The Price of Your Soul: Neural Evidence for the Non-Utilitarian Representation of Sacred Values

Gregory S. Berns[1*], Emily Bell[1], C. Monica Capra[1], Michael J. Prietula[2], Sara Moore[1], Brittany Anderson[1], Jeremy Ginges[3], Scott Atran[4]

[1]Department of Economics, Emory University, 1602 Fishburne Dr, Atlanta, GA 30322
[2]Goizueta Business School, Emory University, Atlanta, GA 30322
[3]Department of Psychology, New School for Social Research, 80 Fifth Ave, New York, NY 10011
[4]Institut Jean Nicod, CNRS-Ecole Normale Supérieure, 29, Rue d'Ulum, 75005 Paris, France

ABSTRACT

Sacred values, such as those associated with religious or ethnic identity, underlie many important individual and group decisions in life, and individuals typically resist attempts to trade-off their sacred values in exchange for material benefits. Deontological theory suggests that sacred values are processed based on rights and wrongs irrespective of outcomes, while utilitarian theory suggests they are processed based on costs and benefits of potential outcomes, but which mode of processing an individual naturally uses is unknown. The study of decisions over sacred values is difficult because outcomes cannot typically be realized in a laboratory, and hence little is known about the neural representation and processing of sacred values. We utilized an experimental paradigm that used integrity as a proxy for sacredness and which paid real money to induce individuals to sell their personal values. Using functional magnetic resonance imaging (fMRI), we found that values that people refused to sell (sacred values) were associated with increased activity in the left temporoparietal junction and ventrolateral prefrontal cortex, regions previously associated with semantic rule retrieval. This suggests that sacred values affect behavior through the retrieval and processing of deontic rules and not through a utilitarian evaluation of costs and benefits.

INTRODUCTION

Sacred values include fundamental religious beliefs, core constructs of national and ethnic identities, and moral norms. These values motivate many important individual and group decisions in life. Decisions bounded by them range from purchasing consumer goods such as kosher foods, patronizing Christian businesses, investing in socially responsible mutual funds, to deciding whom to marry. Disagreements over sacred values also contribute to many political and military conflicts and may also underlie some acts of political violence (1, 2). Thus, understanding how sacred values are represented and processed in the human mind has far reaching implications for policymakers.

By definition, personal sacred values are values for which individuals resist trade-offs with other values, particularly economic or materialistic incentives (3). The nature of sacred values is, in large part, defined by the way in which individuals engage them in decisions, but virtue theory suggests two very different ways in which sacred values might be processed (4). Sacred values could be either deontological in nature (5) or they could be utilitarian (6, 7). Deontic processing is defined by an emphasis on rights and wrongs, whereas utilitarian processing is characterized by costs and benefits. Similarly, deontic processing tends to be absolute and independent of outcomes, while utilitarian processing depends on the relative valuation of outcomes. Utility theory has emerged as a normative framework for the latter (8, 9), and when applied to decisions over sacred values, suggests that the expectation of consequences for violating these values is a deterrent to certain behaviors (10). Lexicographic preferences, in which an agent infinitely prefers one thing to another, has also been used to model sacred values within the utilitarian framework (11). In contrast, the deontic approach suggests that sacred values derive from rules that circumscribe certain actions independently of expected outcomes or prospects of success, and that we act in accordance with them because they are the right thing to do (3, 10).

Until recently, researchers had to rely solely on the decision-makers' self-reports via ratings, rankings, or rationale, which are often taken as evidence in favor of one theory or the other. Although insightful, these reports may be influenced by the context of being studied and what are perceived as socially acceptable reasons for doing things. Because it is very difficult to provide realistic outcomes for sacred decisions in an experiment, it is hard to measure behavioral parameters that would allow one to infer the structure of the sacred value decision space. Functional magnetic resonance imaging (fMRI) has emerged as a viable tool to measure brain regions associated with different aspects of decision making, and the growing literature on the neural correlates of moral judgment has demonstrated that deontic and utilitarian processing are associated with different brain regions (12-17).

While the previous literature has elucidated which brain regions become active when individuals engage either deontic or utilitarian reasoning, the question remains as to how individuals naturally represent sacred values when not forced into a particular framework of decision making. For example, one might consider the permissibility of killing a human being (the sanctity of human life being a common sacred value) in terms of rights and wrongs or in terms of consequences (e.g. legal sanctions or the effect on the victim's family), and either mode can be imposed by a particular experimental situation. Here, we use fMRI to determine whether sacred values are naturally represented as deontological rules or utilitarian values, unconstrained by a choice.

Clearly, this project presents some methodological challenges. It is not possible, nor desirable to request participants in an experiment to make actual decisions that would violate their values. However, we can use an element of sacredness that captures many of the key characteristics in the laboratory: integrity. Here, integrity refers to an individual's consistency of values and actions. For example, although we cannot (and do not wish to) test whether an individual is willing to kill an innocent human being, we can test their willingness to sign a document that says they would. Although signing such a document does not bind the person to that action, it creates an inconsistency between value and action that signals a loss of integrity. It is reasonable to assume that if something is truly sacred, then an individual would maintain their integrity for that value and not sign such a document. What if they were offered money to sign? It then becomes a trade-off between the monetary gain and the cost in personal integrity. In such a scenario, the amount of money required is one measure of integrity for a particular value, as is their willingness to set a price in the first place.

