

**TITLE: ATTENDING TO MORAL VALUES**

Rumen Iliev, Sonya Sachdeva, Daniel M. Bartels, Craig Joseph,

Satoru Suzuki and Douglas Medin

**CONTENTS**

1. Introduction
2. Moral values in the laboratory
  - 2.1 Contingent valuation
  - 2.2 Tradeoffs
  - 2.3 Omission bias and quantity insensitivity
  - 2.4 Tragic vs. taboo tradeoffs and the flexibility of moral decision making
  - 2.5 Presentation order: Judgment in the context of previously-viewed options
  - 2.6 “Pseudo-sacred” values?
  - 2.7 Summary
3. A cognitive perspective on sacred values
  - 3.1 Protected values and attention
  - 3.2 Stroop effects
  - 3.3 Further cognitive consequences: Anchoring effects
  - 3.4 Summary
4. Attentional influences and the acceptability of tradeoffs
  - 4.1 Framing the question
  - 4.2 Framing the response alternatives
  - 4.3 Judgment in the context of one vs. two options
  - 4.4 The attraction effect: Choosing in the context of three options
  - 4.5 Summary
5. General discussion
  - 5.1 Reconciling the rigidity and flexibility of strong moral values

### **ABSTRACT**

There has been an upsurge of interest in *moral* decision making, which appears to have some distinctive properties. For example, some moral decisions are so strongly influenced by ideas about how sacred entities are to be treated, that they seem to be relatively insensitive to the costs and benefits entailed (e.g., “do not allow companies to pollute the earth for a fee, even if pollution credits reduce pollution”). One interpretation of such decisions is that sacred values motivate rigid decision processes that ignore outcomes. This, however, seems paradoxical in that those who are most offended by acts of pollution, for example, likely care more about pollution than others do. Our analysis of the literature on moral decision making (including our own studies) suggests a framework based on a "flexible view," where both actions and outcomes are important, and where attentional processes are intimately involved in how the decision maker conceptualizes the problem, how actions and outcomes are weighted, and how protected values are translated into judgments. We argue that understanding the cognitive processes underlying morally-motivated decision making offers one method for solving the puzzle of why such deeply entrenched commitments (the rigid view) vary widely in their expression across contexts (the flexible view).

## 1. INTRODUCTION

The distinction between two levels of existence for human beings - the sacred versus the profane - has long been a major theme in sociology and anthropology, but has only recently begun to receive attention in the area of decision making. The profane or secular subsumes most of ordinary life, where people have relative freedom to choose or decide what to do. The other level, sometimes overlapping with the first one, but qualitatively distinct, is characterized in terms of special meanings, norms, restrictions and obligations described as sacred or divine (Durkheim, 1912/1954; Eliade, 1959).

A long tradition in decision making, influenced largely by an economic perspective, has focused on the domain of secular existence, where people are considered to be autonomous agents, maximizing some benefit function without being affected by rules or mandates external to the market. As a consequence, phenomena such as ideologies, religious beliefs, virtues, moral values and ethical positions have received relatively little attention in the field. Recently, however, there has been a major upsurge in interest in moral judgment and morally motivated decision making (e.g., Greene & Haidt, 2002; Hauser, 2006; Nichols, 2004; Prinz, 2007).

In one important line of research, Baron, Tetlock and their associates have examined how protected or sacred values affect decision making (e.g., Baron & Green, 1996; Fiske & Tetlock, 1997; McGraw & Tetlock, 2005; Ritov & Baron, 1999; Tetlock, 2003; Tetlock, Kristel, Elson, Green, & Lerner, 2000; Baron & Ritov, 2004; present volume). Tetlock et al. (2000) define sacred values as “any value that a moral community implicitly or explicitly treats as possessing infinite and transcendental significance that precludes comparisons, tradeoffs, or indeed any other mingling with bounded or secular values.” When participants were presented with proposed tradeoffs of sacred values for money, Tetlock et al. (2000) observed strong cognitive, affective (e.g. moral outrage) and behavioral (moral cleansing) responses.

Similarly, Baron and Spranca (1997) define protected values as "... those that resist trade-offs with other values, particularly with economic values". They suggest that protected values are deontological rules and prescriptions concerning actions and inactions (e.g. "do not cause harm"), rather than consequentialist principles that focus on an assessment of probable costs and benefits. As Baron and Ritov (this volume) note, protected values are frequently linked to deontological rules that focus people's attention on acts at the expense of attending to the consequences of those acts, sometimes resulting in "omission bias."

The existence of protected values is of particular concern for utility-based theories, which broadly rely on the method of proposing tradeoffs between goods to assess their relative desirability. These theories assume that decision makers add up the costs and benefits and choose the option that provides the best overall value. However, moral prohibitions against trading of sacred values for secular values implies that these domains cannot be mixed and that no matter how much one increases the utility of the competing alternative, a person holding a protected value will never agree to the exchange (e.g. "you can't put a price on a human life"). In other words, it appears that a person with a sacred or protected value acts as if it has infinite utility.

Infinite utilities pose a conceptual challenge and it is not even clear that the idea is coherent. As Baron and Spranca (1997, see also Baron & Ritov, this volume) have pointed out, for example, if people have absolute values with infinite utility, they should spend all of their time, efforts and resources to maximize this value and neglect all else. This doesn't seem plausible. If a parent assigns infinite value to their child and devotes all their resources to protecting their child from harm, they nonetheless will need to pay some attention to their own health and welfare if only to be able to continue in that role. For example, they presumably would not break into a drug store that is closed and risk being jailed in order to obtain a thermometer to check to see if their child is running a fever.

If people have more than one sacred or protected value, they immediately face contradictions. In some situations they will be frozen in the position of Buridan's Donkey, where they will be unable to choose between two equally and infinitely good or equally and infinitely bad options. Such theoretical concerns make the domain of sacred or protected values an intriguing and challenging area for decision researchers.

Our focus in studying sacred values is on understanding the cognitive processes associated with decision making involving sacred or protected values. We do not doubt that sacred values may reflect inter-personal goals and that, in some contexts, they may be more effective than cost/benefit strategies (Frank, 1988). It seems equally clear that there are real world circumstances where sacred values involve much more than posturing, as in the case of suicide bombers, the self-immolation of monks or for soldiers who throw themselves on top of grenades to save their buddies (Atran, Axlerod, & Davis, 2007).

Although part of our research involves studying these contexts (Ginges, Atran, Medin, & Shikaki, 2007), the studies we review and present here undermine the simple assignment of sacred values to deontological rules and secular values to consequentialism. Instead, our data suggest that people with protected values care both about whether actions are good or bad in themselves (e.g., harmful actions are often judged impermissible) and about the consequences associated with those actions. This dual focus leads to a striking malleability in judgment and decision making and to both hypo-sensitivity and hyper-sensitivity to consequences. The key factors appear to be those that are associated with attentional processes.

