7-2008

Harnessing Our Inner Angels and Demons: What We Have Learned About Want/Should Conflicts and How That Knowledge Can Help Us Reduce Short-Sighted Decision Making

Katherine L. Milkman
University of Pennsylvania

Todd Rogers

Max H. Bazerman

Follow this and additional works at: https://repository.upenn.edu/oid_papers

Part of the [Experimental Analysis of Behavior Commons](#), and the [Other Psychology Commons](#)

Recommended Citation

Milkman, K. L., Rogers, T., & Bazerman, M. H. (2008). Harnessing Our Inner Angels and Demons: What We Have Learned About Want/Should Conflicts and How That Knowledge Can Help Us Reduce Short-Sighted Decision Making. *Perspectives on Psychological Science*, 3 (4), 324-338. <http://dx.doi.org/10.1111/j.1745-6924.2008.00083.x>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/12
For more information, please contact repository@pobox.upenn.edu.

Harnessing Our Inner Angels and Demons: What We Have Learned About Want/Should Conflicts and How That Knowledge Can Help Us Reduce Short-Sighted Decision Making

Abstract

Although observers of human behavior have long been aware that people regularly struggle with internal conflict when deciding whether to behave responsibly or indulge in impulsivity, psychologists and economists did not begin to empirically investigate this type of want/should conflict until recently. In this article, we review and synthesize the latest research on want/should conflict, focusing our attention on the findings from an empirical literature on the topic that has blossomed over the last 15 years. We then turn to a discussion of how individuals and policy makers can use what has been learned about want/should conflict to help decision makers select far-sighted options.

Disciplines

Experimental Analysis of Behavior | Other Psychology

08-020

Harnessing Our Inner Angels and Demons:

**What We Have Learned About *Want/Should*
Conflicts and How That Knowledge Can
Help Us Reduce Short-Sighted Decision
Making**

**Katherine L. Milkman
Todd Rogers
Max H. Bazerman**

Copyright © 2007 by Katherine L. Milkman, Todd Rogers, and Max H. Bazerman

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Harnessing Our Inner Angels and Demons: What We Have Learned About *Want/Should* Conflicts and How That Knowledge Can Help Us Reduce Short-Sighted Decision Making

Katherine L. Milkman, Todd Rogers, and Max H. Bazerman
Harvard University

Although observers of human behavior have long been aware that people regularly struggle with internal conflict when deciding whether to behave responsibly or indulge in impulsivity, psychologists and economists did not begin to empirically investigate this type of *want/should* conflict until recently. In this paper, we review and synthesize the latest research on *want/should* conflict, focusing our attention on the findings from an empirical literature on the topic that has blossomed over the last 15 years. We then turn to a discussion of how individuals and policy makers can use what has been learned about *want/should* conflict to help decision makers select far-sighted options.

Author's Note: We are grateful for the insightful feedback we received from John Beshears on this paper. We also appreciate the advice we have received while conducting research in this area from Nava Ashraf and George Loewenstein. This research was conducted while Todd Rogers was supported by a National Science Foundation Graduate Research Fellowship. The authors would like to thank Harvard Business School for funding support.

Recently, one of the authors of this paper found herself struggling with the choice of what to order for dinner at her favorite neighborhood Italian restaurant. With great difficulty, she debated whether to indulge her craving for the restaurant's sinfully delicious cheesy pizza or to behave virtuously and order a light salad with grilled chicken and raspberry vinaigrette. The battle raged in her head until the waiter hovering over her cleared his throat to signal that it was time for her to place an order. In the end, she chose the pizza over the salad. The option she *wanted* more won out over the option she felt she *should* select. However, we all face this type of internal conflict frequently, and we do not always succumb to our immediate desires at the expense of our long-term interests. Despite the impulse this author feels most afternoons to sit in front of the television instead of going to the gym, she almost always does what is healthier in the long-run rather than what would be more enjoyable in the short-run and heads out the door in her exercise clothes. When her spouse wants to see a new action film and she would prefer to watch the latest romantic comedy, despite the tug she feels to behave selfishly, she compromises and does what will give her spouse the most pleasure about half of the time (although he might debate the precise ratio).

The types of internal conflicts described above are familiar to most people. In fact, the metaphor that individuals behave as if they have two selves – a *want* self fighting for whatever will bring more short-term pleasure and a *should* self representing an individual's long-term interests – is so common that its cinematic representation has become hackneyed. Films and television shows frequently depict internal conflict by showing a character with a whispering angel (the *should* self) perched near one ear and a fiery devil (the *want* self) at the other offering competing recommendations.

Evidence that storytellers have been aware of *want/should* conflict for millennia can be found in Homer's *The Odyssey*. In this epic tale, the hero, Ulysses, fears that like many sailors before him, he will be lured by the desires of his *want* self to his death at the hands of the Sirens – sea nymphs whose sweet songs lead sailors to wreck their ships on the rocks surrounding the nymphs' island. Before encountering the Sirens, Ulysses instructs his crew to plug their ears and bind him tightly to his vessel so he will be able to listen to the Sirens' song without the power to turn his ship. In this way, Ulysses' *should* self arranges for the desires of his *want* self to be kept in check so disaster will not befall him.

Scholars have theorized for decades about various types of internal conflict people face between doing what is best for their long-term interests and what will bring them the most immediate pleasure (Freud, 1923/1961; Strotz, 1956; Ainslie, 1975, 1992; Sen, 1977; Thaler and Shefrin, 1981; Schelling, 1984; Shefrin and Thaler, 1988; Loewenstein, 1996; Bazerman, Tenbrunsel and Wade-Benzoni, 1998), but the last fifteen years have witnessed the publication of the first surge of empirical studies examining the implications of what Bazerman, Tenbrunsel and Wade-Benzoni (1998) call *want/should* conflict. In this paper we review the literature on *want/should* conflict, focusing the majority of our attention on recent developments, and we discuss the implications of this stream of research for policy makers and individuals. We argue that the recent empirical literature on *want/should* conflict has identified ways in which individuals and policy makers can design decision contexts in order to facilitate the selection of *should* options. To set the stage for our discussion, we summarize research on the cognitive processes underlying *want/should* conflict and present a formal definition of relative *want* and

should options. We then review the latest empirical research on *want/should* conflict. After summarizing what is now understood about intrapersonal conflict and discussing when *want/should* conflict most often leads to decision errors, we focus our attention on an exploration of how this new knowledge can be applied. Specifically, we discuss how individuals seeking to increase their chances of making *should* choices and policy makers hoping to improve the odds that *should* policies will be adopted might be able to learn from studies of *want/should* conflict. We conclude with an assessment of opportunities for future research.

COGNITIVE PROCESSES UNDERLYING WANT/SHOULD CONFLICT

It is important to understand what cognitive processes underlie the conflict people feel when deciding whether to base choices on what they *want* to do or what they feel they *should* do because this knowledge will help us identify situational factors that are likely to affect the outcomes of intrapersonal conflicts, which, in turn, will help us determine how we can facilitate more *should* decisions. Bazerman, Tenbrunsel and Wade-Benzoni (1998) propose that individuals experience many decisions as if a *want* self and a *should* self coexist within them, and these selves are susceptible to conflicting preferences. The *want* self is driven by the desires people affectively feel in the moment when a decision will take effect, while the *should* self is guided by more deliberative feelings about what ought to be done given a person's long-term interests. A number of economic models have been proposed to explain intrapersonal conflict by assuming that people are actually controlled by multiple agents with different preferences pitted against one another (Thaler and Shefrin, 1981; Read, 2001; Fudenberg and Levine, 2006).

In contrast to these models of competing internal agents, other research has focused on how different parts of the brain are triggered by different contexts. For example, recent brain imaging research (McClure, Laibson, Loewenstein and Cohen, 2004) has confirmed that two systems in the brain are involved in decision making, one of which is preferentially activated by decisions associated with the receipt of immediate rewards (what Bazerman et al. (1998) would term the *want* self) and one of which is activated uniformly by decisions involving long-term and short-term rewards (what Bazerman et al. (1998) would refer to as the *should* self). Loewenstein (1996) has argued that intrapersonal conflicts stem from changes in the conditions under which decisions are made. He proposes that visceral factors, such as emotions and psychological cravings like hunger, often overwhelm people at the moment of a decision, and that these visceral factors are the source of observed differences between the *should* preferences people often articulate when in a deliberative state (“I would like to lose weight”) and the *want* preferences they often exhibit when making choices in a more visceral state (“I’ll take that cheeseburger with extra fries”).

