Running head: TOM NETWORK ASSOCIATED WITH METAETHICS

Theory of Mind network activity is associated with metaethical judgment: An item analysis

Jordan Theriault¹, Adam Waytz², Larisa Heiphetz³, Liane Young⁴

¹Northeastern University
Department of Psychology
Boston, MA, 02115
USA

²Northwestern University
Kellogg School of Management
Evanston, IL, 60208
USA

³Columbia University
Department of Psychology
New York, NY, 10027
USA

⁴Boston College
Department of Psychology
Chestnut Hill, MA, 02467
USA

Corresponding Author:
Jordan Theriault
jordan_theriault@northeastern.edu

The theory of mind network (ToMN) is a set of brain regions activated by a variety of social tasks. Recent work has proposed that these associations with ToMN activity may relate to a common underlying computation: processing prediction error in social contexts. The present work presents evidence consistent with this hypothesis, using a fine-grained item analysis to examine the relationship between ToMN activity and variance in stimulus features. We used an existing dataset (consisting of statements about morals, facts, and preferences) to examine the variability in ToMN activity elicited by moral statements, using metaethical judgments (i.e. judgments of how objective/subjective morals are) as a proxy for their predictability/support by social consensus. Study 1 validated expected patterns of behavioral judgments in our stimuli set, and Study 2 associated by-stimulus estimates of metaethical judgment with ToMN activity, showing that ToMN activity was negatively associated with objective morals and positively associated with subjective morals. Whole brain analyses indicated that these associations were strongest in bilateral temporoparietal junction (TPJ). We also observed additional by-stimulus associations with ToMN, including positive associations with the presence of a person (across morals, facts, and preferences), a negative association with agreement (among morals only), and a positive association with mental state inference (in preferences only, across 3 independent measures and behavioral samples). We discuss these findings in the context of recent predictive processing models, and highlight how predictive models may facilitate new perspectives on both metaethics and the nature of distinctions between social domains (e.g. morals vs. preferences).

Keywords:

Theory of Mind; Predictive Coding; Mixed Effects; Item Analysis; Morality; Metaethics

1. Introduction

Theory of mind refers to the ability to represent internal mental states (Premack & Woodruff, 1978). The Theory of Mind network (ToMN) is a set of brain regions that are active during mental inference and social cognition, with core regions in this network—medial prefrontal cortex (MPFC), precuneus (PC), and bilateral temporoparietal junction (RTPJ/LTPJ)—showing activation during a variety of social tasks (for review, see Amodio & Frith, 2006; Schurz et al., 2014, 2017; Van Overwalle, 2009). These tasks include reading comics and stories about beliefs (Ciaramidaro et al., 2007; Dodell-Feder et al., 2011; Fletcher, 1995; Gallagher et al., 2000; Saxe & Kanwisher, 2003; Saxe & Powell, 2006; Young, Dodell-Feder, et al., 2010), reading stories about moral violations (Young, Camprodon, et al., 2010; Young et al., 2007; Young & Saxe, 2009), watching social animations (Blakemore, 2003; Gobbini et al., 2007), taking others' perspectives (Ruby & Decety, 2003; Vogeley et al., 2001), making strategic decisions in economic games (Kircher et al., 2009), inferring personal traits (Harris et al., 2005; Ma, Vandekerckhove, Hoeck, et al., 2012), forming impressions (Baron et al., 2011; Bhanji & Beer, 2013; Cloutier et al., 2011; Ma, Vandekerckhove, Hoeck, et al., 2012; Mende-Siedlecki et al., 2013; Mende-Siedlecki & Todorov, 2016; Mitchell et al., 2005; Park & Young, 2020; Schiller et al., 2009), and even listening to narratives, in which case the similarity of representations in the ToMN reflects shared experiences across listeners (Yeshurun et al., 2017). The ToMN is also differentially activated by domains of social information—e.g. in bilateral TPJ, moral statements elicit more activity than statements about facts and preferences (Theriault et al., 2017; also, see Jenkins & Mitchell, 2010). However, knowing that social tasks and stimuli activate the ToMN can only advance a scientific understanding of the ToMN so far: it tells us that particular tasks and stimuli with particular features activate the network, but it leaves us to triangulate between

24 this noisy pattern of associations to understand why it was produced. Ideally, psychologists and
25 neuroscientists would like to understand a more fundamental problem: what underlying
26 computational process is responsible for the activity observed in these socially-sensitive regions?

27 Answering this question completely is beyond the scope of any one paper, but our aim
28 here is to provide empirical support for one recent and promising hypothesis, which has
29 proposed that ToMN activity (and in particular, activity in TPJ) reflects the updating of
30 predictions in social contexts (Koster-Hale & Saxe, 2013; also, see Kim et al., 2020). That is,
31 assuming that a brain uses prior experience to form and issue predictions about incoming sensory
32 signals (e.g. sights, sounds; Rao & Ballard, 1999), it is hypothesized that the brain performs this
33 same fundamental process in social settings, only at a higher level of abstraction. That is, the
34 brain has been hypothesized to issue predictions about both the observable behaviors enacted by
35 others (as, to an observer, these behaviors are combinations of more basic sensory information)
36 and the latent mental states that motivate behavior (which are not observable, but could be
37 inferred from sensory signals and prior knowledge). Critically, this hypothesis implies that the
38 brain can filter sensory signals efficiently, as predictable signals can be ignored as
39 uninformative, and bottom-up information processing can be limited to encoding *prediction*
40 *error* (i.e. the difference between the sensory signal as predicted, and the sensory signal as
41 actually received).

42 This predictive hypothesis for the ToMN is derived from more general predictive coding
43 frameworks (e.g. Barrett & Simmons, 2015; Chanes & Barrett, 2016; A. Clark, 2013, 2015;
44 Denève & Jardri, 2016; Friston et al., 2016, 2017; Hohwy, 2013; Hutchinson & Barrett, 2019;
45 Joiner et al., 2017; Koster-Hale & Saxe, 2013; Rao & Ballard, 1999; Shadmehr et al., 2010;
46 Spratling, 2017; Van de Cruys et al., 2014), which suggest that this computational process—

issuing predictions, encoding prediction error, and using prediction error (i.e. information) to form new predictions—is the general computation performed at all hierarchical levels of the brain. That is, each cortical area is thought to issue predictions, filter predictable signals, and pass prediction error up the cortical hierarchy. Signals, here, refer to either incoming sensory signals (e.g. light through the retina, pressure on skin) or their multimodal compressions (e.g. the unique combination of sights, sounds, smells, etc. that constitute a human social interaction). These predictive coding frameworks are grounded in principles of data compression (i.e. information theory), where signals are only informative to the extent that they reduce uncertainty, allowing prediction error to serve as a common currency of encoded information (Shannon & Weaver, 1949/1964).

Returning to the context of mental inference and social cognition: if your brain's predictions about someone's words, behaviors, etc. (i.e. predictions about compressed representations of raw sensory signals) are perfectly accurate, then your brain's predictions about that person's latent mental state need not be updated. It has been proposed that encoding prediction error or updating predictions related to latent mental states may elicit BOLD activity in the ToMN (Koster-Hale & Saxe, 2013), meaning that if there is no information to encode, then ToMN activity is expected to be low, relative to cases where unpredictable mental state information must be encoded. Thus, ToMN activity is expected to be relatively lower when signals (observable behavior, words, etc.) are predictable, and ToMN activity is expected to be relatively higher when these signals are unpredictable (Koster-Hale & Saxe, 2013).

Preliminary work has been consistent with this hypothesis (e.g. Schuwerk et al., 2017; for review, see Koster-Hale & Saxe, 2013); however, it has also been constrained by methods that prevent a simultaneous examination of pre-existing social predictions (i.e. predictions formed

outside of the experimental setting), and fine-grained within-subject statistical analyses. Both factors are important if our results are to generalize beyond our experimental context and into more naturalistic settings. If social predictions are not pre-existing (i.e. they are formed in the context of the experiment), then the effect of violating these predictions may be artificially emphasized. Further, by analyzing within-subject variability, we can base inferences on naturally varying features of our stimuli-set (e.g. Westfall et al., 2017; see Section 4.3 for further discussion) as opposed to grounding inference in *a priori* categories, which may homogenize real variance among the examples they contain (e.g. S. A. Gelman & Rhodes, 2012).

These factors introduce some ambiguity into interpretations of prior work. For example, several studies have introduced characters to participants in short vignettes, leading them to form an initial impression and set of expectations. Following these introductions, ToMN activity increased when characters were described as holding contradictory beliefs (Saxe & Wexler, 2005), as engaging in contradictory behaviors (Dungan et al., 2016), or as possessing contradictory traits (Ma, Vandekerckhove, Baetens, et al., 2012; Mende-Siedlecki et al., 2013), relative to initial expectations. However, later work raised the possibility that ToMN activity in these designs may be contingent on the in-lab formation of initial expectations: when characters were not initially familiarized to participants, unpredictable beliefs (e.g. someone believing plants will burst into flame if watered) did not elicit increased ToMN activity (relative to predictable beliefs; Young, Dodell-Feder, et al., 2010). Likewise, in studies using videos of behavior, ToMN activity is contingent on the motivation to engage in mental inference: when actors in a video perform unconventional actions (e.g. switching a light on with their knee when their hands are free, as opposed to switching it on while carrying a heavy load), ToMN activity increases (Brass et al., 2007; de Lange et al., 2008), but this effect fails to replicate when

participants are not explicitly instructed to attend to the actor's intentions (Ampe et al., 2014). Other studies have leveraged pre-existing social predictions by using characters that were already familiar to participants, either as actual friends (Park & Young, 2020) or as known political figures (Cloutier et al., 2011), but within these studies fine-grained analyses of within-subject variability could not be performed (e.g. Park & Young (2020) examined associations with by-subject averages of RTPJ activity, but their design precluded trial-level analyses).

The present work addresses both of these issues. We leverage pre-existing mental inferences, by using moral statements that were pretested to confirm or contradict (to varying degrees) people's preexisting moral beliefs. Further, to leverage analyses of within-subject variability we used item analyses (Bedny et al., 2007; Dodell-Feder et al., 2011; Donnet et al., 2006; Westfall et al., 2017), a powerful (but often overlooked) method for examining by-stimulus variability and its relationship to stimuli features. In particular, unlike by-subject analyses, item analyses allow covariates of interest to be collected in independent samples of participants, allowing us to use larger samples to estimate by-stimulus averages of features of interest. Item analyses can also make use of modern multilevel modeling statistical methods, using information about the structure of the dataset to improve estimates (Baayen et al., 2008; Barr et al., 2013; Judd et al., 2012; Westfall et al., 2014, 2017). To address both issues simultaneously, would require stimuli that leverage pre-existing predictions about beliefs and, at the same time, can be characterized by stimuli features that approximate each stimuli's average predictability across participants. Fortunately, moral statements satisfied both of these constraints, with metaethical judgments serving as a proxy for moral predictability.

1.1. Metaethical variability within the moral domain. In the present work, we performed a secondary analysis of an existing dataset (Theriault et al., 2017), to address the

116 social prediction hypothesis (Koster-Hale & Saxe, 2013). The existing dataset consisted of fMRI
 117 data that was previously collected and analyzed to test differences in ToMN activity elicited by
 118 factual, preferential, and moral statements (Theriault et al., 2017). Moral beliefs are
 119 methodologically suited to our aims¹, as (1) they are important to impression formation (i.e.
 120 individuals are motivated to predict what moral beliefs others hold), meaning *predictions are*
 121 *formed outside the lab*, and (2) there is some consensus across individuals (within a culture) that
 122 some moral beliefs should be endorsed and others opposed, meaning that there is *structured and*
 123 *generalizable (across people) variability in predictions about moral beliefs*. Regarding point (1):
 124 as moral beliefs are formed outside the lab, held strongly, and occupy a central position in
 125 individuals' identities (Heiphetz et al., 2018; Strohminger & Nichols, 2014, 2015), we expected
 126 that participants would be motivated to predict what moral beliefs generic people (i.e. speakers
 127 for whom participants have no prior knowledge; as opposed to known friends or politicians, as in
 128 prior work; Cloutier et al., 2011; Park & Young, 2020) were likely to hold. By contrast,
 129 compared to morals, factual statements contained less information about others' beliefs/mental
 130 states, and so were not expected to elicit high ToMN activity. Similarly, preferences are by
 131 definition more person-specific (Heiphetz et al., 2014), meaning that high-precision predictions
 132 about preferences cannot be made the context of a generic person². We return to the topic of

¹ Moral beliefs are also somewhat unique in the empirical literature examining social predictions, as prior work has generally focused on predictions in the context of impression formation (e.g. Cloutier et al., 2011; Mende-Siedlecki et al., 2013; Park & Young, 2020).

² Of course, the boundaries between morals, facts, and preferences are porous and, despite our experimental operationalization of these categories, they should not be taken to represent natural kinds. For example, people may *prefer* pleasure to pain, but if participants are asked whether “pleasure is better than pain” is a moral belief, a fact, or a preference, they may report a preference for pleasure over pain as more fact-like. For this reason, we verified (in Study 1) how our stimuli were interpreted (as fact-like, moral-like, and preference-like). In the same vein, readers may be able to imagine a preference that should be shared by many or most people—however, that such counterexamples can be generated only underscores that morals, facts, and preferences are socially constructed categories. As the present study is concerned with variability among moral claims, it is enough for our purposes to claim that people are generally more motivated and able to deploy predictions about a generic other's moral beliefs

predictive precision in the general discussion (section 4.2), but for present purposes it is sufficient to say that high-precision predictions carry more potential to generate prediction error, and low-precision predictions carry less, as a precise prediction has a tightened probabilistic distribution and makes deviations from the mean more informative (Feldman & Friston, 2010; Kim et al., 2020; Van de Cruys et al., 2014). Initial analyses were consistent with this hypothesis at the categorical level, with moral statements eliciting greater ToMN activity than either facts or preferences (particularly within PC and bilateral TPJ; Theriault et al., 2017). Preferences also elicited greater ToMN activity than facts in this prior work.

Point (2), that some moral consensus exists (within a culture) and varies across moral beliefs, allowed us to examine variability *within* the moral domain. Among moral statements, some are more predictably endorsed than others, and in a moral context, metaethical judgments act as a proxy for this predictability. Metaethical judgments are judgments about what kind of information a moral statement conveys—e.g. whether it is objective (more fact-like, less preference-like), or subjective (more preference-like, less fact-like). Metaethical judgments vary across stimuli (i.e. some views are considered more or less objective/subjective than others; Beebe, 2014; Goodwin & Darley, 2008, 2012; Heiphetz & Young, 2017; Sarkissian et al., 2011; Theriault et al., 2017; Wright et al., 2013) and are highly associated with social consensus (Ayars & Nichols, in press; Beebe, 2014; Goodwin & Darley, 2012; Heiphetz & Young, 2017). That is, moral statements that elicit widespread agreement (e.g. slavery is wrong) are considered to be more objective (and less subjective) than others (e.g. eating meat is wrong). Study 1 verifies that this metaethical variability exists among moral stimuli in our dataset.

compared to their preferences. We are *not* claiming that people are completely unmotivated to predict (or are incapable of predicting) the preferences of generic others.

Putting this into the context of prediction: when someone makes a moral assertion that is judged by others (on average) as objective, the assertion should be predictable on the basis of social consensus (i.e. most people would endorse it). Further, in the context of information theory (Shannon & Weaver, 1949/1964, discussed above)—where signals are only informative to the extent that they reduce uncertainty—objective moral statements may actually communicate *less information*, as they are statements that most people would already be predicted to believe. By contrast, when someone makes a moral assertion that is judged as subjective, the assertion should be less predictable on the basis of social consensus (i.e. most people would not endorse it). Again, in the context of information theory, subjective moral statements communicate *more information*, as they are statements that most people would not be predicted to believe. In the present work, we exploited this theoretically-grounded equivalence between objectivity/subjectivity and information in our existing dataset, examining the relationship between ToMN activity and judgments of moral objectivity in a diverse sample of stimuli (Theriault et al., 2017).

1.2. Present Work. The present work used an existing dataset (Theriault et al., 2017) to examine the relationship between by-stimulus ToMN activity and metaethical judgments (e.g. how objective, or ‘fact-like’, moral claims are; or alternatively, how subjective, or ‘preference-like’, moral claims are). As multiple behavioral measures could not be collected in the scanner (agreement was collected on a 4-point scale in the scanner, but not used in any analysis), we examined the relationship between ToMN activity and by-stimulus ratings collected in independent online samples. Study 1 consisted of an online sample (N = 49) where participants read 72 statements (each designed to be read as a statement about facts, morals, or preferences), and rated each on the extent that they agreed/disagreed with it, and the extent that it was “about

facts”, “about morality”, and “about preferences”. This online sample accomplished three goals. First, it validated our stimuli design (confirming that our facts were more fact-like than they were moral-like or preference-like). Second, consistent with prior work (Ayars & Nichols, in press; Beebe, 2014; Goodwin & Darley, 2012; Heiphetz & Young, 2017), comparisons between subgroups confirmed that moral statements supported by a social consensus (positive-consensus moral statements) were perceived as more objective than moral statements that were opposed to, or ambiguous with respect to, a social consensus (negative-consensus/no-consensus). Finally, Study 1 provided estimates of by-stimulus metaethical judgments to be used in Study 2: a maximal mixed effects models was fitted to the subjects and stimuli of the online sample (Barr et al., 2013), and by-stimulus estimates of metaethical judgments were extracted (best linear unbiased predictors, i.e. BLUPs; Baayen et al., 2008).

