

12-1-2010

Special Problems for Democratic Government in Leveraging Cognitive Bias: Ethical, Political, and Policy Considerations for Implementing Libertarian Paternalism

J. Aaron Brown
Georgia State University

Follow this and additional works at: http://scholarworks.gsu.edu/philosophy_theses

Recommended Citation

Brown, J. Aaron, "Special Problems for Democratic Government in Leveraging Cognitive Bias: Ethical, Political, and Policy Considerations for Implementing Libertarian Paternalism." Thesis, Georgia State University, 2010.
http://scholarworks.gsu.edu/philosophy_theses/76

This Thesis is brought to you for free and open access by the Department of Philosophy at ScholarWorks @ Georgia State University. It has been accepted for inclusion in Philosophy Theses by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

SPECIAL PROBLEMS FOR DEMOCRATIC GOVERNMENT IN LEVERAGING
COGNITIVE BIAS: ETHICAL, POLITICAL, AND POLICY CONSIDERATIONS FOR
IMPLEMENTING LIBERTARIAN PATERNALISM

by

J. AARON BROWN

Under the Direction of Eddy Nahmias

ABSTRACT

Humans have now amassed a sizable knowledge of widespread, nonconscious cognitive biases which affect our behavior, especially in social and economic contexts. I contend that a democratic government is uniquely justified in using knowledge of cognitive biases to promote pro-democratic behavior, conditionally justified in using it to accomplish ends traditionally within the scope of government authority, and unjustified in using it for any other purpose. I also contend that the government ought to redesign institutional infrastructure to avoid triggering cognitive biases where it is not permitted intentionally to manipulate such biases and to optimize the effects of such biases where permissible. I shall use the United States of America as an example throughout, but my conclusions apply equally to any democracy which values the political autonomy of its populace.

INDEX WORDS: Cognitive bias, Libertarian paternalism, Democracy, Autocracy, Autonomy, Public policy, Richard Thaler, Cass Sunstein.

SPECIAL PROBLEMS FOR DEMOCRATIC GOVERNMENT IN LEVERAGING
COGNITIVE BIAS: ETHICAL, POLITICAL, AND POLICY CONSIDERATIONS FOR
IMPLEMENTING LIBERTARIAN PATERNALISM

by

J. AARON BROWN

A Thesis Submitted in Partial Fulfillment of Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2010

Copyright by
James Aaron Brown
2010

SPECIAL PROBLEMS FOR DEMOCRATIC GOVERNMENT IN LEVERAGING
COGNITIVE BIAS: ETHICAL, POLITICAL, AND POLICY CONSIDERATIONS FOR
IMPLEMENTING LIBERTARIAN PATERNALISM

by

J. AARON BROWN

Committee Chair: Eddy Nahmias

Committee: Andrew I. Cohen

William Edmundson

Electronic Version Approved:

Office of Graduate Studies

College of Arts and Sciences

Georgia State University

December 2010

Acknowledgements

I would like to offer thanks to my committee for their thoughtful and swift feedback, to Owen D. Jones for invaluable consultation in the early stages of the project, and to Patricia Rackliffe for being a rockstar registrar when bureaucracy goes wrong. Thanks also to my mother, Debra H. Brown, who has listened to me rant about this topic far more than any non-specialist should. Last, but certainly not least, I would like to thank my wife, Jennifer L. H. Brown, for her support, patience, and love.

Table of Contents

| | |
|---|----|
| <u>Acknowledgements</u> | iv |
| <u>Introduction</u> | 1 |
| I. Theory of Democratic Government | 3 |
| II. Cognitive Biases | 6 |
| III. Three Problems in Governmental Leveraging of Cognitive Biases | 10 |
| IV. Special Justification for Using Cognitive Biases to Promote Pro-Democratic Behavior | 17 |
| V. Conditional Justification in Domains Traditionally Reserved to Government Authority | 25 |
| VI. Unjustified in Other Domains & the Duty to Re-Evaluate Existing Biases | 30 |
| VII. Conclusion | 33 |
| <u>Sources</u> | 36 |

Introduction

Humans have now amassed a sizable knowledge of widespread, nonconscious cognitive biases which affect our behavior and decisions, especially in social and economic contexts. These psychological “shortcuts” can bypass reflective consideration, distorting our reasoning and often leading to decisions which do not advance our own interests. We may, for example, irrationally fear unlikely dangers like plane crashes, which kill only one person in 5,862, over much more likely ones like heart disease, which kills one person in six (National Safety Council, p. 37). Time after time, we reach for the remote instead of our running shoes, even though we know that we would, in the long run, be better off if we chose to exercise instead. Cognitive bias research looks for regularities in how we make these sorts of misjudgments.

Government can, and in America already does, use cognitive biases to change how citizens act, most commonly with an eye towards promoting the wellbeing of the populace. Richard Thaler and Cass Sunstein, leading proponents of this technique, argue that using cognitive biases to promote the best interests of the people is not paternalism as traditionally understood but is rather akin to “nudging” people in directions which are in their own long term interests (Thaler & Sunstein, 2003; 2008). Paternalism has historically referred to systems in which the government makes decisions for the citizens and penalizes noncompliance on the assumption that government is in a better position to judge what is best than are the affected individuals themselves. However, with a few notable exceptions like Medicare, Americans have generally rejected the so-called “nanny state,” and there is a strong current of libertarian rhetoric in American politics. Thaler and Sunstein contend that leveraging cognitive biases can promote the wellbeing of the citizenry without running afoul of ideological opposition to paternalism as traditionally understood by simply making it *more likely* that individuals will choose the

government-preferred course of action rather than mandating it (Thaler & Sunstein, 2003, 177-8). They call their approach “libertarian paternalism” (Thaler & Sunstein, 2003, title), a nod to the disparate philosophies they seek to reconcile.

Libertarian paternalism is not without its problems, though. Especially in a democracy, where the government must remain responsive to the desires of the people, using nonconscious psychological effects to manipulate the behavior and preferences of the citizenry is a perilous undertaking. In addition to damaging the general autonomy American libertarians consider a core value of a free society, bias-leveraging policies also run the very real risk of damaging the political autonomy of the voting populace. However, now that knowledge of cognitive biases exists, policymakers cannot simply ignore it without abdicating their responsibility to govern efficiently and effectively. As a practical matter, the mere fact that the tools of cognitive bias are available puts pressure on policymakers to use them. Second, some cognitive biases change not just how people choose but also what they actually *want*. A democratic government must respond to the desires of its populace, but the issue becomes much more complicated if those desires are themselves a result of biases leveraged by the government. Finally, the potential for overreaching is very real; regulations which set the cost of choosing a non-preferred option too high could easily lose the non-coercive openness that makes libertarian paternalism an attractive approach in the first place.

Toward the goal of minimizing or avoiding these three issues, I shall present here ethical considerations and policy guidelines for policymakers working with cognitive biases. I contend that a democratic government is especially justified in using knowledge of cognitive biases to promote pro-democratic behavior, conditionally justified in using it to accomplish ends traditionally within the scope of government authority, such as protecting public safety, and

unjustified in using it for any other purpose, especially in ways that would distort the political process. I also contend that the government ought, on the one hand, to restructure institutional infrastructure to avoid triggering cognitive biases where it is not permitted intentionally to manipulate such biases, as in the hypothetical case of a default policy that a particular religious text be used when swearing in witnesses at trial. On the other hand, the government also ought to optimize the effects of bias where permissible, as it has in the case of 401(k) enrollment regulations (see §II, below). I shall use the United States as an example throughout, but my conclusions will apply equally, with minor adjustments to accommodate differing democratic architecture, to any democracy which values the political autonomy of its populace.

I. Theory of Democratic Government

In a democracy, the authority of government, by which I mean the justified power to impose obligations upon the populace, derives from the consent of the governed (Locke, 1689, §171). Citizens of a democracy voluntarily surrender certain freedoms to the government in exchange for the benefits that government can confer, such as protection from violence and theft. This contractarian view of democratic authority is uncontroversially one of the founding principles that shaped the birth of American government (Declaration of Independence; Federalists Nos. 37, 51). The ongoing consent of the populace is maintained via regular elections and referenda, ensuring that the government remains responsive to the will of its citizens.

Vital to the proper functioning of a democratic system, and hence to its legitimacy,¹ is a “feedback loop” between the actions of government and the actions of the populace. While a

¹ By “legitimacy,” I mean here justification for imposition of obligation by the government, but see §IV, below, for a discussion of legitimacy as a scalar (rather than binary) property.

degree of stability is necessary for government to perform its role, responsiveness to the will of the people is the touchstone of democracy. The citizenry must be politically autonomous, which means that they must be able to change government when government fails to act as desired, and that they must have sufficient information about the government's activities to do so in a way that is likely to bring about the results they seek. In the absence of this critical feedback relationship by which the citizenry continuously corrects the course of government action, a democratic government can claim no authority. Indeed, a system in which the people either cannot affect or cannot know the actions of government could hardly be called a "democracy" at all.

