Making Regulators Reasonable: Do Procedural Rationality Requirements Fix Cognitive Biases?

A Preprint

Benjamin M. Chen
Faculty of Law
University of Hong Kong
Hong Kong
benched@hku.hk

Brian Libgober
School of Global Policy and Strategy
University of California - San Diego
La Jolla, CA
blibgober@ucsd.edu

July 13, 2021

Abstract

This article uses factorial experiments to assess the comparative efficacy of quantitative cost-benefit analysis and qualitative reason-giving for reducing well-established cognitive biases. Qualitative reason-giving was made a universal procedural requirement for US federal regulations by the Administrative Procedures Act (APA) of 1946. Requiring regulators to explain their policies is not controversial. President Reagan added the requirement of cost-benefit analysis ("CBA") to most "major" rules in 1981. Since then, Presidents of both parties have expanded the role of impact assessments. Multilateral organizations such as the World Bank, OECD, and European Union also support this regulatory reform. Despite increasing adoption, critics argue that CBA burdens regulators and prevents them from defending hard-to-value public goods like the environment. Supporters describe it as a neutral, procedural tool that leads to a more rational policy. Some explicitly argue that CBA is the cure to failings in human reasoning identified by cognitive psychology. This claim, however, has received limited empirical scrutiny until now. We find that reason-giving produces significant reductions in framing effects and increased scope-sensitivity, while CBA significantly reduces partisan motivated-reasoning. The "cocktail" treatment of CBA with reason-giving consistently causing the most considerable improvement in cognitive bias, although these are usually not substantially different from reason-giving alone. Cumulatively, these findings suggest the rationality benefits of CBA are overstated, especially given their higher costs. By contrast, the rationality benefits of qualitative explanation are under-appreciated, although these too could be more or less costly. Our work points to the need for further empirical evaluations of procedural reforms to executive branch policymaking, especially those that appear "neutral" but substantially alter preferences.

Keywords cost-benefit analysis · procedural rationality · cognitive psychology · framing · partisan-reasoning · scope-sensitivity · administrative burden · factorial experiments · survey experiments

1 Introduction

Within his first month of office, President Reagan launched an ambitious program of regulatory reform. It included the most significant systemic overhaul to the regulation-making process since the Administrative Procedures Act of 1946 (Dudley 2021). Its central component was Executive Order 12,291. EO 12,291 gave the Office and Management and Budget a veto over regulations made by the Environmental Protection Agency, the Department of Education, and many other agencies.[1] Before this order, OMB had played an advisory

[1]Indeed, the only agencies excluded from the order were the so-called "independent agencies." Even so, these agencies were encouraged to comply with Reagan's order.

role in assessing the advisability of regulations. After, OMB analysts would evaluate regulatory impact assessments that agencies attached to their proposals. In particular, they would scrutinize a regulation's alleged costs and benefits and could send the rule back to the drawing-board if found wanting. And indeed, the Reagan administration rejected hundreds of rules through OMB review.

Conceptually, the Reaganite reform presented a significant change in how agencies established the advisability of their regulation. Under the APA, the open process of public deliberation ensured the quality of regulatory policy. The process would begin with an agency making a proposed version of its regulation available and a preface explaining the document's meaning and purpose. An interested public could then offer constructive feedback in the form of comment letters. The possibility of judicial review and fire-alarm oversight by Congress incentivized agencies to put effort into their justifications and attend to high-quality public feedback (Gailmard and Patty 2016). Overnight, Reagan changed what agencies needed to produce to get their regulation made and upheld. No longer would agencies primarily need to create a colorable written justification. Now, they would additionally require a quantitative analysis that could get past the bean counters at OMB.

Presidents like control. Often, they feel they have less of it than they should (Neustadt 1991). The need for control is one explanation why subsequent Presidents have issued executive orders expanding cost-benefit analysis, even relatively left-wing Presidents like Clinton or Obama. Yet, the explanation is only partial. Indeed, we do not think it is even the most important reason, given the proliferation of BCA in regulatory regimes outside the US federal government (STAROŇOVÁ 2010; Dunlop and Radaelli 2016). Since 2006, 35 OECD nations have incorporated some kind of economic impact analysis into their regulation-making processes (Francesco, Radaelli, and Troeger 2012). For the past several decades, international organizations like the OECD and the World Bank have heavily promoted the structuring of cost-benefit analysis into the rulemaking process ("Global Indicators of Regulatory Governance: Worldwide Practices of Regulatory Impact Assessments" 2018, 1). Another critical explanation for the rise of the cost-benefit state is ideological, particularly a line of thinking that traces its origin to the American Progressive movement and ultimately the Enlightenment (Sunstein 2018). Public administration is rational administration (Van Riper 1983). Making choices with a positive net benefit is, at least according to prevailing economic conceptions, the definition of rationality. Conversely, taking actions whose costs outweigh the benefits is ipso facto irrational. Surfacing which side of the ledger a policy choice seems like the only prudent thing to do.

Although benefit-cost analysis has broad and seemingly neutral appeal, scholars and policymakers must remain cautious. The process places burdens, both administrative and substantive, on regulators (Moynihan, Herd, and Harvey 2015). It confers strong structural powers on the President, at least in the influential US model. These are reasons enough to evaluate the claim that the procedure improves decision-making. Indeed, the diagnosis justifying this intervention has a surprisingly thin evidentiary record. The most substantial evidence comes from case studies, which are appropriate but also limited. Quantitative evidence of rationality problems, such as it exists, relies on extrapolation from meaningfully different decision-making contexts, often survey or lab studies in psychology. Given the difficulties of directly observing inputs and outputs to the policymaking process, especially in executive agencies, it is not surprising that empiricists have struggled to reach these questions directly. How does one observe that a governmental decision is irrational? Even so, supporters of CBA argue this limited evidence suffices to justify the costly intervention. Yet, one must bear in mind that this fix is layered atop an existing set of processes that aim to solve the same or similar problems. At the US federal level, the most significant pre-existing process is notice-and-comment. The loosely structured reason-giving requirement burdens regulators less and involves less centralized and politicized control than CBA. An extensive literature has argued that BCA might introduce systemic biases, in particular against hard-to-measure goods and in favor of easily identified costs, can ignore distributive consequences, or leads to missing easily anticipated or rare downstream effects (Shapiro 2011). As Libgober (2020) argues, the notice-and-comment process can at least theoretically give regulators substantial flexibility in remaining sensitive to distributive concerns, intense preferences, or other policy goods.

In this paper, we conduct an experimental evaluation of a mechanism assumed to underly and justify one of the most significant structural reforms to the governance of the US political economy, with broad applicability to many other national and sub-national contexts. In particular, we take literally the claims of influential scholars of institution design such as Sunstein (2000), who maintain that the primary justification for benefit-cost analysis is to reduce cognitive bias. We seek to test whether quantitative analysis of policy decision problems yields any improvement regarding cognitive bias compared with qualitative reason-giving. We do so using large-n factorial survey experiments on a nationally representative sample.

Examing three well-established cognitive biases (gain-loss framing, partisan reasoning, and scope insensitivity), we find that quantitative analysis only appears to have some efficacy concerning partisan motivated reasoning.

Perhaps surprisingly, qualitative justification (i.e., "giving reasons") does more to reduce cognitive biases associated with insensitivity to scope or, remarkably, risk-avoiding behavior in the face of gains and risk-seeking behavior in the face of potential losses. We do not find that quantitative and qualitative analyses typically have valuable synergies. If anything, we find the reverse. Asking people to do quantitative analysis and give reasons may debias them somewhat more than asking them to do one or the other task before making a policy choice. The additional debiasing is less than the sum of its parts. Digging deeper, we find further weaknesses in the quantitative treatment. The gains associated with assignment to quantitative analysis do not appear to depend on actually plausibly completing the task. Instead, the mechanism of efficacy likely relates to decreasing the perceived salience of partisan signals. Less burdensome alternatives might also achieve these same ends.

More broadly, we see this work as helping to unbundle a complex set of interventions and benefits with vast stakes but limited means of direct evaluation. In particular, rather than supposing that more thought and analysis produces improvements to reasoning, our contribution shifts the focus to specific procedures and specific systemic errors in reasoning. There remain many questions about how our findings might change if participants had incentives to avoid mistakes, were relatively more expert, or worked in groups instead of alone. These are all features of actual regulation-making that are hard (but not impossible) to simulate in the lab. We consider the evaluation of the role of these features a vital part of a continued, meaningful dialogue between political science, public administration, and law about the proper design of policymaking institutions in light of modern tools of quantitative social science.

2 Theoretical Background

As theoretical preliminaries, we briefly introduce procedural rationality and cognitive bias. In particular, we focus first on how historical reforms to regulatory policymaking processes have sought to improve bureaucratic rationality. Second, we discuss how the experimental identification of cognitive biases has influenced scholarly understandings of the policymaking process at executive agencies.

2.1 Making Regulators Reasonable: Procedural Tools

Most policymaking bodies have formal procedural rules that determine how that body comes to a collective decision. Scholars have offered numerous, varied, and non-exclusive interpretations of such procedures' purpose. Political theorists, for example, have often justified directly democratic processes as devices for building legitimacy and reciprocity behind collective decisions. Formal theorists have argued that procedures in legislatures allow stable and predictable equilibria (Shepsle 1979; Eguia and Shepsle 2015), perhaps at the cost of tilting policy selection in ways that structurally favor one group of elected officials over another (Weingast and Marshall 1988; Cox and McCubbins 1993). By contrast, scholars of administration have long emphasized the role of procedures in improving bureaucratic rationality (Wilson 1887; Weber 1946).

Indeed, the history of policy interventions in administration in the United States is also a history of attempts to make bureaucrats more rational (e.g., Van Riper 1983; Gormley 1989; Mashaw 2018). The Pendleton Act is conventionally understood as having had the purpose of improving the rationality in hiring, firing, and promotion decisions (Moynihan 2004). Although there is a rich literature in social science developing the political and economic logic for why civil service reforms happene when they did (Skowronek 1982; Johnson and Libecap 1994; Theriault 2003), with some even arguing that Progressives had less to do with the Pendleton Act than one might think (Postell 2017), there is little doubt that the increasingly popular Progressive influenced the subsequent expansion of the civil service to cover the majority of federal workers. Indeed, scholars have shown that executive agencies in the early parts of the 20th century gained increasing policymaking authority and power (Carpenter 2001; Ernst 2014). Yet, their decision-making processes remained deeply problematic in many respects. One of the most important of which was the limited analytical and evidentiary basis for their decisions (Gellhorn 1986). In the 1930s, as New Deal Democrats demanded new and bold experiments in regulating the economy, demands for further interventions designed to promote consistency and rationality in agency decisions began to reach a fever pitch (Shapiro 1986; Ernst 2014).