If sacred values are represented in a utilitarian manner, then prior neuroeconomic research suggests that they should be associated with increased neural activity in brain regions associated with the calculation of utility. Because sacred values are preferred above all else, these values should elicit the highest activity in regions processing utility. The most likely regions include ventromedial prefrontal cortex (VMPFC) (18, 19), striatum/nucleus accumbens (20-22), and parietal cortex (23, 24). Alternatively, if sacred values are represented as deontic rules, then brain regions associated with the processing of moral permissibility (rights and wrongs) should show increased activity, with the temporoparietal junction (TPJ) and the MPFC being the most commonly implicated regions (17, 25). Additionally, because rights and wrongs tend to be absolute and stored as rules, deontic processing should also be associated with the retrieval and processing of semantic knowledge. The literature on brain regions involved in semantic retrieval is extensive; however, common themes have emerged with respect to rule retrieval. Several studies have shown that the ventrolateral prefrontal cortex (VLPFC) supports storage or retrieval of semantic forms of stimulus-response contingencies, which are usually referred to as rules (26-30). Moreover, "the VLPFC may be more broadly involved in the retrieval and selection of representations that help to guide and constrain action through stored knowledge" (31). Thus, if sacred values are naturally processed in a utilitarian manner, one would expect to see activity in VMPFC, striatum, and parietal cortex; or, if they are processed deontologically, one would expect to see activity in the TPJ and VLPFC.

METHODS

The experimental task was designed to measure the neural responses to statements of personal values that ranged from the mundane to the sacred and test the utilitarian versus deontic representation of sacredness. As a proxy for sacredness, we measured integrity by an individual's willingness to accept real money to sign a document contradicting one's personal values (see ESM for details). The task was divided into four phases (1-3 in the scanner). To prevent strategic responding, the instructions for the fourth (auction) phase were given only immediately prior to that phase. First, in the Passive Phase, participants were presented with value statements phrased in the second person, one at a time. No decision was required. Statements ranged from values that were thought to represent mundane preferences – e.g. "You are a dog person" and "You are a Pepsi drinker" – to values describing attributions or acts that were thought to be sacred – e.g. "You believe in God" and "You are willing to kill an innocent human being." Every statement also had a complement – e.g. "You are a cat person," "You are a Coke drinker," "You do not believe in God," and "You are not willing to kill an innocent human being." A total of 62 pairs of statements (124 individual statements) was presented in random order. The purpose of this phase was to record a choiceless response to each statement and see whether utility or deontic systems dominate processing. Next, in the Active Phase, complementary statements were presented together, and for each pair, the participant had to choose one of the statements. In the third phase, the Hypothetical Phase, each statement chosen in the active phase was presented with a hypothetical offer of money to disavow the choice they had made in the active phase. For example, if someone previously chose "You believe in God," then the offer was, "Is there a dollar amount that you would accept to disavow your belief in God for the rest of your life?"

In the fourth phase, the Auction Phase, participants were given the opportunity to sell their answers from the active phase for real money. Using the Becker-DeGroot-Marshak (BDM) auction mechanism, participants were instructed to specify an "ask" price for each of the statements they chose in the active phase (32). The price could range from $1 to $100. The BDM auction is generally accepted be an incentive-compatible mechanism to reveal an individual's willingness to pay for something. Here, we use it as a willingness to accept. Submitting an ask price of $1, for example, means that the individual is willing to accept any amount of money and is assured of receiving some amount, which, on average, would be $50. In contrast, an ask price of $100 means that there would only be a 1% chance of receiving money. Importantly, they could also opt out of the auction for any or all items, choosing not to alter their answer. Each item was classified based on whether the participant submitted a price in the auction (Bid) or opted out of the auction (Optout). Submitting a price meant that the participant was willing to exchange this item for money. Opting out meant that the value was non-negotiable or that the amount we offered for the value was not high enough. This provided us with a classification for sacred and non-sacred values.

After all of the ask values were obtained, the participant rolled a pair of 10-sided dice for each of the items for which an ask price was put. If the dice roll was greater than their ask price, they received the value of the dice roll for that item. Their final payment was the average of all items sold and not sold. At the end of the auction phase, the participant received a printout of their chosen statements (active phase), which they did not sell in the auction, and the new statements, which were the complements of the statements sold in the auction. The printout had to be signed. Prior to the auction, participants knew they would have to sign the final document of their personal values. In this manner, the signing of the final document provided an additional incentive to reveal true value. To determine the temporal stability of their personal values, participants repeated the active phase choices through an online survey 6 to 14 months following their fMRI scan. At that point, they were also asked to indicate how they arrived at their decision: rights and wrongs, costs and benefits, or neither. The latter responses were subsequently used to create a functional localizer for deontic and utilitarian processing in the brain.

Neuroimaging data were collected using a 3 Tesla Siemens Magnetron Trio whole body scanner (Siemens Medical Systems, Erlangen, Germany). Functional data consisted of thirty-three axial slices that were sampled with a thickness of 3.5 mm and encompassing a field of view of 192 mm with an inplane resolution of 64 x 64 (T2* weighted, TR = 2000ms, TE = 30ms). Each participant completed 4 runs with 62 trials each, whose length depended on participants' decision time (2 runs of Passive, 1 run each of Active and Hypothetical). The Auction was done outside of the scanner.

FMRI data were analyzed using SPM5 (Wellcome Department of Imaging Neuroscience, University College London) using a standard 2-stage random-effects regression model. Data were subjected to standard preprocessing, including motion correction, slice timing correction, normalization to an MNI template brain and smoothing using an isotropic Gaussian kernel (full-width half-maximum = 8mm). Statistical thresholds were determined based on the estimated smoothness of the 2^{nd}-level contrasts. Using the AlphaSim routine in AFNI, we estimated the combination of height and extent thresholds that yielded a whole-brain FDR < 0.05. Using a voxel level threshold of $P<0.005$, the extent threshold that yielded a cluster level alpha of 0.05 was determined to be k≥53. A 40% gray matter mask was applied to all contrasts before using these thresholds.