Hence government itself potentially poses a threat to its own legitimacy. Maintaining the feedback relationship of course requires that the government not physically stop people from voting, but it has more subtle implications as well. Dictatorships may hold elections but so tightly control the available information, and thus the opinions of the people, that the results are effectively meaningless. Self-styled democracies may permit only a portion of the populace to be counted for political representation, as did antebellum America, or use intimidation (legal or otherwise) to control the vote. Voting districts may be drawn to artificially concentrate or dilute the political power of specific groups of citizens. These sorts of anti-democratic actions detract from the legitimacy of the resulting government by depriving its populace of political power or political autonomy. Because the sheer scope and power of government make it relatively easy to dramatically skew election results in a number of different ways, interference in the political process ought to be undertaken with great care to avoid political distortions.

A robust democratic government can, however, also intervene to help strengthen its connection to the people, and hence its own legitimacy. Government can encourage an

informed, participatory populace by, for instance, providing records and documents at its websites, or by making it easy for voters to register. Of course, it is critical to any such undertaking that the resulting regulations be politically neutral. While fine-tuning of the democratic architecture will inevitably result in some political power shifts over time, such shifts ought to come as a natural result of the strengthening bonds between populace and government in the form of providing information and opportunity for political expression, not as a result of manipulation of the system for political advantage.

Government could also take a more active role, going beyond simple promotion and invitation. Elections could consist of individual poll workers systematically finding every citizen, giving them courses in civics and economics, providing them with candidates' records, and demanding that the citizen vote or be severely punished. Of course, such a system is wildly inefficient. Thus there is some balancing to be done between resources expended on encouraging informed participation and those expended on the other functions of government, such as national defense or the court system. The duty of the government to ensure that it remains responsive to the populace, and vice versa, is but one obligation among many, albeit a particularly important one. Hence, the most compelling methods of promoting pro-democratic behavior are those which are cheap, simple, and easily implemented.

Adjustments to existing democratic infrastructure make excellent targets for efforts to efficiently strengthen the feedback loop between government and governed. Rather than establishing new programs or overhauling the current system, policymakers can look for opportunities to fine-tune what already exists. These sorts of efficiency-enhancing tweaks have been thoroughly embraced already in other contexts via the movement into economic analysis of

law, spearheaded by Richard Posner (1973).² By considering the economic interests of the parties involved in all manner of situations, the reasoning goes, policymakers can predict the parties' decisions and adjust them by altering incentives. In fields like contract law and tort law, where participants are often economically sophisticated, the approach has proven quite successful. In other contexts, however, there is reason to doubt its efficacy.

Classical economics is founded on an assumption of perfectly rational self-interest maximization which is increasingly being challenged by new research into systematic human irrationality. Loosely grouped under the umbrella of “cognitive bias,” these effects are predictable and statistically consistent distortions of judgment and decision making, when compared to the predicted behavior of perfectly rational self-interest maximizers. A system constructed on the assumption of perfect rationality will not function as designed when its components are irrational. Where such irrationality is predictable, it is incumbent upon the government to account for it in designing institutional infrastructure. In the next section, I shall discuss cognitive bias in more detail and offer several examples. Then, I shall set out three distinct problems cognitive biases pose for policymakers and, finally, offer guidelines for avoiding or minimizing the potential negative effects of policies which leverage cognitive biases.

II. Cognitive Biases

In the last few decades, there has been an explosive new wave of interdisciplinary research into complex human behaviors, characterized broadly as “behavioral biology” (Jones, 2005) or “the science of human nature” (Fowler & Schreiber, 2008). Leaps in computing, genetic analysis, and neuroimaging have opened new avenues of research which are challenging

² See <http://encyclo.findlaw.com/tablebib.html> for an extremely thorough overview of the history of the law and economics movement worldwide.

traditional assumptions about human behavior. Among the most interesting results is the revelation that humans are not nearly as rational as previously presumed. Instead, we seem subject to a host of psychological “shortcuts” which sidestep reflection and, in many cases, lead to seemingly irrational or non-optimal choices.

Researchers are beginning to see consistent patterns in exactly *how* people get things wrong, and they call these regularities cognitive biases. First identified in 1972 during studies on intuitive number representation (Kahneman & Tversky, 1972), cognitive biases are becoming increasingly important in our understanding of human behavior and decision making. Interest in cognitive biases is especially strong in the realm of economics, which classically relies on the assumption that agents in economic interactions act as rational utility maximizers, calculating expected benefits and detriments (usually in terms of monetary value), then selecting the course of action most likely to maximize the former and minimize the latter. However, research in many contexts has shown that this is simply not the case, and irrational decisions are not always the result of stupidity or poor information.

One example is the “status quo” or “default” bias, in which individuals are likely to accept a given default option, even where choosing a non-default option would confer significant economic advantage. In a seminal study, Madrian and Shea (2002) examined the correlation between enrollment in 401(k) programs and whether employees are enrolled automatically. Classical economics predicts that enrollment in a retirement program ought to be about the same at different companies, regardless of whether employees are included or excluded by default, because each individual will choose the option that confers the greatest benefit. Note that the process of enrollment or disenrollment imposes minimal cost, usually requiring only a form and perhaps five minutes’ time. Instead, Madrian and Shea found marked differences in enrollment,

with an overwhelming (up to 66%) preference for the default option, whether it be enrollment or exclusion (Madrian & Shea, 2002, pp. 1162-4), and other experiments have since duplicated these results (Choi, 2002; Choi *et al.*, 2003). Thaler and Sunstein (2008, p. 35) attribute this difference to what they call the “yeah, whatever” heuristic, whereby many individuals simply do not bother with even miniscule costs if the status quo does not demand immediate attention.

Another example of cognitive bias is “loss aversion,” in which individuals react more strongly to losses than to equivalent gains, and the closely related “endowment effect,” in which an individual requires a higher price to part with an object than he would pay to obtain it, despite classical economics’ prediction in both cases that the object ought to have the same value whether one possesses it or not. Kahneman *et al.* (1991), for instance, performed an experiment targeting the endowment effect in which half the participants were given a coffee mug adorned with the logo of the university. The other participants were asked what prices they would be willing to pay to acquire a mug, while those who had a mug were asked what payment they would accept to part with it. Over many trials, participants consistently required roughly twice as much to give up a mug as they would be willing to pay to acquire one (Kahneman *et al.*, 1991). Interestingly, the endowment effect has since been demonstrated in both chimpanzees (Brosnan *et al.*, 2007) and capuchins (Chen *et al.*, 2006), suggesting that it is deeply entrenched in human psychology.

Yet another example of cognitive bias is the “framing” effect, in which the way information is presented can strongly influence the choices individuals make. A paradigmatic study is Redelmeier *et al.*’s (1993) paper on patients’ decisions about treatment. Redelmeier and his colleagues presented two sets of individuals with different but logically identical descriptions of treatment efficacy and then asked them to choose whether to undergo the treatment. The first

group was told that 90% of those who had received the treatment were *alive* after five years. The second group was instead told that 10% of those who had received the treatment were *dead* after five years. Perhaps unsurprisingly, patients to whom the information was presented in terms of death rate were much less likely to undertake the procedure than those to whom the exact same statistical information was presented in terms of survival rate (Redelmeier, 1993).

Still another example of cognitive bias is “anchoring,” in which those who are asked to estimate a figure are heavily influenced by a suggested starting point. A favorite experiment of Thaler and Sunstein (2008) is to ask their students to take the last three digits of their phone numbers, add two hundred, and then estimate the year in which Attila the Hun sacked Europe.³ Clearly the number generated from the student’s phone number is completely unrelated to world history, yet Thaler and Sunstein consistently find that students with high starting points will estimate the event to be roughly three hundred years later than students with low starting points (2008, pp. 23-4). Similarly, attorneys and other negotiators have known for quite some time that it is wise to initially demand a far larger sum than one actually expects to receive, especially when arguing before a jury of laypeople who may not know how much the normal award usually is. By starting with an astronomically high number, the negotiator can make large sums demanded later seem reasonable by comparison to the anchor amount, resulting in a better deal than he would have gotten had he initially asked for exactly as much as he hoped to receive.

Though this list is far from exhaustive, the final example of cognitive bias I shall discuss is the effect tactile experience can have on judgment and perception, which I shall call the “haptic bias.” Ackerman *et al.* (2010) find evidence that the sense of touch can affect decisions in six different kinds of complex social interactions, three of which I shall discuss briefly here. In one experiment, subjects were asked to evaluate a job candidate based on a resume which was

³ Thaler and Sunstein (2008) assure the reader the correct answer is 411 A.D.

given to them on either a heavy clipboard or a lighter one. Candidates presented on heavy clipboards were consistently considered more serious about their application for the job and better qualified. In another, people were asked to rate the ease of a social interaction after assembling a puzzle composed of either rough or smooth material. Social interactions rated after the rough puzzle were reported as being generally more difficult. In a third, negotiators seated in hard chairs drove harder bargains than their soft-chaired counterparts (Ackerman et al., 2010, pp. 1713-4).

