

Behavioral economics: implications for regulatory behavior

James C. Cooper · William E. Kovacic

Published online: 25 January 2012

© The Author(s) 2012. This article is published with open access at Springerlink.com

Abstract Behavioral economics (BE) examines the implications for decision-making when actors suffer from biases documented in the psychological literature. This article considers how such biases affect regulatory decisions. The article posits a simple model of a regulator who serves as an agent to a political overseer. The regulator chooses a policy that accounts for the rewards she receives from the political overseer—whose optimal policy is assumed to maximize short-run outputs that garner political support, rather than long-term welfare outcomes—and the weight the regulator puts on the optimal long run policy. Flawed heuristics and myopia are likely to lead regulators to adopt policies closer to the preferences of political overseers than they would otherwise. The incentive structure for regulators is likely to reward those who adopt politically expedient policies, either intentionally (due to a desire to please the political overseer) or accidentally (due to bounded rationality). The article urges that careful thought be given to calls for greater state intervention, especially when those calls seek to correct firm biases. The article proposes measures that focus rewards to regulators on outcomes rather than outputs as a way to help ameliorate regulatory biases.

Cooper and Kovacic completed part of this paper while at the Federal Trade Commission. The views expressed are ours alone and do not necessarily represent those of the Federal Trade Commission or any other individual commissioner.

J. C. Cooper

Law & Economics Center, George Mason University School of Law, Arlington, VA, USA

W. E. Kovacic (✉)

George Washington University Law School, Washington, DC, USA

e-mail: wkovacic@law.gwu.edu

Keywords Antitrust · Behavioral economics · Public choice · Administrative law · Competition policy

JEL Classification K23 · K21 · L40 · D72 · D73 · D03

1 Introduction

Behavioral economics (BE) examines the implications for decision-making when actors suffer from biases documented in the psychological literature. BE scholars replace the assumption of rationality with one of “bounded rationality,” in which consumers’ actions are affected by their initial endowments, their tastes for fairness, their inability to appreciate future costs, their lack of self-control, and the general use of flawed heuristics. BE is increasingly visible in a modern antitrust scholarship that is skeptical of entry as a means to correct anticompetitive markets; doubts the inherent instability of cartels; disputes the efficiency of bundling; and challenges claims that restrictions on intra-brand price competition promote inter-brand competition in non-price dimensions (e.g., [Stucke 2007](#); [Tor 2002](#); [Tor and Rinner 2011](#)).

There are serious questions about BE’s import for public policy ([Klick and Mitchell 2006](#); [Ginsburg and Moore 2010](#); [Wright 2007](#); [Wright and Stone 2011](#)), but that is not our focus. Rather, in the spirit of the public choice—which was revolutionary by peeling back the curtain of government policy to examine how self-interested decision makers would act—we withdraw the veil of policy-making to study how bounded rationality affects regulatory decisions. Several BE scholars recognize this possibility (e.g., [Jolls et al. 1998](#); [Bennett et al. 2010](#)), but generally treat it as a second-order problem compared to the likely biases that beset consumers.¹ We aim to fill this gap in the literature.

We posit a simple model of a regulator who serves as an agent to a political overseer. The regulator chooses a policy that accounts for the rewards she receives from the political overseer—whose optimal policy is assumed to maximize short-run outputs that garner political support, rather than long-term welfare outcomes—and the weight the regulator puts on the optimal long run policy. We use this model to explore the effects of bounded rationality on policymaking, with particular emphasis on competition and consumer protection policy. We find (unsurprisingly) that flawed heuristics (e.g., availability, representativeness, optimism, and hindsight) and myopia are likely to lead regulators to adopt policies closer to the preferences of political overseers than they would otherwise. The effect of status quo and confirmation biases is less clear, and depends on initial policy positions, the order and veracity of information flows, and the regulator’s priors (or first piece of evidence). We conjecture that confirmation bias may create a weak tendency to adopt politically expedient policies given that the first piece of evidence a regulator views on a matter likely will be a call to action.

¹ One notable exception is [Rachlinski and Farina \(2002\)](#), who compare the positive and normative implications of both BE and public choice theory for institutional design. They focus on the interaction among Congress, the executive, courts, and regulators. We provide a formal treatment, allowing public choice theory and BE simultaneously to shape policy outcomes.

Unlike the case of firms that face competition, we argue that the incentive structure for regulators is likely to reward those who adopt politically expedient policies, either intentionally (due to a desire to please the political overseer) or accidentally (due to bounded rationality). These incentives are likely to yield a cadre of regulators who focus excessively on outputs rather than outcomes. Our analysis suggests that careful thought be given to calls for greater state intervention, especially when those calls are directed at firm biases. We also conjecture that instituting internal and/or external teams to review policy, and focusing rewards on outcomes rather than outputs, can help ameliorate regulatory biases.

Section 2 provides a simple model of regulation and explores the implications of bounded rationality in policy making. Section 3 considers how much regulators are likely to suffer biases. Section 4 suggests ways to design the decision-making structure to ameliorate these biases. Section 5 concludes.

