants in the Moral Similarity condition were credibly more likely to think it was wrong to eat pigs than participants in the true Control condition, $b_{\text{Moral}} = 0.51$, 95% Posterior Uncertainty Interval [0.02, 1.00]. Participants in the Somatic Similarity condition were slightly more likely to think that it was wrong to eat pigs than participants in the Control condition, $b_{\text{Somatic}} = 0.25$, 95% Posterior Uncertainty Interval [-0.22, 0.74], but the posterior estimate of this effect was also compatible with effects in the opposite direction (Fig. 8). These relationships were also assessed after

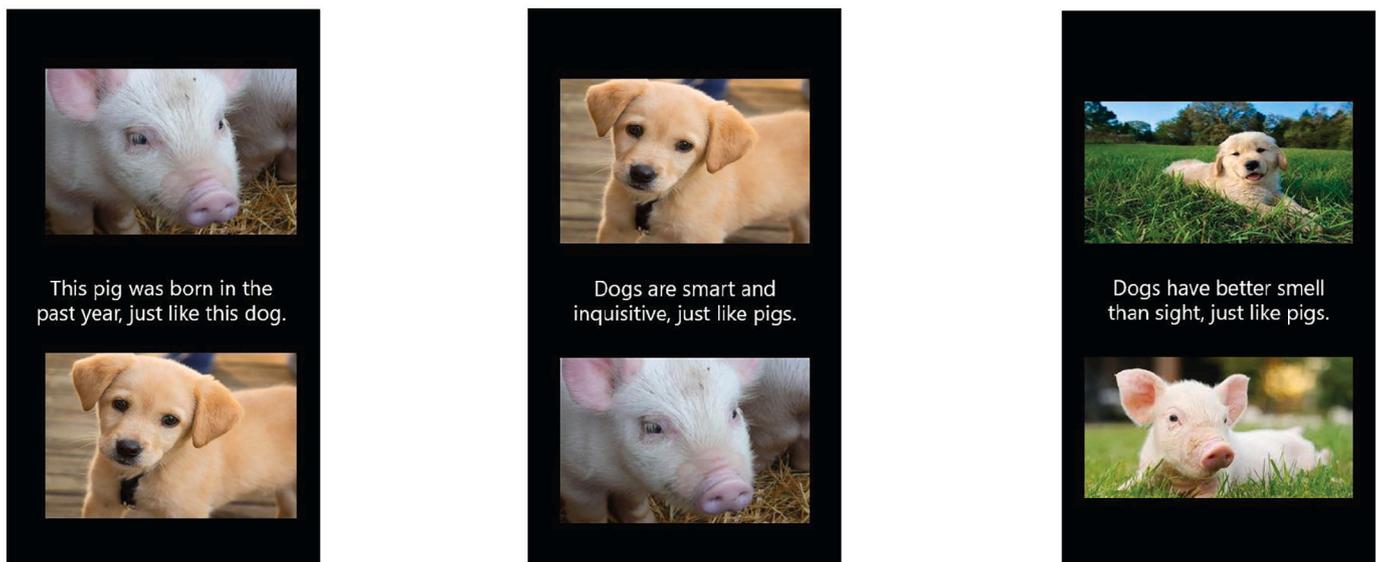


Fig. 6. Sample stimuli used in Experiment 3. Participant saw three memes from the Control (left), Moral Similarity (center), or Somatic Similarity (right) conditions in a between-subjects design. For brevity, the table shows representative examples from each condition.

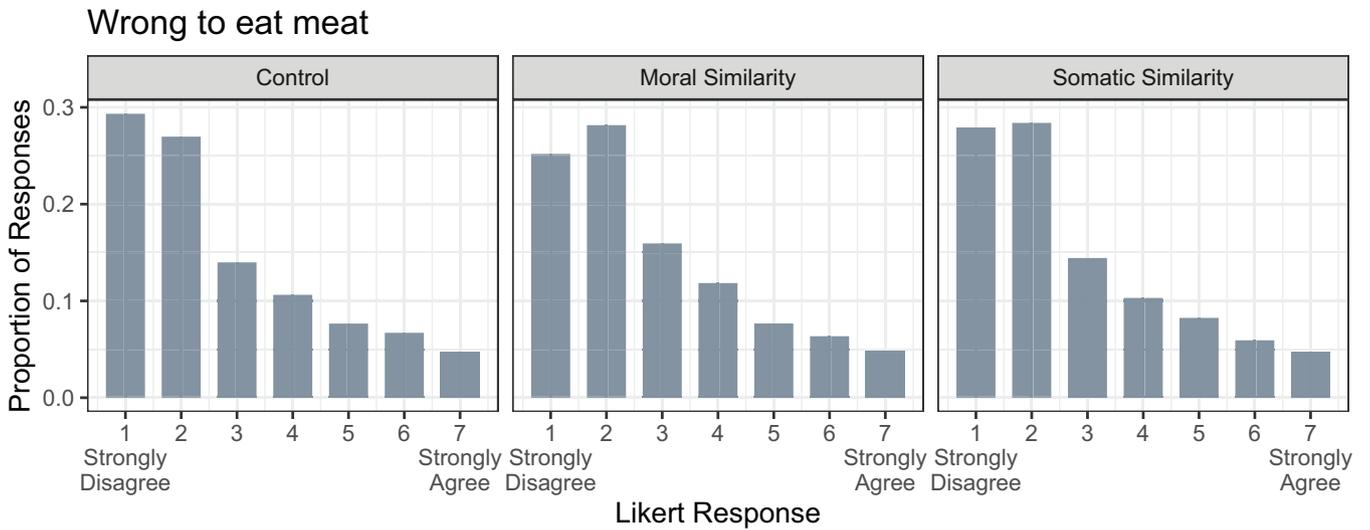


Fig. 7. Beliefs about the wrongfulness of eating meat generally in Experiment 3. Higher Likert scale choices indicate less favorable attitudes about eating meat in general.

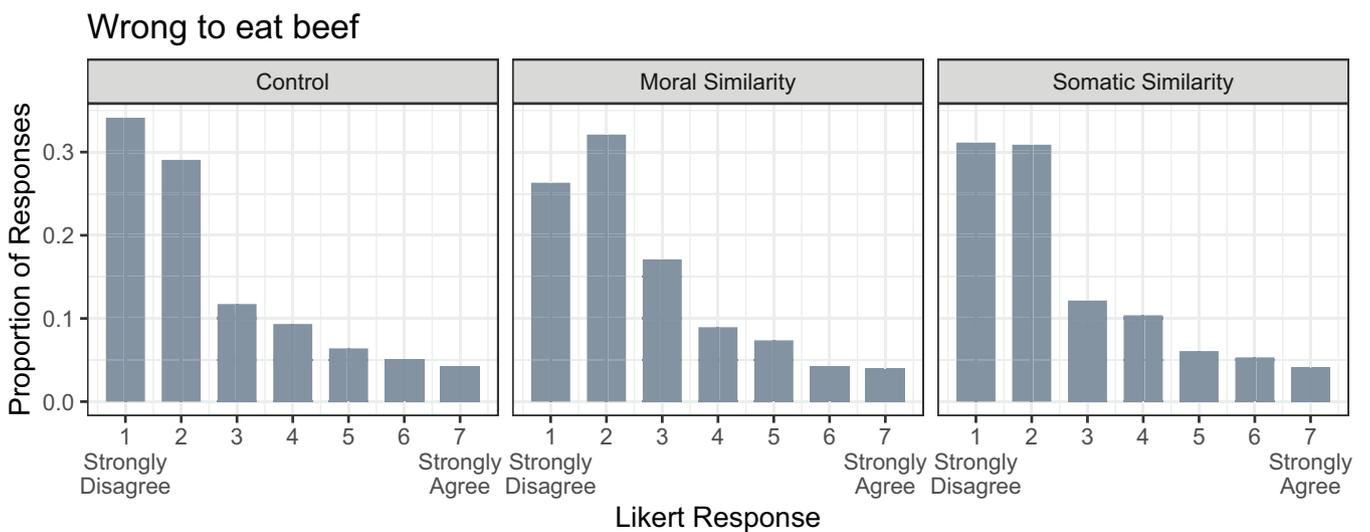
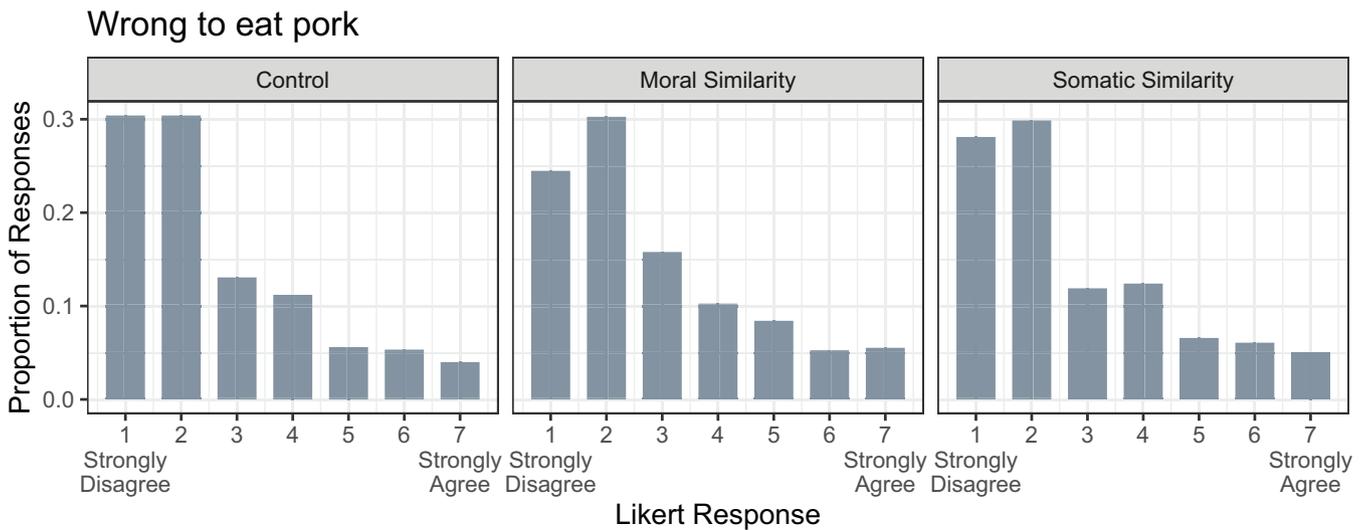


Fig. 8. Beliefs about the wrongfulness of eating pigs and cows in Experiment 3. Higher Likert scale choices indicate less favorable attitudes about eating pork and beef.

interacting pretest beliefs and Condition, but the inclusion of the interaction did not materially change the observed condition effects (see Figs. 4 and 5 in the SOM).