It has long been recognized that human information processing resources are limited and that even the most basic decisions cannot take into account all the information that is potentially relevant (Simon, 1957). The same holds for perceptual processes; any realistic visual scene contains many objects, but the number of objects that one can simultaneously select, analyze and

keep track of, is generally limited to only a few (Franconeri, Alvarez, & Enns, 2007). There has been a great deal of research on attentional processes that determine which objects and/or features are selected for further conscious scrutiny. The overall results suggest that what sensory information reaches perceptual awareness is influenced by both bottom-up and top-down processes (e.g., see Egeth and Yantis, 1997, for a review, and Ling & Carrasco, 2006, for recent empirical results). Visual stimuli that are physically strong (e.g., a bright red flower blossoming out of pale leaves) and/or representationally strong (e.g., one's own name) tend to capture attention relatively automatically - bottom-up capture of attention. What sorts of stimuli capture attention may also be influenced by experience, learning, and the current cognitive goal (e.g., an odd-looking object might capture attention when one is looking for an oddball; a red circle might capture attention when one is looking for a red object). In addition to bottom-up attention capture, one can voluntarily select information pertaining to behaviorally relevant objects that do not attract attention (e.g., looking for a gray pen dropped in a garden of bright-colored flowers) - top-down control of attention. Bottom-up capture directs attention to stimuli that are strong and/or with special meaning; top-down control allows one to flexibly attend to a stimulus of choice.

In our review we will be concerned with both aspects of attention. That is, we examine both how strong moral values affect what attracts attention as well as factors related to how information is selected voluntarily. In some cases we use measures of attention that derive from the attention literature (e.g. the Stroop effect) but in other cases the information being selected is more abstract and conceptual than has been the focus in the field of attention. Even in the latter case, however, we think that attention nonetheless provides a useful framework. We argue that in morally relevant tasks attention acts as a selection mechanism to filter some of the information presented, enhancing some parts of it and suppressing others.

## **Overview**

The rest of this chapter is organized as follows. First, we provide a more detailed background of previous work motivating our present line of research. Then, we review a series of recent studies on cognitive processes associated with sacred or protected values, and after that extend this processing approach to other studies demonstrating context sensitivity as a function of attention allocation. Finally, we try to reconcile the rigidity of sacred or protected values with their apparent flexibility, arguing that cognitive processing is an important factor which should be taken into account.

## **2. MORAL VALUES IN THE LABORATORY**

### **2.1 Contingent valuation**

Researchers initially encountered protected values when they attempted to use contingent valuation as a method for determining the value of various goods (e.g. a nature preserve) for which there were no pre-existing markets (see Ritov & Kahneman, 1997 for a review). Policy makers, trying to find the best solutions of different tradeoff problems, would give questionnaires to the public and ask them to trade off one good for another, or to trade off a good for money. A typical question would ask a person the amount of tax cut that they would be willing to accept as compensation for the destruction of a nearby nature preserve or how much they would be willing to pay to protect a lake in Ontario, for example. These estimates could then be used for policy decisions involving these goods (e.g. how much Exxon should have to pay to compensate for destroying a beach and wildlife from an oil spill). The results, however, were too incoherent to be used for policy purposes. People appeared to be willing to pay no more to “save all the lakes in Ontario” than to save “one lake in Ontario” (Ritov & Kahneman, 1997) and they said that “no amount” of money would be enough to compensate for the loss of certain environmental goods (see also McGraw & Tetlock, 2005). For the latter scenarios, involving willingness to accept compensation, people interpreted the money as a bribe and rejected it on

moral grounds. These kinds of responses made many contingent valuation studies useless for the purposes of comparisons to other goods, determining compensation, or for policy setting.

## **2.2 Tradeoffs**

When asked about protected values in the lab, people also appear to reject some types of tradeoffs. A typical measure used by Baron and colleagues is to ask participants for their opinion about a particular tradeoff involving harm (e.g., cutting acres of old-growth forest) and to give them three options:

- A. I do not oppose this.
- B. This is acceptable if it leads to some sort of benefits that are great enough.
- C. This is not acceptable, no matter how great the benefits.

Subjects who choose option C are considered to hold a protected value on the issue. For a broad range of scenarios, ranging from dolphins dying in nets used for tuna fishing to euthanasia, a fair number of participants choose option C. For example, in our own studies with Northwestern University undergraduates, approximately 1/3 of students say that abortions are not acceptable, no matter how great the benefits.

Laboratory studies have also found other distinctive properties associated with protected values (see Baron and Ritov, this volume, for a review). First, they are intrinsically linked to morality, and are perceived as objective, absolute prohibitions, regardless of what a particular person thinks. Second, challenging protected values may be associated with the experience of anger and denial of the need for tradeoffs (Baron & Spranca, 1997). Tetlock's research has shown that the mere contemplation of forbidden tradeoffs undermines the sacredness of the value and may result in a sense of moral contamination. Participants who were asked to consider such



tradeoffs later engaged in moral cleansing to compensate for the contemplation (Tetlock et al., 2000). For example, participants who were asked to evaluate forbidden tradeoffs (such as human life for money, for example) later showed an increased willingness to sign up for organ donation programs.

The moral status of sacred values can produce aversive reactions towards others who challenged their sacredness in the course of making a decision. Tetlock et al. (2000) presented their participants with scenarios in which a hospital administrator was trying to choose between saving the life of little Johnny or saving \$1 million for the hospital. The participants not only disliked more the administrator who chose the money over the child's life, but they were also sensitive to the amount of time this decision process took. When the decision favoring the money was slow and difficult the administrator was liked even less, suggesting that the longer a person contemplates a tradeoff that challenges a sacred value, the more morally corrupt they are perceived to be.

### **2.3 Omission bias and quantity insensitivity**

For a decision theorist, one of the most interesting properties associated with protected values is known as *quantity insensitivity*. Participants who endorse a protected value for an issue tend to focus on the permissibility of acts and tend to give less weight to the act's consequences. This means they may judge two similar acts being equally impermissible regardless of the difference in the total harm done. For example, an ardent environmentalist might say that clear cutting of 1 acre of old growth forest is no less wrong than clear cutting 5 acres of old growth forest. Nonconsequentialist decision principles like these, which seem consistent with a kind of lay deontology, are difficult to explain in terms of traditional decision theory, which presumes (broadly) that people's choices are a reflection of desired outcomes.