First, a functional localizer for deontic vs. utilitarian processing was obtained from the Active phase. Using the responses from the follow-up survey which characterized the mode of processing for each trial, a first level model consisting of three conditions was created: 1) right/wrong; 2) cost/benefit; 3) neither. Regions involved in deontic processing were identified by the contrast of [right/wrong – cost/benefit], and utilitarian processing regions were identified by [cost/benefit – right/wrong]. Second, the ROI's obtained from the functional localizer were then used to mask the contrasts in the Passive phase. For the first-level model of the Passive phase, each statement was categorized based on the participant's subsequent choice in the Active phase (Chosen vs. Not Chosen) and whether they submitted an ask value during the auction (Bid vs. Optout). This created four conditions: 1) Chosen/Optout; 2) Chosen/Bid; 3) Not Chosen/Optout; 4) Not Chosen/Bid. Clearly, items that were Not Chosen could not be sold, but because the auction was to switch from the Chosen to the Not Chosen item, they were implicitly part of the choice process. One participant was excluded from the analysis because the participant submitted bids of $1 for all items, and thus no contrasts could be formed. The contrast from the first-level main effect of Optout-Bid was input into a second-level model and then masked with the [right/wrong – cost/benefit] and [cost/benefit – right/wrong] maps from the functional localizers. Within each ROI, the average activation across subjects was tested for significance with a t-test. Third, to identify additional regions that might contribute to sacred values, we performed a whole-brain analysis on the Passive phase of Optout vs. Bid. Mean differential activation between Optout-Bid in ROIs identified from this contrast were also correlated with subject-specific attributes, including subject age, religiosity, liberalism/conservatism, and activism. Finally, to evaluate the possibility of alternative interpretations of the activation patterns, four additional models were tested. Each of these models included a specific aspect of the stimulus as either a condition or covariate in the first-level model, and thus controlled for it as a "nuisance" variable: legal doctrine, syntax of statement, statement length, and semantic richness.

RESULTS

Participants (N=43) exhibited a wide range of choices with some participants auctioning nearly all of their answers while others very few (mean=58.6%, range: 8.1-100.0%). The aggregate distribution of ask values was bimodal, with most being either $1 or Optout and a declining tail between the two extremes (see ESM). There was an approximately ordinal and concave relationship between the fraction of individuals submitting bids to change an answer for a particular item and the fraction of individuals who stated they would hypothetically accept money to change their behavior (Figure 1). Follow-up 6-14 months after the initial experiment showed a high degree of stability of sacred values (96.4% consistent), and that sacred values were more stable than non-sacred values [paired $t(27)=7.81$, $P<0.001$]. In addition, 73.2% of the time participants selected "right/wrong" as the rationale for choosing a sacred value. In comparison, only 27.8% of the time was the “right/wrong” rationale used to explain choice for items that were not sacred (i.e., those for which a bid was submitted in the auction).

To determine whether a stimulus naturally prompted deontological or utilitarian processing, we examined the brain activity during the passive phase. Being the first and unconstrained phase of the experiment, and the only one in which items were presented individually, yielded a window into the natural processing of these values. Before analyzing this phase, a functional localizer for deontic vs. utilitarian processing was obtained from the Active phase. Using the responses from the follow-up survey which characterized the mode of processing for each trial, regions involved in deontic processing were identified by the contrast of [right/wrong – cost/benefit], and utilitarian processing regions were identified by [cost/benefit – right/wrong]. Only the left TPJ was identified as a deontic region, but the OFC, and left and right inferior parietal lobules, were identified as utilitarian (Figure 2). Second, the regions obtained from the functional localizer were then used to mask the contrasts in the Passive phase. For the first-level model of the Passive phase, each statement was categorized based on the participant’s subsequent choice in the Active phase (Chosen vs. Not Chosen) and whether they submitted an ask value during the auction (Bid vs. Optout). This created four conditions: 1) Chosen/Optout; 2) Chosen/Bid; 3) Not Chosen/Optout; 4) Not Chosen/Bid. The main effect of Optout-Bid was input into a second-level model and then masked with the deontic and utilitarian maps from the functional localizers. The left TPJ was significantly more active to Optout items compared to Bid items [$t(30)=3.19$, $P=0.0034$]. Both the left and right parietal ROIs were more active for the Bid items compared to the Optout items [$t(30)=2.59$ and 2.04, $P=0.015$ and 0.05, respectively], but the OFC was not [$t(30)=0.13$, $P=0.42$]. To determine if there were additional regions associated with processing sacred values not identified by the functional localizer, we also performed a whole-brain analysis of the Passive phase contrast Optout-Bid, which additionally included the left VLPFC, dorsomedial PFC, and right amygdala (Figure 3).

DISCUSSION

These results provide strong evidence that when individuals naturally process statements about sacred values, they use neural systems associated with evaluating rights and wrongs (TPJ) and semantic rule retrieval (VLPFC) but not systems associated with utility. The involvement of the TPJ is consistent with the conjecture that moral sentiments exist as context-independent knowledge in temporal cortex (14, 33). Both the left and right TPJ have been associated with belief attribution during moral judgments of third parties (17). Our results show that it is also involved in the evaluation of personal sacred values without decision constraints. Thus, one explanation for the reduction in morally prohibited judgments when the

TPJ is disrupted by transcranial magnetic stimulation (25) is because disruption impairs access to personal deontic knowledge.

The involvement of the left VLPFC in personal sacred values is also consistent with the conjecture that deontic rules are retrieved and processed as semantic knowledge as opposed to utility calculations. Although the VLPFC has been historically implicated in language function, more recent neuroimaging work has demonstrated that the particular area we identified as processing sacred values is associated with semantic rule retrieval and processing (27-29). Importantly, the function of the left VLPFC is not restricted to verbal or written rules. In a study of road signs, the anterior division of the left VLPFC was found to be most closely associated with rule retrieval (28). We observed the same division more active when participants processed sacred values. Similar results were also found in a language-based study of rules, with this region being implicated in top-down retrieval of semantic knowledge (27).