Study 2 examined the association between by-stimulus metaethical judgments and ToMN activity in DMPFC, VMPFC, PC, and bilateral TPJ regions of interest (ROIs). We used maximal mixed effects models to estimate by-stimulus ToMN activity from a sample of 25 participants and 72 stimuli (Baayen et al., 2008; Barr et al., 2013; Westfall et al., 2017), then compared these estimates with by-stimulus estimates of item features, extracted from Study 1. We also explored additional by-stimulus relationships with item feature estimates extracted from seven other online samples, measuring judgments of valence and arousal (N = 42), mental imagery (N = 46), whether a person was present in the scenario (N = 48), and whether a scenario evoked mental state information generally (N = 48), about others’ mental states (N = 44), or about one’s own mental states (N = 46). For our core analysis, comparing ToMN activity and by-stimulus estimates of metaethical judgment, we controlled for a variety of syntactic and semantic features that may act as confounds. Overall, our findings demonstrate that the ToMN is responsive to

variance among moral statements: more objective moral claims elicit less ToMN activity, and more subjective moral claims elicit greater ToMN activity.

As the present work was a secondary analysis of an existing dataset, we have tried to remain appropriately conservative in our statistical analysis, testing specific relationships (corrected for multiple comparisons) only when justified by omnibus testing (see Tables S7, S9, & S10, in the online supplemental materials). These steps occasionally lead to non-traditional forms of reporting analyses (e.g. grouping ToMN ROIs and analyzing VMPFC separately in Study 2), but these steps were taken to ensure that generalizations were conservative and justified.

2. Study 1

Study 1 validated our stimulus set in an online sample, confirming that moral statements supported by a social consensus were perceived as more objective (Ayars & Nichols, in press; Beebe, 2014; Goodwin & Darley, 2012; Heiphetz & Young, 2017). The stimulus set consisted of 72 statements about facts, morals, and preferences, and were designed to fit within consensus subcategories, eliciting either *positive-consensus* (where most people would agree), *negative consensus* (where most people would disagree), or *no-consensus* (where neither agreement nor disagreement was strong; Figure 1; also see 2.1.2). Using an independent online sample allowed us to extract by-stimulus estimates of metaethical judgments in a sample approximately twice the size of the fMRI sample. To measure metaethical judgment, we asked participants to rate each statement on the extent that it was about facts, about morals, and/or about preferences (Theriault et al., 2017). This method has several advantages: (a) it validates stimulus conditions, as facts should be rated as most fact-like, morals as most moral-like, and preferences as most preference-like; and b) it avoids artificially imposing relationships among ratings (unlike bipolar methods,

where the design requires that for a statement to be more fact-like, it must necessarily be less preference-like). Our method allows any correlations among ratings to emerge independently, without imposing them by design.

2.1. Method

2.1.1. Participants. Participants were recruited online using Amazon Mechanical Turk (AMT) at an approximate rate of \$6/hour, in line with standard AMT compensation rates. The final sample consisted of 49 adults (25 female, 23 male, 1 unspecified; $M_{\text{Age}} = 33.5$ years, $SD_{\text{Age}} = 10.7$ years, $\text{Range}_{\text{Age}} = 19\text{--}59$ years), after excluding two participants for failing an attention check that asked them to describe any statement they had read. The majority of our sample had either entered or completed college/university (maximum educational attainment of high school = 16.3%; some college/university = 36.7%; completed college/university = 38.8%; completed graduate degree = 8.2%). The majority of the sample was Caucasian (White/Caucasian = 83.7%; Black/African American = 6.1%; Asian = 10.2%; Pacific Islander = 2.0%; Other = 2.0%), and Non-Hispanic (Hispanic = 8.2%; Non-Hispanic = 91.8%). The Boston College Institutional Review Board approved studies 1 and 2, and each participant provided consent before beginning.

2.1.2. Procedure. Participants read a series of statements (e.g. “It is irresponsible for airlines to risk the safety of their passengers”; see Appendix A for all statements), and, for each, rated a) their agreement (“To what extent do you disagree/agree; 1–7, “completely disagree”—“completely agree”), and b) the extent that the statement was about facts, about morals, *and* about preferences (*Rating-type*: fact-like/moral-like/preference-like; “To what degree is this statement about ... [facts, morality, preferences]”; 1–7, “not at all”—“completely”). The order of rating-types was counterbalanced across participants. Participants were instructed that they would complete a “statements task”, where they would “read short statements and decide

whether you agree or disagree with them". Statements were designed to be interpreted as either facts, morals, or preferences, and were evenly divided between categories ($n_{\text{Fact}} = 24$, $n_{\text{Moral}} = 24$, $n_{\text{Preference}} = 24$), but participants were not explicitly alerted to this element of design. Each category included three consensus subcategories: a) *positive-consensus*, where most people would agree with the statement ($n = 6$); b) *negative-consensus*, where most people would disagree ($n = 6$); and c) *no-consensus*, where there would be no strong positive or negative consensus ($n = 12$). No-consensus statements (as opposed to controversial statements) were used because the feature of interest was social consensus; in other words, no-consensus statements were intended to elicit a unipolar, non-skewed distribution of agreement. By contrast, controversial statements (e.g. "abortion is wrong") would presumably produce a bimodal distribution of agreement, introducing strong individual differences that could decrease the power of our item analyses. The no-consensus subcategory was also larger relative to other subcategories on account of an uninformative distinction that was irrelevant to the final design: six no-consensus facts were true and six were false. Statements did not contain any mental state markers (e.g. "She thinks," "He believes").

	Positive-consensus	No-consensus	Negative-consensus
Fact	Airplanes have wings that enable the plane to lift upwards.	The very first waffle cone was invented in Chicago, Illinois, at a state fair.	Cockroaches are a type of cold-blooded reptile related to snakes.
Moral	It is irresponsible for airlines to risk the safety of their passengers	It is unethical for businesses to promote sugary products to children.	It is wrong to harm cockroaches just because humans find them disgusting.
Preference	Going through airport security is an unpleasant experience.	Any ice cream flavor tastes better when served in a crunchy waffle cone.	Cockroaches are delicious to eat because of their hard and crunchy shell.

Figure 1. Sample stimuli. Statements varied in content (fact/moral/preference) and agreement (positive-consensus/no-consensus/negative-consensus). See Appendix A for the full text of all stimuli.

2.1.3. Statistical methods. Studies 1 and 2 used mixed effects analyses to model crossed by-subject and by-stimulus random effects (Baayen et al., 2008; Judd et al., 2012; Westfall et al., 2014). In traditional models (e.g. ANOVA) these two sources of variance cannot be modeled simultaneously, meaning that we would be forced to average across stimuli (or across participants), and limit our conclusions to the *exact* stimuli (or participants) that were tested (Baayen et al., 2008; H. H. Clark, 1973; Judd et al., 2012; Westfall et al., 2017). Mixed effects analysis also allow for the estimation of BLUPs (best linear unbiased predictors; Baayen, 2008; Baayen et al., 2008). BLUPs are by-stimulus estimates of metaethical judgments for each stimuli, and are preferable to simple by-stimulus averages for two reasons: a) by-stimulus BLUPs are independent from by-subject variance, meaning that estimates were specific to the scenarios (and could be compared with by-stimulus estimates of ToMN activity in Study 2); and b) by-stimulus BLUPs incorporate the sample distribution into the estimate (i.e. they are semi-pooled estimates; A. Gelman et al., 2012), meaning that they anticipate regression to the mean and mitigate against outliers. Analysis was conducted in R (R Core Team, 2016), using the *lme4* package (Bates et al., 2015), and *p* values for fixed effects were calculated using the Satterthwaite approximation of degrees of freedom, implemented in the *lmerTest* package (Kuznetsova et al., 2017).

2.1.4 Data and software sharing. De-identified raw behavioral data and code to reproduce all analyses and figures are available at <https://osf.io/cx4dp/>.

2.2. Results

2.2.1. Validating agreement sub-categories. Agreement ratings were fit with a maximal mixed effects model. As fixed effects, the model included main effects of category (fact/moral/preference), consensus (positive-/no-/negative-consensus), and their interaction. As random effects across subjects, the model included random intercepts, and random slopes for category, consensus, and their interaction. As random effects across stimuli, random intercepts were included. Condition means are presented in Table S1 of the online supplemental materials.

```
lmer(agreement ~ 1 + category*consensus
      + (1 + category*consensus | subject)
      + (1 | stimuli))
```

We observed some differences in agreement across categories, but critically, differences among consensus sub-categories were consistent with our design. Main effects were significant for both category, $F(2, 72.15) = 7.35, p = .001$, and consensus, $F(2, 55.11) = 55.1, p < .001$, but not their interaction, $F(4, 66.5) = 0.35, p = .843$. In follow-up contrasts, agreement was greater for facts, $z = 3.27, p = .003$, and morals, $z = 3.33, p = .003$, relative to preferences, whereas agreement did not differ between facts and morals, $z = 0.01, p = .999$ (p values were corrected for 3 comparisons; $\alpha_{\text{familywise}} = .05$; single-step method; multcomp package Hothorn et al., 2008). Differences among consensus sub-categories were consistent with our design: agreement was greater for positive-consensus statements, relative to no-consensus statements, $z = 6.48, p < .001$, and negative-consensus statements, $z = 10.26, p < .001$, and agreement was greater for no-consensus statements relative to negative-consensus statements, $z = 6.80, p < .001$. Although preferences elicited less agreement in general, there was no significant interaction between category and consensus, meaning that differences between consensus sub-categories were comparable across facts, morals, and preferences.

2.2.2. Examining fact-/moral-/preference-like ratings across consensus sub-categories. An initial, maximal mixed effects model failed to converge using restricted maximum likelihood (REML) estimation (after 10,000 iterations):

```
lmer(rating ~ 1 + rating_type*category*consensus
      + (1 + rating_type*category*consensus | subject)
      + (1 + rating_type | stimuli))
```

Given this, we simplified the model, removing consensus sub-categories from by-subject random effects. As fixed effects, the model included main effects of rating type (fact-/moral-/preference-like), category (fact/moral/preference), consensus (positive-/no-/negative-consensus), and all interactions. As random effects across subjects, the model included random intercepts, and random slopes for rating type, category, and their interaction. As random effects across stimuli, the model included random intercepts and random slopes for rating type.

```
lmer(rating ~ 1 + rating_type*category*consensus
      + (1 + rating_type*category | subject)
      + (1 + rating_type | stimuli))
```

Consistent with prior work (Ayars & Nichols, in press; Beebe, 2014; Goodwin & Darley, 2012; Heiphetz & Young, 2017), positive-consensus moral statements were perceived as more objective (fact-like) than other moral statements not supported by a social consensus. The model produced a significant 3-way interaction, $F(8, 63.0) = 4.06, p < .001$ (for condition means see Table S1 of the online supplemental materials). Contrasts compared consensus sub-categories within each category x rating-type grouping (p values corrected for 27 comparisons; $\alpha_{\text{familywise}} = .05$; single-step method). Among morals, positive-consensus statements were perceived as more fact-like than no-consensus, $z = 5.70, p < .001$, and negative-consensus statements, $z = 6.26, p < .001$. Also, among facts, positive-consensus and no-consensus statements were perceived as more fact-like than negative-consensus statements: $z = 4.72, p < .001$, and $z = 5.43, p < .001$,

respectively (although these differences within facts did not replicate in Study 2; see section 3.2.1). Among preferences, there were no significant differences between consensus categories (Figure 2b).

Figure 2. Study 1 metaethical judgments. Participants rated each scenario on the extent that it was fact-like, moral-like, and preference-like (1–7; “not at all” – “completely”). (a) Collapsing across consensus subcategories, ratings were consistent with our *a priori* categories: facts were largely fact-like (left), preferences were largely preference-like (right), and morals were largely moral-like (center). Moral statements were also perceived as largely preference-like (a pattern explored further in Theriault et al., 2017). (b) Across consensus subcategories, positive-consensus morals were perceived as more fact-like than no-consensus or negative-consensus morals (center; blue). Note that variance across items was markedly greater for moral statements than for facts or preferences (dots represent item averages for each rating; 72 stimuli x 3 ratings). Error bars represent 95% confidence interval. For condition means, see Table S1 of the online supplemental materials.

2.3. Discussion

Consistent with prior work (Ayars & Nichols, in press; Beebe, 2014; Goodwin & Darley, 2012; Heiphetz & Young, 2017), participants rated positive-consensus moral statements (i.e. moral statements supported by a social consensus) as more fact-like (i.e. more objective) than no-consensus and negative-consensus moral statements. Study 2 examined whether variability in moral objectivity was related to activity in the ToMN, given the hypothesis that regions in this network are associated with processing prediction error in socially relevant contexts (Koster-Hale & Saxe, 2013). According to this hypothesis, moral statements that are predictable on the basis of social consensus (i.e. statements rated as objective) should elicit less ToMN activity, and moral statements that run counter to the social consensus (i.e. statements rated as less objective) should elicit greater ToMN activity.

3. Study 2

Study 2 examined the relationship between the perceived objectivity of moral statements, and activity evoked in ToMN ROI, including right/left temporoparietal junction (R/LTPJ), dorsal-/ventro-medial prefrontal cortex (DMPFC/VMPFC), and precuneus (PC). We identified ToMN ROIs using a functional localizer that has been frequently used in prior work (contrasting false beliefs with false photographs), allowing us to compare our present findings with prior work suggesting these same regions encode mental state information (Dodell-Feder et al., 2011). By-stimulus independent ratings of moral objectivity were collected and modeled in Study 1, and compared to by-stimulus estimates of ToMN activity that were collected and modeled in Study 2. In addition, other independent ratings of item features (e.g. the extent stimuli evoked mental state inferences) were collected, modeled, and explored as well.

3.1 Method

3.1.1. Participants. Participants were a community sample, recruited through an online posting and paid \$65. The final sample consisted of 25 right-handed adults (12 female, 12 male, 1 unspecified; $M_{\text{age}} = 27.0$ years, $SD_{\text{age}} = 5.2$ years; $\text{Range}_{\text{Age}} = 18\text{--}35$ years). Two more participants were recruited but not analyzed due to excessive movement, identified during spatial preprocessing. Of these 25 participants, two completed only a subset of the scan session runs: one completed five of six runs due to experimenter error, and for the other a movement artifact rendered only the first three runs usable. We were unable to collect post-scan ratings for one of the 25 participants. The majority of our sample had either entered or completed college/university (maximum educational attainment of high school = 4%; some college/university = 28%; completed college/university = 52%; completed graduate degree = 12%; no response = 4%). The majority of the sample was Caucasian (White/Caucasian = 56%; Black/African American = 20%; Asian = 20%; no response = 4%), and Non-Hispanic (Hispanic = 4%; Non-Hispanic = 92%; no response = 4%). All participants were native English speakers with no reported history of learning disabilities, previous psychiatric or neurological disorders, or a history of drug or alcohol abuse.

3.1.2. Procedure. Participants completed the study in a single session. Twenty participated at Harvard University's Center for Brain Science Neuroimaging Facility, and five at the Massachusetts Institute of Technology's Martinos Imaging Center. Scanning parameters and equipment were identical between sites (see 3.1.4). In the scanner, participants read statements and rated their agreement with each (ratings were consistent with consensus subcategories; see Table S2). Statements were shown across six runs (12 per run; items were randomized; conditions were counterbalanced to appear equally in each run). Participants read each statement (6 s), rated their agreement (4 s), and waited during fixation (12 s). Agreement was provided

with a button box (1–4; “Strongly Agree”–“Strongly Disagree”). A thumb press indicated “Don’t Know”, which was coded as an empty cell. This option was provided to avoid confusion, particularly for no-consensus facts where the answer was generally unknown (across our complete sample, 68.5% of “don’t know” responses were for no-consensus facts, followed by 7.7% for no-consensus preferences). Stimuli were presented in white text on a black background using a projector, viewable through a mirror mounted on the head coil. The experimental protocol was run on an Apple Macbook Pro using Matlab 7.7.0 (R2008b) with Psychophysics Toolbox. Each experimental run was 4 min 52 s long, totaling 29 min 12 s across six runs. The in-scanner experiment was preceded by a structural scan (6 min 3 s) and a functional localizer (two 4 min 46 s runs; Dodell-Feder et al., 2011; see 3.1.5). The total scan time was 68 min 8 s due to a second study not reported here involving responses to moral dilemmas (29 min 12 s); runs for both studies were interleaved, so that stimuli in the present work were equally likely to appear early or late in the session, across participants. Post-scan, participants rated all statements (i.e. on the extent to which each was fact-/moral-/preference-like) on an Apple Macbook Pro and completed a brief demographics questionnaire.

3.1.3. Stimuli and measures. Stimuli were identical to those described in Study 1 (see Appendix A). Ratings of additional item features were collected in separate online samples. These included questions used in a prior item analysis of the ToMN (Dodell-Feder et al., 2011) as well as measures of arousal and valence (Kron et al., 2013; see Appendix B). In these separate online samples, participants were asked one of the following questions: *Mental Imagery* ($n = 46$; “To what extent did you picture or imagine what the statement described as you read?”; 1–7–; “Very Little”–“Very much”), *Person Present* ($n = 48$; “Does this statement mention people or a person?”; 0–1; “No”–“Yes”), *Valence* ($n = 42$; the difference between 8-point positive and

negative unipolar scales; Kron et al., 2013), *Arousal* ($n = 42$; the sum of both 8-point positive and negative unipolar scales; Kron et al., 2013); and *Mental States* ($n = 48$; “To what extent did this statement make you think about someone’s experiences, thoughts, beliefs, and/or desires?”; 1–7; “Very Little”–“Very Much”). Given that prior work has established the sensitivity of the ToMN to mental state information (Saxe & Kanwisher, 2003), and given that this *Mental States* measure was ambiguous with respect to whose mental states should be considered (either your own, or the mental states of others), we asked an additional sample of participants two more specific questions: *Mental States (of Others)*, ($n = 44$; “To what extent did this statement make you think about the experiences, thoughts, beliefs, and/or desires OF OTHER PEOPLE?; 1–7; “Very Little”–“Very Much”) and *Mental States (of Self)*, ($n = 46$; “To what extent did this statement make you think about YOUR OWN experiences, thoughts, beliefs, and/or desires?”; 1–7; “Very Little”–“Very Much”).