The slight positive shift caused by the Somatic Similarity intervention could at least partially account for the mixed findings in Experiments 1 and 2. Consistent with this possibility, the shared variance in participants' beliefs about the moral wrongness of eating pigs and eating dogs was nominally higher in both the Moral ($R^2 = .16$) and Somatic ($R^2 = .16$) Similarity conditions, relative to the Control condition, $R^2 = .12$. Furthermore, we found that relative to the Control and Somatic Similarity conditions, participants in the Moral Similarity condition rated dogs and pigs as more similar along key dimensions (namely, the capacity to reason). The details of these analyses are in [Appendix B](#).

Comparing the effect of each intervention on beliefs about eating cows revealed a similar pattern of responses: Participants in the Moral Similarity condition were more likely than those in the Control condition to think it was wrong to eat cows, but this effect appears to be smaller and is also compatible with effects in the opposite direction, $b_{\text{Moral}} = 0.30$, 95% Posterior Uncertainty Interval [-0.14, 0.73]. The effect of the Somatic Similarity condition relative to a true Control condition was considerably smaller, $b_{\text{Somatic}} = 0.08$, 95% Posterior Uncertainty Interval [-0.37, 0.53] and indistinguishable from noise. Finally, we observed a similar pattern in the shared variance in participants' beliefs about cows and dogs, $R^2_{\text{Moral}} = .14$, $R^2_{\text{Somatic}} = .12$, $R^2_{\text{Control}} = .10$. These results suggest participants in *both* the Moral and Somatic Similarity conditions either explicitly or implicitly increased the coherence of their beliefs relative to the Control condition (see Exploratory Analyses in [Appendix B](#)).

Altogether, the preregistered analyses of Experiment 3 suggest that the Somatic Similarity condition in Experiments 1 and 2 may have exerted a small effect on participants' responses, thus diminishing differences with the Moral Similarity condition relative to what would have been found with a true baseline condition. In Experiment 3, when we compared the Control condition and the Moral Similarity condition, we observed a credible effect of an intervention that emphasized the moral similarity between pigs and dogs. Experiment 3 thus provides the clearest evidence that a coherence-based intervention focused on the moral similarities of two animals, one which is typically consumed and another which is not, can subtly shift beliefs. This finding is also notable given that we did not highlight the inconsistency in participants' beliefs directly – we removed the sentence "you find it acceptable to eat pigs... but wrong to eat dogs." And yet participants were still more likely to judge that it was wrong to eat pigs after comparing them to dogs along moral dimensions.

All three experiments included a comparison between Somatic and Moral Similarity interventions. Moreover, all three experiments—including those inconsistent with our hypotheses—were preregistered and are reported here, allowing us to create a meaningful meta-analytic estimate of the condition effect on beliefs about the ethicality of eating pigs and cows ([Vosgerau, Simonsohn, Nelson, & Simmons, 2019](#)). This analysis indicated coherence-based interventions highlighting moral similarity consistently led participants to judge it was more wrong to eat pigs, $b_{\text{Moral vs. Somatic}} = 0.20$, 95% Posterior Uncertainty Interval [0.05, 0.35], Odds-Ratio = 1.22 (see [Fig. C1](#) in [Appendix C](#)). Altogether, this analysis indicates that not only do Moral Similarity coherence manipulations influence beliefs about everyday moral behaviors relative to a true control condition, they are also slightly more effective than interventions emphasizing somatic similarities of dogs and pigs.

6. General discussion

Three experiments suggest that confidence in an everyday moral belief—that it is morally permissible to eat pigs—may be shifted by brief coherence-based interventions. We found that drawing an explicit comparison between morally relevant capacities (e.g., intelligence, emotionality) of an animal that is frequently eaten (pigs) and an animal whose consumption is typically forbidden (dogs) shifted beliefs about the permissibility of eating pigs. Further, we found that highlighting morally relevant capacities did not more globally shift beliefs about eating meat, although this intervention may have subtly changed beliefs about eating some animals perceived as being very similar to pigs (i.e., cows).

In ethics, networks of beliefs are assessed by how they cohere, by how well beliefs in the network hang together. Successful ethical theories are evaluated by how well they can account for the intuitions we have about various situations, in addition to how well they respect overarching moral principles. Theories deemed inconsistent are often rejected on these grounds alone (see [Daniels, 1979](#), for a discussion of so-called "reflective equilibrium"). Coherence-based arguments and the process of reflective equilibrium are a common tactic for assessing ethical theories and changing moral beliefs in philosophy and the law. Given that people appear to have many moral beliefs that do not cleanly cohere with each other in the domain of animal ethics, we predicted that highlighting these inconsistencies either explicitly (Experiments 1 and 2) or more implicitly (Experiment 3) would lead people to increase the coherence of their beliefs. This prediction aligns with methodology in philosophy and the law: theories are routinely evaluated by drawing out their implications to evaluate their plausibility. Our results suggest that this rhetorical strategy typical in ethics and the law (e.g. [Holyoak & Powell, 2016](#); [Norcross, 2004](#)), may indeed be effective for even strongly held moral beliefs that affect routine behaviors (in this case, what people eat). We found that though we used a very minimal, naturalistic intervention, participants slightly shifted the confidence in their beliefs about the permissibility of eating pigs.

The effects we observed were subtle, but even a slight impact of a fleeting experimental manipulation on an enduring network of moral beliefs is notable. Other research has indicated that adults' moral convictions are difficult to manipulate, particularly for highly salient everyday issues (e.g., [Ciuk and Rottman, 2021](#)). While one recent experiment has documented a successful reduction in meat consumption, the intervention consisted of a lengthy class discussion about an essay arguing for the ethics of vegetarianism ([Schwitzgebel, Cokelet, & Singer, 2020](#)), and is thus impractical as a widespread intervention that could be used with the general public.

The effects we observed are also notable in the context of broader cognitive biases like motivated reasoning which can lead people to reject the premise that pigs and other farm animals are highly sentient and intelligent, thus de-mentalizing them to preserve coherence ([Bastian, Loughnan, Haslam, & Radke, 2012](#); [Loughnan, Haslam, & Bastian, 2010](#); [Piazza et al., 2015](#)). Indeed, motivated reasoning may have led our coherence-based intervention to backfire for some people: comparing dogs and pigs could have led some people to shift their beliefs about the impermissibility of eating dogs rather than their beliefs about the permissibility of eating pigs. Yet, we observed that the coherence intervention led people to adopt a moral belief that could impact their everyday behavior – by becoming less confident that it is permissible to eat pigs they might also now be less likely to engage in a practice they previously thought was acceptable. This is notable because, in some

ways, it might be easier for people to adopt the belief that it is permissible to eat dogs, *even if they wouldn't themselves engage in the practice*, because taking on this belief would not imply they ought to change their behavior.

The results of our studies highlight several promising avenues for future research. First, our work raises the question of what other kinds of evidence, arguments, or other stimuli lead people to become vegetarians. For example, what information do vegetarians *think* caused them to become vegetarian? Most vegetarians in America were first omnivores, indicating that some arguments, information, or emotional appeals (at some point) changed their minds about eating meat. It is possible that coherence-based interventions and related arguments subtly shifted these beliefs given their presence in many online forums, but longitudinal research and descriptive studies are necessary to resolve this question. Second, future work could investigate whether lengthier arguments typical of ethics and the law are more persuasive than the brief memes presented in this work (see [Feinberg et al., 2019](#)) – we would expect they would be but further research is necessary to establish this. Third, further research should examine the question of whether coherence-based arguments are similarly effective in impacting specific, everyday moral beliefs outside the domain of animal ethics. It is possible that animal ethics presents a unique case, particularly as refraining from meat is often viewed as supererogatory, even for vegetarians ([Hussar & Harris, 2010](#)). As decisions about meat consumption are determined by myriad factors, only some of which are relevant to moral values and beliefs about animal minds (see [Amiot & Bastian, 2015](#); [Loughnan et al., 2014](#); [Plous, 1993](#); [Ruby, 2012](#)), shifting beliefs about the similarity of dogs and pigs is likely only part of the network of reasons that could

change perceptions of the moral status of non-human animals. For these reasons, other commonplace moral beliefs may be easier to change using coherence-based interventions (but see [Miske et al., 2019](#)). Finally, the unexpected finding that the Somatic Similarity condition induced small coherence-shifts also raises the question of how deep similarities need to be in order to produce coherence-based effects. It is possible that even shallow comparisons will lead people to shift their moral beliefs in various domains, such that comparatively superficial similarity will also shift moral beliefs. This possibility is compatible with the finding that people tend to retrieve relatively shallow, surface similar sources when performing analogical inference (e.g., [Gentner, Rattermann, & Forbus, 1993](#)).

The present research provides evidence that coherence-based interventions highlighting moral similarities can subtly shift confidence in even intransigent moral beliefs. This work demonstrates that coherence-based arguments similar to those found online, and more broadly representative of argumentative strategies in the law and philosophy, may indeed be an effective means for shifting confidence in moral beliefs that impact people's routine behaviors.

