Winter December, 2015

Nudge vs. boost: agency dynamics under "libertarian paternalism"

Michael D Ryall
Dr Ralph Hertwig, *Max Planck Institute*

Nudge vs. boost: agency dynamics under "libertarian paternalism"

R. Hertwig
Max Planck Institute

M. D. Ryall*
University of Toronto

January 1, 2016

Abstract

Thaler and Sunstein (2008) advance the concept of "nudge" policies – non-regulatory and non-fiscal mechanisms designed to enlist people's cognitive biases so as to achieve the desired policy ends. A core assumption is that policy makers engage biases to advance the interests of the nudged individual. We analyze a model of dynamic policy making in which the policy maker's preferences are not always aligned with those of the individual. One novelty of our setup is that one of the policy maker's options is to train the individual to remove the individual bias once and for all. We refer to this as a "boost" policy. Our main result identifies conditions under which nudges have option value – i.e., although it is in the immediate best interests of *both* the policy maker and individual to boost, the policy maker may aver in order to leave open the possibility of future nudges.

1 Introduction

In their influential extension of the lessons of psychology and behavioral economics to public policy, Thaler and Sunstein (2008) explore various ways in which policy makers might design non-regulatory and non-fiscal measures – which they label "nudges" – to help people to make better decisions. Nudges come in two essential forms: "educative" (e.g., disclosure requirements, warnings, labels, reminders) and "non-educative" (e.g., default rules, ordering of

*The authors thank the Berlin Inst. of Advanced Study, all participants of the 13th Blankensee-Colloquium, and Florian Ederer. We also thank our research assistant, Natalia Uborceva. Ryall is grateful for funds provided by the Social Sciences and Humanities Research Council of Canada.

items on a website, or cafeteria design). Non-educative nudges enlist cognitive or motivational biases (e.g., inertia, procrastination, loss aversion; see Rebonato, 2012) to steer individual behavior in a particular direction. Among Thaler and Sunsteins proposals, the manipulative potential of such policies make them both interesting and controversial.[1]

In Thaler and Sunstein (2008) and elsewhere, the manipulation issue is avoided by assuming that policy makers dutifully structure non-educative nudges: i) to achieve ends consistent with individual preferences; and ii) to permit individuals to opt-out if they so desire. This approach to policy design is known as "libertarian paternalism." Even if unobjectionable as so proposed, why take at face-value the promise of self-restraint on the part of policy makers adopting it? Economics has a rich tradition of examining principal-agent problems in a wide variety of settings. Why not extend this tradition to the premise of guileless libertarian paternalists? Potential agency issues may be particularly problematic in this context not only because the policy is by its very nature hidden: the nudge *works* because individuals lack the cognitional machinery required to detect it. A recent survey indicates very little in the way of explicit theoretical work examining this issue (Schnellenbach and Schubert, 2014).

In this paper, we respond with a model to analyze policy formulation in which: i) goal mismatch between policy makers and subjects is admitted; and ii) subjects suffer from a cognitive bias. A key novelty of our approach is that policy makers are endowed with the ability to *eliminate an individual bias permanently and without cost.* We refer to such options as"boost" policies (e.g., an efficacious program of competence-enhancement). By most standards, including those advocated by libertarian paternalists themselves, freeing people from the burden of their cognitive impediments at zero cost is an unambiguous social good. Yet, as we demonstrate, when the goals of policy makers and individuals are at odds, the former may get their way precisely by denying the latter the cognitive tools required to make individually rational decisions. We begin by examining policy in a static setting across a range of goal alignment scenarios,

[1] Although, see Camerer et al. (2003).

then proceed to analyze the dynamic case.

The essential feature of the boost policy in our model is its durability: once an agent learns how to avoid a particular type of bias, that ability is retained and available to be enlisted in future decisions of a similar nature. In our analysis, the policy maker need simply invoke the policy, in which case the boost takes effect instantaneously and at zero cost. Of course, boosting in the real world can be time-consuming and resource-intensive[2] Thus, if policy makers forgo boosts under our idealized conditions, it bodes ill for the prospects of boosts in the real world. It is worth mentioning that, although they do not consider agency issues on the part of policy makers, several authors have begun to explore a variety of learning and information transmission issues in the presence of cognitive biases (Allcott, 2011; Costa and Kahn, 2013; Marteau et al., 2012; Carlin et al., 2013).

Our main result is to demonstrate that nudge policies have option value to policy makers facing the possibility of goal conflict with individuals when new policy problems arise. Specifically, we describe conditions under which the individual falls prey to agency problems on the part of the policy maker, who implements inefficient nudges – even against her own (i.e., the policy maker's) immediate interests – in order to retain the option of using nudges in the future. The dynamic tradeoff that arises here is a result of the asymmetry inherent in the two approaches: nudging today leaves open the option of boosting tomorrow, but the converse is not true. In our model, an individual faces a sequence of problems in which uncertainty exists with respect to goal alignment. On some policy issues, policy maker and individual agree on desired ends. On others, they do not. All of these problems fall within the ambit of a particular type of bias. Thus, in those periods in which boosting is optimal, the policy maker may yet nudge so as not to lose the option to do so in future cases in which goals are in conflict (i.e., in which nudging can be used against the interests of the individual).

[2] Although, low-cost boosts are often available too (Hertwig and Grune-Yanoff, 2016).

Nudging toward the social good? "Drawing on some well-established findings in social science, we show that in many cases, individuals make pretty bad decisions – decisions they would not have made if they had paid full attention and possessed complete information, unlimited cognitive abilities, and complete self-control." So is the starting premise in Thaler and Sunstein (2008, p. 5). In their view, the cognitive and motivational deficiencies displayed by people are both pervasive and consequential."We often make mistakes because we rely too much on our Automatic System" (p. 21); our "biased assessments of risk can perversely influence how we prepare for and respond to crises, business choices, and the political process" (p. 25); we are "unrealistically optimistic even when the stakes are high" (p. 32), and we "adopt ... the 'yeah, whatever' heuristic" (p. 35). As if that isnt enough, we also suffer from self-control problems such as "smoking, alcohol, a failure to exercise, excessive borrowing, and insufficient savings" (p. 42).

Accepting these shortcomings as given, Thaler and Sunstein argue that nudging, under the constraints of libertarian paternalism, is both an efficacious and ethical approach to policy formulation. A non-educative nudge in the context of libertarian paternalism is a policy in which: i) cognitional and motivational biases are enlisted in the service of policy goals; ii) the policy goals are consistent with the preferences of those subject to the nudge; and iii) individuals are easily able to opt out. Thus, the libertarian paternalist is a policy maker tasked with designing institutional mechanisms that create decision-making contexts – "choice architectures" (pp. 11-13) – that steer behavior toward actions that improve welfare at the level of the individual and the group.[3] Policies that preserve autonomy and freedom of choice, pass the libertarian test. The idea is to create a "choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives" (p. 6). Thus, nudging is intended as a "relatively weak, soft and nonintrusive type of paternalism" (p. 6).

The attraction of this approach is, in part, its promise to generate efficient

[3]The policy domain may be public (government-related) or private (business-related).

policy mechanisms driven by science and evidence rather than political dogma. Under nudging, behavior is influenced without recourse to more forceful injunctions (e.g., regulations) or incentives (e.g., taxes). The notion of evidence-based, libertarian paternalism is both novel and appealing. It has already had substantial influence on public policy debates around the globe. For instance, Cass Sunstein served as advisor on regulatory affairs under President Obama, and Richard Thaler advised the British government's Behavioural Insights Team (BIT).

At the same time, libertarian paternalism has also been subject to critical debates. These involve a wide range of issues, including the true pervasiveness and predictability of human irrationality; the ethical implications of nudging; the lack of conceptual clarity around nudging; the extent to which libertarian paternalism is different from other forms of paternalism; whether the policy proposals follow from theoretical claims about the limitations of human rationality and self-control; and the effectiveness of nudges as proposed and implemented. Examples of more recent contributions in this line include: Sugden (2008); Bovens (2009); Berg and Gigerenzer (2010); Allcott (2011); House of Lords (2011); Grüne-Yanoff (2012); Rebonato (2012); Wright and Ginsburg (2012); Handel (2013); Grüne-Yanoff and Hertwig (2015).

