

Are corporations people too? The neural correlates of moral judgments about companies and individuals

Mark Plitt^{1,2}, Ricky R. Savjani^{1,3}, and David M. Eagleman^{1,4}

¹Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA

²Interdisciplinary Program in Cognitive Science, Rice University, Houston, TX 77005, USA

³College of Medicine, Texas A&M Health Science Center, Bryan, TX 77807, USA

⁴Department of Psychiatry, Baylor College of Medicine, Houston, TX 77030, USA

To investigate whether the legal concept of “corporate personhood” mirrors an inherent similarity in the neural processing of the actions of corporations and people, we measured brain responses to vignettes about corporations and people while participants underwent functional magnetic resonance imaging. We found that anti-social actions of corporations elicited more intense negative emotions and that pro-social actions of people elicited more intense positive emotions. However, the networks underlying the moral decisions about corporations and people are strikingly similar, including regions of the canonical theory of mind network. In analyzing the activity in these networks, we found differences in the emotional processing of these two types of vignettes: neutral actions of corporations showed neural correlates that more closely resembled negative actions than positive actions. Collectively, these findings indicate that our brains understand and analyze the actions of corporations and people very similarly, with a small emotional bias against corporations.

Keywords: Corporations; Moral decision-making; MVPA; Theory of mind.

As individuals organize into groups, the collective unit gains its own representation. For example, corporations such as Microsoft and AT&T come to embody a distinct entity, separate from the people who comprise them. Little is known about how our brains process information about collective units such as corporations. As an organization comes to form an identity, a question arises: are corporations and their actions regarded as social beings or as inanimate objects? Though this question has been largely ignored in the cognitive neuroscience literature, it has become an important topic in legal systems and public opinion. For example, the American legal

system has extended the rights of individuals to corporations (“corporate personhood”) and held corporations, as a collective unit, liable (“corporate liability”). In a controversial decision in *Citizens United v. Federal Election Commission* (2010), the United States Supreme Court granted corporations the right to free speech. This decision stemmed from the court’s opinion that the rights of individual citizens, ensured under the First and Fourteenth Amendments, extend to corporations as well (Kang, 2010). More recently, the United States Supreme Court further extended corporate rights, allowing for-profit organizations to be exempt from a law that violates the

Correspondence should be addressed to: David M. Eagleman, Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA. E-mail: david@eaglemanlab.net

We thank the Baylor College of Medicine Computational and Integrative Biomedical Research (CIBR) center for providing software licenses and access to cluster computing resources. We also thank two anonymous reviewers for improving the interpretation and discussion of our findings.

We thank the Initiative on Neuroscience and Law (NeuLaw.org) for funding.

The authors do not declare any financial or other conflicts of interest.

Mark Plitt and Ricky R. Savjani contributed equally to this study.

owners' religious beliefs (*Burwell v. Hobby Lobby*, 2014).

The recent legal and popular attention to this topic motivates a further question: does corporate personhood parallel an underlying similarity in the way human brains represent corporations and individuals? Indirect evidence suggests that the mechanisms that determine our reactions to the behavior of large groups and to the behavior of individuals could overlap or even be identical. For example, people are equally likely to reciprocate favorably if given positive reinforcement from either a corporation (e.g., corporate sponsorship) (Harvey, Kirk, Denfield, & Montague, 2010) or an individual (e.g., cooperation in game theory) (Rilling et al., 2002). Further, overlapping networks of brain activity are engaged whether the reciprocity is toward corporations (Harvey et al., 2010) or people (Rilling et al., 2002).

On the other hand, social reasoning areas of the brain are observed to decrease in activity if an entity is objectified. For example, the superior medial prefrontal cortex (SMPFC), a key component in social cognition (Amodio & Frith, 2006), is less active in impersonal moral or non-moral reasoning than personal moral reasoning (Greene, 2001) and, interestingly, while viewing dehumanizing images of homeless people and drug addicts (Harris & Fiske, 2006) or highly sexualized images of women (Cikara, Eberhardt, & Fiske, 2011).

Collectively, these studies raise an open question of whether people objectify corporations or view them as a social being. No studies to date have directly compared how the brain responds when judging positive, negative, or neutral actions of corporations and the actions of individuals. To this end, we performed a vignette-based functional magnetic resonance imaging (fMRI) experiment. While in the scanner, participants rated their emotional response to short vignettes about the actions of people or corporations. In these vignettes, people or corporations performed pro-social actions (e.g., donating to charity), anti-social actions (e.g., lying or breaking the law), or neutral actions (e.g., buying a printer). To establish a control condition, a third category of vignettes described objects.

The brain invokes reciprocal inhibition, allowing people to think either socially or objectively (Jack et al., 2012). Moral emotional decision-making tasks have been shown to strongly activate circuits involved with theory of mind (ToM) (Greene & Haidt, 2002; Hein & Singer, 2008; Jack et al., 2012; Knabb, Welsh, Erickson, & Felthous, 2009; Krauss, 2010; Mar, 2011; Moll et al., 2002; Moll, de Oliveira-Souza, & Zahn, 2008), representing the thoughts and actions of other people. During non-social tasks, however, regions of

the ToM network deactivate, and regions associated with working memory and objective reasoning such as the dorsolateral prefrontal cortex and large portions of the parietal lobe including the intraparietal sulcus become more active (Greene, 2001; Jack et al., 2012). Given this apparent mutually exclusive framework, we hypothesized the actions of corporations to evoke inherent social reasoning processes. Specifically, we hypothesized that participants would not only rate the actions of corporations and people similarly but also that the neural responses to reading vignettes of these two types would be very similar. Further, we hypothesized that vignettes about both people and corporations would elicit greater activity in the ToM network than vignettes about objects.

RESULTS

Behavior

All 40 participants who successfully completed the experiment were kept in the behavioral analysis. In the scanner, participants correctly answered questions about the vignettes with an accuracy of 87%, indicating appropriate attentiveness and comprehension of the vignettes. Vignette reading times for the three conditions of *agency* were as follows: *person* $\mu = 12.7 \pm 4.6$ s, *company* $\mu = 14.0 \pm 4.9$ s, and *object* $\mu = 13.2 \pm 5.3$ s. Although the difference between reading times was significant ($F = 9.06$, $p < .001$), this difference is less than a single TR (2 s), suggesting the difference should not significantly affect our results. Participants answered in line with expectations about the intended *valences* of the vignettes (**Figure S1**).