These and other cognitive biases present significant opportunities for government to non-coercively, and perhaps drastically, increase efficiency with comparatively minor alterations to existing institutional infrastructure, but the undertaking is not without its complications. Cognitive biases could be leveraged in inefficient ways or, worse, in ways that harm the political autonomy of the citizens. In the sections that follow, I shall set out three problems faced by policymakers who wish to use cognitive biases to advance government interests and offer some guidelines to consider when evaluating whether particular uses of cognitive biases are problematic.

III. Three Problems in Governmental Leveraging of Cognitive Biases

Leveraging widespread psychological predispositions is a subtle, often entirely subliminal, method of altering behavior and thus seems, at least to some degree, to subvert the individual autonomy on which democracy so critically depends. One can fight consciously perceived attempts at psychological manipulation and even violent coercion, but a cognitive bias which operates without the knowledge of the individual affected can be literally irresistible, at least at the conscious level. Hence we must consider when, if ever, such leverage is appropriate

in a government which respects the political autonomy of its citizens, as a democratic government must. There are three distinct problems with which policymakers must grapple when crafting laws and regulations which take advantage of cognitive biases.

The first problem we face is that we cannot unlearn what we have discovered and return to our faulty assumptions of perfect human rationality, yet we also cannot indiscriminately manipulate citizens' cognitive biases without running the risk of seriously diminishing individual autonomy. A cognitive bias by definition acts without the knowledge of the affected individual (and sometimes in spite of the individual's knowledge, as the bias blind spot shows; see §V, below), so when policymakers consider using cognitive biases to alter people's behavior, they are unambiguously considering intentional manipulation of individual decisions. Moreover, it is not overt manipulation of the sort which government ordinarily uses to accomplish its ends. Instead, tweaking cognitive biases is a very subtle method of accomplishing desired behavior.

Such action falls into an uncomfortable grey area between coercion and promotion, leaving its role in governmental regulation ill-defined. We know, for instance, that the American government may not force any individual at gunpoint to exercise regularly, but we also know that the government may make it clear that regular exercise is a good idea through public service announcements, school programs, and the like.⁴ Governmental inducement of subconscious fear to promote regular exercise, however, would not fit cleanly into either of these categories. Hence we must be cautious in implementing policies which leverage cognitive biases, lest the government overstep its authority.

Yet the government also cannot ignore the discoveries made in the emerging sciences of human nature without abdicating its responsibility to govern effectively, and simply as a

⁴ Politically attuned readers and parents of young children will know already that the current First Lady, Michelle Obama, has made the fight against childhood obesity a top priority. See <http://www.letsmove.gov/> for an example.

practical matter, the fact that such tools are increasingly available puts increasing pressure on regulators to adopt techniques which use them. Already, the current Presidential administration is using cognitive biases in crafting administrative rules. The 2009 budget, for instance, included a provision adjusting retirement savings procedure to enroll individuals automatically, thus leveraging the default bias in exactly the way Madrian and Shea (2002) suggest.⁵ Indeed, Cass Sunstein himself, one of the primary architects of the bias-leveraging approach, was in January of 2010 appointed head of the Office of Information and Regulatory Affairs, a subdivision of the Office of Management and Budget which is tasked with promoting administrative efficiency.⁶ Hence, we face the dilemma of knowledge: government can neither ethically nor practically ignore the tools presented by cognitive biases, yet neither can it use such tools without regard for the effects on the autonomy of the citizenry.

Further, it is often simply impossible to avoid triggering biases in one way or another. Many government programs require, for instance, that there be some default state which everyone is presumed to occupy until they act to change it, even if the default is negative, like *not* being an organ donor or *not* being registered to vote, as is generally the case in America. Once an entitlement program is in place, its beneficiaries will fight attempts to rescind their benefits with tenacity strengthened by loss aversion. Government offices must have furniture, and whether such furniture is hard or soft will affect the attitudes and responses of people in that environment. In short, any attempt to refrain entirely from triggering cognitive biases is doomed to failure.

⁵ See Foer & Scheiber, 2009 for both an overview of current efforts and a very thorough political history of the subject.

⁶ For more information, see OIRA's website, http://www.whitehouse.gov/omb/inforeg_default/

Thaler and Sunstein (2003) compare this problem to that of a cafeteria worker who just happens to notice the eating habits of the students. She notes that the majority of the kids will take cake, pudding, or other sweets if those come first in the dessert section of the line, but the majority will choose healthier options like fruit if they come first instead. She knows that the sweets are bad for the kids, and she knows an easy, cheap action she can take to make more of them choose the healthier option. Hence, there are three available courses of action. First, the cafeteria worker can put the fruit at the front of the dessert section, thus encouraging healthier eating. Second, she can put the sweets in the front of the dessert section, encouraging unhealthy eating. Third, she can alternate or randomize the order (Thaler & Sunstein, 2003, p. 175).

Note that none of these courses of action selects a dessert *for* anyone. Every individual student is still completely free to select the dessert he finds most appealing, healthy or not. Instead, the cafeteria worker can alter the statistical distribution of student choice through a minor change in presentation. Some will choose differently on different days, and inevitably some will choose their preferred option, healthy or not, no matter how the desserts are arranged. What the cafeteria worker can change are the choices of those without strong preferences either way, the pliable middle. Thus she can influence without coercing.

Thaler and Sunstein conclude that the only ethically acceptable option for the cafeteria worker is to put the fruit first. There is no way to avoid influencing what the students will select because the desserts must go in one order or another;⁷ the cafeteria worker can only select *how* to influence them. Putting the unhealthy desserts first would be outright malicious, intentionally inflicting harm on the students. Randomizing or alternating, on the other hand, is simply

⁷ Though Thaler and Sunstein expand on this scenario throughout their later work *Nudge* (2008), including options like setting up a separate table for unhealthy desserts or charging a bit extra for them, the average cafeteria worker (and, by analogy, policymaker) rarely has control over that sort of comparatively large-scale organization. Because this discussion focuses on cheap, simple actions, the various embellishments Thaler and Sunstein add to their hypothetical in *Nudge* are not explored here.

abdicated responsibility for the students' nutrition. Analogously, then, the government has an ethical responsibility to use knowledge of behavioral regularities to structure laws and policies in a way that efficiently promotes the wellbeing of the citizenry, especially where it can be done cheaply and easily by making small alterations to existing institutional infrastructure (Thaler & Sunstein, 2003, pp. 176-7). Thus, Thaler and Sunstein argue that if a government knows that it can use the default bias to encourage saving for retirement simply by changing the way a regulation is implemented, it is ethically obligated to do so. They call this approach "libertarian paternalism" (Thaler & Sunstein, 2003, p. 176).

There is, however, a second problem. If our preferences are so much a result of these minor quirks of environment or preceding stimuli, it is uncertain whether pre-existing preferences can be said to exist at all. Ideally, a democratic government ought to respond swiftly and efficiently to the desires of its citizens, but if those desires are inchoate or easily manipulated, it becomes much less clear to what extent the government is obligated to respond to or avoid affecting them. As an example, recall Redelmeier's (1993) work with patient preferences. Patients were much more likely to express a desire to undergo a medical procedure when presented with information about its efficacy in terms of *survival* rate, as opposed to the complementary *death* rate statistics. A doctor who herself had high confidence in the procedure and thought her patient ought to do it could choose to present the information in terms of survival, thus subtly nudging the patient towards the preference she feels is in the patient's best interest.

Hence, cognitive biases can actually change what people *want*, whether it be a decision about medical care or government. As an example, consider the tumultuous history of the inheritance tax. When an individual with an estate worth more than a certain amount dies, the

government assesses a tax on the estate, known historically as “death duties” (see, e.g., Beatty, 1907). In the middle of the 1990s, when the inheritance tax again became an issue of political contention, savvy anti-tax activists mounted a campaign against it which included relabeling it the “death tax.” As a result of this effort, in 1997 the government enacted a law progressively raising the value of estates to which the tax applied and lowering the rate at which the estates were taxed.⁸ By framing the inheritance tax as a penalty for dying rather than a tax on large inheritances, skilled political manipulators were able to dramatically increase public opposition. While there were of course a great many factors contributing to the success of this effort, America’s struggle with the inheritance tax is an excellent example of the framing effect being used effectively to shape citizens’ preferences in a pitched political battle. The government can use this and other cognitive biases to shift preferences in other contexts as well.