Open practices

The experiments in this article earned Open Materials, Open Data, and Preregistered badges for transparent practices. Materials, data, and preregistrations are available at <https://osf.io/uufnu/> (Experiment 1), <https://osf.io/zmr4s/> (Experiment 2), and <https://osf.io/wgczf/> (Experiment 3).

Appendix A. Experiment 2 exploratory analyses

Analytic Strategy for Exploratory Analyses

To reduce the likelihood of model overfitting in our exploratory analyses, we perform k -fold cross-validation, a common technique in machine learning for assessing how well a model fits unobserved data. Cross-validation can proceed in several ways (see [Arlot, Celisse, et al., 2010](#)), but broadly all cross-validation techniques involve (1) training a model on a dataset and (2) testing how well the model fits data the model has not seen. In our exploratory analyses, we assessed the fit of a set of models using the change in their expected log predictive density Δelpd ([Vehtari, Gelman, & Gabry, 2017](#)). As a rule of thumb, models which only differ slightly in their Δelpd (e.g., < 4) perform similarly at predicting observations the model has not been trained on.

We assessed people's beliefs about the ethics of eating pigs using several other measures to better understand the impact and mechanisms underlying any effects of highlighting moral similarity. Specifically, we included a scale examining people's commitment to eating meat and their perceptions of the similarity between dogs and other animals (e.g., pigs, chickens). To examine the possible effects of Condition on these dependent variables, we performed exploratory model selection by assessing the out-of-sample predictions of a series of regression models ([Vehtari et al., 2017](#)).

Table A1

Bayesian mixed-effects cumulative regression models exploration of meat commitment, similarity ratings, and specific meat consumption attitudes. The table displays estimates of the change in the expected log posterior predictive density (Δelpd) and the error in this estimate.

Dependent variable	Model	Δelpd	SE
Meat commitment	Intercept + Pre-Intervention Beliefs	–	–
Meat commitment	+ Condition	–7.7	3.6
Similarity rating	Intercept	–	–
Similarity rating	+ Condition	–1.2	13.1
Similarity rating	+ Condition \times Animal	–21.2	14.3
Specific meat consumption	Intercept + Similarity	–	–
Specific meat consumption	+ Similarity \times Condition	4.3	19.5

Note: Bolded row indicates preferred model.

We first examined how the Moral Similarity intervention affected commitment to eating meat, finding that the most predictive model did not include Condition as a predictor. This result is consistent with our preregistered analyses, which suggested that an intervention highlighting moral similarity did not credibly impact beliefs toward eating meat. Next, we tested the hypothesis that highlighting moral similarities would make people more likely to think these animals were similar to each other. To test this hypothesis, we regressed Similarity Judgments on Condition and the interaction between Condition and Animal (i.e., the animal being compared). The second model included the interaction term to test the hypothesis that the effect of Condition on Similarity ratings could vary as a function of the animal participants were comparing to dogs. For example, it may be that an intervention highlighting moral similarity caused people think that pigs and dogs, but not fish and dogs, were similar to each other. This analysis indicated that the most predictive model did not include Condition as a predictor. Thus, the intervention did not affect how similar people thought pigs were to dogs.

Finally, we anticipated that participants who viewed an animal as similar to a dog would be more likely to think it is wrong to eat that animal. This possibility would be confirmed an interaction between Condition and Similarity ratings of a given animal to a dog – people who viewed a given animal as similar to a dog may be more affected by the Moral Similarity intervention than those who did not. Contrary to our hypothesis, including the interaction between Condition and Similarity Judgments did not improve out-of-sample prediction, even though the posterior of the interaction coefficient excluded zero on training data, $b = 1.34$, 95% Posterior Uncertainty Interval [0.41 to 2.52] (Yarkoni & Westfall, 2017). Altogether, exploratory analyses provide further confirmation of the results of our preregistered analyses.

Appendix B. Experiment 3 exploratory analyses

Exploratory analyses of similarity judgments. We predicted that the Moral Similarity condition would cause participants to judge that pigs and dogs were more similar to each other relative to the Somatic Similarity condition. In Experiment 2, we tested this prediction by directly asking participants to judge how similar pigs are to dogs following the intervention. However, because Experiment 2 lacked a true control condition and because the Somatic Similarity condition highlighted some dimensions on which dogs and pigs are similar (for instance, in their perceptual capacities), it is possible that participants judged dogs and pigs as being similar to each other for different reasons. Consequently, in Experiment 3 we addressed this issue by asking participants to judge the similarity of dogs and pigs along several different dimensions.

We made an a priori prediction that the Moral Similarity would affect participants' similarity judgments relative to the Control and Somatic Similarity conditions, but we did not have explicit hypotheses about *which* dimensions this intervention might affect. (We suspected the Moral Similarity condition would be unlikely to impact responses to the DNA question, for example, but more likely to impact responses to the questions about the capacity to reason and suffer). Thus, we investigated how each condition affected participants similarity judgments in a series of exploratory analyses by performing k -fold cross-validation and assessing the out-of-sample predictions of a set of models. The prior distributions for the reduced and full models are specified in the supplement [Appendix D](#). For brevity, we will only report the primary findings of this set of analyses, but our analysis code and data are available for the interested reader.

Our analyses suggest the model which included the interaction between Condition and Similarity Dimension (i.e., the full model) improved out-of-sample prediction relative to the reduced model which did not include this interaction, $\Delta\text{elpd} = 45.30$, $SE = 14.40$, indicating the effect of condition differed as a function of the similarity dimension in question. Inference on posterior parameter estimates in exploratory modeling should proceed with considerable caution, but an inspection of the raw and conditional effects plot ([Figs. B1 and B2](#)) suggests that the Moral Similarity manipulation appears to shift similarity judgments along one particular dimension; the Moral Similarity condition increased the belief that pigs and dogs have the same capacities for reason. To speculate, the null similarity effect in Experiment 2 could be the result of the fact that the Moral Similarity intervention increasing perceptions of similarity along a particular dimension rather than wholesale.