We then evaluated the differences in reported emotional ratings and intensity as a function of agency. Emotional ratings of vignettes significantly differed across the agency categories, showing that participants describe their emotional responses to vignettes differently for each of the three agency categories ($\chi^2 = 589.15$, $p < 10e-116$; **Figure 1A**). All pairwise comparisons (e.g., person vs. company) showed significant differences in emotional ratings as well. These post hoc comparisons were made using a Bonferroni correction for the repeated tests ($p < .0025$). Since object vignettes were not designed to have valence manipulation, we only compared intensity ratings for company and person trials. There was no main effect of *agency* category on emotional intensity ($F = .02$, $p = .88$; **Figure 1B**). There was, however, a significant main effect of *valence* on emotional intensity ratings ($F = 62.7$, $p < .001$) and a significant interaction




Figure 1. Participants reported different emotional ratings to person and company vignettes. **A.** Histograms of emotional responses for all three *agency* manipulations. Emotional responses significantly differed across agent categories ($\chi^2 = 589.15$, $p < 10e-116$). All pairwise comparisons were also significant (Bonferroni corrected $p < .01$). **B.** Ratings of emotional intensity are plotted against intended *valence* category for person and company trials separately. Intensity ratings show a main effect of vignette *valence* ($F = 62.7$, $p < .001$) and an interaction of vignette *valence* and *agency* ($F = 3.36$, $p < .05$). There was no main effect of *agency* on emotional intensity.

between the *valence* and *agency* for emotional intensity ($F = 3.36$, $p < .05$; Figure 1B). In this crossover interaction, we find that pro-social actions performed by people evoke a more intense positive emotional response than pro-social actions by companies, and anti-social actions by companies elicit more intense negative emotional responses than the same actions performed by people. This finding gives evidence for a negativity bias toward companies on which we elaborate in a later portion of the results (Figure 4).

fMRI

To determine regions responsible for the processing of *agency*, we contrasted GLM β estimates for person versus object and company versus object (Figure 2 and Table S2). Areas that showed greater activation to companies or persons more than to objects included the bilateral superior temporal pole (TP), bilateral angular gyrus (AG), posterior cingulate cortex (PCC), SMPFC, and bilateral cerebellum. An additional region in the midbrain responded with greater activity to person than to object vignettes. Several regions in the left hemisphere showed greater activity

to objects than to either person or company vignettes: the left inferior frontal gyrus (L IFG), the left inferior parietal lobule (L IP), and the left inferior temporal gyrus (L ITG). In addition, left pars opercularis (L PO) showed greater activity to objects than to person vignettes. Blood-oxygen-level dependent (BOLD) traces drawn from these regions reveal that the significance of the TP, AG, and possibly SMPFC are due to an increase in BOLD activity for persons and companies, while the significance of L PO and L ITG are due to an increase in BOLD activity for object vignettes. In the PCC, however, a depression of activity in the object condition drives significance, and in the L IFG and L IP there is a depression of the BOLD response for the person and company conditions.

No regions survive our non-parametric permutations testing method for family-wise error (FWE) correction in the direct comparison of persons and companies.

To ferret out any more subtle differences between persons and companies, we performed a multi-voxel pattern analysis (MVPA). We attempted to decode corporation and person trials using a 2-class support vector machine (SVM) combining all valences and feeding the 9 time points around the button press to

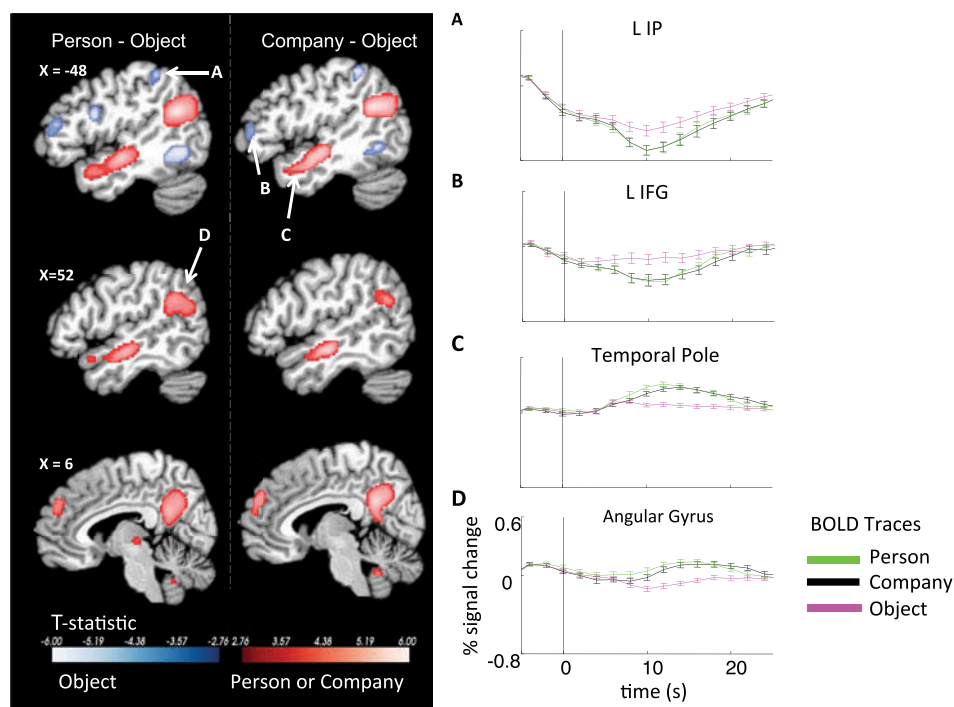


Figure 2. Vignettes about persons and companies elicit similar activity within social networks when compared to stories about objects. The left column of brain images depicts the contrasts of person > object (red colors) and object > person (blue colors). The right column of brain images depicts the contrasts of company > object (red colors) and object > company (blue colors). All contrasts are shown with family-wise error (FWE) correction at $p < .05$. BOLD traces drawn from example regions of interest (ROIs) are shown to the right. The person and company BOLD activity show similar patterns in all ROIs. Time = 0 s corresponds to the vignette onset.

the SVM. However, our overall prediction accuracy was only 54.93% (chance = 50%, $p = .0082$, χ^2), indicating *agency* was difficult to decode. Thus, we ran a separate 3-class SVM in an attempt to decode emotional *valence* separately in person and company conditions. The overall prediction accuracy for both person and company trials was computed by using all 9 time points around the button press in an SVM, and the *valence* was decoded. This yielded a prediction accuracy for person trials of 50.60% (chance = 33.33%, $p < 7.6e-11$, χ^2) and a prediction accuracy for the company trials of 44.07% (chance = 33.33%, $p < 2.8e-13$, χ^2). We then computed a time series analysis decoding using each time point separately. This showed that the peak prediction accuracy occurred 4 s before the vignette ended (Figure 3A). Although the time series analysis shows similar decoding accuracy for person and company trials, the confusion matrices for these SVMs provides further evidence for differential emotional processing of these two trial types as seen in the analysis of the emotional responses (Figure 3B). The prediction accuracies for negative, neutral, and positive vignettes within the person condition were of similar

magnitudes. However, for company vignettes, the SVM was much more successful in classifying positive vignettes than negative or neutral company vignettes. SVM sensitivity maps for the 3-class SVMs for company and person trials at 4 s before the participant finished reading the vignette are shown in Figure S2.