The apparent insensitivity to consequences associated with protected values may also be characterized as a relative “hypersensitivity” to the permissibility of actions, per se, that sometimes results in what Baron and colleagues refer to as “omission bias” (Baron, Spranca, & Minsk, 1991, Baron and Ritov, this volume). Because harmful actions are perceived as violating moral prohibitions by people with protected values, they may prefer not to engage in these actions, even when they mitigate a larger risk. For example, when faced with a situation where the only way to save 20 species of fish downstream is to open a dam that will cause the extinction of 2 species of fish in a river, a fair number of participants with protected values say they would not open the dam (even though not doing so results in a loss of 20 species). This preference for (more) harmful omission over (less) harmful action represents “omission bias”. Even when people with protected values are willing to act, they show less willingness to make tradeoffs than do people for whom the value is not sacred or protected. That is, people with protected values show a relative hypersensitivity to actions that is more commonly referred to as “quantity insensitivity.”

In summary, there is ample evidence from laboratory studies that sacred or protected values can be absolute and resistant to any tradeoffs. This evidence has come from a range of converging measures. Nonetheless, as we suggested earlier there is also evidence that these values may, at the same time, be at least somewhat flexible and context-dependent. In the next two sections we will review some of this evidence.

#### **2.4 Tragic vs. taboo tradeoffs and the flexibility of moral decision making**

Phillip Tetlock has proposed a sacred value protection model, where a tradeoff between two sacred values in conflict (*tragic tradeoff*) is more permissible than a tradeoff between sacred and secular values (*taboo tradeoff*). In the example with the little boy who needed an expensive surgical operation, the hospital administrator was judged more positively if he traded the life of

one boy for the life of another boy, compared to the situation where he traded the boy's life for money. Participants in the tragic tradeoffs condition were also more positive if the administrator made a slow rather than quick decision, no matter which boy he chose.

If taboo tradeoffs are morally outrageous while tragic tradeoffs are relatively acceptable, then a straightforward strategy for enhancing a taboo tradeoff is either to frame it as a regular tradeoff, or to present it as a tragic tradeoff. For example, Tetlock (2005) showed that participants who were firmly opposed to markets for body organs would change their position if the apparent taboo tradeoff was elaborated in different ways. First, the experimenter added the information that the transactions would be allowed only to save lives that otherwise would be lost, putting the entire tradeoff in the domain of sacredness. Second, participants were told that the poor would receive financial aid when they needed a transplant, and that they would be prevented from selling their organs because of financial pressures, alleviating moral concerns about fairness and framing the tradeoff in secular, monetary terms. After receiving this additional information, 40% of the participants who had held sacred values against body organ markets changed their minds.

### **2.5 Presentation order: Judgment in the context of previously-viewed options**

Further evidence on the flexibility of protected values comes from an experiment looking at order effects on judgments of permissibility. Nine scenarios involving moral dilemmas adapted from Bartels (2008) were selected, based on the mean endorsement scores they received in Bartels' original research. They included three types of dilemmas. The first was a standard one, where the question is whether an agent should conduct a harmful act in order to bring about some greater good (e.g., the warden of a POW camp tells a hostage that he can save himself and a number of other hostages from execution if he appeases the warden by killing one of them). The second was a vivid version, where the action was described in detail (e.g. "You are handed a

knife, *and your fellow hostages scream in terror, pleading for their lives*" (emphasis added), and the third was a large consequences version, where the good that the action brings about is much larger than in the standard version. Two separate groups of participants rated their approval of each of actions: one read the vivid scenarios first, then the standard, followed by the large consequences scenarios and the other was given the scenarios in the opposite order.

If these sorts of tradeoff scenarios operate to trigger content dependent deontological rules, then order of presentation should not matter. If, as Hauser (2006) suggests, participants feel bound to be consistent in their strategies, the order should matter and the strategies used in the first few scenarios should carry over into later scenarios (e.g. the participants who see the large consequence scenarios first should be more consequentialist on the standard scenarios than those given the vivid scenarios first). The final possibility is that initial scenarios set up contrast effects (e.g. "I acted to save 20 people but this only involves 5," or "this one isn't as gruesome as the first few."). In this case the first few scenarios should affect later judgments but in the opposite direction from adopting a consistent strategy across scenarios.

Order of presentation affected responses to all three types of scenarios and took the form of contrast effects. Exposure to vivid choice scenarios, which were designed to direct attention to highly morally salient (and affectively charged) features of the situation, appears to have had the effect of making participants more willing to perform the action posed by the scenario, whether the scenario was described neutrally or with information about large consequences. Conversely, exposure to scenarios including information about large consequences appears to reduce endorsement of action in the standard scenarios, and even more so in the vivid scenarios.

These findings further illustrate the malleability and context dependence of moral reasoning found in the experiments of other researchers. Results like this are in sharp contrast with the presumable absoluteness and stability of these values and might even question their existence.

## **2.6 “Pseudo-sacred” values?**

One logical response to claims of infinite utility or an unwillingness to mix the sacred and the secular is to challenge them. There is a body of research suggesting that sacred values may only be “pseudo-sacred” (e.g. Thompson & Gonzalez, 1997), that they may represent only a form of “posturing” (Baron & Spranca, 1997) and that people with sacred values will, in fact, make tradeoffs when they are indirect (e.g. Irwin & Scattone 1997; Tetlock, 2000) or when pushed (Baron & Leshner, 2000). For example, Baron and Leshner asked participants to think about possible counterexamples to the absoluteness of their protected values. This reduced people’s endorsement of protected values (Baron & Leshner, 2000). Furthermore, people with protected values were sensitive to the probability of harm associated with actions that violate these protected values. When the probability that an action would cause harm was sufficiently small, participants with a protected value against this harm behaved similarly to participants who did not have the protected value. Altogether, these results raise the possibility that sacred values are akin to a self-presentation strategy aimed at posturing or claiming the high moral ground but having little depth beyond those goals. In contrast, we think these results are part of the puzzle of how sacred values can be both rigid and flexible, and that they should not be used to argue that these values do not exist at all.

## **2.7 Summary**

There is strong empirical evidence for the principal features of sacred or protected values, suggesting that they focus the moral agent on the permissibility of actions and away from the consequences associated with those actions. There is perhaps equally strong evidence that these values are malleable and context specific. The challenge is to reconcile these distinct patterns of

findings. We have attempted to do so by employing a range of procedures aimed at elucidating cognitive processes associated with moral values.

### **3. A COGNITIVE PERSPECTIVE ON SACRED VALUES**

The apparent contradiction between the absoluteness of sacred or protected values and their sensitivity to contextual factors raises a range of important questions. We have approached these questions from a cognitive perspective, focusing on the processing of morally relevant information.