This leaves room for what Thaler and Sunstein call “choice architects” (Thaler & Sunstein, 2008, pp. 11-13), those who are responsible for structuring the way in which decisions are presented, to tweak not merely behavior but actual preferences. While Thaler and Sunstein argue that such structuring of decision making is non-coercive, it is not clear that they are correct. In fact, Redelmeier’s (1993) focus was to draw attention to serious issues of medical ethics raised by how a doctor chooses to present information to a patient when seeking informed consent. In the same way that a patient’s consent to a medical procedure may lack legitimacy when given as a result of preferences manipulated by a doctor, so too can a populace’s consent to a set of laws lack legitimacy when given as a result of preferences manipulated by the government. It would, for example, be inappropriate for a democratic government to make the

⁸ Currently, in 2010, the estate tax is actually not in force as a result of this odd compromise bill and its successors. In 2011, the tax will return. See Robbins, 2004 for a historical account and an overview of the changes implemented in 1997 and subsequent legislation.

same sort of lexicographical change popularized by opponents of the inheritance tax because doing so would alter preferences in a way that advances a distinct political agenda,⁹ about which more in the sections that follow. Policymakers must be very careful that they do not use cognitive biases to interfere with citizens' preferences in a way that diminishes political autonomy, a danger Thaler and Sunstein neglect to consider in any depth.

The third and final problem is that virtually all actions can be characterized as economic actions, and though Thaler and Sunstein recognize the breadth of costs which can be inflicted by regulators following the libertarian paternalist approach, they fail to provide criteria by which to judge whether a given burden on individual choice is impermissibly non-libertarian. They characterize their theory as acceptably libertarian because it does not block or close off options, but as a practical matter, it is not possible for government to completely block an option, only to impose a very high cost for selecting it. We could, for instance, characterize the investment of time and energy in robbing a bank, as well as the attendant risks of prison and death, as economic expenditures on choosing a governmentally disfavored resource allocation. The law does not (and, indeed, cannot) absolutely block or close off this option but instead simply places so great a cost on bank robbing that it is, economically speaking, an uncertain and spectacularly inefficient undertaking.

Similarly, it is conceivable that the government might make it so difficult to choose a course of action it wishes to discourage that it would effectively require that everyone "stick to the plan." For instance, choosing not to participate in a 401(k) program might require a trip to a single understaffed office in Washington, D.C. and a \$100,000 filing fee. Perhaps policymakers

⁹ Although this sort of political influence often creeps into government policy anyway. See, for example, former President George W. Bush's June 8, 2006 Statement of Administration Policy, which refers to the inheritance tax as the "death tax" nine times in five paragraphs. Available online from the American Presidency Project at <http://www.presidency.ucsb.edu/ws/index.php?pid=24912>

might even require a period of incarceration, or characterize the filing fee as a criminal fine. Indeed, Thaler and Sunstein acknowledge that the degree to which regulations are “libertarian” exists along a continuum, and “there is a thin line between non-libertarian paternalists and libertarian paternalists who impose high costs, procedural or substantive, on those who reject the plan” (Thaler & Sunstein, 2008, p. 30). In the sections which follow, I shall offer some more concrete guidelines for policymakers to consider when evaluating whether a proposed policy crosses this line.

Hence we have three problems. First, government cannot ignore research on cognitive biases without abdicating its responsibility to govern effectively, yet it must also avoid using these tools blindly lest it seriously damage the autonomy of its populace and, hence, its own legitimacy. Second, cases where cognitive biases make it reasonably easy to manipulate actual preferences, in addition to mere behavior, force us to question to what extent these preferences exist at all independently of the environmental circumstances which shape them, further complicating the problem of leveraging cognitive biases without unduly influencing the citizenry. Finally, because leveraging cognitive biases can impose a range of costs for noncompliance, policymakers must be careful not to overreach, lest libertarian paternalism become paternalism outright. In what follows, I shall offer a special justification for using cognitive biases to promote pro-democratic behavior and set out guidelines to help policymakers avoid these problems when crafting laws and regulations in other domains as well.

IV. Special Justification for Using Cognitive Biases to Promote Pro-Democratic Behavior

Policymakers are especially well-justified in using cognitive biases to promote pro-democratic behavior. This is not only because it is precisely the potential weakening of

democracy that makes leveraging cognitive bias an unusually serious threat to a self-governing nation but also because humans are subject to a host of psychological predispositions which continuously threaten to undermine egalitarian self-rule. Human psychology is not well-adapted to promote democracy on a large scale. In fact, our psychology is not adapted to most forms of large-scale social organization because early hominid societies simply did not have the means to reach the requisite size. Until quite recently, evolutionarily speaking, human group membership rarely climbed above roughly 50 individuals (Dunbar, 1993, p. 681), meaning there was virtually no selective pressure in favor of nation building, let alone international and global organization. Large societies simply conferred no advantage; in fact, a very large society can prove disastrous, as evidenced by the collapse of the Mayan and Easter Island civilizations, among others. Nonetheless, the basic human instinct for group living, combined with a complex system of cultural exchange and ongoing technological advances, has pushed our numbers ever upwards towards these sorts of large-scale organizations.

Hence, though natural selection has strongly affected our social predispositions, there is no reason to believe that it has favored democratic organization. To the contrary, Somit and Peterson (2005) argue persuasively that because our evolutionary heritage comes to us from a long line of hominids living in small groups following the dictates of an alpha individual, natural selection has for tens of millions of years favored the psychological traits (whatever those may be) that promote power-concentrating social organization. This makes authoritarian government the “default option” (Somit & Peterson, 2005, p. 9). This theory has significant underpinning in the relevant literature on comparative psychology; it has long been known that primate social structure is generally hierarchical, often with a single individual at the apex of the hierarchy (see Bernstein, 2004; for macaques see Smith, 1993; for baboons see Hausfater, 1982; for bonobos

see Furuichi, 1997; but see de Waal, 1989 on the internal fluidity of this hierarchy in chimpanzees). It is thus reasonable to infer from such widely shared traits in our near evolutionary relatives that a similar set of social predispositions exists in humans, and indeed, social psychologists have identified some already.

Most relevant here is a strong tendency towards obedience to authority, even where an individual may be uncomfortable with the task assigned. The paradigmatic example is Milgram's (1974) study on obedience, in which subjects were told they were participating as "teachers" in an experiment on learning which required them to administer electric shocks to a "learner" (actually a member of the research team) in another room. When the learner got the answers wrong, the teacher was instructed by an experimenter (the authority figure) to increase the voltage. Though no one was really being shocked, the learner would cry out as though in more pain each time the voltage increased. Contrary to the expectations of Milgram, as well as three expert panels asked to predict the subjects' behavior, 26 of the 40 subjects were willing to proceed all the way to the maximum voltage, labeled "XXX Danger: Extreme Intensity Shock." Milgram notes that subjects "were observed to sweat, tremble, stutter, bite their lips, groan, and dig their fingernails into their flesh" (Milgram, 1974, p. 375), yet they continued. Researchers have shown similar effects in other contexts as well, such as Frank's (1944) study which showed that experimenters could lead subjects to consume great quantities of soda crackers, or Sheridan and King's (1972) study in which more than half participants were willing to shock a defenseless puppy at the maximum voltage when instructed by experimenters to do so.

Further evidence for Somit and Peterson's (2005) theory is provided by the tendency of individuals to follow the crowd, even where the crowd is clearly wrong. Asch (1956) presented evidence that subjects would make serious errors in estimating the length of a line if they heard

other “subjects” (experimenters in disguise) make the error first, even though they consistently performed well at the task in isolation. Sunstein (2003) documents comparable results from 17 other countries, including nations as diverse as Germany, Japan, Kuwait, and Zaire. In a similar experiment, Sherif (1934) exposed subjects in a darkened room to a point of light which appeared to move due to an illusion called the autokinetic effect. When asked individually to estimate the distance the light had moved, subjects’ answers were of course quite different because the light did not actually move. When asked to estimate the distance in small groups, however, subjects quickly converged on common estimates (though estimates varied widely between different groups). Further, Sherif’s subjects would stick to the group answers in later experiments *even when asked to estimate the distance without group input*. Thaler and Sunstein (2008, p. 55) offer many other examples of peer-promoted conformity, including a tendency for teenage girls who see that their peers are having children to get pregnant themselves and a propensity for college students to mimic the study habits of their roommates.

Thus, maintaining a democracy is a constant battle against “baser instincts” favoring homogeneous obedience to authority. In the complex system of human cultural exchange, these predispositions can favor a variety of governmental forms, from dictatorship to oligarchy to plutocracy, but for simplicity’s sake, I shall refer generally to such forms of government as “autocracy” to highlight the concentration of political power which distinguishes them from democracy. While autocratic propensities may have promoted survival and reproduction millions of years ago in the so-called “environment of evolutionary adaptation,”¹⁰ we must now fight them to maintain the organizational system we have come to prefer. The mere fact that an autocratic system helped our ancestors survive does not confer political legitimacy upon it. Our

¹⁰ See Tooby & Cosmides, 2005 for a more thoroughgoing introduction to the significance of evolutionary history in the study of cognitive biases.

sense of fairness, along with centuries of complex social and political debate, has led us to value self-government in which each individual has an equal vote, and we have, over time, developed sophisticated systems of governance which suppress, redirect, and channel instincts favoring autocratic organization.