To determine the regions that may underlie this difference in decoding accuracy for person and company emotions, especially the neutral vignettes, we returned to the GLM analysis and collapsed person and company vignettes into one group and then separated the *valence*. Collectively the positive versus neutral and negative versus neutral contrasts yielded the same regions as the positive versus negative contrasts, so we will consider only the latter comparison (Figure 4). The right dorsolateral prefrontal cortex, right anterior cingulum, PCC, precuneus, mid-cingulum, right posterior parietal lobe, and bilateral cerebellar tonsil showed greater activation to positively *valenced* vignettes. Only a small cluster of the superior temporal gyrus showed greater activity to negative vignettes. BOLD traces in these regions show no significant differences between the company and person trials within either the positive or negative

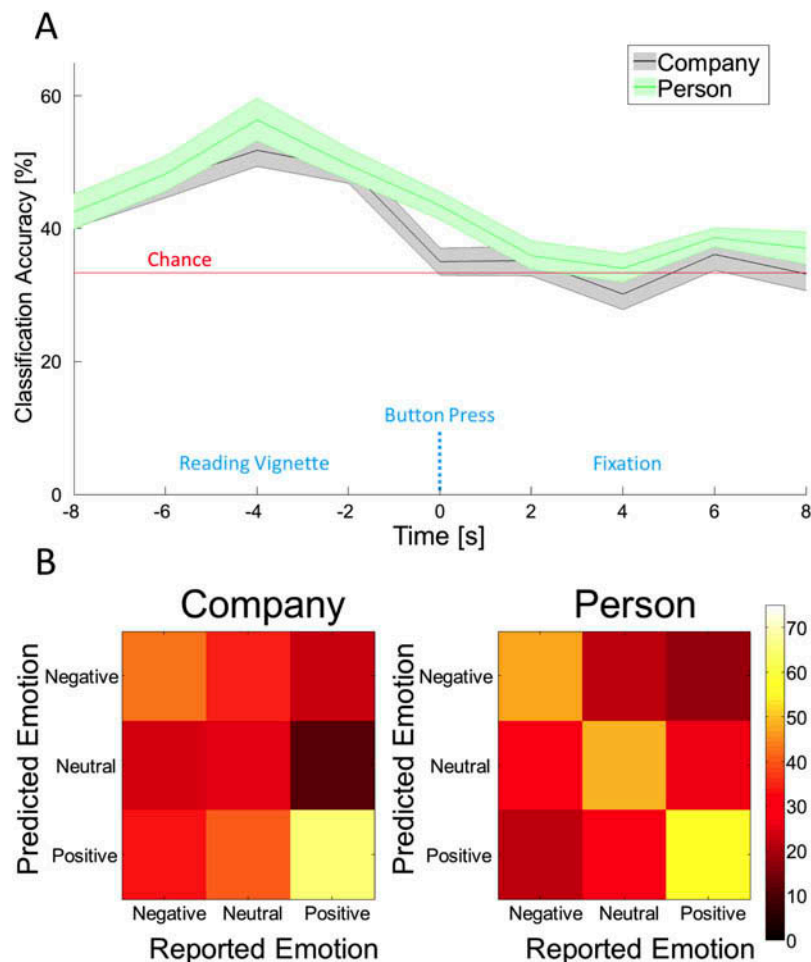


Figure 3. Neural decoding of negative and neutral emotions is weak in Company trials but not in Person trials. **A.** Time series analysis of decoding accuracy is similar in person and company trials. Peak prediction accuracy occurs 4 seconds before participants finished reading vignettes. Shaded regions indicate ± 1 SEM. Horizontal red line indicates chance level performance (33.33%). **B.** Confusion matrices for person and company trials suggest a difference in emotional processing. Columns indicate the true *valence* condition of the trial, and the rows represent the SVM-predicted *valence* of the trial. Colors indicate the percent of trials that were classified with that label. The diagonals of the graph indicate the proportion correct in each trial type.

trial subtypes. Within the neutral vignettes, however, person vignettes showed consistently higher BOLD activity than company vignettes. The BOLD activity from neutral person vignettes resembled the BOLD activity for positive vignettes, and the BOLD activity for neutral company vignettes resembled BOLD activity from negative vignettes, providing additional evidence for a negativity bias toward companies on a neural level.

DISCUSSION

We found that the networks that underlie our ability to understand the actions of companies and people substantially overlap. That is, corporations are

neurally represented as social beings rather than inanimate objects. Vignettes about corporations and people gave rise to largely similar neural responses, but there appears to be important differences in the emotional processing of these stories. Varying emotional responses between these vignette subtypes taken together with neural decoding using SVM classification and region of interest analysis suggests a negativity bias toward corporations. Not only did we see a bias against companies when rating positive and negative vignettes, but the neural response during neutral vignettes suggested a predisposition to judge companies more negatively. The neural response to a neutral vignette about a company was more similar to a negative vignette