2 A simple model of regulation

In this section we present a simple model, where a regulator chooses a policy, π^R , to maximize her utility. In our model, the regulator ultimately must approve a market intervention.² Utility is a function both of the distance (in policy space) the chosen policy is from the policy that the regulator believes is best for society, π^* , and political rewards—future promotions, increased budgets, fewer hearings, good press—from political overseers for choosing policies that maximize their chances of reelection, π^{os} . We assume that π^{os} is a politically expedient policy that maximizes short-run outputs (or resource use). This follows from the observation that constituents (i.e., people) are likely to suffer from various biases that cause them to demand short-sighted policies. Politicians exploit these biases to provide short-run solutions to non-problems that are likely to facilitate reelection.³ In this manner, the political overseers in our model act as transparent middlemen, who convert boundedly rational constituents' preferences into policy demands. Jolls et al. (1998), for example, explain Superfund as an appeal to biased voters by “availability entrepreneurs.” A more recent example consistent with politicians taking advantage of the “availability bias” includes the recent German decision to abandon nuclear power in the wake of the Japanese nuclear crisis. Therefore, π^{os} could include an investigation of market manipulation, an anti-price gouging law, or even price controls imposed to cap surging gasoline prices. Each policy allows the politician to appear to take steps to counteract high gasoline prices. Such moves increase the probability of reelection relative to the politician who took no action and tried to explain that retail gasoline prices were a function of supply and demand on world markets for crude oil and warned that attempts artificially to lower prices by regulation are likely to cause harm.

² Rachlinski and Farina (2002) distinguish between career bureaucrats and agency heads in terms of expertise. We focus on agency leaders who must approve interventions.

³ Of course, this assumes that politicians lack biases when estimating policies that will reelect them. DellaVigna (2009, p. 364), for example, argues that politicians are likely to be rational and constituents biased.

Formally, we write:

$$U = U[\theta \Delta \pi^*; (1 - \theta)\phi(\Delta \pi^{os})] \quad (1)$$

In (1), $\Delta \pi^*$ measures the distance the adopted policy, π^R , is from π^* . The regulator receives disutility as this distance increases, and places weight $\theta \leq 1$ on $\Delta \pi^*$. The term $\Delta \pi^{os}$ is the distance between π^{os} and the regulator's chosen policy, π^R . The function $\phi(\cdot)$ translates increases in $\Delta \pi^{os}$ into punishments for the regulator, such as lower budgets and more hearings. The regulator places a weight $(1 - \theta)$ on disutility due to political punishments. For tractability in what follows, we arbitrarily assume that the policy space from which the regulator chooses π^R is bounded from below by π^* , and from above by π^{os} . The distance between these end points is exogenously determined and assumed to be strictly positive. Thus, higher values of π^R translate into more populist policies. We adopt the following functional form of (1):

$$U = S - \frac{\theta}{2}(\pi^R - \pi^*) - \frac{(1 - \theta)}{2}\phi(\pi^{os} - \pi^R), \quad (2)$$

where S is the level of utility that could be realized if $\pi^R = \pi^* = \pi^{os}$. Solving the first order condition of (2) for π^R yields:

$$\pi^R = \lambda(\theta\pi^* + (1 - \theta)\phi\pi^{os}), \quad (3)$$

where $\lambda = \frac{1}{\theta + (1 - \theta)\phi}$.

Equation 3 shows that the regulator's optimal policy choice is a weighted average of the one she believes best for society's long-run interests and the one that garners the most short-run political support. In this manner, departures from the long-run efficiency benchmark π^* may enter policy through two channels: indirectly via political overseers' preferences, and directly through a biased regulator's preferences.

From (3) we see that a regulator will adopt the optimal policy choice if either she places no weight on political rewards ($\theta = 1$) or if the politician cannot translate constituency support into support (or punishment) for the regulator ($\phi = 0$). On the other hand, the regulator who weighs political advancement heavily ($\theta \approx 0$) or is subject to a powerful political overseer (ϕ large) is more likely to choose the politically expedient policy. In what follows, we use this simple framework to examine how commonly cited biases might affect policy outcomes.

2.1 Flawed heuristics and myopia

If regulators suffer from the biases that plague consumers, they are likely to use heuristics—or mental shortcuts—to estimate the optimal long-run policy choice, π^* . These shortcuts save time but may yield systematic decision-making errors. Much experimental research has documented the existence of several flawed heuristics, which are likely to bias regulators against accounting for long-run considerations when forming policy (Korobkin and Ulen 2000; Jolls 2007).

First, the *availability heuristic* leads people to overemphasize recent, particularly salient events when estimating the overall prevalence of those events. A person who recently witnessed his neighbor's house burn down, for example, is likely to overestimate the odds that his own house will burn in the future. Second, the *representativeness heuristic* causes people to form unduly high estimates of posterior probabilities by ignoring low baseline probabilities and small sample sizes. The canonical example of this bias comes from the [Kahneman and Tversky \(1982\)](#) experiment in which consumers are asked whether it was more likely that a hypothetical woman was a bank teller or a feminist bank teller. The former category clearly contains the latter, but most subjects placed a higher probability on the latter category. Third, people suffering from the *hindsight bias* tend to overestimate the ex ante probability of an event occurring, given that it has actually occurred. For example, a jury considering a negligence case may be too likely to find the defendant's actions were unreasonable ex ante because they know an accident resulted. Finally, *optimism bias* causes individuals to underestimate their own probability of experiencing a bad outcome.

Apart from flawed estimation strategies for unlikely events, regulators also may suffer from *myopia*, which arises from cognitive inability to process life-cycle costs or from self-control problems. [Gabaix and Laibson \(2006\)](#) model myopic consumers who cannot understand a durable good's full life-cycle cost. They find that even under competitive conditions, firms may lack incentives to educate myopic individuals and will offer low up-front prices and supra-competitive aftermarket prices in equilibrium. [Rachlinski and Farina \(2002\)](#) posit that myopia works not only in a temporal dimension but also across subject matters; when designing policies, experts may focus narrowly on their specific areas of responsibility and ignore spillover effects.

In a similar vein, some have examined the inability of actors to bind themselves to take future action that maximizes long-run utility, as viewed from the initial period. [Laibson \(1997\)](#) develops a theory of *time-inconsistent or hyperbolic discounting* to explain why consumers consistently contribute less to their retirement accounts than a rational actor model would predict. [Della Vigna and Malmendier \(2004\)](#) present a formal model in which agents plagued by hyperbolic discounting make multi-period decisions and derive an equilibrium in which firms may be able to take advantage of consumers who underestimate their future will power. They find support for their model in observed contracts and consumer behavior in the credit and health club markets.