Although activation of the left TPJ and VLPFC for sacred values is consistent with a deontic rule retrieval process, it could be explained by properties of the stimuli, as opposed to how participants processed the stimuli. One possibility is that the items deemed sacred (Optout), are governed by legal doctrine and thus the left VLPFC activity simply reflected the retrieval and processing of a legal rule (e.g. it is illegal to kill people). To test this, we created another model in which we added a separate category for items governed by a legal doctrine. An item was coded as a legal doctrine if it or its complement were illegal, either by U.S. or international law (Geneva Conventions). The original analysis was then repeated on the remaining items. Using the original ROI's, we found that the effect of Optout-Bid was still significant in the left VLPFC [t(30)=2.15, P=0.040]. Thus, even for items not governed by any rule of law (e.g. believing in God), if the individual did not sell it, it was retrieved and processed as a rule. We also tested the possibility that sacred values involve concepts, like God, which have more meanings than mundane concepts like dogs and cats. Semantic richness (SR) refers to the amount of semantic information contained in, or associated with, a concept in semantic memory (34) and has been previously associated with activation in VLPFC (35-37). To test the possibility that SR may be partially confounded with our measures of sacredness, we formulated an alternative model that controlled for the SR of the statements. There was no significant correlation between the SR of the stimulus and the fraction of individuals submitting bids to change their answers (R^2=0.045, P=0.053), and when SR was included as a control variable in the Passive phase model, significance remained and changed only slightly for the ROI's. This suggests that sacredness was not confounded with SR. We also tested alternative models that controlled for the length and syntax of each statement, none of which greatly changed the significance of the activations in the ROIs.

Only the left and right inferior parietal lobules showed the opposite activation pattern, with greater activity to the Bid versus Optout items, which also coincided with cost/benefit decisions (Figure 2). This suggests that these regions activate for items that have a measurable utility or value. This is consistent with prior evidence implicating the parietal cortex in utility-based decisions (23, 24). The other region most likely to encode utilitarian values is the VMPFC and striatum (18-22), but we did not observe a significantly greater activation to Bid versus Optout items in these regions during the Passive phase.

The auction mechanism was a unique aspect of our experiment and suggests a new way to quantify sacred values that is not solely dependent on self-report, but there are assumptions behind its use. First, we assume that individuals take the auction seriously. As noted above, signing a document does not bind one to the action that one is signing. It is therefore somewhat surprising that most people didn't sell all of their choices. The fact that participants took money for some items and not others suggests that they were adequately motivated to express their preferences through their choices. The upper limit of $100, however, placed a boundary constraint on the auction, which when averaged over all items, yielded a low value per item in expectation (the framing of the auction explicitly instructed participants to value each item in the $1-$100 range, and all participants' questions about the auction pertained to how to earn the most money). Although $100 may have been insufficient to buy some answers, this could be true for any amount of money offered. The distribution of bids, however, suggests that this was not the case (see ESM). Although the distribution was dominated bimodally by $1 and Optout, the ask values showed a declining frequency toward the $100 boundary which could be fit by a gamma distribution. This indicates a decreasing marginal exchange value, and a higher upper limit would not have made a significant difference in items that were not auctioned.

Our experiment dovetails with a large literature on the neural correlates of moral judgment (12-17). However, it differs in that it initially measured the natural mode of processing sacred values in a way that was relatively unconstrained by a choice framework. This is particularly important for the scientific study of sacred values, because one cannot ethically place volunteers in real situations that would test such values. However, recent findings in neuroeconomics have demonstrated that "choiceless" brain responses are predictive of future actions (38, 39). Here, we find tantalizing evidence of this for sacred values too. We also found that the difference in VLPFC activation between Optout and Bid items correlated with the individual's level of involvement in organizations (Figure 4). This suggests that neural markers for sacredness extend to real-world decisions of group membership. Moreover, when sacred values were contradicted by their opposites, we observed a significant increase in amygdala activation, which suggests the presence of an arousal response and is consistent with the hypothesized role of emotion, especially negative ones, when sacred values are violated (1-3, 10).

Our results complement existing research in sacred values and may have implications for policymakers (1, 2), although further research in conditions that emulate policymaking environments will be required to make the case. Economic, foreign, and military policies are typically based on utilitarian considerations. More specifically, it is believed that those who challenge a functioning social contract should concede if an adequate trade-off is provided (e.g. sanctions or other incentives). However, when individuals hold some values to be sacred, they fail to make trade-offs, rendering positive or negative incentives ineffective at best. Our results suggest that individuals naturally retrieve sacred values as deontic rules, not as representations of utility, providing the first neurobiological evidence for what has been previously conjectured (3).

ACKNOWLEDGEMENTS

Supported by grants from the Air Force Office of Scientific Research (AFOSR) through the Office of Naval Research (ONR N000140910912) and NSF (0827313).

REFERENCES

1. Atran S, Axelrod R, Davis R. Sacred barriers to conflict resolution. Science. 2007;317:1039-40.
2. Ginges J, Atran S, Medin D, Shikaki K. Sacred bounds on rational resolution of violent political conflict. Proc Natl Acad Sci U S A. 2007;104(18):7357-60.
3. Baron J, Spranca M. Protected values. Organizational Behavior and Human Decision Processes. 1997;70(1-16).
4. Casebeer WD. Moral cognition and its neural constituents. Nat Rev Neurosci. 2003;4:841-6.
5. Kant I. Groundwork for the Metaphysics of Morals. Denis L, editor. Toronto: Broadview Press; 1785/2005.
6. Bentham J. The Principles of Morals and Legislation. Amherst: Prometheus Books; 1780/1988.
7. Mill JS. Utilitarianism. Fourth edition ed. London: Longmans, Green, Reader, and Dyer; 1871.
8. von Neumann J, Morgenstern O. Theory of Games and Economic Behavior. Princeton: Princeton University Press; 1944.
9. Nash JF. Equilibrium points in n-person games. Proc Natl Acad Sci USA. 1950;36:48-9.
10. Tetlock PE. Thinking the unthinkable: sacred values and taboo cognitions. Trends in Cognitive Sciences. 2003;7:320-4.
11. Spash CL, Hanley N. Preferences, information and biodiversity preservation. Ecological Economics. 1995;12:191-208.
12. Greene JD, Nystrom LE, Engell AD, Darley JM, Cohen JD. The neural basis of cognitive conflict and control in moral judgment. Neuron. 2004;44:389-400.