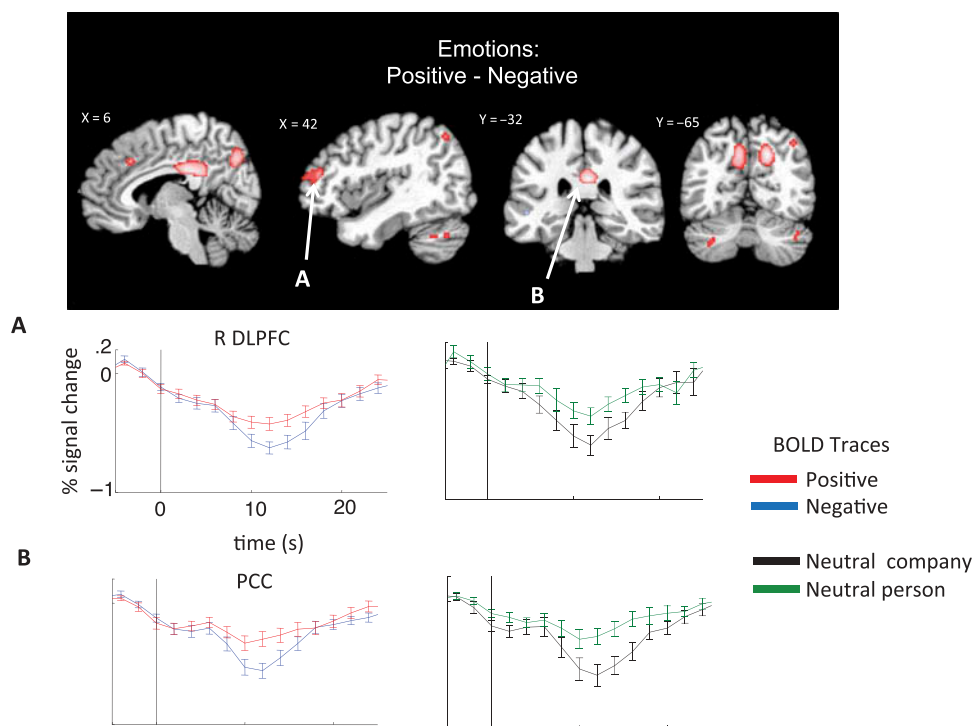


Figure 4. No differences are seen between company and person *agency* when participants read positive and negative vignettes. However, neutral person trials consistently elicit significantly greater activity than neutral company trials. A contrast of all positive > negative trials revealed large medial regions such as the anterior cingulate cortex (ACC), posterior cingulate cortex (PCC), and the precuneus (FEW < .05). Traces are drawn from the labeled ROIs. The plots on the left depict for each ROI negative and positive BOLD activity averaged for person and companies. The plots on the right depict for each ROI neutral vignette BOLD activity for person and companies. Time = 0 s corresponds to vignette onset. Positive and negative as well as neutral company and person BOLD traces shown are significantly different by repeated measures ANOVA.

than to a positive vignette. This bias was not found in the neutral person vignettes.

When contrasted with objects, people and company both elicited activity in a widely cited “mentalizing network” of brain regions responsible for representing the actions and thoughts of others. A meta-analysis of 63 ToM studies finds a strikingly similar network to our agent > object contrasts (Figure 2), including the SMPFC, PCC/precuneus, bilateral TPJ, bilateral STS, and bilateral TP (Mar, 2011). A similar network of regions appears crucial for social cognition and moral decision-making (Greene & Haidt, 2002; Hein & Singer, 2008; Moll et al., 2008).

Activation of this mentalizing network suggests that people apply similar mechanisms of moral reasoning and perspective taking when evaluating the actions of corporations and people alike. The brain can operate under two distinct states with antagonistic networks: a “social reasoning” network (e.g., interacting with other people) and “mechanical reasoning” (e.g., doing physics problems) (Jack et al., 2012). Our results show that both person vignettes and

corporation vignettes activate this social reasoning network, while vignettes about objects activate the mechanical reasoning network. In other words, people do not treat corporations as physical entities but rather as social beings. Further, this social reasoning network allows people to form representations of the others, giving rise to decisions about morality and justice (Robertson et al., 2007).

Other neuroimaging evidence has further characterized this mentalizing brain network beyond the simple social versus non-social distinction. Activation of the thalamus and ventromedial/orbitofrontal activation has been implicated in particularly negative moral emotions (e.g., disgust) (Moll et al., 2005). Additionally the cerebellar regions seen in our agent > object contrasts (Figure 2) have been shown to be functionally connected to the other regions of the mentalizing network (Buckner, Krienen, Castellanos, Diaz, & Yeo, 2011). The sum of the functions within these regions gives rise to the ability to understand the personal and social behavior in both individuals and corporations alike. Our strongest

evidence of corporate personhood stems from activation of the SMPFC in both corporations and individuals when each was contrasted to objects (Figure 2). The medial prefrontal cortex is critical to many social functions such as discriminating emotion, processing reward and punishment, representing and updating the value of future outcomes, and predicting the mental states of other organisms and cartoons (Amodio & Frith, 2006; Mitchell, Banaji, & Macrae, 2005). Crucially, corporation vignettes did not down-regulate the SMPFC, as dehumanization studies have found (Cikara et al., 2011; Harris & Fiske, 2006). Our results show that vignettes involving corporations activate this frontal region just as person vignettes do.

Our MVPA investigated whether emotional judgments about corporations and individuals could be decoded in patterns of activity distributed spatially across the brain. We found that prediction accuracies were highest in company vignettes for positive trials, whereas in people all three emotions were decoded with similar accuracies (Figure 3). This suggests that the brain might respond uniquely to pro-social behavior of a company. Alternatively, the neutral and negative vignettes might elicit indistinguishable neural responses in company trials. Further evidence for the latter interpretation was seen in the GLM analysis (Figure 4). In regions captured by the positive versus negative contrasts, neither positive nor negative vignette BOLD traces differ between people and corporations; however, the neutral traces drawn from these regions are significantly different. The person neutral traces resemble the positive traces, while the company neutral traces resemble the negative traces. Collectively, these results give neural evidence for a negativity bias toward corporations that we see in our analysis of participants' emotional ratings. Such biases toward corporations have been previously noted in the literature outside of neuroscience. For example, participants in psychology studies view unethical behavior by corporations as more predictive of a corporation's future behavior than refraining from unethical behavior or displaying pro-social behavior (Folkes & Kamins, 1999). The bias we find here provides a basis for further investigation of emotional processing in decision-making about corporations.

It is possible that in an alternative experiment, the ToM network could be activated for inert objects. Several studies have been able to ascribe human behavior to non-human objects (e.g., walking with point-light patterns) (Beauchamp, Lee, Haxby, & Martin, 2003; Grossman & Blake, 2001; Grossman et al., 2000), and the ToM network activates when

these inanimate objects engage in social interactions (Martin & Weisberg, 2003). This evidence suggests that the brain is capable of abstracting constructs and developing high-level frameworks of social interaction. Our evidence supports this model, suggesting that corporations can also be abstracted and mentalized as social beings.

Regions that displayed greater activation to objects (L IFG, L IP, L ITG, and L PO) than to either people or companies have been seen a variety of studies. For example, similar regions are activated in non-social reasoning including mechanical reasoning (Jack et al., 2012), non-moral reasoning (Greene, 2001), motion discrimination and multisensory integration during tactile stimulation (Pasalar, Ro, & Beauchamp, 2010; van Boven, Ingeholm, Beauchamp, Bikle, & Ungerleider, 2005), visual and non-visual working memory (Pessoa, Gutierrez, Bandettini, & Ungerleider, 2002) (for review see D'Esposito, Postle, & Rypma, 2000), and semantic and syntactic reasoning (Costafreda et al., 2006; Friederici, Ruschemeyer, Hahne, & Fiebach, 2003; Hirshorn & Thompson-Schill, 2006; Moss et al., 2005; Poldrack et al., 1999). The left IP and in particular the left intraparietal sulcus have additionally been implicated in tasks involving object manipulation. Imagined, observed, and pantomimed tool use as well as point-light displays of tools preferentially activates the left intraparietal sulcus (Beauchamp et al., 2003; Moll et al., 2000). Vignettes from the object condition that contain information about tools and other objects (e.g. scissors, staplers, tractors) may be responsible for the observed activation of this region.

