

Do Administrative Procedures Fix Cognitive Biases?

Benjamin Minhao Chen* and Brian Libgober†

*University of Hong Kong, Hong Kong SAR

†Northwestern University, Evanston, IL, USA

Address correspondence to the author at benched@hku.hk.

Abstract

This article uses survey experiments to assess whether administrative procedures fix cognitive bias. We focus on two procedural requirements: qualitative reason-giving and quantitative cost-benefit analysis (CBA). Both requirements are now firmly entrenched in the U.S. federal regulation-making. Multilateral organizations such as the World Bank, OECD, and EU have encouraged their broad diffusion across many national contexts. Yet CBA, in particular, remains controversial. Supporters of CBA claim it leads to more rational regulation, with Sunstein (2000) asserting that CBA can reduce cognitive biases. By contrast, we argue that procedures should be conceptualized as imperfect substitutes subject to diminishing marginal benefits. To test and illustrate this argument, we examine how each procedure individually and cumulatively modulates gain–loss framing, partisan-motivated reasoning, and scope insensitivity in a nationally representative sample. We find that one or both procedures decreased cognitive bias. CBA was the most helpful against partisan reasoning whereas reason-giving did little. Both procedures seem comparably effective against the other biases, although in each case only one of the two procedures produced reductions distinguishable from zero. We only find substantial synergies between the two procedures with respect to gain–loss framing. For the other biases, the combination of both procedures did not generate significant reductions over and above that achieved by the more effective procedure alone. We hypothesize that procedures will only fix cognitive biases if they disrupt bias-inducing mental processes, and we reconcile this proposition with our findings. We conclude by relating our framework to debates about the design of administrative procedures and describe a research agenda based upon rationality-improving procedures.

Introduction

When does procedure improve decision making? This fundamental question has attracted attention from many disciplines and produced important insights. Yet, public bodies are often slow to reform their decision processes. Indeed, systemic procedural reform has occurred very rarely in the history of the US federal government. Arguably, the most significant such change in recent decades has been a series of executive orders mandating a cost-benefit test for major regulations (Sunstein 2018). These orders require administrative agencies to identify the costs and benefits of competing rules and quantify them if possible. Agencies must select the rules that maximize net social benefits and submit them for review by the Office of Management and Budget, accompanied by regulatory impact analyses. The structured analytical approach imposed by these executive orders stands in contrast to the open-ended reason-giving requirements of the 1946 Administrative Procedure Act. Initially branded an anti-regulatory power grab by a conservative President, cost-benefit analysis (CBA) has been affirmed by successive presidential administrations and is now entrenched in federal policymaking (Revesz and Livermore 2008).

Despite the depth and sophistication of the literature on CBA, the benefits of CBA remain speculative and anecdotal. The need to fill this lacuna is urgent, since CBA and its variants are recommended by influential multilateral organizations

such as the World Bank, the OECD, and the European Union, and have diffused to other national systems (Wiener 2013). Previous studies found that CBA improved perceptions of policymaker legitimacy (Stiglitz 2017) and induced greater support for life-saving regulation (Chen 2020). That said, the principal argument for CBA is that it produces substantively better decisions, not merely better-accepted ones (Adler and Posner 2006; Revesz and Livermore 2008). Why would CBA lead to better policy choices? CBA proponents claim that the procedure can protect policymaking from failings in human reasoning and judgment (Mackie, Worsley, and Eliasson 2014; Sunstein 2000). For example, Cass Sunstein—former Administrator of OIRA and an energetic promoter of CBA—claims that CBA can serve as an antidote for “cognitive biases” (Sunstein 2000). The term cognitive bias here refers to a reasoning failure that manifests predictably as a behavioral deviation from the axioms and predictions of rational choice. For instance, people tend to overestimate the probability of events that they can more easily recall, a phenomenon known as availability bias (Tversky and Kahneman 1973). People also are more willing to take risks when the available options are presented in terms of avoiding losses rather than achieving gains (Tversky and Kahneman 1981). It is plausible that CBA can help fix such cognitive biases, but the proposition has not received much empirical scrutiny until now.¹ Does CBA really produce the benefits touted by some of its most fervent advocates? If so, how do those benefits compare with those of other procedures such as reason-giving?

¹Our review of the 422 publications citing Sunstein (2000) did not turn up any research testing the core claim made in that article, which has also been repeated in several other venues.

We are grateful to Rick Brown for excellent research assistance, and to Jamie Druckman, Daniel Carpenter, Jed Stiglitz, the editors Susan Miller and Mary K. Feeney, and three anonymous reviewers for valuable feedback and suggestions. Publication made possible in part by support from the HKU Libraries Open Access Author Fund sponsored by the HKU Libraries.

In this article, we articulate two principles for thinking about the impact of procedural interventions on bureaucratic rationality: imperfect substitutability and declining marginal benefit. By imperfect substitutability, we mean that while various procedural requirements could improve public decision making, some of them might be more suited to certain kinds of rationality failures than others. By declining marginal benefit, we mean that additional procedures usually bring lesser incremental improvements. These principles imply that we should expect the use of CBA in tandem with reason-giving to help with cognitive bias more than either procedure alone, but not necessarily much more, and perhaps not enough to justify the marginal costs. In particular, we should not generally expect CBA to do much more to fix cognitive bias when reason-giving is already a well-established requirement. At the same time, these principles raise the possibility that CBA can mitigate certain kinds of cognitive biases that reason-giving cannot.

To illustrate these principles, we design and implement several factorial experiments testing how CBA and reason-giving perform against three well-known cognitive biases: gain-loss framing, partisan-motivated reasoning, and scope insensitivity. We expect CBA and reason-giving to perform similarly, if inconsistently, in reducing cognitive bias, and that is what we find. Partisan-motivated reasoning is the only bias we study for which one procedure is clearly no substitute for the other. CBA significantly reduced support for a policy attributed to co-partisans, whereas reason giving had almost no effect on partisan judgments of policies based on their sources.² For gain-loss framing and scope sensitivity, we find evidence that procedures can temper cognitive bias but no evidence that one procedure is superior to the other. Reason-giving appears somewhat stronger against scope sensitivity while CBA appears somewhat stronger against gain-loss framing. In each case, only the stronger procedure significantly improved cognitive bias in a statistical sense. Even so, there is at least some evidence that the less-effective procedure helped too, and we cannot reliably conclude that one procedure is inferior to the other for these two biases. We only find substantial complementarity between the two procedures with respect to gain-loss framing. For the other two biases, layering the less-effective procedure over the more-effective one did not produce substantial improvements. However, the complementarity we identify between procedures accounts for the apparent efficacy of CBA against gain-loss framing. Unless accompanied by reason giving, CBA did relatively little against gain-loss framing bias. Finally, although the reductions in cognitive bias are statistically significant in many cases, gain-loss framing and partisan-motivated reasoning clearly persisted no matter which procedures are employed.³ The “fix” that procedures provide for cognitive biases appears quite imperfect in these experiments.

When interpreting these results, one should acknowledge what the experiments can and cannot show. Our experimental subjects are a nationally representative sample of ordinary citizens and not a group of elite policymakers. Moreover,

the implementation of CBA and reason-giving in our experimental instruments is much simpler and less taxing than the procedural hurdles policymakers encounter when they set out to regulate. These discrepancies, among others, caution against generalizing beyond the experimental setting. We do not take our experiments to measure the precise effect of procedural interventions on real policymaking anymore than [Tversky and Kahneman \(1983\)](#) can be taken as estimating the true impact of gain-loss framing on actual public health strategies. Rather, the experiments illustrate our thematic argument: that once administrative procedures are conceived as multifactorial interventions for improving decisional outcomes, they should also be analyzed as such.⁴ Adopting this perspective leads us to propose the principles of imperfect substitutability and declining marginal contribution, and the experiments we present substantiate their relevance.

Ultimately, we hope that the principles and findings discussed here can stimulate further debate about the costs and benefits of administrative procedure. Instead of assuming that more deliberation or accountability necessarily enhances the rationality of policymakers, we can enrich our understanding of procedural interventions by thinking more deeply about how they counteract the many diverse threats to good decision making in the public sphere. For example, it is worth theorizing about why CBA reduced partisan-motivated reasoning in our experiment but reason giving did not. Our discussion starts from the premise that procedure mitigates cognitive biases by disrupting the mental processes giving rise to them. We then draw on existing research in cognitive psychology about the nature of thought processes that induce bias and to generate finer hypotheses about which procedures will be effective against each bias. In particular, we suggest that CBA is more powerful than reason giving at combating biases that operate by distorting our evaluation of information or evidence. Otherwise, reason giving appears just as good at combating biases that stem from the use of affective heuristics. Thinking about why procedures do or do not fix cognitive biases in the experimental setting can generate theories that have broader relevance and validity and thereby inform the design of procedures for improving public decision making.

Background and Context

We briefly survey both prior procedural interventions aimed at improving bureaucratic rationality and the social science literature on cognitive biases. Our treatment of prior interventions emphasizes the American case, but we also discuss the diffusion and influence of the American model in comparative contexts. We argue that qualitative reason giving and quantitative CBA are distinct reforms which seek to make regulators act more rationally. However, at least in the United States, reason giving has been implemented more widely and with less controversy than CBA. Further, rationality is a normative ideal and also a contested concept. In discussing the social-scientific literature on cognitive biases, we explain and

²However, these benefits seem to flow from trying to do CBA rather than from actually doing it correctly.

³Because there is no natural benchmark for scope sensitivity, we cannot say that this bias clearly persists. There is, however, good reason to believe that respondents are still not sufficiently attentive to the magnitude of the benefits and costs at stake.

⁴[Bagley \(2019\)](#) is an excellent, recent piece of scholarship that conceptualizes administrative procedures as multifaceted policy interventions attended by trade-offs. Although the tenor of [Bagley \(2019\)](#) is similar to ours, our argument and perspective are more descriptive than normative, focus on the design and evaluation of procedural interventions, and are supported by a number of experiments.

justify our attention to this particular family of deviations from rational choice theory.

Making Regulators Reasonable: Procedural Tools

The history of policy interventions in administration in the United States is also a history of attempts to make bureaucrats more rational (e.g. Mashaw 2018). The first significant federal intervention in administration, the Pendleton Act, is conventionally understood as having had the purpose of improving rationality in hiring, firing, and promotion decisions (Moynihan 2004). In the early 20th century, executive agencies gained increasing policymaking authority and power (Carpenter 2001; Ernst 2014). However, their decision-making processes remained deeply problematic. One of the most important festering problems was the limited analytical and evidentiary basis for agency decisions (Gellhorn 1986). In the 1930s, as New Deal Democrats embarked on new and bold experiments in regulating the economy, demands for further interventions to promote consistency and rationality in agency decisions reached a fever pitch (Ernst 2014; Shapiro 1986).