The American model of democracy, on which the bulk of this analysis focuses, is the height of Enlightenment conceptual engineering. Far from instinctive, the American Constitution is a complex product of rational debate and meticulous revision over hundreds of hours, informed by years of experience under the ultimately failed Articles of Confederation. Democracy has changed and developed over time to match our notions of what we want from government, such as extending voting rights to those who do not own land, but at base, we are still committed to egalitarian self-rule. Thus we desire a form of social organization different from that to which we are most strongly psychologically predisposed.

Democracy is justified in undertaking the functions of government because the authority of a democratic government is derived from the consent of the governed (Locke, 1689, §171), a sharp move away from the autocratic default. Both rewards for prosocial behavior and punishments for antisocial behavior are set by the populace to whom they will apply (albeit indirectly, in the case of American-style representative democracy at issue here), thereby justifying the lawful exercise of coercive power by the state. However, the entirety of a given nation's populace never participates fully in every vote, all the time.¹¹ Inevitably, there are many who are waylaid by other obligations or who simply do not bother. Further, every voter necessarily makes her decision based on limited information, consequently limiting her ability to

¹¹ The American Presidency Project calculates participation in Presidential elections since 1972 to be between 49% and just over 55% of eligible voters. Even the highest turnout in American history, during the election of 1876, was only 81.8%, leaving out nearly a fifth of those able to vote. See <http://www.presidency.ucsb.edu/data/turnout.php>

produce the outcome she most prefers. In short, the self-imposed governance of a democracy derives its authority from an abstract principle of well-informed popular participation which is, in reality, never fully achieved. Thus the legitimacy of a democratic government exists along a continuum defined by how well it approximates this ideal.

Democratic government, then, is uniquely justified in using cognitive biases to promote pro-democratic behavior because such use helps counteract our inherited autocracy-promoting tendencies, facilitating informed participation and, hence, bolstering the legitimacy of the resulting government, all without unduly or coercively interfering with the political autonomy of the citizenry. We can, in effect, “fight fire with fire” by using cognitive biases to blunt the effects of psychological predispositions which Somit and Peterson (2005) fear might undermine democracy. Take the default bias as an example. A tendency to accept the default is a tendency to acquiesce to the authority of those who set the default,¹² which could easily contribute to the perpetuation of autocratic rule. In America, one must register in order to vote, making nonregistration the default and hence encouraging passive acquiescence to the political status quo.¹³ The information needed to register a voter is quite basic and the form takes very little time, yet almost 30% of America’s voting-age populace is currently unregistered (U.S. Census Office, 2010, p. 1). This raises a barrier to political participation and ultimately diminishes the legitimacy of the government. Voter registration drives may help, but any such undertaking

¹² Indeed, the electoral advantages of incumbency are well known to political scientists. See, e.g., Erikson, 1971. Some jurisdictions even go so far as to note on the actual ballot which candidate is the incumbent.

¹³ Unless one is a resident of North Dakota. See <http://www.nd.gov/sos/electvote/voting/vote-history.html> for a history of this interesting legislative quirk.

necessarily must fight the inertia of the default bias because, by default, citizens are unregistered.¹⁴

Here, then, is an excellent opportunity to fight fire with fire by leveraging the default bias the *other way*. Policymakers could, for example, automatically register voters when collecting information for other purposes, such as state-issued identification, and provide an opt-out for those who prefer to remain unregistered. While privacy advocates and small government activists could legitimately complain that compulsory registration is intrusive, inclusion of an easily chosen opt-out leaves an alternative for those who prefer to remain unregistered while nonetheless increasing registration across the populace as a whole. Leveraging bias is an especially attractive technique to increase voter participation because it avoids the “noise” problems of compulsory voting, in which voters who are merely avoiding a penalty rather than expressing legitimate political preferences¹⁵ can skew election results. Those who would otherwise check random names on the ballot are free to go about their business as usual on election day if they like.

Importantly, government can also help diminish the effects of biases which affect political decisions. Druckman (2001) offers evidence that expert advice and discussion can help to mitigate framing effects, which is particularly important in an era when politicians are increasingly adopting the techniques of advertisers who have honed framing manipulation to a fine, and very profitable, art. The government can and does maintain organizations of independent experts, such as the President’s Council on Bioethics or the National Oceanic and

¹⁴ Note that it is not voluntary non-participation which detracts from the legitimacy of the government, but rather it is systemic barriers to participation. Of course, individuals must be free to abstain from the political process if they so desire, but being registered to vote does not require that one vote; it merely provides the opportunity to do so. Being unregistered, on the other hand, denies a citizen even the opportunity to participate. Hence, barriers to registration detract from democratic legitimacy in a way that voluntary non-participation does not.

¹⁵ Known to compulsorily voting Australians as a “donkey voter,” http://en.wikipedia.org/wiki/Donkey_vote

Atmospheric Administration, to advise both government and populace on matters of political importance, which, in addition to keeping the public informed, helps mitigate framing effects which might be imposed by other sources. As a somewhat simpler example of cognitive biases the government can mitigate, Koppell and Steen (2004) offer evidence that simply being listed first on a ballot gives a candidate an advantage of about 3.5% of the vote (though the effect diminishes significantly when the candidates are well known). The government could, at relatively little added expense, produce multiple versions of ballots to balance this small deviation. Indeed, the increasing prevalence of electronic voting reduces the cost of including alternate ballot designs to nearly zero, making this an excellent target for efforts to efficiently diminish undesirable bias effects.

Hence, democratic government is most strongly justified in using cognitive biases when it acts to strengthen its bond with its citizens because it is fighting fire with fire, leveraging psychological propensities in novel ways to moderate the effects of other, undesirable psychological propensities. However, the fact that the government is uniquely justified in leveraging cognitive bias in the context of promoting pro-democratic behavior does not give it *carte blanche* to indiscriminately manipulate political behavior generally. Note that it is critical that the government refrain from manipulating voter viewpoint or voting tendency via cognitive biases or otherwise, about which more in §VI. Note also that the political and efficacy considerations set forth in §V below ought to be weighed in this context as well, though a pro-democratic purpose may tip the scales in favor of permitting the use of cognitive biases where such use might otherwise be deemed imprudent.

V. Conditional Justification in Domains Traditionally Reserved to Government Authority

In domains traditionally reserved to government authority, such as public safety or enforcement of contracts,¹⁶ we find the most problematic situation because use of cognitive biases has no clearly defined place in the policymakers' toolbox. In the context of governmental interests justifying violent coercion, we are faced with the question of whether using cognitive biases constitutes prevention or punishment. If the former, then such use falls cleanly into the same category as neighborhood watches and anonymous tip lines. If the latter, such use may constitute a violation of the Due Process Clause of the Fifth Amendment, punishing individuals neither convicted nor even accused of a crime. In the context of governmental interests which justify only encouragement, we must determine whether using cognitive biases constitutes promotion or coercion. These distinctions are unfortunately muddled further still by the youth of the relevant sciences. We can, however, begin to formulate some guidelines going forward, which I shall separate into issues of politics and issues of efficacy.

With regard to issues of politics, remember that problems arise in leveraging cognitive biases to accomplish governmental goals because such biases are involuntary and routinely subconscious. Therefore, the first condition which must be met in using cognitive biases to advance governmental goals is transparency. Transparency in a democracy generally helps ensure that governmental action really is the intent of the populace, but transparency is especially important where the governmental action being undertaken is manipulative at the unconscious level. Indeed, researchers have identified a cognitive bias *against recognizing that one is being or has been affected by a cognitive bias*, known as the "bias blind spot" (Pronin & Kugler, 2006).

¹⁶ An exploration of exactly which domains have traditionally been considered within the purview of the American government's authority is far too involved and technical a task to be undertaken here, but most of the relevant domains are so well-established as to be obvious, such as protection of property interests or national defense. Legal analysis of historical precedent is the appropriate avenue of inquiry here, though see §VI, below, for a more detailed discussion of precedent as a criterion for evaluating bias-leveraging regulation.

Thus it is critical to communicate clearly to the populace what sort of bias is being leveraged, how it is being done, and what purpose the government seeks to advance by doing so.¹⁷

For example, policymakers might decide that furniture intended for visitors to IRS offices ought to be soft, using the haptic bias to nudge people towards simply paying the expected amount when they might otherwise contest small details. While this is likely to have only a small effect, anti-tax advocates or civil rights watchdogs could reasonably complain that such a regulation inappropriately distorts the perceptions of taxpayers visiting the offices by causing some of them to pay more in taxes than is legally required. No matter what the ultimate decision may be, the policy ought to be subjected to the scrutiny of the voting public, as with virtually all policies of a democracy.¹⁸ Transparency is especially important here, however, because the average IRS visitor is probably unlikely to notice that the chairs are particularly soft and even less likely to associate that fact with a willingness to accept less favorable terms in negotiations. Without governmental transparency, leveraged cognitive biases could easily be all but invisible to the populace, leaving them with no opportunity to shape the relevant policies.