To ensure that effects were not driven by semantic/syntactic differences across stimuli, several item characteristics were collected using Coh-Metrix 3.0 (Graesser et al., 2004; McNamara et al., 2014). These included features such as word length, reading ease, noun concreteness, familiarity, and imageability, among others (see Appendix B).

3.1.4. fMRI imaging and analysis. Scanning was performed using a 3.0 T Siemens Tim Trio MRI scanner (Siemens Medical Solutions, Erlangen, Germany) and a 12-channel head coil at the Center for Brain Science Neuroimaging Facility at Harvard University and at the Massachusetts Institute of Technology’s Martinos Imaging Center. Thirty-six slices with 3mm isotropic voxels, with a 0.54mm gap between slices to allow for full brain coverage, were collected using gradient-echo planar imaging (TR = 2000 ms, TE = 30 ms, flip angle = 90°, FOV = 216 x 216 mm; interleaved acquisition). Anatomical data were collected with T1-weighted

multi-echo magnetization prepared rapid acquisition gradient echo image (MEMPRAGE) sequences (TR = 2530 ms, TE = 1.64 ms, FA = 7°, 1mm isotropic voxels, 0.5mm gap between slices, FOV = 256 x 256 mm). Data processing and analysis were performed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>) and in-house Matlab modeling scripts (https://github.com/lypsychlab/younglab_scripts). The data were motion-corrected, realigned, normalized onto a common brain space (Montreal Neurological Institute, MNI), spatially smoothed using a Gaussian filter (full-width half-maximum = 5 mm kernel), and high-pass filtered (128 Hz).

3.1.5. ToMN localizer task. An independent functional localizer task identified ToMN ROIs (Dodell-Feder et al., 2011). The task consisted of ten stories about mental states (*false-belief*) and ten about physical representations (*false-photograph*), presented across two runs. Stories were matched in complexity across conditions; see <http://saxelab.mit.edu/superloc.php> for the complete set. Each story appeared (10 s) and was followed by a statement about it, rated true or false (4 s). Typically, to increase power, this contrast is used to select ROIs individually for each participant. However, this approach also means that ROI coordinates cannot be reported in normalized space. Alternatively, we could select ROIs using the peak voxels of a whole brain random effects contrast (belief > photograph) across all participants. Both approaches returned the same pattern of results (with respect to the significance of relationships between ToMN activity and metaethical judgments), and so in the interest of providing replicable coordinates we used the latter approach, defining each ROI as a 9mm-radius sphere around the peak voxel (for coordinates see Table S3 of the online supplemental materials). The localizer contrast was thresholded at $p < .001$ (Woo et al., 2014), and $k = 10$ (to remain consistent with thresholds used

in prior work; Dodell-Feder et al., 2011; see Table S3 for k values, and additional clusters identified by the contrast).

3.1.6. ROI analysis. BOLD activity for each functional ROI was estimated using a boxcar regressor, beginning with the appearance of the text, and ending after the agreement rating (10 s total). The time-window was adjusted for hemodynamic lag so that data were collected at 4–14 seconds from onset (Dodell-Feder et al., 2011). To model activity in each ROI, we transformed BOLD activity at each time point of the experimental task into percent signal change ($PSC = \text{raw BOLD magnitude for (condition - fixation) / fixation}$), centering each run at mean PSC.

3.1.7. Whole brain correlation analysis. Whole-brain analyses were performed by first estimating beta maps for each item and then correlating beta maps with estimates of metaethical judgments (derived from Study 1, see 3.1.8). For each subject, three models correlated beta estimates with fact-like, moral-like, and preference-like ratings. Subject-level beta maps of each correlation were entered into separate second-level analyses across subjects. Each second-level contrast was cluster-corrected by permutation (5000 samples) to achieve a familywise error rate of $\alpha = .05$, thresholding voxels at $p < .001$ (as recommended by Woo et al., 2014). Permutation tests were performed using SnPM 13 (<http://warwick.ac.uk/snmp>; Nichols & Holmes, 2002).

3.1.8. Statistical methods. As in Study 1, mixed effects analyses were used to model behavioral responses and PSC. Model specification and simplification is described below (3.2.2). For ToMN ROIs, BLUPs (best linear unbiased predictors) were extracted from the fitted models and compared to behavioral BLUPs extracted from the online sample in Study 1 (see Table S4 of the online supplemental materials) and additional online samples (see section 3.1.3, and Appendix A). The advantages of mixed effects models are described in detail in section 2.1.3.

We began all modeling by including all random effects (Barr et al., 2013), but if this model failed to converge using REML estimation, then correlations between random effects were temporarily removed, the model was inspected, and variance components with zero unique variance were removed (Bates et al., 2015). Any exceptions to this procedure are noted in the results section, and full models are reported in the text in formula syntax.

3.1.9 Data and software sharing. De-identified ROI timecourse data and code to reproduce all analyses and figures are available at <https://osf.io/cx4dp/>, as well as raw fMRI data and by-stimulus item estimates in standardized BIDS format (Gorgolewski et al., 2016). Whole-brain unthresholded T maps, corresponding to analyses in Section 3.2.5 and Figure 4, are available at <https://neurovault.org/collections/ZZMRPKDV/>.

3.2. Results

3.2.1. Replicating Study 1 behavioral results. First, we modeled metaethical judgments in our Study 2 sample to ensure that the patterns were consistent with those observed in Study 1. The maximal mixed effects model used to model behavioral data in Study 1 (2.2.2) did not provide a convergent solution using REML estimation in the Study 2 behavioral data, most likely because of Study 2's smaller sample of participants. There were no parameters which could clearly be dropped (i.e. parameters with zero unique variance), so by-stimulus random effects were dropped instead, as we simply wanted to confirm that ratings were similar to those observed in Study 1. The resulting model included, as fixed effects, main effects of rating type (fact-/moral-/preference-like), category (fact/moral/preference), consensus (positive-/no-/negative-consensus), and all interactions. As random effects across subjects, the model included random intercepts, and random slopes for rating type, category, and their interaction.

```
lmer(rating ~ 1 + rating_type*category*consensus
      + (1 + rating_type*category | subject))
```

We performed contrasts for all comparisons of interest (p values corrected for 27 comparisons; $\alpha_{\text{familywise}} = .05$; single-step method). As in Study 1, participants rated positive-consensus morals as more fact-like than both no-consensus morals, $z = 6.82$, $p < .001$, and negative-consensus morals, $z = 7.92$, $p < .001$, and no significant difference emerged between fact-like ratings for no-consensus and negative-consensus morals, $z = 2.33$, $p = .378$. Unlike in Study 1, there was no significant difference in fact-like ratings within facts across consensus sub-categories (for condition means see Table S5 of the online supplemental materials).

3.2.2. Model fitting and by-stimulus estimates of ToMN activity. Initially, a maximal mixed effects model, predicting PSC, was fit across all functional ROIs:

```
lmer(PSC ~ 1 + ROI*category
      + (1 + ROI*category | subject)
      + (1 + ROI | stimuli))
```

However, this maximal model failed to converge using REML and had to be simplified (see section 3.1.8). The final simplified model included, as fixed effects, main effects of ROI (DMPFC /VMPFC/ PC/ RTPJ/ LTPJ), category (fact/moral/preference), and all interactions. As random effects across subjects, the model included random intercepts, and random slopes for ROI, category, and the interactions of VMPFC and LTPJ with the moral category. As random effects across stimuli, the model included random intercepts, and random slopes for VMPFC.

```
lmer(PSC ~ 1 + ROI*category
      + (1 + ROI + category + moral*VMPFC + moral*LTPJ | subject)
      + (1 + VMPFC | stimuli))
```

These simplifications from the maximal model are informative with respect to relations among ROIs (for model details, see Table S6 of the online supplemental materials). For by-stimulus random effects, we observed high correlations between DMPFC (i.e. intercept), PC, RTPJ, and LTPJ. These correlations preclude calculating unique variance terms for each ROI,

but, at the same time, they suggest that the response of these regions cannot be distinguished within our set of stimuli. Thus, the mixed effects analysis provided a data driven rationale for a more conservative analysis, treating these regions as a network and avoiding overfitting the data to every ROI. At the same time, as by-stimulus variance for VMPFC can be estimated separately, the data licensed separate analyses for this region. As all analyses below concern the by-stimulus estimates of ToMN activity, by-stimulus estimates of DMPFC, PC, RTPJ, and LTPJ are averaged and collectively referred to as the ToMN. Given this, ROI interactions reported below examine differences between the ToMN, and VMPFC.

3.2.3. Among moral statements, ToMN activity is positively associated with preference-like ratings, and negatively associated with fact-like/moral-like ratings. Our core question was whether ToMN activity was related to by-stimulus ratings of moral objectivity. By-stimulus estimates of metaethical ratings (fact-like, moral-like, and preference-like ratings) were extracted from Study 1 (section 2.2.2), and, separately, by-stimulus estimates of ToMN activity (PSC; centered and normalized) were extracted from the Study 2 fMRI data (section 3.2.2). A linear model predicted by-stimulus metaethical judgments on the basis of ToMN activity (for model details, see Table S7 of the online supplemental materials), including main effects and interactions of PSC (estimated from the model in 3.2.2), category (fact/moral/preference), and rating type (fact-/moral-/preference-like), as well as interactions between these terms and ROI (ToMN/VMPFC).

```
lm(rating ~ 1 + PSC*category*rating_type +
    ROI:PSC + ROI:PSC:category + ROI:PSC:rating_type +
    ROI:PSC:rating_type:category)
```

The association between metaethical judgments and ToMN activity differed between categories (facts, morals, and preferences): morals showed a distinct pattern of association. We

observed a 4-way interaction³, $F(4, 405) = 5.73, p < .001$ between PSC, category, rating-type, and ROI. Follow-up ANOVAs identified significant 4-way interactions between morals and facts, $F(2, 270) = 7.42, p < .001$, and between morals and preferences, $F(2, 270) = 6.64, p = .002$, but not between preferences and facts, $F(2, 270) = 0.01, p = .986$.

Within morals, ToMN activity was related to metaethical judgments: increased ToMN activity was associated with increases in preference-like ratings, and decreases in both fact-like and moral-like ratings. Following up a 3-way interaction (within morals) between PSC, ROI, and rating type $F(2, 135) = 6.03, p = .003$, further 3-way interactions distinguishing between rating types demonstrated distinctions between preference-like and fact-like ratings, $F(1, 90) = 13.8, p < .001$, between preference-like and moral-like ratings, $F(1, 90) = 5.55, p = .022$, but not between fact-like and moral-like ratings, $F(1, 90) = 0.81, p = .371$. Thus, among moral statements, BOLD activity in ToMN and VMPFC shows one relationship with preference-like ratings, and another for moral-like and fact-like ratings. These relationships also differed slightly between ROIs, as 2-way interactions between PSC and ROI were significant among preference-like ratings, $F(1, 45) = 4.57, p = .038$, and fact-/moral-like ratings, $F(1, 92) = 8.80, p = .004$.

Contrasts tested the strength of each relationship, among morals, between ToMN/VMPFC activity and metaethical judgment. ToMN activity was negatively related to fact-/moral-like ratings, $B = -1.01, \beta = -0.46, t(140) = 4.32, p < .001$, and positively related to preference-like ratings, $B = 0.94, \beta = 0.43, t(140) = 2.85, p = .020$. VMPFC activity was also negatively related to fact-/moral-like ratings, $B = -0.28, \beta = -0.13, t(140) = 3.18, p = .007$, and

³ A sensitivity analysis (estimated by simulation using the *simr* package; Green & MacLeod, 2016) indicated that this 4-way interaction between PSC, category, rating-type, and ROI could detect a minimum effect size 35% below the effect size observed in the present work, while retaining ~80% power (Arend & Schäfer, 2019; Bloom, 1995). All fixed effects in the model were multiplied by .65, and Monte Carlo simulation was used to compare the model above to an alternative omitting the 4-way interaction, $power = 83.40\%$, $95\%CI = [80.95\%, 85.66\%]$, 1000 simulations.

showed a marginal positive association with preference-like ratings, $B = 0.31$, $\beta = 0.14$, $t(140) = 2.48$, $p = .055$ (p values corrected for 4 comparisons; $\alpha_{\text{familywise}} = .05$; single-step method). Thus, among moral statements, ToMN activity was negatively related to fact-/moral-like ratings, and positively related to preference-like ratings, with both relationships present, but weaker in VMPFC (Figure 3a)⁴.

Among both facts and preferences, ToMN activity was not related to metaethical judgments. Within facts, averaging across ToMN and VMPFC, PSC did not interact with rating-type, $F(2, 135) = 0.23$, $p = .792$, and PSC showed no significant main effect, $F(1, 135) = 0.01$, $p = .991$. Likewise, within preferences, averaging across ToMN and VMPFC, PSC did not interact with rating type, $F(2, 135) = 0.61$, $p = .550$, and PSC showed no significant main effect, $F(1, 135) = 2.60$, $p = .109$ (Figure 3b).

⁴ As in footnote 3, a sensitivity analysis (Arend & Schäfer, 2019; Bloom, 1995) was conducted using *simr* (Green & MacLeod, 2016), to identify the minimum effect size we could detect at ~80% power. The negative relationship between ToMN activity and fact-/moral-like ratings could be reduced by 35% from our observed effect, while attaining 80.80% power, 95%CI = [78.22%, 83.20%], 1000 stimulations. The positive relationship between ToMN activity and preference-like ratings could be reduced by 40% from our observed effect, while attaining 82.80% power, 95%CI = [80.32%, 85.09%], 1000 stimulations. Associations with VMPFC activity were less well-powered, and should be considered with less confidence. The association between VMPFC activity and fact-/moral-like ratings could only be reduced 10% while attaining 79.50% power, 95%CI = [76.86%, 81.96%], 1000 stimulations, and the association between VMPFC activity and preference-like ratings could only be reduced 15% while attaining 81.00% power 95%CI = [78.43%, 83.39%], 1000 stimulations.

Figure 3. Behavioral–BOLD relationships. By-stimulus estimates of metaethical judgments were extracted from Study 1 and compared with by-stimulus estimates of PSC (percent signal change), extracted from a model of all ROIs in Study 2. ToMN includes averaged estimates for DMPFC, PC, RTPJ, and LTPJ (all by-stimulus random effects were perfectly correlated). (a) Within moral statements, PSC for ToMN was positively related to preference-like ratings, and negatively related to fact-/moral-like ratings. These relationships were present, but weakened in VMPFC. (b) Within facts and preferences, there was no relationship with PSC. Shaded areas represent 95% confidence intervals. We also performed a supplemental analysis using agreement and metaethical judgments collected within participants, which showed a pattern of results across ROIs consistent with these results (Figure S1).

3.2.4. The ToMN–metaethical judgment association remains significant after controlling for reaction time and semantic/syntactic features of stimuli. The identified relationship between metaethical judgments and BOLD activity could be driven by variance on some dimension outside of experimental interest, e.g. reaction time, or semantic/syntactic features of the stimuli. Along with reaction time, 13 semantic/syntactic features were collected

for each stimulus (e.g. reading ease, noun concreteness; Graesser et al., 2004; McNamara et al., 2014; see Appendix B), and added as fixed effects to the model of ToMN activity identified in 3.2.2 (along with their interactions with ROI and category). Beginning with a maximal model (omitting correlations among random effects), non-significant confounds were dropped from the model step-wise. The final model controlled for stimulus differences in concreteness, words before the main verb, and familiarity (but notably, not reaction time, which was not related to ToMN activity; for model details and full analysis, see Table S9 of the supplemental online materials). By-stimulus estimates of ROI activity were extracted from the final model, and the analyses in 3.2.3 were repeated. The results were consistent with the initial findings, although the interaction between ToMN and VMPFC was no longer significant. Among moral statements, positive ToMN activity was associated with increased preference-like ratings, $B = 1.02$, $\beta = 0.25$, $t(70) = 2.81$, $p = .013$, and decreased fact-like/moral-like ratings, $B = -0.97$, $\beta = -0.23$, $t(70) = 3.79$, $p < .001$ (p values corrected for 2 comparisons; $\alpha_{\text{familywise}} = .05$; single-step method). These relationships did not change if reaction time, and its interactions with category and ROI, were reintroduced to the behavioral model.

3.2.5. Whole brain analysis: bilateral TPJ activity is negatively associated with fact-like (objectivity) ratings. A whole brain random effects analysis of moral statements provided context for our analysis of ToMN ROIs, showing that by-stimulus BOLD activity was negatively related to fact-like ratings, and positively related to preference-like ratings in overlapping regions of bilateral TPJ (Figure 4). We performed three whole brain correlation analyses, testing the relationship between average by-stimulus PSC (modeled in Study 2) and by-stimulus behavioral estimates of fact-like, moral-like, and preference-like ratings (modeled in Study 1). Preference-like ratings were positively correlated with activity in bilateral TPJ (peak MNI coordinates: right

[54, -60, 34]; left [-36, -70, 48]), and fact-like ratings were negatively correlated in overlapping regions of bilateral RTPJ (peak MNI coordinates: right [44, -68, 46]; left [-44, -62, 48]). Peak activation for associations with fact-like ratings were slightly dorsal to the functionally defined ROI positions [peak RTPJ_{localizer}; [52, -60, 24]; peak LTPJ_{localizer} [-56, -56, 28]; see Table S3 in the online supplemental materials); however, their overlap was particularly noticeable in RTPJ (Figure 4). Surprisingly, fact-like ratings were negatively correlated with activity in the superior and bilateral middle frontal gyri, and positively correlated with activity in the parietooccipital sulcus. No voxels showed negative associations with preference-like ratings, and no voxels correlated (positively or negatively) with moral-like ratings. Thus, fact-like ratings are negatively correlated with activity in several regions, but correlations between BOLD activity and both fact-like and preference-like ratings were present in bilateral TPJ.