The Administrative Procedure Act of 1946 (APA) was the culmination of these debates. As with the Pendleton Act, the political science literature suggests that there were self-serving political interests behind the APA (McNollgast 1999). Even so, improving bureaucratic rationality was a frequently articulated goal of reform, and one cannot deny the salience of this motive for some elite decision-makers.⁵ The APA prescribes uniform standards for agency rulemaking. Agencies may make rules through a formal, trial-like procedure (5 USC 556) or through an informal, notice-and-comment procedure (5 USC 553). Generally speaking, most regulations today are promulgated through notice-and-comment rulemaking. Under 5 USC 553, agencies must furnish notice of the proposed rulemaking, “give interested persons an opportunity to participate in the rule making through submission of written data, views, or arguments,” and “incorporate in the rules adopted a concise general statement of their basis and purpose.” 5 USC 706 further instructs courts to “hold unlawful and set aside agency action, findings, and conclusions” that are “arbitrary [and] capricious.” The constraints imposed on agencies by the APA are relatively onerous. Congress may set out the background and purpose of a law in a statutory preamble. However, it is not required to do so. Moreover, while judges give opinions elaborating upon the basis for their decisions, “a lower court g[iving] the wrong reasons for a correct decision is not by itself a justification for reversal” (Mashaw 2018). In a sense, the APA requires agencies to go beyond judicial and legislative standards of reasonability when making rules for the public.

For several reasons, the immediate impact of the APA on regulatory behavior was limited. In the 1940s and 1950s, agencies mostly proceeded through adjudication rather than rulemaking. That is, they rendered determinations in individual cases rather than formulating norms applicable to the general public. In that context, courts were not punctilious about the justification requirement (Mashaw 2018). However,

powerful social movements in the 1960s and 1970s gave rise to broad delegations of power to agencies (Jones, Theriault, and Whyman 2019) and a correspondingly more prominent role for agencies as national policymakers. Courts responded by giving teeth to the strictures of the APA (Stephenson 2006). To ensure their regulations survive judicial review, agencies have built increasingly voluminous records of what they are doing and why. There is a lively dispute about the ills of excessive reason-giving. Some argue that too much explanation can slow down or “ossify” the rulemaking process (McGarity 1992) while others insist that such concerns are more theoretical than real (Yackee and Yackee 2010). Still, the animating idea behind the APA—that agencies should provide reasons when they regulate—is largely uncontroversial.

Uncontroversial, however, is not a label one could use to describe the rise of CBA, a kind of regulatory impact assessment used in rulemaking. DeMuth (2016) traces the rise of regulatory impact analysis to budgeting systems implemented in Robert McNamara’s Defense Department which were later extended to other agencies and state governments. Although these policy-planning systems were ultimately abandoned, similar quantitative approaches were mandated by the National Environmental Policy Act (NEPA) and the “quality of life” reviews applied to environmental, health, and safety regulations by President Nixon. These reviews did not require quantification of every cost and benefit, nor was an executive veto power built into the system. Indeed, even the environmental impact statements required by NEPA did not substantively restrict agency action.

President Reagan’s EO 12,291 represented a watershed in the use of CBA within executive agencies. In particular, EO 12,291 required agencies to select the least costly regulatory alternative identified by the agency’s CBA unless that option was prohibited by law. It also empowered the Office of Management and Budget to indefinitely delay the publication of regulations it did not find acceptable. Some regarded this development as an anti-regulatory power-grab. The read of Morrison (1986) is typical: “[the] system of OMB control imposes costly delays that are paid for through the decreased health and safety of the American public.” Even so, there were more than a few defenders of the new economic approach, even on the left (DeMuth 2016).

Despite the initial controversy over Reagan’s executive order, Democratic presidents since Reagan have retained the centralized regulatory review process and entrenched it. President Clinton’s Executive Order 12,866, which replaced Reagan’s EO 12,291, essentially affirmed the cost-benefit paradigm while making allowances for “distributive impacts” and “equity” and moving away from “least-cost” toward “maximum net-benefit” (Sunstein 2018). President Obama identified “human dignity” and “fairness” as difficult or impossible to quantify *desiderata*. Importantly, Obama’s executive order also encouraged the review of existing regulations. President Biden has recently reaffirmed the Clinton and Obama Executive Orders while urging attention to “disadvantaged, vulnerable or marginalized communities.” Although there are some differences between the approach of antiregulatory Republican and proregulatory Democratic administrations, Sunstein (2018) emphasizes that these differences are subtle and have more to do with how administrations exercise their powers rather than institution of centralized review itself. At the same time, vigorous debate about CBA has continued,

⁵To take one of many examples, both the Pound and Landis Commission reports, which influenced the APA, discuss concerns about rationality at length.

with some arguing that CBA is manipulable for political ends (Cole 2012) and other times a smoke screen for decisions made on other grounds (Adler and Posner 1999).

Procedural interventions are of course not limited to the United States. De Francesco (2021) relates the diffusion of regulatory innovations through advanced industrial democracies. According to one account, CBA has diffused more readily than APA-like procedures, especially in parliamentary systems (Baum, Jensen, and McGrath 2016). Multilateral organizations such as the World Bank, EU, and OECD have promoted regulatory impact analyses by, for example, funding the modernization of public governance, establishing transnational working groups such as the Directors and Experts on Better Regulation, and publishing naming-and-shaming reports on the progress of member states in implementing “good regulatory governance” reforms (De Francesco, Radaelli, and Troeger 2012). The World Bank has extended such reporting beyond advanced industrial democracies to countries in the Global South which may thereby face increasing pressure to overhaul their regulatory processes. Against this backdrop, enhancing public deliberation and reason-giving and mandating regulatory impact analyses are two models that developing nations could emulate.

Cognitive Biases

Although American bureaucratic reforms often invoke the ideal of rationality, rationality is a contested concept with multiple meanings. Simon (1976), for example, distinguishes between substantive and procedural rationality. A choice is substantively rational if it furthers the decision-maker’s ends. A choice is procedurally rational if it is made following the proper steps. Bureaucratic reforms along the procedural axis are frequently intended to produce benefits on the substantive axis. Put differently, many who design administrative rules believe that adherence to procedure will produce better public outcomes, however defined. Within social science, the dominant account of substantive rationality is predicated on the von Neumann and Morgenstern axioms which permit mathematical analysis of decision problems using utility functions, expectation operators, and the tools of calculus. There are multiple and rich vocabularies for describing how human decision-making deviates from the prescriptions of classical decision-making theory. There is also a vast literature examining the role of affect or emotions with respect to cognition (Redlawsk 2002; Brader and Marcus 2013; Ladd and Lenz 2008; Lodge and Taber 2013; Marcus, Neumann, and MacKuen 2000), reporting violations of the axioms of von Neumann-Morgenstern expected utility theory, and developing alternative frameworks such as bounded rationality.

CBA and reason-giving are procedural interventions that might help bureaucratic decision-makers approximate the ideal of an expected utility maximizer by forcing them to confront alternatives that might otherwise be obscured by fear, laziness, motivated reasoning, and intuitive thinking, among others. While there are many elements and factors that could upset the classical model of rational choice, whether and how they manifest is often speculative. Perhaps for this reason, scholars interested in the design of administrative institutions have tended to focus on systematic, predictable deviations from rationality, such as framing effects (Tversky and Kahneman 1981). Framing is one of the paradigmatic examples of cognitive bias. In a ground-breaking article in *Science*, Tversky and Kahneman

demonstrated that whether risky options were described as potential gains or potential losses had a large, consistent effect on the respondents’ choices: respondents were more willing to gamble when contemplating losses rather than gains. This phenomenon was interpreted as a cognitive bias because the loss and gain frames are informationally equivalent. The shift in preferences or judgments is induced merely by the way the choices are described.

Tversky and Kahneman’s conception of cognitive biases precipitated a lively research agenda to catalog and explain them. In many cases, cognitive bias seems to arise from the use of heuristics to optimize mental resources in the face of complexity. Particularly relevant for politics and public administration are biases caused by motivated reasoning, especially partisan motivated reasoning. Motivated reasoning occurs when individuals’ desires for a particular outcome shape their beliefs and judgments (Lodge and Taber 2000). Political ideology or identity can cause individuals to seek out corroborating information, rate confirming information more highly, and resist disconfirming evidence (Druckman, Leeper, and Slothuus 2018).

The so-called cognitive argument for CBA holds that CBA can mitigate the impact of cognitive biases on policymakers (Sunstein, 2000, 2002; Christiansen 2018). Sunstein (2000), for example, enumerates six cognitive biases CBA might help alleviate. The biases he identifies tend to involve overdependence on highly available information, inappropriate reliance on social cues and priors, and difficulties in grasping the magnitude of benefits, costs, and probabilities, especially in an integrative way. Although some institutional scholars are enthusiastic about the promise of procedural reforms to enhance substantive rationality, research in cognitive psychology gives reason for pause. Lerner and Tetlock (1999) describes conditions under which accountability devices do not attenuate cognitive biases and may even exacerbate them. As summarized by them, “predecisional accountability to an unknown audience will have no effect on bias if, even after increased attention to one’s decision process, no new ways of solving the problem come into awareness.”⁶ Accountability can worsen bias to the extent that one choice appears easier to justify than another. Accountability can also incentivize greater risk-taking by political actors who have to answer to an electorate (Sheffer and Loewen 2019). While there have been quite a number of studies looking at the relationship between accountability and bias, few examine CBA as an accountability device or explore the comparative or cumulative efficacy of accountability devices (Lerner and Tetlock 1999).

Experimental Approach

Broadly speaking, our goal is to assess the efficacy of two procedural devices for improving human decision making. Survey experiments identify cognitive biases through random manipulations embedded in the instrument fielded to respondents. Suppose Y_i is a dichotomous policy choice and τ_i is a bias-inducing manipulation expected to influence that choice, contrary to the predictions of classical theory. The existence of a cognitive bias can then be demonstrated by fitting

⁶Bolsen, Druckman, and Cook (2014) find, however, that urging individuals to evaluate policy “in an even-handed way” and telling them they will be asked “to justify the reasons for [their] judgment” substantially decreased the influence of partisan cues on attitudes toward the 2007 Energy Act.

