



Psychology, Behavioral Economics, and Public Policy

ON AMIR
Yale University

oamir@ucsd.edu

DAN ARIELY
MIT

ALAN COOKE
University of Florida

DAVID DUNNING
Cornell University

NICHOLAS EPLEY
URI GNEEZY
University of Chicago

BOTOND KOSZEGI
University of California, Berkeley

DONALD LICHTENSTEIN
University of Colorado, Boulder

NINA MAZAR
MIT

SENDHIL MULLAINATHAN
Harvard University

DRAZEN PRELEC
MIT

ELDAR SHAFIR
Princeton University

JOSE SILVA
University of California, Berkeley

Abstract

Economics has typically been the social science of choice to inform public policy and policymakers. In the current paper we contemplate the role behavioral science can play in enlightening policymakers. In particular, we provide some examples of research that has and can be used to inform policy, reflect on the kind of behavioral science that is important for policy, and approaches for convincing policy-makers to listen to behavioral scientists. We suggest that policymakers are unlikely to invest the time translating behavioral research into its policy implications, and researchers interested in influencing public policy must therefore invest substantial effort, and direct that effort differently than in standard research practices.

Keywords: public policy, psychology, behavioral economics

Imagine waking up one morning, turning on the radio and hearing on the NPR news that the president of the US has issued the following statement:

“After consulting my chief psychologist, I am confident that the reframing proposed in the new well-being policy will increase happiness by 34% at almost no cost.”

Although this statement is unlikely to be uttered by any publicly elected official during our lifetime, this type of statement represents our hope that psychologists and behavioral economists will one day become more central and substantive contributors to public policy. Our hope is motivated by two observations. First, the past two centuries the study of human behavior has yielded many important and counterintuitive insights. Second, despite this accumulating knowledge of human nature and behavior, these findings rarely find their way to one of their most important potential applications—public policy.

The failure of psychology and behavioral science more generally to influence public policy is particularly painful and frustrating in light of the success of its sibling, economics, as the basis for policy recommendations. It is not that economics has nothing to offer policy—economics indeed provides policy-makers with vital tools. Rather, the success of economics clearly demonstrates that policy-makers are willing to look to academic fields for guidance in setting their policies. Given this general willingness to accept advice, it is unfortunate that behavioral scientists are not providing their own perspectives. One striking example of this failure involves one of the most famous experiments in social science—Zimbardo’s Stanford prison experiment (1971). Every undergraduate student in psychology knows of the terrifying behaviors of the participants in this experiment. Yet, the simple link from these results to the military and civilian penal system has yet to be made. The painful consequences of not doing so are abundantly clear.

Another frustrating aspect regarding the role of psychology in policy is that economics in general, and its application to public policy in particular, is tightly linked to psychology—considering and predicting how people will change their behavior as a consequence of policy or incentive changes inherently involves psychology. Although it is easy to see that economics is at its core a psychological science, it is also easy to see that it is based on naïve psychology that has not always stood up to empirical scrutiny. From this perspective, psychologists may play a larger role in public policy by helping economics develop more valid assumptions and intuitions for the science that would underlie policy in the future.

In this paper we review some examples that illustrate the importance of using more realistic psychological assumptions in policy, and consider the possible ways to increase the impact of behavioral research on public policy by outlining the hurdles and possible avenues of such influence.

Example 1: Savings

In the face of continuously declining savings rates in the US, economic research has been emphasizing the importance of retirement savings programs to the future welfare of workers

(Ando and Modigliani, 1963). Perhaps as a consequence of this emphasis, many savings plans offered today aim to enable people to act upon their best interests and save money for their retirement period. The troubling finding, however, is that despite the obvious economic benefit of saving for retirement too many employees without defined benefits type of pension plans do not save nearly enough for their future (Benartzi and Thaler, 2004).