Recent research on Construal Level Theory (CLT) suggests that different ways in which people construe the world may underlie *want/should* conflict. CLT posits that the temporal proximity of an event (and perhaps other factors as well) systematically affects how it is construed (Trope & Liberman, 2003; Liberman, Sagristano, Trope, 2002). Events in the distant future are construed at a high level, which means they are associated with schematic, abstract, and purpose-focused qualities, while events in the near future are construed at a low level and are associated with concrete, specific, detail-focused qualities. To give an example, a high level construal of a salad would focus on the

salad's healthfulness and its likelihood of increasing longevity, while a low level construal of a salad would focus on its taste and its likelihood of leaving you hungry. It has been suggested that the long-term focus of the *should* self is triggered by high level construal, while the short-term focus of the *want* self is triggered by low level construal (Rogers and Bazerman, in press; Trope and Liberman, 2000). Rogers and Bazerman (in press) and Kivetz and Tyler (2007) have found evidence that a preference for *should* options is associated with high level construal. Similarly, Fujita, Trope, Liberman and Levin-Sagi (2006) have demonstrated that the preferences of the *should* self are more likely to win out over those of the *want* self when options are construed at a high level rather than at a low level. Together, these findings suggest that the different ways in which our minds operate when options are construed at high versus low levels may underlie *want/should* conflict. In other words, the *want* self may represent the outcomes favored by a low level construal of the world, while the *should* self may represent the outcomes favored by a high level construal of the world.

A FORMAL DEFINITION OF RELATIVE WANT AND SHOULD OPTIONS

In order to consider the applications of research on *want/should* conflict, we must also have an understanding of what types of options lead individuals to experience such conflict when making decisions. We propose a definition for options that are associated with the preferences of the *want* and *should* selves. When compared to one another, some options are preferred by the *should* self (e.g., salads, documentary films, trips to the gym, etc.), while others are preferred by the *want* self (e.g., ice cream cones, action films, skipping the gym, etc.). Given two options, we define one option as having relatively

more *want* and fewer *should* characteristics than a second option if and only if the following two conditions hold:

- 1) The instantaneous utility obtained from the *want* option is greater than the instantaneous utility obtained from the *should* option.
- 2) The sum of the utility (discounted at a standard rate, $\delta = 1 - \epsilon$) that will be derived from the *want* option in all future periods is less than the sum of the utility that will be derived from the *should* option in all future periods.

It is important to note that this definition does not classify whether a *want* or *should* option is optimal. The optimal choice between *want* and *should* options requires summing the short-run and long-run utility that would be gained from each option and selecting whichever provides more discounted net utility. Although *should* options have more long-run benefits than *want* options, in many cases the short-run benefits of a *want* option may be significant enough to outweigh the long-run benefits of a *should* option.

To put our definition of *want* and *should* options in context, consider some examples. First, think of foods. According to our definition of things associated with the *want* and *should* selves, pizza is a *want* good relative to salad, which is a *should* good. Pizza gives most consumers more instant gratification than salad while it is being consumed (yum, grease!). However, the future utility gained from eating the salad is higher because consumers will likely be much healthier, slimmer and happier if they manage to resist the temptation to eat pizza each night instead of salad. Now, think of films. An action-packed blockbuster is likely to be far more of a *want* film and less of a *should* film than a History Channel documentary. Tying this back to our definition of relative *want* and *should* options, a blockbuster is typically more exciting to watch than a

history documentary. However, a history documentary is more likely than a blockbuster film to contain information that will serve you well in the future, enriching your conversations, helping you impress your boss, and potentially even leading you to make better decisions. Finally, think of potential ballot issues. Voting for a tax on gasoline is a *should* behavior (it will help reduce emissions, thereby improving air quality and the environment in the future) but voting against it is something that many citizens *want* to do (“I love filling up my gas guzzling SUV on the cheap!”). Voting for the *should* policy is likely to bring you more happiness in the future because it will increase your chances of breathing cleaner air and slowing global warming. It will also allow you to feel pride whenever you recall your vote because you will know that you contributed to a good cause. However, voting for the *want* policy will increase the odds that you will be able to enjoy your SUV in the short-run.

RECENT EMPIRICAL RESEARCH ON WANT/SHOULD CONFLICT

Recent empirical research on *want/should* conflict has primarily focused on identifying situational factors that cause individuals to reverse their preferences over *want* versus *should* options. In this section, we review the literature on different conditions known to induce *want/should* preference reversals. The knowledge gained from the empirical studies we discuss in this section will be the cornerstone of our upcoming discussion of how individuals and policy makers may be able to design decision making contexts that facilitate more *should* choices and why it makes sense for them to do so.

Intertemporal Choice

The majority of past research on *want/should* conflict has been focused on the domain of intertemporal choice. To summarize, it has been hypothesized and confirmed in a variety

of contexts that when making choices that will take effect in the future, people are considerably more likely to favor *should* options over *want* options than when making decisions that will take effect immediately. This observation holds when decisions are made in the domains of money (Thaler, 1981; Ainslie and Haendel, 1983; King and Logue, 1987; Kirby and Herrnstein, 1995; Kirby and Marakovic, 1996; Kirby, 1997; McClure, Laibson, Loewenstein and Cohen, 2004;), exercise (Della Vigna and Malmendier, 2006), and film choice (Read, Loewenstein and Kalyanaraman, 1999; Milkman, Rogers, and Bazerman, 2007), among others. For example, to most people, the idea of going to the gym tomorrow is much more palatable than the idea of going to the gym this minute. Similarly, the idea of starting a diet next week sounds much more appealing than the idea of starting a diet today. In this sub-section, we will briefly discuss how *want/should* conflict has been modeled in the context of intertemporal choice before reviewing the empirical research in this area.

Researchers began modeling dynamically inconsistent preferences as early as 1956 with non-standard time discount functions (Strotz, 1956). To fit the observation that people have great difficulty passing up a large reward in the present for a larger reward tomorrow, while they have considerably less difficulty passing up the same large reward tomorrow for that larger reward in two days, economists have modeled individuals as possessing a discount rate for utility that is extremely high in the short-run but relatively low in the long-run (Strotz, 1956; Loewenstein and Prelec, 1992; Ainslie, 1992; Laibson, 1996). A leading model of impulsiveness in the economic literature is Laibson's (1996) quasi-hyperbolic time discount model, which models the extreme short-run drop in valuation that has been observed in people's time preferences by adding a

discount factor of β , which is much less than 1, to all but the first time period of a traditional discrete-time exponential discount model. Specifically, Laibson's model assumes individuals place no discount on immediate utility but discount all future utilities by β in addition to the traditional, exponential discount rate, δ , which is much greater than β . For partial reviews of the economic literature on dynamic inconsistency, see Loewenstein and Thaler (1989), Ainslie (1992), O'Donoghue and Rabin (1999), and Frederick, Loewenstein, and O'Donoghue (2001).

Combining Laibson's model with the definition of *should* and *want* options discussed in Section II produces the prediction that people will often reverse their preferences over relative *want* and *should* items like an action film and a documentary when the delay between selection and consumption switches between short and long. For example, consider two movies that you could rent for tomorrow night: Steven Spielberg's *want* action blockbuster "Jurassic Park" or Steven Spielberg's *should* Academy Award winning film about the Holocaust "Schindler's List". If "Schindler's List" is a *should* film relative to "Jurassic Park" and if it is also your optimal film choice because it provides more net utility, when choosing which film to watch tomorrow, you will rent "Schindler's List." This is because the utility flows from both movies will be uniformly downweighted by β when you consider which film to watch tomorrow, since all utility flows from each movie will be achieved in the future.

However, when tomorrow arrives and you are faced with the decision of which film to watch today, you may choose to watch "Jurassic Park" instead of "Schindler's List" as a result of the high discount (β) you place on future utility flows and the fact that

the immediate rewards from watching “Jurassic Park” exceed the immediate rewards from watching “Schindler’s List (by the definition of relative *want* and *should* options).