Table 1 Summary of Design for Three Studies

	Gain-Loss Framing*	Partisan Reasoning	Scope Insensitivity
Platform		Lucid—Nationally Representative Sample	
Structure		2 ³ (Imbalanced) Factorial	
Dates	March 2–11, 2022	April 7–13, 2021	April 11–29, 2021
Sample sizes	<i>n</i> = 1,975	<i>n</i> = 4,679 [‡]	<i>n</i> = 1,171
Textual prompt	See figure 2(a)	See figure 2(b)	See figure 2(c)
Outcome	Prefers less risky Program A	Supports Highway Program	Supports Crop Insurance Program
Cognitive treatment (+)	Policy saves lives	Copartisans support bill	Policy creates 2,500 jobs
Cognitive treatment (–)	Policy prevents deaths	Opposing party supports bill	Policy creates 500 jobs
Cognitive treatment effect	Increase support for less risky policy	Increase support for policy	Leave support unchanged
Null hypotheses			
Assigning reason giving does not change the cognitive treatment effect			
Assigning CBA does not change the cognitive treatment effect			
Assigning CBA and reason giving together does not change the cognitive treatment effect			

*These details correspond to a second experimental run requested by a reviewer. The first run of this experiment conducted between April 5–10, 2021 had *n* = 2,537, but suffered from attrition. Differences in estimates between runs are minor and probably attributable to sampling variation, as discussed in more detail below and in [Supplementary Appendix](#).

‡Platform pricing rules led us to also recruit *n* = 826 nonpartisans, who completed the study but were dropped, as stated in our preregistration.

the following OLS model and showing that the estimate on β_1 is statistically different from 0.

$$Y_i = \beta_0 + \beta_1\tau_i + \gamma \cdot \mathbf{X}_i + \epsilon_i.$$

The variable \mathbf{X}_i refers to a set of nontreatment, independent variables that might account for the variance in individual policy choices. \mathbf{X}_i can be omitted to improve interpretability given random assignment to control and treatment groups. Without these other independent variables, β_0 is the proportion choosing the policy under the control condition and β_1 is the increase in the proportion choosing the policy due to assignment of τ or, more technically, the sample average treatment effect of τ . Here, we are primarily interested in whether procedural interventions can reduce cognitive biases. In other words, we want to know whether a given procedural treatment π has a significant interaction with the bias-inducing treatment τ . We estimate the following OLS model and test whether β_3 is statistically distinguishable from 0.

$$Y_i = \beta_0 + \beta_1\tau_i + \beta_2\pi_i + \beta_3\tau_i\pi_i + \gamma \cdot \mathbf{X}_i + \epsilon_i$$

$$H_0 : \beta_3 = 0$$

With the general empirical approach and estimation strategy laid out, we now turn to the design of our experiments and details on implementation. [Table 1](#) presents a concise summary.

The cognitive biases we study are gain–loss framing, partisan motivated reasoning, and scope insensitivity. These three biases have been shown to be strong and robust. While they are well-documented in the literature on public administration and policy ([Baekgaard and Serritzlew 2016](#); [Olsen 2015](#)), we still know very little about how to counteract them.⁷ To

recapitulate, gain–loss framing is the tendency to choose the gamble rather than the sure thing when the options are described in terms of potential losses and vice versa when the options are described in terms of potential gains. Partisan motivated reasoning, on the contrary, is the tendency to support or oppose policies based on party endorsements. Finally, scope insensitivity is the tendency for choices to be unresponsive to differences in the scale of the problem or its proposed solution ([Kahneman et al. 1999](#)). For gain–loss framing and partisan motivated reasoning, the cognitive bias manifests as a shift in policy choices. By contrast, the cognitive bias in scope insensitivity manifests as an absence of such a shift. Even so, the null hypothesis of no interaction between procedures and the cognitive treatment is the same across biases and procedural interventions.

[Figure 1](#) displays the scenarios used for gain–loss framing, partisan reasoning, and scope insensitivity, respectively. For gain–loss framing, we employed a version of [Tversky and Kahneman \(1981\)](#)'s so-called Asian Disease Problem. The Asian Disease Problem has been replicated many times by the ensuing literature on framing effects ([Druckman 2001](#)) and its adoption here facilitates comparison to these prior studies. The only important difference between the original version and ours is that we did not describe the disease as originating in Asia.

Because there is not an equally canonical example of partisan motivated reasoning, we developed our own. First, we identified a policy proposal that could plausibly be attributed to either Democrats or Republicans. Major infrastructure bills have been discussed as high priorities by recent Republican and Democratic administrations. We performed a pilot study on Amazon Mechanical Turk to test for partisan motivated reasoning in this scenario. We found that 76% of respondents supported the infrastructure bill when it was ascribed to their party as opposed to 44% when the bill was ascribed to the opposite party. This difference evinced partisan motivated reasoning.

⁷See [Battaglio et al. \(2019\)](#) for a survey of the literature on cognitive biases in the context of public decision making. A recent contribution that examines whether justificatory requirements can temper motivated reasoning is [Christensen and Moynihan \(2020\)](#). To the best of our knowledge, no previous study has looked at whether CBA can mitigate cognitive biases.

(a) Gain-Loss Framing	(b) Partisan-motivated Reasoning	(c) Scope Insensitivity
<p>A novel disease has broken out in the United States. If unaddressed, the expected death toll from the disease is 600. Two programs to combat the disease have been proposed. The first program, Program A, uses conventional medicine to combat the disease. Under Program A, {200 people will be saved/400 people will die}. The second program, Program B, uses experimental medicine to combat the disease. Under Program B, there is a 33% chance that {600 people will be saved/0 people will die} and a 66% chance that {0 people will be saved/600 people will die}. Program A will cost \$100 million. Program B will cost \$100 million.</p> <p>{CBA Evaluation/Blank}</p> <p>{Explanation/Blank}</p> <p>Imagine you were in charge of making this decision. Which program would you select?</p>	<p>{Democrats/Republicans} on the House Transportation and Infrastructure Committee have proposed a bill for the construction of a new federal highway. According to the best projections, 10 billion trips will be taken on the highway during its lifetime. The highway will reduce travelling time for these trips by an average of 30 minutes. It is estimated to cost \$15 billion. Committee {Republicans/Democrats}, who were not consulted about the bill, have criticized the project for its lack of specificity.</p> <p>{CBA Evaluation/Blank}</p> <p>{Explanation/Blank}</p> <p>Imagine you had to vote on the highway bill. Would you support or oppose it?</p>	<p>The government is considering expanding the a of crop insurance into jurisdictions that, for historical reasons, were not previously eligible. The expansion would cover 20,000 acres of current farmland and would create approximately {500/2,500} jobs. Taxpayers are expected to incur \$30 million in direct costs.</p> <p>Given the amount of land affected, the expansion is not expected to effect global commodity prices one way or the other.</p> <p>{CBA Evaluation/Blank}</p> <p>{Explanation/Blank}</p> <p>Imagine you were in charge of making this decision. Do you think the government should expand the availability of crop insurance or not?</p>

Figure 1 Description of the Survey Prompts Used in Each Experiment.

Finally, for scope insensitivity, we initially explored the trade-off between forest preservation and employment in the logging industry (Milkman et al. 2012). In a pilot study run on Mechanical Turk, roughly 80% of respondents backed the creation of a natural reserve regardless of whether it saved 10,000 or 100,000 acres of forest. This lack of difference was expected and demonstrated scope insensitivity. The high baseline level of support raised concerns, however, about a ceiling effect. We therefore tested another scenario which presented a trade-off between the public cost of crop insurance and jobs. The proportion of respondents supporting the expansion of crop insurance did not vary much whether the program created 500, 2,500, or 5,000 jobs, ranging from 70% in the first condition, to 73% in the second, to 68% in the third. We ultimately chose to compare 500 against 2,500 jobs because of the visual contrast between the two numbers. At a glance, 500 might look like 5,000 but it cannot easily be mistaken for 2,500.

The bias-inducing manipulations generate the primitive effects for our experiments, but our ultimate interest is in whether these effects are diminished by procedural interventions. Figure 1 shows where the procedural treatments were administered in the survey instruments. Figure 2 provides an illustrative example of how the CBA and reason giving treatments were implemented in the highway bill scenario. Partisan-motivated reasoning was induced by randomly ascribing the highway bill to Republicans or Democrats. Subjects who were assigned to the control procedure (i.e. “no procedure”) expressed their preferences immediately after reading the scenario. All the other subjects had to satisfy one or more procedural requirements before expressing a policy preference. Subjects assigned to CBA had to complete a quantitative evaluation of the cost and benefits of the highway bill. They were informed at the outset of the exercise that the cost-benefit evaluation they were about to

perform was “nonbinding,” meaning that they could “choose to follow or not follow the results of the evaluation.” The evaluation commenced by eliciting subjects’ willingness to pay for one unit of the good to be produced by the policy at issue, here being the reduction of the traveling time of a trip by 30 min. Subjects were then instructed to compute the net benefit of the policy in monetary terms by multiplying their previous answer by the total number of units of good produced, that is, 10 billion, and then subtracting the cost of the project, that is, \$15 billion. Subjects were reminded at the end of the cost-benefit evaluation that they were not bound by its results. Subjects assigned to reason giving were invited to “[e]xplain in a few sentences” their attitudes toward the highway bill. No formal or substantive constraints were imposed on the explanations given by subjects.

For subjects assigned to CBA and reason-giving, cost-benefit evaluation preceded their written explanations. Although subjects not assigned to any procedure and subjects assigned to reason-giving alone did not have to perform the cost-benefit evaluation before deciding whether to support or oppose the highway bill, they were directed to do it after. This additional step equalized the burden across experimental conditions, thereby dampening the risk of differential attrition.⁸ Unless otherwise specified, only subjects who completed the entire instrument were included in our analysis. The CBA and reason-giving treatments were similarly implemented for the other two scenarios. Because gain-loss framing involves

⁸The gain-loss framing experiment was run twice. In the first iteration we did not ask subjects assigned to non-CBA conditions to do a cost-benefit evaluation after expressing a policy choice. In these results, substantially different rates of attrition were observed across procedural conditions. In the second iteration and other experiments, we did ask those not assigned CBA to nevertheless do one after making a policy selection and differential attrition was thereby improved. Our analysis is based on data collected from the second run of the gain-loss framing experiment. Data collected from the first run are presented in [Supplementary Appendix](#) and discussed below as well.