13. Greene J, Haidt J. How (and where) does moral judgment work? Trends in Cognitive Sciences. 2002;6(12):517-23.
14. Moll J, Zahn R, de Oliveira-Souza R, Krueger F, Grafman J. The neural basis of human moral cognition. Nat Rev Neurosci. 2005;6:799809.
15. Moll J, de Oliveira-Souza R. Moral judgments, emotions and the utilitarian brain. Trends Cogn Sci. 2007;11(8):319-21.
16. Koenigs M, Young L, Adolphs R, Tranel D, Cushman F, Hauser M, et al. damage to the prefrontal cortex increases utilitarian moral judgements. Naure. 2007;446:908-11.
17. Young L, Cushman F, Hauser M, Saxe R. The neural basis of the interaction between theory of mind and moral judgment. Proc Natl Acad Sci U S A. 2007;104(20):8235-40.
18. Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. Nature. 2006;441:223-6.
19. Plassmann H, O'Doherty J, Rangel A. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J Neurosci. 2007 Sep 12;27(37):9984-8.
20. Schultz W, Apicella P, Scarnati E, Ljungberg T. Neuronal activity in monkey ventral striatum related to the expectation of reward. J Neurosci. 1992;12(12):4595-610.
21. Knutson B, Rick S, Wimmer GE, Prelec D, Loewenstein G. Neural predictors of purchases. Neuron. 2007;53:147-56.
22. Montague PR, Berns GS. Neural economics and the biological substrates of valuation. Neuron. 2002;36:265-84.
23. Kable JW, Glimcher PW. The neurobiology of decision: consensus and controversy. Neuron. 2009;63:733-45.
24. Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. Nature. 1999;400:233-8.
25. Young L, Camprodon JA, Hauser M, Pascual-Leone A, Saxe R. Disruption of the right temporoparietal junction with transcranial magnetic stimulation reduces the role of beliefs in moral judgments. Proc Natl Acad Sci U S A. 2010;107(5):6753-8.
26. Bunge SA. How we use rules to select actions: a review of evidence from cognitive neuroscience. Cognitive, Affective, & Behavioral Neuroscience. 2004;4(4):564-79.
27. Badre D, Poldrack RA, Pare-Blagoev EJ, Insler RZ, Wagner AD. Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. Neuron. 2005;47:907-18.
28. Souza MJ, Donohue SE, Bunge SA. Controlled retrieval and selection of action-relevant knowledge mediated by partially overlapping regions in left ventrolateral prefrontal cortex. NeuroImage. 2009;46:299-307.
29. Bhanji JP, Beer JS, Bunge SA. Taking a gamble or playing by the rules: dissociable prefrontal systems implicated in probabilistic versus deterministic rule-based decisions. NeuroImage. 2010;49:1810-9.
30. Wallis JD, Anderson KC, Miller EK. Single neurons in prefrontal cortex encode abstract rules. Nature. 2001;411:953-6.
31. Badre D, Wagner AD. Left ventrolateral prefrontal cortex and the cognitive control of memory. Neuropsychologia. 2007;45:2883-901.
32. Becker GS, DeGroot M, Marschak J. Measuring utility by a single-response sequential method. Behavioral Science. 1964;9:226-32.
33. Lieberman MD. Social cognitive neuroscience: a review of core processes. Ann Rev Psychol. 2007;58:529-89.

34. Kounios J, Green DL, Payne L, Fleck JI, Grondin R, McRae K. Semantic richness and the activation of concepts in semantic memory: evidence from event-related potentials. Brain Res. 2009;1282:95-102.
35. Pexman PM, Hargreaves IS, Edwards JD, Henry LC, Goodyear BG. The neural consequences of semantic richness. Psychological Science. 2007;18(5):401-6.
36. Shivde G, Thompson-Schill SL. Dissociating semantic and phonological maintenance using fMRI. Cognitive, Affective, & Behavioral Neuroscience. 2004;4(1):10-9.
37. Thompson-Schill SL, D'Esposito M, Aguirre GK, Farah MJ. Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. Proc Natl Acad Sci U S A. 1997;94:14792-7.
38. Wunderlich K, Rangel A, O'Doherty JP. Economic choices can be made using only stimulus values. Proc Natl Acad Sci U S A. 2010;107:15005-10.
39. Berns GS, Capra CM, Chappelow J, Moore S, Noussair C. Nonlinear neurobiological probability weighting functions for aversive outcomes. NeuroImage. 2008;39:2047-57.