Figure 4. Whole brain behavioral–BOLD correlations, within moral statements. In three separate models, BOLD estimates were correlated with fact-like, moral-like, and preference-like ratings, extracted from Study 1. ToMN ROIs are pictured for reference. Fact-like ratings were negatively related to activity in bilateral TPJ, overlapping with regions showing positive correlations with preference-like ratings. These areas of overlap included areas within the defined TPJ ROIs in both hemispheres. Fact-like ratings were also negatively related to BOLD estimates in superior and middle frontal gyri, and positively related in parietooccipital sulcus. For peak coordinates see Table S8 of the online supplemental materials. For unthresholded T maps, see <https://neurovault.org/collections/ZZMRPKDV/>

3.2.6. Among moral statements, ToMN activity is not associated with the extent of

mental state information. Given that we identified an association, across stimuli, between ToMN activity and judgments of objectivity, we also examined whether a by-stimulus relationship exists between ToMN activity and the extent that a statement is judged to contain information about mental states. Unlike metaethical judgments—where low variability among facts and preferences precluded comparisons across facts, morals, and preferences (Figure 2b)—comparisons across content categories were possible for ratings of mental state inference and other item features. In total, seven additional item features were analyzed in an exploratory analysis: (a) agreement (estimated using Study 1 data); (b) ratings of the presence of *mental states* (either generally, with reference to one’s own mental states, or with reference to the mental states of others); (c) ratings of evoked mental imagery; (d) binary judgments of whether a person was present in the statement; (e) ratings of valence; (f) and ratings of arousal. Each rating was fit with maximal mixed effects model, and behavioral by-stimulus estimates were extracted and compared with by-stimulus estimates of ToMN activity (see Table S10 of the online supplemental materials for model details and ANOVA details; see also Figure S2 for a correlation matrix, displaying linear relationships among all by-stimulus estimates).

Agreement was negatively related to ToMN activity among moral statements, but showed a marginally positively relationship among preferences (Figure 5a). Correcting for three

comparisons ($\alpha_{\text{familywise}} = .05$; single-step method): among facts, the agreement–ToMN relationship was non-significant, $B = 0.02$, $\beta = 0.01$, $t(69) = 0.04$, $p = 1.00$; among preferences, the agreement–ToMN relationship was positive and marginally significant, $B = 1.01$, $\beta = 0.53$, $t(69) = 2.30$, $p = .072$; and among morals, the agreement–ToMN relationship was negative and significant, $B = -1.46$, $\beta = -0.77$, $t(69) = 2.57$, $p = .037$. That high agreement moral statements elicited less ToMN activity is consistent with our behavioral findings, where fact-like moral statements also elicited less ToMN activity (Figure 3), and positive-consensus (high agreement) moral statements were perceived as more fact-like (Figure 2b).

Mental state information (either general, self-oriented, or other-oriented; see 3.1.3), was only positively associated with ToMN activity among preferences, and no significant association was observed among facts or morals (Figure 5b–d). Among preferences, ToMN activity was positively associated with general mental state information, $B = 0.37$, $\beta = 0.31$, $t(66) = 3.39$, $p = .004$, self-oriented mental state information, $B = 0.51$, $\beta = 0.52$, $t(66) = 2.48$, $p = .047$, and other-oriented mental state information, $B = 0.47$, $\beta = 0.47$, $t(66) = 2.77$, $p = .022$. No significant association was observed among facts or preferences on any measure (Figure 5b–d; Table S10). Thus, we find some limited support that ToMN activity is sensitive to mental state information among preferences, but we also find that the relationship does not generalize to morals.

Of the remaining item features, only the presence of a person was positively associated with ToMN activity (Figure 5e): across all categories, if, on average, our online sample agreed that a statement contained a person, then increased ToMN activity was observed, $B = 1.40$, $\beta = 0.35$, $t(68) = 2.28$, $p = .026$. Other measures were not significantly related to ToMN activity (Figure 5f–h; for analyses, see Table S10 of the online supplemental materials).

Figure 5. Exploratory behavioral–BOLD relationships. By-stimulus behavioral estimates of agreement were extracted from Study 1, and from independent online studies for remaining measures. By-stimulus estimates of ToMN activity were extracted from Study 2 (3.2.2), and ToMN and VMPFC estimates were averaged for the figures above, given that no interactions with ROI were observed. (a) Agreement was negatively associated with ToMN activity among moral statements. Mental state inferences were positively associated with ToMN activity among preferences, but not among facts or morals, for each of the three ways the questions were asked: (b) general mental state information, (c) self-oriented mental state information, and (d) other-oriented mental state information. (e) The presence of a person in the statement was positively associated with ToMN activity among facts, morals, and preferences. (f–h) ToMN activity was not significantly associated with mental imagery, arousal, or valence. Shaded areas represent 95% confidence intervals. The full correlation matrix comparing all measures is available in Figure S2, in the online supplemental materials.

4. General Discussion

In the present work, we examined the relationship between by-stimulus judgments of moral objectivity/subjectivity (i.e. metaethical judgments) and activity in the ToMN network. We observed that moral statements that were judged as more objective (i.e. more fact-like and

less preference-like) elicited less ToMN activity, and moral statements judged as more subjective (i.e. less fact-like and more preference-like) elicited greater ToMN activity. This finding was confirmed in both ROI analyses (Figure 3; using an established functional localizer; Dodell-Feder et al., 2011) and a whole brain analysis (Figure 4; for unthresholded T-maps, see <https://neurovault.org/collections/ZZMRPKDV/>). ROI analyses were robust to controls for reaction time and semantic/syntactic features of the stimuli (section 3.2.4). The whole brain analysis implicated bilateral TPJ as locations in which the observed associations were relatively strong compared to other ToMN ROIs, which is consistent with the emphasis of prior work on TPJ as particularly critical region for processing social or high-level prediction error (e.g. Geng & Vossel, 2013; Schuwerk et al., 2017; but, for a brief discussion of MPFC in a predictive context, see Koster-Hale & Saxe, 2013). In addition, we explored relationships between ToMN activity and additional item features as rated by online samples (Figure 5), finding that ToMN activity was negatively associated with agreement among morals, positively associated with mental state inferences among preferences, and positively associated with the presence of a person among morals, facts, and preferences. To our knowledge, ours is the first study to identify a relationship between by-stimulus variance in ToMN activity and features of individual stimuli (prior work has tried and failed to identify significant by-stimulus relationships within the ToMN; Dodell-Feder et al., 2011)⁵. Our study is also one of only a few to leverage pre-existing social expectations in the experimental context (e.g. Cloutier et al., 2011; Park & Young, 2020),

⁵ Dodell-Feder et al (2011) did identify an association between activity in right posterior temporal parietal sulcus (a region not typically considered part of the ToMN) and mental state information (the general measure included in our study; Figure 5b). The effect size was comparable ($\beta = 0.680$), although slightly higher than the effects observed in the present work (e.g. mental state information in general, $\beta = 0.31$, information about one's own mental states, $\beta = 0.47$, and information about others' mental states, $\beta = 0.51$).

and, to our knowledge, the only one to do this while examining by-stimulus variability. Below, we discuss how these results might be integrated into a predictive framework (e.g. Koster-Hale & Saxe, 2013), how they may facilitate new approaches to understanding the social relevance of moral beliefs, and how our methodology for item analysis may be of use in future work.

4.1. Integrating predictive processing accounts and functional associations with ToMN activity. Our core aim was to help address what underlying process is supported by ToMN activity, given the heterogeneity of the tasks and stimuli that activate it. As outlined in our introduction, a variety of social tasks activate the ToMN— e.g. reading stories, watching social animations, taking perspectives, making strategic decisions, inferring traits, forming impressions, reading moral statements, and experiencing shared narratives (for review, see Amodio & Frith, 2006; Schurz et al., 2014, 2017; Van Overwalle, 2009). However, the functional scope of these cortical regions broadens even further when cortex in close proximity to the ToMN is considered. For instance, ToMN regions partially overlap with the default mode network (Buckner, 2012; Mars et al., 2012; Schurz et al., 2017), a network thought to subserve the formation of a predictive model of the external environment (Barrett, 2017; Hassabis & Maguire, 2009). Further, although ToMN regions can be spatially distinguished by peak activation from nearby cortical regions that are responsive to non-social tasks (e.g. attentional reorienting; Scholz et al., 2009; Young, Dodell-Feder, et al., 2010), it remains the case that ToMN regions are in close spatial proximity to regions activated by a variety of other abstract non-social tasks (e.g. the TPJ in particular, by episodic memory encoding/retrieval and language; for review, see Cabeza et al., 2012; Carter & Huettel, 2013), and by sensory events occurring over long temporal scales (recruiting ToMN regions more generally; Baldassano et al., 2017; Richardson & Saxe, 2019). To adequately describe the computational processes underlying

ToMN activity, a computational account should accommodate the functional heterogeneity in these areas.

In one sense, then, the contribution of the present work might seem small. ToMN activity has been previously characterized by a list of associations, and we added one more: an association between ToMN activity and metaethical judgment, where subjective morals elicited more ToMN activity, and objective morals elicited less. However, in another sense, the present work contributes to more fundamental debates about what computational process the ToMN implements. That is, the present work identified an association that is more plausibly incorporated into unifying predictive accounts of ToMN activity (e.g. Koster-Hale & Saxe, 2013), as opposed to an account based around functional modules (Baron-Cohen, 1995; Cosmides & Tooby, 1992; Scholl & Leslie, 2001)—a strong version of which might use the association we have identified to suggest that the brain contains a spatially constrained module for metaethical judgment, as prior work has similarly argued for the existence of functionally localized modules for language, number, and memory retrieval, among others (e.g. Nelson et al., 2012). Perhaps an argument for such a metaethical module could be advanced; but in our opinion, given that metaethical judgment represents judgments about an abstract feature of a moral judgment (which itself, is not clearly localized; Young & Dungan, 2012), interpreting our findings in this way might push against the limits of plausibility.

Likewise, because we did identify an association between ToMN activity and agreement within morals (Figure 5a), and because agreement was strongly correlated with fact-like, moral-like, and preference-like ratings among morals (Figure S2), one could argue that our study is better understood as identifying a relationship between agreement with morals and ToMN activity. Within the present design we cannot rule out this confound, but we want to reemphasize

that this interpretation would move us toward a theoretical framework that embraces functional specificity and brings the heterogeneity of observed effects into conflict. That is, it seems more probable that the distinct and occasionally opposite sets of associations we observed within the ToMN (e.g. a positive association with agreement among preferences, and a negative association with agreement among morals) all fall under the umbrella of some more general function that is consistently associated with ToMN activity but varies in its implementation across social contexts. It is for this reason that we favor a predictive processing account to explain our results. Future work might use a design or stimuli set that unconfounds agreement and predictability, for example, by presenting a conservative living in a liberal city with a widely shared liberal moral belief (e.g. pro-choice beliefs) and a conservative moral belief held by a minority (e.g. pro-life beliefs). In this case, we would predict relatively less ToMN activity in a pro-life participant reading the widely shared (but personally disagreeable) pro-choice belief, compared to the minority (but personally agreeable) pro-life belief.

To this end, rather than concluding that subregions within the ToMN (or surrounding the TPJ) each implement a distinct functional computation (i.e. a “fractionated” modular perspective; see Schuwerk et al., 2017), some have put forward more generalized accounts of ToMN activity, hypothesizing that regions within it subserve more general and fundamental aspects of information processing (Cabeza et al., 2012; Carter & Huettel, 2013; Decety & Lamm, 2007; Geng & Vossel, 2013; Lee & McCarthy, 2016; Schuwerk et al., 2017), such as updating context-dependent predictions when those predictions are violated (Geng & Vossel, 2013)⁶. These accounts dovetail nicely with predictive processing accounts reviewed in the introduction

⁶ Geng and Vossell (2013, p. 2616) are also clear that updating in the context of mental inference may be subserved by distinct sub-regions within the TPJ (as in Scholz et al., 2009) under the umbrella of this more general process of contextual updating.

(e.g. Barrett & Simmons, 2015; Chanes & Barrett, 2016; A. Clark, 2013, 2015; Denève & Jardri, 2016; Friston et al., 2016, 2017; Hohwy, 2013; Hutchinson & Barrett, 2019; Joiner et al., 2017; Koster-Hale & Saxe, 2013, 2013; Rao & Ballard, 1999; Shadmehr et al., 2010; Spratling, 2017; Van de Cruys et al., 2014), where the brain is thought to act as a “hierarchical prediction machine” (A. Clark, 2013), predicting and filtering incoming sensory signals and their multimodal compressions throughout the cortical hierarchy.

Note that this explanation does not dispute the finding that social information reliably elicits activity in spatially constrained cortical regions (Dodell-Feder et al., 2011; Saxe & Kanwisher, 2003; Saxe & Powell, 2006; Saxe & Wexler, 2005; Scholz et al., 2009; Young, Dodell-Feder, et al., 2010, 2010). Indeed, we observed just this, finding positive associations between ToMN activity and the presence of a person within facts, morals, and preferences, and a positive association between ToMN activity and multiple measures of mental state information within preferences (Figure 5b–d). We also observed a marginal positive association between ToMN activity and agreement within preferences (the opposite of the pattern observed among morals; Figure 5a), but this most likely stems from the confounding of agreement and measures of mental state information within our sample of preferences (Figure S2).

A predictive account, however, unlike a modular perspective, suggests that socially-sensitive brain regions are the particular point of collision between bottom-up multimodal compressions of sensory signals (e.g. the sights, sounds, etc., that collectively form the percept of a person) and top-down predictions about latent causes (i.e. the latent mental states that are predicted to produce observed patterns of behavior, movement, etc.). On this predictive account, then, ToMN activity is thought to reflect the updating of predictions about high-level

compressions of sensory information (Bach & Schenke, 2017; Hohwy, 2013; Kilner & Frith, 2008; Koster-Hale & Saxe, 2013; Theriault et al., 2019; also, see Ondobaka et al., 2017).

If this general predictive account were correct, then one would expect to see multiple functional associations within ToMN regions, given that several stimulus features may simultaneously facilitate updating predictions about latent mental states. We observed exactly this: in addition to the association, within morals, between ToMN activity and metaethical judgments, we also observed a positive association between the presence of a person and ToMN activity across facts, morals, and preferences, a negative association between agreement and ToMN activity among morals, and a positive association between mental state inferences and ToMN activity among preferences. On a general predictive interpretation, across facts, morals, and preferences the presence of a person may facilitate predictions about mental states (as their presence provides a referent for mental state predictions), whereas metaethical features may only facilitate predictions about mental states among moral statements (as social consensus is generally more likely to guide predictions in the moral domain than among facts and preferences; see section 4.2). Although our present findings are only *consistent* with this account—and not definitive evidence for it—a general account of ToMN activity may nonetheless be a more attractive way to integrate the heterogeneity of previously observed associations in ToMN regions and the cortex surrounding them, especially when considering the association between ToMN activity and metaethical judgment observed in the present work.

4.2. Morality and metaethics in an informational context. Principles of information theory and predictive processing—e.g. that perfectly predictable signals carry no information—are of obvious use to neuroscience, as they make concrete predictions about how information is represented within the brain. Although the utility of these perspectives may appear less obvious

in the context of morality, they may actually help to organize a variety of insights into the nature and importance of moral beliefs under a coherent theoretical framework. Critically, an understanding of social information grounded in prediction, prediction error, and the precision of predictions, may provide new perspectives on metaethical judgment, the importance of morality to social identity, and the demarcation of social domains (i.e. morals vs. preferences). Below, we review each of these potential contributions in turn.

Objective moral beliefs are moral beliefs that are supported by a social consensus (Ayars & Nichols, in press; Beebe, 2014; Goodwin & Darley, 2012; Heiphetz & Young, 2017), making them predictable. By extension, if objective moral beliefs are predictable (i.e. on the basis of social consensus), then affirming them carries less information. For example, when someone tells you they believe that “drinking and driving is wrong”, then their beliefs roughly match what you would have already predicted, knowing nothing about them. Conversely, disavowing an objective moral belief provides more information about the person who does the disavowing, e.g. hearing someone say “drinking and driving is good” is informative about them. In this way, metaethics may be productively conceptualized (for some purposes) in informational terms: objective and subjective moral claims may roughly correspond to moral claims that are uninformative and informative, respectively.

The centrality of moral beliefs to people’s perceptions of individual identity has been demonstrated in experimental contexts involving hypothetical (Heiphetz et al., 2018; Strohminger & Nichols, 2014) and real (Strohminger & Nichols, 2015) moral changes. For example, changes in a person’s moral beliefs are judged as changing their core identity more than changes in memories or physical abilities. In informational terms, if moral beliefs are generally predicted to remain more consistent over time than memories, then moral changes may

866 be understood as representing an informationally important shift in an individual's personality.
867 That is, we expect memories and physical abilities to change; but we do not expect moral
868 changes to the same degree, making such changes more informative. Of course, this leaves
869 unanswered *why* moral beliefs are predicted to remain more consistent over time, a line of
870 questioning that could be pursued in future work.