The Administrative Procedure Act of 1946 (APA) was the culmination of these debates. As with the Pendelton Act, there is a political science literature showing that the motivations for enacting this reform were more complicated and subtle than one might assume (McNollgast 1999). Again, however, improving bureaucratic rationality was a frequently articulated goal of reform. The APA creates several frameworks for receiving evidence, advice, and demands from stakeholders. These frameworks include informal or "notice and comment" rulemaking, formal rulemaking, and adjudication. An extensive literature in political science now explores

the functioning of these policymaking formats, especially the ubiquitous informal rulemaking (Yackee 2019; Potter 2019; Libgober 2020). Regardless of the framework used, one of the most critical aspects of the APA was the procedural requirement that when agencies make new rules, "the agency shall incorporate in the rules adopted a concise general statement of their basis and purpose" (5 USC 553). This requirement is much more remarkable than it sounds. Legislatures may include in their laws some explanation of their goals. However, they are not required to do so. When they do, their words generally fall far short of a reasoned justification for their policymaking choices. Judges have long given opinions intended to elaborate the basis for their decisions. But as Mashaw (2018) notes, the rationality demands imposed on agencies are *higher* than on judges, since "a lower court g[iving] the wrong reasons for a correct decision is not by itself a justification for reversal." The APA asked agencies to go beyond prior judicial and legislative standards of reasonability in their decisions when making regulations.

For several reasons, the immediate impact of the APA on regulatory behavior was more limited than one might assume In the 1940s and 1950s, agencies largely focused on making individual determinations rather than use of general rulemaking powers and Courts were not overly insistent about the justification requirement (Mashaw 2018). The rise of powerful social movements in the 1960s and 1970s gave rise to broad delegations of policymaking authority to agencies (Jones, Theriault, and Whyman 2019), and a correspondingly larger role for executive agencies as policymakers. Courts responded by putting teeth into the APA's written justification requirement (Stephenson 2006). Agencies have followed suit by increasing the length of the preamble that 5 USC 553 requires. There remain significant debates about just how much reason giving by agencies is actually desirable. Some argue that too much explanation can become a burden that slows down the regulation-making process (Elliott 1992; McGarity 1992). Still, the basic push of the APA toward insisting on reasons to take quasi-legislative action is largely uncontroversial.

Uncontroversial is not a description one could use to describe the rise of cost benefit analysis (often called regulatory impact assessments). Gormley (1989) and DeMuth (2016) both identify the rise of regulatory impact analysis with budgeting systems implemented in Robert McNamara's Defense Department that later were mandated for 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 were fairly flexible and did not require explicit quantification of every cost and benefit, nor was there an explicit veto power built into the process. Indeed, even the environmental impact statements required by NEPA had no direct effect on the ability of an agencty to make a regulation.

President Reagan's EO 12,291 represented a significant shift 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 cost-benefit analysis the agency provided unless that option was prohibited by law. It also empowered the Office of Management and Budget to indefinitely delay the publication of regulations that it did not find acceptable. Some regarded this development as a 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." At the same time, there were more than a few defenders of the approach, even on the left. In particular, DeMuth (2016) emphasizes a growing sentiment that government agencies were prone to capture and over-regulation, which ultimately hurt the broad social movements that by the 1980s had grown to define the Democratic party. DeMuth mentions as an example Alfred Kahn, chair of the Civil Aeronautics Board under President Carter, who pushed for the dismantling of airline regulations, supported by allies such as Ralph Nader and Ted Kennedy.

Despite the initial contorversy of Reagan's approach to CBA, no Democratic President since Reagan has eliminated it. In fact, each has expanded it in important way. What Clinton did. What Obama did. What Biden did. The project of left-leaning Presidents with respect to cost-benefit analysis has tended to be the encouragement of more synoptic and comprehensive reviews. Again, the normative appeal of rational public administration is a powerful driver and Progressive ideology crosses easily between the right and the left.

2.2 Cognitive Biaes

Although reform to American bureaucracy have often sought improvements to rationality, rationality is a contested concept susceptible to multiple meanings (Adcock 2001). Simon (1976), for example, distinguishes between substantive and procedural rationality. A choice is substantively rational if it fits the decision-maker's goals given their constraints. A choice is procedurally rational if it follows the appropriate series of steps. As has been argued, structural interventions in bureaucratic policymaking on the procedural axis are frequented intended to produce benefits on the substantive axis. Put differently, those who design adminsitrative

procedures do so in order to "stack the deck" in favor of policy choices that better fit their goals given their constraints (McCubbins, Noll, and Weingast 1987). Yet deviations from substantive rationality are themselves a complex topic. Within social science, the dominant conception of substantive rationality follows from the von Neumann and Morgenstern axioms, which permits mathematical analysis of decision problems using utility functions, expectation operators, and the tools of calculus. There are multiple and rich vocabularies for describing different ways in which human decision-making might differ from that assumed by classical decision-theory. There are vast literatures discussing the role of "affect" or "emotions" with respect to cognition (Redlawsk 2006; Brader and Marcus 2013; Lodge and Taber 2013; Marcus, Neumann, and MacKuen 2000; Ladd and Lenz 2008), studies about the breakdown of specific axioms of von Neumann Morgenstern expected utility theory such as transitivity[2] (Tversky 1969; Regenwetter, Dana, and Davis-Stober 2011) or the independence axiom[3] (Allais 1953; Machina 1987), and alternative conceptions such as bounded rationality (Simon 1997; Jones 1999).

Quantitative analysis and reason-giving are, hypothetically, tools that might help bring bureaucratic decision-making closer to that of an expected utility maximizer by reducing the role of affective reasoning, bounded reasoning, non-transitivity, consideration of irrelevant alternatives, and still other bugbears. While it is clear that these phenomena lead to *cognitive errors* as compared with the rational-theoretic benchmark, the typical impact of these errors is not obvious with respect to any kind of decision problem.[4] Perhaps for this reason, scholars interested in the design of administrative institutions have tended to focus on "systematic" and "predictable" deviations from rationality, as exemplified in Tversky and Kahneman's seminal article in *Science* (Tversky and Kahneman 1981). In that paper, Tversky and Kahneman introduced the concept of framing effects, the paradigmatic form of *cognitive bias*, by showing that whether a choice was phrased in terms of gains or in terms of losses had an outsized effect on the typical preferences of respondents. Moreover, these preferences were not sensitive to the subject matter: whether about Asian diseases or lotteries, respondents were more likely to take risks in the face of losses and more likely to avoid risks in the face of potential gains. Both before and especially since that article, social scientists have uncovered many and varied forms of cogntive bias. In many cases, the source of cognitive bias seems to be the use of cogntivie shortcuts ("hueristics") or prior mental maps ("schemas") to decrease the cognitive load associated with a more well-reasoned solution. Particularly important for the political science literature are biases associated with the concept of motivated reasoning, especially partisan motivated reasoning. Introduced by Lodge and Taber (2000), motivated reasoning occurs when individuals' desires about the solution to a problem influences their resolution of that problem. Political ideology or identity has been found to provide sufficient incentive for individuals to seek out information that agrees with prior beliefs, rate confirming information more highly, and resist disconfirming evidence (Druckman, Leeper, and Slothuus 2018).

Contemporary analysts have explicitly proposed that cost-benefit analysis is a plausible mechanism for reducing cognitive biases like those identified by Tverskey and Kahneman, as well as problems of motivated reasoning such as identified by Taber and Ldodge (Sunstein 2000, 2002; Rachlinski and Farina 2002). At times, these authors verge on the claim that CBA will reduce all forms of cognitive bias, however there are certain biases that in particular standout. Sunstein (2000) discusses six cognitive biases that he thinks CBA can help treat. Abstracting somewhat, these biases tend to involve the over-reliance on highly availabile information, an inappropriately large role for social cues and priors, and problems attending to the magnitude of benefits, costs, and probabilities, especially in an integrative way.

At the same time that some institutional scholars have expressed great enthusiam about the potential of procedural reforms to improve substantive rationality, the work of cognitive psychologists gives reason for doubt. Lerner and Tetlock (1999) describe conditions under which accountability mechanisms such as reason-giving or CBA can leave cognitive biases unchaged or even woerse. As they write, "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." It can worsen bias to the extent that one choice appears easier to justify. On the other hand, Bolsen, Druckman, and Cook (2014) find that encouraging individuals to evaluate policy "in an evenhanded way" and telling them they will be asked "to justify the reasons for your judgment" did substantially decrease the role of partisan cues in influencing attitudes toward the 2007 Energy Act. At the same time, the literature on accountability mechanisms surely features several studies that look at the role of reason-giving, there are very few studies that consider the

[2]The transitivity axiom is that if X is preferred to Y and Y is preferred to Z then X is preferred to Z

[3]The independence axiom holds that whenever lottery A is preferred to lottery B, a gamble between lottery A and lottery C is preferred to an equivalent gamble between lottery B and lottery C.

[4]It is also not clear that they are normative bads (Persad 2014), although we focus our analysis on the descriptive question of what mitigates cognitive bias.

comparative efficacy of accountability mechanisms (Lerner and Tetlock 1999). This is particularly important in the context of reforms to administrative policymaking, where one accountability mechanism has received universal adoption and is deeply entrenched, while another has only received partial albeit growing adoption.