In your opinion, what dollar value should the government place on reducing the travelling time of a trip by 30 minutes? Please enter your answer in dollar amounts.

Based on your previous answer, what is the benefit of highway, measured in dollar terms? The net benefit of the highway is the dollar value you believe the government should place on reducing the travelling time of a trip by 30 minutes, multiplied by 10 billion, minus \$15 billion.

Imagine you had to vote on the highway bill. Explain in a few sentences whether you would support or oppose it. Remember the evaluation you completed is non-binding. That means you may choose to follow or not follow the results of this evaluation.

Figure 2 An Example of the Two Procedural Treatments as They Appear in the Partisan Reasoning Experiment.

uncertainty and a comparison between two, nonstatus quo alternatives, the cost-benefit evaluation in that experiment comprised of an extensive series of computations. That said, the arithmetic operations consisted only of multiplication and addition.

As a rule, testing for interaction effects requires much larger sample sizes than testing for main effects. Given concerns about statistical efficiency, we adopt a factorial design for all three experiments. There are three randomly varied factors in each experiment: a bias-inducing treatment, reason-giving, and CBA. Each factor can take on two levels. In [figure 3](#), the levels for each factor are represented by “+” and “-”. For the bias-inducing manipulation, the “+” level refers to the one expected to generate a higher level of approval (e.g. same-party endorsement in the partisan-motivated reasoning scenario). For the procedural treatments, that is, reason-giving and CBA, “+” means being subject to the procedure, and “-” means not being subject to the procedure. The factorial structure implies that each respondent has a random and equal chance of being assigned one out of the $2^3 = 8$ possible combinations of factors and levels. We conducted simulations based on the findings in the literature and the results of our pilot studies to determine the targeted sample size for each experiment. For the gain-loss framing and partisan-motivated reasoning scenarios, we picked sample sizes that would give us a greater than 75% chance of detecting a 40% reduction in the magnitude of the relevant cognitive bias. For the scope insensitivity experiment, we picked a sample size that would give us a greater than 80% chance of detecting a 15-point interaction between the bias-inducing and the procedural treatments. More details on the power analysis may be found in our preanalysis plan and [Supplementary Appendix](#).

We fielded the experiments on the Qualtrics platform to nationally representative samples acquired through Lucid. Respondents are recruited nonprobabilistically and are quota sampled to match the United States adult population for age, gender, ethnicity, and region.

Results

We analyze the experimental data to identify whether each procedure can mitigate various cognitive biases, not just in the absolute sense, but also on the margin and in concert. Each experiment presents subjects with a choice between two options. For purposes of analysis, we take the risky policy in the gain-loss framing scenario as the reference policy. Responses choosing the risky policy are coded as 1 whereas responses choosing the sure thing policy are coded as 0. For the other two scenarios, the proposed policy is taken as the reference policy. Support for the proposed policy is coded as 1 whereas opposition is coded as 0. [Figures 4](#) and [5](#) summarize the proportion of respondents choosing the risky option in the gain-loss framing scenario or supporting the proposed policy in the partisan-motivated reasoning and scope insensitivity scenarios. Differences along the y -axis are generated by the bias-inducing treatment. Differences along the x -axis are generated by the procedure that respondents were subject to before expressing their choice. The influence of procedure on choice can be observed by comparing differences along the y -axis across the x -axis. The bottom panel of both figures illustrates the difference-in-difference estimates of the impact of procedure on cognitive bias and the uncertainty around those estimates. Details of how these estimates were derived are given below. We generally report t -statistics rather than p




Figure 3 Design Representation for Three, 2^3 Factorial Experiments.

-values to allow readers to make their own judgments about statistical significance.⁹

To study the effect of CBA or reason-giving on cognitive bias, we estimate the following model:

$$Y_i = \beta_0 + \beta_1\tau_i + \beta_2\pi_i + \beta_3\tau_i\pi_i + \epsilon_i, \quad (1)$$

where τ_i is an indicator variable for the bias-inducing manipulation and π_i is an indicator variable for assignment to the procedure of interest, that is, CBA, reason-giving, or CBA and reason-giving (Dasgupta, Pillai, and Rubin 2015).¹⁰ This model is estimated by ordinary least squares using HC2 robust standard errors (Lu 2016). The regression analysis compares those assigned to a specified procedural intervention against those who were not. The main effect, captured by β_1 , is the magnitude of the cognitive bias in the absence of any procedure. The first-order interaction effect, captured by β_3 , measures the attenuation of cognitive bias by the procedure in question.¹¹ The estimates of β_3 for CBA and reason-giving are plotted in the bottom panel of figure 4.

A key observation from figure 4 is that, for all scenarios, at least one of the two procedures reduced cognitive bias. Take the partisan-motivated reasoning scenario. In the absence of any procedure, the party cue boosted support for the highway bill by 29.7% points. CBA reduced this gap to 18.8% points—a 10.9% point reduction ($t \approx -3.90$)—whereas reason-giving had almost no impact on subjects' susceptibility to the party cue ($t \approx -0.01$). The disparity between CBA and reason-giving is less striking in the other scenarios. Both CBA and reason giving seem to lessen the influence of gain–loss framing,

although the effect of CBA was stronger. In the absence of any procedure, framing the competing options in terms of losses rather than gains resulted in a 28.3% point jump in the proportion of subjects choosing the risky option. CBA pared this effect by 9.5% points ($t \approx -2.26$) and reason-giving by 4.4% points ($t \approx -1.05$). For scope insensitivity, both CBA and reason-giving appear to mitigate cognitive bias, although their relative strengths are now reversed. In the absence of any procedure, the crop insurance scheme enjoyed 0.2% point more support when it created 2,500 as compared to 500 jobs. Reason-giving increased this difference by 10.1% points ($t \approx 1.92$) and CBA by 5.3% points ($t \approx 1.00$). Although statistical significance obtains for only one of the two procedures against framing and scope insensitivity, the magnitudes of the estimates are not that far apart. Neither the framing nor the scope-insensitivity experiments yield conclusive statistical evidence for the superiority of one procedure over the other.

Our discussion of the experimental data has thus far focused on the first-order interaction between a bias-inducing manipulation and a given procedural treatment, pooled across factors. Figure 5 sharpens the analysis by looking at second-order interactions between a bias-inducing manipulation and the two procedural treatments. The panel for gain–loss framing reveals that the effect of CBA on gain–loss framing is driven by a powerful interaction between the two procedures. The other panels do not reveal any noteworthy second-order interaction effects. Even so, one can make several general observations about the absolute, as opposed to marginal, benefit of using both procedures. Figure 5 shows that the combination of both procedures consistently served to mitigate cognitive bias relative to no procedure. Together, CBA and reason-giving shrunk the partisan-motivated reasoning scenario effect from 30.4% points in the no procedure condition to 19.8% points with both procedures ($t \approx -2.67$)

⁹As a reviewer noted, there is a plausible argument that one-sided p -values are more appropriate than two-sided p -values since procedures are not expected to increase cognitive bias. (Some studies have shown, however, that introducing more evidence on an issue can sometimes result in greater polarization (Baekgaard et al. 2019; Kahan et al. 2017)). To avoid confusion, we always use two-sided p -values when we occasionally refer to these values but readers should consider this point when judging the strength of the evidence.

¹⁰The advisability of making covariate adjustments when conducting regression analysis of experimental data is disputed. (Freedman 2008; Lin 2013). Regardless, Supplementary Appendix presents estimates from models that adjust for pretreatment covariates such as education, ethnicity, household income, and party affiliation. Inclusion of these covariates does not qualitatively alter our findings. The most notable change is in the scope experiment, where adding controls increases the significance of the reason-giving estimate from significant at the 90% confidence level to significant at the 95% confidence level.

¹¹In our experiment there are four possible combination of procedures which we can denote $\pi_{0,0}$, $\pi_{1,0}$, $\pi_{0,1}$, and $\pi_{1,1}$. The first index indicates assignment to the first procedure and the second index indicates assignment to the second procedure. To estimate the effect of the first procedure, the model groups $\pi_{0,0}$ and $\pi_{0,1}$, where the first procedure was not assigned, and groups $\pi_{1,0}$ and $\pi_{1,1}$, where the first procedure was assigned. The estimate of the effect of the first procedure, β_3 , is therefore a weighted average of two estimates of the sample average treatment effect (SATE). The first estimate can be obtained by fitting the same model on $\pi_{0,0}$ and $\pi_{1,0}$ and is the SATE *without* the second procedure. The second estimate can be obtained by fitting the same model on $\pi_{0,1}$ and $\pi_{1,1}$ and is the SATE *with* the second procedure. β_3 is the weighted average of these two SATEs.