871 Finally, behavioral and neural distinctions between social domains (e.g. morals vs.
872 preferences) might be explained by differences in how predictive processes are applied, for
873 example, differences in the precision of predictions that are typically made about the domain. In
874 the context of information theory and predictive processing models, prediction error is
875 determined by both the prediction (e.g. a distribution, with a mean) and its precision (e.g. the
876 width of that distribution; Feldman & Friston, 2010; Kim et al., 2020; Van de Cruys et al., 2014).
877 For example, a slight deviation from the mean of a high-precision prediction (i.e. a tight
878 distribution) would create more prediction error than a slight deviation from the mean of a low-
879 precision prediction. Predictions in some social domains may be more precise than in others. For
880 example, in the context of morality, a culture often accepts only one, or a few, beliefs as
881 acceptable (e.g. on the acceptability of slavery), whereas on questions of preferences, many
882 beliefs are acceptable (e.g. on the best flavor of ice cream)⁷. On an information theoretic
883 perspective, this would mean that predictions in the moral domain may be (on average) more

⁷ One might object that this is an unfair comparison, and that there are fewer possible answers to moral questions (e.g. slavery either is, or is not acceptable, whereas any ice cream flavor might be best); however, we think this view is mistaken. A society might organize itself around any number of arbitrary rules for when slavery is, or is not acceptable (e.g. only for people of a certain race, or for people with a certain accent, or not for children, or only when someone has committed a crime, etc.). That questions of morality intuitively appear as binary, whereas questions of preferences do not, may itself be a product of the precision of predictions in the moral domain. That is, if moral beliefs are predicted precisely, then a sharp contrast may be drawn between the few beliefs that fit the prediction and the many beliefs that do not. By contrast, less precise predictions about preferences may make it more intuitive that a variety of beliefs could answer the question of "which ice cream flavor is best", each fitting to a different degree.

precise than predictions in the preference domain. This interpretation that may explain our previous finding that moral statements elicit greater activity than preferences throughout the ToMN (Theriault et al., 2017), as low-precision predictions about the preferences of generic people afford less opportunity for predictions to be violated (and prediction error to be generated). Further, an explanatory framework grounded in predictive processes may provide an avenue for understanding other discrepancies we observed, e.g. that ToMN activity was associated with mental state information among preferences but not among morals (Figure 5b–d). Rather than assuming that mental inference draws on different brain regions for morals and preferences, we propose that common predictive processes support mental inference across diverse social contexts. Of course, future work might more closely examine factors that vary across these contexts (e.g. predictive precision); but nonetheless, we believe that understanding social cognition in information-based terms may be a productive method to integrate several empirical findings in social psychology, both with each other and with emerging and unifying predictive theories of cortical function.

4.3. Probing BOLD activity through item analysis. Finally, it is worth emphasizing the potential benefits of the item analysis approach used in the present work. Social information is complex, and although one analytic approach would be to fully embrace this complexity, this would also complicate investigations of specific features and their influence on BOLD activity. At the other end of the spectrum, computational models can cleanly characterize a process, but also require that the task be somewhat removed from naturalistic contexts (e.g. as in economic games). Item analysis offers a middle path, where researchers can relax constraints on their stimulus set to maximize variance and generalizability (Westfall et al., 2017), but, at the same time, conduct fine-grained analyses across the feature space represented by the population of

stimuli. Furthermore, this approach is advantageous in that stimuli can be normed and reused—e.g. researchers can utilize our stimuli (and our by-stimulus estimates) to test their own hypotheses about the dimensions underlying BOLD activity (see Appendix A for the complete stimulus set and paired by-stimulus estimates, and see <https://osf.io/cx4dp/> for raw functional, anatomical, and behavioral by-stimulus estimates in standardized BIDS format; Gorgolewski et al., 2016). By-stimulus variability is often left unexamined in existing datasets (e.g. in an emotional expression task, amygdala activity varies across emotional faces; Westfall et al., 2017), and future work in neuroimaging may benefit from adopting this underutilized method of independently norming stimuli to examine by-stimulus variance.

4.4. Conclusion. The theory of mind network is broadly involved in social cognition (for review, see Amodio & Frith, 2006; Schurz et al., 2014, 2017; Van Overwalle, 2009), and it has recently been suggested that a more general computational process may account for this activity: processing prediction error in social contexts (Koster-Hale & Saxe, 2013). Prior work has been consistent with this hypothesis, but this work has generally examined social expectations formed inside the lab (e.g. Dungan et al., 2016; Saxe & Wexler, 2005) or under conditions of explicit instruction (Brass et al., 2007; de Lange et al., 2008), and has not conducted more fine-grained analyses of by-stimulus variance. In the present work, we leveraged an existing dataset (Theriault et al., 2017) to examine by-stimulus variability in ToMN evoked by moral statements, focusing on preexisting expectations as operationalized by metaethical judgments. We observed that ToMN activity was negatively associated with moral objectivity (i.e. how fact-like a moral statement was judged to be, by an independent online sample) and positively associated with moral subjectivity (i.e. how preference-like a moral statement was judged to be, by an independent online sample). This finding is consistent with hypotheses derived from predictive

processing models, and, although it is not definitive, it underscores the need for an overarching explanation of ToMN activity that can accommodate the heterogeneity of associations previously observed in these regions and in the cortex surrounding them (e.g. Carter & Huettel, 2013; Geng & Vossel, 2013; Koster-Hale & Saxe, 2013).

Acknowledgements

We thank Fiery Cushman, Drew Linsley, Sean MacEvoy, Jonathan Phillips, James Russell, Ajay Satpute, Rebecca Saxe, and members of the Morality Lab for feedback. This work was supported by the National Science Foundation 5103831 (L.Y.), the John Templeton Foundation 5107321 (L.Y.), the National Science Foundation SMA-1408989 (L.H.), and Natural Sciences and Engineering Research Council of Canada PGSD3-420445 (J.T.). This work was also supported (in part) by an award from the Russell Sage Foundation (L.H.). Any opinions expressed are those of the authors alone and should not be construed as representing the opinions of the Foundation.

References

- Amodio, D. M., & Frith, C. D. (2006). Meeting of minds: The medial frontal cortex and social cognition. *Nature Reviews Neuroscience*, 7(4), 268–277. <https://doi.org/10.1038/nrn1884>
- Ampe, L., Ma, N., Van Hoeck, N., Vandekerckhove, M., & Van Overwalle, F. (2014). Unusual actions do not always trigger the mentalizing network. *Neurocase*, 20(2), 144–149. <https://doi.org/10.1080/13554794.2012.741251>
- Arend, M. G., & Schäfer, T. (2019). Statistical power in two-level models: A tutorial based on Monte Carlo simulation. *Psychological Methods*, 24(1), 1–19. <https://doi.org/10.1037/met0000195>
- Ayars, A., & Nichols, S. (in press). Rational learners and metaethics: Universalism, relativism, and evidence from consensus. *Mind & Language*. <https://doi.org/10.1111/mila.12232>
- Baayen, R. H. (2008). *Analyzing Linguistic Data: A Practical Introduction to Statistics using R*. Cambridge University Press.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Bach, P., & Schenke, K. C. (2017). Predictive social perception: Towards a unifying framework from action observation to person knowledge. *Social and Personality Psychology Compass*, 11(7), e12312. <https://doi.org/10.1111/spc3.12312>
- Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3), 709–721.e5. <https://doi.org/10.1016/j.neuron.2017.06.041>

- Baron, S. G., Gobbini, M. I., Engell, A. D., & Todorov, A. (2011). Amygdala and dorsomedial prefrontal cortex responses to appearance-based and behavior-based person impressions. *Social Cognitive and Affective Neuroscience*, 6(5), 572–581.
<https://doi.org/10.1093/scan/nsq086>
- Baron-Cohen, S. (1995). *Mindblindness: An essay on autism and theory of mind*. MIT Press.
- 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. <https://doi.org/10.1016/j.jml.2012.11.001>
- Barrett, L. F. (2017). The theory of constructed emotion: An active inference account of interoception and categorization. *Social Cognitive and Affective Neuroscience*, 12(1), 1–23. <https://doi.org/10.1093/scan/nsw154>
- Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, 16(7), 419–429. <https://doi.org/10.1038/nrn3950>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48.
<https://doi.org/10.18637/jss.v067.i01>
- Bedny, M., Aguirre, G. K., & Thompson-Schill, S. L. (2007). Item analysis in functional magnetic resonance imaging. *NeuroImage*, 35(3), 1093–1102.
<https://doi.org/10.1016/j.neuroimage.2007.01.039>
- Beebe, J. R. (2014). How Different Kinds of Disagreement Impact Folk Metaethical Judgments. In J. C. Wright & H. Sarkissian (Eds.), *Advances in Experimental Moral Psychology* (pp. 167–187). Bloomsbury Academic. <https://doi.org/10.5040/9781472594150>

- Bhanji, J. P., & Beer, J. S. (2013). Dissociable Neural Modulation Underlying Lasting First Impressions, Changing Your Mind for the Better, and Changing It for the Worse. *Journal of Neuroscience*, 33(22), 9337–9344. <https://doi.org/10.1523/JNEUROSCI.5634-12.2013>
- Blakemore, S.-J. (2003). The Detection of Contingency and Animacy from Simple Animations in the Human Brain. *Cerebral Cortex*, 13(8), 837–844. <https://doi.org/10.1093/cercor/13.8.837>
- Bloom, H. S. (1995). Minimum Detectable Effects: A Simple Way to Report the Statistical Power of Experimental Designs. *Evaluation Review*, 19(5), 547–556. <https://doi.org/10.1177/0193841X9501900504>
- Brass, M., Schmitt, R. M., Spengler, S., & Gergely, G. (2007). Investigating Action Understanding: Inferential Processes versus Action Simulation. *Current Biology*, 17(24), 2117–2121. <https://doi.org/10.1016/j.cub.2007.11.057>
- Buckner, R. L. (2012). The serendipitous discovery of the brain's default network. *NeuroImage*, 62(2), 1137–1145. <https://doi.org/10.1016/j.neuroimage.2011.10.035>
- Cabeza, R., Ciaramelli, E., & Moscovitch, M. (2012). Cognitive contributions of the ventral parietal cortex: An integrative theoretical account. *Trends in Cognitive Sciences*, 16(6), 338–352. <https://doi.org/10.1016/j.tics.2012.04.008>
- Carter, R. M., & Huettel, S. A. (2013). A nexus model of the temporal–parietal junction. *Trends in Cognitive Sciences*, 17(7), 328–336. <https://doi.org/10.1016/j.tics.2013.05.007>
- Chanes, L., & Barrett, L. F. (2016). Redefining the role of limbic areas in cortical processing. *Trends in Cognitive Sciences*, 20(2), 96–106. <https://doi.org/10.1016/j.tics.2015.11.005>

- Ciaramidaro, A., Adenzato, M., Enrici, I., Erk, S., Pia, L., Bara, B. G., & Walter, H. (2007). The intentional network: How the brain reads varieties of intentions. *Neuropsychologia*, 45(13), 3105–3113. <https://doi.org/10.1016/j.neuropsychologia.2007.05.011>
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 1–24. <https://doi.org/10.1017/S0140525X12000477>
- Clark, A. (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12(4), 335–359. [https://doi.org/10.1016/S0022-5371\(73\)80014-3](https://doi.org/10.1016/S0022-5371(73)80014-3)
- Cloutier, J., Gabrieli, J. D. E., O’Young, D., & Ambady, N. (2011). An fMRI study of violations of social expectations: When people are not who we expect them to be. *NeuroImage*, 57(2), 583–588. <https://doi.org/10.1016/j.neuroimage.2011.04.051>
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163–228). Oxford University Press.
- de Lange, F. P., Spronk, M., Willems, R. M., Toni, I., & Bekkering, H. (2008). Complementary systems for understanding action intentions. *Current Biology: CB*, 18(6), 454–457. <https://doi.org/10.1016/j.cub.2008.02.057>
- Decety, J., & Lamm, C. (2007). The Role of the Right Temporoparietal Junction in Social Interaction: How Low-Level Computational Processes Contribute to Meta-Cognition. *The Neuroscientist*, 13(6), 580–593. <https://doi.org/10.1177/1073858407304654>

- Denève, S., & Jardri, R. (2016). Circular inference: Mistaken belief, misplaced trust. *Current Opinion in Behavioral Sciences*, 11, 40–48. <https://doi.org/10.1016/j.cobeha.2016.04.001>
- Dodell-Feder, D., Koster-Hale, J., Bedny, M., & Saxe, R. (2011). FMRI item analysis in a theory of mind task. *NeuroImage*, 55(2), 705–712. <https://doi.org/10.1016/j.neuroimage.2010.12.040>
- Donnet, S., Lavielle, M., & Poline, J.-B. (2006). Are fMRI event-related response constant in time? A model selection answer. *NeuroImage*, 31(3), 1169–1176. <https://doi.org/10.1016/j.neuroimage.2005.08.068>
- Dungan, J. A., Stepanovic, M., & Young, L. L. (2016). Theory of mind for processing unexpected events across contexts. *Social Cognitive and Affective Neuroscience*, 11(8), 1183–1192. <https://doi.org/10.1093/scan/nsw032>
- Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4. <https://doi.org/10.3389/fnhum.2010.00215>
- Fletcher, P. (1995). Other minds in the brain: A functional imaging study of “theory of mind” in story comprehension. *Cognition*, 57(2), 109–128. [https://doi.org/10.1016/0010-0277\(95\)00692-R](https://doi.org/10.1016/0010-0277(95)00692-R)
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O’Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68, 862–879. <https://doi.org/10.1016/j.neubiorev.2016.06.022>
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation*, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912

- Gallagher, H. L., Happé, F., Brunswick, N., Fletcher, P. C., Frith, U., & Frith, C. D. (2000). Reading the mind in cartoons and stories: An fMRI study of ‘theory of mind’ in verbal and nonverbal tasks. *Neuropsychologia*, 38(1), 11–21. [https://doi.org/10.1016/S0028-3932\(99\)00053-6](https://doi.org/10.1016/S0028-3932(99)00053-6)
- Gelman, A., Hill, J., & Yajima, M. (2012). Why We (Usually) Don’t Have to Worry About Multiple Comparisons. *Journal of Research on Educational Effectiveness*, 5(2), 189–211. <https://doi.org/10.1080/19345747.2011.618213>
- Gelman, S. A., & Rhodes, M. (2012). “Two-Thousand Years of Stasis”: How Psychological Essentialism Impedes Evolutionary Understanding. In K. S. Rosengren, S. Brem, E. M. Evans, & G. M. Sinatra (Eds.), *Evolution challenges: Integrating research and practice in teaching and learning about evolution*. Oxford University Press. <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199730421.001.0001/acprof-9780199730421-chapter-1>
- Geng, J. J., & Vossel, S. (2013). Re-evaluating the role of TPJ in attentional control: Contextual updating? *Neuroscience & Biobehavioral Reviews*, 37(10, Part 2), 2608–2620. <https://doi.org/10.1016/j.neubiorev.2013.08.010>
- Gobbini, M. I., Koralek, A. C., Bryan, R. E., Montgomery, K. J., & Haxby, J. V. (2007). Two Takes on the Social Brain: A Comparison of Theory of Mind Tasks. *Journal of Cognitive Neuroscience*, 19(11), 1803–1814. <https://doi.org/10.1162/jocn.2007.19.11.1803>
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339–1366. <https://doi.org/10.1016/j.cognition.2007.06.007>

- Goodwin, G. P., & Darley, J. M. (2012). Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology*, 48(1), 250–256.
<https://doi.org/10.1016/j.jesp.2011.08.006>
- Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., Flandin, G., Ghosh, S. S., Glatard, T., Halchenko, Y. O., Handwerker, D. A., Hanke, M., Keator, D., Li, X., Michael, Z., Maumet, C., Nichols, B. N., Nichols, T. E., Pellman, J., ... Poldrack, R. A. (2016). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, 3(1), 1–9.
<https://doi.org/10.1038/sdata.2016.44>
- Graesser, A. C., McNamara, D. S., Louwerse, M. M., & Cai, Z. (2004). Coh-Metrix: Analysis of text on cohesion and language. *Behavior Research Methods, Instruments, & Computers*, 36(2), 193–202. <https://doi.org/10.3758/BF03195564>
- Green, P., & MacLeod, C. J. (2016). SIMR: an R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493–498.
<https://doi.org/10.1111/2041-210X.12504>
- Harris, L. T., Todorov, A., & Fiske, S. T. (2005). Attributions on the brain: Neuro-imaging dispositional inferences, beyond theory of mind. *NeuroImage*, 28(4), 763–769.
<https://doi.org/10.1016/j.neuroimage.2005.05.021>
- Hassabis, D., & Maguire, E. (2009). The construction system of the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1263–1271.
<https://doi.org/10.1098/rstb.2008.0296>

- Heiphetz, L., Spelke, E. S., Harris, P. L., & Banaji, M. R. (2014). What do different beliefs tell us? An examination of factual, opinion-based, and religious beliefs. *Cognitive Development*, 30, 15–29. <https://doi.org/10.1016/j.cogdev.2013.12.002>
- Heiphetz, L., Strohminger, N., Gelman, S. A., & Young, L. L. (2018). Who am I? The role of moral beliefs in children's and adults' understanding of identity. *Journal of Experimental Social Psychology*, 78, 210–219. <https://doi.org/10.1016/j.jesp.2018.03.007>
- Heiphetz, L., & Young, L. L. (2017). Can only one person be right? The development of objectivism and social preferences regarding widely shared and controversial moral beliefs. *Cognition*, 167, 78–90. <https://doi.org/10.1016/j.cognition.2016.05.014>
- Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199682737.001.0001>
- Hothorn, T., Bretz, F., & Westfall, P. (2008). Simultaneous inference in general parametric models. *Biometrical Journal. Biometrische Zeitschrift*, 50(3), 346–363.
<https://doi.org/10.1002/bimj.200810425>
- Hutchinson, B., & Barrett, L. F. (2019). The power of predictions: An emerging paradigm for psychological research. *Current Directions in Psychological Science*, 28(3), 280–291.
<https://doi.org/10.1177/0963721419831992>
- Jenkins, A. C., & Mitchell, J. P. (2010). Mentalizing under Uncertainty: Dissociated Neural Responses to Ambiguous and Unambiguous Mental State Inferences. *Cerebral Cortex*, 20(2), 404–410. <https://doi.org/10.1093/cercor/bhp109>
- Joiner, J., Piva, M., Turrin, C., & Chang, S. W. C. (2017). Social learning through prediction error in the brain. *Npj Science of Learning*, 2(1). <https://doi.org/10.1038/s41539-017-0009-2>

- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69.
<https://doi.org/10.1037/a0028347>
- Kilner, J. M., & Frith, C. D. (2008). Action Observation: Inferring Intentions without Mirror Neurons. *Current Biology*, 18(1), R32–R33. <https://doi.org/10.1016/j.cub.2007.11.008>
- Kim, M., Park, B., & Young, L. (2020). The psychology of motivated versus rational impression updating. *Trends in Cognitive Sciences*, 24(2), 101–111.
<https://doi.org/10.1016/j.tics.2019.12.001>
- Kircher, T., Blümel, I., Marjoram, D., Lataster, T., Krabbendam, L., Weber, J., van Os, J., & Krach, S. (2009). Online mentalising investigated with functional MRI. *Neuroscience Letters*, 454(3), 176–181. <https://doi.org/10.1016/j.neulet.2009.03.026>
- Koster-Hale, J., & Saxe, R. (2013). Theory of mind: A neural prediction problem. *Neuron*, 79(5), 836–848. <https://doi.org/10.1016/j.neuron.2013.08.020>
- Kron, A., Goldstein, A., Lee, D. H.-J., Gardhouse, K., & Anderson, A. K. (2013). How Are You Feeling? Revisiting the Quantification of Emotional Qualia. *Psychological Science*, 24(8), 1503–1511. <https://doi.org/10.1177/0956797613475456>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(1), 1–26.
<https://doi.org/10.18637/jss.v082.i13>
- Lee, S. M., & McCarthy, G. (2016). Functional Heterogeneity and Convergence in the Right Temporoparietal Junction. *Cerebral Cortex (New York, N.Y.: 1991)*, 26(3), 1108–1116.
<https://doi.org/10.1093/cercor/bhu292>

Ma, N., Vandekerckhove, M., Baetens, K., Van Overwalle, F., Seurinck, R., & Fias, W. (2012).