First, we have been interested in how a decision maker recognizes a situation as relevant to a sacred value. We have already seen that there are straightforward ways to do this, some bordering on circularity - participants say they would not allow some action regardless of the benefits and then are given the opportunity to demonstrate this in a scenario involving tradeoffs. To be sure, researchers have employed converging, noncircular measures (see Bauman and Skitka, this volume; Baron and Ritov, this volume) but we sought another form of converging evidence that did not involve explicit questions about values.

Given that a morally motivated decision maker must recognize the moral relevance of a situation before choosing any particular strategy, we examined whether morally relevant information has a cognitive advantage over neutral information even in the initial stages of processing. That is, does moral information affect bottom up attentional processes? To address this question we looked at attentional correlates of morally relevant content, and its broader interplay with more distant cognitive tasks.

Our second research direction is concerned with allocation of attention and cognitive resources to different aspects of morally relevant information. Once a situation is recognized as morally relevant, the next step is to isolate its important components. Even a short scenario containing relatively few facts relies on much broader pre-existing knowledge, and accessing the

most relevant information may not be an easy task. When participants who care about a particular moral issue or domain encounter a situation in which they have to decide on an action that may produce harm, they might consider this particular action as the most relevant information and base their decision solely on it. That is, people may rely on decision principles that are consistent with a kind of lay deontology. However, when other aspects of the situations are made salient, the same person may come to a different decision. In a series of experiments, we show that participants who hold strong moral values are quite sensitive to cues that prompt one or another aspect of the situation, and can be swayed either to ignore or to focus on the actions versus the consequences of actions.

### **3.1 Protected values and attention**

Morally relevant information likely captures attention. When a person encounters a word that has high personal relevance, such as their own name, this information is more easily and more rapidly processed than less relevant information (Mack & Rock, 1998). When a moral concept is of high personal relevance it may also have processing priority. A person who is strongly committed to a pro-life cause might find the word *abortion* more salient than a person who holds a neutral position when it is encountered in text. Similarly, a pro-choice activist from an on-campus student organization might be drawn to the same word, albeit for very different reasons.

### **3.2 Stroop effects**

A classic task for studying attentional selection has been the Stroop effect (Stroop, 1935; MacLeod, 1992), where participants have to name the font color by which a word is displayed. The standard finding is that the speed of naming the font color of a word is influenced by its semantic properties. The most frequent demonstration of this effect is when the words are

themselves color terms that do not match the color of the font. But semantic effects extend well beyond this case.

Following previous research on the emotional version of Stroop task (Gotlib & McCann, 1984, Williams, Mathews, & MacLeod, 1996) and taboo words Stroop (MacKay, Shafto, Taylor, Marian, Abrams, & Dyer, 2004), we expected that the meaning associated with moral values will capture attention and as a result will interfere with color naming. We presented participants with a list of abortion related words, such as *pregnancy*, *trimester*, *fetus*, and neutral words, such as *river*, *boat* and *grant* (Iliev & Medin, 2006). In addition, we measured the personal importance of the abortion question on a scale ranging from *not important at all* to *extremely important* and the particular position on the issue ranging from *strongly pro-life* to *strongly pro-choice*.

A comparison between the reaction times for neutral versus the abortion-related words showed that, overall, participants were slower to name the color of the abortion-related words. However, using a median split to separate the participants into low versus high importance, the low importance participants did not show any difference between the two types of words. The high importance group, however, took significantly longer to name the color of abortion-related words compared to neutral words. Whether one was strongly pro-life or strongly pro-choice did not affect the magnitude of the Stroop effect.

After the Stroop experiment, the participants were given a surprise memory task, where they were asked to recall any words that they might have seen. There was no difference between the high and low importance groups in the number of neutral words recalled. However, compared to the low importance group the high importance participants recalled significantly more abortion-related words. These findings suggest that the low level capturing of attention is further accompanied by higher level processing related to memory and awareness.

The results from the Stroop task provide evidence that morally relevant information is more difficult to inhibit compared to neutral information. Further, this study suggests that the



perceived importance of the concept, rather than the specific attitude on the issue, influences attention selection. Note that emotional or taboo-word Stroop effects cannot explain this pattern, since abortion related words are not likely to trigger the same emotional valuation across strongly pro-life and strongly pro-choice participants. Relevance rather than similarity in meaning is the key.

### **3.3 Further cognitive consequences: Anchoring effects**

If morally relevant information is processed differently, then this should be reflected in other cognitive tasks as well. Although moral values have been linked to cognitive phenomena such as quantity insensitivity and omission bias, their role has not been explored in more peripheral judgment tasks which do not explicitly ask about moral judgments (for choice tasks, see Irwin & Scattonne , 1997 and McGraw & Tetlock, 2005).

We chose to examine the influence of moral values on anchoring effects. A classical example comes from Tversky and Kahneman (1974). Participants were given an arbitrary number (10 or 65) determined by spinning a wheel of fortune, and then were asked to compare this number to the percentage of African nations in the UN. After the participants answered whether this percentage was higher or lower than the arbitrary number, they asked the further question of exactly what that percentage was. Participants in the high-anchor condition gave an average estimate of 45%, compared to 25% for the low anchor condition. This effect is surprising because the anchor is presumably useless for estimating the correct answer and should have no influence on it.

Most explanations of the anchoring effect refer to attentional processes, characterizing the effect in terms of selective accessibility, activation, and use of information in memory. For example, one explanation makes use of an activation model similar to semantic priming (Strack & Mussweiler, 1997). In this model, the comparison process between the anchor and the target

makes the similarities between the two more accessible, influencing the absolute judgment by retrieving a biased sample of information

Wilson, Houston, Etling, and Brekke (1996) found a basic anchoring effect that was positively correlated with the amount of attention paid to the anchor, despite the fact that the anchor and the target were semantically different. For example, participants were presented with an ID number and were later asked to check whether the ID number was written in blue or red ink (attend-to-color condition), or whether it was lower or higher than 1920 (attend-to-number condition). The participants then had to estimate the number of physicians listed in the phone book. Anchoring effects were observed in both conditions, but the effect was larger in the attend-to-color condition (but see Brewer & Chapman, 2002 for a critique).

Because of its established links to attention, we used the simple version of this type of anchoring task to examine the indirect link between moral values and decision making. Consider the following scenario (Iliev & Medin, 2006):

“In recent research, published in the Netherlands, it was estimated that the percentage of women who have had an abortion in Europe almost doubled in the last  $X$  years. What is your best guess for this percentage in USA nowadays?”