One explanation of why retirement savings are so low is that people understand the value of saving but that they simply prefer to spend their money rather than save. In contrast, behavioral scientists point out that many retirement savings programs make unrealistic assumptions about the mental capabilities of the people facing retirement decisions. Empirical evidence makes it clear that low savings rates are at least partly produced by the complexity in choosing between different savings options (Huberman et al., 2003), the substantial impact of defaults in decision making (Samuelson and Zeckhauser, 1988), the greater impact of losses relative to gains (Kahneman and Tversky, 1979), and the intertemporal asymmetry between the costs and benefits of the decision (Thaler and Shefrin, 1981).

Understanding some of these behavioral antecedents to individuals' poor decision making led Benartzi and Thaler (2004) to the creation of an innovative savings policy called "Save More Tomorrow"—a policy that explicitly builds on the apparently irrational aspects of human decision making rather than assuming them away. In this program individuals are asked to commit to a future saving that would be taken out of an upcoming raise. People may opt out of this program, but the default is to stay enrolled and continue increasing one's savings out of future raises. This program therefore capitalizes on people's tendency to care about the present more than the future (i.e., temporal discounting), to select default options when making decisions (i.e., the status quo bias), and to adapt to their current circumstances (i.e., adaptation).

Example 2: Fertilizers

A second example of the importance of behavioral understandings in overcoming individual shortcomings is based on recent experimental evidence on the usage of fertilizers in Africa (Duflo et al., 2005). In this work the authors examined why poor farmers in Western Kenya fail to use beneficial fertilizers on their crops. The results of multiple controlled experiments showed that fertilizers were effective in increasing the yield of the crops, and that they could dramatically increase the wealth of the farmers. At the same time, even the farmers that were educated about these effects with actual demonstrations did not systematically utilize these fertilizers. Duflo et al. (2005) demonstrated that the main reason for this "irrational behavior" was the farmer's inability to save the income necessary to purchase these fertilizers. The farming environment in Western Kenya is such that the sellers of the fertilizers approached the farmers toward the beginning of the season when the fertilizers would be most effective. This also happens to be the time furthest from the harvest season, when the farmers had the least amount of savings.

Once Duflo et al. (2005) identified savings as the main barrier to usage, they designed a new payment timing for fertilizers and with prepayment were able to dramatically increase the usage of fertilizers. Further increases in the adoption of fertilizers were achieved by coupling the initial sale of the crop and the purchasing of the fertilizers.

Example 3: Police Lineups

A third example of the importance of understanding the ways in which people make decisions concerns police lineups. Police lineups in most places constitute of a witness attempting to recognize one person in a group of potential suspects standing next to each other in a line (hence the name). Although these lineups are popular, accumulating evidence shows that they suffer from the disturbing possibility of false recognition—mistakenly identifying an innocent person as the perpetrator. Using insights about how individuals form judgments, Gary Wells and colleagues recently suggested an improvement to the lineup system whereby suspects would be evaluated sequentially versus simultaneously (Turtle et al., 2003; Wells et al., 2000). In a telling example, Lindsay and Wells (1985) convincingly showed that the probability of a false recognition of an innocent suspect is reduced dramatically when the presentation mechanisms changes. Instead of the common practice of presenting multiple individuals simultaneously (which creates a tendency to identify one of the people as the suspect), a better approach is to present the individuals one at a time, making sequential identification judgments. This improved methodology is already being used in Ontario and New Jersey.

What Kind of Behavioral Science is Important for Policy?

As the examples above illustrate, in principle it is clear that behavioral science has important implications for public policy. What is less clear is what *kind* of behavioral science is best suited to inform public policy. If the kind of behavioral science that is ideal for policy-making were different from the behavioral science commonly practiced, how would it need to be changed in order to better fit this role? In this section we explore how research that informs public policy might be different from the behavioral science as it is usually practiced.

Distinction 1: Vaguely Correct or Precisely Wrong?