Now that we have discussed the theoretical work on dynamic inconsistency, we move to a discussion of the empirical research that has been conducted to test the prevailing theory. Before people were using terms like *should* and *want* to describe multiple-selves conflict, a number of laboratory studies were conducted to examine impulsiveness in the domain of money. In an early study of dynamic inconsistency, Ainslie and Haendel (1983) conducted a laboratory experiment in which subjects were divided into two conditions. In one condition they were given a hypothetical choice between receiving \$50 immediately (the *want* option) or \$100 in six months (the *should* option), and in the second condition they were given a hypothetical choice between receiving \$50 in 12 months or \$100 in 18 months. Standard economic theory suggests that the choices made by subjects in the two conditions should be indistinguishable. However, Ainslie and Haendel found that most subjects chose the smaller reward in the first condition, when it would be available immediately, and the larger reward in the second condition, when all rewards would be delayed. A plethora of other field and laboratory studies have been conducted since using real monetary payoffs as well as hypothetical payoffs, and all have confirmed that people exhibit an abnormally high discount rate between immediate and delayed payoffs (Kirby and Herrnstein, 1996; Kirby and Marakovic, 1996; Kirby, 1997; McClure, Laibson, Loewenstein, and Cohen, 2004; Angeletos, Laibson, Repetto, Tobacman, and Weinberg, 2001). However, many of these studies have been subject to a number of criticisms. First, if subjects believe there is some risk that an experimenter will not follow through on payoffs offered in the future,

measured short-term discount rates may be biased upward. Second, if utility flows result from the things that money can buy rather than directly from money, studies conducted with monetary payoffs should not induce as much short-term utility as studies conducted with real goods as payoffs, so measured short-term discount rates might be biased downward.

Empirical studies of the multiple-selves phenomenon focusing on choices between relative *should* and *want* options in realms besides money followed research on monetary discount rates. Such studies allowed academics to confirm that previously observed anomalies of intertemporal choice were not caused by anything special about the way people value money relative to other goods. One non-monetary domain where multiple studies of *want/should* conflict have been conducted is that of film rentals. Read, Loewenstein and Kalyanaraman (1999) conducted an experiment in which subjects were given an opportunity to rent three films for three future dates from a list of 14 highbrow (*should*) and 10 lowbrow (*want*) movies. Study participants were randomly assigned to one of two experimental conditions. In the first condition, they were asked to simultaneously choose three films from the list of 24 that they would like to rent in the future. In the second condition they were allowed to choose the three films sequentially on the days when their rentals would take place. The authors found that subjects in the sequential choice condition, who were able to select films right before consumption rather than well in advance of consumption, rented significantly more lowbrow movies and fewer highbrow movies than subjects in the simultaneous choice condition. This finding is consistent with the prediction of the theories presented above about intertemporal choice and *should/want* conflict. When choosing for now, subjects in this

study were more susceptible to the whims of their *want* selves, but when choosing for later, subjects were more likely to rent *should* movies.

Milkman et al. (2007) followed up on this research by Read et al. (1999) with a field study of dynamic inconsistency in the domain of online DVD rentals. In order to study intertemporal choice and *want/should* conflict among online DVD rental customers, the authors obtained four months of data on individual customers' rental and return decisions from an Australian online DVD rental company. Milkman et al. (2007) began by classifying the thousands of films offered for rent by this company along a continuous scale from extreme *should* films to extreme *want* films. The authors used two different film classification methods. The first relied on a survey in which respondents were asked about the extent to which different film genres were composed of *should* and *want* movies. Films were then assigned *should minus want* ratings according to the difference between the average *should* and *want* scores respondents gave the genres in which they were classified. The second classification method relied on a survey in which respondents were asked to rate the extent to which 500 different films were *should* and *want* movies. Films' average ratings were then used to estimate a regression equation for predicting a film's *should minus want* score based on its quantifiable characteristics (e.g., genre, year of release, average user rating, etc.), and this equation was used to give all films in the authors' database *should minus want* scores. After creating these scoring systems, Milkman et al. (2007) tested and confirmed the hypothesis that people are more likely to rent DVDs in one order and return them in the reverse order when *should* DVDs (e.g., documentaries) are rented before *want* DVDs (e.g., action films). In addition, the authors predicted and found that *should* DVDs are held longer by customers than *want*

DVDs. Although these field results were consistent with previous theory and laboratory research on intertemporal choice, they are notable because they confirmed that the effects of *want/should* conflict are large enough to significantly affect real world decision making. Perhaps more noteworthy still, Milkman et al. (2007) determined that a movie's *want/should* score is as strong a predictor of how long that movie will be held by an online DVD rental customer as all of the film's other quantifiable characteristics (e.g., genre, year of release, average user rating, etc.) combined.

Another domain in which recent research has been conducted on *want/should* conflict and intertemporal choice is the domain of groceries. Rogers, Milkman and Bazerman (2007) examined dynamic inconsistency in this context by obtaining a year of data detailing the orders placed by the customers of a major North American online grocer. The authors examined how the mix of *should* and *want* goods purchased by the same shoppers differed depending on how far in advance of delivery an order was completed. Goods were assigned *should* and *want* scores based on the average score survey respondents assigned to groceries in their category. In addition to finding that customers spent more when ordering for more immediate delivery (spending is a typical *want* behavior, while saving is a *should* behavior), Rogers et al. (2007) determined that the percentage of extreme *should* groceries in a customer's basket generally increases the further in advance of delivery an order is completed, while the percentage of extreme *want* groceries in a customer's basket generally decreases the further in advance of delivery an order is completed.

Related research has been conducted by Oster and Scott Morton (2005) on the newsstand and subscription prices for leisure magazines (*want* magazines), or magazines

that are fun to read “now,” and investment magazines (*should* magazines), or magazines that provide benefits in the future. These authors recognized that if the type of internal *want/should* conflict discussed in this paper has a significant effect on decision making, magazine prices should reflect the fact that people will rarely plan ahead when it comes to the consumption of leisure magazines (e.g., *People*) but will regularly plan ahead when it comes to the consumption of investment magazines (e.g., *The Economist*). In an efficient market, such behavior should lead the ratio of a magazine’s newsstand price to its subscription price to be considerably larger for a leisure magazine than for an investment magazine. Oster and Scott Morton find that this is the case in their study of approximately 300 American magazines.

Della Vigna and Malmendier (2006) conducted a study that examined gym attendance (a *should* behavior) and found evidence that people exhibit dynamic inconsistency in the domain of exercise. The authors analyzed data on the gym contracts purchased by thousands of gym customers as well as the gym attendance records of those customers. They found that people regularly paid a high fee for the right to visit their gym an unlimited number of times when they could have saved money by paying flat per-visit fees instead. These findings are consistent with a model of consumers who make *should* choices when thinking about the future (i.e., they pay to go to the gym frequently in the future) but systematically reverse their preferences and opt for *want* options when the future eventually becomes the present (i.e., they do not actually go to the gym frequently).

Joint versus Separate Decision Making

Want/should conflict is also evident in reversals of preference that have been observed in joint versus separate decision making. Academics have hypothesized and confirmed in a variety of settings that individuals are more likely to favor *want* options over *should* options when evaluating different possibilities one at a time rather than simultaneously (see Bazerman, et al. (1998), Hsee, Loewenstein, Blount and Bazerman (1999), and Bazerman, Moore, Tenbrunsel, Wade-Benzoni and Blount (1999) for reviews of this literature). For example, the idea of donating to a charity that protects baby polar bears may generate more enthusiasm than the idea of donating to a charity that supports research on malaria when each donation opportunity is considered separately. However, when given a choice between saving baby polar bears and reducing the numbers of human deaths from malaria, most people feel more obliged to save people than polar bears (Kahneman and Ritov, 1994). Bazerman et al. (1998) propose that *should* options are more likely to win out over *want* options in joint evaluation than in separate evaluation because direct comparison forces an individual to rationally weigh the costs and benefits of her choices. In separate evaluation, the short-term, visceral desires of the *want* self are more likely to overwhelm a decision maker than in joint evaluation because in separate evaluation there is no explicit tradeoff to cue the rational weighing of costs and benefits.

In the first explicit study of joint versus separate preference reversals, Bazerman, Loewenstein and White (1992) evaluated the likelihood that participants in a fictional dispute with a neighbor would accept two different settlement options depending on whether those options were presented jointly or separately. In one settlement option,

both disputants would receive the same amount of money. In the second settlement option, both disputants would receive more money than they would have under the terms of the first settlement, but the participant's neighbor would receive a larger payout than the participant. The second option was considered the *should* option because it yielded a higher payoff to the participant and would thus provide more long-term utility, while the first option was considered the *want* option because the participant would not have to experience the short-run, visceral displeasure associated with receiving less than her neighbor. In the separate evaluation condition, the two settlements were presented to participants sequentially and participants were asked to rate the appeal of each option on a continuous scale. Under these circumstances, the average participant rated the option in which she and her neighbor would receive equal payoffs (the *want* option) more favorably than the option in which she would receive less than her neighbor but more money overall (the *should* option). However, in the joint evaluation condition, which pitted the two settlement options directly against one another, the average participant preferred the settlement in which she would receive a higher payoff (the *should* option).