Similarity Ratings

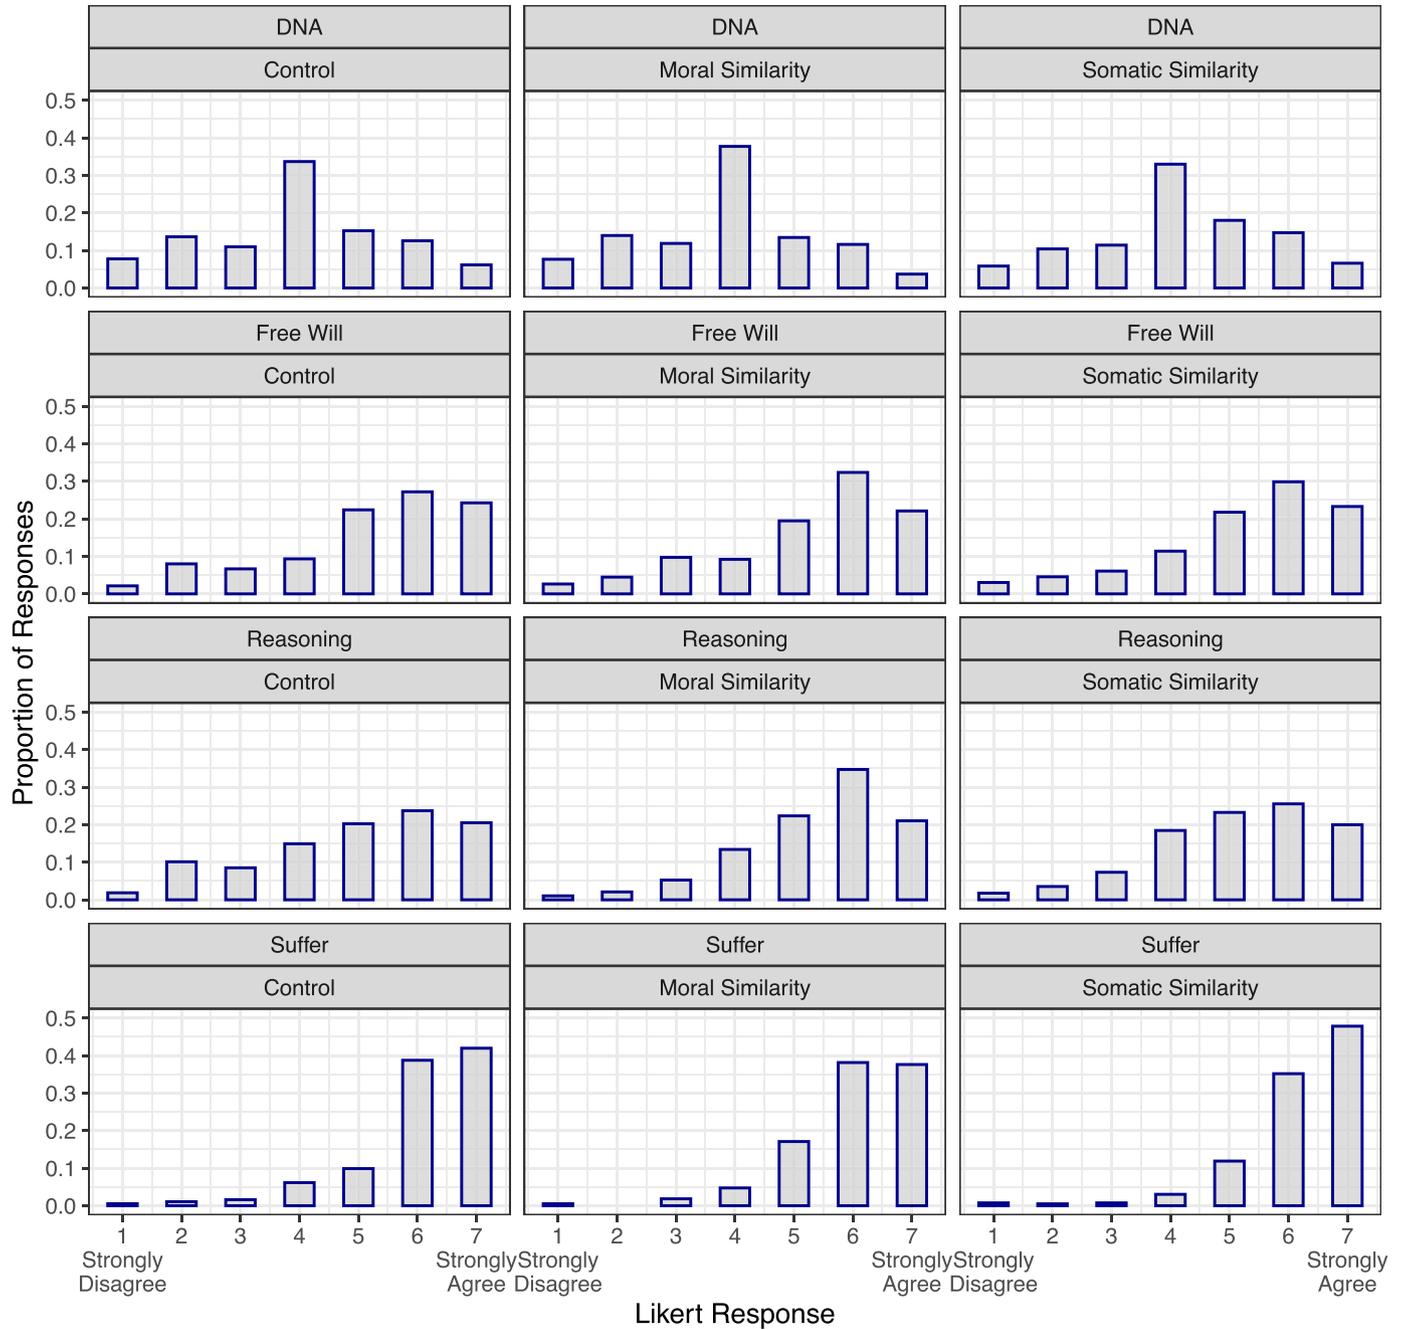


Fig. B1. Beliefs about the similarity between dogs and pigs in Experiment 3 across four dimensions. Higher Likert scale choices indicate higher similarity ratings. This figure suggests similarity was rated similarly across all three conditions except for ratings of the reasoning capacities of dogs and pigs.

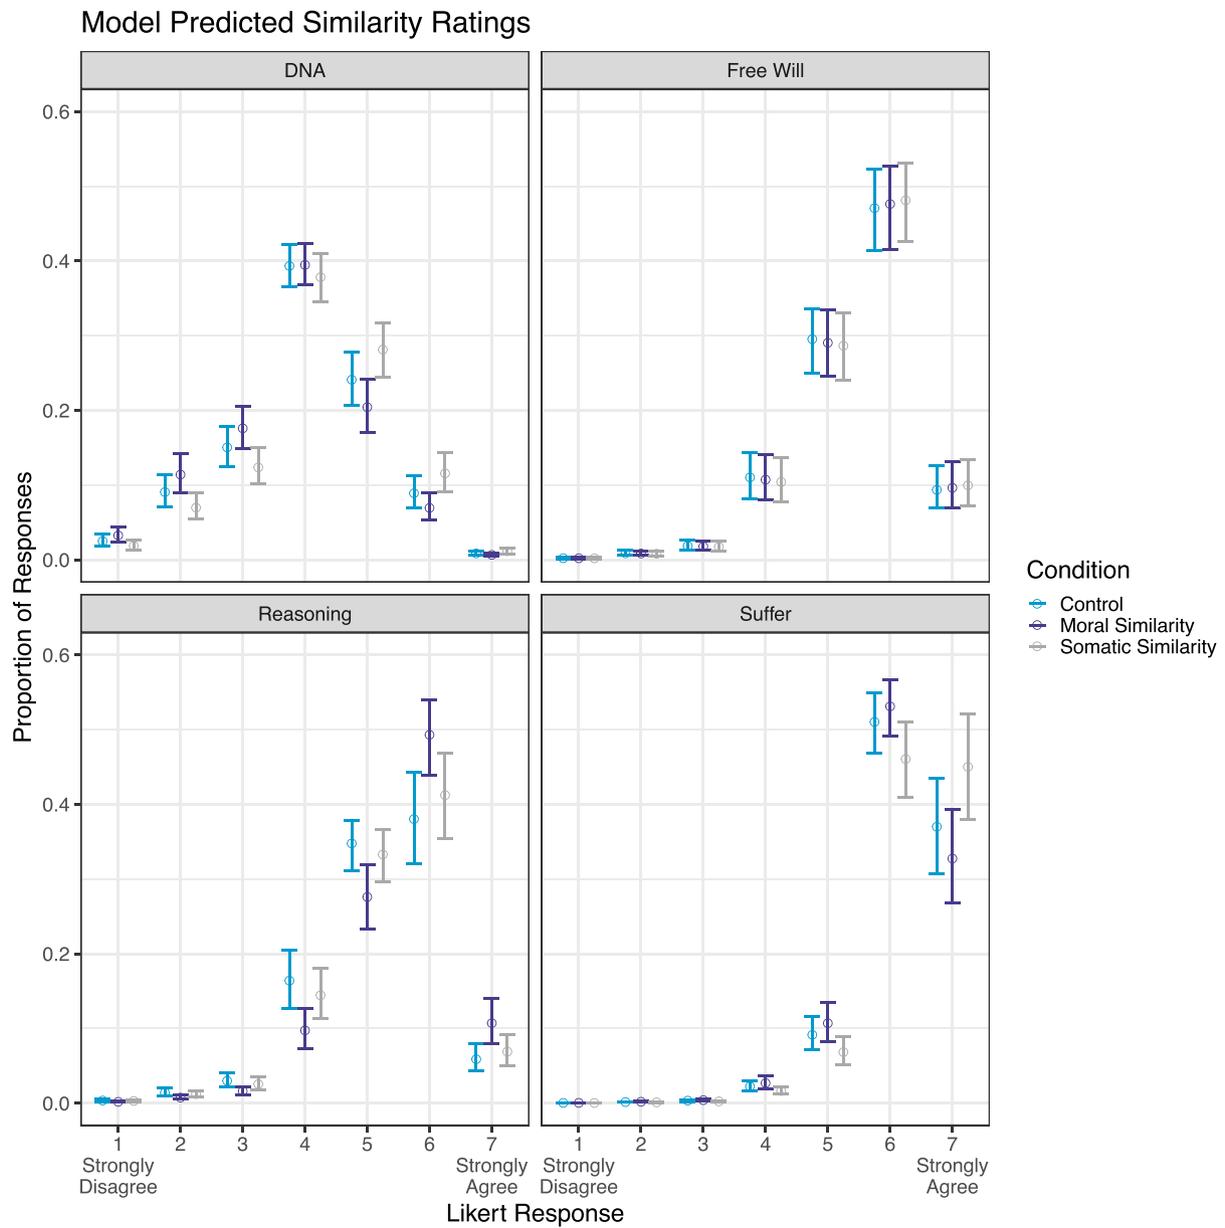


Fig. B2. A conditional effects error-bar plot of beliefs on four dimensions of similarity. As a rough heuristic, when the 95% Posterior Uncertainty Interval from one condition does not overlap the dot of another interval, those intervals are credibly different from each other. Thus, like Fig. B1, this plot indicates that participants in the Moral Similarity condition were more likely to judge that dogs and pigs have similar reasoning capacities. These small differences were robust to informative regularizing priors.

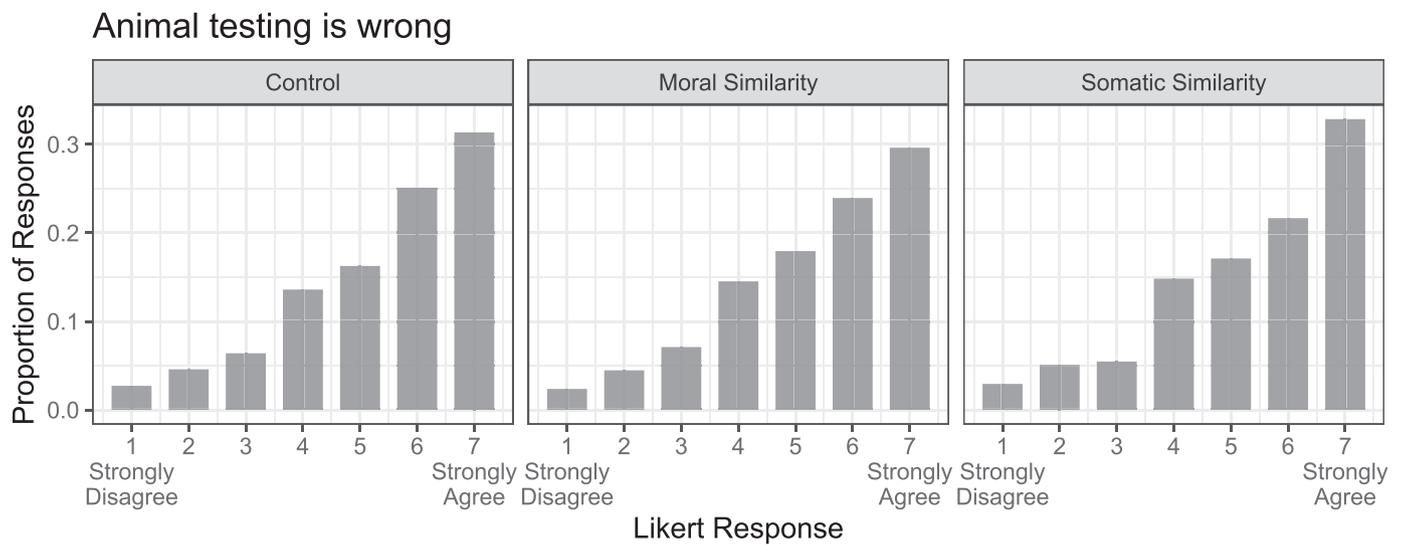


Fig. B3. Beliefs about the wrongfulness of animal testing in Experiment 3. Higher Likert scale choices indicate less favorable attitudes about animal testing.