Instead of  $X$ , participants saw a number between 5 and 100. This number is not informative with respect to the answer to the question, but if sufficient attention is paid to it, people may be influenced by this anchor. Our expectation is that the attention of a person who cares about abortions will be captured more easily by the concept of abortion presented in the scenario, which will result in distraction from the anchor. If, on the other hand, the abortion concept is not that salient, more attention will be dedicated to the anchor itself. Thus, participants who care more about the abortion issue should show smaller anchoring effects than participants for whom the issue is less relevant.

That was what we found. We used the same importance and polarity scales from the Stroop task to see whether and how moral relevance affected anchoring. The low importance

group exhibited reliable anchoring effects, but the high importance group did not. Only attitude importance predicted (the absence of) anchoring—one's specific attitude (whether one was pro-life or pro-choice) was not predictive.

In another experiment we used the protected value measure, where instead of high and low importance groups we compared subjects who found abortions impermissible under any circumstances to those who did not have such extreme values. We ran the same anchoring task, and correlated the answer and the anchor for each of the two groups. The group without protected values on abortions showed the typical anchoring effect and the group with protected values showed virtually no anchoring. These results, combined with the previous anchoring study, suggest that including morally relevant information in a scenario overshadows the salience of the irrelevant anchor, eliminating the basic anchoring effect

### **3.4 Summary**

The experiments described above suggest that morally relevant information is more likely to attract attention and influence performance on peripherally related decision making tasks. If moral values influence attention and the distribution of cognitive processing resources, then by manipulating attentional cues, one might be able to influence some of the properties associated with these values. The studies reported next address this question.

## **4. ATTENTIONAL INFLUENCES AND THE ACCEPTABILITY OF TRADEOFFS**

Any real world tradeoff has many aspects, and attending to some of them and not others may lead to different perception of the tradeoff. In series of experiments, we manipulated the salience of different aspects of tradeoffs, showing that when a decision is framed one way, strong

moral values may lead to a greater focus on acts (consistent with deontology), and when framed differently, values may lead to a more consequentialist focus.

Compared to dispassionate decision makers, those who view a resource as sacred may care more about the consequences realized by the resource and more about how the resource is treated (e.g., actions that actively harm the forest are wrong, simply stated). So we might expect the judgments of morally-motivated decision makers to show larger effects of manipulations that direct attention to the permissibility of harmful actions (and away from consequences) and those that direct attention to consequences (and away from the sometimes harmful antecedent actions).

#### **4.1 Framing the question**

Recall that one of the major properties associated with protected values is quantity insensitivity. Stemming from the deontological nature of protected values, the nature of the act appears to be more important than the particular consequences. For example, Baron and Spranca (1997) suggested that “Quantity of consequences is irrelevant for protected values. Destroying one species through a single act is as bad as destroying a hundred through a single act (p. 5).” They argued that abortion opponents might oppose government spending on family planning programs that carry out abortions in developing countries, even if this spending ultimately reduces the number of abortions performed.

Baron and Spranca (1997) measured quantity (in)sensitivity with the following questions where X is some prohibited action (e.g. abortions):

*Is it equally wrong for X to happen as 2X to happen?*

*Is it worse for 2X to happen than for X to happen?*

If participants are consistently insensitive to quantity, they should answer “yes” to the first question and “no” to the second. However, as decades’ worth of data on framing effects show, logically equivalent rewordings of questions are not necessarily interpreted in the same way. One

explanation for why this occurs is that the interpretation of a situation is dependent on what frame of reference is constructed by the wording of a question. Thus, 75% saved versus 25% lost directs people's attention to different reference points (McKenzie & Nelson, 2003).

The property of quantity insensitivity may be influenced by where the question directs their attention. For example, the reference point activated when people are asked if "is it less wrong for X to happen than for 2X to happen" may direct attention towards the distinction between an act occurring and it not occurring. This would make the number of acts occurring more or less irrelevant because the reference point is the absence of a harmful act. This form of the question may direct attention to the permissibility of the act (and away from consequences), thus eliciting responses that appear non-consequentialist.

However, the logically equivalent statement but reworded as "it is worse for 2X to happen than for X to happen" may direct attention towards the quantitative difference between 2X and X implying that some amount of X will inevitably occur. When comparing outcomes, less harm seems better than more harm. We expected this question to direct attention to consequences (away from the act) and elicit responses from people with protected values that are sensitive to quantity, and thus more consistent with consequentialism .

In a series of experiments we showed there was a large difference in how quantity sensitive participants were depending on how the question was framed (Sachdeva & Medin, 2008). The "less" question (i.e., is it less wrong for X to happen than for 2X to happen) elicited quantity insensitivity. Participants without protected values were sensitive to quantity in only 48% of the cases. Consistent with predictions, this quantity insensitivity was especially pronounced for people with protected values (only 38% of whom were sensitive to quantity). Overall, the "more" question elicited quantity-sensitive responses (in 80% of cases). Additionally, people with protected values were significantly more quantity sensitive (84% of the cases) than participants without protected values (74%).

These results can be taken to indicate that participants with protected values can act in accordance with both consequentialist and deontological principles. Our data suggest that whether people's responses appear deontological or consequentialist is highly dependent on where attention is directed. By manipulating the frame of reference provided, we are able to vary one of the most basic properties of protected values.

#### **4.2 Framing the response alternatives**

A series of studies conducted by ourselves and others has found that when attention is directed towards harmful actions, participants with sacred values appear less consequentialist (less quantity sensitive, more act sensitive) than their more dispassionate peers. For example, Bartels and Medin (2007) examined tradeoffs like the opening of a dam scenario described earlier in two ways. In one condition adapted from Ritov and Baron (1999), participants were first asked “would you open the dam? Y/N” (recall that doing so killed 2 species upstream to save 20 downstream). Then they were asked for an upper threshold value for harm—the largest number of fish killed by the action at which they would open the dam. It is perhaps unsurprising that participants who have protected values (those who indicate on another measure that killing fish species is wrong no matter the consequences) provide lower thresholds, or appear “quantity insensitive” relative to participants without protected values. This is a standard result, replicated many times (see Baron & Ritov, this volume).

A second group of participants was presented with a normatively equivalent elicitation procedure which omitted the initial yes/no question and instead asked about a range of tradeoff values (a procedure derived from Connolly & Reb, 2003). These participants were asked five tradeoff questions per scenario, for example: “Would you open the dam if it caused the death of 2 (6, 10, 14, 18) species of fish?” The prediction was that by holding the action constant and

varying the consequences within-participants, attention would be directed away from the permissibility of harmful action and toward outcomes.