Inconsistencies in spontaneous and intentional trait inferences. *Social Cognitive and Affective Neuroscience*, 7(8), 937–950. <https://doi.org/10.1093/scan/nsr064>

Ma, N., Vandekerckhove, M., Hoeck, N. V., & Overwalle, F. V. (2012). Distinct recruitment of temporo-parietal junction and medial prefrontal cortex in behavior understanding and trait identification. *Social Neuroscience*, 7(6), 591–605.

<https://doi.org/10.1080/17470919.2012.686925>

Mars, R. B., Neubert, F.-X., Noonan, M. P., Sallet, J., Toni, I., & Rushworth, M. F. S. (2012).

On the relationship between the “default mode network” and the “social brain.” *Frontiers in Human Neuroscience*, 6. <https://doi.org/10.3389/fnhum.2012.00189>

McNamara, D. S., Louwerse, M. M., Cai, Z., & Graesser, A. (2014). *Coh-Metrix version 3.0*.

<http://cohmetrix.com>.

Mende-Siedlecki, P., Cai, Y., & Todorov, A. (2013). The neural dynamics of updating person impressions. *Social Cognitive and Affective Neuroscience*, 8(6), 623–631.

<https://doi.org/10.1093/scan/nss040>

Mende-Siedlecki, P., & Todorov, A. (2016). Neural dissociations between meaningful and mere inconsistency in impression updating. *Social Cognitive and Affective Neuroscience*,

11(9), 1489–1500. <https://doi.org/10.1093/scan/nsw058>

Mitchell, J. P., Banaji, M. R., & Macrae, C. N. (2005). General and specific contributions of the medial prefrontal cortex to knowledge about mental states. *NeuroImage*, 28(4), 757–762.

<https://doi.org/10.1016/j.neuroimage.2005.03.011>

- Nelson, S. M., McDermott, K. B., & Petersen, S. E. (2012). In favor of a ‘fractionation’ view of ventral parietal cortex: Comment on Cabeza et al. *Trends in Cognitive Sciences*, 16(8), 399–400. <https://doi.org/10.1016/j.tics.2012.06.014>
- Nichols, T. E., & Holmes, A. P. (2002). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping*, 15(1), 1–25. <https://doi.org/10.1002/hbm.1058>
- Ondobaka, S., Kilner, J., & Friston, K. (2017). The role of interoceptive inference in theory of mind. *Brain and Cognition*, 112, 64–68. <https://doi.org/10.1016/j.bandc.2015.08.002>
- Park, B., & Young, L. (2020). An association between biased impression updating and relationship facilitation: A behavioral and fMRI investigation. *Journal of Experimental Social Psychology*, 87, 103916. <https://doi.org/10.1016/j.jesp.2019.103916>
- Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4), 515–526. <https://doi.org/10.1017/S0140525X00076512>
- R Core Team. (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.R-project.org>
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>
- Richardson, H., & Saxe, R. (2019). Development of predictive responses in theory of mind brain regions. *Developmental Science*, e12863. <https://doi.org/10.1111/desc.12863>
- Ruby, P., & Decety, J. (2003). What you believe versus what you think they believe: A neuroimaging study of conceptual perspective-taking: A PET study of conceptual

- perspective-taking. *European Journal of Neuroscience*, 17(11), 2475–2480.
<https://doi.org/10.1046/j.1460-9568.2003.02673.x>
- Sarkissian, H., Park, J., Tien, D., Wright, J. C., & Knobe, J. (2011). Folk Moral Relativism. *Mind & Language*, 26(4), 482–505. <https://doi.org/10.1111/j.1468-0017.2011.01428.x>
- Saxe, R., & Kanwisher, N. (2003). People thinking about thinking people: The role of the temporo-parietal junction in “theory of mind.” *NeuroImage*, 19(4), 1835–1842.
[https://doi.org/10.1016/S1053-8119\(03\)00230-1](https://doi.org/10.1016/S1053-8119(03)00230-1)
- Saxe, R., & Powell, L. J. (2006). It’s the Thought That Counts: Specific Brain Regions for One Component of Theory of Mind. *Psychological Science*, 17(8), 692–699.
<https://doi.org/10.1111/j.1467-9280.2006.01768.x>
- Saxe, R., & Wexler, A. (2005). Making sense of another mind: The role of the right temporo-parietal junction. *Neuropsychologia*, 43(10), 1391–1399.
<https://doi.org/10.1016/j.neuropsychologia.2005.02.013>
- Schiller, D., Freeman, J. B., Mitchell, J. P., Uleman, J. S., & Phelps, E. A. (2009). A neural mechanism of first impressions. *Nature Neuroscience*, 12(4), 508–514.
<https://doi.org/10.1038/nn.2278>
- Scholl, B. J., & Leslie, A. M. (2001). Minds, modules, and meta-analysis. *Child Development*, 72(3), 696–701. <https://doi.org/10.1111/1467-8624.00308>
- Scholz, J., Triantafyllou, C., Whitfield-Gabrieli, S., Brown, E. N., & Saxe, R. (2009). Distinct Regions of Right Temporo-Parietal Junction Are Selective for Theory of Mind and Exogenous Attention. *PLoS ONE*, 4(3), e4869.
<https://doi.org/10.1371/journal.pone.0004869>

- Schurz, M., Radua, J., Aichhorn, M., Richlan, F., & Perner, J. (2014). Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience & Biobehavioral Reviews*, 42, 9–34. <https://doi.org/10.1016/j.neubiorev.2014.01.009>
- Schurz, M., Tholen, M. G., Perner, J., Mars, R. B., & Sallet, J. (2017). Specifying the brain anatomy underlying temporo-parietal junction activations for theory of mind: A review using probabilistic atlases from different imaging modalities. *Human Brain Mapping*, 38(9), 4788–4805. <https://doi.org/10.1002/hbm.23675>
- Schuwerk, T., Schurz, M., Müller, F., Rupprecht, R., & Sommer, M. (2017). The rTPJ's overarching cognitive function in networks for attention and theory of mind. *Social Cognitive and Affective Neuroscience*, 12(1), 157–168. <https://doi.org/10.1093/scan/nsw163>
- Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33(1), 89–108. <https://doi.org/10.1146/annurev-neuro-060909-153135>
- Shannon, C., & Weaver, W. (1964). *The Mathematical Theory of Communication* (10th ed.). The University of Illinois Press. (Original work published 1949)
- Spratling, M. W. (2017). A review of predictive coding algorithms. *Brain and Cognition*, 112, 92–97. <https://doi.org/10.1016/j.bandc.2015.11.003>
- Strohming, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1), 159–171. <https://doi.org/10.1016/j.cognition.2013.12.005>
- Strohming, N., & Nichols, S. (2015). Neurodegeneration and Identity. *Psychological Science*, 26(9), 1469–1479. <https://doi.org/10.1177/0956797615592381>

- Theriault, J. E., Waytz, A., Heiphetz, L., & Young, L. L. (2017). Examining overlap in behavioral and neural representations of morals, facts, and preferences. *Journal of Experimental Psychology: General*, 146(11), 1586–1605.
<https://doi.org/10.1037/xge0000350>
- Theriault, J. E., Young, L. L., & Barrett, L. F. (2019). The sense of should: A biologically-based model of social pressure. *PsyArXiv*. <https://doi.org/10.31234/osf.io/x5rbs>
- Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de-Wit, L., & Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review*, 121(4), 649–675. <https://doi.org/10.1037/a0037665>
- Van Overwalle, F. (2009). Social cognition and the brain: A meta-analysis. *Human Brain Mapping*, 30(3), 829–858. <https://doi.org/10.1002/hbm.20547>
- Vogeley, K., Bussfeld, P., Newen, A., Herrmann, S., Happé, F., Falkai, P., Maier, W., Shah, N. J., Fink, G. R., & Zilles, K. (2001). Mind Reading: Neural Mechanisms of Theory of Mind and Self-Perspective. *NeuroImage*, 14(1), 170–181.
<https://doi.org/10.1006/nimg.2001.0789>
- Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General*, 143(5), 2020–2045.
<https://doi.org/10.1037/xge0000014>
- Westfall, J., Nichols, T. E., & Yarkoni, T. (2017). Fixing the stimulus-as-fixed-effect fallacy in task fMRI. *Wellcome Open Research*, 1.
<https://doi.org/10.12688/wellcomeopenres.10298.2>

- Woo, C.-W., Krishnan, A., & Wager, T. D. (2014). Cluster-extent based thresholding in fMRI analyses: Pitfalls and recommendations. *NeuroImage*, *91*, 412–419.
<https://doi.org/10.1016/j.neuroimage.2013.12.058>
- Wright, J. C., Grandjean, P. T., & McWhite, C. B. (2013). The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology*, *26*(3), 336–361. <https://doi.org/10.1080/09515089.2011.633751>
- Yeshurun, Y., Swanson, S., Simony, E., Chen, J., Lazaridi, C., Honey, C. J., & Hasson, U. (2017). Same Story, Different Story: The Neural Representation of Interpretive Frameworks. *Psychological Science*, *28*(3), 307–319.
<https://doi.org/10.1177/0956797616682029>
- Young, L. L., Camprodon, J. A., Hauser, M., Pascual-Leone, A., & Saxe, R. (2010). Disruption of the right temporoparietal junction with transcranial magnetic stimulation reduces the role of beliefs in moral judgments. *Proceedings of the National Academy of Sciences*, *107*(15), 6753–6758. <https://doi.org/10.1073/pnas.0914826107>
- Young, L. L., Cushman, F., Hauser, M., & Saxe, R. (2007). The neural basis of the interaction between theory of mind and moral judgment. *Proceedings of the National Academy of Sciences*, *104*(20), 8235–8240. <https://doi.org/10.1073/pnas.0701408104>
- Young, L. L., Dodell-Feder, D., & Saxe, R. (2010). What gets the attention of the temporoparietal junction? An fMRI investigation of attention and theory of mind. *Neuropsychologia*, *48*(9), 2658–2664.
<https://doi.org/10.1016/j.neuropsychologia.2010.05.012>

Young, L. L., & Dungan, J. (2012). Where in the brain is morality? Everywhere and maybe nowhere. *Social Neuroscience*, 7(1), 1–10.

<https://doi.org/10.1080/17470919.2011.569146>

Young, L. L., & Saxe, R. (2009). An fMRI Investigation of Spontaneous Mental State Inference for Moral Judgment. *Journal of Cognitive Neuroscience*, 21(7), 1396–1405.

<https://doi.org/10.1162/jocn.2009.21137>

Appendix A. Experimental Stimuli and best linear unbiased predictors (BLUPs).

Morals					
Positive-consensus	ToM BLUPs	VMPFC BLUPs	Fact-like BLUPs	Moral-like BLUPs	Preference-like BLUPs
The deplorable conditions of Chinese electronics workers should not be ignored.	0.011	1.1269	3.0221	6.2564	3.8453
Driving after drinking heavily is a stupid and selfish way to behave.	-0.5251	0.5227	3.9144	5.9972	4.0798
Professors should not tolerate students cheating on their exams.	0.2148	0.2742	3.2157	5.9815	4.303
It is irresponsible for airlines to risk the safety of their passengers.	-0.7352	-0.0465	3.8749	6.2122	3.2389
Parents should be willing to make sacrifices for the benefit of their baby.	0.1883	1.5678	3.0313	5.7222	4.8858
The goal of sports should be to teach children that respect for others is more important than winning.	0.0277	0.9978	2.0062	4.5673	5.3466
No-consensus					
It is wrong to use animals as disposable space shuttle test pilots.	0.3939	1.7801	2.5669	6.14	4.6066
Dog racing is harmful and exploitative to the dogs being raced.	-0.2061	1.1528	3.3337	5.9298	4.1966
Destroying the habitats of owls through deforestation is deplorable.	-0.229	0.9305	2.497	6.1893	3.9591
It is unjust for businesses to allow apples to rot rather than giving them to the needy.	0.4754	0.8605	2.13	6.0441	4.4336
Music stores should prevent children from buying CDs with violent or sexist lyrics.	0.9306	2.5563	1.9186	5.1095	5.2449
Eating fish is acceptable if they were treated humanely when caught or raised.	-0.0027	0.4265	1.947	5.4241	5.0757
Good Americans buy American cars, such as Hummers.	0.371	2.2236	1.4816	2.6911	6.4033
It is wrong to knowingly buy sandals made using sweatshop labor.	0.0083	0.2	2.1262	6.1149	4.5567
People should help their elderly neighbors clear snow from their driveway.	1.0134	1.664	2.0092	5.6906	4.6681
Harry Potter should be banned from school libraries for idolizing witchcraft.	0.673	2.6323	1.2825	3.9891	5.6957
It is wrong to cheat when playing games such as Monopoly.	-0.0669	-0.0154	2.1859	6.1189	4.2002

It is unethical for businesses to promote sugary products to children.	0.0183	2.0761	2.1744	5.8517	4.3915
Negative-consensus					
It is wrong to harm cockroaches just because humans find them disgusting.	0.6888	1.2232	1.6661	5.3421	4.8423
Universal donors should be obligated to donate their blood.	0.4606	1.284	1.6699	4.8386	4.805
Sport fishing to kill and eat fish is barbaric and evil.	-0.07	1.4002	2.0179	5.5038	5.1079
Private beaches are immoral, as everyone should be able to share the space.	-0.0172	0.6177	1.7373	4.6917	5.2054
Child labor in coffee bean farming is acceptable because it lowers the market price.	-0.0665	1.8448	2.5858	5.7527	4.0372
It is fine for doctors to accidentally kill a small number of patients per year.	0.3621	-0.1727	1.5847	6.0397	3.8442
Preferences					
Positive-consensus					
Using touchscreens is a much more satisfying way to interact with computers.	-0.9189	0.1251	1.8004	1.1998	6.6632
Having a drink every now and then is a good way to relax.	0.6489	1.6982	1.9433	1.7279	6.4285
Professors who play videos make their classes more entertaining.	-0.1557	1.3853	1.9827	1.2996	6.4646
Going through airport security is an unpleasant experience.	0.3005	1.2072	2.372	1.3584	6.3961
Babies that are temperamental are aggravating to spend time around.	-0.5007	0.7931	1.971	1.4424	6.2982
Afterschool programs involving sports are more fun than most of the alternatives available to children.	-0.4045	1.4263	1.9626	1.3536	6.4338
No-consensus					
Gazing at planets through a telescope is a satisfying activity.	-0.1058	0.6164	1.8833	1.3411	6.5164
Dogs are not worth the stress and aggravation it takes to own them.	0.0287	0.648	1.5868	1.4261	6.5202
The "hoots" of owls in the woods make camping more enjoyable.	-0.9634	0.5373	1.6953	1.2657	6.5675
Green apples are too sour to be an enjoyable lunchtime snack.	-0.5099	0.4378	1.7474	1.1547	6.6263
Rock music is pleasing to the ear, and much more agreeable than rap music.	0.5904	1.8393	1.4535	1.195	6.5902
Sitting in a boat and fishing all day long is boring and a waste of time.	0.3776	1.0699	1.4302	1.2174	6.665

Nothing is more awesome than driving in a Hummer.	-0.573	1.0348	1.423	1.1657	6.7799
Because sandals have fewer styles, they are less fun to go shopping for.	-0.985	0.4179	1.8624	1.2667	6.5586
In the wintertime, it is fun to catch snowflakes on the tip of your tongue.	-0.5057	0.6827	1.93	1.2962	6.3668
The Harry Potter books are engaging and delightful to read, even for adults.	0.8586	1.1436	2.0413	1.3028	6.6514
Many games are better than Monopoly, which is incredibly boring.	-0.0913	1.3158	1.4839	1.2129	6.6934
Any ice cream flavor tastes better when served in a crunchy waffle cone.	-1.2491	-1.05	1.6208	1.2306	6.5764
Negative-consensus					
Cockroaches are delicious to eat because of their hard and crunchy shell.	-0.4857	-0.5003	1.6712	1.3205	6.6532
Having blood drawn is a pleasurable experience.	-0.4634	-0.2433	1.5413	1.3137	6.3696
Nothing is more appealing than the smell of rotting fish.	-1.413	-1.4443	1.4348	1.3312	6.5875
While at a hot beach, it is agonizing to dip your toes in the cool water.	-1.3909	-0.9091	1.8264	1.2195	6.301
Drinking coffee is a miserable experience when you are tired and need energy.	-0.2406	0.8269	1.5902	1.178	6.478
Having a doctor listen attentively to your medical concerns is awful.	-0.0912	1.4404	1.4676	1.5881	6.3524
Facts					
Positive-consensus					
Touchscreens are used in a variety of electronics, including smartphones.	-1.3451	-1.1744	6.689	1.1949	1.4992
A breathalyzer is used to determine whether a driver is intoxicated.	-1.3908	-1.3893	6.6404	1.5536	1.375
University professors teach classes but also conduct research.	-0.2091	0.5427	6.3226	1.2109	1.7311
Airplanes have wings that enable the plane to lift upwards.	-1.5847	-0.7983	6.6292	1.1615	1.277
In a full-term human pregnancy, babies spend nine months in a woman's womb.	-1.1268	-0.3927	6.7811	1.3927	1.206
In sports-based afterschool programs children participate in sports such as baseball or basketball to name a few.	-1.5024	-0.9464	6.2851	1.1473	2.0334