The success of the libertarian paternalism program hinges upon an absence of agency problems on the part of the policy maker. While this is always true in some degree, it is especially true in this setting precisely because the nudge is a hidden action on the part of the policy maker. To appreciate the centrality of this premise to the program of libertarian paternalism, consider its most novel aspects. The notion that social influences and pressures can be used to drive behavior is not new.[4] Nor is it original to observe that the provision of information – educative nudges – may often facilitate better decisions.[5] What is genuinely novel is the notion of enlisting people's cognitive and motivational deficiencies so as to help them to make decisions that better versions of themselves would

[4]See e.g., Cialdini (2009).

[5]Thaler and Sunstein (2008, pp. 93-94.).

make. Thus understood, non-educative interventions require (adapted from Rebonato, 2012, p. 32):

1. The identification of a pervasive bias that is attached to a specific decision task. This bias may arise as a property of the environment (e.g., framing of information), limited cognition (e.g., inability to assess probabilities), or motivation (e.g., self-control problems, such as procrastination).
2. An implication that the behavioral consequence of the bias is detrimental to the individual.
3. The feasibility of enlisting the bias (or some other bias) in such a way as to defeat the undesired consequences it normally induces.
4. Knowledge on the part of policy makers regarding the preferences of the affected individuals.
5. Policy goals that are consistent with the preferences of the individuals subject to the policy.

The final entry on the preceding list underlines the central role of the policy maker's intentions. *The Economist* (Anonymous, 2006) expressed this concern as follows: "From the point of view of liberty, there is a serious danger of overreach, and therefore ground for caution. Politicians, after all are hardly strangers to the art of framing the public's choices and rigging its decisions for partisan ends. And what is to stop lobbyists, axe-grinders and busybodies of all kinds hijacking the whole effort." Earlier, Sunstein and Thaler (2003) suggest that an "opt-out right operates as a safeguard against confused or improperly motivated policy makers" (p. 1201). Thaler and Sunstein (2008) also acknowledge the problem by raising the possibility of "evil nudges" (pp. 239-41). As safeguards, they emphasize transparency and regulation ("rules of engagement," p. 240). Ultimately, they advocate trust in democratic systems of checks-and-balances, which afford the electorate the ability to appoint and remove its policy makers.

Unfortunately, the appeal to rational selection of policy makers via the institutions of democracy is credulous given that the act of voting should, in principle, be subject to the same cognitional limitations that libertarian paternalism aims to exploit.[6] As *The Economist* asks, why assume that politicians, their appointees, and professional bureaucrats are immune from the temptation to opportunism? By introducing the possibility of goal divergence between policy makers and their constituents, we explore the conditions under which "evil nudges" arise and formulate some implications with respect to individual (and, by extension, social) welfare. We leave open the question of whether agency problems of this kind are adequately controlled by democratic institutions.[7]

The boost alternative One way of framing the policy problem is as a policy choice between libertarian paternalism, traditional policy options ("hard paternalism"), or doing nothing at all. However, because the root of the problem in the nudge context is a robust bias on the part of the individual, a forth possibility arises: training the individual to avoid the bias and, thereby, equipping him to maximize his own welfare without the help of some social policy maker or other outside party. This option is interesting because, ceteris paribus, it dominates the nudge.[8] Grüne-Yanoff and Hertwig (2015) us the term "boosting" for this option.

The identification of effective training options is at the heart of the "fast-and-frugal heuristics" research program (also known as the "simple heuristics" program; e.g., Gigerenzer et al., 1999; Todd and Gigerenzer, 2012; Gigerenzer et al., 2011). Scholars in this line disagree with advocates of the nudge approach on the extent to which individuals are capable of surmounting their biases. The

[6] The biases of policy makers themselves is a separate issue (see, e.g., Tasic, 2011; Kuehnhanss et al., 2015).

[7] For answers, see, e.g., Besley et al. (2011); Besley and Reynal-Querol (2011); Galasso and Nannicini (2011, 2015).

[8] This assertion should not be controversial as it is consistent with the welfare criterion of nudge advocates themselves, who say that the policy maker's responsibility is "to create a choice architecture that will make it more likely that people will promote their own ends, as they themselves understand them" (Sunstein, 2014, p. 19).

nudge view is that human beings are essentially hostage to a rapid, automatic system of cognition (Thaler & Sunstein, 2008, p. 21). The boost view, in contrast, is that individual decision competence can be improved either directly (by enriching a person's repertoire of skills and decision strategies) or indirectly (through thoughtful changes in the environment; e.g., information representations that are tailored to cognitive algorithms).[9]

Boost policies have two key features: an effort requirement, and generalizability. First, the individual is required to be sufficiently motivated to invest time and effort to acquire and exercise a new skill. At the policy level, additional resources may also be required beyond the individual effort (e.g., instruction materials, instructors, an institutional framework). Second, having acquired the skill, that ability can be exercised in all instances of decision problems prone to the bias it defeats. Thus, under a boost policy, the individual obtains the ability to independently surmount a bias for an entire class of problems – but, at the cost of an up-front resource investment. In contrast, nudge policies impose little in the way of up-front investments (in the sense that individual treatment is not required) – but, at the expense of leaving the individual vulnerable to future instances of the problem. These contrasts create an interesting dichotomy by which to study the strategic behavior of policy makers.

We analyze precisely this choice between boosting and nudging in a dynamic context with the potential for the appearance of policy issues on which the policy maker's goals are not aligned with those of the people whose interests she is meant to represent. By boosting competences, the policy maker ensures that individuals are henceforth able to choose in accord with their own preferences. Yet, if policy makers anticipate that, in future periods, the extant bias steers individuals to decide in accordance with their own preferences, they might choose to withhold such measures. This may be true even when boosting under present circumstances are consistent with policy maker preferences. We say "may" be-

[9]Hoffrage et al. (2000). Also, see Grüne-Yanoff and Hertwig (2015) for a discussion of specific boost examples.

cause the final resolution will depend upon a number of factors, such as the patience of the policy maker, the frequency with which goal mismatch occurs, and relative cost considerations.

The bias we study For the purpose of analyzing a concrete case, we wish to endow the individual in our model with a cognitive bias: i) that appears to be widespread in its effect and, hence, important; and ii) for which there exist treatments by which the bias can be overcome. By these criteria, we settle upon the "the base-rate fallacy," which involves an inability to update probabilities in full accordance with Bayes theorem. First, the base-rate fallacy has been documented in a wide variety of decision domains (e.g. Kahneman and Tversky, 1974; Grether, 1980; Koehler, 1996). Second, well-documented evidence supports methods by which to boost people to reason in a Bayesian fashion (e.g., Hoffrage et al., 2000).

Therefore, in our model, policy makers choose between a nudge, which enlists base-rate fallacy for its desired outcome, and a boost, which aims to overcome the difficulties posed by this bias. Each period a new policy issue arrives for which the policy maker must select her approach. The frequency with which the policy maker and individual see eye-to-eye on the goals of a particular policy is a stochastic feature of the environment. This leads to four possible cases, distinguished by whether or not the goals of the policy maker and individual are aligned, and whether or not the bias is consistent with the goals of the individual. The cases are: I) goals aligned, bias consistent with individual's goals; II) goals aligned, bias inconsistent with individual's goals; III) goals misaligned, bias consistent with individual's goals; IV) goals misaligned, bias inconsistent with individual's goals.

Our setup is such that the individual is always happy to receive boosts, which allow him to maximize his individual welfare in the current period and all future periods. Admitting the possibility that the cost of boosting is greater than nudging, the individual is content to leave his bias in place in Cases I and IV. In Case II, the policy maker and her subject agree that a boost is optimal in the current period. Case III is interesting in the sense that goals are misaligned but there is nothing the policy maker can do to induce the individual to act in accordance with her preferences. Both the nudge and the boost result in the individual choosing against the preferences of the policy maker. In Case IV, the policy maker has an immediate incentive to leave the bias in place (nudge) – against the best interests of the individual. This is the case in which the agency problem manifests. To the extent that misalignment arises with positive probability and the policy maker places some value on future payoffs, nudging is the dominant policy maker choice in both Cases III and IV. Case II is interesting because, although the situation calls for the policy maker to elect a boost consistent with both the policy maker's and individual's best interests, the fact that doing so eliminates the possibility of future opportunism (in Case IV situations) may cause the policy maker to implement a nudge.