The distinction between a theoretical and an applied science is particularly important as it becomes clearer that the accurate answer to many of the questions concerning how people behave in certain situations is “it depends.” In fact, as more knowledge about human behavior accumulates, it becomes evident that situational factors have tremendous impact on the action of the individual, even seemingly minor situational factors. These situational factors include the framing of a stimulus (Tversky and Kahneman, 1981), the relative attribute value of alternatives in a choice set (Huber et al., 1982; Simonson, 1989), the type of response required (Tversky et al., 1988), the number of alternatives provided (Iyengar and Lepper, 2000) or implied (Amir and Ariely, 2004), the temporal nature of the decision (Laibson, 1997), the emotional aspects of the decision (Loewenstein et al., 2001; Slovic et al., 2002; Slovic et al., 2002), the order in which the alternatives are presented (Russo et al., 1996), and a host of other factors too long to list. Under this “it depends” state of the world, behavioral scientists are unable to provide policy makers with unqualified answers to their questions. Yet this is exactly the input that policy-makers need most in order to

better craft policy. Here behavioral scientists face a difficult tradeoff between providing an accurate and vague answer that offers no clear recommendations, or an inaccurate and precise answer that offers a clear policy and might bring about at least some change.

Distinction 2: Theoretical or Applied Research?

A related distinction is whether the best role behavioral scientists could play in the policy arena is to search for general principles or to concoct exact recipes for how to address specific policy issues. Despite Kurt Lewin's claim that "There is nothing so useful as a good theory," and James Maxwell's claim that "There is nothing more practical than a good theory," there is still a large gap between finding general principles and using those principles to prescribe particular policies. Although theories and general principles are clearly useful, it is also clear that policymakers themselves are not going to conduct the research needed to translate these general constructs into specific policies. Behavioral scientists who want their work so translated will need to do that work themselves to bridge the gap between theoretical and applied settings. Taking these extra steps means not only doing more applied work, but also becoming an expert in the particular policy domains (savings, healthcare, taxes, education, police lineups, etc.). Without such expertise, the researcher might not be able to understand important situational nuances (legal limitations, all affected parties, long-term consequences, etc.) and may therefore conduct research that would miss some of the central aspects of the applied domain. Finally, it is naïve to expect policymakers to read academic journals, and the applied research should be disseminated in channels that are easily accessed by policymakers—including the popular press and personal communications.

Distinction 3: What Technology to Use?

A final distinction involves the experimental materials used in research. From a scientific perspective, it is almost always better to pick stimuli that would allow the researcher to directly and unambiguously attribute the effects to the theoretical construct. To achieve this goal, scientific experiments often utilize artificial stimuli that are uncommon or nonexistent in everyday life. For example, the use of simple gambles of the form win \$ x with probability y has been instrumental in exploring decision making under uncertainty, but is clearly more abstract than any of the probabilistic decisions people encounter in their daily lives. Decisions regarding investment portfolios, insurance policies, and lottery tickets are carried out in somewhat different environments than simple gambles. As a consequence, decisions about these classes of stimuli have the potential to play out very differently. To make research in behavioral science more applicable, the stimuli used should reflect the richness of the environment they are meant to represent. Such selection of stimuli will increase the ability to generalize results to the setting of the policy, and increase the face validity of experiments making it easier to "sell" them to policymakers. One downside of more realistic stimuli is that they often simultaneously manipulate multiple factors, making the causal inferences so central to science harder to make.