Myopic regulators who use flawed heuristics are likely to make imprecise estimates of the optimal policy choice, π^* , which they use to form policy, π^R . In what follows we assume that these biases are likely to lead the regulator to favor policies that focus excessively on short-run considerations. Although we recognize that this assumption risks converting our model into a tautology, we argue that this assumption is grounded in reality and more plausible than the alternative that biases lead regulators to focus more sharply on optimal long-run policies.

[Jolls et al. \(1998, p. 1518\)](#), for example, argue that the availability bias gives rise to the "pollutant of the month syndrome," in which highly publicized events lead regulators to pursue overly stringent environmental regulation. Representativeness bias may cause a competition authority to ignore the prevalence of a business practice (e.g., exclusive dealing) in competitive industries and overlook the low prior odds that markets are noncompetitive, when it attempts to judge the probability that the practice will

reduce a market's competitiveness. Similarly, in the quasi-negligence determinations involved in certain consumer protection violations, hindsight bias is likely to cause an agency to look more skeptically on practices that led to harm ex post. Optimism bias may cause regulators to hold an unduly optimistic view of the likely success of a policy choice. More generally, to the extent that regulators are better characterized as "lay" rather than "expert" decisions-makers, [Rachlinski and Farina's](#) (2002, p. 572) warning that representativeness, availability bias, and loss aversion likely will cause Congress to focus on the wrong set of problems and solutions also applies to regulators.⁴

Returning to our formal model, we can write the regulator's estimate of π^* , $\hat{\pi}^* = \alpha\pi^*$, with $\alpha \geq 1$. In this case, she underestimates the gap between π^* and π^{os} , with $\hat{\pi}^*$ as the perceived lower bound in policy space. Consequently, the regulator will choose policy $\tilde{\pi}^R = \frac{1}{\lambda}(\theta\hat{\pi}^* + (1 - \theta)\phi\pi^{os})$, which is closer to π^{os} than π^R , the unbiased policy choice.

Because the effects of bounded rationality and a taste for political rewards work in the same direction, it is not possible in our model to identify these effects separately.⁵ Even an unbiased regulator has an incentive to choose populist policies due to the political rewards that come from instant action, especially with limited time horizons.

2.2 Status quo bias

Various cognitive errors together tend irrationally to wed people to the status quo. First, due to what is known as the *endowment effect*, experimental subjects seem to require more compensation to part with an endowment than they are willing to pay to gain it ([Korobkin and Ulen 2000](#), pp. 107–108). The observed willingness-to-accept (WTA)/willingness-to-pay (WTP) gap suggests that people are more averse to losing what they already possess than rational choice theory predicts.⁶ Second, *loss aversion* suggests that framing of decisions matters, because people place a higher negative value on expected losses than on expected gains of equivalent value. Relatedly, the *omission/commission bias* leads subjects to be concerned more with errors associated with action than inaction. These cognitive shortcomings create inertia to maintain a current course of action rather than take new action that would increase expected utility. Concern about this bias has inspired vigorous debate on proper default rules for consumer choices over retirement plans, insurance coverage, and privacy policies, and has raised questions about the efficiency of Coasian solutions to property rights issues (e.g., [Korobkin and Ulen 2000](#), pp. 1109–1111; [Camerer et al. 2003](#), pp. 1226–1230).

Following this argument, the class of cognitive shortcomings that make regulators reluctant to alter the status quo will tend to make policies "sticky" around initial policies. The direction in which the status quo bias will steer policy is indeterminate theoretically, and will depend on the initial policy endowment and the direction

⁴ [Rachlinski and Farina](#) (2002, p. 562) also contend that specialization from committees and learning from incumbency may ameliorate the effects of these biases.

⁵ This result flows directly from our assumption that biases tend to focus regulators on short-term, rather than long-term, solutions.

⁶ More technically, the endowment effect means that the marginal rate of substitution between two goods is a function of one's point on an indifference curve.




Fig. 1 Regulatory policy with flawed heuristics/myopia

of the changes in π^* and π^{OS} . For example, suppose that new learning suggests a more aggressive policy is needed to maximize consumer welfare, shown in Fig. 1 as a movement in the lower bound of policy space from π_0^* to π_1^* .

A rational regulator would adopt policy π_1^R in response to this change, but a regulator affected by loss aversion would move only to $\tilde{\pi}_1^R$ because she perceives the movement away from π_0^* as more expensive than she would have perceived prior to his endowment at policy π_0^R . Similar changes in π^{OS} would generate similarly muted responses. From this stickiness emerges a path dependency in policy choice, where policies adopted in the past have a lingering impact on future policy adoption.

2.3 Confirmation bias

A large body of experimental research suggests that individuals tend to become irrationally wedded to their early impressions about an initially ambiguous situation.⁷ This *confirmation bias* comes about either because subjects ignore all new evidence once they have made up their minds, or because they erroneously interpret evidence contradicting their beliefs as supporting their beliefs. In regulatory settings, confirmation bias leads to overconfidence in one's estimates of optimal policy. At the micro level, regulators may misread or ignore evidence that conflicts with the theory of a case or policy initiative. Consider the example of an antitrust agency team preparing to challenge a merger as likely to harm competition. The team may tend to interpret documents from the merging parties that, seen objectively, tend to cast the merger in a competitive light as either neutral or supportive of their case. As Jolls (2007, p. 15) notes, several experimental studies find that litigants tend to interpret ambiguous evidence as supporting their view of the case. This leads to the counterintuitive finding that greater information revelation can actually reduce the possibility of settlement. At the macro level, the regulator may misread evidence to confirm priors regarding larger policy choices, such as adopting an interventionist or laissez-faire attitude toward certain business practices.