3 Data and Methods

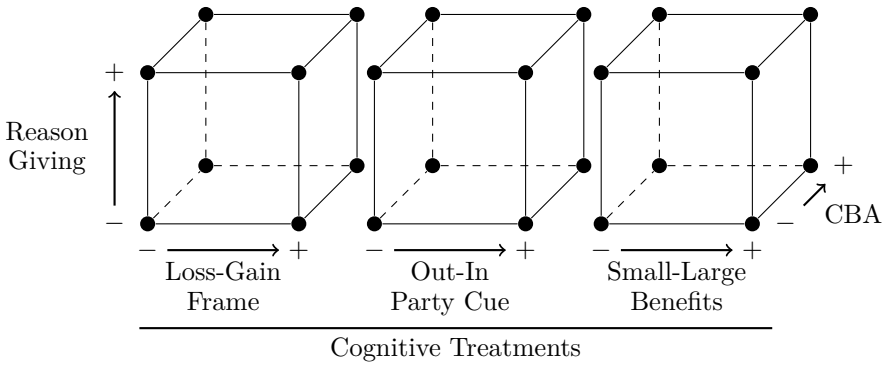

Figure 1: Design Representation

Our goal is to assess the plausibility of two mechanisms commonly billed as having curative properties for human cognition. Survey experiments identify cognitive biases through subtle, random manipulation of textual prompts given to respondents. The difference in average preference for some choice across randomly assigned treatment conditions provides evidence that a cognitive bias exists. Therefore, we are interested in whether the difference in average preferences is substantially smaller using one of the proposed procedures. In essence, we want to estimate an interaction effect between the cognitive bias and the procedure used to guide choice.

There are many cognitive biases, but resources for evaluation are limited. The cognitive biases we focus on are gain-loss framing, partisan motivated reasoning, and scope sensitivity. Prior work has shown that gain-loss frames and partisan cues can produce large effects on choice. The bias involved with scope-sensitivity is different. Here the problem is the *lack* of an impact. Prior work has found that respondents are insensitive to remarkable changes in the scale of the tradeoffs involved in choices. The fact that we include a "non-difference" cognitive bias creates a significant asymmetry in the experiments. Even so, we do include a test of scope sensitivity, because we consider it a "best case" for CBA to address. Put differently, if CBA does not perform well with improving scope sensitivity, then that would heighten skepticism about the procedure's ability to improve any cognitive biases.

We now make some breif remarks on our process for designing the textual prompts used in our surveys. Figure 2 show the prompts we used for gain-loss framing, partisan-reasoning, and scope-insensitivity, respectively. In the case of gain-loss framing, we decided to replicate precisely the language used in Tversky and Kahneman's article in *Science* about two alternative programs to fighting a novel disease originating in Asia. To determine the appropriate sample size for our experiment, we assumed a similarly large effect as they report.

We did not find an equally canonical test of partisan motivated reasoning in our review of the literature. Therefore, we decided to draft our own. Crucial to the text of partisan reasoning was a policy that could plausibly be associated with Democrat or Republican support. Republicans and Democrats alike have discussed infrastructure improvements in recent years. To validate that our textual prime would produce a large difference without any procedures, we performed a pilot study on m-Turk. We found that 76% of copartisans supported the bill, and support dropped to 44% when the policy was proposed by the opposite party, a large difference as expected.

Finally, concerning scope sensitivity, we initially explored a set of tradeoffs describe in Milkman et al. (2012). Their policy prompt involved creating a forest preserve at the cost of some jobs in the logging industry. In a pilot study on m-Turk, 80% of respondents supported the forest preserve regardless of whether it saved 10,000 or 100,000 acres, as expected. The pilot study raised concerns, however, that the baseline level of support was too high. With limited headroom to grow under alternative procedures, it seemed that we would need too large a sample to implement such a study. Given this concern, we switched attention to a less popular

(a) Gain versus loss framing.

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 poeple will die}** . Program A will cost $100 million. Program B will cost $100 million.

(b) Partisan-motivated reasoning

{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.

(c) Scope sensitivity

The government is considering expanding the availability 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.

Given the amount of land affected, the expansion is not expected to effect global commodity prices one way or the other.

Figure 2: Textual prompts

policy. In particular, we developed a textual prompt about expanding crop insurance, as described in the Figure. We did not find a statistically significant difference in support whether program expansion created 500, 2,500, or 5,000 jobs. We ultimately chose to compare 500 and 2,500 jobs because of their visual contrast.

For statistical efficiency reasons, we also adopt a factorial structure to the experiment. Figure 1 presents a concise visual representation of this structure. In each experiment, there are three randomly varied factors: a textual prompt designed to trigger cognitive bias ("cognitive treatment factor"), the reason-giving procedure, and the cost-benefit procedure. We assign each factor at two levels. In Figure 1, these are diagramed with "+" and "-". For the cognitive treatment factor, the "+" level refers to the one expected to generate a higher level of approval (i.e. "same-party sponsor"). For the procedural factors, "+" means using the procedure, and "-" means not. The factorial structure implies that each respondent in the study has a random and equal chance of being assigned one out of the $2^3 = 8$ possible combinations of factors and levels. Table 1 makes the possible assignments more explicit. A major benefit of this experimental structure is that, in cases where factors do not generate meaningful differences, one may collapse factors and increase one's sample size to evaluate the factors that do potentially matter. More concretely, if one finds reason-giving does not lead to differences in a cognitive bias, one has twice the n to analyze CBA alone. We fielded the experiment using a Qualtrics survey and a nationally representative sample acquired through Lucid. Given the online survey format, we assumed that there would be some amount of non-completion for whatever reasons. We, therefore, anticipated that each experiment would have an unbalanced 2^3 factorial structure and that some combinations of treatments would have more units than others.

Both quantitative cost-benefit analysis and qualitative reason-giving could take many different forms in any particular study like ours. One could conceive of these forms as imposing more or fewer constraints on respondents. For example, if one asks a respondent to explain their decision, one could impose minimum length requirements or give guidelines. One could ask individuals to consider particular aspects of the decision in specific. With cost-benefit analysis, there is extensive debate about what factors might go into the analysis. Should one explicitly account for risk? If so, how? What about distributional considerations? In both cases, we have erred toward the unconstrained end. We provide no constraints on explanations nor give any guidance as to what the justification should cover. While the cost-benefit analysis is more structured, we generally ask for a valuation of one unit of a good and ask individuals to multiply this valuation by the number of goods a policy will obtain. We then ask for the respondent to calculate the net benefit before responding. In cases where both CBA and reason-giving are assigned, we always ask for the CBA before the qualitative justification. Similarly, we always ask for the individual's choice after they have completed the procedural task.

	Treatment Factors			
Experiment	**Cognitive Treatment**	**CBA**	**Reason-giving**	
Risk Framing	Lives Saved	No	No	Difference Expected
	Deaths Prevented	No	No	
	Lives Saved	Yes	No	
	Deaths Prevented	Yes	No	
	Lives Saved	No	Yes	
	Deaths Prevented	No	Yes	
	Lives Saved	Yes	Yes	
	Deaths Prevented	Yes	Yes	
Partisan Framing	Copartisan Cue	No	No	Difference Expected
	Opposing Party Cue	No	No	
	Copartisan Cue	Yes	No	
	Opposing Party Cue	Yes	No	
	Copartisan Cue	No	Yes	
	Opposing Party Cue	No	Yes	
	Copartisan Cue	Yes	Yes	
	Opposing Party Cue	Yes	Yes	
Scope Sensitivity	500 Jobs	No	No	Difference Not Expected
	2500 Jobs	No	No	
	500 Jobs	Yes	No	
	2500 Jobs	Yes	No	
	500 Jobs	No	Yes	
	2500 Jobs	No	Yes	
	500 Jobs	Yes	Yes	
	2500 Jobs	Yes	Yes	