Some Possible Directions for Behavioral Policy Research

In this section we would like to point out a few possible directions for policy-oriented behavioral research. Before discussing such possible directions it is important to make a few comments about paternalism. To the extent that behavioral research on policy is successful, policy-makers will be equipped with tools to increase the effectiveness of policies. For example, behavioral research might help create policies that would increase savings, decrease drunk driving, inhibit violence and aggressions, or increase the duration of marriages. Although behavioral research is likely to make such policies more effective, it is still not clear that the government *should* implement them. The question of paternalism, control, and manipulation of the citizens is a complex and delicate issue that is beyond the scope of the current discussion—yet at the same time, the question of paternalism is central to the issue of research into policy because any successful research could potentially increase paternalism. Individuals who have strong anti-paternalistic views may decide at this point that they do not want to increase the potential for paternalism and hence do not want to take part in any research related to policy. This perspective is understandable, but it is worth pointing out that policies uninformed by scientific research are no less paternalistic; they are simply likely to be less effective. To illustrate, consider a hypothetical policy that involves framing of taxation changes as either gains or losses. The use of such framing may invoke the notions of paternalistic interventions, but it is important to note that with or without intention to use framing, and with or without the knowledge of their effectiveness, framing is inherent. The question therefore is not whether to have framing or not, but rather whether it should be considered as a part of the policy or be allowed to be determined outside of the control of the policy making process. Finally, it is important to note that in many cases it is not a question of whether a policy should be implemented or not, but rather *how* a policy should be implemented. Paternalism is obviously a lesser issue in these cases.

Returning to the question of possible directions for behavioral policy research, a beautiful example of research that is highly related to a policy question is the work on the effects of defaults of organ donations (Johnson and Goldstein, 2003), showing that the proportion of people who have organ donor status is over 90% in countries where the default is donation (and people have to opt out if they don't want to be a donor), while the proportion of people who have organ donor status is below 20% in countries where the default is non-donation (and people have to opt in if they want to be a donor). There are many other cases in which the power of defaults can be harnessed to do good—it can be used to help people contribute to their 401K plans, to their Roth accounts, to enroll people in healthcare, gyms, etc. (again with all the problems related to paternalism).

A second example of research that might be highly applicable to policy questions is the research on context effects such as the asymmetric dominance effect (Huber et al., 1982), and the compromise effect (Simonson and Tversky, 1992). The work looking at context effects has repeatedly demonstrated that the alternatives provided in the choice set, even if they are not chosen, can have substantial effects on the options that are chosen. In the domain of policy, these effects could be used to influence the choices individuals make on a range of topics from healthcare plans, to the selection of public officials, and even to convince people that they are not paying much income tax.

A third example could be based on anchoring (Kahneman and Tversky, 1974; Epley and Gilovich, 2001). It has been repeatedly demonstrated that asking people to answer a question about their willingness to pay (for example: would you pay an amount equal to the last two digits of your social security number for this box of chocolate), can have a substantial effect on their true willingness to pay for the good when elicited later using an incentive compatible procedure. In the policy domain, anchoring can be used to “help” people contribute more to charity, increase their savings, etc.

An example that is perhaps more far fetched relates to Duval and Wicklund’s (1972) notion of objective self-awareness. The idea is that when attention is directed toward oneself that induces self-evaluation in relation to standards salient or accessible in the immediate situation, which in turn increases motivation to meet the standard (see also (Carver and Scheier, 1998). One finding within this framework has shown that when students sign an honor code, or when they are simply asked to recall the ten commandments, they are less likely to cheat on a test relative to a control group. Moreover, this effect persists even when the performance on the test has real financial consequences (Mazar et al., 2005). Using related devices such as awareness, recall, or honor code signatures to increase self-awareness could have far reaching implications. One could easily imagine applying this principle to driving (reducing the tendency not to obey traffic rules), to personal tax returns (decreasing tax evasion), and to dishonesty at the workplace. For example, asking individuals to sign their tax returns forms at the beginning rather than the end might increase their self-awareness and decrease deception.

Another example of an application of a well-documented result to the domain of policy involves an examination of the framing of tax reduction on spending. In a recent paper, Epley et al. (2004) examined why the effect of the 2002 tax return on the economy was smaller than anticipated. Based on a series of experiments the authors conclude that if the tax reduction had been framed as a “bonus” rather than a “rebate,” people would have spent significantly more of it. More generally, framing can be used in many situations ranging from framing the propositions citizens vote on during election times, to Medicare prescription options, and even to the question of how to trade-off personal freedom for security.