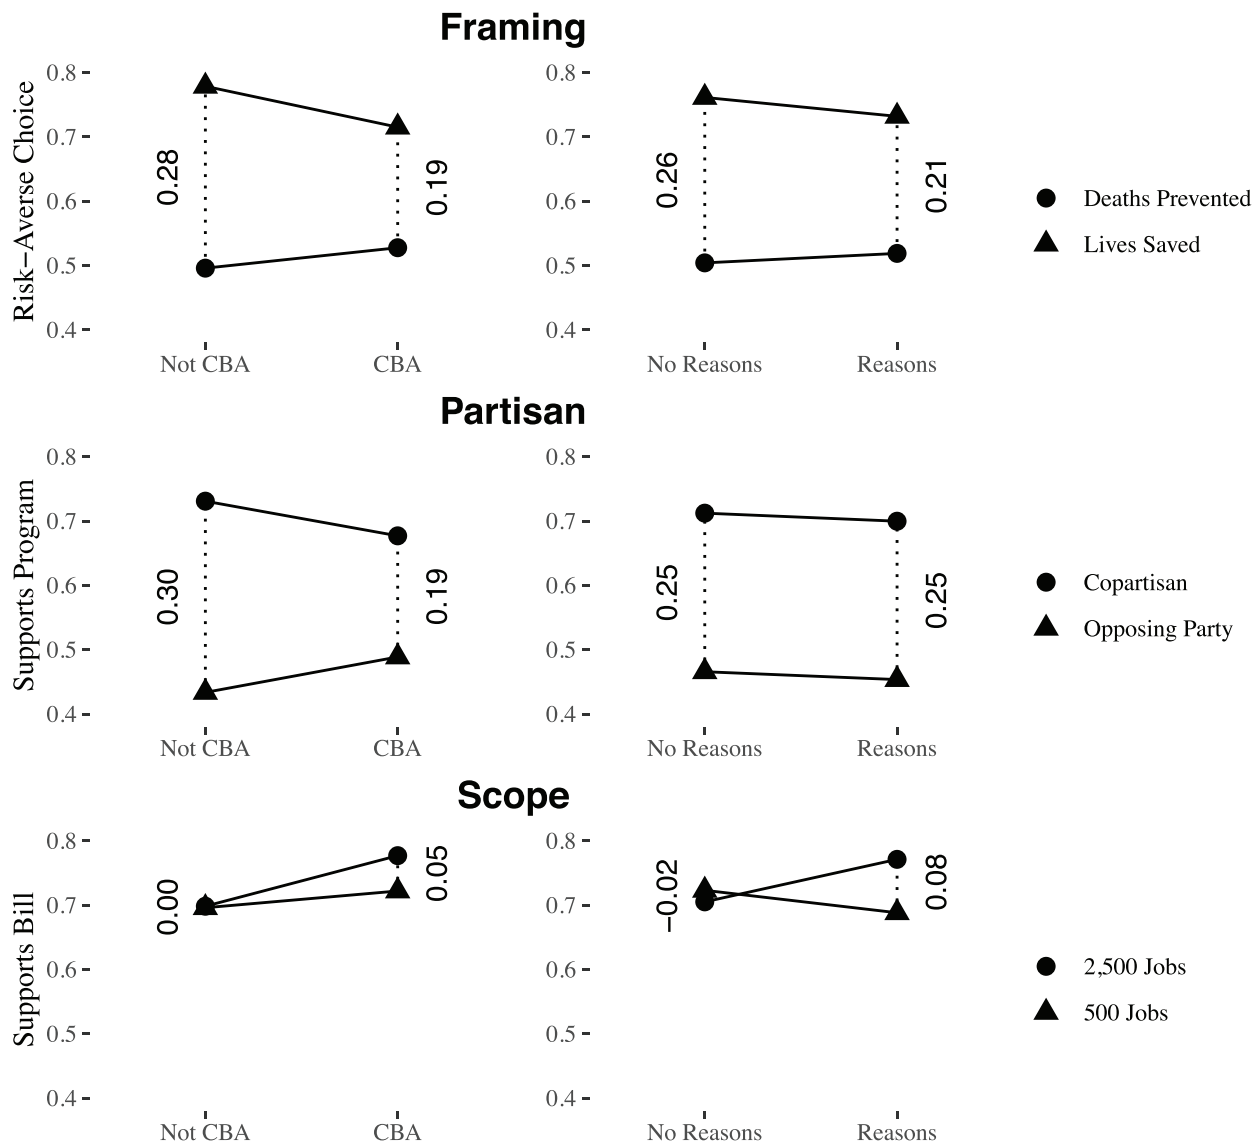


Figure 4 First-Order Interaction of Bias-Inducing Manipulation and Procedural Treatments. Estimates are Measured on a Binary Scale of 0 to 1, Where 1 Corresponds to Choosing the Proposed Policy and 0 to Choosing the Opposition. In the Gain-Loss Framing Scenario, 1 Corresponds to Choosing the Riskier Policy and 0 to Choosing the Alternative.

and the gain-loss framing effect from 26.8% points in the no procedure condition to 12.9% points with both procedures ($t \approx -2.67$). It also heightened sensitivity to scope, widening the gap in policy support by 15.8% points, from -5.0% points to 10.8% points ($t \approx -2.67$).

Figure 5 also shows that while the combination of both procedures brought consistent benefit relative to the no procedure condition, it yielded less reliable marginal benefit over the “best” single procedure. Scope sensitivity was the only bias for which reason-giving was the best single procedure.

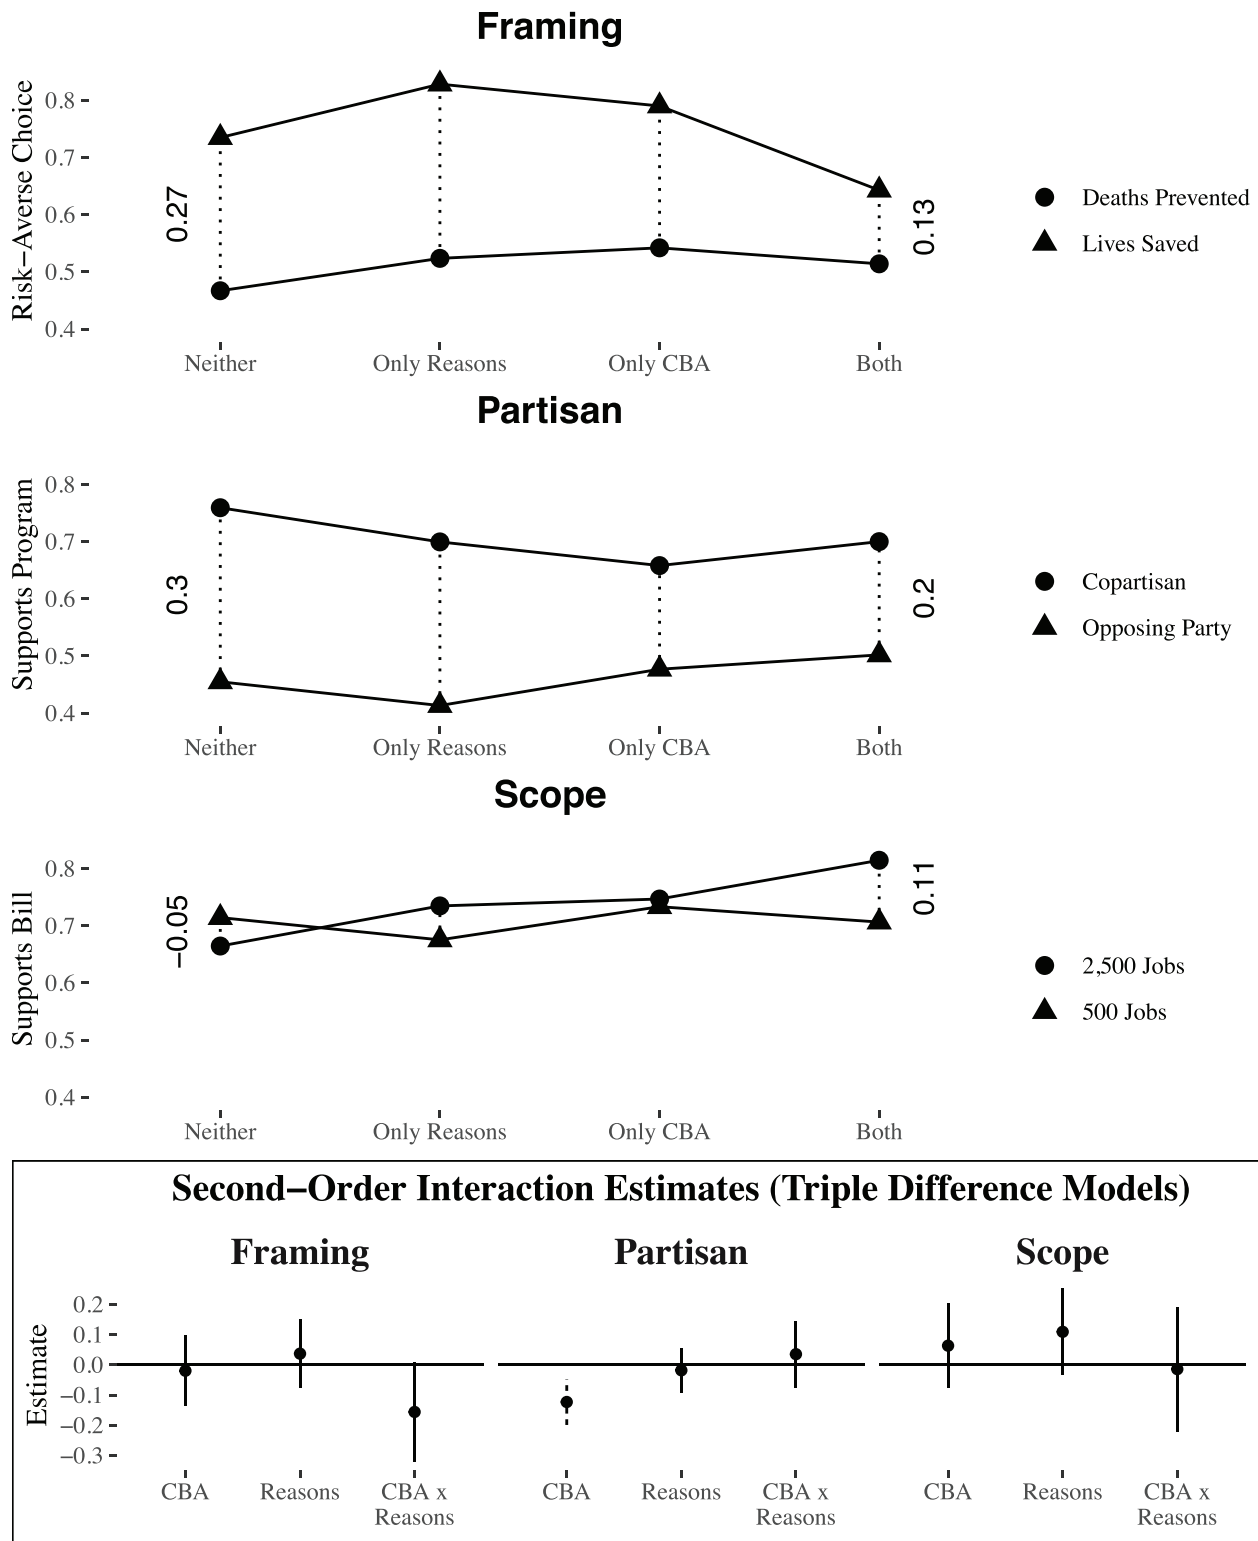


Figure 5 Second-Order Interaction of Bias-Inducing Manipulations and Procedural Treatments. Estimates are Measured on a Binary Scale of 0 to 1, Where 1 Corresponds to Choosing the Proposed Policy and 0 to Choosing the Opposition. In the Gain-Loss Framing Scenario, 1 Corresponds to Choosing the Riskier Policy and 0 to Choosing the Alternative.