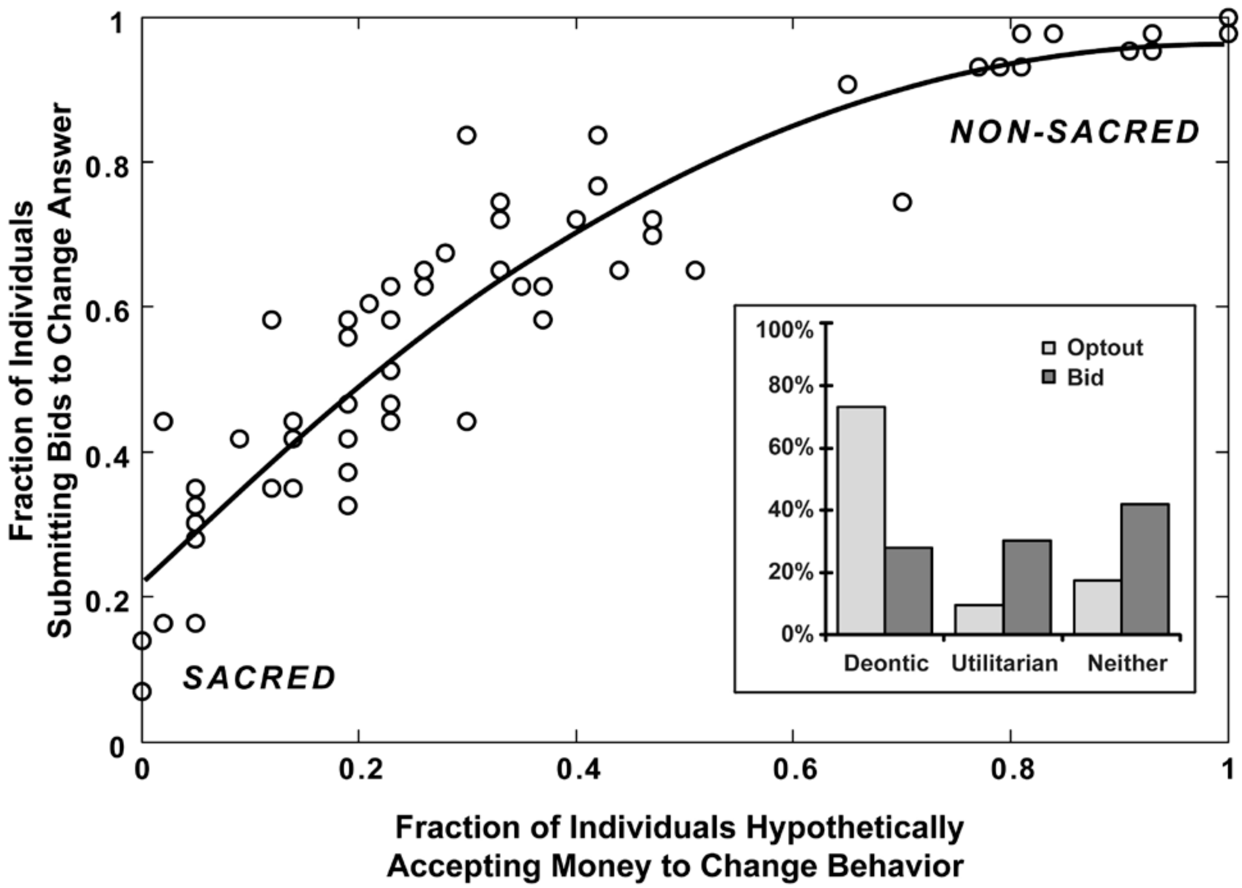

Figure 1. Relationship of auction behavior to hypothetical solicitation of money. Averaged over participants (N=43), each point represents one pair of personal values. The fraction of individuals submitting auction bids to change their answer for a given pair is plotted against the fraction of individuals who indicated that hypothetically they would accept money to change the corresponding behavior. These two measurements are highly correlated (adjusted R^2=0.87 for quadratic). Items in the lower left (*sacred*) had a low percentage of individuals willing to accept either hypothetical or real money. These items pertained to the sanctity of human life, especially children. Items in the upper right (*non-sacred*) had a high percentage of individuals willing to accept both hypothetical and real money to change and represented utilitarian preferences (e.g. Coke vs. Pepsi, dog vs. cat). Despite this correlation, all of the points lie above the diagonal, indicating a hypothetical bias (participants sold more often than they said they would hypothetically). The mode of decision-making (*inset*) was significantly different for statements that were not auctioned (Optout) versus those that were (Bid) [F(2,132)=58.7, *P*<0.001].