Limitations and future directions

Our study gives rise to two additional questions to further probe the neural and emotional responses to the actions of corporations versus individuals. First, our study took a broad approach to observe the effects of three different *agents* (person, company, and object) under three different *valences* (positive, negative, and neutral) in a single fMRI run. Future work will examine negative corporation and person trials to further elucidate our observed negativity bias toward corporations. To this end, we will use multiple short runs to generate robust classifiers for MVPA (Coutanche & Thompson-Schill, 2012) to perform cross-validation on independent runs of data within subject. Second, we assessed here only the emotional responses to different vignettes. Future work will also examine how participants would choose to punish negative behavior and reward positive behavior differently in individuals

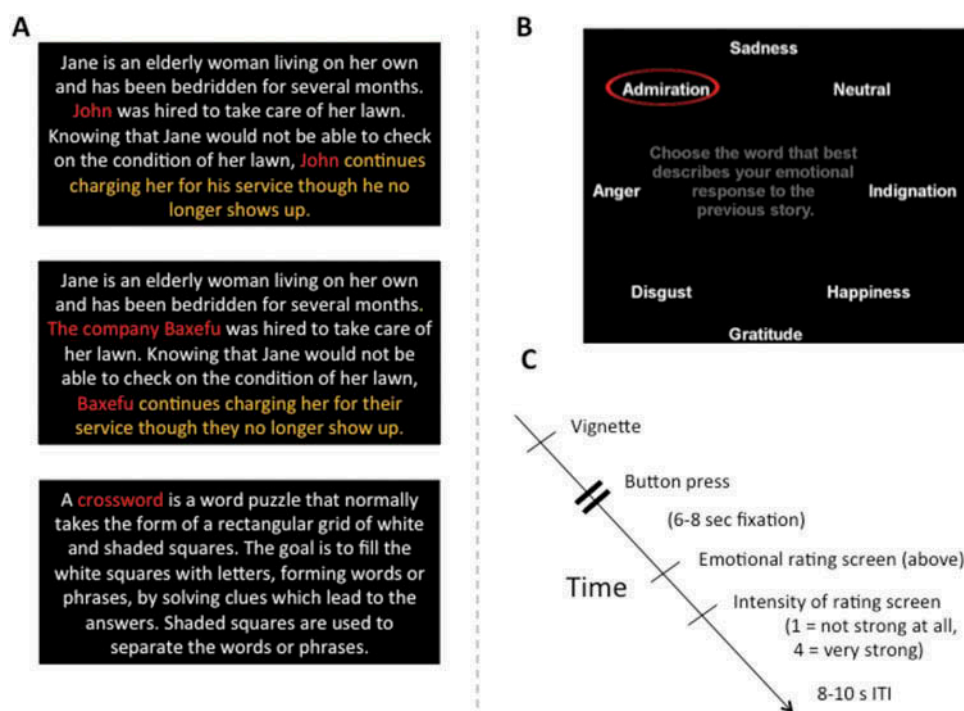


Figure 5. Experimental paradigm. **A.** A representative pair of matched “person” (top) and “company” (middle) vignettes are shown with the *agency* manipulation highlighted in red and the *valence* manipulation highlighted in yellow. Participants did not see any such highlighting, and each participant only saw one vignette from a matched pair. An “object” vignette (below) is also shown. **B.** An example emotion rating screen that participants saw between each trial. **C.** Schematic of a single trial.

and corporations. This study will allow us to differentiate the actions participants take from their emotional response.

It should be noted that company vignettes occasionally mentioned individual people (e.g., as the recipient of the action of a company, see Jane in Figure 5A). The strong activation of the ToM network observed in this experiment could be due to the mentioning of individuals in these vignettes. Future work should fully isolate corporations or the collective unit being studied from the individuals with whom they often interact.

As mentioned earlier, ToM activity has been elicited by other non-human entities such as “socially interacting” basic shapes (Martin & Weisberg, 2003). Therefore, we cannot conclude that corporations or companies represent a special case of ToM network activation. We did not test participants’ judgments or brain activity with any other type of collective, such as charities, educational institutions, or government organizations. It is likely that the same networks will be strongly activated when contemplating the actions of other such collective groups.

Lastly, this study does not provide a scientific justification or grounding for court decisions such as *Citizens United v. Federal Election Commission* (2010). There are numerous political, legal, and economic reasons why corporations ought or ought not to be granted the rights of individuals; these are beyond the scope of this study. Rather, this study concludes that judgments about both individuals and corporations are underpinned by remarkably similar neural mechanisms.

Conclusion

We investigated how participants judge the actions of corporations compared to the actions of individuals. Our results showed that participants elicited the same networks of brain activity in response to the actions of corporations and individuals alike. Both moral ratings and neural responses to these vignettes revealed a slight negativity bias toward corporations in which participants appeared to judge the actions of companies more harshly. Collectively, our results support our hypothesis that corporations are viewed as social beings.

EXPERIMENTAL PROCEDURES

Participants

Forty two adults (19 males) aged 26.86 ± 7.42 years (mean \pm SD) with 17.1 ± 1.92 years of education normal or corrected-to-normal vision and no history of current or past mental or neurological illness were recruited from the Houston, TX, metropolitan area to take part in the experiment after giving written consent in accordance with the Institutional Review Board at Baylor College of Medicine. Two participants were excluded from all analyses for failure to follow instructions during the in-scanner task. An additional 10 participants did not meet head movement standards described later and were removed from all fMRI analyses, leaving 30 participants (15 males, age 28.6 ± 8.18 years, 17.23 ± 2.14 years of education) that were included in the GLM and MVPA analyses. In order to maximize statistical power, all 40 participants (18 males, age 27.33 ± 7.45 years, 17.18 ± 1.91 years of education) who successfully completed the experiment were included in the behavioral analyses.

Paradigm

Passive and active viewing of short, emotionally laden vignettes has successfully been used to investigate moral emotional reactions to situations of moral violations, moral dilemmas, empathic responses, and decisions of culpability (Buckholtz et al., 2008; Casebeer & Churchland, 2003; Greene, 2001; Moll et al., 2002; Schaich Borg, Lieberman, & Kiehl, 2008; Young, Cushman, Hauser, & Saxe, 2007). Although other paradigms, such as the viewing of emotionally laden images and faces, have been used to investigate moral emotions, the vignette paradigm is the only one that lends itself to judgments about an abstract entity such as a company or organization (Decety & Lamm, 2006; Greene & Haidt, 2002; Moll et al., 2008; Thielscher & Pessoa, 2007).