Individuals assigned to reason-giving grew 5.9% points more supportive of the policy when its benefits were multiplied five-fold. Adding CBA increased the level of support by an additional 4.9% points. That increase, however, might be due to chance ($t \approx 0.63$). For the other two biases, CBA was the best

single procedure. Telling respondents assigned to CBA that the policy came from their side of the political aisle generated 18.1% points more support. When assigned both procedures, the bias actually increased to 19.8% points ($t \approx 0.40$). Gain-loss framing was the only bias where the cocktail was clearly

more potent than its strongest ingredient, reducing the bias by a further 11.9% points ($t \approx -1.97$).

In addition to the marginal benefits of both procedures over no procedure and both procedures over the best single procedure, there is the closely related question of complementarity between procedures. Does the total add up to more or less than its parts? To study this, we estimate the following saturated model

$$Y_i\beta_0 + \beta_1\tau_i + \beta_2\pi_{CBA,i} + \beta_3\tau_i\pi_{CBA,i} + \beta_4\pi_{RG,i} + \beta_5\tau_i\pi_{RG,i} + \beta_6\pi_{CBA,i}\pi_{RG,i} + \beta_7\tau_i\pi_{CBA,i}\pi_{RG,i} + \epsilon_i,$$

where τ_i is, as before, an indicator variable for the bias-inducing treatment and $\pi_{CBA,i}$ is an indicator variable for CBA and $\pi_{RG,i}$ is an indicator variable for reason-giving. This model is also estimated by ordinary least squares using HC2 robust standard errors. The main effect, captured by β_1 , is once again the magnitude of the cognitive bias in the absence of any procedure. The first-order interaction effects, captured by β_3 and β_5 , measure the attenuation of cognitive bias by CBA and reason-giving, respectively. $\beta_3 + \beta_5 + \beta_7$ is the effect of both procedures on cognitive bias. The second-order interaction effect, captured by β_7 , indicates the marginal contribution of carrying out both procedures. If β_7 has the same sign as β_3 and β_5 , it indicates that the treatments are complements. If the sign is opposite, it implies that the treatments are substitutes.

The estimates of β_3 , β_5 , and β_7 are plotted in the bottom panel of [figure 5](#). CBA and reason-giving only have strong synergies in the gain–loss framing scenario. Relative to people assigned to no procedure, those assigned to only CBA exhibit 2.0% points less sensitivity to framing while those assigned reason-giving proved 3.7% points more sensitive. Yet the combination of both procedures reduced sensitivity to framing by a sizeable 15.6% points ($t \approx -1.86$). For the other two biases, we have some, albeit weak, evidence of substitutability between the procedures. For scope sensitivity, the combination of the two procedures performed very slightly worse than the sum of its parts. For partisan-motivated reasoning, the combination of both procedures did not even outperform CBA alone. In sum, then, there is evidence that reason-giving enhanced CBA's effectiveness in taming gain–loss framing. There is no evidence of any synergies between reason-giving and CBA when it comes to countering partisan-motivated reasoning or scope insensitivity.

Differential Attrition

We now address the issue of differential attrition. Procedures are burdensome. A respondent who found the procedural intervention too taxing or time-consuming might abandon the study before making a policy choice. Such attrition threatens the internal validity of our findings. Suppose, for instance, that only respondents who have above-average numeracy complete the instrument if they are asked to do a CBA before making a policy choice. Suppose further that those not asked to complete a procedure prior to making a policy choice always finish the survey instrument. Then, any difference between the choices made by those assigned to no procedure and those assigned to CBA could be due to differences between more and less numerate respondents rather than the absence or presence of procedures. Any variance in outcomes could therefore be attributable not only

to random assignment of treatments but also selection on numeracy or other unobservable factors.

Recall that, to minimize differential attrition, we equalized the burden across experimental conditions by requiring respondents assigned to no procedure or to reason-giving alone to also perform a cost-benefit evaluation after they had expressed their policy choices. Since all respondents do a CBA, but some do it after making a policy selection, we can use data on survey completion to probe the issue of differential attrition. Indeed, it appears that the burden equalizing design choice functioned as intended. [Figure 6](#) shows the number of individuals who (a) were assigned to each procedural treatment condition, (b) proceeded to make a policy choice, and (c) completed the experiments. There was almost no differential attrition for the gain–loss framing scenario.¹² The number of respondents who completed the experiment was evenly distributed across all four possible procedural treatment conditions. Burden equalization also muted differential attrition in the partisan-motivated reasoning and scope insensitivity scenarios to some degree. Certainly, the differential attrition would have been much worse if we had ended the experiments as soon as respondents expressed a preference. Nevertheless, respondents in both experiments were still more likely to give up if they were assigned to both CBA and reason-giving as compared with any other procedural treatment condition. The threat posed by differential attrition must therefore be taken seriously.

We begin by looking at the demographics of respondents across procedural treatment conditions for all three scenarios. [Figure 7](#) presents a balance plot comparing demographic covariates across conditions. These covariates include age, education, gender, household income, and race and ethnicity. Each dot represents the standardized difference in covariate means between the no procedure condition—which suffered very little attrition—and one of the other conditions. The largest observed standardized difference in covariate means is about one quarter of a standard deviation. This imbalance occurred for the scope insensitivity experiment, where 54% of the respondents assigned to reason-giving and CBA and who completed the experiment were male. The same figure stands at 41% for respondents who were assigned to no procedure and who also completed the experiment. First, a 13% point difference may look substantial but it is not an unusual observation given that there are 315 standardized differences visualized in the plot. Second, and more critically, there is no evidence of gender moderating either scope insensitivity or the effect of procedure on scope insensitivity (see [Supplementary Tables A6, A10, and A11](#)).

Covariate balance cannot rule out selection on unobservables. There could still be respondent characteristics, unobserved by us and not captured by the demographic covariates considered above, for which the procedures select. One might posit that these characteristics, and not the procedural interventions, are the explanation for any findings. To explore this possibility, we rerun our preceding analysis on all respondents who expressed a policy choice, including those who abandoned the experiment. Inclusion of these respondents worsens differential attrition across procedural treatment conditions. It does not, however,

¹²In the initial run, we did not equalize burdens, and there was differential attrition. An express warning at the start of the experiment that some arithmetic would be required and that calculators were allowed may explain the absence of differential attrition relative to the other two experimental scenarios.




Figure 6 Extent of Attrition Across Experimental Arms. *y*-Axis Reflects Number of Respondents in Each Cell. *x*-Axis Reflects Assignment. Shading Contrasts Counts as Assigned, Counts Based on Expressing a Preference, and Instrument Completion.

change the regression estimates by very much.¹³ It is therefore unlikely that differential attrition accounts for our results.

Discussion

Conceptually, we describe administrative procedures as multifaceted, factorial interventions that, among other purposes, strive to improve the rationality of public decision making. We identify CBA and reason-giving as two of the most important procedural interventions, especially in the US federal regulatory landscape. While some have proposed quantitative regulatory impact analysis as a general remedy for rationality failures in policymaking, our perspective is more nuanced. To a degree, procedures can fix cognitive biases. But, if various procedures are imperfect substitutes with declining marginal benefits, then we should ask when a given procedure provides a better fix than other alternatives. We cannot presume that the benefit of additional procedures always justifies its costs and should consider tailoring procedures to the problems they might solve.

The experiments described earlier illustrate the principles of imperfect substitutability and declining marginal benefit. CBA substantially diminished partisan-motivated reasoning, whereas reason-giving did almost nothing. There is strong evidence

that reason giving diminished scope insensitivity, while CBA appeared to do less. For gain–loss framing, the combination of both interventions was needed for CBA to make a meaningful difference. It was the only bias for which strong complementarity between the procedures were found. For the other two biases, the combination of both procedures was, if anything, mildly antagonistic; there was no evidence of a marginal benefit to doing both procedures over the best single procedure.

While we have designed these experiments to illustrate broader conceptual principles about procedural interventions, it is tempting to think about how these experiments might translate into the practice of public administration and public policy. And indeed, such extrapolation is sometimes necessary and desirable, since it is virtually impossible to discern cognitive mistakes through purely observational studies (Rachlinski and Farina 2002; Dudley and Xie 2022).¹⁴ But even so, caution must be exercised because the experimental setting is quite far removed from actual policymaking. Real-world policymaking involves subject matter experts who might deliberate in groups and who, in all likelihood, have stronger intrinsic–and extrinsic–motivation than the

¹³It may also be worthwhile to consider the differences in estimated treatment effects between the two runs of the gain-loss framing experiment. Figure A1 shows that the confidence intervals around estimates from each run substantially overlap, and the treatment effect estimates from individual runs are within the relatively tighter confidence intervals of the pooled estimates. Table A12 confirms that the interaction of run and estimated treatment effect is never significant.

¹⁴By approximating real world situations, experimental studies such as Baekgaard et al. (2019); Bellé, Cantarelli, and Belardinelli (2018); Christiansen (2018); and Nielsen and Donald (2017) are able to convincingly diagnose pathologies in public decision making. But such fidelity is not always feasible due to resource or logistical constraints. Our experimental design, for example, requires the estimation of first and second-order interactions effects. The sample size required to detect moderate interaction effects is in the thousands. Successfully recruiting thousands of policymakers and having them prepare sophisticated CBAs is not a realistic plan.




Figure 7 Balance Plot to Detect Observable Differences Between Treatment And Control Conditional On Instrument Completion. Estimates are Measured in Standard Deviations.

typical survey respondent. These features of the real-world policymaking might call into question any conclusion drawn here. Still, it should be noted that insofar as there has been experimental research on the impact of expertise, incentive, or group processes on cognitive biases, these biases have proven to be quite prevalent and persistent (English and Mussweiler 2001; Enke et al. 2021; Holmgren et al. 2018; Kertzer et al. 2022; Tversky and Kahneman 1983; Guthrie, Rachlinski,

and Andrew 2000).¹⁵ Given the robustness of the underlying biases to the kinds of differences that could matter, we reasonably expect the mitigating effects of procedure to operate

¹⁵As Kertzer (2020) persuasively argues through a meta-analysis of hundreds of paired experiments, critics sometimes overstate the distinction between elite and lay decision making. Only 12% of the paired studies looked at found a difference in magnitude between elite and ordinary samples, and only 2% exhibited a difference in sign.

quite similarly across individuals and environments, although more research must be done to confirm those expectations.