Exploratory analyses of animal testing judgments. We also included an additional exploratory measure at posttest to examine beliefs not directly related to the consumption of meat. One possibility is that in addition to change their beliefs about the ethicality of eating pigs, participants may also shift their beliefs about the treatment of pigs in other domains. To investigate this possibility, after the intervention, participants were asked to rate how much they agreed with statements about the moral permissibility of animal testing (e.g., "Cosmetics research that uses animal testing cannot be justified and should be stopped." and "We need more regulations governing the use of animals in research."). Fig. B3 shows the proportion of responses at each Likert point in each condition. A Bayesian cumulative mixed model indicated that not only were participants generally against animal testing (a finding inconsistent with our preregistered threshold priors), but that neither the Somatic nor Moral Similarity condition credibly affected participants' judgments about the moral permissibility of animal testing, $b_{Moral} = 0.15$, 95% Posterior Uncertainty Interval [-0.59, 0.29] and $b_{Somatic} = -0.01$, 95% Posterior Uncertainty Interval [-0.43, 0.43]. One possibility these findings suggest is the coherence-based interventions used here are unlikely to increase participants' already strong opposition to animal testing because these beliefs are consistent with their reduced endorsement of eating pigs – that is, this aspect of their network beliefs is already coherent.

Appendix C. Moral vs. somatic similarity interventions in Experiments 1–3

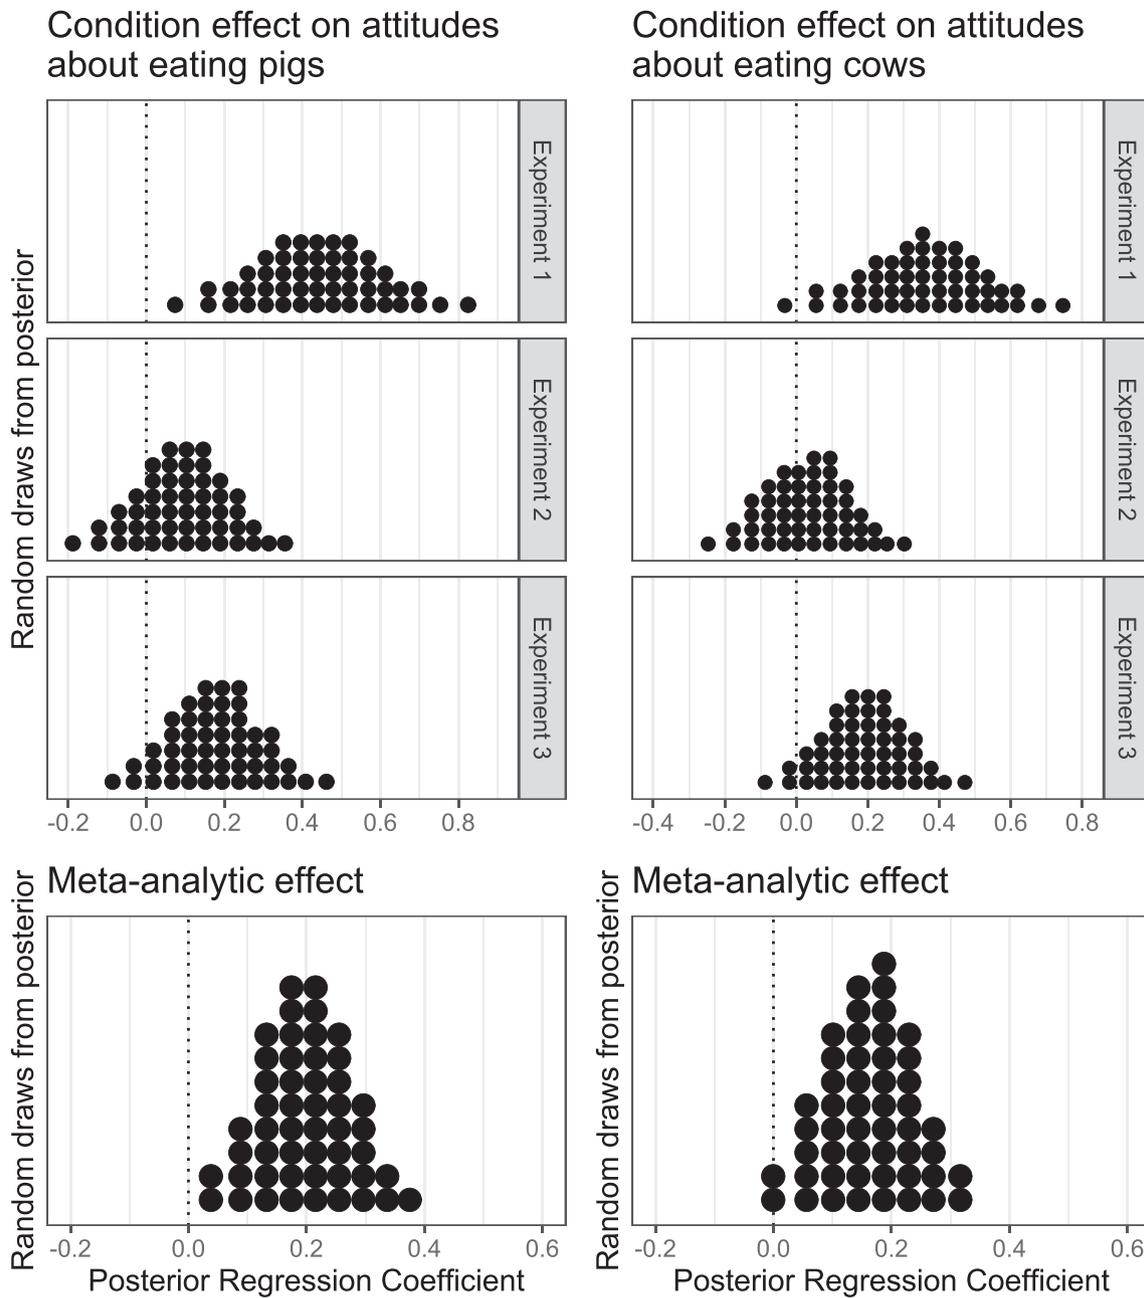


Fig. C1. Dot-plot of condition parameter comparing the Moral and Somatic Similarity conditions in Experiments 1 through 3. Each dot represents a random draw from the posterior distribution. For instance, the figure indicates that in Experiment 3 for every 50 draws, 47 draws are greater than 0. This analysis allows us to more directly quantify how often we would infer the Moral Similarity intervention would change beliefs relative to the Somatic Similarity intervention.

Appendix D. Prior specification

We now provide a walk-through to aid the reader in interpreting priors in cumulative regression models. Readers already familiar with specifying priors in Bayesian cumulative models can skip this section of the Appendix.

Stated most broadly, Bayesian statistical models formulate model parameters (which are unknown) as probability distributions wherein the joint probability distribution of the data, \vec{y} (which is known), and model parameters $\vec{\theta}$ are computed via the prior probability of $\vec{\theta}$ and the probability of \vec{y} given $\vec{\theta}$ (van de Schoot et al., 2021).

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

You can read this formula as follows: the joint probability of y and θ equals the probability of y given θ (i.e., the Likelihood) multiplied by the probability of θ (i.e., the Prior). Sticking with convention, y denotes the data and θ denotes a set of parameters of a distribution (e.g., μ and σ are parameters of the normal distribution). The equivalence stated above can be derived from Bayes' Theorem, which is used to calculate the posterior probability of $\vec{\theta}$ given \vec{y} :

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

You can read this formula as follows: The probability of θ given y (the Posterior) is proportional to the Likelihood and the Prior. For instance, if we wanted to compute the posterior probability of a mean given the data, we would need to compute the probability of the data given the mean and the prior probability of the mean.

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

To model the joint probability distribution of participants' responses [i.e., $p(y, \theta)$], we need to set priors over the possible effects each model parameter θ could have on the response variable y . The remainder of this section will describe the priors we set for each model. Throughout the paper, we fit cumulative models so we must first briefly discuss some details of these models before turning to the priors we set.

Cumulative models assume Y , which is observed, is underlain but an unobserved continuous distribution \tilde{Y} (Bürkner & Vuorre, 2019). These models assume within this continuous space there are K thresholds τ_k (which are similar but not identical to intercepts in a linear model), which partition \tilde{Y} into $K + 1$ response options of Y . For example, a response scale with seven Likert points will have six thresholds. As noted above, Bayesian analyses formulate model parameters as *probability distributions* which are jointly proportional to the prior and the likelihood. Thus, in our model we must set priors over each threshold, in addition to the β and group-level parameters (we explain these parameters below). Here and in subsequent analyses of Experiments 1 and 2, τ parameters specify the cumulative log-odds of choosing a given Likert point in the *Somatic Similarity condition* (the reference group of Experiments 1 and 2); the reference condition in Experiment 3 was the control condition.

Threshold priors. As an illustration, we now describe the threshold priors in our first regression model. The first model we fit is stated below in brms syntax for brevity.

Model 1:

```
General Meat Judgment ~ Condition + (1 + Condition | Scale Question) + (1 | Subject)
```

This is a mixed-effects model which will estimate: (1) a global intercept,⁴ (2) random effect intercepts for Scale Question (i.e. for each question in the General Meat Eating Scale, that question's intercept's deviation from the global intercept), (3) a single global estimate for the effect of Condition, (4) the effect of Condition for each question of the General Meat Eating Scale (more specifically, the degree to which the Condition effect for a given question deviates from the global effect of Condition), and (5) the correlation between the intercept deviations and Condition effect deviations for each question in the General Meat Eating Scale. This is the maximal model given our design (Barr et al., 2013).