Second, when deciding how to use a given bias, policymakers also ought to consider the ease with which a given bias can be overcome by the affected individual, and a useful metric here is the context in which the decision to be affected is being made. Fleeting interactions which would ordinarily command little attention are simultaneously the least likely context in which bias effects will be noticed and the context in which they may have the greatest impact. Some biases, however, such as the anchoring effect, remain effective even in reflective decision

¹⁷ Interestingly, the very blind spot which makes it so important that government acknowledge leveraging of cognitive bias also ensures that such transparency will not significantly detract from the efficacy of the program. Knowing that one is in a situation which will trigger a cognitive bias usually does little to prevent it from being triggered. Again, see Pronin & Kugler, 2006.

¹⁸ I say “virtually all policies” because there are a few narrow exceptions to this rule, such as sensitive national security information. The complex question of precisely which policies may legitimately be withheld from citizen review is a problem left unexplored here.

making. Thus the context in which a bias is expected to act must be weighed in deciding whether to go forward. Generally, easily overcome biases active in a context permitting reflection and consideration may be used more freely than entrenched biases active in a context requiring swift action.

As an example of the importance of context, consider the opposing contexts of a bias leveraged in automobile traffic and, alternatively, the same bias leveraged in the way a government form is structured. In traffic, the affected individual is processing a great deal of information at one time, much of it critical to avoiding injury or death. There are few cognitive resources available for a driver to dedicate to analyzing potential nonconscious psychological effects, and little time in which to do so as well. Filling out a form, however, provides more opportunity for reflection and analysis, which will generally give individuals a better chance to identify and compensate for potential bias effects, though the leveraged bias may nonetheless affect them. Consider also the framing effect as compared to the default bias as an example of the importance of bias entrenchment. The framing effect, as demonstrated by Redelmeier *et al.* (1993), remains effective even during the sort of reflective decision making undertaken by patients whose lives are at stake. The default bias, on the other hand, can be overcome through careful consideration, as demonstrated by the many people Madrian and Shea (2002) found to have enrolled in 401(k) programs even where disenrollment was the default. Thus, the ease with which an ordinary individual may overcome a particular bias leveraged in a given context ought to be considered when crafting policy.

As a final political matter, government must consider whether a fundamental right¹⁹ is implicated by a given use of cognitive biases. As with any other technique, a significant burden on the exercise of a fundamental right would be an impermissible impact of a regulation using cognitive bias. For instance, policymakers might leverage the status quo bias to discourage new gun purchases in the name of protecting public safety, but the Supreme Court has over the last few years increasingly moved towards explicit recognition of the right to bear arms as a fundamental right (see District of Columbia v. Heller, 554 U.S. ____ [2008] and McDonald v. Chicago, 554 U.S. ____ [2010]) and could reasonably take exception to a law leveraging bias to burden that right. Any such regulation would have to be, in line with Supreme Court jurisprudence, narrowly tailored to advance a compelling government interest (see, e.g., Gonzales v. O Centro Espirita Beneficente Uniao do Vegetal, 546 U.S. 418 [2006] for a recent articulation and application of this “strict scrutiny” in the context of free exercise of religion). Cognitive biases emphatically should not be used to expand government power into areas traditionally beyond its reach, about which more in §VI, below.

Matters of efficacy are to be judged by the relevant sciences. While it would be premature and unwise to set a definitive threshold for use in advance of the relevant research, it is safe to say that a given bias must first be scientifically well established before the government ought to make use of it (or attempt to mitigate it). This guideline helps avoid unintended consequences of regulation, as well as expenditures on institutional infrastructure changes which have little effect. Again, the duty of the government to leverage cognitive biases in constructing laws hinges on efficiency. If there is little evidence that leveraging a given bias increases the effectiveness of government intervention, there is of course no duty to use it.

¹⁹ “Fundamental right” is a term of art in American jurisprudence, referring generally to rights found in the Bill of Rights. Among them are the right to free speech, the right to free exercise of religion, and, notably, the implicit right to privacy.

As a second matter of efficacy, leveraged cognitive biases also ought to meet a minimum threshold of reliability. Obviously it is likely to be of extraordinarily little utility to leverage cognitive biases which only work 1% of the time, even if such biases are well-supported in the literature. Leveraged biases also must be widespread in the population if they are to serve as useful regulatory tools, which is one of the reasons evolutionary psychology is so often linked to the study of cognitive biases. Evolutionary psychology finds deeply entrenched predispositions that are likely to be very widespread in the population due to many, many years of selection for the trait.²⁰ If a bias acts rarely or is present in only a small portion of the population, policymakers usually ought to ignore it merely as a matter of efficiency.

Thus there are five factors policymakers should consider when leveraging cognitive biases in domains traditionally reserved to government authority. As political matters, policymakers must be transparent about their use of cognitive biases, consider the ability of individuals to compensate for the bias being leveraged, and refrain from burdening fundamental rights even coincidentally. Unless these conditions are met, the government runs the very real risk of diminishing the political autonomy of its citizens. As matters of efficacy, policymakers must be sure that a given bias is scientifically well established and also that it acts regularly enough in a large enough portion of the populace that it can be usefully employed to accomplish the desired result. Without these assurances, the government could easily waste time and money on inefficient actions or, worse, unwittingly damage the citizens' political autonomy by leveraging biases which are poorly understood. In what follows, I shall discuss the limits of leveraging cognitive biases and the duty of the government to re-evaluate existing policies which trigger cognitive biases.

²⁰ Again, see Tooby & Cosmides, 2005 for a more thoroughgoing introduction to the significance of evolutionary history in the study of cognitive biases.

VI. Unjustified in Other Domains & the Duty to Re-Evaluate Existing Biases

Governmental use of cognitive biases to affect behavior which is not within domains traditionally reserved to government ought to be broadly prohibited to prevent government overreaching. Leveraging of cognitive biases is not a tool for the government to extend its power beyond its current scope, nor should it be. Actions prohibited to government through ordinary means remain prohibited in this context as well, and bias-leveraging regulations affecting citizens' political autonomy are the most objectionable. It is not difficult to envision the use of cognitive biases to subtly and systematically disadvantage a particular political ideology or social class. The anchoring effect, for example, could be used to manipulate the tax burden economically unsophisticated citizens are willing to impose upon themselves by shifting the range of debate up or down. In an outrageous power grab, the ruling party might leverage the default bias in their favor by requiring that all uncast votes be counted as votes for their candidates. More subtly, agencies making information about government actions available to the people could leverage the framing effect to alter popular perception of the data to the advantage of some political faction. Taxes might be presented as yearly gains (refunds) rather than monthly losses (withholding), leveraging loss aversion to diminish opposition to increases. The possibilities for abuse are vast, and they will continue to grow as research goes forward.

Note however that leveraging cognitive biases may ultimately prove so powerful a regulatory tool that government will be able to undertake functions previously thought simply impossible. The scope, pervasiveness, and subtlety of cognitive bias may be brought to bear on problems previously unsolvable yet nonetheless recognized as within the purview of legitimate government authority. Subject to the considerations set forth in §V, above, we ought not prevent

policymakers from pursuing important and formerly impossible goals simply because the government must use the new tools of cognitive bias to accomplish them.

For instance, the government could permissibly leverage cognitive biases towards the goal of completely eliminating transmissible diseases under the state power to protect public health, safety, and welfare. The public health power of the government has long been recognized by the judiciary as extraordinarily broad (see Hodge & Gostin, 2002 for a historical overview of the public health power in the context of mandated vaccines), and though such a task has never been undertaken because of practical limitations, policymakers would not be overstepping by using the tools of cognitive bias in the attempt.²¹ Government may take on unprecedented tasks which are nonetheless within the traditional understanding of government authority, and this should be distinguished from an undertaking which seeks to effect an outcome which is itself outside the traditional scope of government authority.

As a final note, knowledge of cognitive biases imposes a duty on the government to avoid biases where it is not explicitly using them for some permitted purpose and to compensate for cognitive biases where they cannot be avoided. There are likely a great many regulatory and administrative elements which trigger biases simply because the architects did not have the knowledge to avoid them, and as research reveals these effects, policymakers ought to set about correcting them. Biases of course ought to be optimized when, as a result of simple ignorance in administrative design, they are active in non-optimal ways within permissible domains, but there

²¹ Regrettably, I have not devised a way to eliminate transmissible diseases using cognitive biases. I have selected this example for two reasons. First, the power to protect public health, safety, and welfare has a long legal history establishing its great breadth, meaning even public health actions of unprecedented scope may well be upheld by the courts if challenged. Again, see Hodge & Gostin, 2002. Second, it is precisely because it is so difficult to envision this sort of use of cognitive bias that I feel compelled to discuss unprecedented actions which are nonetheless still within the traditional domains of government power. My point is that important tasks ought not be foreclosed simply because they were formerly impossible or even unimaginable without the tools presented by research into cognitive biases.

are two contexts in which policymakers ought to eliminate bias if possible and to minimize its effects if not.