Table 1: Planning Matrix

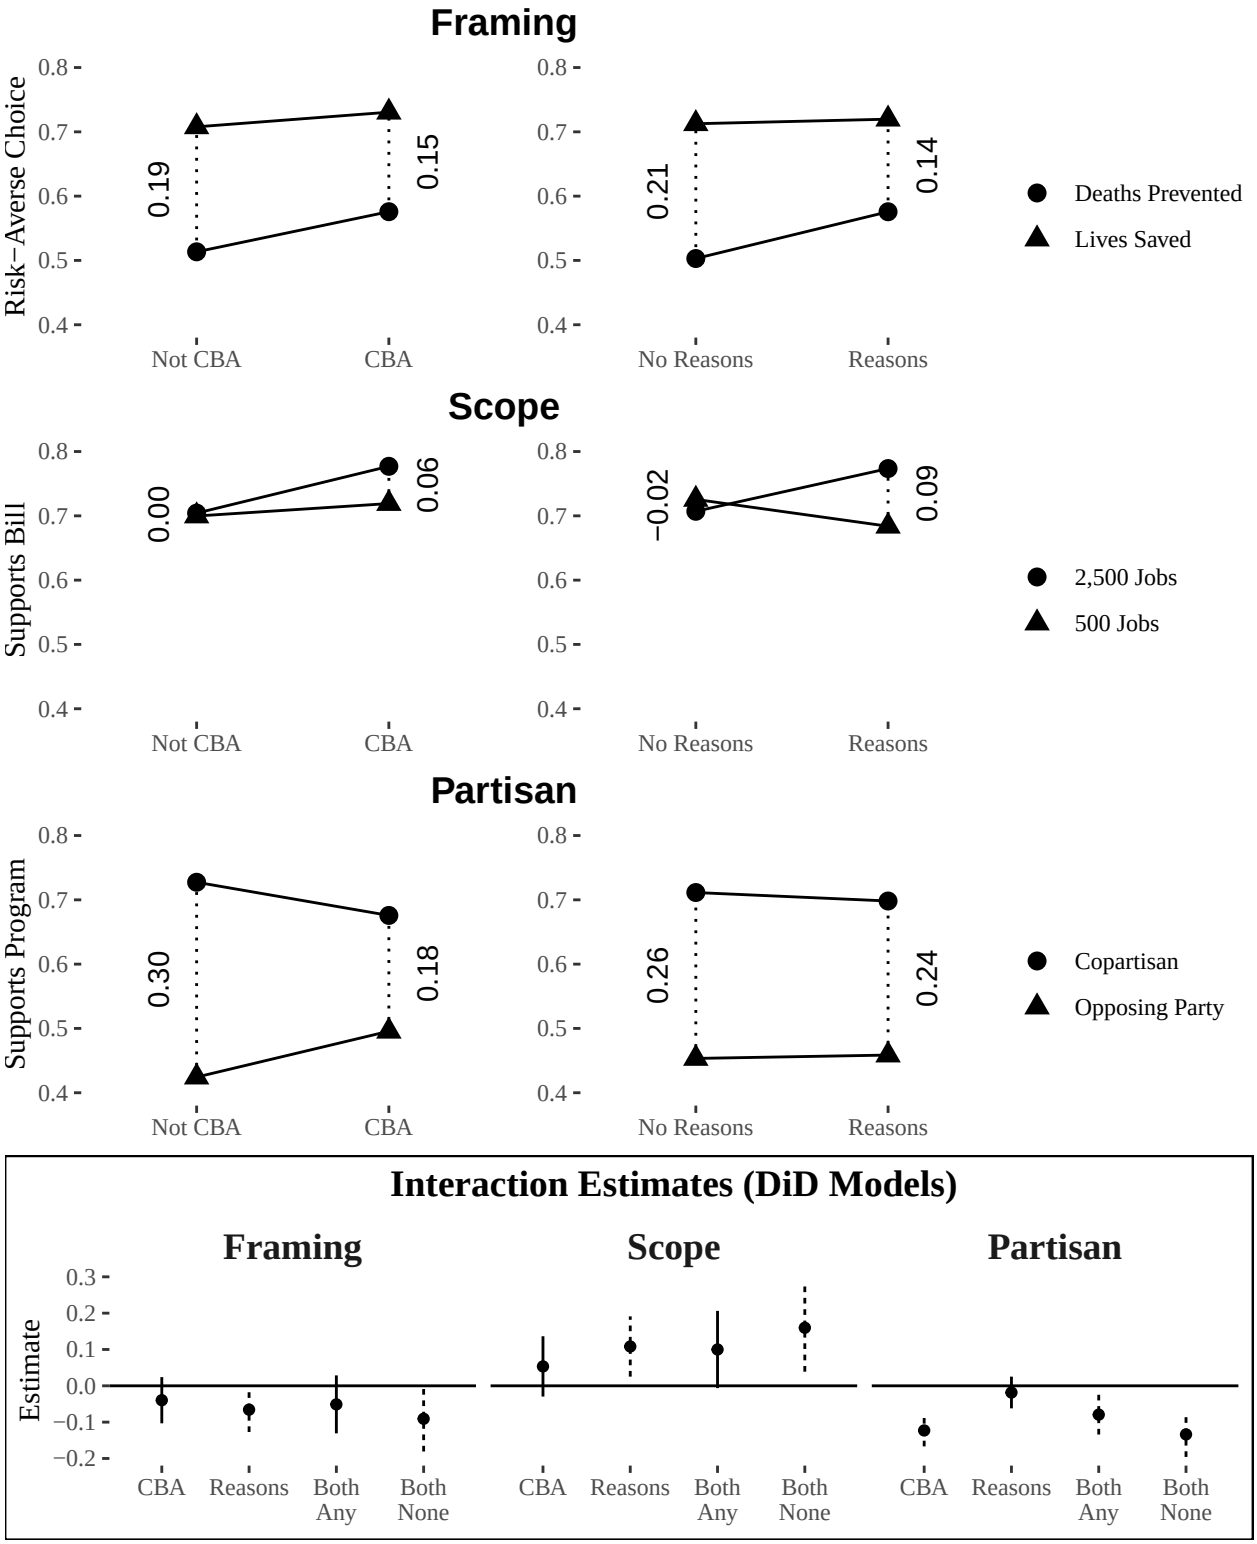

Figure 3: Interaction of Cognitive Treatments and Procedures (Pooled).

4 Results

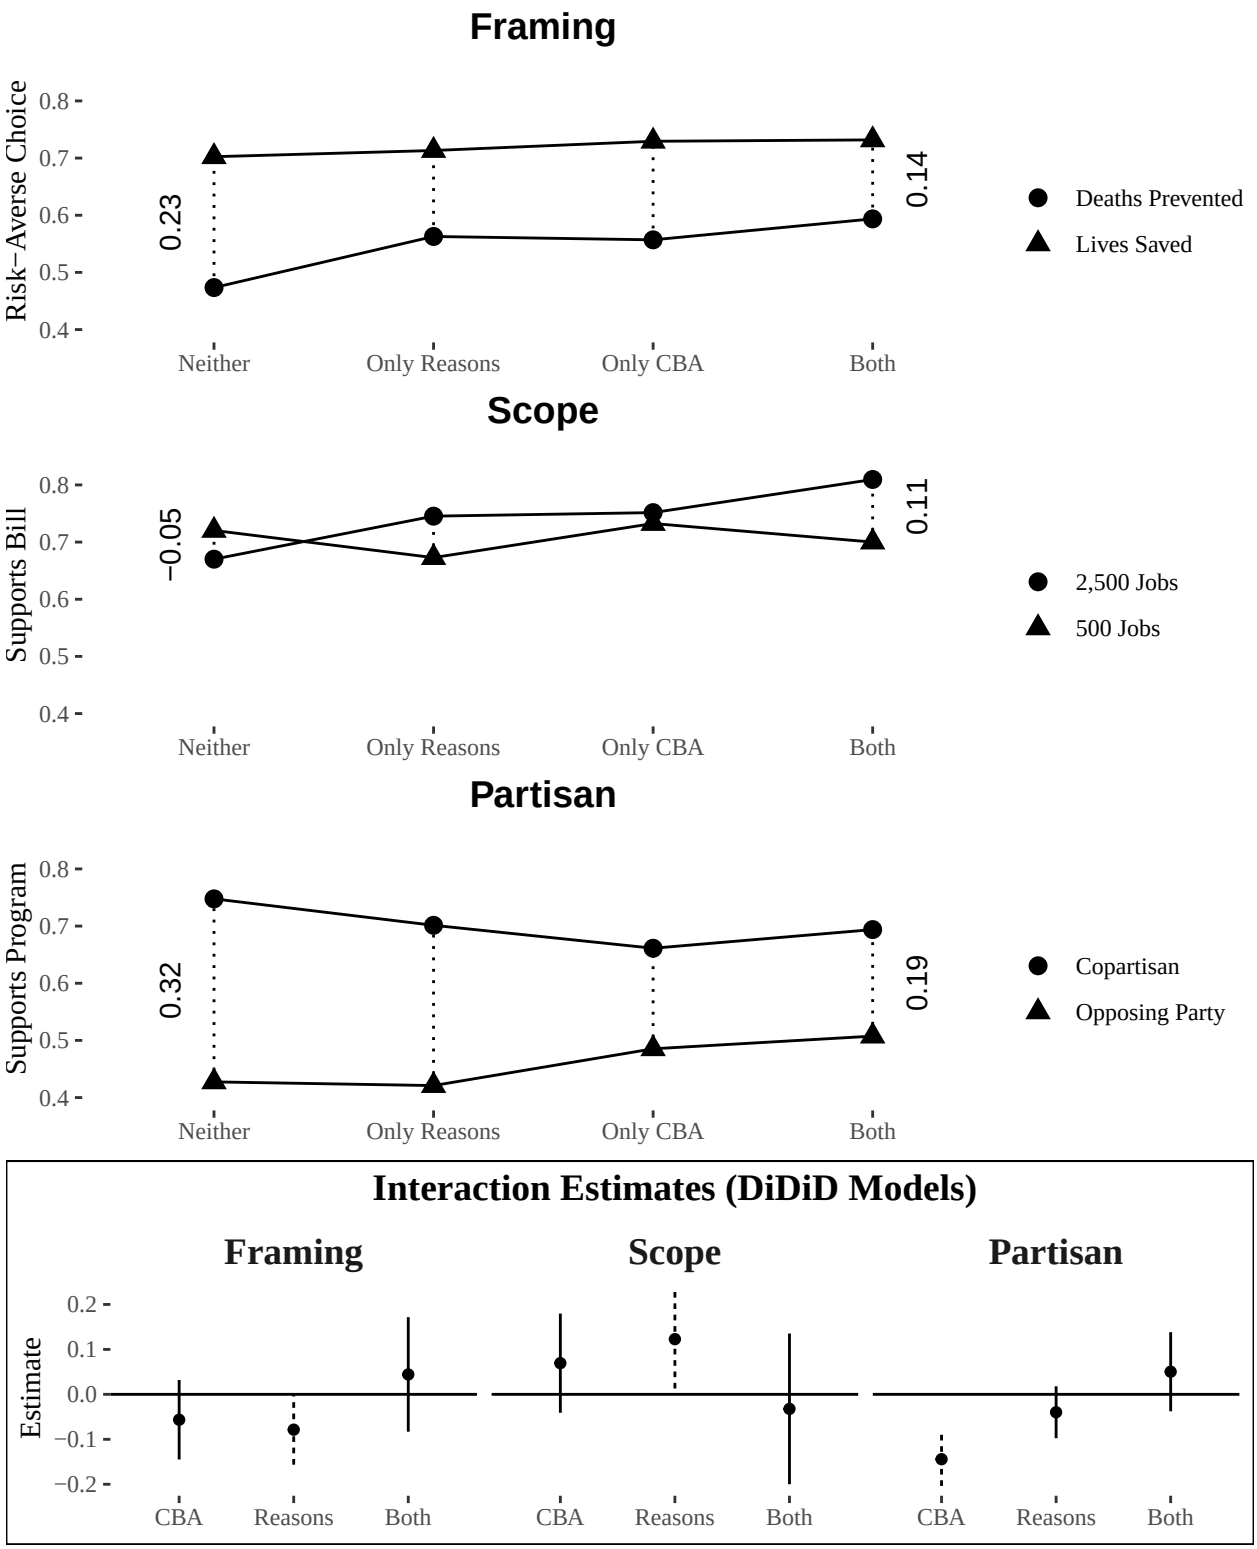

Each experiment presents respondents with a choice between two options. For expositional purposes, we assume that the "risk-averse" choice in the framing experiment is option one, while "supporting the policy" is option one in the other experiments. Figures 3 and **??** present the proportion of individuals choosing option one depending on the assigned procedure and cognitive treatment. Differences along the y-axis reflect the

impact of what we call the "cognitive treatment" that demonstrates a cognitive bias. The x-axis distinguishes between the procedure respondents followed before making a policy choice. Of principal interest is how procedures influence cognitive bias, which one can observe by comparing how the difference on the y-axis differs on the x-axis. The bottom panel of both figures illustrates estimates of the difference-in-difference parameter along with calibrations for uncertainty, details on which we explain below.