In this condition, participants with protected values were *more* willing to make tradeoffs than participants without protected values. To our knowledge, this was the first demonstration of a link between endorsing a constraint (a protected value) and being more willing to make tradeoffs to achieve the best consequences. Baron and Ritov (this volume) conducted an internet experiment and failed to replicate our results. We are uncertain as to why, and more research is needed to elucidate other moderators of these effects. For now, we take our evidence, along with the more versus less results, to suggest that both means (acts) and ends (consequences) matter, and matter a lot, for people with protected values.

#### **4.3 Judgment in the context of one versus two options**

We have just seen how framing a question or the presentation of answer alternatives can be a prompt as to what is relevant in a given situation. Another way to make some aspects more salient is to juxtapose situations, where the relative similarities and differences may serve to highlight some parts more than others.

In further work (Bartels, 2008), we showed that participants with protected values appeared especially sensitive to actions when policy decisions like those above (where a harmful action mitigates a larger risk) were evaluated in isolation. However, when participants were invited to compare omissions (with worse consequences) and harmful actions (with better consequences), they focused on ends, strongly preferring the net-benefit-maximizing actions.

In the first of these studies, participants read about government administrators' decisions to either (i) knowingly do harm to mitigate a greater risk (e.g., "Rich wants to save the children from the disease. He first calculates that administering the vaccine will kill 480 children. Knowing that doing so will kill many children, he chooses to vaccinate the children.") or (ii) to

merely allow the harm to happen (e.g., “Julie does not want to kill any of the children with the vaccine. So, the vaccine is not administered. The 600 children die.”). In this study, participants read about a decision and made two judgments, one about whether the administrator’s decision broke a moral rule, and a judgment of (dis)approval for the decision. When each decision was evaluated in isolation, participants with protected values were less likely to approve of harmful actions. The correlation between judgments of rule violation and disapproval was higher among the items for which participants indicated protected values than those for which they did not. These findings suggest that morally-motivated decision makers evaluate morally-relevant decisions by attending to the relationship between an act and a moral rule.

In the second study, both solutions to the problem (the cost-benefit-driven action and the prohibition-driven omission) were presented to participants on a single page, inviting direct comparison between the decisions. This context was intended to highlight the differences in outcomes (i.e., that 120 more children are saved with vaccination), and under this procedure, people with and without protected values strongly preferred the net-benefit-maximizing action. The results of both Bartels and Medin (2007) and Bartels (2008) demonstrate that adding or subtracting options can shift attention from acts to consequences (see also Unger, 1996), and that what is attended to is a powerful determinant of what is valued and preferred.

#### **4.4 The attraction effect: Choosing in the context of three options**

The above results suggest that participants try to make sense from the information available by searching for available reference points, and when these reference points are readily available, they provide the basis for making a decision. In the next studies we further explored context effects in the form of choice alternatives. To do this we used a popular paradigm in decision making research that elicits what is referred to as the *attraction effect* (Huber, Payne, & Puto, 1982).



Attraction effects are preference shifts induced by an irrelevant alternative. Consider, for example, a choice between two cars A and B, where A has better gas mileage, while B has better performance. Normatively speaking, adding a third car C, which is slightly worse than B on both mileage and performance, should not change preferences, since no sound decision maker will choose option C, which is dominated by car B. If someone prefers car A over B when there are only two options in the choice set, they should keep their preferences when C is added. Nonetheless, a typical empirical finding is that adding the irrelevant option boosts the preferences for the dominating option, in our example, car B. If instead of adding C, we add a car D that is slightly worse than car A in both facets, choices of car A go up relative to the two-choice situation.

All explanations of attraction effects (e.g., those based on theories from Parducci, 1965; Roe, Busemeyer, & Townsend, 2001; Tversky & Simonson, 1993; Weddell, 1991) assume that there must be some type of tradeoff between the competing dimensions, otherwise the effect would not arise. Going back to the car example, if the gas prices go high enough that the only relevant dimension is gas-mileage, adding a third alternative will not be able to shift the strong one-dimensional preferences.

In the same way, if a decision maker holds protected values on one of the dimensions but not the other, then one might expect that they will maximize only on this dimension and no context effects will be observed. On the other hand, if both dimensions involve moral goods, then the situation may entail a tragic rather than a taboo tradeoff. In this case, if the particular context makes one alternative or dimension more salient than the other, we may observe attraction effects. To test this hypothesis we designed scenarios where participants were told that they are in a position of a decision maker who has to allocate money to different charities, varying from environmental programs to building new shelters for the homeless (Iliev & Medin, 2007). For each round they had to choose only one of three programs, and each of the programs

combined two different charities (see Fig. 1). Following the paradigm that elicits attraction effects, one of the programs was better for one charity, but worse for another, the second program was the opposite, and the third was slightly worse on both dimensions than one of the other two. Separately, we measured whether the participants had protected values for each of the charities in question.

---

Insert Figure 1 about here

---

If a participant has no protected values for either of the dimensions, or has them for both, then attraction effects should be possible. If they have a protected value for one dimension but not the other, then that dimension should dominate choice and no context effect should be observed.

We found robust attraction effects, and it made no difference whether one, two or none of the dimensions were protected against tradeoffs. In agreement with our previous studies, these results suggest that morally motivated decision makers take into account much of the information they are presented with, and if the information makes one or another part of the tradeoff more salient, this is reflected in subsequent decisions.

#### **4.5 Summary**

Moral values and information related to them attract attention, yielding large Stroop effects on the one hand and less influence of an irrelevant anchor on the other. In addition, our laboratory studies have repeatedly shown that subtle manipulations can direct attention towards

the action dimension, where we observe relative insensitivity to quantity or direct attention towards outcomes where we observe greater sensitivity to quantity. Altogether our underlying message is that values affect attention and attention affects how values are translated into judgments.

Although we found attentional correlates of moral values, we did not straightforwardly manipulate attention. By varying the reference points to which a situation is compared we emphasized one or another factor to be considered, but we did not directly manipulate cognitive processing. In future work manipulations such as time pressure and memory load might substitute for the more passive types of manipulations that we used here (see Greene, Morelli, Lowenberg, Nystrom, & Cohen, 2008).

A blunt summary of our results is that sacred or protected values are far from rigid; instead, they are “all over the place” with respect to judgment and decision making. A core challenge is to understand how these values can be at once fundamental to personal identity and (sometimes) lead to dramatic actions and sacrifices but at the same time be so flexible that modest experimental manipulations lead to qualitative changes in judgments. Many of our studies have focused on attention to actions versus outcomes (deontology versus consequentialism) and it remains to be shown how actions and outcomes associated with sacred values are or are not integrated in people’s everyday lives. If nothing else, our studies show that one cannot rely on the results of a single task to make claims about the character of sacred values.