First, bias-affecting administrative architecture which significantly affects areas outside the permissible scope of government authority demands first attention. For instance, it is imperative to eliminate or compensate for biases which favor a particular candidate or political party, as well as those which burden a fundamental right. While ignorance excuses the initial establishment of features which impermissibly affect biases, ongoing research demands a government response. Rooting out the unknown bias-triggering elements lurking in government infrastructure is sure to be a time-consuming process, but the response must be exceptionally swift and thorough where these elements affect the political autonomy of the populace.

Second, policymakers may wish to offer citizens legitimately neutral options with the intention that the citizens choose for themselves, free from the influence of cognitive biases. Though the policymakers may be permitted to manipulate biases in the relevant domain, it is possible that they might simply find it prudent to refrain from doing so. In these contexts, the effects of cognitive biases can actually subvert the intent of those who developed the policy. Here, then, it is prudent to remove or minimize the effects of bias in order to align the policy more closely with the intent of its crafters. Reviewing existing regulations for bias-triggering elements in domains in which the government is permitted to leverage cognitive biases, however, is merely a good idea (because it promotes efficiency) rather than a moral imperative (because it does not threaten political autonomy) and, accordingly, ought to be a lower priority than rooting out biases operating in prohibited domains.

In short, policymakers ought not apply the tools of cognitive bias to expand government power beyond its current bounds, but this limitation ought not be construed to forbid ambitious

or unprecedented actions which are nonetheless within the scope of government authority as traditionally understood. Further, the government has a duty to re-evaluate existing policy architecture for aspects which might trigger cognitive biases and to optimize, minimize, or eliminate them, especially biases active in domains in which the government is not permitted intentionally to manipulate them. In the next and final section, I shall offer a summary of my research and conclusions.

VII. Conclusion

Research on cognitive biases has revealed a number of tools which can enable policymakers to increase operational efficiency, promote the general wellbeing of the populace, and strengthen the ties between government and governed. Proponents of leveraging cognitive biases contend that the technique avoids widespread American antipathy towards paternalism while nonetheless altering large scale behavior, but policymakers must be careful to keep this central value in mind when crafting regulations, lest they betray the very core of the libertarian paternalist approach. The temptation to overreach becomes especially dangerous where, as here, the methods of government manipulation at issue are virtually impossible for those affected to resist or even detect.

Leveraging cognitive biases presents three distinct problems. First, knowledge of cognitive biases imposes a duty on policymakers to use them to make government operations as efficient as possible, but the nonconscious nature of bias effects demands caution in choosing where and how to leverage them. Second, biases can affect preferences in addition to behavior, which complicates the democratic government's task of responding to citizens' desires, as those desires could be the result of government-leveraged bias. Finally, policymakers must be careful

that they do not impose too high a burden on those who wish to depart from the government-approved plan, lest the nudge of cognitive bias devolve into a mandate. These are not, however, insurmountable obstacles.

Problems can be minimized or avoided altogether by considering the tools of cognitive bias in a three-part framework. First, government is uniquely justified in leveraging biases to promote pro-democratic behavior because there is reason to believe humans are subject to psychological predispositions which promote instead autocracy-supporting behavior. Second, government is conditionally justified in leveraging biases to accomplish goals traditionally considered within the scope of government authority. When operating in this second part of the framework, the government must be transparent about its use of cognitive biases, consider the ability of affected individuals to overcome leveraged biases, refrain from burdening fundamental rights, use only biases which are scientifically well established, and ensure that the effects of the bias to be leveraged are consistent and widespread enough to be usefully leveraged. Finally, the government is broadly prohibited from leveraging biases to accomplish goals which have not traditionally been within its purview. Note also that the government has a duty to re-evaluate existing institutional architecture for bias-triggering features and optimize, eliminate, or minimize them.

While government may use cognitive biases more freely than violent coercion in most contexts, there are still limits on the range of available implementations. Cognitive biases have vast potential as government tools to efficiently and non-coercively adjust behavior on a very large scale, but policymakers and the people must remain ever mindful of the potential for abuse. The outlines sketched here are but a beginning to the policy considerations which will emerge as the science matures and more biases, as well as new discoveries about old ones, come to light.

Nonetheless, the general principles I have set forth will hopefully serve as a useful starting point for policymakers seeking to implement bias-leveraging regulations in the near future.

Sources

- Ackerman, Joshua M.; Christopher C. Nocera; & John A. Bargh. "Incidental Haptic Sensations Influence Social Judgments and Decisions." *Science*. Vol. 328, pp. 1712-1715. 2010.
- Alford, John R. & John R. Hibbing. "The Origin of Politics: An Evolutionary Theory of Political Behavior." *Perspectives on Politics*. Vol. 2, No. 4, pp. 707-723. 2004.
- Asch, Solomon. "Studies of Independence and Conformity: A Minority of One Against a Unanimous Majority." *Psychological Monographs*. Vol. 70, No. 9. 1956.
- American Presidency Project. "Voter Turnout in Presidential Elections."
<http://www.presidency.ucsb.edu/data/turnout.php> Last Viewed July 11, 2010.
- Beatty, Charles. *A Practical Guide to the Death Duties and to the Preparation of Death Duty Accounts*. London: Effingham Wilson. 1907. Available online at
http://books.google.com/books/download/A_practical_guide_to_the_death_duties_an.pdf?id=LmlJAAAIAAJ&output=pdf&sig=ACfU3U0AL9hBqftL6iM3ybvVhvnYN51EQ
- Bernstein, Irwin S. "Management of Aggression as a Component of Sociality." In *The Origins and Nature of Sociality*, Robert W. Sussman and Audrey R. Chapman, eds. Hawthorne, NY: Aldine de Gruyter. 2004.
- Blumoff, Theodore Y. "On the Potential of Neuroscience: A Comment on Greene and Cohen's 'For the Law, Neuroscience Changes Nothing and Everything.'" *bepress Legal Series*. Working Paper 844. 2005.
- Brosnan, Sarah F., *et al.* "Endowment Effects in Chimpanzees." *Current Biology*. Vol. 17, pp. 1704-1707. 2007.
- Brosnan, Sarah F. *et al.* "Chimpanzee Autarky." *PLoS ONE*. Issue 1. 2008.
- Camerer, Collin F., *et al.* "Neuroeconomics: How Neuroscience Can Inform Economics." *Journal of Economic Literature*. Vol. 43, No. 1, pp. 9-64. 2005.
- Camerer, Collin F. "Wanting, Liking, and Learning: Neuroscience and Paternalism." *University of Chicago Law Review*. Vol. 73, No. 1, pp. 87-110. 2006.
- Chen, M. Keith, *et al.* "How Basic Are Behavioral Biases? Evidence from Capuchin Money Trading Behavior." *Journal of Political Economy*. Vol. 114, No. 3, pp. 517-537. 2006.
- Choi, James, *et al.* "Defined Contribution Pensions: Plan Rules, Participant Decisions, and the Path of Least Resistance." In *Tax Policy and the Economy*, James M. Poterba, ed. Vol. 16, pp. 67-113. 2002.

- Choi, James, *et al.* "Optimal Defaults." *The American Economic Review*. Vol. 93, No. 2, pp. 180-185. 2003.
- Chorvat, Terrence & Kevin McCabe. "The Brain and the Law." *Philosophical Transactions: Biological Sciences*. Vol. 359, No. 1451, pp. 1727-1736. 2004.
- Decety, Jean *et al.* "The Neural Bases of Cooperation and Competition: An fMRI Investigation." *NeuroImage*. Vol. 23, pp. 744-751. 2004.
- Declaration of Independence. 1776. Available online at http://www.archives.gov/exhibits/charters/declaration_transcript.html
- Denno, Deborah. "A Mind to Blame: New Views on Involuntary Acts." *Behavioral Sciences and the Law*. Vol. 21, pp. 601–618. 2003.
- de Waal, Frans. *Peacemaking Among Primates*. Cambridge, MA: Harvard University Press. 1989.
- Dunbar, R. "Coevolution of Neocortex Size, Group Size and Language in Humans." *Behavioral and Brain Sciences*. Vol. 16, No. 4, pp. 681-735. 1993.
- Druckman, James N. "Using Credible Advice to Overcome Framing Effects." *Journal of Law, Economics, and Organization*. Vol. 17, No. 1, pp. 62-82. 2001.
- Erikson, Robert S. "The Advantage of Incumbency in Congressional Elections." *Polity*. Vol. 3, pp. 395-405. 1971.
- The Federalist Papers*. 1787-8. Available online at <http://thomas.loc.gov/home/histdox/fedpapers.html>
- Foer, Franklin & Noam Scheiber. "Nudgeocracy." *The New Republic*. May 6, 2009. Available online at <http://www.tnr.com/article/politics/nudge-ocracy>.
- Fowler, James H. & Darren Schreiber. "Biology, Politics, and the Emerging Science of Human Nature." *Science*. Vol. 322, No. 5903, pp. 912-914. 2008.
- Frank, Jerome. "Experimental Studies of Personal Pressure and Resistance." *Journal of Genetic Psychology*. Vol. 30, pp. 23-64. 1944.
- Freese, Jeremy, *et al.* "The Potential Relevances of Biology to Social Inquiry." *Annual Review of Sociology*. Vol. 29, pp. 233-256. 2003.
- Furuichi, Takeshi. "Agonistic Interactions and Matrifocal Dominance Rank of Wild Bonobos (*Pan paniscus*) at Wamba." *International Journal of Primatology*. Vol. 18, No 6. Pp. 855-875. 1997.