⁷ Babcock and Loewenstein (1997) and Loewenstein and Moore (2004) document this bias in experimental settings involving litigants.

In the simple binary context, confirmation bias will have an asymmetric effect on policy outcomes; regulators with incorrect priors cause more harm than their counterparts who are irrationally wedded to the correct decision. For example, assume that the correct prior odds on vertical restraints is a laissez-faire posture: $\frac{P(A)}{P(C)} < 1$, where $P(A)$ and $P(C)$ are the probabilities that an observed vertical restraint is anticompetitive and procompetitive, respectively. A regulator with strong priors that vertical restraints are anticompetitive (i.e., $\frac{P(A)}{P(C)} < 1$) is likely to misinterpret evidence to confirm this belief, leading to welfare-reducing interventions. On the other hand, a regulator with correct priors may be too pessimistic about the odds that a given vertical practice is anticompetitive, but nonetheless makes the correct policy choice. Given a distribution of policy decisions, at the margin, a regulator with such a bias will bring too few vertical cases, but because her bias is toward the correct decision, overconfidence will have a smaller effect on the efficacy of ultimate policy choices.

If the regulator begins with a truly blank slate, the first information she receives about a matter will shape her bias when she interprets future evidence. In theory, there is no way to identify the direction of the bias. In practice, however, the first information a regulator is likely to see is some form of evidence supporting action (e.g., initiating an investigation or issuing a proposed rule); agency decision-makers are likely to learn of many policy issues only when staff or a political overseer requests intervention. An antitrust decision-maker often learns for the first time of a possible anticompetitive merger or business practice upon reading a staff memorandum that seeks compulsory process to investigate the matter. By their nature, these memoranda present a case for investigation. The investigation's targets are unlikely to present their views to the decision-maker until much later in the process. Similarly, the first piece of evidence can come from political overseers calling for an investigation into a practice. Here, again, the targets have no opportunity to counter the charges until the process is well advanced. If these requests for action become the anchoring point from which the regulator interprets subsequent evidence to estimate π^R , it will lead to an intervention bias even where the regulator weighs long-run welfare heavily (i.e., high θ).

In practice, we doubt that regulators begin with a truly blank slate. More realistically, regulators approach decisions with priors, or a "mental model" of how the world works. This model likely correlates with political beliefs, education, and experience (Wind and Crook 2009). Like the Bayesian "blank-slate" updating framework, the regulator screens out or discounts information that does not conform to her pre-existing world view in a desire to achieve policy outcomes consistent with her model. The key difference is that the regulator does not begin as an empty vessel when faced with each policy choice. Rather, she evaluates even the first piece of information with potential bias. Thus, an antitrust regulator who views markets as generally self-correcting is less likely to find evidence sufficient to support intervention than a counterpart who views business practices skeptically. Unlike the "blank slate" model, even if the first piece of information is a call for action from staff or a political overseer, the skeptical regulator will not process future information to confirm the call for action, but rather through her existing framework. Consequently, an intervention skeptic will require more "pro-intervention" information to reach a decision to intervene than a regulator operating with a truly blank slate.

Publicly stated positions also can anchor policy. Once a regulator takes a position on a particular policy, she will want to filter future matters in a way that verifies her initial wisdom. These public pronouncements will represent the regulator's true belief of π^R blended with the influence of political patrons—for example, promises made during confirmation hearings to assure support—to arrive at a π_0^R . In this way, the policy underlying the time 0 public announcement is not itself biased. Future policy decisions, however, are biased to the extent that future estimates of optimal long-run policy, $\hat{\pi}^*$, stem from confirmation bias anchored on the policy announced in time 0.⁸ As with loss aversion, this anchoring can lead to path dependency in regulatory policy adoption.

Confirmation bias also can reinforce preferences for short-sighted decisions that derive from the flawed heuristics and myopia introduced in Sect. 2.1. As discussed above, a myopic regulator is more likely than a rational regulator to invest suboptimally in policy development. This decision will form the regulator's priors, and non-confirming subsequent evidence will have a minimal effect on changing the regulator's world view. In this manner, myopia and confirmation bias can reinforce each another. When confirmation bias affects a boundedly rational regulator, a decision to initiate litigation or rulemaking due to insufficient consideration of long term costs and benefits will shape the interpretation of subsequent information in a way that tends to depict the decision favorably.

3 Will regulators suffer from biases in the long run?

A threshold question we have not addressed is how much regulators are actually likely to suffer from biases? On one hand, if regulatory institutions operate like firms there are reasons to believe that regulators in large part may escape the cognitive problems that plague consumer decision-making. The consensus within BE scholarship appears to be that firms are unlikely to make systematically biased decisions in the long run. DellaVigna (2009, p. 361) explains how consumers and firms differ:

Experience is the key difference. Unlike individual consumers, firms can specialize, hire consultants, and obtain feedback from large data sets and capital markets . . . Compared to consumers, therefore, firms are less likely to be affected by biases (except for principal-agent problems), and we expect them to be close to profit maximization.

Recent evidence also suggests that consumers who initially display biases can learn to overcome them with marketplace experience.⁹ Is it reasonable to assume that because regulators often are “experts” and face similar problems repeatedly they will be able to make unbiased policy decisions? Even if agency heads who make decisions are political appointees and not true field experts, as Rachlinski and Farina (2002, p. 579) note, agency career staff can provide expertise and experience.

⁸ That is $\hat{\pi}_1^* = f(\pi_0^R(\pi_0^*))$.

⁹ List (2003, 2004) finds evidence that the endowment effect fades as agents become more experienced traders.