A final example of an application of established results relate to the “hot cold empathy gap” (Loewenstein, 1996). This work has demonstrated that when people are in a “cold” and non-emotional state, they are unable to accurately predict how they themselves would behave if they were in a “hot” emotional state. Drawing on personal experience, it is commonly observed that people who go food shopping while hungry usually buy too much food, and moreover that they do not seem to learn from their past experiences. A more controlled examination of this idea was provided by Ariely and Loewenstein (forthcoming) where they asked subjects to indicate the likelihood that when aroused they will have safe sex, and the likelihood that when aroused they will behave themselves immorally in order to secure sexual gratification. The male respondents who answered these questions in a cold state indicated that they were unlikely to take risks of unprotected sex and that they would not engage in morally questionable behavior in order to obtain sexual gratification. On the other hand, when sexually roused, the same participants gave dramatically different responses. Indicating that they would take risks of unprotected sex and engage in morally questionable behavior in order to obtain sexual gratification. Such changes of preferences

in the “heat of the moment” coupled with the intra-personal empathy gap (the inability to fully predict these changes) can have substantial implications for our understanding of the causes of undesirable behaviors such as crime of passion, date-rapes, and unprotected sex, and hence can educate policies that are designed to curb such behaviors. For example, understanding the effects of preferences in the “heat of the moment” should influence what behaviors sex offenders are instructed to avoid, and how sexual education in high schools is designed.

In summary, there are many aspects of research that could be applicable to policy and policymaking. In seeking where to start, one approach is to look for research directions that are likely to create the lowest levels of resistance and opposition from policymakers and their constituents. One type of interventions that are likely to result in little resistance is one that in general, people don’t believe would have any effect on behavior. We term these research directions “small interventions,” and bundle under this title all the possible effects that lay people, including all of those who are not familiar with the behavioral literature, would predict not to have any effects on behavior. There is a long list of such “small interventions,” including most of the findings noted in this section. For example, the data suggest that participants in anchoring experiments don’t believe that answering a hypothetical question about their willingness to pay can actually change their willingness to pay. Anchoring and other interventions that have the feature of unaware-effectiveness would be therefore classified under the “small interventions” category. We reason that if policy makers predict that changes of the small interventions type will lead to no or small effects, or if they predict that others will predict a low effectiveness they would be more likely to embrace such policy-mechanisms.

Some Possible Approaches to Get Policy-Makers to Listen

Perhaps more difficult than identifying relevant policy application from behavioral science research, however, is convincing policy makers to seek the advice of behavioral scientists and to take this advice seriously. The first approach we would like to promote is the grassroots approach. The idea here is that instead of hoping that someone in Washington DC will read research papers in behavioral science or invite researchers to provide advice on policy issues, a better way might be to start at local communities. Starting at the communities researchers live in has the advantages that they know the environment better, are closer geographically, the stakes are lower (which should make it simpler to try something new), and hopefully the bureaucracy is less potent, generating lower hurdles for implementation. Moreover, to the extent that a change in local policy is successful it could be spread by people who are using this policy in their day-to-day lives. One example of a successful grassroots approach is the abovementioned change to the policy of police lineup promoted by Gary Wells and colleagues (Turtle et al., 2003; Wells et al., 2000). Using the grassroots approach, researchers related to this project were individually involved in educating policemen and judges about their findings. Consequently, improved policy was introduced in Ontario and New Jersey, not only getting police to adopt this procedure but also getting judges to start demanding that police use this procedure regularly.