Factorial experiments are a statistically efficient design, however, they do pose some complexities due to the treatment patterns. In these experiments in particular, there are four possible procedures assigned with equal probability by the experiment ({CBA,Reasons},{No CBA,Reasons},{CBA,No Reasons},{No CBA,No Reasons}). In the approach of Figure 3, we look at the effect of being assigned one procedure while ignoring the role of the other procedure. More specifically, the CBA responses are calculated by "pooling" the subjects assigned {CBA,Reasons} and {CBA,No Reasons} while the No CBA responses combine respondents who received the other assignments. The difference-in-difference estimates presented in the bottom panel are calculated using ordinary least squares, as are the standard errors. Hence, the effect sizes correspond exactly to what is visible in the top three panels. The Appendix also includes difference-in-difference estimates after controlling for available background covariates such as education or household income, which do not substantively influence the estimates. Also pictured, since estimated from the same difference-in-difference equations, are the effect of both procedures together versus any alternative procedure (i.e. {CBA,Reasons} versus {No CBA, Rasons},{CBA,No Reasons}, and {No CBA, No Reasons}) and both procedures versus no procedure (i.e. {CBA,Reasons} versus {No CBA,No Reasons}).

The most important observation from 3 is that, in every case, one or both procedures improved the cognitive bias. Taking the partisan experiment first, asking for explanations generated almost no impact on individual's susceptibility to the partisan cue, but the cost-benefit procedure practically halved this susceptibility ($t \approx -4.61$). The other panels seem to suggest less dramatic differences between the two procedures, with reason-giving producing differences-in-differences that look only somewhat larger in absolute magnitude. In fact, the numerical differences are more striking than the visualization suggests because those not assigned reason-giving were more responsive to the cognitive treatment than those not assigned CBA. Multiplying the number of jobs created by the program generated no difference in support for those not assigned CBA while doing the same decreased support by 2% among those not assigned to give reasons. As a result, the difference-in-difference associated with CBA was 6% in the scope experiment, but the difference-in-difference associated with reason-giving was 11%, almost double. Similarly, for the framing experiment, the difference-in-difference for CBA implied a 4% decrease in susceptibility to the cognitive treatment versus a 7% decrease for reasons, again almost double. These large differences in magnitudes matter for statistical significance. There is minimal evidence to reject the hypothesis of no impact to CBA in either the framing or scope experiments ($p \approx 0.305$ and 0.29,respectively. By contrast, the two-sided p-value associated with reason-giving was 0.0795 with respect to framing and 0.0321 with respect to scope. While we do not support an overly rigid approach to significance testing, one-sided p-values are arguably more appropriate here since one does not expect any procedure to exacerbate cognitive bias. In any event, we think there is strong evidence that reason-giving influenced respondents to select choices in a more substantively rational way in response to challenges based on gain-loss framing and magnitudes of benefit. These regressions provide almost no evidence supporting CBA's impact on these same challenges. The situation is the reverse for partisanship.

Figure **??** presents an alternative approach to analyzing the factorial experiments. Here each of the four possible procedures is treated as a separate procedure and not pooled. The bottom panel, which calibrates uncertainty associated with these difference-in-differences, is constructed using a difference-in-difference-in-difference (DiDiD) model estimated using ordinary least squares. In the bottom panel, "Both" reflects the triple-interaction term, so is not the same as what is visible in the top three panels.[5] Figure **??** clarifies several points. First, the largest reduction in susceptibility to the partisan cue came from being assigned CBA alone; assigning CBA and reason-giving together actually generated somewhat larger susceptibility to the partisan cue (although plausibly due to sampling alone). Second, the impact of CBA in the scope experiment was driven by individuals also assigned reason-giving. Visually, at least, CBA seemed to play more of the role of enhancing some underlying efficacy of reason-giving. Third, the differences between procedures in the framing experiment again appear even more subtle than they did in Figure 3. Finally, the bottom panel clarifies that the DiDiD formulation leads to substantively very similar conclusions as the DiD formulation. All estimates of the procedures' impacts on cognitive bias are on the expected side of the line. The same estimates are statistically significant (or not) as in Figure 3. The only additional noteworthy point relates to the triple-interaction term under "Both" in the bottom panel. It is never statistically significant but always

[5]The y-axis difference in the top three panels under "Both" is therefore the sum of the estimates captured in the bottom panel

has the opposite sign of each procedure by itself. The implication is that there is no evidence that these procedures have useful synergies, i.e. are more than the sum of their parts. If anything, the consistency of the wrong sign is limited evidence that the processes interact antagonistically, so they become less than the sum of their parts. Still, the degree of antagonism is mild in most cases, so the cognitive bias is generally less or about the same with both procedures as it is with the assignment of the "best" procedure.

5 Survey Attrition and Treatment Compliance

Two other topics warranting explanation from a data analysis perspective are survey attrition and treatment compliance. Procedures requiring justification of preferences impose costs upon survey respondents. As a result, some respondents may choose to drop out of the study before expressing a preference (attrition). Others may express a preference but put little effort into following the procedures (non-compliance). It is worth noting that these issues are very much central to the debates about the advisability of CBA generally.[6] Yet, these related cost-avoidance strategies by survey respondents create analytical concerns as a threat to validity and reproducibility. One major benefit of the survey platform we use is that we have background demographics on individuals who participated in the study, even if they attrited before making a policy selection or did a poor job complying with the procedures.

Figure 4 shows the number of individuals assigned to each procedural treatment arm and the proportion completed. As expected, random assignment distributed procedures evenly across respondents. Those assigned the "naked" choice procedure hardly attrited in any experimental arm (completion rate ranged from 97-99%). An open-ended request for justification, however, led to some notable attrition (82-89% completion rate), while asking for quantitative justification led to substantial dropout (57-75% completion rate). Assigning both procedures caused an even larger proportion of respondents to quit the survey (54-62%). Unsurprisingly, the biggest declines in completion occurred under the framing experiment, where the calculations were the most complicated.

One concern with attrition is limited statistical power. We consider this concern mild. We chose sample sizes for these experiments based on power calculations that assumed (1) effects under the control condition that we determined from pilot studies and (2) scholarly priors about what sized reductions in bias would have substantive meaning. Many of the effects we found were much larger than those minimal benchmarks. The sample size calculations were, fortunately, conservative, even after attrition. In any event, the procedural combination with the fewest completions (Scope Experiment with Both procedures) *still* had 236 responses. Attrition did not transform a large n into a small n study.

The more serious concern is not attrition per se, but rather *differential attrition* across procedural arms. It is not surprising that different procedures impose different costs on respondents. Yet, it does make the studies susceptible to subtle problems of selection. Traits that lead some individuals to perceive relatively higher costs in following some procedures might relate to their choices in response to the policy prompts. At worst, any apparent reductions in bias from one procedure might purely result from changes in the composition of the pool of respondents, which might be so large as to overshadow the true and, at least potentially, opposite effect.

While differential attrition does pose a substantial internal validity concern, the possibility of such selection effects does not prove their existence. Moreover, we can evaluate the plausibility of such selection mechanisms with data. Figure 5 presents a balance plot leveraging the demographic characteristics provided by the survey platform, including age, education, gender, household income, and race and ethnicity. Each dot shows the standardized difference in the covariate between the control condition of no procedure in one of the experiments, which had almost no attrition, and one of the other procedures. The largest standardized difference in any covariate between control and active treatment in any experiment was about 1/4 of a standard deviation. This imbalance occurred in the scope experiment, with 54% of individuals who expressed a preference after giving reasons and doing the quantitative exercise were male, while only 41% of individuals who completed in the control condition were. While a 13% difference may appear substantial, a simple simulation study can confirm that this is not a hugely surprising difference for 250 draws from a 50-50 trait in the general population. And there are hundreds of standardized differences visible on the chart (315 to be exact). While we cannot rule out the possibility that there are unobserved covariates related to choices that have become imbalanced due to attrition, one presumes that they would have some correlation with education, household income, or some other covariate that we do observe. As the balance problems do not appear severe,

[6] In particular, scholars have argued that CBA in the real world attrites beneficial policies, rather than survey respondents, and that agencies may shirk or even strategically misrepresent in their justifications.

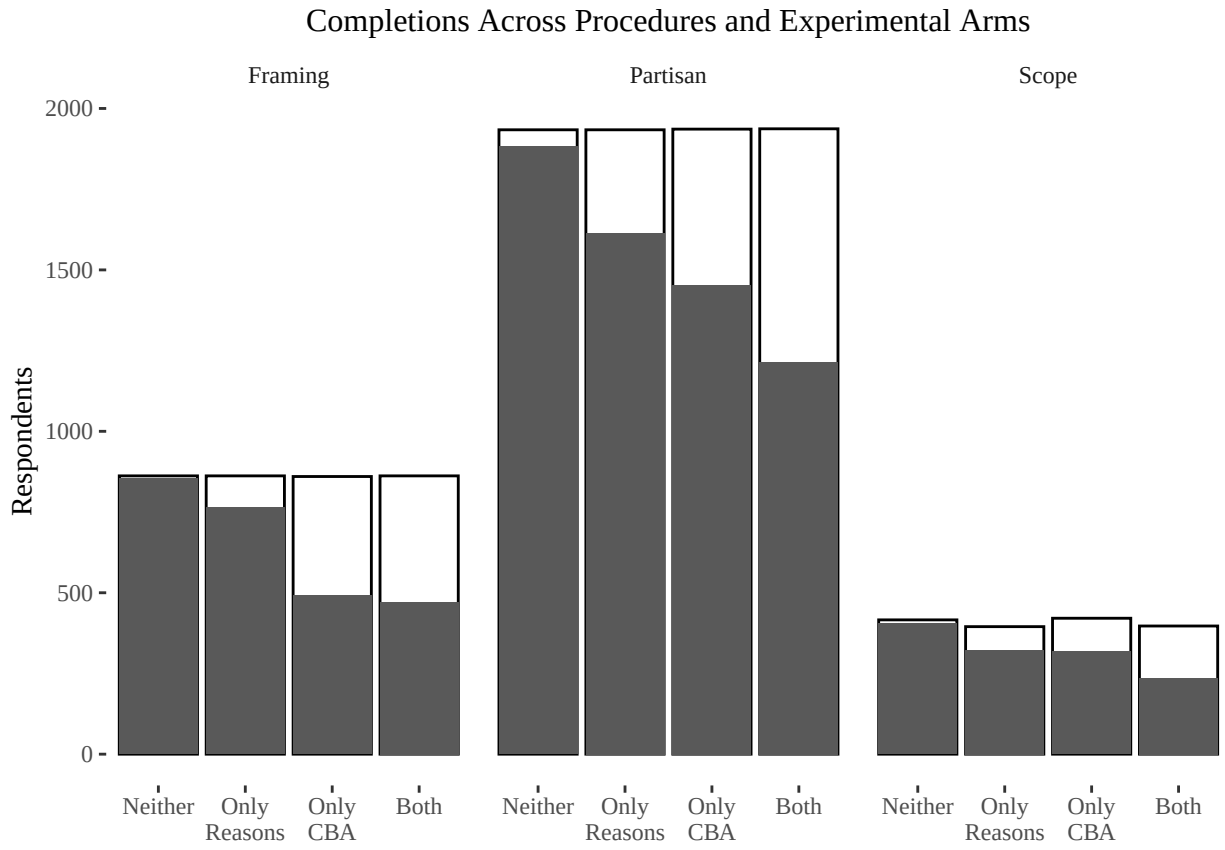

Figure 4: Extent of Attrition.

we decline to present results that attempt to correct for attrition based on observables.[7] Replication does more to prove internal validity than different analysis techniques ever will. Therefore, we consider it best to document the issue as we see it and move along.