A subsequent study by Bazerman, Schroth, Shah, Diekmann, and Tenbrunsel (1994) replicated the basic joint versus separate preference reversal result described above. In this study, participants were asked to consider accepting matched hypothetical job offers. In a typical comparison, one job, the *want* job, would pay the participant and other new hires the same yearly salary. The other job, the *should* job, would pay the participant a higher yearly wage than the first job but less than it would pay other new hires. When evaluated jointly, the vast majority of participants expressed a preference for the *should* job over the *want* job, presumably due to its superior paycheck. When

evaluated separately, however, a higher percentage of participants said they would accept the *want* job, where they would be paid as much as other new hires, than the *should* job, where they would earn more overall but less than other new hires.

Shortly after Bazerman et al.'s (1992) initial study of joint versus separate preference reversals was published, Kahneman and Ritov (1994) conducted a similar study of individuals' willingness to give to different types of charities when donation opportunities were presented jointly versus separately. Kahneman and Ritov (1994) identified a pattern of joint versus separate preference reversals similar to those observed by Bazerman et al. (1992) and Bazerman et al. (1994). In their study of willingness to give, subjects were presented with opportunities to donate to wildlife charities (e.g., a charity that protected kangaroos) or charities supporting human health or safety (e.g., a charity that provided free skin cancer check-ups for farm workers). The wildlife charities fit our definition of *want* options because they were more immediately, emotionally attractive, while the charities promoting human health and safety were more pragmatically important in the long-run by most measures and thus fit our definition of *should* options. In joint evaluation, the *should* options were systematically preferred over the *want* options by study participants, but in separate evaluation, this pattern reversed itself as in the Bazerman et al. (1992) and Bazerman et al. (1994) studies.

In another early study of this phenomenon, Irwin, Slovic, Lichtenstein and McCelland (1993) compared people's willingness to pay for public goods (e.g., air quality improvements) to their willingness to pay for commodities for their personal use (e.g., a new VCR) in joint versus separate evaluation. Consistent with the findings of others, Irwin et al. (1993) determined that in joint comparison, public goods (or *should*

goods) elicited higher willingness to pay from study subjects than commodities for personal use (or *want* goods), but this pattern reversed itself in separate evaluation. Lowenthal also conducted a study of joint versus separate preference reversals in 1993 in which he examined people's willingness to vote for different candidates running for political office. In Lowenthal's study, a candidate boasting the ability to bring 1,000 new jobs to his district and a clean criminal record (the *want* candidate) was pitted against a candidate who had been convicted of a misdemeanor but could produce 5,000 new jobs for his constituents (the *should* candidate). When evaluated jointly in an "election", the *should* candidate earned more votes, but when rated separately on likeability, the *want* candidate was more popular. These studies offer yet more demonstrations of joint versus separate *want/should* preference reversals.

Hsee (1995, 1996, 1998) examined whether joint versus separate preference reversals might be driven by what he termed the "evaluability hypothesis". Hsee's evaluability hypothesis proposes that when people evaluate options with multiple attributes, in joint evaluation it will be possible for them to compare and weigh each attribute appropriately but in separate evaluation the attributes with clearer standards for evaluation (e.g., GPA, SAT scores, etc.) will be overweighed relative to attributes with less clear standards for evaluation (e.g., a score on an unknown test) resulting in joint versus separate preference reversals. Hsee conducted a number of studies to test his evaluability hypothesis. In one study, he demonstrated that people are more likely to prefer a dictionary, A, with a torn cover and twice as many entries as a second dictionary, B, with a immaculate cover when the two are compared jointly but that this preference will reverse itself when the dictionaries are presented separately (Hsee, 1996). Similarly,

Hsee (1998) found that in joint evaluation people would prefer a 40 piece set of china with 9 broken pieces to an undamaged set with just 24 pieces, but in separate evaluation the unmarred china set is more appealing. Hsee conducted similar studies of ice cream cups, CD changers and job applicants, among other things. However, although the pattern of preference reversals Hsee identified is consistent with his evaluability hypothesis, it is also consistent with the idea that *want* items are more often preferred over *should* items in separate evaluation than in joint evaluation. A damaged dictionary or china set creates a negative visceral reaction, so an undamaged good with fewer of the qualities we know we ought to care about is a *want* option relative to the damaged good, which is the *should* option.

Other Moderators of Should/Want Conflict

Although the majority of empirical research on *want/should* conflict to date has focused on intertemporal choice and joint versus separate decision making, these are not the only contexts in which differences in a decision maker's situation can systematically sway whether she is more likely to act on the desires of her *want* self or the pragmatism of her *should* self. Recent studies have examined forces that affect the outcomes of choices between *should* and *want* options in contexts besides those previously discussed, such as extreme cognitive load (Shiv and Fedorkin, 1999) and whether an individual views a choice as isolated or believes she will make similar choices in the future (Khan and Dhar, 2007).

Shiv and Fedorikhin (1999) conducted a study to examine the impact of cognitive load on people's preferences for foods associated with intense positive emotions (*want* foods) versus foods associated with less positive affect but more favorable cognitions

(*should* foods). Participants in their study were randomly assigned to a high or low cognitive load condition. In the high cognitive load condition, subjects were instructed to memorize a seven digit number, and in the low cognitive load condition they were asked to memorize a two digit number. Study subjects were then given a choice between two snacks: a slice of chocolate cake (a *want* snack) and a cup of fruit salad (a *should* snack). Shiv and Fedorikhin hypothesized that subjects in the high cognitive load condition would be more likely to select the chocolate cake than subjects in the low cognitive load condition because fewer of their intellectual resources would be available to help them resist temptation. The authors found that significantly more subjects selected cake over fruit salad in the high cognitive load condition (63%) than in the low cognitive load condition (42%).

Khan and Dhar (2007) evaluate the difference in choices people make between *want* and *should* goods when making a one-shot decision versus the first in a series of similar decisions. The authors conducted three studies in which subjects were given the opportunity to choose between an array of virtue (*should*) and vice (*want*) goods. Subjects in all three studies were randomly assigned to one of two conditions and given the opportunity to select either one magazine, one movie or one snack for immediate enjoyment. In the first treatment condition subjects were told that their choice would be the first of a series of similar choices, the rest of which would be made in the future. In the second condition subjects were told they would be making an isolated choice. Subjects selected a significantly higher proportion of virtuous movies, magazines and snacks when they believed their choice was made in isolation and not the first of a series of similar choices. Khan and Dhar hypothesize that subjects gave into their visceral

desires more frequently when they believed they were making the first in a series of similar choices because they were able to offset the guilt associated with their impulsive behavior by anticipating that their future selves would make more virtuous selections. Consistent with this hypothesis, the authors found that when subjects in the repeated choice condition of their snack experiment believed they would be eating a virtuous snack the following week, the percentage choosing a virtuous snack for immediate consumption decreased.

MAKING OPTIMAL CHOICES

Before turning to a discussion of ways in which the findings from the literature on *want/should* conflict can be applied to help individuals and policy makers, it is important to address the question of when *want/should* conflict leads to decision making errors. As discussed previously, given a choice between a *want* option and a *should* option, a perfectly rational decision maker would select whichever option will yield more exponentially discounted net utility. When a small change to the context in which a decision is being made systematically leads to predictable reversals in people's preferences for *should* versus *want* options, it is safe to assume that one context is inducing a decision error since the same choice typically remains optimal across the contexts. In this section we argue that decision errors that involve favoring *want* options when *should* options are optimal occur more frequently and are more detrimental than errors that involve favoring *should* options when *want* options are optimal, although there is evidence that both types of mistakes occur. As a result, we propose that individuals and policy makers should focus their attention on developing strategies that will help

them and their constituents, respectively, increase their odds of choosing *should* options over *want* options.

By definition, when the difference between the net future utility of a *should* option and a *want* option is greater than the difference between the immediate utility those *want* and *should* options have to offer, the *should* option is the optimal choice. We argue that the majority of decision making errors resulting from *want/should* conflict arise when a *should* option is optimal but a *want* option is selected. Angeletos et al. (2001) provide evidence that this is the case with a study that estimates Americans' average discount rates based on retirement wealth data. According to their calibration, between the present and one year and the future, Americans discount money at a rate of .53, meaning that on average, people view \$1.00 in one year and \$0.53 today as equivalent. It is difficult to argue that such a steep discount rate is optimal in an environment where yearly inflation has averaged 2.7% over the last decade.¹ Indeed, Angeletos et al. report that 55% of respondents in a 1997 survey said they were behind on their savings goals, while only 6% reported being ahead. The fact that America's obesity problem is far larger than America's problem with anorexia (66% of Americans are overweight or obese)² also suggests that people more frequently sub-optimally overweight the desires of their *want* selves relative to those of their *should* selves than visa versa. The preponderance of attention given to self-control problems in the psychology and economics literatures compared with under-indulgence problems (see Appendix A for a comparison) also suggests that self-control problems are more common than under-indulgence problems. Evidence from a survey we conducted with a national

¹ CPI Inflation Calculator. <http://data.bls.gov/cgi-bin/cpicalc.pl> accessed on September 13, 2007.