In the next section, we lay out the details of our nudge vs. boost stage game. In order to help readers interpret our setup, we briefly discuss how it relates to existing behavioral research on the base rate fallacy. Ultimately, our findings cast light on the the broader issue of agency problems in the context of nudge policies. Our hope is that, by focusing upon an important and well-documented cognitive bias, we will facilitate insights into the broader issues. In Section 3 we conduct a thorough analysis of the stage game. With the results of Section 3 in hand, we go on to examine the dynamic case in Section 4. This section contains Proposition 4, the main result in the paper. Section 5 wraps up with a few concluding comments.

2 Stage game setup

The setting involves a policy maker and a individual. The stage game presents the individual with a class of problems that require evidence-based updating of beliefs in order to act optimally. This presents the policy maker with a policy choice – to boost or nudge. The timing is: i) Nature randomly chooses one of two possible states of the world; ii) the policy maker and individual observe a noisy signal of the true state; iii) the policy maker decides on a policy intervention, nudge or boost; iv) the individual chooses an action; and v) payoffs are awarded based upon the individual's action and the true state.

The individual's problem is characterized by two *states*: θ_1 and θ_2. At the start Nature chooses θ_1 with base rate (marginal) probability $0 < p < 1$ and θ_2 with probability $(1-p)$. We can imagine that θ_1 corresponds to a patient having prostate cancer and θ_2 to him not having it. As we describe below, the states affect the alignment of the individual's goals with those of the policy maker. Once the true state is established, Nature generates a *signal*, which takes the values θ_1^* or θ_2^*.[10] The signal is observed by both players. Here, the signal could represent the outcome of a medical test, such as the prostate-specific antigen (PSA) test for prostate cancer.

The *accuracy* of each signal is α, by which we mean:

$$Pr(\theta_1^*|\theta_1) = Pr(\theta_2^*|\theta_2) = \alpha, \text{ which implies}$$
$$Pr(\theta_2^*|\theta_1) = Pr(\theta_1^*|\theta_2) = 1 - \alpha.$$

Thus, α is the conditional probability that a signal matches the true state. In the prostate cancer example, α is the probability that the PSA test indicates cancer conditional on prostate cancer, indeed, being present in the screened person (i.e., it is the tests sensitivity or true positive rate). Note the assumption of symmetry: α is also equal to the probability of a negative test outcome given that the patient does not have cancer (i.e., the specificity or true negative rate). This implies the probability of Type I errors (the hypothetical PSA test rejects the presence of cancer when it is present) and Type II errors (the test indicates the presence of cancer when it is not present) are both equal to $1-\alpha$. It is worth mentioning that these symmetries are included for convenience and simplicity – they do not drive our conclusions. Assume the signal is *informative* in the sense that $\alpha > \frac{1}{2}$.[11]

[10] In general, a "signal" is any observable that is correlated with a variable of focal interest.

[11] The important assumption is that the signal not be random, i.e., $\alpha \neq \frac{1}{2}$. Signals that are

The individual must choose between one of two actions, a_1 and a_2.[12] To help with transparency, assume the individual's preference is to choose a_i when θ_i is the true state. For example, suppose a patient knows he has prostate cancer, but is uncertain as to the actual type. Then, θ_1 could be a form that grows and spreads quickly – in which case, it is in the interest of the patient to seek immediate treatment (e.g., surgery). In contrast, θ_2 could be a form that grows very slowly, unlikely ever to be a threat during the patient's lifetime – in which case, it is in the patient's interest to forgo immediate treatment.

Of specific interest is the alignment of preferences between the policy maker and individual. We wish to consider three scenarios: i) goals fully aligned; ii) goals partially aligned; and iii) goals fully misaligned. These scenarios are represented by our payoff formulation, which is shown in Table 1 (the policy maker's payoff is on the left side of each cell and the individual's on the right). These are, of course, highly stylized. We forgo the added complexity of variable payoff parameters in order to provide a clear illustration of the logic driving the decision dynamics in this setting.

Table 1: Payoffs for three scenarios

	Aligned		Partial		Misaligned	
state	a_1	a_2	a_1	a_2	a_1	a_2
θ_1	1,1	-1,-1	1,1	-1,-1	-1,1	1,-1
θ_2	-1,-1	1, 1	1,-1	-1,1	1,-1	-1,1

The full-alignment scenario (A) may be interpreted as a decision problem in which individual welfare (the individual's payoff) is congruent with public welfare (the policy maker's payoff). It is also consistent with situations in which the policy maker follows the precepts of libertarian paternalism and restricts the use of a nudge intervention to cases in which the policy goal is to maximize

systematically wrong ($\alpha < \frac{1}{2}$) are also informative. Our consideration of $\alpha > \frac{1}{2}$ amounts to a labeling convention.

[12]One of these actions can represent a "do nothing" option.

individual welfare.[13] In the partial-alignment scenario (P), the goals of both parties are matched in one state, θ_1, but are at odds in the other, θ_2. In this scenario, for example, θ_1 could be the presence of a rare and aggressive form of cancer with very high treatment costs but low or possibly no mortality reduction; θ_2 could be the absence of this disease. If a_1 is the decision to seek no treatment and a_2 the decision to seek treatment, then policy maker and individual preferences are aligned under θ_2 but not under θ_1. In the latter state, the individual desires treatment, while the policy maker prefers not to allocate scarce resources to a fairly hopeless case. In the misalignment scenario (M), individual welfare and public welfare are consistently in conflict. Note that, in cases of conflict, a consistent interpretation is that the individual's preferences are aligned with public welfare, which are at odds with those of a policy maker seeking to maximize personal interests (such as acceding to pressures from private interest groups, at the expense of the public good).

In order to choose an action, the individual must compute the expected payoff of each action given the observed signal. If the observed signal is θ_i^*, then the expected payoff computation requires the individual to formulate a belief about the probability that θ_j, $j = 1, 2$, is the true state given this signal, which we denote $\beta(\theta_j|\theta_i^*)$. A competent individual makes this assessment using Bayes' Rule. We denote the individual's belief that the true state is θ_j given observation of the signal θ_i^* as $\beta(\theta_j|\theta_i^*)$. According to Bayes' Rule,

$$\beta(\theta_j|\theta_i^*) = \frac{Pr(\theta_i^*|\theta_j)Pr(\theta_j)}{Pr(\theta_i^*)}. \tag{1}$$

Thus, for example, a well-calibrated individual would believe:

$$\beta(\theta_1|\theta_2^*) = \frac{(1-\alpha)p}{\alpha(1-p) + (1-\alpha)p}$$

Notice that (1) requires integration of information about both the base rates of events (e.g., cancer) as well as about the accuracy of the signal (e.g., positive rate of medical tests). As it turns out, numerous studies conclude that laypeople and professionals alike fail to process these pieces of information properly

[13] Sunstein & Thaler (2003, p. 1193).

(e.g., Bar-Hillel, 1980; Grether, 1980, 1992). Specifically, when revising beliefs, they pay little to no attention to base rates. This observation leads Kahneman and Tversky (1972, p. 450) to conclude, "In his evaluation of evidence, man is apparently not a conservative Bayesian: he is not Bayesian at all." This bias – base-rate neglect – is the one we examine here. In our setting, an individual suffering from the bias of base-rate neglect bias observes the signal and updates beliefs in the following, incomplete fashion:

$$\beta(\theta_j|\theta_i^*) = Pr(\theta_i^*|\theta_j). \tag{2}$$

That is, the biased individual fails to adjust the accuracy of the signal (the probability of the signal given the true state) to account for the base-rate likelihood of the state. The biased individual mistakenly takes the probability of the state given the signal to be equal to the accuracy of the signal. For example, he believes:

$$\beta(\theta_1|\theta_2^*) = (1-\alpha) \text{ and } \beta_N(\theta_2|\theta_2^*) = \alpha.$$

After the policy maker and individual observe the signal but before the individual acts, the policy maker must choose a policy option – to boost or nudge. We assume that, if the policy maker selects the Boost option, then a treatment is applied to the individual that equips him to compute beliefs according to (1). Alternatively, choice of the Nudge option leaves the individual prone to his bias, which results in beliefs being formulated according to (2). When helpful, we distinguish the post-policy beliefs with subscripts corresponding to these options: β_B and β_N, respectively.

Assume Boost policies require *of the policy maker* an incremental cost of $0 \leq c < 1$ relative to their nudge counterparts, the idea being that the former require more effort and resources than simply permitting the bias to hold sway. When $c = 0$, Boosting is strictly dominant from the perspective of the individual. Assuming that $c < 1$ implies that boosting is never ruled out solely on the basis of cost.