A second approach for getting policymakers to embrace behavioral science involves influencing policy via economics. The idea here is to use the established path from economics to policy—attempting to modify economics to be more descriptively accurate, and from there influencing policy. A prime example for this type of approach is prospect theory (Kahneman and Tversky, 1979), which formalized the idea that judgments and preferences were reference dependent, and has since spurred many applications. In one example Ariely et al. (2004) provide experimental evidence for the dependence of consumers' maximum willingness to pay (WTP) on the prices they expect to see in the marketplace—challenging the assumption that demand (WTP) is an independent force from production (supply) (see also, Amir et al., 2005). Their results show that as the price distribution for products increases in magnitude (i.e., a shift in the supply curve), so does consumers' willingness to pay (i.e., shifting the demand curve). They then go further and illustrate how neoclassical economists, who assume that the forces of supply and demand are independent, will be led astray when they calculate the effects of policy changes, such as taxation, on consumption. In particular, they show that the assumption of independence will overestimate the effects of taxation, and that this overestimation will increase as the dependency of supply on demand increases. If these results were to hold more generally, and if this dependency were to be incorporated in the economics models attempting to estimate the effects of policy changes, the estimation might be more accurate.

A third approach for increasing the influence of behavioral science involves influencing policy via law. As in the example of the eye-witness research (e.g., Wells et al., 2000), or the recently evolving field of behavioral-law-and-economics (Sunstein, 2000; Jolls and Sunstein, 2004), legal academia influences both judges, lawyers (i.e., grassroots) and regulators, and thus may potentially provide the right bridge for the existing gap between behavioral research and policy. For example, Jolls and Sunstein (2004) consider the potential to correct behavioral biases through corrective regulation. However, as may be suggested by the narrow scope of overlap between behavioral research and the field of law, it is possible that the applicability of behavioral science to law might be limited in its scope.

The final and most challenging approach to induce policymakers to listen is to directly do research on policy. As behavioral scientists we are very used to pilot testing our ideas—knowing all too well that we cannot possibly think about all the possible things that could go wrong with our design, and use the pilot data to validate or modify our thinking. Moreover, we are also painfully aware of the effort and cost of running experiments and use pilot testing to minimize the waste of money and time. It is amazing, to say the least, there are virtually no policy pilot experiments. If anything, we would expect there to be many more pilot testing in policy given the complexity of the environment, the high uncertainty, and in particular given the incredible cost of implementing policy. How is it that the government cuts taxes by billions of dollars without any pilot test? Why not give the residents of Rhode Island (just as an example) one of four levels of tax cuts for a year or two and see the effect? Wouldn't this be much more efficient and beneficial in the long run? The main point of this type of behavioral policy research is that in many cases it is hard to make inferences from particular studies to a real policy question and that the only way to truly determine the effectiveness of policies is to engage in policy testing as an experimental endeavor. Obviously this idea is going to be difficult for policy makers to accept since it

is so different from the way they currently go about making policy decisions, but we can dream about a day in which the Congress will debate the experimental design of a newly proposed policy-experiment to test the effects of increased funding to higher educational institutes on welfare.

Summary

There is no question that the insights from research in psychology and behavioral economics could be very useful in informing policy decisions. If the designers of the prison systems would have been more familiar with the work of Zimbardo (1971), the travesty at the Abu Ghraib Prison (as well as in others) might have been prevented. Despite the importance of behavioral insights, the picture we draw here on the relationship between behavioral science and public policy is not a very optimistic one. In fact, the obstacles facing behavioral researchers who want to influence policy are substantial.

Because of these obstacles we highly recommend that behavioral scientists who want to take this path choose policy domains they are personally passionate about—hopefully this passion will carry them throughout the process and give them the required energy. A second advantage of general interest in a particular domain stems from the idea that in order to conduct experiments that can inform policies the experimental setup must take into account the factors that are most relevant to the policies in question. Without domain-specific knowledge academic researchers are likely to miss some of the important elements. Thus, it is clear to us that to influence policies individual researchers have to acquire specific knowledge and expertise in the policy domain.