Stimuli

We constructed 75 vignettes. 30 of these described the actions of people (“person” vignette; mean 45.63 words), of which 10 were about pro-social or positive actions (such as donating to charity), 10 were about anti-social or negative actions (such as lying or breaking the law), and 10 were of a neutral *valence* (such as painting a room or buying a printer). Each of these 30

Person vignettes was matched with a vignette about a company (“company” vignette; mean 44.73 words) with only minor details of the vignette changed for plausibility. The names of persons in the vignettes were randomly chosen on each trial from a list of 15 popular male names, and company names were randomly chosen on each trial from a list of 15 names generated from an online company name generator (www.company-name-generator.com). The additional 15 vignettes (“object” vignette; mean 47.6 words) were adapted descriptions from wikipedia.org of nouns generated from a random noun generator (www.desiquintans.com/noungenerator.php), see Figure 5A for example vignettes. Participants were asked to think of each of the vignettes as independent events. We refer to the company, person, and object vignettes as the *agency* axis, and the positive, negative, or neutral vignettes as the *valence* axis. The length of the vignettes did not significantly differ across the *agency* ($p = .92$) or *valence* axis ($p = .74$). See **Supplemental Experimental Procedures** for a list of all 75 vignettes.

Experimental procedure

While in the MRI scanner, participants read a total of 45 vignettes, 15 each of the company, person, and object conditions, during a single EPI run. The company and person vignettes were matched such that half of the participants were randomly presented with the company version of the vignette and the other half was presented with the person version. This was done to ensure that vignettes were matched without having any individual participant read the same story twice with different *agency*. The company and person vignettes included five positive *valence*, five negative *valence*, and five neutral *valence* vignettes per *agency*.

During each trial, the participants first saw a randomly drawn vignette from the available pool of 45. Once the participant finished reading the vignette, he/she pressed a button to move forward. This vignette screen was followed by a fixation cross for 6–8 seconds. After fixation, participants were prompted to choose the word that best described their emotional response to the previous vignette from a list of eight words (“happiness,” “disgust,” “anger,” “admiration,” “sadness,” “indignation,” “gratitude,” and “neutral”) placed at random, evenly spaced positions along a circle (Figure 5B). These emotions were selected from lists of basic and moral emotions in a previous study investigating disgust (Moll et al., 2005). The list of emotions in our current study does not reflect the

full extent of moral or basic emotions in Moll et al. (2005) but reflects the range of emotions expected from the stimuli in the current study. Subjects used two buttons to scroll through choices before confirming that choice with a third button press. Following this emotional rating, the participants rated the intensity of their emotional reaction on a scale of 1–4, with 4 being the most intense. On 25% of trials, participants additionally had to answer a true/false question about the content of the previous vignette to ensure sustained attending and understanding. Each trial was followed by a fixation period of 8–10 seconds. See Figure 5C for a schematic of an experimental trial.

fMRI data acquisition and pre-processing

High-resolution T1-weighted scans were acquired on a Siemens 3.0 Tesla Trio (Erlangen, Germany) scanner using an MPRage sequence. Functional run details were as follows: hyperscan echo-planar imaging, gradient recalled echo; repetition time (TR) = 2000 ms; echo time (TE) = 40 ms; flip angle = 90°; 64 × 64 matrix, 34 axial slices of 4 mm, yielding functional 3.4 mm × 3.4 mm × 4.0 mm voxels.

fMRI GLM data analysis

General Linear Model analysis was performed using SPM8 (www.fil.ion.ucl.ac.uk/spm/software/spm8) with motion artifact removal using the Art toolbox (www.nitrc.org/projects/artifact_detect). Images were created using Mango (<http://ric.uthscsa.edu/mango>).

Motion Correction was carried out by co-registering data to a mean functional volume. Images in which head motion exceeded a cutoff (>1 mm of translation or rotation between consecutive TRs) were regressed out of the model. Images that were outliers (>3 standard deviations from mean) in global brain activation were also regressed out. Additionally, any participant whose head movement exceeded 3 mm at any point in the scan was removed from analysis. Ten participants were completely removed from the GLM analysis for head movement exceeding 3 mm, leaving 30 total participants analyzed. Note that we did not exclude the 10 participants with excessive head motion from the behavioral analysis, as head motion only impacted image quality.

The average of the motion-corrected images was co-registered to each individual's structural MRI using a 12-parameter affine transformation. EPI images were spatially normalized to the MNI template

(2 mm × 2 mm × 2 mm voxels) by applying a 12-parameter affine transformation, followed by a non-linear warping using basis functions (Kao, Davis, & Gabrieli, 2005). Images were then smoothed using a 6-mm isotropic Gaussian kernel and highpass filtered in the temporal domain (filter width of 128 s).

To identify regions of interest of increased activation, we performed a general linear model (GLM) regression. Regressors were defined from the onset times and durations of all vignettes, emotional rating screens, and questions (separated by *valence* and *agency* condition). Additionally, the timing of participants' button presses and head movement parameters were included in the GLM as effects of no interest to account for motor responses and head movements. The events were convolved with SPM's canonical HRF (characterized by 2 gamma functions) to create the regressors used for analysis. After performing the regressions, we formed 12 contrasts of β values; FWE correction to a p -value of .05 was performed for all contrasts by non-parametric methods (10,000 permutations) using SnPM8 (<http://warwick.ac.uk/tenichols/snpm>).

Multi-voxel pattern analysis

In addition to the standard GLM univariate analysis, we employed MVPA using a linear SVM to search for patterns of activity that could successfully decode *agency* and *valency* separately. The MVPA could also provide insight on how brains might differentially process emotional content when considering the actions of persons as compared to companies.

MVPA preprocessing

The MVPA analysis was performed independently of the GLM. As such, we started with the raw data and built a separate preprocessing pipeline for MVPA analysis. The raw functional data for each of the 30 participants were preprocessed using AFNI (<http://afni.nimh.nih.gov/afni>) (Cox, 1996). The preprocessing pipeline consisted of linear slice timing correction, epi to anatomical alignment, anatomical to Talairach alignment via non-linear registration (3dQwarp), re-sampling to 3 × 3 × 3 mm, and volume registration to the first volume. We then performed additional preprocessing steps using the PyMVPA python package (<http://www.pympva.org>) (Hanke, Halchenko, Sederberg, Hanson, et al., 2009; Hanke, Halchenko, Sederberg, Olivetti, et al., 2009). We carried out fifth-order polynomial de-trending and

z-scoring (normalizing) on each voxel independently over time. Lastly, we used a Talairach probability atlas to select only grey matter voxels with probability greater than 20% for each aligned functional data-set.