We solve the stage game by assuming that the individual chooses the subjectively optimal action given his beliefs, as influenced by the policy makers choice of policy direction and, aware of this, the policy maker chooses a policy direction that maximizes her own expected payoff. Note that this approach raises consistency issues (especially when we move to the dynamic case) with respect to the sustainability of inaccurate beliefs. Is it reasonable to assume that, faced with persistent evidence that one's expectations fail to capture what actually happens in the real world, an individual will continue to proceed in block-headed fashion to make the same error over and over again? Presumably, the answer has something to do with the frequency of the decisions and the egregiousness of the error. That said, resolving this issue is beyond the scope of this paper. We assume the bias is persistent as long as the Boost option is not invoked. Once it is, the bias is forever eliminated.

A concrete example: cancer screening In the context of base-rate neglect, what would boosting versus nudging amount to in the case of a person awaiting the result of a prostate cancer screening? To appreciate the respective policies, let us briefly explain the prostate-specific antigen (PSA) test. Medical treatments commonly have both benefits and cause harm, and this is also true for cancer screening. Although the major alleged benefit of PSA screening–the reduction of prostate-cancer mortality–is relatively obvious, the harms associated with treatment are subtler and include over-diagnosis, over-treatment, and psychological distress due to false-positive test results. These risks are particularly high for the PSA test, with a reported sensitivity (true positive rate) of approximately 21 percent and a specificity (true negative rate) of approximately 94 percent. The numbers become even more meaningful once one supposes that this screening test is given to 1,000 men from the general population. Assuming a base rate of prostate cancer in the population of about 6.3 percent, 81 percent of the positive test results represent false positives (for numbers and arguments reported here see Arkes and Gaissmaier, 2012).

Designing the choice architecture, the policy maker is assumed to have two options. One is to simply reveal the signal and its accuracy to the the individual, who is then left to assess the probability that he has prostate cancer when it returns a positive indication (this is known as the "positive predictive value" of the test, abbreviated PPV). Assuming the individual neglects the low base rates in computing the PPV, a policy maker could nudge him to seek further diagnosis and treatment by leaving the bias in place: here, base-rate neglect results in an overestimation of the probability of having cancer (given a positive PSA test). A policy maker who deems the risk of over-diagnosis and over-treatment to be an acceptable price for the potential reduction of mortality due to prostate cancer can thus exploit the base-rate fallacy. The benefit of PSA screening for reducing prostate-cancer mortality is very much disputed (see Andriole et al., 2012; Arkes and Gaissmaier, 2012).

Alternatively, the policy maker may choose to foster the statistical competence of the individual. For example, Gigerenzer and Hoffrage (1995) and Hoffrage et al. (2000) demonstrate that statistics expressed as natural frequencies improve Bayesian reasoning skills of experts and non-experts alike. Natural frequencies are non-normalized frequencies that still carry implicit information about base rates, and therefore, reduce the number of computations required to arrive at the correct answer. Individuals' statistical reasoning skills can thus be boosted by either representing statistical information in terms of natural frequencies or by teaching to translate probabilities into natural frequencies (representation training), thus also enabling people to generalize their skills to new Bayesian problems (Sedlmeier and Gigerenzer, 2001). Applied to the PSA test, a policy maker can foster statistical competence by educating the individual to translate probability information into natural frequency form. So equipped, the individual is better able to compute the correct positive predictive value (PPV) and, hence, to decide for himself whether to take additional diagnostic and treatment steps. Moreover, once obtained, this training is applicable to any decision problem requiring similar belief updating.

3 Policy interventions in static settings

In this section, we analyze each of the three scenarios elaborated in Table 1. In each case, we summarize the analysis with a table depicting how the players' strategies and outcomes (with respect to their goals) depend upon the parameters of the problem. We begin with the partial alignment. Once this scenario is examined, the other scenarios are straightforward (Sections 3.2 3.3). We interpret and discuss these results in Section 3.4.

For the purposes of the stage game analysis, assume $c = 0$ (i.e., Boost is the *dominant* policy) and that the policy maker opts for nudging when indifferent between policy options (consistent with the "option value" of nudging that will arise in the dynamic case).

3.1 Partial goal alignment

We proceed by analyzing what happens under each signal, beginning with θ_1^*. Suppose the policy maker observes θ_1^* and adopts the Boost policy. Which action does the individual take? Referring to the payoff table, her subjective expected payoff to selecting a_1 is

$$E_B(a_1) = \beta_B(\theta_1|\theta_1^*)(1) + \beta_B(\theta_2|\theta_1^*)(-1) = \beta_B(\theta_1|\theta_1^*) - \beta_B(\theta_2|\theta_1^*).$$

Similarly, her expected payoff to a_2 is $E_B(a_2) = -\beta_B(\theta_1|\theta_1^*) + \beta_B(\theta_2|\theta_1^*)$. The individual strictly prefers a_1 if $E_B(a_1) > E_B(a_2)$, a_2 if $E_B(a_1) < E_B(a_2)$, and is indifferent otherwise. When the individual is indifferent, assume she takes the action favored by the social policy maker (a_1). Thus, given θ_1^* and the Boost policy, a_1 is chosen if

$$\beta_B(\theta_1|\theta_1^*) - \beta_B(\theta_2|\theta_1^*) \geq -\beta_B(\theta_1|\theta_1^*) + \beta_B(\theta_2|\theta_1^*), \text{ or}$$

$$\beta_B(\theta_1|\theta_1^*) \geq \beta_B(\theta_2|\theta_1^*).$$

Given the parameters in our setup, this condition is given by

$$\frac{\alpha p}{\alpha p + (1-\alpha)(1-p)} \geq \frac{(1-\alpha)(1-p)}{\alpha p + (1-\alpha)(1-p)}, \text{ or}$$
$$p \geq 1 - \alpha. \tag{3}$$

The interpretation is straightforward. Given partial goal alignment, the individual wishes to select a_1 when the true state is θ_1. If the signal is θ_1^* (e.g., positive PSA test), then the conditional probability that θ_1 (e.g., prostate cancer) is, indeed, the true state rises with increasing accuracy of the signal α. For example, in the case of a perfectly informative signal ($\alpha = 1$), this condition will be met (by assumption, $p > 0$). Of course, the probability that θ_1 is the true state given the signal is also a function of the base rate of θ_1, p. The way this condition fails is when the base-rate probability is sufficiently low relative to the accuracy of the signal. That is, the greater the accuracy of the signal, the lower must be the base rate in order for a well-calibrated individual to conclude a_1 is *not* the optimal choice.

Now, suppose, the policy maker adopts a nudge intervention. In this case, the individuals belief remains marred by bias, and the subjective expected payoffs to a_1 and a_2 are, respectively,

$$E_N(a_1) = \beta_N(\theta_1|\theta_1^*)(1) + \beta_N(\theta_2|\theta_1^*)(-1) = 2\alpha - 1, \text{ and}$$
$$E_N(a_2) = \beta_N(\theta_1|\theta_1^*)(-1) + \beta_N(\theta_2|\theta_1^*)(1) = 1 - 2\alpha.$$

Thus, a_1 is chosen under the nudge intervention if

$$\beta_N(\theta_1|\theta_1^*) \geq \beta_N(\theta_2|\theta_1^*), \text{ or,}$$
$$\alpha \geq \frac{1}{2}. \tag{4}$$

The assumption that α is informative implies that (4) is automatically met (given θ_1^*).

In light of these considerations, which intervention does the policy maker implement? Recall, under partial alignment, the policy maker wants the individual to choose a_1 regardless of the state. Given θ_1^* and the assumption that

α is informative, the policy maker has no incentive to expend the additional resources required to boost the individual's analytical competence. Under an nudge intervention, the policy maker – with θ_1^* and α reported – opts for a_1 because condition (4) is then satisfied. If the parameters of the problem are such that condition (3) also holds, then the policy maker's invention results in behavior of the individual that is, indeed, consistent with his interests. If, however, condition (3) fails (i.e., when the base-rate/marginal probability of θ_1, p, is sufficiently low), then the nudge does not square with individuals interests.