The analogy between firms and regulatory institutions extends only so far. The feedback mechanism that facilitates learning differs significantly between firms and regulators. Unlike the marketplace, which quickly gives firms feedback in the form of prices, profits, and output, the link between public policy decisions and outcomes is more attenuated. Accurate measurement of a policy's welfare effects can be difficult and is somewhat rare. Even when effects are determined, the lag from a policy choice to policy execution can be long; cases and rule making can take several years from their initiation through final appeals in the courts. As the link between a decision and feedback weakens, the ability to learn diminishes.

Another consideration is that the costs for the regulator for being wrong are low compared to those of a firm. Generally speaking, a company that makes systematic mistakes is more likely to exit than a regulatory body that continually adopts welfare-reducing policies. The head of such an agency may continue to enjoy rewards as long as she increases outputs on margins that political overseers care about—action that maximizes the probability of reelection. When competition among regulatory bodies occurs, it likely assumes the form of activities (e.g., protection of jurisdictional boundaries, the output of observable policy interventions) that have no necessary correlation with positive welfare outcomes.

The weak connection between welfare and regulatory rewards can yield a cadre of regulators who are biased toward short-run politically expedient policies. It is obvious that the regulator with a low θ will tend to enjoy policy rewards. The regulator with a high θ , but who suffers from myopia or confirmation bias, however, may accidentally receive more support than an unbiased regulator with a similar taste for long-run welfare maximization. A biased regulator who sees herself as independent from political control nonetheless may bring a large number of cases that grab headlines but reduce welfare. Although viewed as welfare maximizing policy by the regulator, an objective observer would see these as politically expedient policies, which please the political overseer. Thus, regulators with short-term biases—due to high political weighting and bias—are likely to be over represented in the population of regulators.

In sum, it appears that the regulatory feedback structure makes it unlikely that regulators will learn to overcome their biases. This distinction between regulatory and market feedback is significant: marketplace performance directly measures consumer benefit from actions, whereas regulatory outputs have no necessary relationship to consumer welfare. Even if rewards were tied more closely with outcomes, the time lags and measurement problems discussed above make it difficult to link clear regulatory failures to specific regulators, especially given short tenures.

These observations warrant caution in assuming the efficacy of intervention to correct perceived firm biases. For example, [Stucke \(2007\)](#) and [Tor \(2002\)](#) argue that because firms may overestimate their chances of successful entry into a market, anti-trust authorities should place less confidence on entry as a means to ameliorate potential anticompetitive effects. [Tor and Rinner \(2011\)](#) propose that the rule of reason be applied more stringently to minimum resale price maintenance agreements to account for the possibility that biases lead firms to overestimate the profit-reducing effects of price competition. These normative prescriptions for enhanced intervention implicitly assume that the decisions of regulators to intervene will not themselves stem from biases. When these biases guide policy choices, it is not evident that regulators will

be able to intervene successfully to “correct” firms’ entry and vertical contracting decisions. Comparing the limitations of private and public decision making, market feedback is likely to correct biases more effectively than regulatory feedback.

The role of markets in correcting firm biases also suggests an important role for antitrust enforcement. Although regulatory biases may exacerbate already-existing tendencies for antitrust authorities to intervene in markets too often, a competitive market is a necessary condition for firms to correct their biases. If firms are to learn from the market, mistakes need to be costly. A rational firm generally will earn higher profits than a biased firm, conditional on the competitiveness of the market in which it operates.¹⁰ However, the feedback from poor decisions—and hence the incentive to correct biases—is stronger for a biased firm operating in a competitive market than one facing little competition. Sound competition enforcement directly ameliorates welfare losses from the illegal exercise of market power and indirectly may improve market performance in the spirit of Hayek by enhancing the information flows that firms need to identify and correct biases.

How much biases survive depends upon the costs of being wrong. When it is difficult to detect flawed decisions and the costs of being wrong are small, there is little incentive to invest in correcting biases. We next discuss how improved accountability—linking outcomes to rewards—can help ameliorate this problem.

4 Possible correcting mechanisms

As Sect. 2 shows, not all biases predict the same policy drift; some tend to temper others. For example, availability, representativeness, optimism, and myopia probably push regulators to adopt more politically expedient policies. Although we posit that information coming to regulators is likely to be calls for intervention, the effect of confirmation bias is uncertain, depending on the regulators’ priors, and the order and relative veracity of information flows. Finally, the status quo bias can stabilize policies around a focal point, but there is no a priori means to determine whether policy is stuck closer to a long-run optimal policy or a politically expedient one.

Imagine a regulator who suffers from all the biases we have cataloged and inherits an agency with a relatively interventionist policy stance. Assume also that her priors are weighted toward long-run welfare maximizing policies. The biases of availability, representativeness, myopia, and overconfidence could pull the regulator’s estimate of the optimal policy closer to the political overseers’ preferred position. Given her priors, however, the regulator would tend overly to discount calls from staff and political overseers calling for intervention, and to misinterpret evidence that supports intervention as evidence in favor of non-intervention. At the same time, because the status quo is assumed to be one of relative intervention, she may hesitate to disturb the agency’s current posture.

The point is that we cannot predict how a boundedly rational regulator suffering from a collection of biases will act. We can say that her actions will have no systematic

¹⁰ [Armstrong and Huck \(2010\)](#) describe instances where “irrational” firms may earn higher profits than rational ones.

correlation with optimal long-run policy; if they happen to mesh, it is due only to chance. Accordingly, if there is value to predictable policy, the elimination of regulatory bias is desirable. This section explores the potential of two strategies for reducing the harm associated with these biases: insulation and de-biasing.

4.1 Insulation

Through “insulation,” a “choice architect” eliminates (or makes more difficult to choose) poor alternatives (e.g., smoking, fatty foods, payday borrowing) from the choice set, thereby protecting the biased decision-maker from bad decisions. Much BE literature advocates this type of paternalistic solution in the context of flawed consumer decision-making. It is unclear how insulation would work to eliminate regulatory bias.