Threshold priors for Model 1:

$$\tau_1 \sim \mathcal{N}(-2.51, 0.50)$$

$$\tau_2 \sim \mathcal{N}(-1.73, 0.50)$$

$$\tau_3 \sim \mathcal{N}(-1.09, 0.50)$$

$$\tau_4 \sim \mathcal{N}(-0.41, 0.50)$$

$$\tau_5 \sim \mathcal{N}(0.60, 0.50)$$

$$\tau_6 \sim \mathcal{N}(1.73, 0.50)$$

The threshold priors for our first model are listed above. We will now explain what these priors mean. We set the first threshold prior, τ_1 , as a normal distribution, \mathcal{N} . Normal distributions have two parameters μ and σ – the mean and standard deviation of the distribution. Recall that all regression parameters (here, τ_1) in Bayesian models are themselves distributions, so regression parameters themselves have parameters (often referred to as hyperparameters). Because we are performing cumulative regression, the mean and standard deviation of this normal distribution are the cumulative log-odds of selecting that Likert response option. So, if we convert these log-odds values to values on the response scale,⁵ the 95% posterior uncertainty interval ranges from 2% to 20% for τ_1 . The mean of this distribution is 7.5%, though the mean is of less relevance given how diffuse the prior distribution is. You can calculate these probabilities using the formula $\frac{\exp(x)}{1+\exp(x)}$. In essence, this is a formalization of the prior expectation that few participants will **Strongly disagree** that it's morally acceptable to eat meat – national public polls (Newport, 2012) and everyday experience indicate as much. The reader likely shares this expectation, but rather than informally interpreting regression models while bearing this expectation in mind, Bayesian models formally build this expectation into the regression model. To consider another threshold prior, τ_5 , from Model 1, the 95% posterior uncertainty interval on the response scale ranges from 2% to 46%. The mean of this distribution is 25% – our expectation is more participants will **Somewhat agree** that it's morally acceptable to eat meat. To clarify another common point of confusion, each τ represents the *cumulative* log-odds of choosing a given Likert point, and thus to compute τ_k we subtract τ_{k-1} . So to calculate the probability of selecting **Somewhat agree**, we use the formula $\frac{\exp(\tau_5)}{1+\exp(\tau_5)} - \frac{\exp(\tau_4)}{1+\exp(\tau_4)}$.

What do these priors mean? The threshold priors in the first model weakly bias model posteriors toward the agree-end of the scale because prior to running the study, we know people generally think it is morally acceptable to eat meat (Newport, 2012). However, these threshold priors also allow for possible, if unlikely Likert choices. Thus, our threshold priors in this first model allow the data to primarily—if not entirely—determine a given threshold's posterior. Consistent with this, when we estimated threshold parameters using Maximum Likelihood Estimation (MLE), the coefficients were often similar to our Bayesian model, though sometimes MLE coefficients were more exaggerated (e.g., the MLE model predicted that 0% of participants would select "Strongly disagree"). The full set of threshold priors are plotted below to illustrate their relative informativeness. For brevity, prior specifications for all remaining models are not plotted.

⁴ As noted above, in fact, the model is estimating six thresholds rather than a single global intercept, but we adopt the terminology of linear mixed-effects modeling for simplicity.

⁵ Note that this uncertainty is *jointly* determined by σ of τ_1 and the group-level prior $\sigma \sim \text{Half-}\mathcal{N}(1, 3)$, which we provide more detail about below.

We'll now pause to address a few questions that are often raised about Bayesian data analysis – more detailed explanations addressing these issues can be found in comprehensive textbooks (Kruschke, 2014; McElreath, 2016). First, a curious reader might wonder *how* the mean and standard deviation of these Gaussian distributions were chosen, or the impact choosing other means and standard deviations, or how choosing altogether different distributions (e.g., a Cauchy distribution, a Student-*t* distribution etc.) would impact the resulting posterior distributions. In essence, we chose priors by first simulating various possible priors and judging what priors were sufficiently informative that we could account for existing data on people's attitudes about eating meat while not biasing the model in such a way that it fails to learn from the data we've collected. In practice though, nothing material about the outcome of our analyses hinges on the particular threshold priors we chose – weaker priors, uninformative priors, or priors with subtly different central tendencies do not substantially impact threshold posteriors. This is because the sample sizes across all three studies is sufficiently large that the likelihood primarily determines the posteriors for threshold parameters.

The idea that priors are worth specifying at all but (often) do not substantially impact the resulting posterior distribution is a common point of confusion. After all, if these threshold priors do not materially impact the posterior why include them at all – why introduce bias into the model? The rationale is straightforward: providing structure to high-dimensional models like mixed-effects models (1) substantially improves the ability to fit more complex random effects structures, (2) addresses convergence failure (a common issue; Barr et al., 2013), and (3) substantially improves the speed with which these models can be fit (Schad, Betancourt, & Vasishth, 2020).

Priors on β parameters. Now we will discuss how we set priors on β parameters. Although threshold priors are important for model convergence and efficiency in fitting the model, β parameters are of considerably more theoretical importance. First, let us clarify what this parameter means in the present model – the β parameter is the log-odds of choosing a lower (or higher) Likert point in the Moral Similarity condition. More plainly, it is an estimate of the difference between the Moral Similarity condition and Somatic Similarity condition. We focus our discussion on interpreting this parameter because it most succinctly captures the aim of our studies. How did we specify priors on this parameter?

$$\beta_{\text{Condition}} \sim \mathcal{N}(0.00, 1.00)$$

In Experiment 1, the initial β prior we set is very weak, a normal distribution with a log-odds of zero and a standard deviation of 1, meaning it will exert very little influence on the posterior distribution. This prior considers large conditions differences either in the predicted or unpredicted direction credible (i.e., 95% Posterior Uncertainty Interval [Odds-Ratios from 0.14 to 7.39]). However, the mean of this prior is zero, which weakly biases the model against thinking there is a difference between the two conditions. Thus, this prior is *contrary* to our hypothesis. This is the norm in Bayesian statistical modeling (Kruschke, 2014): researchers often use priors that bias predictor estimates toward zero because samples routinely overestimate population-level effect sizes (Open Science Collaboration, 2015; Yarkoni & Westfall, 2017). Using priors to bias estimates toward zero is known as *regularization* and can yield more realistic (more regular) estimates. The motivation to use priors to regularize parameter estimates is similar to the motivation underlying the λ parameter in Ridge regression, Lasso regression, Elastic-net regularization, and other kinds of penalization for model complexity commonly used in machine learning. Routine overfitting in social psychology (Open Science Collaboration, 2015) prompted our decision to regularize and interrogate our models. We used multiple priors with different degrees of informativeness to formally address this issue; we examined how different prior beliefs—skepticism, uncertainty, and so on—impact the inferences we might draw (see Experiment 1 for an example of how different prior beliefs can influence the inferences drawn from data we've collected).

Group-level priors. Now, we'll briefly mention the priors we set for group-level effects. σ parameters specify group-level effects (e.g., the variability in intercepts across subjects) and, as we have set them, are functionally flat. Ω represents a prior for correlations between correlated observations, where $lkj = 1$ indicates that every correlation coefficient between a given participant's responses is (approximately) equally likely. For example, a correlation coefficient of $r = 0.1$ between two-responses from the same participant is as credible as a correlation coefficient of $r = 0.9$. *LKJ* distributions with larger values place a lower credence on strong correlations (i.e. Pearson r values closer to -1 and 1) in participants' responses.