- Greene, Joshua & Jonathan Cohen. "For the Law, Neuroscience Changes Nothing and Everything." *Philosophical Transactions: Biological Sciences*. Vol. 359, No. 1451, pp. 1775-1785. 2004.
- Haselton, M.G.; D. Nettle; & P. W. Andrews. "The Evolution of Cognitive Bias." In *Handbook of Evolutionary Psychology*, D. M. Buss, ed. Pp. 724–746. 2005.
- Hausfater, Glenn; Jeanne Altmann; & Stuart Altmann. "Long-Term Consistency of Dominance Relations Among Female Baboons (*papio cynocephalus*)." *Science*. Vol. 217, pp. 752-755. 1982.
- Hodge Jr., James G. & Lawrence O. Gostin. "School Vaccination Requirements: Historical, Social, and Legal Perspectives." *Kentucky Law Journal*. Vol. 90, pp. 831-890. 2002.
- Hoffman, Morris B. "The Neuroeconomic Path of the Law." *Philosophical Transactions: Biological Sciences*. Vol. 359, No. 1451, pp. 1667-1676. 2004.
- Jones, Owen D., *et al.* "Economics, Behavioral Biology, and Law." Vanderbilt Public Law Research Paper No. 08-38; Vanderbilt Law and Economics Research Paper No. 08-19. *Supreme Court Economic Review*. Vol. 19. 2011.
- Jones, Owen D. "Law and Biology: Towards an Integrated Model of Human Behavior." *Journal of Contemporary Legal Issues*. Vol. 8, pp. 167-208. 1997.
- Jones, Owen D. & Timothy H. Goldsmith. "Law and Behavioral Biology." *Columbia Law Review*. Vol. 105, No. 2, pp. 405-502. 2005.
- Kahneman, D. & Tversky, A. "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology*. No. 3, pp. 430–454. 1972.
- Kahneman, Daniel; Jack L. Knetsch; & Richard H. Thaler. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives*. Vol. 5, No. 1, pp. 193-206. 1991.
- Kahneman, Daniel & Shane Frederick. "Representativeness Revisited: Attribute Substitution in Intuitive Judgment." In *Heuristics and Biases: The Psychology of Intuitive Judgment*, Thomas Gilovich, Dale Griffin, and Daniel Kahneman, eds. Cambridge: Cambridge University Press. Pp. 51–52. 2002.
- Kaplow, Louis & Steven Shavell. "Moral Rules, the Moral Sentiments, and Behavior: Toward a Theory of an Optimal Moral System." *Journal of Political Economy*. Vol. 115, No. 3, pp. 494-514. 2007.
- Koppell, Jonathan G. S. & Jennifer A. Steen. "The Effect of Ballot Position on Election Outcomes." *Journal of Politics*. Vol. 66, pp. 267-81. 2004.

- Krawczyk, Daniel. "Contributions of the Prefrontal Cortex to the Neural Basis of Human Decision Making." *Neuroscience and Biobehavioral Reviews*. Vol. 26, pp. 631–664. 2002.
- Locke, John. *Second Treatise of Civil Government*. 1689. Available online at <http://www.constitution.org/jl/2ndtreat.htm>
- Madrian, Brigitte C. & Dennis F. Shea. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *The Quarterly Journal of Economics*. Vol. 116, No. 4, pp. 1149-1187. 2004.
- Milgram, Stanley. *Obedience to Authority: An Experimental View*. New York: Harper and Row. 1974.
- Morse, Stephen J. "The Moral Metaphysics of Causation and Results." *California Law Review*. Vol. 88, No. 3, pp. 879-894. 2000.
- Morse, Stephen J. "Uncontrollable Urges and Irrational People." *Virginia Law Review*. Vol. 88, No. 5, pp. 1025-1078. 2002.
- National Safety Council. *Injury Facts Book, 2010 Edition*. 2010. Available online in relevant part at http://www.nsc.org/news_resources/Documents/nscInjuryFacts2011_037.pdf
- North Dakota Office of the Secretary of State. "Voter Registration in North Dakota." <http://www.nd.gov/sos/electvote/voting/vote-history.html> Last Viewed July 7, 2010.
- O'Hara, Erin Ann. "How Neuroscience Might Advance the Law." *Philosophical Transactions: Biological Sciences*. Vol. 359, No. 1451, pp. 1677-1684. 2004.
- Pesendorfer, Wolfgang. "Behavioral Economics Comes of Age: A Review Essay on 'Advances in Behavioral Economics.'" *Journal of Economic Literature*. Vol. 44, No. 3, pp. 712-721. 2006.
- Posner, Richard A. *Economic Analysis of Law, Fifth Edition*. New York: Aspen Law & Business. 1998; originally 1973.
- Posner, Richard A. "The Decline of Law as an Autonomous Discipline: 1962-1987." *Harvard Law Review*. Vol. 100, No. 4, pp. 761-780. 1987.
- Pronin, Emily & Matthew B. Kugler. "Valuing Thoughts, Ignoring Behavior: The Introspection Illusion as a Source of the Bias Blind Spot." *Journal of Experimental Social Psychology*. Vol. 43, No. 4, pp. 565-578. July 2006.
- Redelmeier, Donald; Paul Rozin; & Daniel Kahneman. "Understanding Patients' Decisions: Cognitive and Emotional Perspectives." *Journal of the American Medical Association*. Vol. 270, pp. 72-76. 1993.

- Robbins, Gary. "Estate Taxes: An Historical Perspective." Washington, D.C.: The Heritage Foundation. January 16, 2004. Available online at <http://www.heritage.org/Research/Reports/2004/01/Estate-Taxes-An-Historical-Perspective>
- Roughgarden, Joan, *et al.* "Reproductive Social Behavior: Cooperative Games to Replace Sexual Selection." *Science*. Vol. 311, pp. 965-969. 2006.
- Sheridan, C. L. & R. G. King. "Obedience to Authority With an Authentic Victim." *Proceedings of the American Psychological Association*. Pp. 156-66. 1972.
- Smith, D. G. "A 15-Year Study of the Association Between Dominance Rank and Reproductive Success of Male Rhesus Macaques." *Primates*, Vol. 34, pp. 471-80. 1993.
- Smith, Kevin B. "Representational Altruism: The Wary Cooperator as Authoritative Decision Maker." *American Journal of Political Science*. Vol. 50, No. 4, pp. 1013-1022. 2006.
- Somit, Albert & Steven A. Peterson. *The Failure of Democratic Nation Building: Ideology Meets Evolution*. New York: Palgrave Macmillan. 2005.
- Sunstein, Cass R. *Why Societies Need Dissent*. Cambridge: Harvard University Press. 2003.
- Supreme Court of the United States. Gonzales v. O Centro Espirita Beneficente Uniao do Vegetal. 546 U.S. 418. 2006.
- Supreme Court of the United States. District of Columbia v. Heller. 554 U.S. _____. 2008.
- Supreme Court of the United States. McDonald v. Chicago. 554 U.S. _____. 2010.
- Thaler, Richard H. & Cass R. Sunstein. "Libertarian Paternalism." *The American Economic Review*. Vol. 93, No. 2, pp. 175-179. 2003.
- Thaler, Richard H. & Cass R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven: Yale University Press. 2008.
- Tingley, Dustin. "Evolving Political Science: Biological Adaptation, Rational Action, and Symbolism." *Politics and the Life Sciences*. Vol. 25, No. 1-2, pp. 23-41. 2007.
- Tooby, J. & Cosmides, L. "Conceptual Foundations of Evolutionary Psychology." In *The Handbook of Evolutionary Psychology*, D. M. Buss, ed. Pp. 5-67. Hoboken, NJ: Wiley. 2005.
- Tversky, A. & Kahneman, D. "The Framing of Decisions and the Psychology of Choice." *Science*. Vol. 211, pp. 453-458. 1981.

U. S. Census Office. "Voting and Registration in the Election of November 2008." Issued May 2010. Online at <http://www.census.gov/prod/2010pubs/p20-562.pdf>

Wikipedia. "Donkey Vote." http://en.wikipedia.org/wiki/Donkey_vote Last Viewed June 17, 2010.

Wilson, Edward O. *Consilience: The Unity of Knowledge*. New York, NY: Alfred A. Knopf, Inc., 1998.

Wilson, Edward O. *Sociobiology: The New Synthesis, 25th Anniversary Edition*. Cambridge: Harvard University Press, 2000; originally 1975.

Zak, Paul J. "Neuroeconomics." *Philosophical Transactions: Biological Sciences*. Vol. 359, No. 1451, pp. 1737-1748. 2004.