In practice, legislatures and courts limit regulatory discretion. For example, legislation directing agency action (e.g., rule-making or law enforcement) defines the scope—sometimes narrowly—of the agency’s authority to act. Some scholars have written on the need of legislatures to implement *ex ante* statutory constraints to curb *ex post* opportunism by regulators.¹¹ Courts also police agencies to assure fealty to congressional intent, the Administrative Procedure Act, and constitutional strictures. In theory, legislatures and courts could use the same tools to eliminate poor choices from the regulator’s set. To the extent that insulation strategies bind regulators to certain future actions to maximize welfare, they avoid suffering the issue of “consumer sovereignty” that plagues choice architecture directed at myopic consumers.

To employ these *ex ante* and *ex post* restraints to insulate agencies from poor regulatory choices, one must establish objectively “correct” choices. Doing so requires an unbiased choice architect. For the reasons detailed in Sect. 2, there is no reason to believe that legislatures could systematically identify and limit poor choices from the regulator’s set. Instead, they are likely to encourage them. As [Rachlinski and Farina \(2002, p. 578\)](#) note, the *seriatim* nature of judicial decision making, focused on particular parties and unique facts, is likely to diminish a court’s ability to escape cognitive bias through learning, particularly when dealing with complex regulations. Because courts stake out positions in their published decisions, they may be especially susceptible to confirmation bias anchored on their previous rulings.

Even if one could find an unbiased architect, the informational demands to establish the optimal future regulatory choice set may be insurmountable. Deciding to block or restrict certain regulatory paths is a far more complex task than placing the fruit in front of the fries in the cafeteria. Any insulation strategy involving *ex post* review by

¹¹ For example, [Knieps and Weib \(2007\)](#) note:

Since it is well-known that regulatory authorities cannot be forced into welfare-maximising behavior, the question arises whether opportunistic behavior can be excluded by the design and implementation of adequate hostages . . . Regulatory authorities as part of the bureaucracy cannot be fined for inadequate behavior. As a consequence, it is only by means of a statutory constraint that opportunist behavior by the regulatory authority can be disciplined.

courts almost necessarily would engage judges in ranking regulatory choices based on normative criteria. This contradicts established legal doctrines which direct courts to focus on agency adherence to procedural mandates, as well as constitutional and legislative restraints, rather than assess the wisdom of regulatory policy choices.¹²

4.2 De-biasing

A second way to limit the impact of regulatory bias is to take steps to enable boundedly rational agents to make decisions as rational agents. De-biasing can include the complete elimination of biases or measures that lead naive regulators (i.e., those who are unaware of their biases) to account for their biases when making choices. For example, [Jolls \(2007\)](#) notes that in some experiments, showing jurors pictures of Tiger Woods before deliberation may act to eliminate unconscious racial bias. Similarly, taking advantage of the availability heuristic by publicizing the most severe adverse effects associated with smoking may help counteract over-optimism bias in smoking decisions. Below, we examine two de-biasing strategies in the regulatory context: routine adversarial review of policies and making regulators more accountable for outcomes.

4.2.1 Adversarial review

As explained above, regulatory institutions are unlikely to feature competition that corrects biases. One way to generate needed feedback is to establish an internal adversarial process. A serious internal critique of regulatory proposals can help punish irrational policy choices before they take effect. To some extent, this replicates the negative feedback the marketplace provides firms that act irrationally. [Jolls and Sunstein \(2006\)](#) discuss research suggesting that outside directors on corporate boards help to ameliorate overly optimistic inside directors. As [Jolls \(2007\)](#) notes, some studies show that litigants who are forced to consider a case from their adversaries' perspectives are less likely to suffer from undue optimism or confirmation bias.

An agency could implement this "B-Team" approach by assigning staff members to act as defense counsel in a proposed case. This would provide an alternative narrative to explain documentary and statistical evidence. The adversary team may be biased as well, but anchoring would be in the opposite direction, forcing the two biases to compete. As the staff prepares a recommendation to intervene, the B-Team contemporaneously could assemble an internal memorandum that musters the best arguments against the case. The two internal teams also could take part in a mock-trial. A complementary method, used by the FTC, is to have the Bureau of Economics provide a recommendation separate from the attorney case handlers. The welfare-centered approach of economics provides a perspective distinct from the orientation of attorneys who have investigated a case or researched a proposed rule. Making economists subordinate to the lawyers in charge of cases or rulemaking mutes the de-biasing effect.

¹² See *Chevron v. Nat. Resources Defense Council, Inc.*, 467 U.S. 837 (1984); *Olsen v. Nebraska*, 313 U.S. 236, 243 (1941); *W. Coast Hotel Co. v. Parrish*, 300 U.S. 379 (1937); *Nebbia v. New York*, 291 U.S. 502, 537–538 (1934).

Ex ante review by external parties also may help to counteract biases in regulatory decision-making. The Office of Management and Budget must approve proposed rules and certain research projects (though not cases) before an agency can proceed.¹³ Ex ante peer review by a panel of experts, as now used in scientific policymaking, could also be imported into competition policy-making. Some statutes that require agency reports compel the agencies to consult with various other regulatory entities that are likely to have expertise in the subject area.¹⁴

In rulemaking, ex post judicial review may force agencies to engage in this type of de-biasing. Established legal doctrines require agencies to open their decision-making process to public participation, which acts to solicit multiple opposing viewpoints for most proposed regulations. Agencies also must show that they considered the record established during the rulemaking and articulate a plausible nexus between the rule ultimately adopted and the record evidence.¹⁵ As [Rachlinski and Farina \(2002, pp. 588–589\)](#) explain, “Having to assess the force of criticisms coming from a variety of perspectives, and craft a persuasive response to those criticisms that are (or may be viewed by a reviewing court as) significant, helps an agency to step outside of the decision-making process.” This ex post review, however, does not apply to internal decisions to engage in law enforcement or less formal policy initiatives.