MVPA trial and time selection

Because participants' functional data were structurally aligned to Talairach space, we combined all participants and all trials into one large data-set. This greatly increased the ability to train a classifier, as each participant individually only saw five trials of each *valence* (positive, negative, neutral) for each *agency* of interest (person and company). By combining all participants, our pattern analysis could be performed on an average of 139 trials for each *agency*-specific *valence* (see **Table S1**). Not all trials for each participant were kept. We used artifact detection to identify individual trials that entailed too much motion (>4 STD in fMRI signal intensity or >1 STD in motion) and rejected those specific trials from the MVPA analysis. We also ensured that each participant had at least three trials in each category after removal of artifact trials. As a result of this criterion, we rejected one participant who did not have the adequate number of trials, leaving a total of 29 participants on which we performed the MVPA analysis. Further, because participants had variable reading times, we aligned to the time the participant finished reading the vignette (indicated by button press) with four volumes (8 s) on either side.

MVPA pattern classification

We constructed several linear SVMs using PyMVPA (LinearCSVMC, implemented via LIBSVM) to classify *valency* (i.e., positive, negative, vs. neutral vignettes) and *agency* (person vs. company vignettes) separately. For all classifiers, we used a one-way ANOVA to select for 5000 task-related features to pass to the SVM. We used leave-one-participant-out cross-validation to train and test the SVM for all analyses. The strength of the margin for each fold was estimated from the normal of the selected features. For the first classifier, we used all nine volumes surrounding the button press to create a 2-class SVM to decode on a trial-by-trial basis the *agency* (i.e., person or company). Separately, we used a 3-class SVM to decode *valency* (i.e., positive, negative, or neutral) within each *agency*. To dissect the 3-class SVM further, we computed the confusion matrices of the *valency* classification. Finally, we conducted a temporal analysis and created a set of 3-class SVMs to

decode *valency* for each of the 9 time points surrounding the button press separately. This allowed us to determine *when* the decoding accuracy was strongest. We also plotted the sensitivities (SVM weight vectors) for each voxel and projected the sensitivities overlaid on the Talairach template. In addition to overall prediction accuracy, the sensitivity maps allowed us to infer which areas in the brain are implicated in classifying *valency*.

Supplementary material

Supplementary contents are available via the 'Supplementary' tab on the article's online page (<http://dx.doi.org/10.1080/17470919.2014.978026>).

Original manuscript received 20 May 2014
Revised manuscript accepted 13 October 2014
First published online 10 November 2014

REFERENCES

- Amodio, D. M., & Frith, C. D. (2006). Meeting of minds: The medial frontal cortex and social cognition. *Nature Reviews Neuroscience*, *7*, 268–277. doi:10.1038/nrn1884
- Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2003). fMRI responses to video and point-light displays of moving humans and manipulable objects. *Journal of Cognitive Neuroscience*, *15*, 991–1001. doi:10.1162/089892903770007380
- Buckholz, J. W., Asplund, C. L., Dux, P. E., Zald, D. H., Gore, J. C., Jones, O. D., & Marois, R. (2008). The neural correlates of third-party punishment. *Neuron*, *60*, 930–940. doi:10.1016/j.neuron.2008.10.016
- Buckner, R. L., Krienen, F. M., Castellanos, A., Diaz, J. C., & Yeo, B. T. T. (2011). The organization of the human cerebellum estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*, 2322–2345. doi:10.1152/jn.00339.2011
- Burwell v. Hobby Lobby*, 573 U.S. (2014).
- Casebeer, W., & Churchland, P. (2003). The neural mechanisms of moral cognition: A multiple-aspect approach to moral judgment and decision-making. *Biology and Philosophy*, *18*, 169–194. doi:10.1023/A:1023380907603
- Cikara, M., Eberhardt, J. L., & Fiske, S. T. (2011). From agents to objects: Sexist attitudes and neural responses to sexualized targets. *Journal of Cognitive Neuroscience*, *23*, 540–551. doi:10.1162/jocn.2010.21497
- Citizens United v. Federal Election Commission*, 558 U.S. 310 (2010).
- Costafreda, S. G., Fu, C. H. Y., Lee, L., Everitt, B., Brammer, M. J., & David, A. S. (2006). A systematic review and quantitative appraisal of fMRI studies of verbal fluency: Role of the left inferior frontal gyrus. *Human Brain Mapping*, *27*, 799–810. doi:10.1002/hbm.20221