Experimentalists often distinguish between effects due to assigning treatment and those actually due to receiving treatment. The difference is between the impact of being told to take an aspirin and actually taking aspirin. In this context, we ask whether those who follow the assigned procedure behave differently than those who do not. Although subtle, there is an essential connection between the treatment compliance question and the counter-factual effect of attrition. If one could force those who dropped out to complete, they would presumably be more likely to produce non-compliant responses to our procedural tasks than those who did not drop out. Instrumental variables methods are the primary tool used to distinguish the effect of compliance from being assigned treatment. The application of these methods to factorial experiments has recently received some attention from methodologists (Blackwell and Pashley 2020).

In applying the IV approach, the first step is determining whether non-compliance is one or two-sided. Already here, we run into difficulties. Our discussion focuses on compliance with CBA and ignores the complexities of interactions with the reason-giving procedure. We can define "treatment" with CBA in this context as producing arithmetic that uses the respondent's valuation on a particular good to calculate (a) the benefit or (b) the net benefit of the policy described. Two-sided non-compliance is conceivable here: it is hard to rule out the possibility that respondents are doing cost-benefit calculations before making a policy selection even when we do not ask them to do so or collect records from them.[8] We do not have any way of knowing what their treatment status was, leading to the uncomfortable position of needing to make an implausible one-sided compliance assumption to proceed with the analysis.

[7] Our read of the methodological literature suggests that there is limited agreement about the best thing to do in such cases, especially given the factorial structure of the experiment.

[8] It is even harder, perhaps impossible, to rule out the possibility that respondents are providing reasoned justifications to themselves before making a selection even though we do not ask them to.

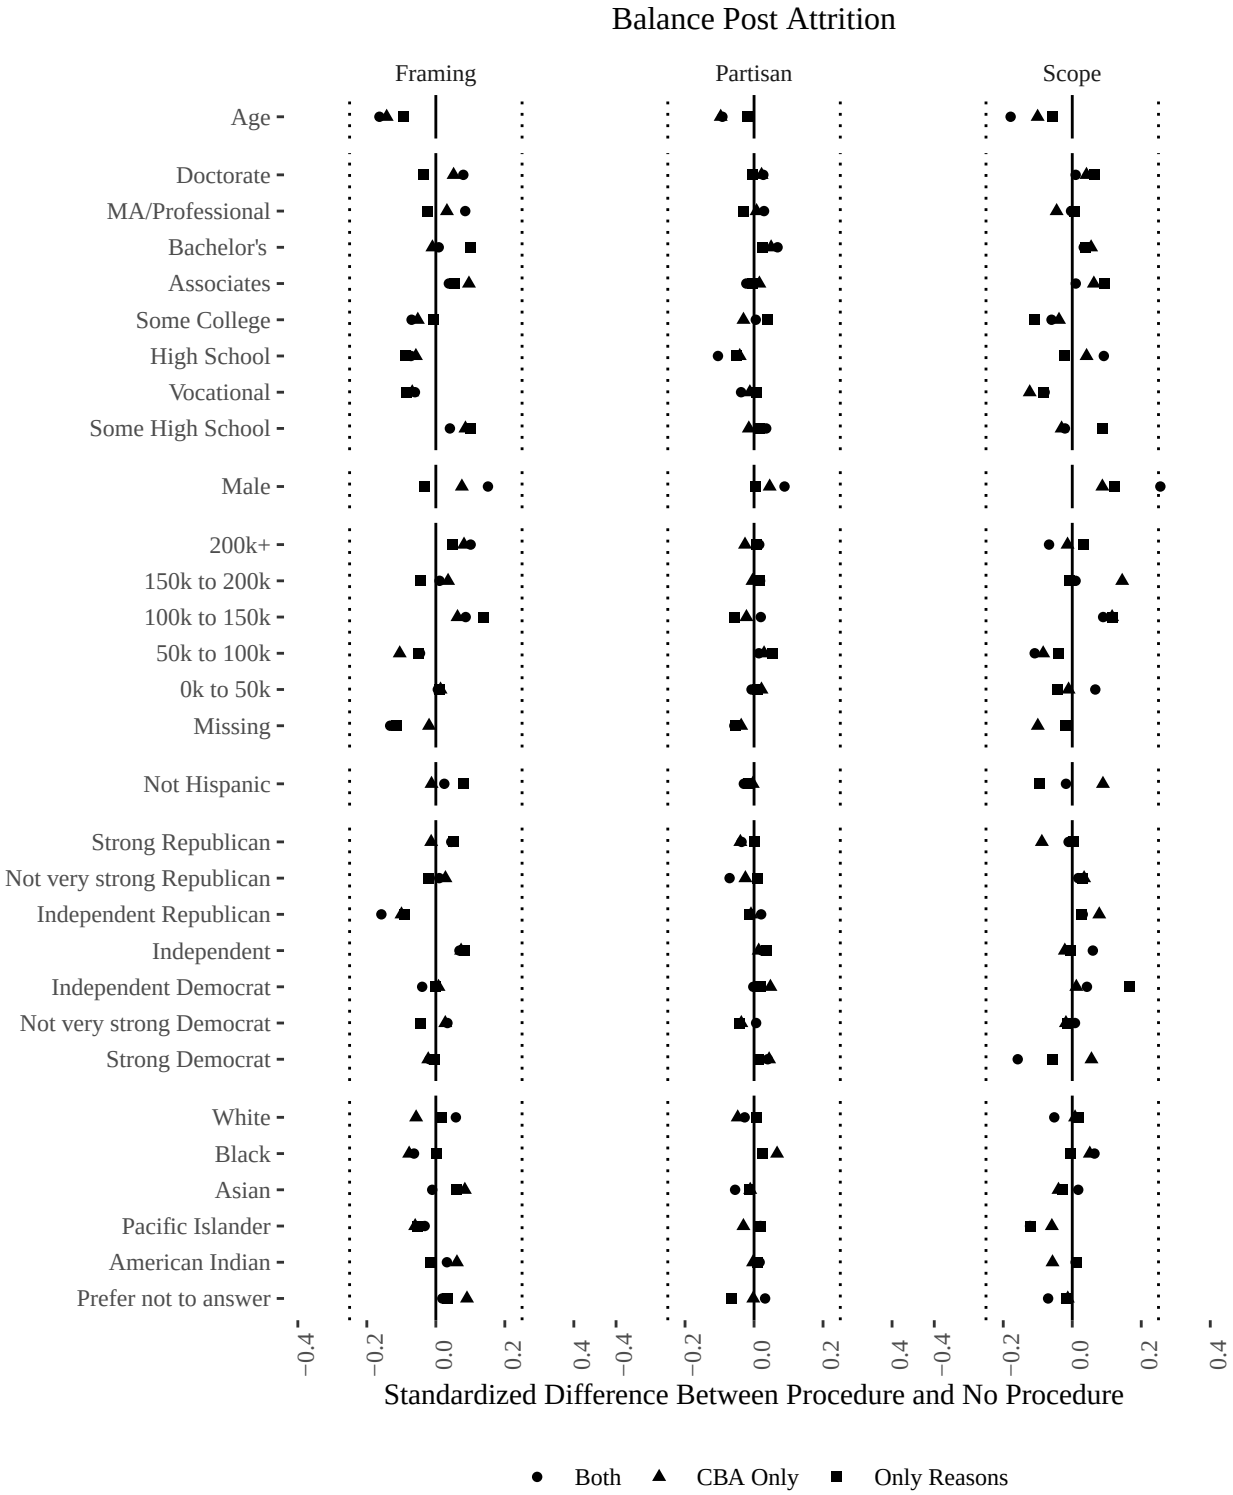
Balance Post Attrition
Framing
Partisan
Scope
Age
Doctorate
MA/Professional
Bachelor's
Associates
Some College
High School
Vocational
Some High School
Male
200k+
150k to 200k
100k to 150k
50k to 100k
0k to 50k
Missing
Not Hispanic
Strong Republican
Not very strong Republican
Independent Republican
Independent
Independent Democrat
Not very strong Democrat
Strong Democrat
White
Black
Asian
Pacific Islander
American Indian
Prefer not to answer
−0.4
−0.2
0.0
0.2
0.4
Standardized Difference Between Procedure and No Procedure
Both
CBA Only
Only Reasons

Figure 5

Even assuming we were willing to make that assumption, a further difficulty relates to IV's exclusion restriction. To illustrate the problem, we leverage a set of questions from the partisan and scope experiments. After respondents had made their policy selection in these experiments, individuals in the control condition were to perform a cost-benefit analysis.[9] Using this information, we can build Figure 6, which decomposes the previously identified CBA effect by the plausibility of the analysis the individual produced at some point in the experiment. Those assigned CBA did it before selecting a policy, while those assigned control did it after choosing a policy.[10] Noteworthy here is that the individuals who did not produce a valid CBA were the major drivers of the apparent reduction in partisan sensitivity. While there is less obvious evidence against the exclusion restriction in the scope experiment, these concerns are severe enough that we decline to estimate the CACE through IV.

Nevertheless, we think it is worth considering the average effect of assigning cost-benefit analysis to those who did it. Suppose the assumption that whether an individual produces a plausible CBA is not dependent on when exactly they are asked to do so (i.e. before or after making a policy choice).[11] In that case, the average treatment effect on the treated is already visible in Figure 6. Individuals assigned CBA who produced a plausible analysis were twice as likely to approve a policy when it created five times the jobs as those who made a plausible analysis but were assigned control. Similarly, partisan favoritism decreased 5% through assigning CBA to those who produced plausible analyses. The magnitude of the benefit through CBA among the compliers in the scope expeirment was similar to the effect of being assigned reason-giving reported earlier. The magnitude of the benefit through CBA among compliers in the partisan experiment was much smaller, however, as the larger part of the effect passed through those not doing plausible calculations. Although we cannot make the same claim at estimating the average treatment effect on the treated for the framing experiment, it is worth noting that those assigned CBA who did a plausible analysis were not sensitive to the gain or loss framing.