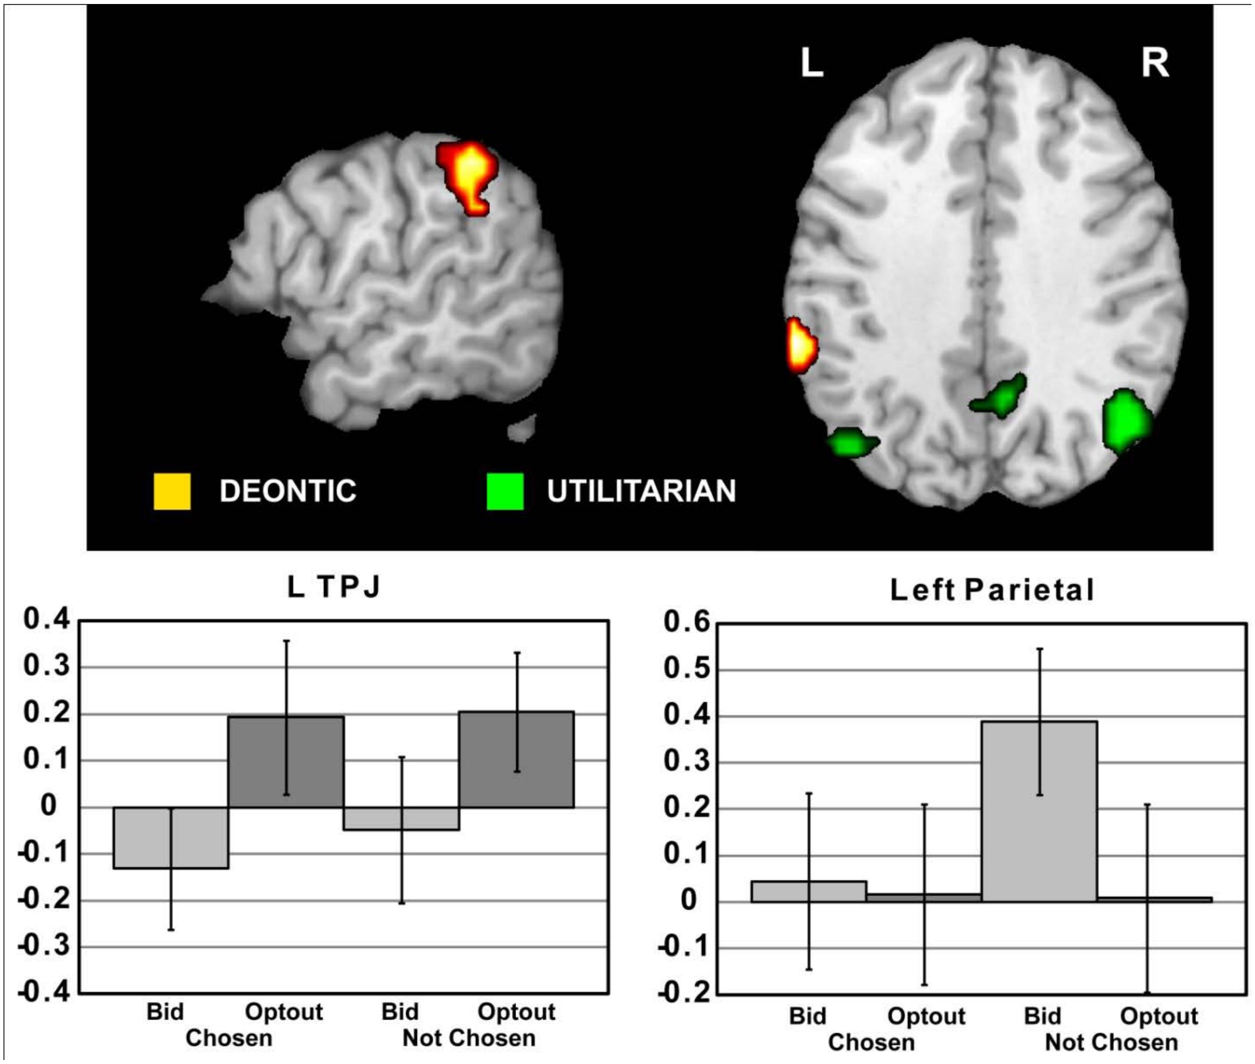

Figure 2. Functional localizer for brain regions with differential activity to deontic (*yellow*) and utilitarian (*green*) processing (P<0.005, extent≥53). Regions were classified as deontic when the participant indicated that their choice was based on rights and wrongs, and utilitarian when their choice was based on costs and benefits. These regions were then applied to the Passive phase activation in which each statement had been presented individually in the absence of choice. Each statement was categorized based on whether the participant sold a particular personal value during an auction held after the brain imaging session (*Bid*) or opted out of the auction for that value (*Optout*). At the time of imaging, participants did not know that they were going to have the opportunity to sell these values for real money. The left temporoparietal junction (MNI coordinates: -63, -39, 42) showed significantly greater activity for the Optout statements than the Bid statements (T=3.19, P=0.003), indicating that these were processed in the deontic region. Both the left and right inferior parietal lobule (MNI coordinates: -45, -72, 46 and 48, -66, 35) had the opposite pattern (T=4.09, P=0.001), which was driven primarily by the Not Chosen Bid statements, indicating that these statements were processed in utilitarian regions. Vertical scale on bar graphs are estimated beta values for the individual conditions ± s.e.m. across all subjects.