- Coutanche, M. N., & Thompson-Schill, S. L. (2012). The advantage of brief fMRI acquisition runs for multi-voxel pattern detection across runs. *Neuroimage*, *61*, 1113–1119. doi:10.1016/j.neuroimage.2012.03.076
- Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*, 162–173. doi:10.1006/cbmr.1996.0014
- D'Esposito, M., Postle, B. R., & Rypma, B. (2000). Prefrontal cortical contributions to working memory: Evidence from event-related fMRI studies. *Experimental Brain Research*, *133*, 3–11. doi:10.1007/s002210000395
- Decety, J., & Lamm, C. (2006). Human empathy through the lens of social neuroscience. *The Scientific World Journal*, *6*, 1146–1163. doi:10.1100/tsw.2006.221
- Folkes, V., & Kamins, M. (1999). Effects of information about firms' ethical and unethical actions on consumers' attitudes. *Journal of Consumer Psychology*, *8*, 243–259. doi:10.1207/s15327663jcp0803_03
- Friederici, A. D., Ruschmeyer, S. A., Hahne, A., & Fiebach, C. J. (2003). The role of left inferior frontal and superior temporal cortex in sentence comprehension: Localizing syntactic and semantic processes. *Cerebral Cortex*, *13*, 170–177. doi:10.1093/cercor/13.2.170
- Greene, J., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences*, *6*, 517–523. doi:10.1016/S1364-6613(02)02011-9
- Greene, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, *293*, 2105–2108. doi:10.1126/science.1062872
- Grossman, E., Donnelly, M., Price, R., Pickens, D., Morgan, V., Neighbor, G., & Blake, R. (2000). Brain areas involved in perception of biological motion. *Journal of Cognitive Neuroscience*, *12*, 711–720. doi:10.1162/089892900562417
- Grossman, E. D., & Blake, R. (2001). Brain activity evoked by inverted and imagined biological motion. *Vision Research*, *41*, 1475–1482. doi:10.1016/S0042-6989(00)00317-5
- Hanke, M., Halchenko, Y. O., Sederberg, P. B., Hanson, S. J., Haxby, J. V., & Pollmann, S. (2009). PyMVPA: A python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics*, *7*, 37–53. doi:10.1007/s12021-008-9041-y
- Hanke, M., Halchenko, Y. O., Sederberg, P. B., Olivetti, E., Frund, I., Rieger, J. W., ... Pollmann, S. (2009). PyMVPA: A Unifying Approach to the Analysis of Neuroscientific Data. *Front Neuroinform*, *3*, 3.
- Harris, L. T., & Fiske, S. T. (2006). Dehumanizing the lowest of the low: Neuroimaging responses to extreme out-groups. *Psychological Science*, *17*, 847–853. doi:10.1111/j.1467-9280.2006.01793.x
- Harvey, A. H., Kirk, U., Denfield, G. H., & Montague, P. R. (2010). Monetary favors and their influence on neural responses and revealed preference. *Journal of Neuroscience*, *30*, 9597–9602. doi:10.1523/JNEUROSCI.1086-10.2010
- Hein, G., & Singer, T. (2008). I feel how you feel but not always: The empathic brain and its modulation. *Current Opinion in Neurobiology*, *18*, 153–158. doi:10.1016/j.conb.2008.07.012
- Hirshorn, E. A., & Thompson-Schill, S. L. (2006). Role of the left inferior frontal gyrus in covert word retrieval: Neural correlates of switching during verbal fluency. *Neuropsychologia*, *44*, 2547–2557. doi:10.1016/j.neuropsychologia.2006.03.035
- Jack, A. I., Dawson, A. J., Begany, K. L., Leckie, R. L., Barry, K. P., Ciccio, A. H., & Snyder, A. Z. (2012). fMRI reveals reciprocal inhibition between social and physical cognitive domains. *Neuroimage*, *66C*, 385–401.
- Kang, M. S. (2010). After Citizens United. *Indiana Law Review*, *44*, 243–254.
- Kao, Y.-C., Davis, E. S., & Gabrieli, J. D. E. (2005). Neural correlates of actual and predicted memory formation. *Nature Neuroscience*, *8*, 1776–1783. doi:10.1038/nn1595
- Knabb, J. J., Welsh, R. K., Erickson, S. K., & Felthous, A. R. (2009). Neuroscience, moral reasoning, and the law. *Behavioral Sciences & the Law*, *27*, 219–236. doi:10.1002/bsl.854
- Krauss, R. (2010). Neuroscience and Institutional Choice in Federal Sentencing Law. *Yale Law Journal*, *120*, 367–378.
- Mar, R. A. (2011). The neural bases of social cognition and story comprehension. *Annual Review of Psychology*, *62*, 103–134. doi:10.1146/annurev-psych-120709-145406
- Martin, A., & Weisberg, J. (2003). Neural foundations for understanding social and mechanical concepts. *Cognitive Neuropsychology*, *20*, 575–587. doi:10.1080/02643290342000005
- Mitchell, J. P., Banaji, M. R., & Macrae, C. N. (2005). The link between social cognition and self-referential thought in the medial prefrontal cortex. *Journal of Cognitive Neuroscience*, *17*, 1306–1315. doi:10.1162/0898929055002418
- Moll, J., De Oliveira-Souza, R., Eslinger, P. J., Bramati, I. E., Mourao-Miranda, J., Andreiuolo, P. A., & Pessoa, L. (2002). The neural correlates of moral sensitivity: A functional magnetic resonance imaging investigation of basic and moral emotions. *The Journal of Neuroscience*, *22*, 2730–2736.
- Moll, J., De Oliveira-Souza, R., Moll, F. T., Ignacio, F. A., Bramati, I. E., Caparelli-Daquer, E. M., & Eslinger, P. J. (2005). The moral affiliations of disgust: A functional MRI study. *Cognitive and Behavioral Neurology*, *18*, 68–78. doi:10.1097/01.wnn.0000152236.46475.a7
- Moll, J., De Oliveira-Souza, R., Passman, L. J., Cunha, F. C., Souza-Lima, F., & Andreiuolo, P. A. (2000). Functional MRI correlates of real and imagined tool-use pantomimes. *Neurology*, *54*, 1331–1336. doi:10.1212/WNL.54.6.1331
- Moll, J., de Oliveira-Souza, R., & Zahn, R. (2008). The neural basis of moral cognition: Sentiments, concepts, and values. *Annals of the New York Academy of Sciences*, *1124*, 161–180. doi:10.1196/annals.1440.005
- Moss, H. E., Abdallah, S., Fletcher, P., Bright, P., Pilgrim, L., Acres, K., & Tyler, L. K. (2005). Selecting among competing alternatives: Selection and retrieval in the left inferior frontal gyrus. *Cerebral Cortex*, *15*, 1723–1735. doi:10.1093/cercor/bhi049
- Pasalar, S., Ro, T., & Beauchamp, M. S. (2010). TMS of posterior parietal cortex disrupts visual tactile multi-sensory integration. *The European Journal of*

- Neuroscience*, 31, 1783–1790. doi:10.1111/j.1460-9568.2010.07193.x
- Pessoa, L., Gutierrez, E., Bandettini, P. A., & Ungerleider, L. G. (2002). Neural correlates of visual working memory: FMRI amplitude predicts task performance. *Neuron*, 35, 975–987. doi:10.1016/S0896-6273(02)00817-6
- Poldrack, R. A., Wagner, A. D., Prull, M. W., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. E. (1999). Functional specialization for semantic and phonological processing in the left inferior prefrontal cortex. *Neuroimage*, 10, 15–35. doi:10.1006/nimg.1999.0441
- Rilling, J. K., Gutman, D. A., Zeh, T. R., Pagnoni, G., Berns, G. S., & Kilts, C. D. (2002). A neural basis for social cooperation. *Neuron*, 35, 395–405. doi:10.1016/S0896-6273(02)00755-9
- Robertson, D., Snarey, J., Ousley, O., Harenski, K., Bowman, F. D., Gilkey, R., & Kilts, C. (2007). The neural processing of moral sensitivity to issues of justice and care. *Neuropsychologia*, 45, 755–766. doi:10.1016/j.neuropsychologia.2006.08.014
- Schaich Borg, J., Lieberman, D., & Kiehl, K. A. (2008). Infection, incest, and iniquity: Investigating the neural correlates of disgust and morality. *Journal of Cognitive Neuroscience*, 20, 1529–1546. doi:10.1162/jocn.2008.20109
- Thielscher, A., & Pessoa, L. (2007). Neural correlates of perceptual choice and decision making during fear-disgust discrimination. *Journal of Neuroscience*, 27, 2908–2917. doi:10.1523/JNEUROSCI.3024-06.2007
- van Boven, R. W., Ingeholm, J. E., Beauchamp, M. S., Bickle, P. C., & Ungerleider, L. G. (2005). Tactile form and location processing in the human brain. *Proceedings of the National Academy of Sciences*, 102, 12601–12605. doi:10.1073/pnas.0505907102
- Young, 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, 8235–8240. doi:10.1073/pnas.0701408104