[9] We declined to do the same for the framing experiment because the cost-benefit on that policy problem is particularly complex, and we were worried about the number of questions leading to attrition, a concern that 4 justifies to some extent.

[10] It is tempting to try and use this to infer the treatment status of those assigned to control. For example, one might entertain an assumption that respondents who do plausible CBA after making a policy choice would have also done a plausible job if they had been asked to do so before (and if not, then not). Closer reflection will lead one to realize that implicitly one would then be assuming no compliers, which is not a reasonable assumption for estimating complier average causal effects.

[11] This assumption is supported by the fact that the proportion of individuals producing valid CBAs is constant across treatment conditions in each experimental arm.

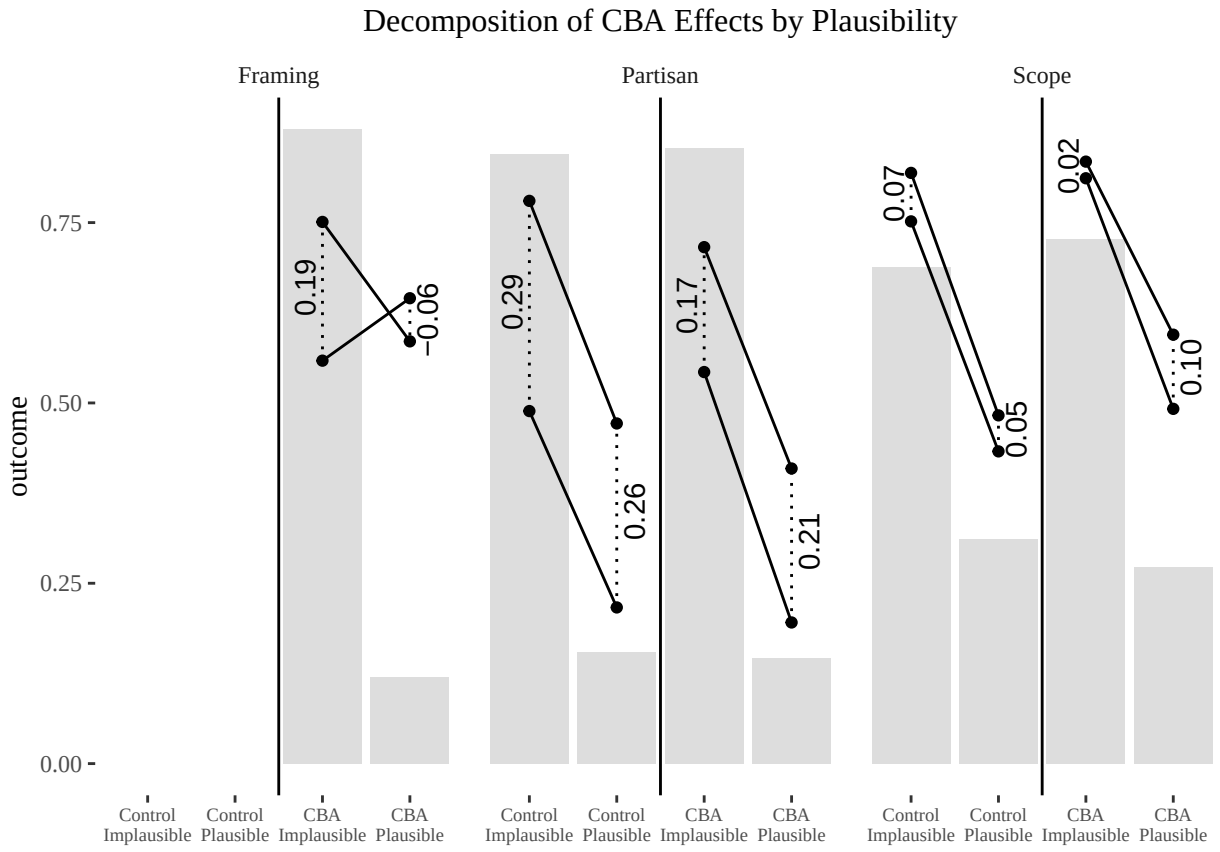

Figure 6

6 Discussion

	All Respondents			Only Compliers
	CBA	Reasons	Both	CBA
Framing	-	✓	✓	-
Scope	-	✓	✓	✓
Partisan	✓	-	✓	-

Table 2: **Summary of Findings**. Checkmarks indicate where we are able to reject the null of no interaction between procedure and cognitive treatment.

Scholars and policymakers have proposed interventions designed to improve rationality as if all rationality problems were fundamentally similar. If that were true, it would follow that any intervention that can help with one cognitive failing must help with all. Our experimental findings profoundly challenge this assumption: different procedures produce different improvements for different cognitive failings. Table 2 summarizes the numerous findings of our studies. The key takeaway: reason-giving did much more than CBA for framing effects and scope sensitivity. In contrast, reason-giving did practically nothing to decrease partisan-motivated reasoning, and CBA had a big effect. We also examined separately how cognitive bias changed among those who did a plausible quantitative analysis relative to those that did not perform a plausible analysis. We found that doing the CBA appeared to eliminate sensitivity to framing and might increase sensitivity to scope but did not reduce sensitivity to the partisan identity of the proponent of the policy. All reductions identified in the study as statistically significant appear substantively meaningful. Even small, systemic improvements in policymaker rationality should translate into important welfare improvements. As with all survey-style experiments, how these findings translate into the real world remains somewhat mysterious. However, large

effects in a lab setting seem to be a precondition for believing that these interventions will have a meaningful impact in the real world.

These findings raise several important questions. Are these interventions in policymaking substitutes or compliments? Is rationality more improved through both procedures than through one alone? The answer to that particular question seems to be a qualified yes. The combined process generally obtained improvements in cognitive bias that were qualitatively and statistically similar to the most effective treatment. For this reason, if one were to guess which treatment would be most effective at ameliorating a novel cognitive bias that we do not test, one would most reasonably select the compound procedure. Since something like the compound procedure now governs most regulatory policymaking at the US federal level, the findings may leave one encouraged.

Yet, there are caveats. First, the triple-interaction of CBA, reason-giving, and cognitive bias always has the "wrong sign." The triple-interaction term measures the degree of synergy between procedures. For example, if the triple-interaction were negative in the framing effect experiment, as some might have expected, it would imply that CBA brought some reduction in framing effects, reason-giving brought some more. Still, the combination of both CBA and reason-giving was even greater than the sum of its part. The problem is that the estimates in Table **??** (4) imply that CBA with reason-giving is *less* than the sum of its parts. Indeed, no experiment found useful synergies between the procedures. Instead, reason-giving and CBA seem to have an antagonistic effect for cognitive bias. Admittedly, these wrongly-signed estimates were never statistically significant. Still, in all three cases, the estimate was substantial, in some cases enough to wipe out the implied benefit of the less useful second procedure and in others to substantially reduce it. Second, and relatedly, there is no evidence of the combined process performing significantly better than the best-performing single procedure against any particular cognitive bias. Put differently, one process always drove the combined effect against a specific cognitive bias. Third, there is no free lunch. Procedural rationality requirements have costs. Generally, one would not select the most invasive procedure unless it produced significantly and substantially more significant benefits than a less invasive procedure.

Therefore, the primary justification for the use of combined CBA and reason-giving procedures appears rooted in the precautionary principle. The exact nature of the cognitive biases at play in policymaking is varied. They may include the deviations from rationality explored here to differing degrees, depending on the policy at play or the regulators involved. They likely have still other problems not explored. The results confirm that the combined use of reason-giving and CBA is a plausible mechanism for reducing such biases. However, the results are also consistent with CBA not having any plausible use as a rationality improvement device except for highly-politicized policies and particularly susceptible to partisan-motivated reasoning. The actual assignment of procedures to regulators does not track the kinds of biases that are particularly likely to be at play but rather institutional and historical factors. In particular, reason-giving applies to all, and CBA applies to agencies that are not independent.

A second and related question is what to make of the difference in behavior caused by compliance with the CBA procedure. The fact that the efficacy of treatment would depend on compliance is seldom surprising. The shape of dependency uncovered in these experiments is problematic for CBA. There is extensive debate about how incentivized regulators are to judge costs and benefits accurately. To the extent that they have incentives to produce facially valid analyses, there are also concerns about perverse incentives to tip the scales in other ways, perhaps through inflating statistics or identifying more benefits or costs. CBA only achieves equivalent cognitive benefits as reason-giving when respondents comply, otherwise, it seems to do nothing in both the scope and framing experiments. By contrast, CBA was better at helping respondents reason despite partisan source cues, where giving reasons failed. Here, however, actually doing the CBA there did not seem to bring any real improvement. As a result, one could believe that in the area of greatest strength for CBA, a less costly procedure that was equally good at focusing attention would produce similar cognitive benefits, perhaps with fewer compliance problems.

If anything, these experiments point to the under-appreciated benefits of asking for written justification as a device for improving rationality. Indeed, our experiments provided a very unrestrictive invitation to justify one's decisions. There were no length requirements. We did not ask to consider the merits of each policy. The minimum amount of cognitive effort necessary to complete the task was lower than the case of CBA. Even so, the opportunity to provide a reasoned justification did a great deal to ameliorate cognitive bias. In most experiments, an invitation to give reasons did more than an invitation to quantify, and in some experiments, more than actually doing a plausible quantitative analysis.

References

Adcock, Robert. 2001. "Measurement validity: A shared standard for qualitative and quantitative research." *American Political Science Association* 95 (3): 529–46. http://journals.cambridge.org/abstract{_}S0003055401003100.

Allais, Maurice. 1953. "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine." *Econometrica: Journal of the Econometric Society*. JSTOR, 503–46.

Blackwell, Matthew, and Nicole E. Pashley. 2020. "Noncompliance and Instrumental Variables for 2ˆK Factorial Experiments," 1–53. https://mattblackwell.org/files/papers/factorial-iv.pdf.