Altering the decision-making structure also may ameliorate biases. Some evidence involving corporate governance suggests that multimember boards with heterogeneous priors will act more rationally than a unitary decision maker.¹⁶ Public agencies with bi-partisan boards (e.g., the FTC) may be less susceptible to biases than executive branch agencies with one decision-maker.

4.2.2 Greater accountability: focus on outcomes rather than outputs

The moral hazard that exists with respect to regulatory decision-making serves to maintain biases. Creating an institutional framework that makes regulators more accountable for the welfare effects of their policies may help provide the feedback needed to correct or ameliorate biases. As noted earlier, this is easier said than done. Measuring a policy’s welfare effects can be a daunting task. Nonetheless, if agencies devoted additional resources to perform ex post evaluations of interventions, with budgets and other rewards (e.g., fewer hearings, good publicity) tied more closely to these results rather than outputs, regulators might internalize more of the costs associated with their biases rather than merely imposing them upon consumers. As a complement to this policy, longer tenures for regulators would make it more difficult for them to obfuscate their connection with a failed policy.

¹³ [Rachlinski and Farina \(2002, pp. 597–598\)](#) criticize OMB review as merely a means to assure conformity with the President’s political agenda, rather than to satisfy objective welfare criteria.

¹⁴ The Postal Accountability and Enhancement Act, 39 U.S.C. Sec. 101 et seq., required the FTC to consult with GAO, the USPS, and the Treasury Department for its report.

¹⁵ See, e.g., *Business Roundtable v. Securities & Exchange Comm’n*, 647 F.3d 1144, 1148 (D.C. Cir. 2011) (internal citations omitted).

¹⁶ Of course, any gains from de-biasing must be weighed against the decision-making costs associated with a multi-member decision-making processes.

A focus on outcomes also can act as a commitment device to mitigate the effects of myopia. Consider a regulator who in period 1 faces a choice to announce a case or initiate a rule making prematurely. The publicity surrounding the announcement and the impression that “something is being done” garners the regulator immediate political benefits of b_1 . Because action is premature, however, it will provide benefits of only V_L to consumers in period 2. Alternatively, if the regulator waits until period 2 to decide on the policy—thus allowing more time for research and data collection—the decision will create benefits $V_H > V_L$ in period 2.¹⁷ Thus, if the regulator chooses to consume benefits b_1 in period 1, she will cause consumers to suffer harm $V_H - V_L$ in period 2. We assume that δ is the discount rate and that the regulator places weight θ on lost welfare $V_H - V_L$ and $\phi(1 - \theta)$ on b_1 , because immediate action is more likely to garner political support.

Following Della Vigna and Malmendier (2004), suppose that the regulator has a time inconsistent, quasi-hyperbolic discount rate: $\beta\delta^t$, where δ^t is the constant discount rate between adjacent future periods and $\beta < 1$ is the hyperbolic component that increases preferences for current over near-future consumption. Note that as t increases, the present-bias effect of β diminishes rapidly. Thus, the regulator’s time 0 utility can be written as: $U_0 - \beta \sum_{t=1}^{\infty} \delta^t U_t$. If we assume risk neutrality and substitute the payoffs the regulator will obtain in periods 1 and 2 for choosing the politically expedient policy (described above), the present value of the regulator’s utility in period 0 becomes:

$$U_0 + [\beta\delta(1 - \theta)\phi b - \beta\delta^2\theta(V_H - V_L)], \tag{7}$$

which simplifies to:

$$U_0 + [\beta\delta(1 - \theta)\phi b - \delta\theta(V_H - V_L)], \tag{8}$$

where U_0 is the baseline utility in time 0, and the expression

$\beta\delta[(1 - \theta)\phi b - \delta\theta(V_H - V_L)]$ is the net present value of choosing to commence the case prematurely in period 1. To maximize lifetime utility, the regulator in period 0 wants the regulator in period 1 to commence the case in period 1 only if:

$$\frac{(1 - \theta)\phi b_1}{\delta\theta} > (V_H - V_L). \tag{9}$$

At the beginning of period 1, however, $U_0 = (1 - \theta)\phi b_1$, and $U_1 = \beta\delta\theta(V_H - V_L)$. Thus, the same regulator will now face payoffs from commencing the case of:

$$(1 - \theta)\phi b_1 - \beta\delta\theta(V_H - V_L). \tag{10}$$

Accordingly, the regulator’s period 1 self will commence the action provided that:

$$\frac{(1 - \theta)\phi b_1}{\delta\theta\beta} > (V_H - V_L). \tag{11}$$

¹⁷ This decision may include taking no action.

Comparing the two decision rules, we see that hyperbolic discounting can lead to time inconsistent decisions in which the regulator at time zero does not want to see the case commence in period 1, but changes her mind when period 1 arrives with because the immediate benefits are now more attractive due to a more heavily discounted period 2 cost. Because $\frac{(1-\theta)\phi b_1}{\delta\theta\beta} > \frac{(1-\theta)\phi b_1}{\delta\theta}$ for $\beta < 1$, hyperbolic discounting clearly increases time-inconsistent decisions.

If the regulator had a time-consistent discount rate (i.e., $\beta = 1$), the appropriate rule for commencing action in period 1 would be the same as the rule for the hyperbolic-discounting regulator in time 0 (see above): $\frac{(1-\theta)\phi b_1}{\delta\theta} > (V_H - V_L)$. As expected, with or without hyperbolic discounting, the probability of bringing the case prematurely is a positive function of the political benefit from current action and the relative weight that the regulator places on political rewards versus social welfare. Hyperbolic discounting, however, implies both a greater propensity for time inconsistency and well as increased incentives to commence the case prematurely relative to an unbiased regulator.