The experimental setup to answer policy questions should also be considered. For example, research stimuli should have high face validity. The experimental design should closely resemble reality, even at a cost of ability to pinpoint the causes of the results. While some may argue that this is only a cosmetic issue, it is still crucial. Using ecologically valid stimuli is also instrumental in creating more precise recommendations to policy makers. An additional surface feature relates to the ways researchers present themselves. We find it hard to imagine that one day the President of the US will consult his or her psychological advisor (or at least publicly admit to doing this). The popular image of psychology usually conjures up the images of psychotherapy, Freud, and the leather couch—and as such does not necessarily create a positive image for policy. Psychologists can potentially improve their position by calling themselves behavioral scientists, or coming up with a new and even more impressive title (behavioral policy science?).

Although the overall picture we draw may seem daunting in its implications for how difficult it is for behavioral researchers to truly influence policy, the battle is not yet lost and as more researchers join this initiative, treading this path is likely to become easier. On a practical level, we have two specific recommendations. The first is for behavioral scientists to participate in the policy making maelstrom of Washington DC in the same way as economists. This is not easy, but being willing to spend a few years in DC and taking the steps to do so is likely to yield progress. The second is to start small. Instead of imagining that Congress will read our latest paper and decide to change their policy, researchers should

try to approach local institutions around their community (local government, school boards, local police etc.), as they are more likely to adopt changes. The likelihood of a grassroots movement succeeding at this point seems to us to be much higher.

References

- Amir, O. and D. Ariely. (2004). "The Pain of Deciding: Indecision, Procrastination, and Consumer Choice Online." Unpublished manuscript.
- Amir, O., D. Ariely, and Z. Carmon. (2005). "The Locus and Appropriateness of Monetary Evaluations: Why Monetary Assessments do not Reflect Predicted Utility," Unpublished manuscript.
- Ando, A. and F. Modigliani. (1963). "The 'Life Cycle' Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review* 53, 55–84.
- Ariely, D., B. Koszegi, and N. Mazar. (2004). "Price-Sensitive Preferences," Unpublished manuscript.
- Ariely, D. and G. Loewenstein. (forthcoming). "In the Heat of the Moment The Effect of Sexual Arousal on Sexual Decision Making," *Journal of Behavioral Decision Making*.
- Benartzi, S. and R. Thaler. (2004). "Save More Tomorrow: Using Behavioral Economics to Increase Employee Savings," *Journal of Political Economy* 112(1), 164–187.
- Carver, Charles S. and Michael F. Scheier. (1998). *On the self-Regulation of Behavior*. New York, NY, USA: Cambridge University Press.
- Duflo, E., M. Kremer, and J. Robinson. (2005). "Understanding Technology Adoption: Fertilizer in Western Kenya, Preliminary Results from Field Experiments," Working paper MIT
- Duval, Shelley and Robert A. Wicklund. (1972). *A Theory of Objective Self Awareness*. New York, NY, USA: Academic Press.
- Epley, N., L. Idson, and D. Mak. (2004). "Rebate or Bonus? The Impact of Income Framing on Spending and Saving," Unpublished manuscript.
- Epley, N. and T. Gilovich. (2001, Sep). "Putting Adjustment Back in the Anchoring and Adjustment Heuristic: Differential Processing of Self-Generated and Experimenter-Provided Anchors," *Psychological Science* 391–396.
- Fish, S. (2004). Why We Built the Ivory Tower. *The New York Times*, May 21st.
- Huber, J., J. Payne, and C. Puto. (1982). "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis," *Journal of Consumer Research (pre-1986)* 9(1), 90.
- Huberman, G., S. Iyengar, and W. Jiang. (2003). "Defined Contribution Pension Plans: Determinants of Participation and Contribution Rates." Working paper, Columbia Business School.
- Iyengar, S. S. and M. R. Lepper. (2000). "When Choice is Demotivating: Can One Desire too Much of a Good Thing?" *Journal of Personality and Social Psychology* 79(6), 995–1006.
- Johnson, E. J. and D. Goldstein. (2003, Nov 21). "Do Defaults Save Lives?" *Science* 1338–1339.
- Jolls, C. and C. R. Sunstein. (2004). "Debiasing Through Law". Unpublished manuscript.
- Kahneman, D. and A. Tversky. (1979). "Prospect Theory: An Analysis of Decisions Under Risk," *Econometrica* 47, 263–291.
- Laibson, D. (1997). "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 62, (May), 443–477.
- Lindsay, R. C. L. and G. L. Wells. (1985). "Improving Eyewitness Identification from Lineups: Simultaneous Versus Sequential Lineup Presentations," *Journal of Applied Psychology* 70, 556–564.
- Loewenstein, G. (1996). "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Processes* 272.
- Loewenstein, G. F., C. K. Hsee, E. U. Weber, and N. Welsh. (2001). "Risk as Feelings," *Psychological Bulletin* 267–286.
- Mazar, Nina, On Amir, and Dan Ariely. (2005). "Almost Honest: Internal and External Motives for Honesty," MIT Sloan School of Management.
- Russo, E. J., V. H. Medvec, and M. G. Meloy. (1996). "The Distortion of Information During Decisions," *Organizational Behavior and Human Decision Processes* 66(1), 102–110.