No-consensus					
Saturn's moon, Titan, is the only moon known to have clouds.	-0.8656	-0.2341	6.7416	1.1645	1.2553
The dog breed, Basenji, is the world's only barkless dog breed.	-1.9031	-1.3962	6.7146	1.2293	1.3857
Of all types of birds, owls are the ones that can see the color blue.	-1.2773	-0.0252	6.5033	1.1704	1.221
Newtown Pippin was the first apple variety exported from the US.	-0.9766	0.02	6.5508	1.2877	1.3402
The first CD made for commercial release was the rock CD: "Born in the USA."	-0.7217	0.7578	6.3238	1.2247	1.4223
There are more fish species in the Amazon River than in the Atlantic Ocean.	-1.3589	-1.3786	6.2895	1.2685	1.2493
Hummer trucks were first marketed to civilians in 1990.	-0.9593	0.3892	6.6263	1.1715	1.3217
The oldest sandals in the world were found in Oregon's Paisley Caves.	-1.5318	-1.2682	6.5087	1.2753	1.4208
A town in North Dakota holds the world record for the tallest snowman.	-0.8145	-0.6858	6.6517	1.1539	1.2155
The author J.K. Rowling has two younger siblings, one brother and one sister.	-0.226	0.5963	6.6338	1.2753	1.1995
Monopoly pieces were made from wood, not metal, during WWI.	-0.699	0.4126	6.6678	1.1593	1.2906
The very first waffle cone was invented in Chicago, Illinois, at a state fair.	-0.8793	-0.0489	6.5895	1.2161	1.3635
Negative-consensus					
Cockroaches are a type of cold-blooded reptiles related to snakes.	-1.7082	-1.072	5.6398	1.2107	1.4144
In humans, the liver pumps blood throughout the body.	-1.8077	-1.992	5.6126	1.2347	1.2591
Fish are able to live outside of water for an extended time.	-1.1876	-0.5137	5.7055	1.4063	1.2456
The sand on beaches is usually transported there from nearby deserts.	-0.8853	0.0665	6.11	1.2452	1.2927
Coffee beans grow particularly well in freezing cold climates, such as Alaska and Russia.	-1.4569	-1.0238	5.8752	1.1975	1.4958
Medical students at hospitals are able to perform surgeries with little to no training.	-0.2611	0.9317	4.6601	1.734	1.6692

ToM BLUPs average estimates for DMPFC, PC, RTPJ, and LTPJ, as all by-stimulus random slopes were perfectly correlated. For model details, see Tables S4 and S6 of the supplemental online materials.

Appendix B. List of covariates, with descriptions.

Semantic/syntactic measures (2.2.6)		
Question name	Source	Description
Word count	Coh Metrix 3.0	Number of words in statement.
Flesch reading ease	Coh Metrix 3.0	Measures reading difficult through the average sentence length and number of syllables per word. Higher scores indicate more difficulty.
Anaphor reference	Coh Metrix 3.0	Measures the number of times a single idea is referenced by counting the use of anaphors (e.g. pronouns: he, she, it; ellipsis markers: did, was).
Intentional verb incidence	Coh Metrix 3.0	Measures intentional information by counting verbs categorized as intentional by Wordnet ratings (Fellbaum, 1998; Miller et al., 1990).
Causal verb incidence	Coh Metrix 3.0	Measures causal information by counting verbs categorized as causal by WordNet ratings.
Causal verb ratio	Coh Metrix 3.0	Measures the cohesion of causal events to actors through the ratio of causal particles (e.g. because, if) to causal verbs. Higher scores indicate increased cohesion and easier readability.
Noun concreteness	Coh Metrix 3.0	Measures concreteness of content words (e.g. chair is high in concreteness, democracy is low) using the mean concreteness ratings of content words, taken from human ratings in the MRC Psycholinguistics Database (Coltheart, 1981).
Noun familiarity	Coh Metrix 3.0	Measures the familiarity of content words using the mean familiarity ratings of all content words, taken from human ratings in the MRC Psycholinguistic Database.
Noun imageability	Coh Metrix 3.0	Measures the imageability of content words using the mean familiarity ratings of all content words, taken from human ratings in the MRC Psycholinguistic Database.
Negation density	Coh Metrix 3.0	Provides a measure of syntactic complexity (i.e. working memory load) through the count of negative expressions in the text (e.g. not, un-).
Number of modifiers	Coh Metrix 3.0	Provides a measure of syntactic complexity (i.e. working memory load) through the mean number of modifiers per noun phrase.
Left embeddedness	Coh Metrix 3.0	Provides a measure of syntactic complexity (i.e. working memory load) through the mean number of words before the main verb in a sentence.
Reaction time	In-scanner N = 25	The time from the appearance of the in-scanner agreement rating prompt to the input of a response by the participant.
Online Item Features		
Agreement	Study 1 (N = 49)	“To what extent do you agree / disagree with this statement?” (1-7; “strongly disagree”-“strongly agree”).
Valence	Online sample (N = 42)	Valence was the difference between unipolar positive and negative ratings (Kron et al., 2013), described below: <i>Instructions:</i> “Please rate your feelings regarding this statement using the following two scales. An extreme

		<p>unpleasant rating means you feel completely unpleasant, unhappy, annoyed, unsatisfied, melancholic, or despaired. An extreme pleasant rating means you feel completely pleased, happy, satisfied, content or hopeful.”</p> <p><i>Ratings:</i> Negative valence (1-8; “no unpleasant feelings”- “strong unpleasant feelings”) and positive valence (1-8; “no pleasant feelings”- “strong pleasant feelings”).</p>
Arousal	Online sample (N = 42)	<p>Arousal was the sum of unipolar positive and negative ratings, described above.</p> <p>Recent work has demonstrated that summed unipolar valence ratings are highly correlated with physiological measures of arousal, and may be superior to separately measuring arousal (Kron et al., 2013).</p>
Mental imagery	Online sample (N = 46)	<p>“To what extent did you picture or imagine what the statements described as you read?” (1-7; “Very Little”- “Very Much”; Dodell-Feder et al., 2011).</p>
Mental state	Online sample (N = 48)	<p>“To what extent did this statement make you think about someone’s experiences, thoughts, beliefs and/or desires?” (1-7; “Very Little”- “Very Much”; Dodell-Feder et al., 2011).</p>
Mental States (of Others)	Online sample (N = 44)	<p>To what extent did this statement make you think about the experiences, thoughts, beliefs, and/or desires OF OTHER PEOPLE? (1-7; “Very Little”- “Very Much”)</p>
Mental States (of Self)	Online sample (N = 46)	<p>“To what extent did this statement make you think about YOUR OWN experiences, thoughts, beliefs, and/or desires?” (1-7; “Very Little”- “Very Much”)</p>
Person present	Online sample (N = 48)	<p>“Does this statement mention people or a person?” (“Yes” / “No”).</p>

Coh Metrix ratings are calculated using an online tool at <http://cohmetrix.com> (Graesser et al., 2004; McNamara et al., 2014). In online samples, participants who did not correctly answer a catch question (asking them to describe any of the 72 statements they had read) were excluded from analysis. This caused some variability in N across covariates.

Supplemental Materials

Figure S1. We ran a supplemental analysis, predicting percent signal change (PSC) in each ROI, evoked by fact/moral/preference statements using within-subject by-trial agreement ratings (collected in-scanner), and within-subject by-trial fact-/moral-/preference-like ratings (collected after the scan session). Due to the lack of data, and the ordinal nature of the unaveraged by-trial behavioral measures, we performed a multilevel Bayesian analysis, using *brms* in R. Marginal effects for each behavioral measure, within facts, morals, and preferences, and for each ROI are plotted below, and the models are available in our online data repository (<https://osf.io/cx4dp/>). On a visual inspection, the marginal effects within the moral domain appear either to be consistent with our reported findings, or to show no effect. In our paper, we found a negative association between ToMN activity and fact-/moral-like ratings, and a positive association with preference-like ratings (Section 3.2.3; Figure 3). In this analysis, for fact-like ratings (top left), VMPFC (red) shows a negative association, $b = -0.005$, 95%CI [-0.009, -0.001]. For moral-like ratings (top right) VMPFC (red) and LTPJ (yellow) both appear to show negative associations, although neither association excluded zero from a 95% credible interval: VMPFC, $b = -0.004$, 95%CI [-0.009, 0.0002]; LTPJ, $b = -0.004$, 95%CI [-0.009, 0.0005]. For preference-like ratings (bottom left), both VMPFC (red) and DMPFC (blue) appear to show a positive association with ToMN activity; and indeed, these associations both exclude zero from a 95% credible interval: VMPFC, $b = 0.005$, 95%CI [0.002, 0.008], DMPFC, $b = 0.007$, 95%CI [0.0008, 0.013]. The only discrepancy with the results reported in the main paper body was within agreement ratings (bottom right), which showed a positive association with VMPFC among morals, $b = 0.006$, 95%CI [0.00001, 0.012]—the opposite of what we observed in our reported results. We speculate that collecting measurements on a 4-point scale, or the collection of this data within the scanner settings (as opposed to outside of it) may account for this relationship. However, it is also worth noting that a concern in the main paper was that agreement and measures of objectivity were potentially confounded among morals, and for this reason, we find the opposite pattern of results observed here for agreement promising, rather than concerning.

Figure S2. Correlation table of by-stimulus estimates.

Table S1. Study 1 condition means.

Metaethical judgments				
Model: Rating ~ Category * Rating-type * Consensus + (1+ Category * Rating-type ID) + (1 + Rating-type Item)				
Category	Rating-type	Consensus		
		Positive-consensus	No-consensus	Negative-consensus
		Mean (SE)	Mean (SE)	Mean (SE)
Morals	About Facts	3.24 (0.21) ^A	2.13 (0.18) ^B	1.83 (0.21) ^B
	About Morals	5.78 (0.28)	5.46 (0.23)	5.33 (0.28)
	About Preferences	4.28 (0.28)	4.81 (0.25)	4.59 (0.28)
Facts	About Facts	6.59 (0.17) ^A	6.59 (0.17) ^B	5.53 (0.20) ^B
	About Morals	1.31 (0.24)	1.19 (0.18)	1.35 (0.24)
	About Preferences	1.58 (0.21)	1.29 (0.16)	1.36 (0.21)
Preferences	About Facts	2.03 (0.20)	1.68 (0.17)	1.57 (0.20)
	About Morals	1.40 (0.24)	1.26 (0.18)	1.32 (0.24)
	About Preferences	6.46 (0.22)	6.60 (0.17)	6.43 (0.22)
Agreement ratings				
Model: Agreement ~ Category * Consensus + (1+ Category ID) + (1 Item)				
Category		Consensus		
		Positive-consensus	No-consensus	Negative-consensus
		Mean (SE)	Mean (SE)	Mean (SE)
Fact		6.41 (0.40)	4.76 (0.31)	2.46 (0.42)
Moral		6.05 (0.40)	4.90 (0.29)	2.66 (0.39)
Preference		4.98 (0.40)	4.08 (0.28)	1.58 (0.40)

Mean estimates and standard errors are derived from contrasts within the models described in 2.2.1 and 2.2.2.

Superscripts denote significant differences within each row (p values corrected for 27 comparisons; single-step method; $\alpha_{\text{familywise}} = .05$; single-step method).

Table S2. Study 2 in-scanner agreement ratings.

Category	Consensus		
	Positive-consensus	No-consensus	Negative-consensus
	Mean (SD)	Mean (SD)	Mean (SD)
Facts	3.64 (0.06)	2.68 (0.16)	1.47 (0.11)
Morals	3.64 (0.07)	2.81 (0.06)	1.89 (0.11)
Preferences	2.93 (0.07)	2.56 (0.06)	1.46 (0.12)

Mean and standard error (across participants). All comparisons, within content categories, were significant at $p < .001$.

Table S3. Theory of Mind network ROI coordinates.

Region	x	y	z	T score	K
R Temporoparietal Junction (RTPJ)	52	-60	24	10.55	3398
L Temporoparietal Junction (LTPJ)	-56	-56	28	9.69	3210
Precuneus (PC)	0	-52	40	10.81	2263
Ventromedial Prefrontal Cortex (VMPFC)	0	44	-20	7.69	533
Dorsomedial Prefrontal Cortex (DMPFC)	0	58	22	5.62	480
Additional significant clusters from localizer contrast					
Region	x	y	z	T score	K
R Cerebellum	32	-80	-38	7.42	425
R Inferior Frontal Gyrus	58	22	12	6.31	85
L Cerebellum	-18	-74	-34	6.30	597
L Thalamus	-8	-12	8	5.77	78
L Cerebellum	-6	-50	-50	5.70	54
R Superior Frontal Gyrus	8	38	54	5.64	60
R Middle Frontal Gyrus	30	30	52	5.60	158
R Occipital Pole	10	-96	4	5.57	626
L Middle Frontal Gyrus	-40	6	52	5.31	52
R Middle Frontal Gyrus	46	8	50	5.23	48
L Inferior Frontal Gyrus	-52	22	10	4.93	26
L Superior Frontal Gyrus	-10	36	56	4.92	61
R Lingual Gyrus	8	-48	-46	4.83	33
R Frontal Pole	16	52	36	4.74	75
L Frontal Pole	-2	62	4	4.72	19
Cingulate Gyrus	2	-16	40	4.61	17
L Occipital Fusiform Gyrus	-28	-88	-14	4.47	18
L Frontal Orbital Cortex	-42	22	-14	4.40	50
L Occipital Fusiform Gyrus	-12	-86	-16	4.30	32
Subcallosal Cortex	-4	18	-6	4.16	33
L Lateral Ventricle	-4	4	2	3.81	12

ROIs were a 9mm sphere around the reported coordinates. T scores represent difference scores in the false belief > false photograph contrast, in a random effects analysis across all subjects ($df = 24$). All coordinates are reported in MNI space. SPM contrast file is available at <https://osf.io/cx4dp/>.

Table S4. Study 1 behavioral rating model details.

ToM network activity										
Model:										
Rating ~ Category * Rating-type + (1+ Category * Rating-type ID) + (1 + Rating-type Item)										
REML criterion at convergence: 32392.5										
Dummy coded control conditions: Facts (category) & Fact-like ratings (rating-type)										
Random effects structure (by-subject)										
	Variance	St.Dev	Correlations							
			Intercept	Moral	Pref	Moral-like	Pref-like	M* M-like	M* P-like	P* MI-like
Intercept	0.08	0.89								
Moral	2.52	1.59	-0.86							
Preference	2.37	1.54	-0.89	0.94						
Moral-like	1.94	1.39	-0.95	0.84	0.90					
Preference-like	2.24	1.50	-0.97	0.88	0.94	0.98				
M*Moral-like	7.23	2.69	0.84	-0.89	-0.87	-0.88	-0.87			
M*Pref-like	6.85	2.62	0.73	-0.82	-0.84	-0.77	-0.77	0.74		
P*Moral-like	2.34	1.53	0.91	-0.94	-1.00	-0.91	-0.95	0.89	0.85	
P*Pref-like	7.70	2.78	0.89	-0.90	-0.98	-0.94	-0.94	0.88	0.86	0.98
Random effects structure (by-stimulus)										
	Variance	St.Dev	Correlations							
			Intercept	Moral-like						
Intercept	.301	.549								
Moral-like	.349	.590	-0.58							
Pref-like	.750	.866	-0.90	0.27						
Residual										
	Variance	St.Dev								
	1.10	1.05								
Fixed Effects										
	Name		B (SE)		t(df)				p	
	Intercept		6.32 (0.17)		t(101.1) = 36.85				< .001	***
	Moral		-3.99 (0.28)		t(88.2) = 14.25				< .001	***
	Preference		-4.58 (0.27)		t(90.0) = 16.70				< .001	***
	Moral-like		-5.06 (0.24)		t(79.0) = 21.38				< .001	***
	Preference-like		-4.94 (0.28)		t(98.6) = 17.59				< .001	***
	Moral*Moral-like		8.24 (0.42)		t(66.1) = 19.40				< .001	***
	Preference*Moral-like		4.63 (0.28)		t(91.9) = 16.32				< .001	***
	Moral*Preference-like		7.23 (0.45)		t(86.2) = 15.93				< .001	***
	Preference*Preference-like		9.72 (0.47)		t(82.8) = 20.57				< .001	***

St.Dev = standard deviation. *** $p < .001$; ** $p < .01$; * $p < .05$; † $p < .10$

Table S5. Study 2 behavioral rating means.