Now, suppose the policy maker observes θ_2^*. Once again, let us begin with the boost intervention and analyze the individuals behavior. Following the same line of reasoning as the preceding case, a_1 is chosen if

$$\beta_B(\theta_1|\theta_2^*) \geq \beta_B(\theta_2|\theta_2^*).$$

Given our setup, this is satisfied if

$$\frac{(1-\alpha)p}{\alpha(1-p)+(1-\alpha)p} \geq \frac{\alpha(1-p)}{\alpha(1-p)+(1-\alpha)p}, \text{ or}$$
$$p \geq \alpha. \tag{5}$$

That is, a well-calibrated individual chooses a_1 when the marginal probability of θ_1 being true is sufficiently high relative to the accuracy of the signal θ_2^*.

Similar to the preceding case, when the policy maker implements the nudge intervention, the individual chooses a_2. Thus, if condition (5) holds, the policy maker has an incentive to boost. Burdened with the bias, the individual chooses a_2; properly calibrated, however, he chooses a_1. Since, the costs of boosting are lower than the gain to inducing a_1 ($c = 0 < 1$), the policy maker implements the boost intervention. Alternatively, if condition (5) fails, then the individual chooses a_2 regardless of the policy intervention. Therefore, the policy maker opts for the nudge and suffers the inevitable. The following proposition summarizes the analysis up to this point (note that policy makers are indifferent on the boundaries, e.g., $p = 1 - \alpha$).

Proposition 1 *The equilibrium outcomes under scenario P, given signal θ_i^*, base-rate probability (p), and signal accuracy (α) are summarized in* Table 2. *Each cell in the table reports: policy maker's policy choice, individual's action, and whether the action taken is, in fact, consistent with the individual's welfare (yes or no).*

Table 2: Partially aligned preferences

	Base-rate vs. signal accuracy		
Signal	$0 \leq p \leq 1-\alpha$	$1-\alpha \leq p \leq \alpha$	$\alpha \leq p \leq 1$
θ_1^*	Nudge, a_1, no	Nudge, a_1, yes	Nudge, a_1, yes
θ_2^*	Nudge, a_2, yes	Nudge, a_2, yes	Boost, a_1, yes

3.2 Full goal alignment

We now turn to the scenario of full alignment, and thus a situation in which the policy maker aims to meet the individual's welfare in every state. The analysis is straightforward. Because everyone's goals are aligned, the policy maker implements a boost strategy whenever a well-calibrated individual would arrive at a better decision. There are two such cases. First, when the signal is θ_1^* (e.g., positive PSA test) and the base rate for the indicated state (e.g., prostate cancer) is too small relative to the signal accuracy ($0 \leq p \leq 1-\alpha$). The boost intervention is used for the analogous reason when the signal is θ_2^* (note that the condition $\alpha \leq p \leq 1$ is equivalent to $0 \leq (1-p) \leq 1-\alpha$, where $(1-p)$ is the base-rate probability for θ_2). Proposition 2 summarizes the cases.

Proposition 2 *The equilibrium outcomes under scenario A, given signal θ_i^*, base-rate probability (p), and signal accuracy (α) are summarized in* Table 3.

3.3 Misalignment

Finally, we turn to the misalignment scenario. This case may be interpreted as consistent with situations in which individual welfare is at odds with public

Table 3: Fully aligned preferences

Signal	Base-rate vs. signal accuracy $0 \leq p \leq 1-\alpha$	$1-\alpha \leq p \leq \alpha$	$\alpha \leq p \leq 1$
θ_1^*	Boost, a_1, yes	Nudge, a_1, yes	Nudge, a_1, yes
θ_2^*	Nudge, a_2, yes	Nudge, a_2, yes	Boost, a_1, yes

welfare. Alternatively, it is also consistent with situations in which the policy maker is guileful and wishes to use her authority to promote personal gain at the expense of the public good. Once again, the analysis viz. the case of partial alignment is straightforward. The only time the policy maker has an incentive to boost is when she wants the individual to make well-informed decisions. Since the individual's goals are always at odds with those of the policy maker and boosted decisions always optimize the former, the policy maker is *never* willing to implement a boost policy.

The results are summarized in Proposition 3. Nudge interventions are effective when the base-rate probabilities of the state being signaled are low relative to the accuracy of the information signaling it. Whenever this results in behavior that deviates from the policy makers goals, she will nudge and, therefore, lead the individual away from his goal. Note the implication that, whenever a misalignment scenario with these characteristics arises, the policy-maker never expends resources on fostering unbiased decisions via boost-oriented programs.

Proposition 3 *The equilibrium outcomes under scenario M, given signal θ_i^*, base-rate probability (p), and signal accuracy (α) are summarized in* Table 4.

Table 4: Misaligned preferences

Signal	Base-rate vs. signal accuracy $0 \leq p \leq 1-\alpha$	$1-\alpha \leq p \leq \alpha$	$\alpha \leq p \leq 1$
θ_1^*	Nudge, a_1, No	Nudge, a_1, yes	Nudge, a_1, yes
θ_2^*	Nudge, a_2, yes	Nudge, a_2, yes	Nudge, a_1, no

3.4 Discussion of stage game results

Our findings characterize the outcome of the policy interventions in terms of the relationship between the accuracy of the signal and base-rate probabilities. Base-rate neglect causes the individual to misjudge the true state of the world in a systematic fashion. The policy maker may remedy this problem by fostering the individual's competence for probabilistic reasoning via a boost intervention. Arriving at more accurate assessments of the true state of the world, the boosted individual acts differently than he would have under biased assessment. A necessary condition is that the likelihood ratio ($Pr(\theta_i)/Pr(\theta_i^*)$) be "sufficiently" different from one.

Proposition 1 says that the policy maker forsakes attempting to foster the individual's reasoning competence in most cases, opting instead for nudge policies. In these cases, the policy maker *always* extracts the desired behavior from the individual. The policy maker's choices coincide with the individual's interests in all but one case (Table 2). This case arises when the signal suggests the individual's and policy maker's goals are aligned (θ_1^*) and the error probability ($1 - \alpha$) is high relative to the base-rate probability (p). In the medical example, the nudged individual would choose further diagnosis and treatment even though it is not individually optimal to do so.

In the partial-alignment scenario, the signal may induce biased decisions in conflict with the policy maker's preferences. When teaching the individual to make well-calibrated assessments leads to decisions consistent with those preferences, the policy maker opts for the Boost policy. In this scenario, one such case arises when the signal indicates the individual's and policy maker's goals are misaligned (θ_2^*) and the accuracy (α) is low relative to the base-rate probability (p).

The decision is consistent with the individual's preferences when the base-rate probability of the state of the world *not* suggested by the signal is sufficiently high relative to the accuracy of the signal. For illustration, the accurate appreciation of a medical test for a very rare disease could persuade a patient to forgo aggressive intervention (which may conform with the policy maker's goal). Sometimes, it takes more than a "nudge" to induce the behavior desired by the policy maker.

In the full-alignment scenario, the joint behavior is always in the best interests of both parties. This is elaborated in Proposition 2. For example, if the individual is about to decide whether to undertake and pay directly for advanced education himself, he would choose to do so in exactly the same cases as the policy maker. It is worth noting that, even under full alignment, the policy maker will rely on nudging for a potentially large share of the parameter values. This is because the cognitive bias does not necessarily invite bad decisions; otherwise nudging would loose its rationale. When the bias does not result in bad decisions, it is efficient to leave people to their biases (ignoring such issues as individual dignity or a preference for good judgement). Thus, the full-alignment scenario is by no means tantamount to a world in which policy makers always foster individuals decision-making competences. Rather investments in boosting competencies are only made when the benefits justify it. Other things equal, this is more often the case than in the partial-alignment scenario. Of course, here, the individual always agrees that the investment should, in fact, be made.[14]

Proposition 3 concerns a scenario in which the parties' goals are always in conflict. Here the policy maker will never turn to boosting interventions. From the policy maker's perspective, competent individuals are always undesirable. This brings us to an important insight: any time, individuals encounter a boosting intervention–which require, by definition, the awareness and some minimum engagement (motivation) of the targeted person (Grüne-Yanoff and Hertwig, 2015) – they can legitimately infer that the policy makers have their individual

[14]This need not be true. For example, the policy maker may decide in the interest of public welfare to create a more competent citizenry by always investing in boosting interventions. Such a policy would not be individually rational in our setting. Full alignment represents a world in which one would expect to see symmetric investments in competence improvement in the sense that (for a range of parameter values) boost interventions would be undertaken regardless of the state signaled.

interests at heart. This agreement of intervention policy and individual goals can be brought about by a number of factors, ranging from agreement between public and individual welfare to the policy maker's strict commitment to using nudges only when they conform to individuals goals. Whatever the reason, however, boosts offer strong signals to individuals that resources are being allocated in agreement with their interests. Nudges, in contrast, offer a more ambiguous indicator to individuals, assuming that the nudge is even visible to them (in at least some settings nudge policies can be considered "hidden persuaders"; Smith, 2013). Across the three scenarios analyzed, nudges represent a mixed indicator because they can be employed in cases in which policy makers and individuals have divergent interests.