² Medline Plus: Weight Control. <http://www.nlm.nih.gov/medlineplus/weightcontrol.html> accessed on September 13, 2007.

sample of paid participants also supports this view. 66% of respondents reported making errors that involved favoring their *want* selves over their *should* selves when it was suboptimal to do so more frequently than mistakenly favoring their *should* selves over their *want* selves.

In addition to being the more common error, the mistake of sub-optimally undervaluing *should* options relative to *want* options is an error that can often lead to far more detrimental consequences than the opposite mistake. For instance, consider the following potential outcomes of self-control problems: under-saving for retirement, becoming a drug addict, becoming an alcoholic, failing to earn a high school diploma, paying late fees. Now consider the outcomes of under-indulgence problems in the same domains: over-saving for retirement, missing out on the thrills associated with drug and alcohol use, receiving “excess” education, paying bills so early that some interest is foregone. These examples suggest that errors in judgment that stem from overweighting *want* options relative to *should* options often have far more severe consequences than errors caused by overweighting *should* options relative to *want* options, which is another reason why we believe policy makers and individuals interested in reducing the negative effects of sub-optimal decision making should focus on finding strategies to increase the odds that people will make *should* choices.

Thus far, we have focused our attention on the most common mistake associated with *want/should* conflict – the error that leads people to overweight the desires of their *want* selves. However, as discussed above, sometimes people make the opposite error. That is to say, they underweight their short-term, *want* desires and do what will provide more future utility, even when doing so is sub-optimal. By definition, it is optimal to

choose a *want* option when the immediate utility from that option exceeds the immediate utility from a *should* option by more than the *should* option's net future utility exceeds that of the *want* option. However, sometimes people make *should* choices when a *want* choice would make them better off overall, like dieting when they are already at or below a healthy weight or over-saving for retirement instead of buying a nicer house. Wilson, Schooler, Hodges, Klaaren, and LaFleur (1993) found that a simple change in the instructions given to participants both increased the odds that they would make a *should* choice and decreased their future satisfaction with their choice, suggesting that the same manipulation led people to overweight the value of a *should* option and made them worse off. Wilson and his co-authors gave college students the choice between two posters: a print of an impressionist painting (the *should* choice) and a humorous poster (the *want* choice). Half of the participants were first asked to think about why they liked or disliked each poster, while the other half were not asked to introspect in this way. The experimenters found that those who introspected were more likely than others to choose the *should* poster, and they were also less satisfied in the future with the poster they selected. This study's findings illustrate that under certain conditions people may mistakenly choose *should* options when choosing a *want* option would be optimal.

Kivetz and Simonson (2002a) have demonstrated that people are aware that they sometimes choose *should* options when they would be better off choosing *want* options. In a series of studies, they found that a large proportion of people are willing to pre-commit to future *want* options that are objectively inferior to available *should* options. For example, Kivetz and Simonson found that 28% of subjects waiting in an airport would choose a bottle of wine valued at \$50 (the *want* option) as a prize over \$55 in cash

(the *should* option). They argue that this is evidence that some people are willing to ‘pre-commit to indulgence’ because they know they will sub-optimally under-indulge otherwise.

While there is evidence that people occasionally make the sub-optimal decision to under-indulge in *want* options, it seems that people are considerably more likely to over-indulge in *want* options. In addition, we have argued that over-indulgence in *want* options is often associated with larger penalties than under-indulgence in *want* options. For these reasons we focus our discussion of the implications of research on *want/should* conflict on how decision makers and policy makers can increase the rate at which they and their constituents, respectively, select *should* options when experiencing internal conflict.

APPLICATIONS OF RESEARCH ON WANT/SHOULD CONFLICT

Understanding the conditions in which people select *should* options over *want* options and vice versa as well as when these choices are likely to be sub-optimal can help us think about strategies for solving important problems that result from flawed decisions. With a better understanding of the conditions that affect whether people lean towards *want* or *should* options, individual decision makers may be able to help themselves make more choices that are optimal but contradict what they *want* to do, and policy makers may also be able to help facilitate more optimal *should* decisions. In this section we describe recent empirical research demonstrating how people can and often do take steps to increase the likelihood that they will follow the advice of their *should* selves. This work suggests that people are eager to find ways to better manage their intrapersonal conflicts and demonstrates how some individuals have effectively managed to help

themselves make more *should* choices. We also discuss empirical work with implications for how policy makers may be able to design decision contexts that facilitate the selection of *should* options.

It is important to note that in this section we do not discuss ways in which incentives can be manipulated to change the likelihood that people will select *should* options. Many policies have been proposed and implemented, which favor *should* options by changing incentives. For example, placing “sin” taxes on cigarettes and alcohol and outlawing heroine are ways of increasing the likelihood that people will make *should* choices by raising the costs associated *want* choices. While such policies may have many benefits (see Gruber and Koszegi, 2004 for example), we focus our attention on discussing ways that the *should* self can be given a leg up without directly manipulating the costs associated with *want* options. Unlike incentive manipulations, the methods we discuss for increasing the likelihood that people will make *should* decisions do not require policy makers to decide what choices are in all people’s long-run best interest. Instead, we examine the conditions that will help each individual do what is in her own long-term best interest. The choices of individuals who do not face *want/should* conflict in a given domain will not be altered by the methods we propose for increasing the incidence of *should* decision making, whereas everyone who drinks alcohol would be affected by an increase in its price, regardless of whether drinking less alcohol represents a *should* choice for them.

Commitment Devices

Some researchers interested in intertemporal *want/should* conflict have begun to explore the question of whether people are willing to take measures to prevent themselves from

acting on *want* impulses instead of doing what they feel they *should*. Preventative measures taken to restrain the *want* self are frequently referred to as “commitment devices,” which typically require an individual to commit in the present to making a *should* choice rather than a *want* choice in the future. Some examples of commitment devices that many people are familiar with include piggy banks, which people’s *should* selves use to prevent their *want* selves from dipping into their savings; diet treatment centers, which people’s *should* selves enroll in to prevent their *want* selves from overeating; and pills like Antabuse, which people’s *should* selves take so their *want* selves will face the prospect of a painful illness if they consume an alcoholic beverage.

In an early study of commitment devices, Wertenbroch (1998) used supermarket scanner data to conduct a paired comparison of the quantity discounts applied to a matched sample of 30 virtue (*should*) and 30 vice (*want*) grocery products. He found that, on average, vice foods are subject to steeper quantity discounts than virtue foods, and the demand for virtue goods is less price sensitive than the demand for vice goods. These two findings suggest that consumers are aware of their impulsivity and that their *should* selves take steps at the time of purchase (which is in advance of consumption) to prevent their *want* selves from having the opportunity to binge in the future. In other words, people are willing to pay more to buy smaller packages of vice foods in order to avoid having too many such foods around to tempt their impulsive, *want* selves when they sit down to eat. Buying small packages of *want* foods helps people commit to eating less junk food than they might otherwise, and people are willing to pay a price for this commitment device.

One interpretation of the study of gym contracts and gym attendance conducted by Della Vigna and Malmendier in 2006, which we discussed previously, is that it provides evidence that people place a positive value on commitment devices. Della Vigna and Malmendier (2006) found that people often pay for gym memberships that entitle them to unlimited gym visits despite the fact that they would save money if they simply paid per-visit usage fees. It is very possible that Della Vigna and Malmendier observed this pattern of behavior because when thinking about the future, people naively over-predict how frequently they will make the *should* decision to go to the gym. An alternative explanation, however, is that people are sophisticated about their dynamic inconsistency and sign up for unlimited gym visit contracts in order to increase the likelihood that their future selves will go to the gym. In other words, reducing the marginal cost of a gym visit to zero may serve as a commitment device, which increases gym attendance and thus has a positive value.