Experimental methods can be employed to develop and test theories that apply and extend more broadly than the specific context in which they can be evaluated. Here, we have articulated rough-and-ready general principles for thinking about procedural interventions, but going further we can use similar experimental methods to test ideas about when procedures will succeed and when they will fail. Indeed, although the number of biases explored through our experiments is limited, we can already formulate some tentative hypotheses about mechanisms. To the extent that procedures fix cognitive biases, they do so by altering the cognitive processes that give rise to biased decisions. The literature commonly explains cognitive bias by reference to a dual process theory of reasoning. The two processes are conventionally termed System 1 and System 2 thinking (Stanovich 1999). On standard accounts, System 1 thinking is fast, intuitive, and automatic, whereas System 2 thinking is slow, deliberate, and controlled (Evans and Stanovich 2013). System 1 thinking manifests in “[e]xperience based decisionmaking”, and System 2 thinking, in “consequential decisionmaking” (Evans and Stanovich 2013). Gain–loss framing has been attributed to an affective heuristic typical of System 1 thinking (Cassotti et al. 2012; Gosling and Sylvain 2019). People are emotionally attracted to certain gains and emotionally averse to certain losses. We therefore select the gamble over the sure-thing when the options are framed as losses, and we choose the sure-thing over the gamble when the same alternatives are framed as gains. Scope insensitivity has also been characterized as the product of an affective heuristic (Hsee and Rottenstreich 2004; Kahneman et al. 1999). People arrive at valuations through feeling or by calculation. When we rely on feelings, we evoke mental representations of the prototypical example of the thing to be valued. Our valuations then become insensitive to scope because they are determined by how we affectively value the prototype. Partisan-motivated reasoning too has been traced to affective processes (Bolsen, Druckman, and Cook 2014; Lodge and Taber 2000; Redlawsk 2002). We may have accuracy or directional goals when forming opinions. That is, we may be motivated to reach the true or correct conclusion or to confirm a specific, preconceived position. Partisanship can activate directional goals that then color how we evaluate new evidence or information.

Because the biases studied here are said to originate in System 1 thinking, one conjecture is that insofar as procedure reduced cognitive bias, it did so by compelling System 2 thinking. But unlike gain–loss framing and scope insensitivity, partisan-motivated reasoning also contaminates System 2 thinking. When intuitive reasoning leads to the desired conclusion, the directional goal is achieved and the decision maker stops. But when it does not, the decision-maker transitions into a more deliberate mode of reasoning, seeking out favorable information and discounting contrary evidence (Kahan et al. 2017; Redlawsk 2002). Hence, it is not enough, in the case of partisan-motivated reasoning, for procedure to evoke more conscious and reflective thought. It also must neutralize the decision maker’s directional goals. This aspect of partisan-motivated reasoning could explain the contrasting impact of CBA and reason-giving. To all appearances, the requirement to give reasons did not make respondents less directionally oriented. This is perhaps because respondents had a vague but plausible reason to fall back on when justifying

their partisan-motivated choice.¹⁶ The cost-benefit evaluation, on the contrary, demanded attention to quantitative information about outcomes and seems to have thereby oriented respondents towards accuracy. If this interpretation is correct, then CBA not only facilitates a more structured assessment of policy, it also transforms how people process the available evidence or information. In so doing, it alleviates partisan-motivated bias in a way that reason-giving cannot.

To further explore how CBA alters cognitive processes, it is also worth asking whether the debiasing property of CBA depends on it being done properly. Recall that everyone in these experiments did a cost-benefit evaluation of some kind. The only question is whether they performed the analysis before or after making a policy choice. We code these analyses as either plausible or implausible depending on whether the calculations are internally consistent.¹⁷ Respondents are then divided into subgroups based on the plausibility of their completed evaluations. The proportion of respondents returning a plausible evaluation fluctuated across experiments and was never high (see table A9). Nonetheless, there were at least several hundred plausible evaluations in each experiment, and they were distributed quite evenly across experimental conditions. The results by subgroup are plotted in figure 8.

Visually, it seems that any effect that CBA has on cognitive bias does not always depend on the plausibility of the analysis itself. Implausible evaluations reduced gain–loss framing and partisan-motivated reasoning to the same, if not greater, degree than did plausible evaluations. The exception is scope insensitivity, where plausible evaluations generated more support for the policy when it created five times the number of jobs and implausible evaluations did not. Given limitations of statistical power, these observations are suggestive rather than conclusive. But they imply that in the case of partisan-motivated reasoning, it is perhaps the schema or mindset invoked by CBA that helps attenuate cognitive bias. On the contrary, the completion of a plausible evaluation appears to be crucial for CBA to mitigate scope insensitivity. It remains unsettled whether it is the amount of mental exertion or getting the arithmetic right that ultimately matters. The two are correlated and our experiments were not designed to tease them apart. But it seems that merely thinking in terms of costs and benefits is not enough to counteract scope insensitivity.

Conclusion

Administrative procedures can be characterized as costly, multifactorial interventions that promise to improve rationality in the public sphere. In this article, we have examined qualitative reason-giving and quantitative regulatory impact analyses as two of the most important procedural requirements in the modern administrative state. While some reformers have extolled the benefits of CBA in disciplining

¹⁶Christensen and Moynihan (2020) finds politicians to be more resistant to debiasing interventions than ordinary people. In particular, asking politicians for justifications worsened motivated reasoning, a phenomenon the researchers suggest is facilitated by the ability of experienced politicians to conjure up alternative considerations that marginalize the relevance of inconvenient data.

¹⁷By internally consistent, we mean that the benefit of the policy was correctly monetized given subjects’ own valuations of one unit of policy good, for example, jobs, lives, or commuting time. For example, if a respondent valued the creation of one job at \$10,000 and the benefit of a program creating 500 jobs at \$5 million, their analysis is coded as plausible, while any other valuation such as \$0.5 billion, or “a lot” is not.





Figure 8 CBAs Impact on Cognitive Bias Conditional on Completing Internally Consistent (i.e. Plausible) Analysis.

bureaucratic decision making, the principles of imperfect substitutability and diminishing marginal contribution suggest that procedural innovations have to be evaluated based on the threats they are supposed to neutralize and the range of existing alternatives. We have shown experimentally that administrative procedures can mitigate some of the best-known cognitive biases discussed in the literature. But the efficacy of any one procedure depends on the bias that one has in mind.

One tentative hypothesis, grounded in the literature on cognitive psychology, is that qualitative reason-giving can engender greater deliberation but may be of scant use if the bias at issue also infects the way one assesses new information or evidence. Quantitative CBA, on the contrary, might be more powerful in disrupting latent prejudices and orientations.

The practical implication of imperfect substitutability and diminishing marginal contribution is that a one-size-fits-all

solution to failures of rationality is not only too blunt but may also turn out to be unnecessarily burdensome. Admittedly, the findings reported here do lend support to the notion that more procedure produces better decisions. But if the experiments are anything to go by, better is not always that much better. Moreover, if one were to engage in a CBA of CBA, our discussion thus far has been mostly limited to the benefit side of the ledger. We have not paid much attention to the other side of the ledger, but the costs of CBA are hardly trivial. Indeed, the direct cost of producing such analyses easily runs into the millions of dollars (Boardman et al. 2018), and perhaps much higher if one takes into account indirect costs such as welfare losses from delayed regulation. The costs of CBA are therefore quite likely to exceed those of reason-giving, which, according to these experiments, may also do an impressive amount to reduce cognitive biases.

Future work can build on our experiments in several discrete directions. A natural question is how enhancing the verisimilitude of the experiments would change our findings. Possible dimensions for increased realism include elite samples of real regulators and policymakers,¹⁸ group discussion and deliberation, and some form of penalty for making biased decisions. While we do not expect the results to flip, we believe that further research in this direction is certainly desirable and worthwhile. There are also many other cognitive biases we have not examined, some rooted in phenomena besides affective or directional reasoning. Expanding the research horizon to include them could shed greater insight into the mechanisms underlying when and why procedures mitigate cognitive biases.

Finally, the experimental methods we propose here are useful for evaluation in a broad and conceptual sense but could also be used to think more creatively about the design of procedures. Full-fledged CBA can be decomposed into a number of steps, including identification of benefits and costs, quantification of the number of units created, monetization of the value of these units, aggregation of the benefits and costs into total costs and total benefits, and comparison of benefits and costs. Depending on the national or policy context, some or all of these steps may be omitted. It is possible that most of the decisional benefits come from the first two steps—identification and quantification—whereas much of the economic costs and political controversy surrounding CBA stem from the latter three. Adopting an experimental approach to the study of administrative procedures can help us design cheaper, more robust, and less onerous procedures to make policymaking smarter.

Supplementary Material

Supplementary data is available at the *Journal of Public Administration Research and Theory* online.

Funding

University of Hong Kong Seed Fund for Basic Research for New Staff.

Conflict of Interest

No known conflict of interests.

¹⁸As noted in an earlier footnote, this may not be easy to accomplish given the sample sizes required for the experiments to be informative and meaningful.