Thus, we have arrived at three major theoretical results: First, in settings where policy makers and individuals have **partially aligned** goals, there is one case in which the policy maker's optimal policy is boosting. Thus, nudging cannot be the only tool in the policy maker's toolbox if efficiency is desired. Equally important, there is a condition under which nudging fails to serve the individual's interest, while meeting the policy maker's, thus the theory posits situations in which individuals demand more investment into correcting their biases than their policy makers are willing to provide. Second, in **full alignment** settings, there are more cases (than in partial-alignment settings) in which the optimal choice of the policy maker is to boost. That is, to the extent that nudging is restricted to circumstances under which individual and public welfare coincides, investments into boosting are required for a greater breadth of parameters. Third, in settings of **misalignment**, the optimal choice of policy maker is to always nudge, never investment in boosting. This means, if individual welfare and the policy makers interests diverge – for whatever reason – nudging is the policy chosen by the policy maker. From the point of the view of the individual, this turns nudging into a double-edged sword, making it hard to infer from policy whether the intentions of the policy maker coincide with those of the individual.

4 Policy interventions in dynamic settings

Now, let us turn to a dynamic setting. The reason dynamics are important is the inherent asymmetry between boost and nudge interventions. Once boosted, the individual is enduringly endowed with the ability to make well-calibrated, self-interested choices – irrespective of whether they comport with the policy maker's objectives. This means that once a person has received a successful boost, the nudge associated with the corrected bias forever loses its effectiveness. This is not true in the reverse case – nudging does not eliminate the effectiveness of future boosts. How does this tension resolve?

Assume the policy maker and individual face a sequence of problems, one in each period $t = 1, 2, \ldots$. We now focus only on the partial-alignment scenario for two reasons. First, the full alignment and misalignment scenarios are trivial. In the latter, the policy maker never boosts (Table 4) and, in the former, the policy maker waits until a boost is indicated in the period according to Table 3, then boosts and then abstains from further interventions. Therefore, the partial-alignment scenario is the interesting one.

To represent a sequence of distinct policy problems, we assume that the base-rate probability of θ_1 (p) changes every period. We attach t subscripts to indicate components of the model that vary over time – in the first instance, the *problem* p_t. Assume that p_t is selected from some distribution $F(\cdot)$, where F is absolutely continuous and strictly increasing on $[0, 1]$. This implies the existence of a density function $f(p_t) = \frac{d}{dp_t}F(p_t)$ that is strictly positive on $(0, 1)$. Little in the way of additional insight will be gained by varying α, so let us assume signal accuracy is a stable element of the setting.

According to Proposition 1, problems can be classified into one of three focal categories defined by the relationship of their base-rate probabilities to the signal. Let $A \equiv [0, 1-\alpha], B \equiv [1-\alpha, \alpha]$, and $C \equiv [\alpha, 1]$ denote the intervals that correspond to the key ranges identified in Proposition 1. Then, for $j = A, B, C$, define:

$$\sigma_j \equiv \int_j f dp, \qquad \bar{p}_j \equiv \sigma_j^{-1} \int_j p f dp. \tag{6}$$

The σ_js correspond to the probability that the problem generated in a period falls within the Category j range. Note that $\sigma_A = F(1-\alpha)$, etc. The $\bar{p}_j$s correspond to the conditional expected value of p *given* Category j. From our earlier assumptions about F, these are nonzero and, moreover, $\bar{p}_A \leq \bar{p}_B \leq \bar{p}_C$. Assume the problem and state generation mechanisms are stochastically independent. The only case in which the policy maker ever has an incentive to boost, according to Proposition 1 is when $p_t \in C$ *and* the signal is θ_2^*. Under this setup, the probability of this situation arising at the start of a period is $\sigma_C[\alpha - (2\alpha - 1)p_t]$.

Each period, the policy maker observes p_t and signal θ_{jt}^*, $j \in \{1, 2\}$. She then implements either a nudge or a boost. The individual responds to the policy maker's intervention to the best of his cognitive ability (i.e., according to Proposition 1). The policy maker faces a non-trivial decision of whether or not to boost in periods where $p_t \in C$ and the signal is θ_2^*. Within such a period, the best policy without regard to future effects is to boost. Once empowered, however, an individual always has the ability to conduct proper Bayesian analyses. Consequently, in future periods $k > t$ with $p_k \in A$ and signal θ_2^* – that is, in which the policy maker prefers to nudge – the base-rate neglect nudge is no longer available.

To proceed, let K denote the event: "the individual is *competent*," (i.e., was subjected to a boost intervention in some previous period during the history of play). Alternatively, let U denote the event: "the individual is miscalibrated," (i.e. a boost intervention has never been implemented). Then, we can define a *state* $\omega = (x, y, z)$, where $x \in \{U, K\}$ indicates the knowledge of the individual at the start of this period (based upon whether a boost event has ever occurred in any previous event), $y \in \{A, B, C\}$ is the problem category (described above), and $z \in \{\theta_1^*, \theta_2^*\}$. Let Ω denote the set of all (twelve) states. To keep terminology straight in the dynamic setting, we refer to the focal sets of problems (the key intervals within which the p_ts fall) as problem *categories* and the ω's as *states*.

The policy maker observes an element of Ω at the start of each period. Given a state $\omega_t \in \Omega$ in period t, the policy maker's action in that period determines a probability distribution on Ω, the *transition probabilities*, according to which the next period's state is generated. If the policy maker observes ω_t and implements $a_t = N$, then $Pr(\omega_{t+1}|a_t, \omega_t)$ is the resulting transition probability that ω_{t+1} occurs in $t+1$. Note that current period payoffs and the probability of next period's state only depend upon this period's state plus the players' actions. Therefore, this is a finite stochastic game, a key implication of which is the existence of Markov Perfect equilibria.

Assume the policy maker's *discount factor* is $\delta \equiv (\frac{1}{1+r})$, where $r \in [0, 1]$ is her *discount rate*. A policy maker with a larger discount rate places lower value on future payoffs relative to present ones and, hence, is considered more "impatient" relative to policy makers with smaller rs. Let $V(\omega_t)$ denote the net present value to the policy maker of being in state ω_t given a dynamically optimal nudge intervention from t on. The net present value to the policy maker today is equal to the payoff in the present period plus the discounted value of next period's value function for whatever state arises. Consequently, in each period, the policy maker's problem can be solved by maximizing the sum of today's payoff plus the discounted expected value of the states in the next period. This also allows us to simplify the notation by dispensing with t subscripts unless absolutely necessary to avoid ambiguity. Thus, in any period, the net present value of being in state ω_t can be written as:

$$V(\omega_t) = \max_{a_t \in \{N,B\}} \pi(a_t, \omega_t) + \delta \sum_{\omega_{t+1} \in \Omega} Pr(\omega_{t+1}|a_t, \omega_t) V(\omega_{t+1}). \tag{7}$$

That is, the net present value of being in state ω_t, $V(\omega_t)$, is the payoff in the current period plus the expected discounted value of next period's state (given the continuance of a value-maximizing nudge intervention). The payoff, $\pi(a_t, \omega_t)$, is the expected payoff in the present period given ω_t and policy maker policy a_t.

Now we come to the key question: What is the policy maker's optimal dynamic behavior? As noted, the only situation in which a boost intervention makes the policy maker better off is when the state implies a badly calibrated individual, the present problem is in Category C, and the signal is θ_2^*. If a boost intervention has already been implemented (event K), then the policy makers optimal choice is to always nudge – it is costless and, in any case, nothing the policy maker does from this point on affects the individual's decision. If it is not optimal to boost in this state, then it is never optimal to do so. Therefore, a boost occurs only once – the first time $\omega_t = (U, C, \theta_2^*)$ – or not at all.