- Samuelson, W. and R. J. Zeckhauser. (1988). "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty* 1(March), 7–59.
- Simonson, I. (1989). "Choice Based On Reasons: The Case Of Attraction And Comprom," *Journal of Consumer Research* 16(2), 158.
- Simonson, I. and A. Tversky, (1992). "Choice in Context: Tradeoff Contrast and Extremeness Aversion," *Journal of Marketing Research* 29(3), 281–295.
- Slovic, P., M. Finucane, E. Peters, and D. G. MacGregor. (2002). "The Affect Heuristic." In T. Gilovich (ed.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York, NY, US: Cambridge University Press pp. 397–420 .
- Slovic, P., D. Griffin, and A. Tversky. (2002). "Compatibility Effects in Judgment and Choice," In T. Gilovich (ed.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York, NY, US: Cambridge University Press pp. 217–229.
- Sunstein, C. R. (2000). *Behavioral Law and Economics*. (ed.), Cambridge University Press.
- Thaler, R. H. and H. M. Shefrin. (1981). "An Economic Theory of Self-Control," *Journal of Political Economy* 89(April), 392–406.
- Turtle, J. W., R. C. L. Lindsay, and G. L. Wells. (2003). "Best Practice Recommendations for Eyewitness Evidence Procedures: New Ideas for the Oldest Way to Solve a Case," *The Canadian Journal of Police and Security Services* 1, 5–18.
- Tversky, A. and D. Kahneman. (1981). "The Framing of Decisions and the Psychology of Choice," *Science* 211(4481), 453–458.
- Tversky, A. and D. Kahneman. (1974). "Judgment Under Uncertainty: Heuristics and Biases," *Science* 185, 1124–1131.
- Tversky, A., S. Sattath, and P. Slovic. (1988). "Contingent Weighting in Judgment and Choice," *Psychological Review* 95(3), 371–384.
- Wells, G. L., R. S. Malpass, R. C. L. Lindsay, R. P. Fisher, J. W. Turtle, and S. Fulero. (2000). "From the Lab to the Police Station: A Successful Application of Eyewitness Research," *American Psychologist* 55, 581–598.
- Zimbardo, P. G. (1971). The power and pathology of imprisonment. Congressional Record. (Serial No. 15, October 25, 1971). Hearings before Subcommittee No. 3, of the Committee on the Judiciary, House of Representatives, Ninety-Second Congress, First Session on Corrections, Part II, Prisons, Prison Reform and Prisoner's Rights: California. Washington, DC: U.S. Government Printing Office.