Bolsen, Toby, James N. Druckman, and Fay Lomax Cook. 2014. "The Influence of Partisan Motivated Reasoning on Public Opinion." *Political Behavior* 36 (2): 235–62. https://doi.org/10.1007/s11109-013-9238-0.

Brader, Ted, and George E. Marcus. 2013. *Emotion and Political Psychology*. July. https://doi.org/10.1093/oxfordhb/9780199760107.013.0006.

Carpenter, Daniel P. 2001. *The Forging of Bureaucratic Autonomy: Reputations, Networks, and Policy Innovation in Executive Agencies, 1862-1928*. Princeton, New Jersey: Princeton University Press.

Cox, Gary W., and Mathew D. McCubbins. 1993. *Legislative Leviathan: Party Government in the House*. Berkeley, CA: University of California Press.

DeMuth, Christopher. 2016. "Can the administrative state be tamed?" *Journal of Legal Analysis* 8 (1): 121–90. https://doi.org/10.1093/jla/law003.

Druckman, James N., Thomas J. Leeper, and Rune Slothuus. 2018. "Motivated Responses to Political Commnuications." In *The Feeling, Thinking Citizen*, edited by Howard Lavine and Charles S. Taber, 125–50.

Dudley, Susan E. 2021. "Milestones in the Evolution of the Administrative State." *Daedalus* 150 (3): 33–48. https://doi.org/10.1162/daed_a_01858.

Dunlop, Claire A., and Claudio M. Radaelli. 2016. "The politics and economics of regulatory impact assessment." In *Handbook of Regulatory Impact Assessment*, 3–18. Edward Elgar Publishing. https://doi.org/10.4337/9781782549567.00011.

Eguia, Jon X., and Kenneth A. Shepsle. 2015. "Legislative Bargaining with Endogenous Rules." *The Journal of Politics* 77 (4): 1076–88. https://doi.org/10.1086/682389.

Elliott, E. Donald. 1992. "Re-Inventing Rulemaking." *Duke Law Journal* 41: 1490.

Ernst, Daniel R. 2014. *Tocqueville's Nightmare*. Vol. 15. 1. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199920860.001.0001.

Francesco, Fabrizio de, Claudio M. Radaelli, and Vera E. Troeger. 2012. "Implementing regulatory innovations in Europe: The case of impact assessment." *Journal of European Public Policy* 19 (4): 491–511. https://doi.org/10.1080/13501763.2011.607342.

Gailmard, Sean, and John W. Patty. 2016. "Participation, process and policy: the informational value of politicised judicial review." *Journal of Public Policy* 1 (28). Cambridge University Press. https://doi.org/10.1017/S0143814X16000106.

Gellhorn, Walter. 1986. "The Administrative Procedure Act: The Beginnings." *Virgina Law Review* 72 (2): 219–33.

"Global Indicators of Regulatory Governance: Worldwide Practices of Regulatory Impact Assessments." 2018. Washington, D.C.: World Bank Group. http://documents.worldbank.org/curated/en/905611520284525814/Global-Indicators-of-Regulatory-Governance-Worldwide-Practices-of-Regulatory-Impact-Assessments.

Gormley, William T. 1989. *Taming the bureaucracy : muscles, prayers, and other strategies*. Princeton, New Jersey: Princeton University Press.

Johnson, Ronald N., and Gary D. Libecap. 1994. *The Federal Civil Service System and the Problem of Bureaucracy*. https://doi.org/10.7208/chicago/9780226401775.001.0001.

Jones, Bryan D. 1999. "Bounded Rationality." *Annual Review of Political Science*, no. 2: 297–321.

Jones, Bryan D., Sean M. Theriault, and Michelle Whyman. 2019. "Causes of the Great Broadening: The Role of Social Movements." In *The Great Broadening*, 109–30. https://doi.org/10.7208/9780226626130-007.

Ladd, Jonathan Mc Donald, and Gabriel S. Lenz. 2008. "Reassessing the role of anxiety in vote choice." *Political Psychology* 29 (2): 275–96. https://doi.org/10.1111/j.1467-9221.2008.00626.x.

Lerner, Jennifer S, and Philip E Tetlock. 1999. "Accounting for The Effects of Accountability Distilling the Wisdom of Crowds: Prediction Markets versus Prediction Polls View project" 125 (2): 255–75. https://www.researchgate.net/publication/13203036.

Libgober, Brian. 2020. "Strategic Proposals, Endogenous Comments, and Bias in Rulemaking." *Journal of Politics*.

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*, edited by Arthur Lupia, Mathew D. McCubbins, and Samuel L. Popkin, 183–213. New York, NY.

———. 2013. *The rationalizing voter*. Cambridge University Press.

Machina, Mark J. 1987. "Choice Under Uncertainty: Problems Solved and Unsolved." *Journal of Economic Perspectives* 1 (1): 121–54. https://doi.org/10.1257/jep.1.1.121.

Marcus, George E., W. Russell Neumann, and Michael MacKuen. 2000. *Affective intelligence and political judgment*. University of Chicago Press.

Mashaw, Jerry L. 2018. *Reasoned Administration and Democratic Legitimacy*. Cambridge University Press. https://doi.org/10.1017/9781108355827.

McCubbins, Matthew D., Roger G. Noll, and Barry R. Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics, and Organization* 3 (2): 243–77. https://doi.org/10.2307/764829.

McGarity, Thomas O. 1992. "Some Thoughts on "Deossifying" the Rulemaking Process." *Duke Law Journal* 41 (6): 1385–1462.

McNollgast. 1999. "The Political Origins of the Administrative Procedure Act." *Journal of Law, Economics, & Organization* 15 (1): 180–217. https://doi.org/10.1093/jleo/15.1.180.

Milkman, Katherine L., Mary Carol Mazza, Lisa L. Shu, Chia Jung Tsay, and Max H. Bazerman. 2012. "Policy bundling to overcome loss aversion: A method for improving legislative outcomes." *Organizational Behavior and Human Decision Processes* 117 (1). Elsevier Inc.: 158–67. https://doi.org/10.1016/j.obhdp.2011.07.001.

Morrison, Alan B. 1986. "OMB interference with agency rulemaking: the wrong way to write a regulation." *Harvard Law Review* 99 (5). JSTOR: 1059–74.

Moynihan, D., P. Herd, and H. Harvey. 2015. "Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen-State Interactions." *Journal of Public Administration Research and Theory* 25 (1): 43–69. https://doi.org/10.1093/jopart/muu009.

Moynihan, Donald 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. https://doi.org/10.1353/jph.2004.0005.

Neustadt, Richard E. 1991. *Presidential power and the modern presidents: The politics of leadership from Roosevelt to Reagan*. Simon; Schuster.

Persad, Govind. 2014. "When, and how, should cognitive bias matter to law." *Law & Ineq.* 32. HeinOnline: 31.

Postell, Joseph. 2017. *Bureaucracy in America: The administrative state's challenge to constitutional government*. University of Missouri Press.

Potter, Rachel A. 2019. *Bending the Rules: Procedural Politicking in the Bureaucracy*. Chicago: University of Chicago Press.

Rachlinski, Jeffrey J., and Cynthia R. Farina. 2002. "Cognitive Psychology and Optimal Government Design." *Cornell Law Review* 87: 549.

Redlawsk, David P., ed. 2006. *Feeling Politics*. Palgrave Macmillan.

Regenwetter, Michel, Jason Dana, and Clintin P. Davis-Stober. 2011. "Transitivity of Preferences." *Psychological Review* 118 (1): 42–56. https://doi.org/10.1037/a0021150.

Shapiro, Martin. 1986. "APA: Past, Present, Future." *Virgina Law Review* 72: 447.

Shapiro, Stuart. 2011. "The evolution of cost–benefit analysis in US regulatory decisionmaking." In *Handbook on the Politics of Regulation*, edited by David Levi-Faur. Edward Elgar Publishing.

Shepsle, Kenneth A. 1979. "Institutional Arrangements and Equilibrium in Multidimensional Voting Models." *American Journal of Political Science* 23 (1): 27. https://doi.org/10.2307/2110770.

Simon, Herbert A. 1976. "From substantive to procedural rationality." *25 Years of Economic Theory*, 65–86. https://doi.org/10.1007/978-1-4613-4367-7_6.

———. 1997. *Models of Bounded Rationality*. MIT Press.

Skowronek, Stephen. 1982. "Building a New American State." Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511665080.

STAROŇOVÁ, KATARÍNA. 2010. "Regulatory Impact Assessment: Formal Institutionalization and Practice." *Journal of Public Policy* 30 (1): 117–36. https://doi.org/10.1017/S0143814X09990201.

Stephenson, Matthew C. 2006. "A Costly Signaling Theory of Hard Look Judicial Review." *Administrative Law Review* 58 (4): 753–813.

Sunstein, Cass R. 2018. *The cost-benefit revolution*. MIT Press.

Sunstein, Cass R. 2000. "Cognition and Cost-Benefit Analysis." *The Journal of Legal Studies* 29 (S2): 1059–1103. https://doi.org/10.1086/468105.

———. 2002. "The cost-benefit state: the future of regulatory protection." In. American Bar Association.

Theriault, Sean M. 2003. "Patronage, the Pendleton Act, and the Power of the People." *The Journal of Politics* 65 (1): 50–68. https://doi.org/10.1111/1468-2508.t01-1-00003.

Tversky, Amos. 1969. "Intransitivity of preferences." *Psychological Review* 76 (1): 31–48. https://doi.org/10.1037/h0026750.

Tversky, Amos, and Daniel Kahneman. 1981. "The framing of decisions and the psychology of choice." *Science* 211 (4481). https://doi.org/10.1126/science.7455683.

Van Riper, Paul P. 1983. "The American administrative state: Wilson and the founders-An unorthodox view." *Public Administration Review*. JSTOR, 477–90.

Weber, Max. 1946. *From Max Weber: Essays in Sociology*. Edited by H.H. Gerth and C. Wright Mills. New York, NY: Oxford University Press.

Weingast, Barry R., and William J. Marshall. 1988. "The Industrial Organization of Congress: or, Why Legislatures, Like Firms, Are Not Organized as Markets." *Journal of Political Economy* 96 (1): 132–63.

Wilson, Woodrow. 1887. "The Study of Administration." *Political Science Quarterly* 2 (2): 197–222.

Yackee, Susan Webb. 2019. "The Politics of Rulemaking." *Annual Review of Political Science* 22. https://doi.org/10.1146/annurev-polisci-050817-092302.