A framework that tied rewards to outcomes, rather than outputs, would act to increase θ because the welfare loss would not only provide internal disutility but would reduce future rewards. In theory, if the new θ exactly counterbalanced β , the regulator would act as if she were completely de-biased. A focus on outcomes in period 2 can help make the period 0 regulator credibly commit to bringing a case only when $\frac{(1-\theta)\phi b_1}{\delta\theta} > (V_H - V_L)$. Of course, this de-biasing strategy does not lead to optimal long run policies. As long as the regulator values political rewards, she will always have some incentives to take premature action. Only strategies that also reduce political influence (ϕ) or increase tastes for long-run welfare (θ) can reduce the public choice influence on policy making.

5 Conclusion

Much BE research prescribes increased regulatory intervention to constrain consumer choice in response to consumer biases and to expand use of competition law to correct consumer harm that arises from biased firm behavior. If regulators, who are human after all, suffer from the same biases, our analysis suggests a greater skepticism of these calls for increased intervention. Although regulators are likely to have an edge over consumers in terms of experience and expertise, it is unclear that they will be able to intervene in ways that systematically improve welfare. The model we present shows that political pressure will cause rational regulators to choose policies that are not optimal from a consumer standpoint, and that in many circumstances regulatory bias will exacerbate this tendency. We also suggest special caution when trying to correct firm behavior. Regulatory bias seems likely to endure longer than firm bias because the market provides a stronger corrective feedback mechanism than exists in regulatory policy-making. To the extent that we can de-bias regulators—either through greater use of internal and external adversarial review or by creating a closer nexus between outcomes and rewards—they will become more effective in devising welfare-enhancing interventions designed to correct biases of consumers or firms.

Acknowledgments We thank participants in the CRRRI Eastern and Western Conferences, and an anonymous referee for helpful comments. We are grateful to Angela Diveley for superb research assistance.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Armstrong, M., & Huck, S. (2010). Behavioral economics as applied to firms: A primer. *Competition Policy International*, 6(1), 3–45.
- Babcock, L., & Loewenstein, G. (1997). Explaining bargain impasse: The role of self-serving biases. *Journal of Economic Perspectives*, 11(1), 109–126.
- Bennett, M., et al. (2010). What does behavioral economics mean for competition policy?. *Competition Policy International*, 6(1), 111–137.
- Camerer, C., et al. (2003). Regulation for conservatives: Behavioral economics and the case for “asymmetric paternalism”. *University of Pennsylvania Law Review*, 151, 1211–1253.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2), 315–372.
- Della Vigna, S., & Malmendier, U. (2004). Contract design and self-control: Theory and evidence. *Quarterly Journal of Economics*, 119(2), 353–402.
- Gabaix, X., & Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, 121(2), 505–540.
- Ginsburg, D. H., & Moore, D. W. (2010). The future of behavioral economics in antitrust jurisprudence. *Competition Policy International*, 6(1), 89–98.
- Jolls, C. (2007). Behavioral law and economics. In P. Diamond & H. Vartiainen (Eds.), *Behavioral economics and its applications* (pp. 115–144). Princeton, NJ: Princeton University Press.
- Jolls, C., & Sunstein, C. R. (2006). Debiasing through law. *Journal of Legal Studies*, 35(1), 199–242.
- Jolls, C., Sunstein, C. R., & Thaler, R. (1998). A behavioral approach to law and economics. *Stanford Law Review*, 50(5), 1471–1549.
- Kahneman, D., & Tversky, A. (1982). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 84–98). New York: Cambridge University Press.
- Klick, J., & Mitchell, G. (2006). Government regulation of irrationality. *Minnesota Law Review*, 90(6), 1620–1663.
- Knieps, G., & Weib, H.-J. (2007). Reduction of regulatory risk: A network economic approach. Institut für Verkehrswissenschaft und Regionalpolitik discussion paper No. 117.
- Korobkin, R. B., & Ulen, T. S. (2000). Law and behavioral science: Removing the rationality assumption from law and economics. *California Law Review*, 88(4), 1051–1144.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2), 443–477.
- List, J. A. (2003). Does market experience eliminate market anomalies?. *Quarterly Journal of Economics*, 118(1), 41–71.
- List, J. A. (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, 72(2), 615–625.
- Loewenstein, G., & Moore, D. A. (2004). When ignorance is bliss: Information exchange and inefficiency in bargaining. *Journal of Legal Studies*, 33, 37–58.
- Rachlinski, J. J., & Farina, C. R. (2002). Cognitive psychology and optimal government design. *Cornell Law Review*, 87(2), 549–615.
- Stucke, M. E. (2007). Behavioral economics at the gate: Antitrust in the twenty-first century. *Loyola University Chicago Law Journal*, 38(3), 513–592.
- Tor, A. (2002). The fable of entry: Bounded rationality, market discipline, and legal policy. *Michigan Law Review*, 101(2), 482–568.
- Tor, A., & Rinner, W. J. (2011). Behavioral antitrust: A new approach to the rule of reason after *Leegin*. *University of Illinois Law Review*, 2011(3), 805–864.

- Wind, J., & Crook, C. (2009). From mental models to transformation: Overcoming inhibitors to change. *Rotman Magazine*.
- Wright, J. D. (2007). Behavioral law and economics, paternalism, and consumer contracts: An empirical perspective. *New York University Journal of Law & Liberty*, 2(3), 470–511.
- Wright, J. D., & Stone, II, J. E. (2011). Misbehavioral economics: The case against behavioral antitrust. *Cardozo Law Review* (forthcoming).