In a more controlled study of commitment devices, Ariely and Wertenbroch (2002) examined whether college students would opt to assign themselves deadlines for three papers they were required to hand in by the end of an academic semester. Students who assigned themselves deadlines were committing to complete one or more papers before the last minute (a *should* behavior) rather than procrastinating for as long as possible (a *want* behavior). The authors found that when students were given the option to assign themselves deadlines, 73% elected to impose deadlines on themselves that would require them to turn in one or more of their papers before the last day of class. This indicated that many students were aware of their self-control problems and placed a

positive value on a commitment device that would prevent them from doing what they *wanted* to do rather than what they felt they *should*.

In 2006, Ashraf, Karlan and Yin conducted a study of people's willingness to take up commitment devices in the domain of savings. The authors partnered with a bank in the Philippines, which contacted a group of its former clients to offer them a choice between two savings products. The first savings product was a standard savings account, but the second was a commitment savings account from which money could not be withdrawn until the depositor reached a self-selected, pre-determined date or savings goal. Money in both accounts earned the same rate of interest, and thus people only had an incentive to place their money in the commitment savings account if they wanted to prevent their *want* selves from impulsively withdrawing and spending funds set aside for the future by their *should* selves. The study by Ashraf et al. (2006) had two important findings. The first was that a significant percentage of people (28%) were willing to give up the freedom to withdraw money from a savings account at their discretion without accepting any compensation for this sacrifice, indicating that they placed a positive value on the commitment device. The second finding was that individuals who were given the opportunity to use a commitment savings product saved 81% more over the course of a year than individuals who were not offered this product, indicating that people are able to save more when their *want* selves are kept in closer check. These findings suggest that those who understand the implications of dynamic inconsistency may be able to use their knowledge to design policies with the potential to benefit society by increasing savings rates, among other things.

In a similar study, Thaler and Benartzi (2004) presented randomly selected employees at several different companies with the opportunity to enroll in a savings plan called “Save More Tomorrow” or SMarT. This savings plan invited employees to pre-commit to automatically placing 50% of the proceeds of their future pay raises in an investment savings account. Thaler and Benartzi found that people who were offered the opportunity to participate in this plan saved dramatically more than those who were not invited to participate. SMarT’s design capitalized on several different psychological and economic principles to increase its appeal and effectiveness. For example, the plan offered participants the opportunity to purchase expert-recommended investments (reducing the cognitive costs of deciding what to invest in), invited employees to save only from future pay increases (thus capitalizing on dynamic inconsistency), and automatically deposited employees’ money in savings (taking advantage of the power of defaults). Because Thaler and Benartzi did not isolate individual features of SMarT, their research does not reveal which specific characteristics of the plan increase savings rates. However, Thaler and Benartzi’s work suggests that offering people the opportunity to pre-commit to savings products may have the potential to increase savings rates.

The research we have discussed thus far in this section provides evidence that many people are willing to pay a price to take up commitment devices in order to help their deliberative *should* selves overcome the impulsive desires of their *want* selves. These findings suggest that policy makers may be able to improve people’s welfare by making more commitment devices available, and that people may be able to increase their own happiness by seeking out and using commitment devices. An example of a commitment device that policy makers might want to consider making available to

consumers, which could prove enormously beneficial, is a program that would allow people to sign up to be prevented from buying cigarettes. Unlike many other types of commitment devices that people might be willing to pay for, it would be difficult for any non-governmental group to offer consumers the opportunity to sign up for such a program because of the coordination of innumerable small businesses that would be necessary to make such a “do not sell” list effective. However, there may be many people who would value the opportunity to commit to being prevented from buying cigarettes in the future. Although some types of commitment devices such as Christmas Clubs and diet treatment centers have arisen naturally in response to demand, there are other types of commitment devices that it may be difficult or impossible for private companies to provide due to coordination problems. Governments may be able to help people struggling with *want/should* conflict by creating commitment devices when coordination problems prevent private companies from offering such products. In addition, governments may be able to help individuals struggling with *want/should* conflict by devoting funds to educating them about the availability of commitment devices.

Facilitating *Should* Decisions without Restricting Choice

Another important line of research on the applications of *want/should* conflict has specifically examined how policy makers might be able to facilitate people’s selection of and support for *should* choices without restricting their choice at any point in time.

While research on commitment devices asks whether people value the opportunity to commit their future selves to making more *should* decisions and suggests that one way to increase *should* decision making is to make more commitment devices available, other

research has examined how policy makers can increase the odds that people will make *should* decisions without restricting choice. For example, how could what has been learned from past research on *want/should* conflict be harnessed to increase donations to charity and to increase support for policies with important long-term benefits but short-term costs without tying the hands of people's future selves? Specifically, we rely on what we know about the conditions that favor the preferences of the *should* self over those of the *want* self and discuss how *should* choices may be facilitated by manipulating the conditions in which decisions are made.

Breman (2007) conducted a study to investigate whether donations to a large charity could be increased by extending the period of time separating a person's decision to give and her actual payment date. Following Thaler and Benartzi's "Save More Tomorrow" study, Breman called her plan "Donate More Tomorrow." Breman (2007) found that of people who donated on a monthly basis to a large charity and had agreed to increase their monthly donations, those who were asked to increase their donations beginning in two months were willing to increase their monthly donations by 32% more than those who were asked to increase their donations beginning in one month. Thaler and Benartzi's "Save More Tomorrow" study did not isolate the effect of delaying the time until a *should* option (saving money) would be implemented on people's willingness to take up that *should* option. However, Breman's "Donate More Tomorrow" study did isolate the effect of a delay to implementation on people's willingness to do what they felt they *should* (donating money to a charity).

Rogers and Bazerman (in press) set out to isolate the effect of delayed implementation on people's support for a general group of *should* policies. They first

identified a set of policies that people report feeling they *should* support but do not *want* to support. These policies were identified using two different methods. For some policies, participants read a description of *want/should* conflict and were then asked whether a given policy was a *want* policy or a *should* policy. A policy was considered to reflect the interests of the *should* self, as opposed to those of the *want* self, when a significant majority of participants reported that it was a *should* policy. For other policies, participants were asked to evaluate the extent to which they felt they *should* support the policy and the extent to which they *wanted* to support the policy. A policy was classified as a *should* policy when participants reported feeling significantly more strongly that they *should* support it than that they *wanted* to support it. The following five *should* policies were selected using these methods: a policy that would increase a participant's savings rate (a modified version of "Save More Tomorrow"); a policy that would increase a participant's donations to charity (a modified version of "Donate More Tomorrow"); a policy that would increase the price of gasoline to reduce pollution; a policy limiting the number of fish that could be caught in the ocean to reduce overharvesting (thus increasing seafood prices); and a policy that would move a participant's employer to a more profitable region (where the participant would not want to live).

Consistent with Breman's (2007) findings, Rogers and Bazerman (in press) found that people report stronger support for *should* policies when these policies will be implemented in the distant rather than the near future. Consistent with Construal Level Theory (Trope and Liberman, 2003), which was described previously, Rogers and Bazerman determined that this "future lock-in effect" was partly mediated by how people

construe a policy option. For example, when a *should* policy (e.g., an increase in the price of gas to reduce pollution) will be implemented in the distant future, it is more strongly associated with its abstract, superordinate, goal-relevant attributes (e.g., this will reduce pollution), whereas when it will be implemented in the near future it is more strongly associated with its concrete attributes and tangible implications (e.g., this will increase in the price of gas).

As Rogers and Bazerman (in press) and Breman (2007) demonstrated, delaying the time to implementation may be a useful strategy for policy makers trying to bolster support for policies that people feel they *should* support but do not *want* to support. An important example of a contemporary issue that could benefit from this strategy, which Rogers and Bazerman call “future lock-in,” is the issue of how to reduce domestic consumption of fossil fuels and other materials that contribute to global climate change. While the vast majority of citizens agree that the United States needs to do more to reduce its contribution to this global problem (see Gallup Polls on pollingreport.com/enviro.htm), most proposed initiatives face stiff opposition. The opposition often comes from both producers of goods who are concerned about the impact of new policies on their profits and consumer groups concerned about the short-term costs of the policies. By advocating for reforms that would go into effect in the more distant future, policy-makers might be able to leverage the benefits of the future lock-in effect to increase the proportion of people who support *should* reforms, as well as the strength of their support for such reforms.

Entirely independent of the systematic changes in how much support people will give *should* legislation designed to take effect in the distant future, an added benefit of

delaying the time to a policy's implementation is that it gives affected parties more time to optimally prepare for the legislation's impact. For example, passing stricter automobile fuel-efficiency legislation that would take effect in seven years would have two practical benefits over identical legislation that would take effect sooner. First, vehicle owners could enjoy up to seven more years of value out of the vehicles they currently own, while replacing them with more efficient vehicles when, or closer to when, their current vehicles are ready for replacement. Second, future implementation of fuel-efficiency legislation would allow producers to gradually increase their capacity to manufacture more efficient vehicles.