Category	Rating-type	Consensus		
		Positive-consensus	No-consensus	Negative-consensus
		Mean (SE)	Mean (SE)	Mean (SE)
Morals	About Facts	3.14 (0.24) ^A	2.30 (0.22) ^B	2.01 (0.24) ^B
	About Morals	6.35 (0.18) ^A	5.88 (0.17) ^B	5.87 (0.18) ^B
	About Preferences	4.23 (0.31)	4.31 (0.30)	4.34 (0.31)
Facts	About Facts	6.59 (0.19)	6.65 (0.18)	6.38 (0.19)
	About Morals	1.55 (0.21)	1.29 (0.20)	1.57 (0.21)
	About Preferences	1.85 (0.22) ^A	1.37 (0.21) ^B	1.68 (0.22) ^{AB}
Preferences	About Facts	2.52 (0.23)	1.95 (0.21)	2.14 (0.23)
	About Morals	2.11 (0.25)	1.74 (0.24)	1.90 (0.25)
	About Preferences	6.51 (0.19) ^A	6.58 (0.18) ^B	6.32 (0.19) ^{AB}

Mean and standard error are estimated using contrasts within the model defined in 3.2.1. Superscripts denote significant differences within each row (p values corrected for 27 comparisons; single-step method; $\alpha_{\text{familywise}} = .05$; single-step method).

Table S6. Study 2 ToM network mixed effects model.

ToM network activity										
Model:										
PSC ~ Category * ROI +										
(1+ Moral + Preference + VMPFC + PC + RTPJ + LTPJ + Moral*(VMFPC+LTPJ) ID) +										
(1 + VMPFC Item)										
REML criterion at convergence: 2271										
Dummy coded control conditions: Facts (category) & DMPFC (ROI)										
Random effects structure (by-subject)										
	Variance	St.Dev	Correlations							
			Intercept	VMPFC	PC	RTPJ	LTPJ	Moral	Pref	M* VMPFC
Intercept	.008	.091								
VMPFC	.010	.101	0.01							
PC	.008	.090	-0.60	-0.13						
RTPJ	.011	.105	-0.57	-0.20	0.71					
LTPJ	.007	.085	-0.46	-0.22	0.54	0.53				
Moral	.003	.053	-0.27	0.24	-0.34	-0.30	-0.41			
Pref	.002	.042	-0.07	0.18	-0.14	0.07	-0.11	0.52		
M*VMPFC	.002	.042	-0.14	-0.08	-0.43	-0.18	-0.41	0.62	0.66	
M*LTPJ	.002	.041	-0.33	0.30	0.59	0.35	0.64	-0.08	0.37	-0.30
Random effects structure (by-stimulus)										
	Variance	St.Dev	Correlations							
			Intercept							
Intercept	.002	.048								
VMPFC	.004	.064	.03							
Residual										
	Variance	St.Dev								
	.071	.027								
Fixed Effects										
	Name	B (SE)	t(df)	p						
	Intercept	-0.147 (0.023)	t(46) = 6.27	< .001 ***						
	Moral	0.166 (0.023)	t(128) = 7.05	< .001 ***						
	Preference	0.160 (0.022)	t(169) = 7.11	< .001 ***						
	VMPFC	0.084 (0.029)	t(51) = 2.93	.005 **						
	PC	0.005 (0.024)	t(47) = 0.20	.846						
	RTPJ	0.005 (0.026)	t(42) = 0.18	.858						
	LTPJ	0.105 (0.023)	t(45) = 4.55	< .001 ***						
	Moral*VMPFC	-0.035 (0.030)	t(97) = 1.18	.241						
	Moral*PC	-0.059 (0.022)	t(844) = 2.68	.007 **						
	Moral*RTPJ	-0.113 (0.022)	t(844) = 5.12	< .001 ***						
	Moral*LTPJ	-0.068 (0.024)	t(162) = 2.88	.004 **						
	Pref*VMPFC	-0.074 (0.029)	t(114) = 2.56	.012 *						
	Pref*PC	-0.123 (0.022)	t(844) = 5.56	< .001 ***						
	Pref*RTPJ	-0.152 (0.022)	t(844) = 6.89	< .001 ***						
	Pref*LTPJ	-0.11 (0.022)	t(844) = 5.05	< .001 ***						

St.Dev = standard deviation. *** $p < .001$; ** $p < .01$; * $p < .05$; † $p < .10$

Table S7. Study 2 ToM-behavioral analysis.

DV: Metaethical judgments (Study 1 BLUPs)			
Term		<i>F statistic</i>	<i>p</i>
PSC x 2 (ROI: ToM/VMPFC) x 3 (rating-type (fact-/moral-/preference-like) x 3 (category: fact/moral/preference))		$F(2, 405) = 5.73$	< .001 ***
PSC x 2 ROI x 3 rating-type x 2 (fact/preference)		$F(2, 270) = 0.01$	.986
PSC x 2 ROI x 3 rating-type x 2 (fact/moral)		$F(2, 270) = 7.42$	< .001 ***
PSC x 2 ROI x 3 rating-type x 2 (moral/preference)		$F(2, 270) = 6.64$	.002 **
Within moral statements			
PSC x 2 ROI x 3 rating-type		$F(2, 135) = 6.03$	.003**
PSC x 2 ROI x 2 (fact-like/preference-like)		$F(1, 90) = 13.77$	< .001 ***
PSC x 2 ROI x 2 (fact-like/moral-like)		$F(1, 90) = 0.81$	.371
PSC x 2 ROI x 2 (moral-like/preference-like)		$F(1, 90) = 5.47$	.022 *
Within preference-like ratings			
PSC x ROI		$F(1, 45) = 4.57$	.038 *
Within fact-/moral-like ratings			
PSC x ROI		$F(1, 92) = 8.80$	.004 **
Model: rating-type + (PSC x ROI) + (PSC x ROI x preference-like)	<i>B (SE)</i>	<i>t statistic</i>	<i>p</i>
Intercept (Fact-like rating)	2.57 (0.11)	$t(137) = 22.80$	< .001 ***
Moral-like	3.17 (0.14)	$t(137) = 23.00$	< .001 ***
Preference-like	1.80 (0.17)	$t(137) = 10.62$	< .001 ***
PSC (within fact-like/moral-like)	-1.01 (0.23)	$t(137) = 4.32$	< .001 ***
PSC x preference-like (within ToM)	1.94 (0.40)	$t(137) = 4.81$	< .001 ***
PSC x ROI (interaction for VMPFC, within fact-/moral-like)	0.72 (0.23)	$t(137) = 3.11$	.002 **
PSC x preference-like x ROI (interaction for VMPFC, for preference-like)	-1.35 (0.40)	$t(137) = 3.35$	.001 **
Contrasts:			
Fact-/moral-like-ToM relationship	-1.01 (0.23)	$t(140) = 4.32$	< .001 ***
Preference-like-ToM relationship	0.94 (0.33)	$t(140) = 2.85$	.020 *
Fact-/moral-like-VMPFC relationship	-0.28 (0.09)	$t(140) = 3.18$	.007 **
Preference-like-VMPFC relationship	0.31 (0.12)	$t(140) = 2.48$	.055 †

Contrast *p* values corrected for 4 comparisons; single-step method; $\alpha_{\text{familywise}} = .05$; single-step method. *** *p* < .001; ** *p* < .01; * *p* < .05; † *p* < .10

Table S8. Whole brain correlation peak coordinates.

Contrast	Name	Cluster Size	Peak T	x	y	z
Fact-like rating (negative)	M Superior frontal gyrus	1037	7.60	-8	24	58
			6.76	-18	20	62
			5.40	16	44	48
	L Middle frontal gyrus	658	7.22	-28	10	50
			6.24	-42	12	44
			5.33	-48	28	24
	L Angular gyrus	398	6.53	-44	-62	-48
			5.60	-38	-72	48
			5.23	-34	-62	48
	L Middle temporal gyrus	142	6.51	-58	-32	-16
			5.57	-66	-26	-12
	R Angular gyrus	412	6.43	44	-68	46
			5.76	50	-56	32
			5.74	52	-62	40
	R Middle frontal gyrus	239	5.58	40	20	44
			5.14	32	28	42
	L Medial caudate nucleus	132	5.18	-14	14	10
			5.16	-12	6	14
Fact-like rating (positive)	M Parietooccipital sulcus	145	4.85	8	-78	36
			4.83	-6	-80	18
			4.14	6	-84	44
Preference-like rating (positive)	R Angular gyrus	124	5.12	54	-60	34
			4.84	50	-62	42
			3.95	44	-68	46
	L Angular gyrus	102	4.98	-36	-70	48
			4.36	-30	-64	54
			4.14	-30	-76	46

First level models produced a beta map for each item, for each participant. For each participant, 3 models, predicting by-stimulus estimates were created, with fact-like, moral-like, and preference-like ratings as respective predictors. Beta maps for ratings from each model were entered into a random effects analysis across all participants. Permutation tests (5000 samples) were used to achieve a cluster-corrected familywise error rate of $\alpha = .05$ in each contrast, while thresholding voxels at $p < .001$. Permutation testing was performed using SnPM 13 (<http://warwick.ac.uk/snpm>; Nichols & Holmes, 2001). All coordinates reported in MNI space.

Table S9. Study 2 ToM-behavioral analysis: controlling for semantic/syntactic features

Model:			
PSC ~ Category*ROI*NounConcreteness + ROI*LeftEmbeddedness + NounFamiliarity + (1+ Moral + Preference + VMPFC + PC + RTPJ + LTPJ + Moral*(VMFPC+LTPJ) ID) + (1 + VMPFC Item)			
REML criterion at convergence: 2378.2			
Dummy coded control conditions: Facts (category) & DMPFC (ROI)			
DV: Metaethical judgments (Study 1 BLUPs)			
PSC corrected for syntactic/semantic features			
Term		<i>F statistic</i>	<i>p</i>
PSC x 2 (ROI: ToM/VMPFC) x 3 (rating-type (fact-/moral-/preference-like) x 3 (category: fact/moral/preference)		$F(4, 405) = 0.86$	.486
PSC averaged across ROI			
PSC x 3 rating-type x category		$F(4, 198) = 9.54$	<.001 ***
PSC x 3 rating-type x category (fact/preference)		$F(2, 132) = 0.61$	.542
PSC x 3 rating-type x category (fact/moral)		$F(2, 132) = 9.81$	<.001 **
PSC x 3 rating-type x category (moral/preference)		$F(2, 132) = 14.17$	<.001 **
Within moral statements			
PSC x 3 rating-type		$F(2, 66) = 9.89$	<.001 ***
PSC x 2 (fact-like/preference-like)		$F(1, 44) = 18.20$	<.001 ***
PSC x 2 (fact-like/moral-like)		$F(1, 44) = 0.001$	.972
PSC x 2 (moral-like/preference-like)		$F(1, 44) = 13.97$	<.001 ***
Model: rating-type + PSC + (PSC x preference-like)	<i>B (SE)</i>	<i>t statistic</i>	<i>p</i>
Intercept (Fact-like rating)	2.18 (0.14)	$t(67) = 15.31$	<.001 ***
Moral-like	3.18 (0.19)	$t(67) = 16.50$	<.001 ***
Preference-like	2.61 (0.21)	$t(67) = 12.72$	<.001 ***
PSC (within fact-like/moral-like)	-0.97 (0.26)	$t(67) = 3.79$	<.001 ***
PSC x preference-like	1.99 (0.44)	$t(67) = 4.48$	<.001 ***
Contrasts:	<i>B (SE)</i>	<i>t statistic</i>	<i>p</i>
Fact-/moral-like-ToM relationship	-0.97 (0.26)	$t(70) = 3.79$	<.001 ***
Preference-like-ToM relationship	1.02 (0.36)	$t(70) = 2.81$	.013 *

Contrast *p* values corrected for 2 comparisons; single-step method; $\alpha_{\text{familywise}} = .05$; single-step method. *** $p < .001$;

** $p < .01$; * $p < .05$; † $p < .10$

Table S10. Model simplification for item features.

DV: Agreement (Study 1 BLUPs)			
Term		F statistic	p
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		$F(2, 135) = 1.27$	.284
PSC x 2 (ROI)		$F(1, 137) = 0.30$	.587
PSC averaged across ROI			
PSC x 3 (category)		$F(2, 66) = 5.92$	.004 **
PSC x 2 (category: moral/fact)		$F(1, 44) = 3.71$	.061 †
PSC x 2 (category: moral/preference)		$F(1, 44) = 11.78$	.001 **
PSC x 2 (category: fact/preference)		$F(1, 44) = 2.32$	.135
Model: PSC x 3 (category)	B (SE)	t statistic	p
Contrast: PSC within Facts	0.02 (0.50)	$t(69) = 0.04$	1.00
Contrast: PSC within Morals	-1.46 (0.59)	$t(69) = 2.59$	.037 *
Contrast: PSC within Preferences	1.01 (0.44)	$t(69) = 2.30$	.072 †
DV: Mental State General (Online Sample BLUPs)			
Term		F statistic	p
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		$F(2, 135) = 1.36$	.261
PSC x 2 (ROI)		$F(2, 137) = 0.04$	.844
PSC averaged across ROI			
PSC x 3 (category)		$F(2, 66) = 5.92$	.004 **
PSC x 2 (category: moral/fact)		$F(1, 44) = 5.07$	.029 *
PSC x 2 (category: moral/preference)		$F(1, 44) = 12.33$	.001 **
PSC x 2 (category: fact/preference)		$F(1, 44) = 1.06$	.309
Model: PSC x 3 (category)	B (SE)	t statistic	p
Contrast: PSC within Facts	0.20 (0.12)	$t(66) = 1.65$	.275
Contrast: PSC within Morals	-0.24 (0.14)	$t(66) = 1.70$	.253
Contrast: PSC within Preferences	0.37 (0.11)	$t(66) = 3.39$	.004 **
DV: Mental State Self-Oriented (Online Sample BLUPs)			
Term		F statistic	p
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		$F(2, 135) = 0.43$	.689
PSC x 2 (ROI)		$F(1, 137) = 0.12$	.720
PSC averaged across ROI			
PSC x 3 (category)		$F(2, 66) = 3.96$	.024 *
PSC x 2 (category: moral/fact)		$F(1, 44) = 0.44$	.511
PSC x 2 (category: moral/preference)		$F(1, 44) = 6.96$	.011 *
PSC x 2 (category: fact/preference)		$F(1, 44) = 4.16$	.047 *
Model: PSC x 3 (category)	B (SE)	t statistic	p
Contrast: PSC within Facts	-0.12 (0.23)	$t(66) = 0.55$	.928
Contrast: PSC within Morals	-0.37 (0.27)	$t(66) = 1.37$	.436
Contrast: PSC within Preferences	0.51 (0.21)	$t(66) = 2.48$	.047 *
DV: Mental State Other-Oriented (Online Sample BLUPs)			
Term		F statistic	p
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		$F(2, 135) = 0.52$	.596
PSC x 2 (ROI)		$F(1, 137) = 0.09$	.763
PSC averaged across ROI			
PSC x 3 (category)		$F(2, 66) = 3.48$	.037 *
PSC x 2 (category: moral/fact)		$F(1, 44) = 2.34$	.133
PSC x 2 (category: moral/preference)		$F(1, 44) = 9.30$	.004 **
PSC x 2 (category: fact/preference)		$F(1, 44) = 0.83$	.367
Model: PSC x 3 (category)	B (SE)	t statistic	p
Contrast: PSC within Facts	0.23 (0.19)	$t(66) = 1.20$	.547
Contrast: PSC within Morals	-0.26 (0.22)	$t(66) = 1.19$	.552
Contrast: PSC within Preferences	0.47 (0.17)	$t(66) = 2.77$	.022 *
DV: Person Present (Online Sample BLUPs)			
Term		F statistic	p
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		$F(2, 135) = 0.61$	.546
PSC x 2 (ROI)		$F(1, 137) = 0.22$	.637
PSC averaged across ROI			
PSC x 3 (category)		$F(2, 66) = .04$	.962
Model: DV ~ category + PSC	B (SE)	t statistic	p

Main effect: PSC (<i>within</i> moral/preference/fact)		1.49 (0.65)	<i>t</i> (68) = 2.28	.026 *
DV: Mental Imagery (Online Sample BLUPs)				
Term		<i>F statistic</i>	<i>p</i>	
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		<i>F</i> (2, 135) = 0.31	.732	
PSC x 2 (ROI)		<i>F</i> (1, 137) = 0.07	.797	
PSC averaged across ROI				
PSC x 3 (category)		<i>F</i> (2, 66) = 0.39	.680	
Model: DV ~ category + PSC		<i>B (SE)</i>	<i>t statistic</i>	<i>p</i>
Main effect: PSC (<i>within</i> moral/preference/fact)		-0.13 (0.11)	<i>t</i> (68) = 1.15	.253
DV: Arousal (Online Sample BLUPs)				
Term		<i>F statistic</i>	<i>p</i>	
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		<i>F</i> (2, 135) = 0.11	.892	
PSC x 2 (ROI)		<i>F</i> (1, 137) = 0.26	.611	
PSC averaged across ROI				
PSC x 3 (category)		<i>F</i> (2, 66) = 0.35	.703	
Model: DV ~ category + PSC		<i>B (SE)</i>	<i>t statistic</i>	<i>p</i>
Main effect: PSC (<i>within</i> moral/preference/fact)		-0.05 (0.06)	<i>t</i> (68) = 0.77	.443
DV: Valence (Online Sample BLUPs)				
Term		<i>F statistic</i>	<i>p</i>	
PSC x 3 (category: fact/moral/preference) x 2 (ROI: ToM/VMPFC)		<i>F</i> (2, 135) = 0.13	.868	
PSC x 2 (ROI)		<i>F</i> (1, 137) = 0.03	.857	
PSC averaged across ROI				
PSC x 3 (category)		<i>F</i> (2, 66) = 0.03	.857	
Model: DV ~ category + PSC		<i>B (SE)</i>	<i>t statistic</i>	<i>p</i>
Main effect: PSC (<i>within</i> moral/preference/fact)		0.51 (0.47)	<i>t</i> (68) = 1.09	.281
*** <i>p</i> < .001; ** <i>p</i> < .01; * <i>p</i> < .05; † <i>p</i> < .10				

*** $p < .001$; ** $p < .01$; * $p < .05$; † $p < .10$