## **5. GENERAL DISCUSSION**

Our studies show that sacred or protected values involve more than the application of deontological rules and more than the neglect of consequences. Indeed sometimes we find that sacred values are associated with hyper-sensitivity to consequences. Related work by Carmen Tanner and her associates suggests that deontological and consequentialist orientations are not logical opposites but rather are orthogonal or even weakly correlated (Hanselmann & Tanner, in press; Tanner, Iliev & Medin, 2007).

In many respects, these results should not be surprising. The same person who would feel offended by an offer to purchase his daughter or his dog and refuses to think of them as secular goods, cares a great deal about what happens to both daughter and dog. Deontological rules may be in the service of achieving value-related goals.

The bulk of our studies rely on data from undergraduates. One might question the depth of their protected values or at least whether our results on attention would hold for people whose dedication to some cause spans decades or saturates their lives. We know relatively little about how strong moral values evolve over time and experience. Still, we will make some speculations about the generality of our results.

Questions about generality are necessarily empirical questions. (One nice example of real world relevance are Baron's studies using scenarios concerning decisions to vaccinate children and connecting them with policy decisions (see also Connolly & Reb, 2003). First, we suspect that the Stroop results demonstrating attention capture by words related to protected values would be equally robust or even more robust for those people who have stronger commitment to their protected values. Second, framing effects also may be robust. For example, we have found that the more versus less framing affects sensitivity to quantity in a sample of Palestinians who endorse jihadist actions ("Does God love the martyr less who kills one of the infidels than the martyr who kills ten? versus "Does God love the martyr more who kills ten of the infidels than the martyr who kills one?"). Other interviews with members of Hamas suggest that the an act of

suicide martyrdom is not readily set aside for other moral imperatives such as taking care of a sick parent but, at the same time, is readily given up if a roadside bomb will accomplish the same goal (Atran, 2004). The chapter by Ginges (this volume) also shows clear context (framing) effects among Middle East populations whose lives are organized by sacred values.

Nonetheless, cautions about generalizability should remain salient. The fact that we observe framing effects in the context of individual interviews with a sample of Palestinian refugees does not imply that a framing manipulation will also carry the day in social contexts and everyday lives where people are exposed to competing frames and frameworks for understanding events. Druckman and Nelson (2003), for example, showed that framing effects in more realistic contexts are highly susceptible to exposure to additional information during within-group discussions. It may be that activists have so much information at their disposal that simply bringing up the topic activates a large ensemble of actions, attitudes and associations that might tend to swamp or re-frame any effects of new information and its framing.

An undergraduate college student saying that destruction of old grown forest is unacceptable no matter how great the benefits might make the same claim as a Greenpeace activist, but their real world behavior is likely to be very different. Although both have strong opinions on the issue, for one of them the link between moral values and moral action would be much stronger than for the other. When presented with similar situations, they might find different factors to be most relevant, or may respond differently to external attentional manipulations. In short, despite the surface similarity in attitudes they may think and behave in dramatically different ways.

We close with some speculations on the malleability and rigidity of sacred or protected values that take the form of two analogies. One is that sacred values may be seen as requiring a catalyst in order to be translated into decision behavior. The Northwestern undergraduate environmentalist may differ from the Greenpeace activist solely in terms of the latter having

been exposed to a catalyst. Catalytic events may range from things like media coverage, the presence of some triggering event such as some Supreme Court decision, belonging to a soccer team whose members share the same attitudes and life circumstances (e.g. Atran & Sageman, 2006), or chronic, saturated experiences such as citizens of Gaza facing seemingly endless delays at checkpoints (Atran, 2004).

The second similar analogy is that values related to personal identity may be like biological needs. We do not organize our lives around food, water, or air, nor do we assign them infinite utility, at least in ordinary circumstances. But we cannot live without them. We cannot imagine someone giving up our right to air for even 30 minutes, no matter what reward we are offered. And we can readily imagine circumstances, such as being trapped under water, where the quest for air would become a central concern. To give up air, water and food is to give up life.

To give up sacred values may have similar consequences for personal identity. Northwestern undergraduates say they would not sell their wedding ring, even for a price well in excess of its market value (Medin, Schwartz, Blok, & Birnbaum, 1999). They say it is about meaning, not money. No doubt these undergraduates could be challenged on this (e.g. “What if you needed the money to pay for a kidney transplant for your spouse?”) and we imagine they would yield, but it is hard to see this as undermining this sacred symbol.

In ordinary life our sacred values are not irrelevant but they are rarely challenged or unambiguously directly relevant. In point of fact, few of us are in a position to open a dam to save fish species at the cost of other species and almost never is the situation so unambiguous (e.g. “What other possible consequences of opening the dam should we be worrying about?”). Presidents and paramedics may represent exceptions but presidents have advisors and paramedics have protocols.

For either analogy the key theoretical or conceptual questions center around how different settings are seen as relevant to moral values, how the social context interacts with these values and how settings and social contexts map between moral values and actions. Asking individuals about unlikely hypothetical scenarios is only one step on a long journey to understanding.

## REFERENCES

- Atran S. & Sageman M., (2006) Connecting the Dots, *Bulletin of the Atomic Scientists*, August 2006.
- Atran, S., Axelrod, R., & Davis, R. (2007) Sacred Barriers to Conflict Resolution. *Science*, 317, 1039-1040.
- Atran, S., (2004, October). Tuning out Hell's Harpist, Paper Presented to the Permanent Monitoring Panel on Terrorism, World Federation of Scientists, Geneva.
- Baron, J., & Greene, J.D. (1996). Determinants of insensitivity to quantity in valuation of public goods. *Journal of Experimental Psychology: Applied*, 2, 107-125.
- Baron, J., & Leshner, S. (2000). How serious are expressions of protected values? *Journal of Experimental Psychology-Applied*, 6, 183-194.
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and Human Decision Processes*, 94, 74-85.
- Baron, J., & Spranca, M. (1997). Protected values. *Organizational Behavior and Human Decision Processes*, 70, 1-16.
- Bartels, D. M. (2008). Principled moral sentiment and the flexibility of moral judgment and decision making. *Cognition*, 108(2):381-417.
- Brewer, N.T. & Chapman, G.B. (2002) The fragile basic anchoring effect. *Journal of Behavioral Decision Making*, 15, 65-77.