A danger of passing policies designed to take advantage of the future lock-in effect is that future legislatures could overturn them. This danger is not as damning as one might first suppose, however, because initially passing a policy cognitively differs from overturning an existing one. Once a policy has been chosen for the future, people anticipate its instatement, and the policy gradually comes to be viewed as the default or status quo (e.g., "This fuel efficiency requirement has been on the horizon for years, and I don't want to go backwards by overturning it."). Past research has demonstrated the power of defaults (Choi, Laibson, Madrian, & Metrick, 2003; Johnson & Goldstein, 2003) and people's aversion to changing what they perceive to be the status quo (Samuelson & Zeckhauser, 1988; Kahneman, Knetsch & Thaler, 1991). While policies that will be implemented in the distant future do risk being overturned, overall, taking advantage of the future lock-in effect could be an effective political strategy for increasing support for policy options that are perceived as *should* options but cannot gain enough support to be implemented immediately.

Another attractive aspect of leveraging the future lock-in effect in the realm of public policy is that it would not actually require changing the time to implementation of many pieces of legislation. Many policies are already designed to go into effect in the distant future. Rather than changing the time to implementation of proposed *should* policies that are to be implemented in the distant future, policy makers could increase support for them simply by changing people's temporal focus when thinking about the policy. Specifically, emphasizing the distant future implementation of a proposed *should* policy (as opposed to emphasizing the near future decision to support or oppose the policy) could harness the future lock-in effect. Rogers and Bazerman (in press) found empirical support for this idea when they asked a national sample of subjects about how favorably they would view a policy that would increase the price of gas by \$.53 in two years, but which they would vote on in a few months. All participants in their study read a full description of the policy. Half of the participants also read text emphasizing when the policy would be implemented, while the other half read text emphasizing when the policy would be voted on. Participants who read the text emphasizing the distant future implementation indicated that they supported the policy significantly more and were significantly more likely to vote for the policy than participants who read the text emphasizing the near future vote.

The aim of leveraging the future lock-in effect is to increase people's support for *should* policies. This approach to policy design could be objected to on the grounds that it is paternalistic to try to influence people to choose some options (*should* options) more often than others (*want* options). In response to this objection, we argue that leveraging the future lock-in effect is consistent with a philosophy Sunstein and Thaler call

“libertarian paternalism” (2003). Libertarian paternalism is a term that describes policies designed to encourage welfare-promoting choices, which do not eliminate a decision maker’s freedom of choice. Moreover, note that the effectiveness of the strategy we propose actually requires that individuals face some internal conflict when weighing their options. In the absence of *want/should* internal conflict induced by a policy option, a delay to the time when a policy will be implemented would have no psychological effect on an individual’s support for that policy (although it might have a rational effect if the additional time to implementation changes the policy’s costs and benefits). To illustrate this point, Rogers and Bazerman (in press) showed that when a policy is not widely seen as favoring the interests of the *should* self, as opposed to those of the *want* self, distant future implementation does not affect support for it. This suggests that the future lock-in effect is even less paternalistic than many libertarian paternalistic policies like setting defaults to favor welfare-maximizing options. Whereas libertarian paternalism endorses strategically facilitating people’s selections of options that policy makers decide are welfare-promoting, future lock-in facilitates people’s selections of options that they, internally, believe are better for them in the long-run (*should* options).

The policy applications of *want/should* conflict discussed in this section are just a few of many that have yet to be fully explored. For example, selection of *should* options could be encouraged by scheduling decision making during low cognitive load times relative to high cognitive load times (Shiv and Fedorkihn, 1999), or by structuring decision contexts so that people evaluate options jointly rather than separately (Bazerman, Loewenstein, and White, 1992). Ultimately, many of the most important problems facing the world today are exacerbated by myopic decision making (e.g.,

climate change, under-saving for retirement, deficit spending, obesity, etc.). Solutions to these problems will require far-sighted and patient decision makers who select and support options that serve the interests of their *should* selves. This makes applications of research on *want/should* conflict critically important and useful today, and we believe the research to date on this area offers valuable insights to policy makers interested in finding ways to help people maximize their long-term welfare.

CONCLUSION

In this paper, we have reviewed fifteen years of research on *want/should* conflict and discussed what we believe to be the most important applications of this work. We have argued that the results of recent empirical studies of *want/should* conflict have the potential to help individuals and policy makers by arming them with insights about how to increase the chances that they and their constituents, respectively, will favor options that are in their best interest. In addition, we have offered specific prescriptions for how research on *want/should* conflict can be used to facilitate *should* decision making.

We believe there are many promising opportunities for future research on *want/should* conflict. Although better models of the sources of *want/should* conflict are needed, it seems to us that the two most important questions for academics to investigate in this area are: (1) what other moderators besides those that have already been explored affect whether individuals lean towards *want* or *should* options, and (2) what other mechanisms known to favor *should* choices over *want* choices besides delayed implementation have the potential to help individuals and policy makers increase support for *should* options. The more we know about what factors moderate people's preferences for *want* versus *should* options and about how these factors can be used to design

decision making contexts that favor the preferences of the *should* self, the more advice we will be able to give individuals and policy makers about how to solve problems that result from impulsive, short-sighted decision making.

REFERENCES

- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82, 463-509.
- Ainslie, G. (1992). *Picoeconomics: The interaction of successive motivational states within the individual*. New York: Cambridge University Press.
- Ainslie, G., & Haendel, V. (1983). The motives of the will. In E. Gottheil, K. A. Druley, T. E. Skoloda, & H. M. Waxman (Eds.), *Etiologic aspects of alcohol and drug abuse* (pp. 119–140). Springfield, IL: Charles C Thomas.
- Angeletos, G., Laibson, D., Repetto, A., Tobacman, J., & Weinberg, S. (2001). The hyperbolic consumption model: Calibration, simulation, and empirical evaluation, *Journal of Economic Perspectives*, 15(3), 47 – 68.
- Ashraf, N., Karlan, D., & Yin, W. (2006). Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics* May 2006, 121(2), 635-672.
- Bazerman, M., Loewenstein, G. & White, S.B. (1992). Reversals of preference in interpersonal decision making: The difference between judging an alternative and choosing between multiple alternatives. *Administrative Science Quarterly*, 37, 220-240.
- Bazerman, M.H., Moore, D.A., Tenbrunsel, A.E., Wade-Benzoni, K.A., & Blount, S. (1999). Explaining how preferences change across joint versus separate evaluation. *Journal of Economic Behavior and Organization*, 39, 41-58.
- Bazerman. M. H., Schroth, H., Pradhan, P., Diekmann, K., & Tenbrunsel. A. (1994). The inconsistent role of comparison others and procedural justice in reactions to

- hypothetical job descriptions: Implications for job acceptance decisions,
Organizational Behavior and Human Decision Processes, 60, 326-352.
- Bazerman, M. H., Tenbrunsel, A. E., & Wade-Benzoni, K. (1998). Negotiating with Yourself and Losing: Making Decisions with Competing Internal Preferences. *Academy of Management Review*, 23(2), 225-241.
- Breman, A. (2006). Give more tomorrow: Evidence from a randomized field experiment. Unpublished manuscript.
- Choi, J., Laibson, D., Madrian, B. & Metrick, A. (2003). Optimal Defaults. *The American Economic Review Papers and Proceedings*, 93(2), 180-185.
- Frederick, S., Loewenstein, G.A., & O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2), 351 – 401.
- Freud, S. (1961). (Originally published in 1923). The ego and the id. In J. Strachey (Ed. and Trans.), *The standard edition of the complete psychological works of Sigmund Freud*, 19, (pp. 3 – 66). London, UK: Hogarth Press.
- Fudenberg, D. & Levine, D. (2006). A dual self model of impulse control. *American Economic Review*, 96(2006), 1449-1476
- Fujita, K., Trope, Y., Liberman, N., & Levin-Sagi, M. (2006). Construal levels and self-control. *Journal of Personality and Social Psychology*, 90, 351–367.
- Gallup Polls. (2006, March 13-16). *PollingReport.com, Environment polls*. Retrieved October 12, 2006, from the PollingReport.com website:
pollingreport.com/enviro.htm

- Gruber, J. & Köszegi, B. (2004). Tax incidence when individuals are time inconsistent: The case of cigarette excise taxes. *Journal of Public Economics*, 88(9-110), 1959-1988.
- Hsee, C. K. (1995). Elastic justification in decision making: How task-irrelevant but tempting factors influence decisions. *Organizational Behavior and Human Decision Processes*, 62, 330-337.
- Hsee, C. K. (1996). The evaluability hypothesis: An explanation of preference reversals between joint and separate evaluations of alternatives. *Organizational Behavior and Human Decision Processes*, 46, 247-257.
- Hsee, C. K. (1998). Less is better: When low-value options are valued more highly than high-value options. *Journal of Behavioral Decision Making*, 11, 107-121.
- Hsee, C.K., Loewenstein, G.F., Blount, S. & Bazerman, M.H. (1999). Preference reversals between joint and separate evaluations of options: A review and theoretical analysis. *Psychological Bulletin*, 125(5): 576-590.
- Irwin, J. R., Slovic, P., Lichtenstein, S., & McClelland, G. 1993. Preference reversals and the measurement of environmental values. *Journal of Risk and Uncertainty*, 6(1), 5-18.
- Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, 302, 1338-1339.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5, 193-206.
- Kahneman, D., & Ritov, I. 1994. Determinants of stated willingness to pay for public goods: A study in the headline method. *Journal of Risk and Uncertainty*, 9: 5-38.