Proposition 4 *When $c = 0$, a Boost policy is optimal in period t if and only if $\omega_t = (U, C, \theta_2^*)$ and*

$$\delta \leq \frac{1}{1 + Pr(A, \theta_1^*) - Pr(C, \theta_2^*)}, \tag{8}$$

or,

$$r \geq Pr(A, \theta_1^*) - Pr(C, \theta_2^*). \tag{9}$$

Proof of Proposition 4. See the Appendix. ■

As is conventional in analyses of this kind, Proposition (4) is stated in terms of a condition on the patience of the policy maker. Presenting the result in this way highlights the tension faced by the policy maker between cashing-in on today's payoff versus maximizing those in the future. When $\omega_t = (U, C, \theta_2^*)$, the policy maker's period-t payoff is maximized by boosting the individual. This, however, imposes an opportunity cost: the inability to nudge the individual in future periods. As we know from Proposition (1), there are two cases in which the knowledge of the individual matters from the policy maker's point-of-view. The first is when $p \in A$ and the signal is θ_1^*. Here, a well-calibrated individual chooses a_2, against the policy maker's preferences, while a biased one chooses a_1, consistent with them. By choosing to boost in period t, the benefits of nudging in future periods when this case arises is forgone. However, in the case in which $p \in C$ and the signal is θ_2^*, not boosting results in an immediate loss of payoff, as the biased individual chooses a_2 rather than a_1 (the latter of which would be chosen by a boosted subject). In all other combinations of signals and events, the individual's choice is unaffected by policy treatment.

Thus, conditions (8) and (9) say that the policy maker will boost when the rate at which she discounts future payoffs is "sufficiently high" i.e., when greater weight is being placed on current vs. future payoffs. How low? The answer depends upon the probabilities with which the cases of interest arise. Higher values of $Pr(A, \theta_1^*)$ mean that leaving open the possibility of nudging has greater future value. Hence, the greater this probability, the lower must be the policy maker's discount factor for her to choose a boost. Alternatively, the greater is $Pr(C, \theta_2^*)$, the higher the value δ consistent with choosing a boost (because (C, θ_2^*) occurs more frequently anyway). Note that, while we have assumed $c = 0$, it is not hard to show that the greater the cost of a boost, the less attractive it is and, hence, the lower must be δ in order for boosting to be optimal for the policy maker.

Corollary 1 *When $c = 0$, a Boost policy is optimal in period t if and only if $\omega_t = (U, C, \theta_2^*)$ and*

$$r \geq \Big((1-\alpha) + (2\alpha - 1)\bar{p}_A\Big)\sigma_A - \Big(\alpha - (2\alpha - 1)\bar{p}_C\Big)\sigma_C. \tag{10}$$

Corollary (1) restates Proposition (4) in terms of the primitives defined in (6). This is helpful because, e.g., $Pr(A, \theta_1^*)$ is the joint probability of facing a problem $p_t \in A$ *and* a signal θ_1^*. The terms in the large parentheses are the respective (from left to right) expected probabilities with which θ_1^* and θ_2^* are observed in any given period. Condition (10) accords with our intuition. Greater probabilities of receiving future problems in category A, or of observing the signal θ_1^* when faced with a problem from this category, imply a greater expected opportunity cost of losing the ability to nudge in those periods by boosting today. Thus, the need to offset these with a larger discount rate, which shifts weight to the current period. Alternatively, greater probabilities of receiving a category C problem, or of observing the signal θ_2^* when faced with such a problem, imply a lower opportunity cost to boosting today.

Thus, in the extreme case of perfectly informative signals ($\alpha = 1$), condition (10) becomes

$$r \geq \bar{p}_A \sigma_A - \bar{p}_C \sigma_C.$$

That is, all that matters are the relative probabilities of each category weighted by the conditional probabilities of seeing θ_1^* when faced with a problem from those categories. When signals are completely uninformative ($\alpha = \frac{1}{2}$), condition (10) becomes

$$r \geq \frac{1}{2}(\sigma_A - \sigma_C).$$

In other words, if the signal is useless, the probabilities of seeing one versus the other is also useless.

5 Conclusions

The nudge approach to policy advocated by Thaler and Sunstein (2008) is important not only as a compelling call to policy makers to take (some) results of psychology and behavioral economics seriously but, perhaps even more so, because it advances a specific framework for doing so, libertarian paternalism. Recent attempts by real-world policy makers to adopt this framework speak to its inherent attractiveness. These attempts notwithstanding, relatively little is presently known about the efficacy of constructing social policy based upon a reliance on such biases (House of Lords, 2011). At a higher level still, one also faces issues surrounding the ethics of such policies and, relatedly, the special agency problems that may arise in this context. That actual policy makers have already made preliminary incursions into libertarian paternalism suggests the importance of fleshing out these research issues.

This paper explores the issue of agency at the policy maker level in the context of readily available policy options that are capable of eliminating the decision biases of those subject to the policy maker's plans. This places the agency problem in stark relief since empowering individuals to be liberated from biased judgement is an uncontroversial social good. However, as we clearly show, the asymmetric nature of boosting versus nudging – in the sense that nudging always leaves open the option for boosting, but the converse is not true – means that nudging has option value to the policy maker. Thus, a sufficiently patient policy maker may prefer not to heal a bias, even when doing so is in the immediate interests of the policy maker, in order to leave open future opportunities to employ nudges in the future. In such situations, the emphasis of libertarian paternalism leans much more heavily in the direction of paternalism.

Our results raise a number of related issues. For example, a common complaint in Western democracies is that politicians, who seek reelection every few years, consider only the short-term consequences of the policies. This implies that the discount rate of the policy makers is very low indeed. However, a short-run horizon on the part of the policy maker in our setting may, counterintuitively, imply the adoption of policies that are actually efficient over the long-term (with boosted constituents becoming some *other* politician's problem later on down the road). Another example is our restriction to the policy makers of the capability to correct bias. One imagines that the discovery of bias-correcting programs implies the creation of markets for such services. What are the policy dynamics in the presence of privately offered bias-correction services? Which segments of the population are more likely to seek private de-biasing? What are the policy implications when libertarian paternalism affects some segments more than others? How would one interpret the imposition of regulations imposing "quality requirements" on educational programs designed to circumvent the biases exploited by paternalistic policy makers? These are a but handful of issues raised by our results.

A Proof of Proposition 4

The necessity of $\omega_t = (U, C, \theta_2^*)$ is true by Proposition 1: a change in any element implies that a nudge yields the highest expected payoff in period t. In order for a boost to be dynamically optimal, it must result in the highest expected payoff in the present period.

To see sufficiency, assume $\omega_t = (U, C, \theta_2^*)$. The policy maker can either Boost or Exploit. Suppose she Boosts. From Proposition 1, her payoff is $1 - c$ for certain in the present period. Her expected payoff is then constant in all future periods because, from then on, K is true. The expected payoff per period is:

$$\begin{aligned}
\mathbb{E}(K) \equiv\ & Pr(K, A, \theta_1^*)(-1) + Pr(K, B, \theta_1^*)(1) + Pr(K, C, \theta_1^*)(1) \\
& + Pr(K, A, \theta_2^*)(-1) + Pr(K, B, \theta_2^*)(-1) + Pr(K, C, \theta_2^*)(1), \\
& = -Pr(A, \theta_1^*) + Pr(B, \theta_1^*) + Pr(C, \theta_1^*) - Pr(A, \theta_2^*) - Pr(B, \theta_2^*) + Pr(C, \theta_2^*),
\end{aligned}$$

because $Pr(K) = 1$ forever once the Boost is administered. If it is optimal for the policy maker to Nudge, then she receives a payoff of -1 in the present period and then she Nudges forever. If so, the expected payoff per future period is:

$$\begin{aligned}
\mathbb{E}(U) \equiv\ & Pr(U, A, \theta_1^*)(1) + Pr(U, B, \theta_1^*)(1) + Pr(U, C, \theta_1^*)(1) \\
& + Pr(U, A, \theta_2^*)(-1) + Pr(U, B, \theta_2^*)(-1) + Pr(U, C, \theta_2^*)(-1), \\
& = Pr(A, \theta_1^*) + Pr(B, \theta_1^*) + Pr(C, \theta_1^*) - Pr(A, \theta_2^*) - Pr(B, \theta_2^*) - Pr(C, \theta_2^*).
\end{aligned}$$