- Connolly, T., & Reb, J. (2003). Omission bias in vaccination decisions: Where's the "omission"? Where's the "bias"? *Organizational Behavior and Human Decision Processes*, 91, 186-202.
- Druckman, J.N. & Nelson, K.R. (2003), Framing and deliberation: How citizens' conversations limit elite influence. *American Journal of Political Science*, 47, 729-745.
- Durkheim, E. (1954) *The Elementary forms of the Religious Life*. Translated by J. W. Swain. New York: The Free Press.
- Egeth, H.E. & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Reviews of Psychology*, 48, 269-297.
- Eliade, M., (1959) *The Sacred and the Profane: The Nature of Religion*. Translated by W. Trask, London: Harcourt Brace Jovanovich.
- Fiske, A. P., & Tetlock, P. E. (1997). Taboo trade-offs: Reactions to transactions that transgress the spheres of justice. *Political Psychology*, 18, 255-297.
- Franconeri, S., L., Alvarez, G. A., & Enns, J. T. (2007). How many locations can you select? *Journal of Experimental Psychology: Human Perception & Performance* 33(5), 1003-1012.
- Frank, R. H. 1988. *Passions Within Reason: The Strategic Role of the Emotions*. New York: Norton.
- Ginges J., Atran, S., Medin, D., & Shikaki, K. (2007) Sacred Bounds on Rational Resolution of Violent Political Conflict., *Proceedings of the National Academy of Sciences* 104, 7357
- Gotlib, I.H., & McCann, C.D. (1984). Construct accessibility and depression: An examination of cognitive and affective factors. *Journal of Personality and Social Psychology*, 47, 427-439.
- Greene, J.D., Morelli, S.A., Lowenberg, K., Nystrom, L.E., Cohen, J.D. (2008) Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, 107, 1144-1154

- Hanselmann, M. & Tanner, C. (2008). Taboos and conflicts in decision making: Sacred values, decision difficulty, and emotions. *Judgment and Decision Making*, 3, 51-63.
- Hauser, M. D. (2006) *Moral minds: How nature designed our universal sense of right and wrong*. New York: Harper Collins
- Huber J., Payne J., & Puto C. (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9, 90-98.
- Iliev, R. & Medin, D. L. (2006) Cognitive Concomitants of Protected Values, Annual Meeting of the Psychonomic Society, Houston.
- Iliev, R. & Medin, D., ( 2007) Context Effects in the Moral Domain, Annual meeting of the Society for Judgment and Decision Making, Long Beach.
- Irwin, J. & Scattoni, J., (1997), Anomalies in the valuation of consumer goods with environmental attributes, *Journal of Consumer Psychology*, 6, 339-363.
- Ling, S. & Carrasco, M., (2006). Sustained and transient covert attention enhance the signal via different contrast response functions, *Vision Research*, 46, 1210-1220.
- Mack, A., & Rock, I. (1998). *Inattention blindness*. Cambridge, MA: MIT Press.
- MacKay, D. G., Shafto, M., Taylor, J. K., Marian, D. E., Abrams, L., & Dyer, J.R. (2004). Relations between emotion, memory and attention: Evidence from taboo Stroop, lexical decision, and immediate memory tasks. *Memory & Cognition*, 32, 474-488.
- MacLeod, C. M. (1992). The Stroop task: The "gold standard" of attentional measures. *Journal of Experimental Psychology: General*, 121, 12-14.

- McGraw, P. & Tetlock, P.E. (2005). Taboo trade-offs, relational framing, and the acceptability of exchanges. *Journal of Consumer Psychology*, 15(1), 35-38.
- McKenzie, C. R. M., & Nelson, J. D. (2003). What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin and Review*, 10(3), 596-602.
- Medin, D. L., Schwartz, H. C., Blok, S. V., & Birnbaum, L. A. (1999). The semantic side of decision making. *Psychonomic Bulletin & Review*, 6(4), 562-569
- Nichols, S. 2004. *Sentimental Rules: On the Natural Foundations of Moral Judgment*. New York: Oxford University Press.
- Pardicci, A. (1965). Category Judgment: A Range-Frequency model, *Psychological Review*, 72(6), 407-418.
- Prinz, J. (2007). *The emotional construction of morals*. Oxford: Oxford University Press
- Ritov, I. and Kahneman, D. (1997). How people value the environment: Attitudes versus economic values. In Bazerman, M., Messick, D., Tenbrunsel, A., and Wade-Benzoni, K. (Eds.). *Psychological Perspectives to Environmental and Ethical Issues* (pp. 33-51). The New Lexington Press.
- Ritov, I., & Baron, J. (1999). Protected values and omission bias. *Organizational Behavior and Human Decision Processes*, 79, 79-94.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108, 370-392.

- Sachdeva, S. & Medin, D.L. (2008). Is It More Wrong to Care Less: The Effects of More and Less on the Quantity Sensitivity of Protected Values. Proceedings of the 30<sup>th</sup> Annual Meeting of the Cognitive Science Society, CogSci2008. Hillsdale, NJ: Earlbaum.
- Simon, H. A. (1957). *Models of man: Social and rational*. New York: John Wiley and Sons.
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27, 76-105.
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27, 76-105.
- Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, 73, 437-446.
- Stroop, J.R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643-662.
- Tanner, C., Medin, D.L. & Iliev, R. (2007). Influence of deontological vs. consequentialist orientations on act choices and framing effects: When principles are more important than consequences. *European Journal of Social Psychology* online
- Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in Cognitive Sciences*, 7, 320-324.
- Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology*, 78, 853-870.
- Thompson, L., & Gonzalez, R. (1997). Environmental Disputes: Competition for Scarce Resources and Clashing of Values. In M. Bazerman, D. Messick, A. Tenbrunsel, & K. W.

- Benzoni (Eds.), *Environment, Ethics, and Behavior*. San Francisco: New Lexington Press, pp. 75–104.
- Tversky, A. & , Kahneman, D. (1974). Judgment under uncertainty: Heuristics and Biases. *Science*, 185, 1124-1131.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, 39, 1179–1189.
- Unger, P. (1996) *Living high and letting die: Our illusion of innocence*. New York: Oxford University Press
- Wedell, Douglas. (1991) 'Distinguishing among models of contextually induced preference reversals.' *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 17-767-778.
- Williams, J. M. G., Mathews, A., & MacLeod, C. (1996). The emotional Stroop task and psychopathology. *Psychological Bulletin*, 120, 3–24.
- Wilson, T. D., Houston, C. E., Etling, K. M., & Brekke, N. (1996). A new look at anchoring effects: Basic anchoring and its antecedents. *Journal of Experimental Psychology: General*, 125, 387-402.

## Figure Caption

*Figure 1. A hypothetical committee member had to choose to sponsor only one of the three plans presented. Both plans A and B are better on one dimension, but worse on the other, while plan C is slightly worse than B on both dimensions.*

Figure 1.