References

- Adler, M., and E. Posner. 1999. Rethinking cost-benefit analysis. *Yale Law Journal* 109 (2): 165–246.
- Adler, Matthew, and Eric Posner. 2006. *New foundations of cost-benefit analysis*. Cambridge, MA: Harvard University Press.
- Baekgaard, M., J. Christensen, Casper Mondrup Dahlmann, Asbjørn Mathiasen, and Niels Bjørn Grund Petersen. 2019. The role of evidence in politics. *British Journal of Political Science* 49 (3): 1117–40.
- Baekgaard, M., and S. Serritzlew. 2016. Interpreting performance information: motivated reasoning or unbiased comprehension. *Public Administration Review* 76 (1): 73–82.
- Bagley, N. 2019. “The procedure fetish. *Michigan Law Review* 118 (3): 345–401.
- Battaglio, R. Paul, P. Belardinelli, N. Bell’e, and P. Cantarelli. 2019. Behavioral public administration Ad Fontes. *Public Administration Review* 79 (3): 304–20.
- Baum, J. R., C. Jensen, and R. McGrath. 2016. Constraining a shadowy future: Enacting APAs in parliamentary systems. *Legislative Studies Quarterly* 41 (2): 471–99.
- Bellé, Nicola, P. Cantarelli, P. Belardinelli. 2018. Prospect theory goes public: Experimental evidence on cognitive biases in public policy and management decisions. *Public Administration Review* 78 (6): 828–40.
- Boardman, Anthony, David Greenberg, Aidan Vining, and David Weimer. 2018. *Cost-benefit analysis: Concepts and practice*. New York, NY: Cambridge University Press.
- Bolsen, Toby, J. Druckman, F. L. Cook. 2014. The influence of partisan motivated reasoning on public opinion. *Political Behavior* 36 (2): 235–62.
- Brader, Ted, and George E. Marcus. 2013. Emotion and political psychology. In *The Oxford handbook of political psychology*, ed. Leonie Huddy, David O. Sears, and Jack S. Levy, 165–204. New York, NY: Oxford University Press.
- Carpenter, Daniel. 2001. *The forging of bureaucratic autonomy*. Princeton, New Jersey: Princeton University Press.
- Cassotti, M., M. Habib, N. Poirel, A. Aite, O. Houdé, and S. Moutier. 2012. Positive emotional context eliminates the framing effect in decision-making. *Emotion* 12 (5): 926–31.
- Chen, B. M. 2020. The expressiveness of regulatory trade-offs. *Georgia Law Review* 55: 1029.
- Christensen, J., and D. Moynihan. 2020. Motivated reasoning and policy information: Politicians are more resistant to debiasing interventions than the general public. *Behavioural Public Policy* 1–22.
- Christiansen, A. 2018. On the cognitive argument for cost-benefit analysis. *Ethical Theory and Moral Practice* 21 (2): 217–30.
- Cole, D. 2012. Law, politics, and cost-benefit analysis. *Alabama Law Review* 64: 55–89.
- Dasgupta, T., N. Pillai, and D. Rubin. 2015. Causal inference from 2^k factorial designs by using potential outcomes. *Journal of the Royal Statistical Society: Series B* 77 (4): 727–753.
- De Francesco, F., C. M. Radaelli, and V. E. Troeger. 2012. Implementing regulatory innovations in Europe: The case of impact assessment. *Journal of European Public Policy* 19 (4): 491–511.
- De Francesco, Fabrizio. 2021. The diffusion of regulatory governance innovations: A research synthesis. In *Handbook of policy transfer, diffusion and circulation*, ed. O. Porto de Oliveira, 443–62. Northampton, MA: Edward Elgar Publishing.
- DeMuth, C. 2016. Can the administrative state be tamed?. *Journal of Legal Analysis* 8 (1): 121–90.
- Druckman, J. 2001. The implications of framing effects for citizen competence. *Political Behavior* 23 (3): 225–56.
- Druckman, James, Thomas Leeper, and Rune Slothuus. 2018. Motivated responses to political communications. In *The feeling, thinking citizen*, ed. Howard Lavine and Charles Taber, 125–50. New York, NY: Routledge.
- Dudley, S., and Z. Xie. 2022. Nudging the nudger: Toward a choice architecture for regulators. *Regulation & Governance* 16 (1): 261–73.

- Englich, B., and T. Mussweiler. 2001. Sentencing under uncertainty: Anchoring effects in the courtroom. *Journal of Applied Social Psychology* 31 (7): 1535–51.
- Enke, B., U. Gneezy, B. Hall, D. Martin, V. Nelidov, T. Offerman, and J. van de Ven. 2021. Cognitive biases: Mistakes or missing stakes?. *The Review of Economics and Statistics* 1–45. doi: 10.1162/rest_a_01093.
- Ernst, Daniel. 2014. *Tocqueville's nightmare*. New York, NY: Oxford University Press.
- Evans, J., and K. E. Stanovich. 2013. Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science* 8 (3): 223–41.
- Freedman, D. A. 2008. On regression adjustments to experimental data. *Advances in Applied Mathematics* 40 (2): 180–93.
- Gellhor, W. 1986. The administrative procedure act: The beginnings. *Virginia Law Review* 72 (2): 219–33.
- Gosling, C., and M. Sylvain. 2019. Is the framing effect a framing affect?. *Quarterly Journal of Experimental Psychology* 72 (6): 1412–21.
- Guthrie, C., J. J. Rachlinski, and J. W. Andrew. 2000. Inside the judicial mind. *Cornell Law Review* 86: 777.
- Holmgren, M., A. Kabanshi, J. E. Marsh, and P. Sörqvist. 2018. When a + b < a: Cognitive bias in experts' judgment of environmental impact. *Frontiers in Psychology* 9: 823.
- Hsee, C. K., and Y. Rottenstreich. 2004. Music, pandas, and muggers: On the affective psychology of value. *Journal of Experimental Psychology: General* 133 (1): 23–30.
- Jones, Bryan, Sean Theriault, and Michelle Whyman. 2019. Causes of the great broadening: The role of social movements. In *The Great Broadening*, 109–30. Chicago, IL: University of Chicago Press.
- Kahan, D. M., E. Peters, E. C. Dawson, and P. Slovic. 2017. Motivated numeracy and enlightened self-government. *Behavioural Public Policy* 1 (1): 54–86.
- Kahneman, Daniel, Ilana Ritov, David Schkade, Steven J. Sherman, and Hal R. Varian. 1999. Economic preferences or attitude expressions? An analysis of dollar responses to public issues. In *Elicitation of Preferences*, ed. Baruch Fischhoff and Charles F. Manski, 203–42. Dordrecht: Springer Netherlands.
- Kertzer, J., M. Holmes, B. LeVeck, and W. Carly. 2022. Hawkish biases and group decision making. *International Organization* 76 (3): 513–48.
- Kertzer, Joshua. 2020. Re-assessing elite-public gaps in political behavior. *American Journal of Political Science* 66: 539–53. doi: 10.1111/ajps.12583.
- Ladd, J. M., and G. S. Lenz. 2008. Reassessing the role of anxiety in vote choice. *Political Psychology* 29 (2): 275–96.
- Lerner, J. S., and P. E. Tetlock. 1999. Accounting for the effects of accountability. *Psychological Bulletin* 125 (2): 255–75.
- Lin, W. 2013. agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *The Annals of Applied Statistics* 7 (1): 295–318.
- Lodge, Milton, and Charles S Taber. 2000. Three steps toward a theory of motivated reasoning. In *Elements of reason: Cognition, choice, and the bounds of rationality*, ed. Arthur Lupia, Mathew D. McCubbins, and Samuel L. Popkin, 183–213. New York, NY: Cambridge University Press.
- Lodge, Milton, and Charles S Taber. 2013. *The rationalizing voter*. New York, NY: Cambridge University Press.
- Lu, J. 2016. On randomization-based and regression-based inferences for 2^k factorial designs. *Statistics & Probability Letters* 112: 72–8.
- Mackie, P., T. Worsley, and J. Eliasson. 2014. Transport appraisal revisited. *Research in Transportation Economics* 47: 3–18.
- Marcus, George E., W. Russell Neumann, and Michael MacKuen. 2000. *Affective intelligence and political judgment*. Chicago: University of Chicago Press.
- Mashaw, Jerry L. 2018. *Reasoned administration and democratic legitimacy*. New York, NY: Cambridge University Press.
- McGarity, T. O. 1992. Some thoughts on “deossifying” the rulemaking process. *Duke Law Journal* 41 (6): 1385–462.
- McNollgast. 1999. The political origins of the administrative procedure act. *Journal of Law, Economics, & Organization* 15 (1): 180–217.
- Milkman, K. L., M. C. Mazza, L. L. Shu, C. J. Tsay, and M. H. Bazerman. 2012. Policy bundling to overcome loss aversion: A method for improving legislative outcomes. *Organizational Behavior and Human Decision Processes* 117 (1): 158–67.
- Morrison, A. B. 1986. OMB interference with agency rulemaking: The wrong way to write a regulation. *Harvard Law Review* 99 (5): 1059–74.
- Moynihan, D. P. 2004. Protection versus flexibility: The Civil Service Reform Act, Competing Administrative Doctrines, and the Roots of Contemporary Public Management Debate. *Journal of Policy History* 16 (1): 1–33.
- Nielsen P. A., and Donald P. Moynihan. 2017. How do politicians attribute bureaucratic responsibility for performance? Negativity bias and interest group advocacy. *Journal of Public Administration Research and Theory* 27 (2): 269–83.
- Olsen, A. L. 2015. Citizen (Dis) satisfaction: An experimental equivalence framing study. *Public Administration Review* 75 (3): 469–78.
- Rachlinski, Jeffrey J., and Cynthia R. Farina. 2002. Cognitive psychology and optimal government design. *Cornell Law Review* 87: 549–615.
- Redlawsk, D. 2002. Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making. *The Journal of Politics* 64 (4): 1021–44.
- Revesz, Richard L., and Michael A. Livermore. 2008. *Retaking rationality: How cost-benefit analysis can better protect the environment and our health*. New York, NY: Oxford University Press.
- Shapiro, M. 1986. APA: Past, Present, Future. *Virginia Law Review* 72 (2): 447.
- Sheffer, L., and P. J. Loewen. 2019. Accountability, framing effects, and risk-seeking by elected representatives: An experimental study with American local politicians. *Political Research Quarterly* 72 (1): 49–62.
- Simon, Herbert A. 1976. From substantive to procedural rationality. In *25 years of economic theory*, ed. T. J. Kastelein, S. K. Kuipers, W. A. Nijenhuis, and G. R. Wagenaar, 65–86. Boston, MA: Springer US.
- Stanovich, Keith E. 1999. *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Psychology Press.
- Stephenson, M. C. 2006. A costly signaling theory of hard look judicial review. *Administrative Law Review* 58 (4): 753–813.
- Stiglitz, E. H. 2017. Cost-benefit analysis and public sector trust. *Supreme Court Economic Review* 24 (1): 169–96.
- Sunstein, C. 2000. Cognition and cost-benefit analysis. *The Journal of Legal Studies* 29 ((S2): 1059–103.
- Sunstein, Cass. 2002. *The cost-benefit state: The future of regulatory protection*. Chicago, IL: American Bar Association.
- Sunstein, Cass. 2018. *The cost-benefit revolution*. Cambridge, MA: MIT Press.
- Tversky, A., and D. Kahneman. 1973. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology* 5 (2): 207–32.
- Tversky, A., and D. Kahneman. 1981. The framing of decisions and the psychology of choice. *Science* 211 (4481): 453–8.
- Tversky, A., and D. Kahneman. 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* 90 (4): 293–315.
- Wiener, J. B. 2013. The diffusion of regulatory oversight. In *The Globalization of Cost-Benefit Analysis in Environmental Policy*, ed. Michael A. Livermore and Richard L. Revesz, 123–41. New York, NY: Oxford University Press.
- Yackee, J. W., and S. W. Yackee. 2010. Administrative procedures and bureaucratic performance: Is federal rule-making ‘Ossified’?. *Journal of Public Administration Research and Theory* 20 (2): 261–82.