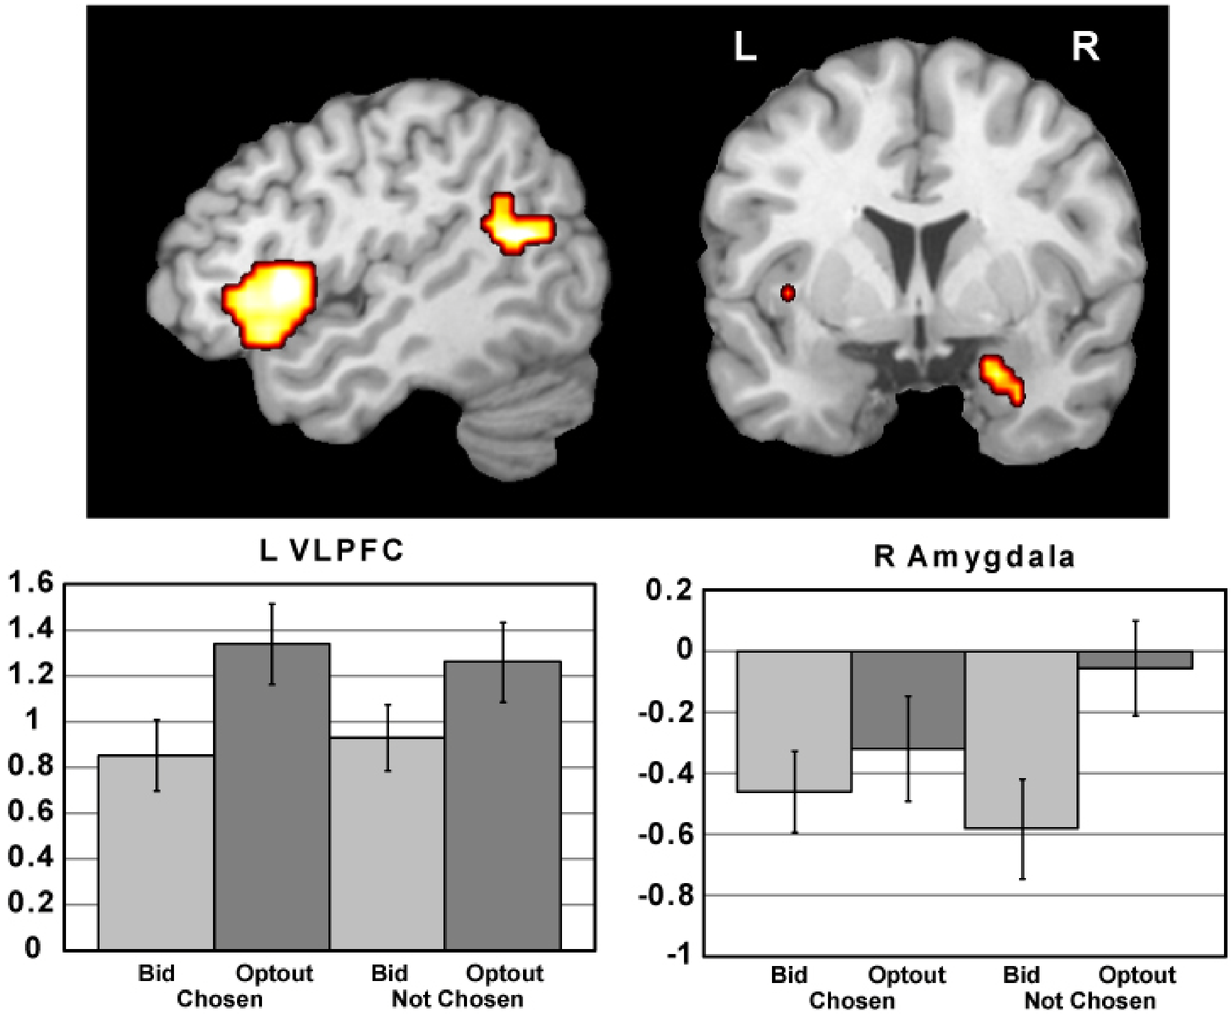

Figure 3. Additional regions identified in which sacred (Optout) items resulted in greater activation than non-sacred (Bid) items during the passive phase (P<0.005, extent≥53). These included the left ventrolateral prefrontal cortex (L VLPFC) and the right amygdala. Relative to the other three conditions, only the Not Chosen Optout items resulted in more amygdala activation. The latter statements represent the most repugnant items to the individual (those not chosen and not auctioned) and would be expected to provoke the most arousal, which is consistent with the idea that when sacred values are violated they induce outrage (1-3).

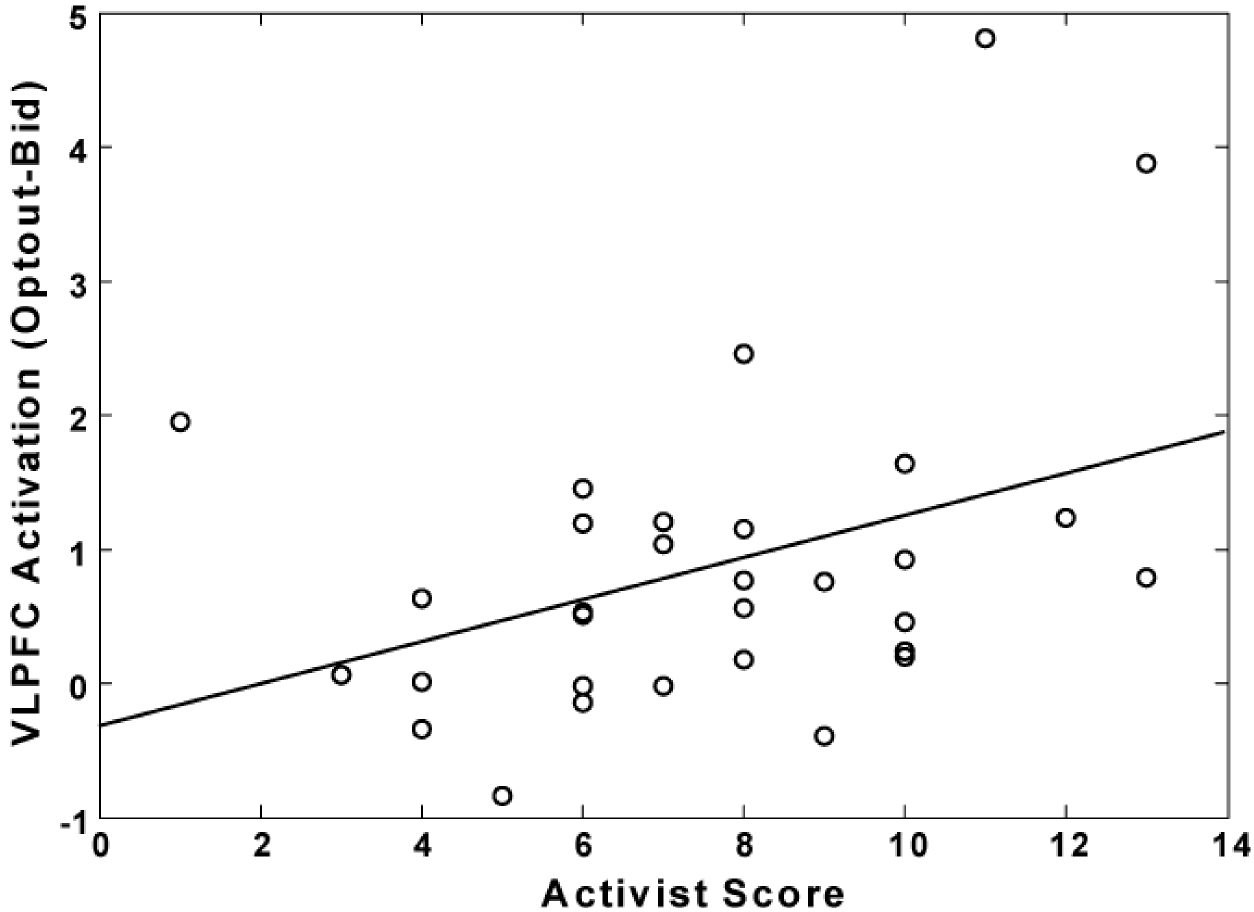

Figure 4. Difference in left ventrolateral prefrontal cortex (VLPFC) activation to sacred items (Optout) relative to non-sacred (Bid) items as function of each participants level of involvement in group activities (N=31). The activist score was calculated as the sum of ratings for membership in ten types organizations. Participants rated their involvement as 0 (don't belong), 1 (inactive member), or 2 (active member) for each organization: religious, sports/recreational, art/music/educational, labor union, political party, environmental, professional, humanitarian/charitable, consumer, other. There was a significant positive correlation between the overall level of organization involvement and the average difference in VLPFC activation to sacred and non-sacred items (R^2=0.39, *P*=0.032). This suggests that individuals who have stronger semantic representations of sacred values are more likely to act on their beliefs.