$$\text{All } \sigma \sim \text{Half} - \mathcal{N}(13)$$

$$\Omega \sim lkj(1.25)$$

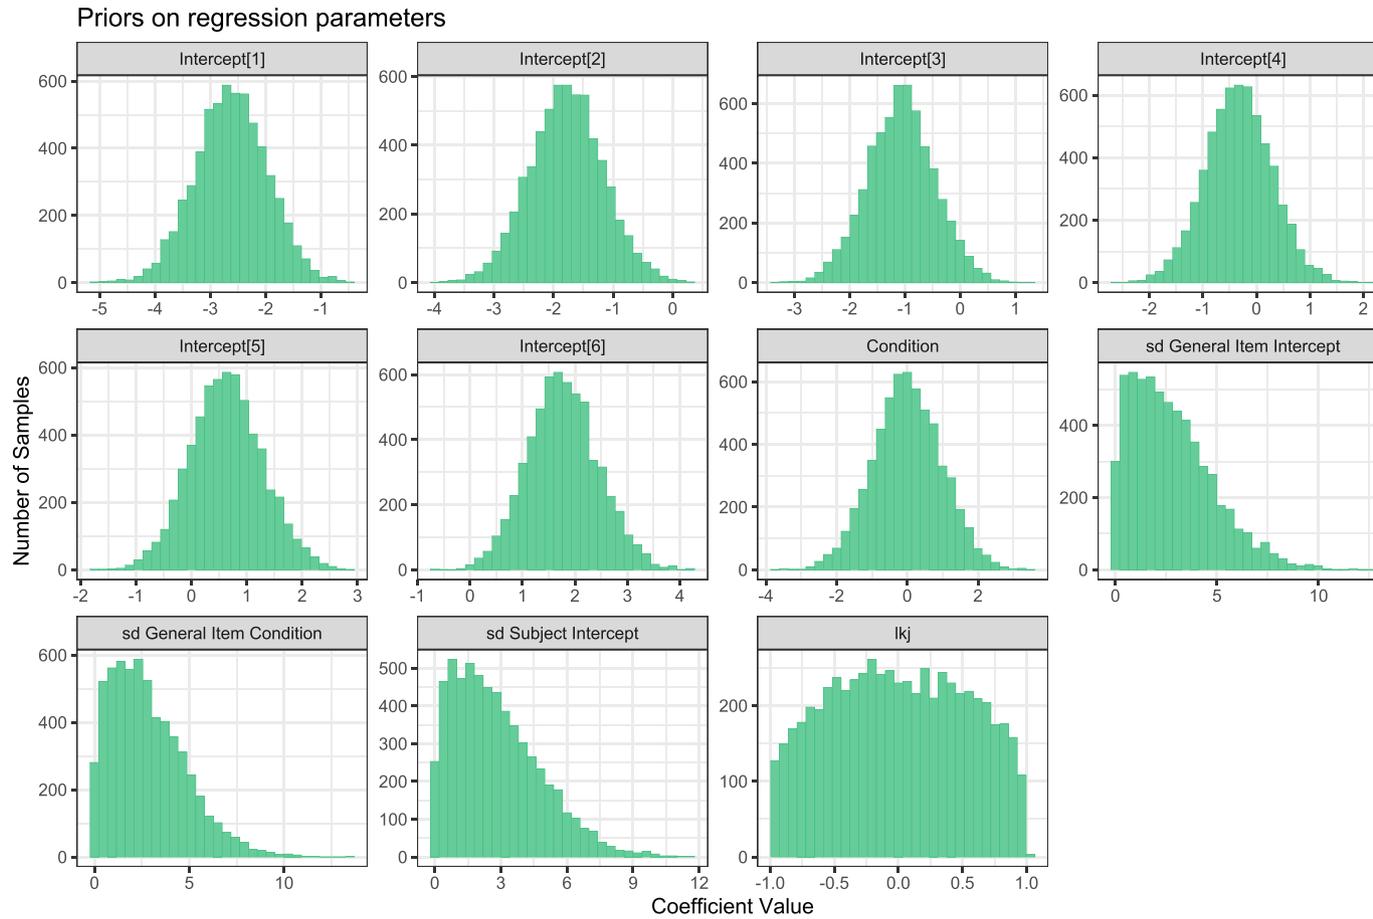


Fig. D1. Priors over model parameters in Model 1 in Experiment 1. All parameters in log-odds except the *LKJ* coefficient, which is a Pearson r .

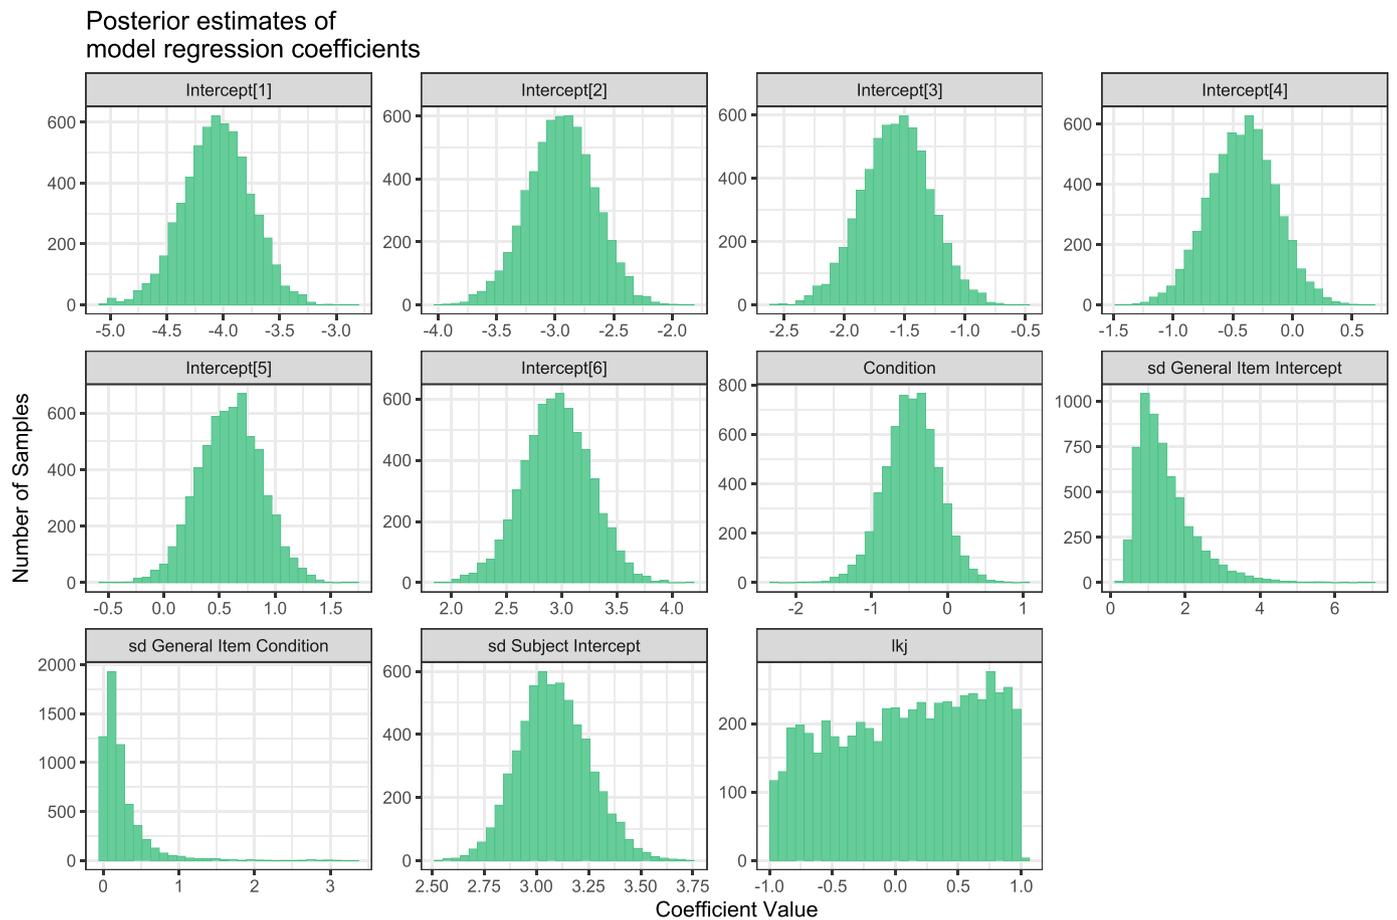


Fig. D2. Posterior model regression coefficients in Model 1 in Experiment 1. All parameters in log-odds except the lkj coefficient, which is a Pearson r .

Summary and Diagnostics. Priors in subsequent models were determined in an analogous way to what we’ve described above – (1) We provided weak priors to guide the model in estimating the most likely Likert point selections (i.e., thresholds) and (2) we regularized β posteriors to provide more conservative effect size estimates.

For this and all subsequent analyses, we performed diagnostic checks on the sampling behavior of our models with, for example, trace plots to ensure Markov chains mixed appropriately (see Fig. D2). Additional diagnostic checks can be found in the RData file at <https://osf.io/uufru/> by using the plot(model name) function.

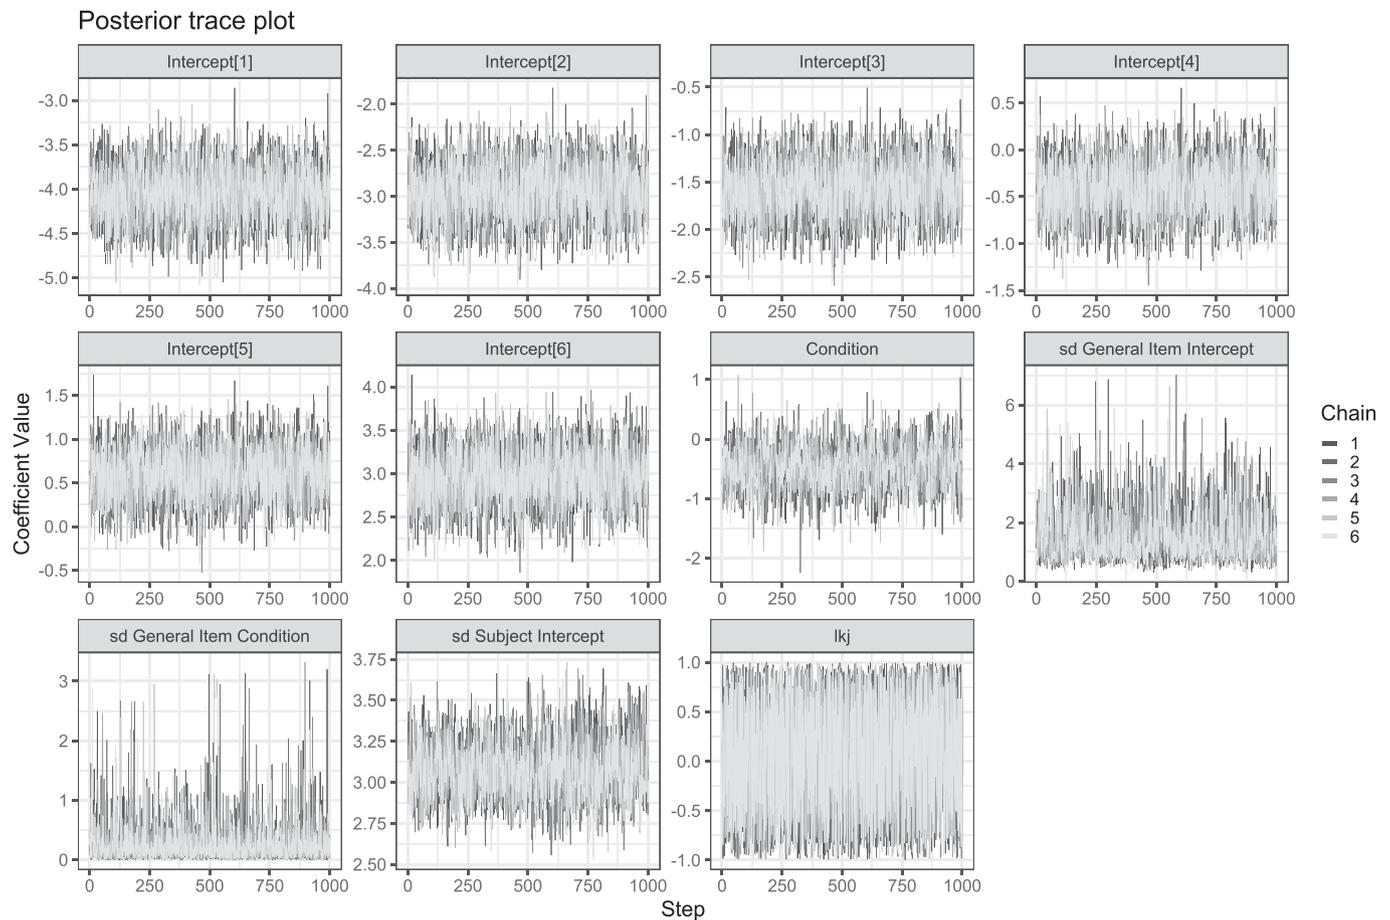


Fig. D3. Trace plot of model posterior regression coefficients in Model 1 in Experiment 1. The plot indicates acceptable chain convergence. All parameters are shown in log-odds except the lkj coefficient, which is a Pearson r coefficient.