- King, G.R. & Logue, A.W. (1987). Choice in a self-control paradigm with human subjects: effects of changeover delay duration. *Learning and Motivation*, 18, 421-438.
- Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards, *Journal of Experimental Psychology*, 126, 54 – 70.
- Kirby, K.N. & Herrnstein, R.J. (1995). Preference reversals due to myopic discounting of delayed rewards, *Psychological Science*, 6(2), 83 – 89.
- Kirby, K.N. & Marakovic, N.N. (1996). Delayed-discounting probabilistic rewards rates decrease as amounts increase, *Psychonomic Bulletin and Review*, 3(1), 100 – 104.
- Kivetz, R. & Keinan, A. (2006). Repenting hyperopia: An analysis of self-control regrets. *Journal of Consumer Research*, 33, 273-282.
- Kivetz, R. & Simonson, I. (2002a). Self control for the righteous: Toward a theory of precommitment to indulgence, *Journal of Consumer Research*, 29(2), 199 – 217.
- Kivetz, R. & Simonson, I. (2002b). Earning the right to indulge: Effort as a determinant of customer preferences toward frequency program rewards. *Journal of Marketing Research*, 39(2), 155 – 170.
- Kivetz, Y., & Tyler, T. R. (2007). Tomorrow I'll be me: The effect of time perspective on the activation of idealistic versus pragmatic selves. *Organizational Behavior and Human Decision Processes*, 102, 193-211.
- Khan, U and Dhar, R (2007). Where there is a way, is there a will? The effect of future choices on current preferences. *Journal of Experimental Psychology – General*, 136(2), 277 – 288.

- Laibson, D. (1996). Hyperbolic discount functions, undersaving and savings policy, *NBER working paper w5635*.
- Laibson, D.I. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(May), 443–77.
- Liberman, N., Sagristano, M., & Trope, Y. (2002). The effect of temporal distance on level of construal. *Journal of Experimental Social Psychology*, 38, 523–535.
- Loewenstein, G. F. (1996). Out of Control: Visceral Influences on Behavior. *Organizational Behavior & Human Decision Processes*, 65(3), 272-292.
- Loewenstein, G.F. & Prelec, D. (1992). Anomalies in Intertemporal Choice: Evidence and an Interpretation, *Quarterly Journal of Economics*, 107, 573-597.
- Loewenstein, G. & Thaler, R. (1989). Anomalies: Intertemporal choice, *The Journal of Economic Perspectives*, 3(4), 181 – 193.
- Lowenthal, D. (1993). Reversals of preference in candidate choice. *Unpublished manuscript*, Carnegie-Mellon University, Pittsburgh, PA.
- McClure S. M., Laibson D. I., Loewenstein G, & Cohen J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306, 503-507.
- Milkman, K. Rogers, T. & Bazerman, M.H. (2007). Film rentals and procrastination: A study of intertemporal reversals in preferences and intrapersonal conflict. *HBS Working Paper 07-099*.
- O'Donoghue, T. & Rabin, M. (1999). Doing it now or later, *The American Economic Review*, 89(1), 103 – 124.

- Oster, S. & Scott Morton, F.M. (2005). Behavioral biases meet the market: the case of magazine subscription prices, *Berkeley Electronic Journals Economic Analysis and Policy, Advances*, 5(1).
- Read, D., Loewenstein, G., & Kalyanaraman, S. (1999). Mixing virtue with vice: Combining the immediacy effect and the diversification heuristic. *Journal of Behavioral Decision Making*, 12, 257–273.
- Read, D. (2001). Intrapersonal dilemmas. *Human Relations*, 54(8), 1093 – 1117.
- Rogers, T. & Bazerman, M.H. (in press). Future Lock-in: Future implementation increases selection of ‘should’ choices. *Organizational Behavior and Human Decision Processes*.
- Rogers, T., Milkman, K.L., & Bazerman, M.H. (2007). I'll Have the Ice Cream Soon and the Vegetables Later: Decreasing Impatience over Time in Online Grocery Orders. *HBS Working Paper 07-078*.
- Samuelson, W. F., & Zeckhauser, R. (1988). Status Quo Bias in Decision Making. *Journal of Risk and Uncertainty*, 1, 7-59.
- Schelling, T. C. (1984). *Choice and Consequence: Perspectives of an Errant Economist*. Cambridge, MA: Harvard University Press.
- Sen, A. K. (1977). Rational fools: A critique of the behavioral foundations of economic theory. *Philosophy and Public Affairs*, 6, 317-345.
- Shefrin, Hersch M. and Richard H. Thaler (1988). The Behavioral Life-Cycle Hypothesis. *Economic Inquiry*, 26(4), 609 – 643.

- Shiv, B. & Fedorikhin, A.. (1999). Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making. *Journal of Consumer Research*, 26(3), 278 – 292.
- Strotz, R. H. (1956). Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies*, 23, 165-180.
- Sunstein, C. R., & Thaler, R. H. (2003). Libertarian Paternalism Is Not an Oxymoron. *University of Chicago Law Review*, 70(Fall), 1159–99.
- Thaler R.H. (1981). Some empirical evidence on time inconsistency. *Review of Economic Studies*, 23,165–180.
- Thaler, R. H., & Benartzi, S. (2001). Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving. *Journal of Political Economy*, 112, S164-S187.
- Thaler, R. H., & Shefrin, H. (1981). An Economic Theory of Self Control. *Journal of Political Economy*, 89, 392-406.
- Trope, Y. & Fishbach, A. (2000). Counteractive self-control in overcoming temptation. *Journal of Personality and Social Psychology*, 79(4), 493-506.
- Trope, Y., & Liberman, N. (2000). Temporal construal and time-dependent changes in preference. *Journal of Personality and Social Psychology*, 79, 876–889.
- Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review*, 110, 403–421.
- Wertenbroch, K. (1998). Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice. *Marketing Science*, 17(4), 317 – 337.

Wilson, T.D., Lisle, D., Schooler, J., Hodges, S.D., Klaaren, K.J., & LaFleur, S.J. (1993).

Introspecting about reasons can reduce post-choice satisfaction. *Personality and Social Psychology Bulletin*, 19, 331 – 339.

APPENDIX A

While the below reference lists are not exhaustive, they represent the preponderance of the literature we are aware of on self-control problems and under-indulgence problems.

References to Literature about Self-Control Problems

Strotz, 1956; Ainslie, 1975; Thaler and Shefrin, 1981; Thaler, 1981; Ainslie and Haendel, 1983; Schelling, 1984; King and Logue, 1987; Shefrin and Thaler, 1988; Loewenstein and Thaler, 1989; Ainslie, 1992; Loewenstein and Prelec, 1992; Kirby and Herrnstein, 1995; Kirby and Marakovic, 1996; Laibson, 1996; Kirby, 1997; Wertenbroch, 1998; O'Donoghue and Rabin, 1999; Trope and Fishbach, 2000; Angeletos et al., 2001; Read, 2001; Thaler and Benartzi, 2001; Frederick et al., 2002; Gruber and Kozegi, 2004; McClure et al., 2004; Oser and Scott Morton, 2005; Ashraf et al., 2006; Fudenberg and Levine, 2006

References to Literature about Under-Indulgence Problems

Wilson et al., 1993; Kivetz and Simonson, 2002a; Kivetz and Simonson, 2002b; Kivetz and Keinan, 2006