Therefore, assuming optimal play in the future,

$$\begin{aligned}
V(U, C, \theta_2^*, B) &= (1 - c) + \delta \sum_{t=0}^{\infty} \delta^t \mathbb{E}(K), \\
&= (1 - c) + \frac{\delta}{1 - \delta} \mathbb{E}(K),
\end{aligned}$$

and,

$$\begin{aligned}
V(U, C, \theta_2^*, E) &= -1 + \delta \sum_{t=0}^{\infty} \delta^t \mathbb{E}(U), \\
&= -1 + \frac{\delta}{1 - \delta} \mathbb{E}(U),
\end{aligned}$$

A Boost is, therefore, optimal when $V(U, C, \theta_2^*, B) - V(U, C, \theta_2^*, E) \geq 0$, or

$$(2 - c) + \frac{2\delta}{1 - \delta}(Pr(C, \theta_2^*) - Pr(A, \theta_1^*)) \geq 0. \tag{11}$$

Setting $c = 0$ and solving (11) for δ results in

$$\delta \leq \frac{1}{1 + Pr(A, \theta_1^*) - Pr(C, \theta_2^*)}. \tag{12}$$

B Proof of Corollary 1

Proof. For the corollary, it remains to elaborate $Pr(A, \theta_1^*)$ and $Pr(C, \theta_2^*)$. Let us take these in turn. $Pr(A, \theta_1^*)$ is the probability that $p \in [0, 1-\alpha]$ *and* $Pr(\theta_1^*)$. Fixing p, the latter is

$$\begin{aligned}
Pr(\theta_1^*) &= Pr(\theta_1^*|\theta_1)Pr(\theta_1) + Pr(\theta_1^*|\theta_2)Pr(\theta_2) \\
&= \alpha p + (1 - \alpha)(1 - p) \\
&= (1 - \alpha) + (2\alpha - 1)p.
\end{aligned}$$

Therefore,

$$\begin{aligned}
Pr(A, \theta_1^*) &= \int_0^{1-\alpha} [(1 - \alpha) + (2\alpha - 1)p]\, f(p)dp \\
&= (1 - \alpha)\int_0^{1-\alpha} f(p)dp + (2\alpha - 1)\int_0^{1-\alpha} pf(p)dp, \\
&= (1 - \alpha)\sigma_A + (2\alpha - 1)\bar{p}_A\sigma_A, \\
&= \big((1 - \alpha) + (2\alpha - 1)\bar{p}_A\big)\sigma_A,
\end{aligned}$$

where the third equality follows from the definitions in (6). Similarly,

$$\begin{aligned}
Pr(\theta_2^*) &= Pr(\theta_2^*|\theta_1)Pr(\theta_1) + Pr(\theta_2^*|\theta_2)Pr(\theta_2) \\
&= \alpha + (1 - 2\alpha)p.
\end{aligned}$$

Which implies,

$$Pr(C, \theta_2^*) = \big(\alpha + (1 - 2\alpha)\bar{p}_C\big)\sigma_C.$$

Substituting back into (12):

$$\delta \leq \frac{1}{1 + \big((1-\alpha) + (2\alpha - 1)\bar{p}_A\big)\sigma_A - \big(\alpha - (2\alpha - 1)\bar{p}_C\big)\sigma_C}. \qquad (13)$$

■

References

Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics 95*(9), 1082–1095.

Andriole, G. L., E. D. Crawford, R. L. Grubb, S. S. Buys, D. Chia, T. R. Church, M. N. Fouad, C. Isaacs, P. A. Kvale, D. J. Reding, J. L. Weissfeld, L. A. Yokochi, B. O'Brien, L. R. Ragard, J. D. Clapp, J. M. Rathmell, T. L. Riley, A. W. Hsing, G. Izmirlian, P. F. Pinsky, B. S. Kramer, A. B. Miller, J. K. Gohagan, and P. C. Prorok (2012). Prostate Cancer Screening in the Randomized Prostate, Lung, Colorectal, and Ovarian Cancer Screening Trial: Mortality Results after 13 Years of Follow-up. *JNCI Journal of the National Cancer Institute 104*(2), 125–132.

Anonymous (2006). Soft paternalism: the state is looking after you. *The Economist Online*.

Arkes, H. and W. Gaissmaier (2012). Psychological research and the prostate-cancer screening controversy. *Psychological Science 23*(6), 547–553.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica 44*(3), 211–233.

Berg, N. and G. Gigerenzer (2010). As-if behavioral economics: Neoclassical economics in disguise? *History of Economic Ideas 18*(1), 133–166.

Besley, T., J. G. Montalvo, and M. ReynalQuerol (2011). Do Educated Leaders Matter?*. *The Economic Journal 121*(554), F205–227.

Besley, T. and M. Reynal-Querol (2011). Do democracies select more educated leaders? *American political science review 105*(3), 552–566.

Bovens, L. (2009). *The ethics of nudge.* Springer Netherlands.

Camerer, C., S. Issacharoff, T. O'donoghue, and M. Rabin (2003). Regulation for Conservatives: Behavioral Economics and the Case for" Asymmetric Paternalism". *University of Pennsylvania law review*, 1211–1254.

Carlin, B., S. Gervais, and G. Manso (2013). Libertarian paternalism, information production, and financial decision making. *Review of Financial Studies 26*(9), 2204–2228.

Cialdini, R. (2009). *Influence: Science and practice*. Boston: Pearson Education.

Costa, D. and M. Kahn (2013). Energy conservation nudges and environmentalist ideology: evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association 11*(3), 680–702.

Galasso, V. and T. Nannicini (2011). Competing on good politicians. *American Political Science Review 105*(1), 79–99.

Galasso, V. and T. Nannicini (2015). So closed: Political selection in proportional systems. *European Journal of Political Economy 40*(3), 260–273.

Gigerenzer, G., R. Hertwig, and T. Pachur (Eds.) (2011). *Heuristics: The foundations of adaptive behavior*. Oxford: Oxford University Press.

Gigerenzer, G. and U. Hoffrage (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological review 102*(4), 684.

Gigerenzer, G., P. M. Todd, and the ABC Group (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Grether, D. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics 95*(3), 537–557.

Grether, D. (1992). Testing Bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization 17*(1), 31–57.

Grüne-Yanoff, T. (2012). Old wine in new casks: libertarian paternalism still violates liberal principles. *Social Choice and Welfare 38*(4), 635–645.

Grüne-Yanoff, T. and R. Hertwig (2015). Nudge versus boost. How coherent is policy and theory?

Handel, B. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *The American Economic Review 103*(7), 2643–2682.

Hoffrage, U., S. Lindsey, R. Hertwig, and G. Gigerenzer (2000). Communicating statistical information. *Science 290*, 2261–2262.

House of Lords (2011). Behavior Change (HL paper 179). Technical report, House of Lords, London.

Kahneman, D. and A. Tversky (1974). Subjective probability: A judgment of representativeness. In *The Concept of Probability in Psychological Experiments*, pp. 25–48. Springer Netherlands.

Koehler, J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and brain sciences 19*(1), 1–53.

Kuehnhanss, C., B. Heyndels, and K. Hilken (2015). Choice in politics: Equivalency framing in economic policy decisions and the influence of expertise. *European Journal of Political Economy 40*, 360–374.

Marteau, T., G. Hollands, and P. Fletcher (2012). Changing human behavior to prevent disease: the importance of targeting automatic processes. *Science 337*(6101), 1492–1495.

Rebonato, R. (2012). *Taking liberties: A critical examination of libertarian paternalism*. Palgrave Macmillan.

Schnellenbach, J. and C. Schubert (2014). Behavioral Politiacal Economy. *CESifo Working Paper no. 2698*.

Sedlmeier, P. and G. Gigerenzer (2001). Teaching Bayesian reasoning in less than two hours. *Journal of Experimental Psychology: General 130*(3), 380.

Smith, N. (2013). Choice without awareness: ethical and policy implications of defaults. *Journal of Public Policy & Marketing 32*(2), 159–172.

Sugden, R. (2008). Why incoherent preferences do not justify paternalism. *Constitutional Political Economy 19*(3), 226–248.

Sunstein, C. (2014). *Why Nudge?: The Politics of Libertarian Paternalism.* Yale University Press.

Sunstein, C. and R. Thaler (2003). Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review*, 1159–1202.

Tasic, S. (2011). Are Regulators Rational? *Journal des Economistes et des Etudes Humaines 17*(1).

Thaler, R. and C. Sunstein (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.

Todd, P. and G. Gigerenzer (2012). *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.

Wright, J. and D. Ginsburg (2012). Behavioral law and economics: Its origins, fatal flaws, and implications for liberty. *Nw. UL Rev. 106*, 1033.