Appendix E. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2021.104160>.

References

- Amiot, C. E., & Bastian, B. (2015). Toward a psychology of human-animal relations. *Psychological Bulletin*, *141*(1), 6–47.
- Arlot, S., Celisse, A., et al. (2010). A survey of cross-validation procedures for model selection. *Statistics Surveys*, *4*, 40–79.
- Barkan, R., Ayal, S., Gino, F., & Arieli, D. (2012). The pot calling the kettle black: Distancing response to ethical dissonance. *Journal of Experimental Psychology: General*, *141*(4), 757–773.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.
- Bastian, B., Loughnan, S., Haslam, N., & Radke, H. R. (2012). Don't mind meat? The denial of mind to animals used for human consumption. *Personality and Social Psychology Bulletin*, *38*(2), 247–256.
- Bürkner, P.-C. (2017). Brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*(1), 1–28.
- Bürkner, P.-C., & Charpentier, E. (2018). Modeling monotonic effects of ordinal predictors in regression models (PsyArXiv).
- Bürkner, P.-C., & Vuorre, M. (2019). Ordinal regression models in psychology: A tutorial. *Advances in Methods and Practices in Psychological Science*, *2*(1), 77–101.
- Campbell, R., & Kumar, V. (2012). Moral reasoning on the ground. *Ethics*, *122*(2), 273–312.
- Ciuk, D. J., & Rottman, J. (2021). Moral conviction, emotion, and the influence of episodic versus thematic frames. *Political Communication*. <https://doi.org/10.1080/10584609.2020.1793847>.
- Dancy, J. (1984). On coherence theories of justification: Can an empiricist be a coherentist? *American Philosophical Quarterly*, *21*(4), 359–365.
- Daniels, N. (1979). Wide reflective equilibrium and theory acceptance in ethics. *The Journal of Philosophy*, *76*(5), 256–282.
- Feinberg, M., Kovacheff, C., Teper, R., & Inbar, Y. (2019). Understanding the process of moralization: How eating meat becomes a moral issue. *Journal of Personality and Social Psychology*, *117*(1), 50–72.
- Fessler, D. M., Arguello, A. P., Mekdara, J. M., & Macias, R. (2003). Disgust sensitivity and meat consumption: A test of an emotivist account of moral vegetarianism. *Appetite*, *41*(1), 31–41.
- Festinger, L. (1962). *A theory of cognitive dissonance* (vol. 2). Stanford University Press.
- Gelman, A., & Carlin, J. (2014). Beyond power calculations: Assessing Type-S and Type-M errors. *Perspectives on Psychological Science*, *9*(6), 641–651.
- Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, *25*(4), 524–575.
- Gruen, L. (2017). *The moral status of animals*. Stanford Encyclopedia of Philosophy. Retrieved from <https://plato.stanford.edu/entries/moral-animal/>.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, *108*(4), 814–834.
- Harris, M. (1998). *Good to eat: Riddles of food and culture*. Waveland Press.
- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin*, *142*(11), 1179–1203.
- Horne, Z., Powell, D., & Hummel, J. (2015). A single counterexample leads to moral belief revision. *Cognitive Science*, *39*(8), 1950–1964.
- Horne, Z., Powell, D., & Spino, J. (2014). Belief updating in moral dilemmas. *Review of Philosophy and Psychology*, *4*(4), 705–714.
- Huckfeldt, R., & Sprague, J. (2000). Political consequences of inconsistency: The accessibility and stability of abortion attitudes. *Political Psychology*, *21*(1), 57–79.
- Hussar, K. M., & Harris, P. L. (2010). Children who choose not to eat meat: A study of early moral decision-making. *Social Development*, *19*(3), 627–641.
- Kelly, D. (2011). *Yuck!: The nature and moral significance of disgust*. MIT press.
- Kohlberg, L., & Kramer, R. (1969). Continuities and discontinuities in childhood and adult moral development. *Human Development*, *12*(2), 93–120.

- Kruschke, J. (2014). *Doing bayesian data analysis: A tutorial with r, jags, and stan*. Academic Press.
- 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.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*(11), 2098–2109.
- Loughnan, S., Bastian, B., & Haslam, N. (2014). The psychology of eating animals. *Current Directions in Psychological Science, 23*(2), 104–108.
- Loughnan, S., Haslam, N., & Bastian, B. (2010). The role of meat consumption in the denial of moral status and mind to meat animals. *Appetite, 55*(1), 156–159.
- May, J. (2018). *Regard for reason in the moral mind*. Oxford University Press.
- McElreath, R. (2016). *Statistical rethinking: A Bayesian course with examples in R and Stan*. Boca Raton, FL: CRC Press.
- Miske, O., Schweitzer, N., & Horne, Z. (2019). *What information shapes and shifts people's attitudes about capital punishment?*
- Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., & Wagenmakers, E.-J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review, 23*(1), 103–123.
- Mullen, E., & Skitka, L. J. (2006). Exploring the psychological underpinnings of the moral mandate effect: Motivated reasoning, group differentiation, or anger? *Journal of Personality and Social Psychology, 90*(4), 629–643.
- Newport, F. (2012). *In US, 5% consider themselves vegetarians* (Gallup Wellbeing).
- Nichols, S. (2004). *Sentimental rules: On the natural foundations of moral judgment*. Oxford University Press.
- Norcross, A. (2004). Puppies, pigs, and people: Eating meat and marginal cases. *Philosophical Perspectives, 18*(1), 229–245.
- Norcross, A., et al. (2008). Off her trolley? Frances kamm and the metaphysics of morality. *Utilitas, 20*(1), 65.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science, 349*(6251). <https://doi.org/10.1126/science.aac4716>. Retrieved from.
- Petty, R. E., Wheeler, S. C., & Tormala, Z. L. (2003). Persuasion and attitude change. In T. Millon, & M. J. Lerner (Eds.), *Handbook of psychology* (pp. 353–382). John Wiley & Sons.
- Piazza, J., Ruby, M. B., Loughnan, S., Luong, M., Kulik, J., Watkins, H. M., & Seigerman, M. (2015). Rationalizing meat consumption. The 4Ns. *Appetite, 91*, 114–128.
- Plous, S. (1993). Psychological mechanisms in the human use of animals. *Journal of Social Issues, 49*(1), 11–52.
- Popper, K. (2005). *The logic of scientific discovery*. Routledge.
- 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.
- Ridge, M., & McKeever, S. (2020). Moral particularism and moral generalism. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy (Winter 2020 ed.)*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archive/s/win2020/entries/moral-particularism-generalism/>.
- Rozin, P., Markwith, M., & Stoess, C. (1997). Moralization and becoming a vegetarian: The transformation of preferences into values and the recruitment of disgust. *Psychological Science, 8*(2), 67–73.
- Ruby, M. B. (2012). Vegetarianism. A blossoming field of study. *Appetite, 58*(1), 141–150.
- Sánchez-Meca, J., Marín-Martínez, F., & Chacón-MoscOSO, S. (2003). Effect-size indices for dichotomized outcomes in meta-analysis. *Psychological Methods, 8*(4), 448.
- Schad, D. J., Betancourt, M., & Vasishth, S. (2020). Toward a principled Bayesian workflow in cognitive science. *Psychological Methods*.
- van de Schoot, R., Depaoli, S., King, R., Kramer, B., Kaspers, M., & Tadesse, M. (2021). Bayesian statistics and modelling. *Nature Review Methods Primers, 1*(1).
- Schwitzgebel, E., Cokelet, B., & Singer, P. (2020). Do ethics classes influence student behavior? Case study: Teaching the ethics of eating meat. *Cognition, 203*, 104397.
- Schwitzgebel, E., & Cushman, F. (2015). Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition, 141*, 127–137.
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy & Public Affairs, 229–243*.
- Skitka, L. J. (2010). The psychology of moral conviction. *Social and Personality Psychology Compass, 4*(4), 267–281.
- Slovic, P. (2007). "If I look at the mass I will never act": Psychic numbing and genocide. *Judgment and Decision making, 2*(2), 79–95.
- Twitter. (2021) (n.d.). Retrieved from <https://twitter.com/peta>.
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing, 27*(5), 1413–1432.
- Vosgerau, J., Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2019). 99% impossible: A valid, or falsifiable, internal meta-analysis. *Journal of Experimental Psychology: General, 148*(9), 1628–1639.
- Wilks, M., & Phillips, C. J. (2017). Attitudes to in vitro meat: A survey of potential consumers in the United States. *PLoS One, 12*. e0171904.
- Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science, 12*(6), 1100–1122.