WHAT BIASED RULEMAKING LOOKS LIKE

BRIAN LIBGOBER

Abstract. Each year, administrative agencies issue thousands of rules with substantial social and economic ramifications. Yet, many observers remain wary of administrative policymaking. One significant reason is that rulemaking has a "democracy deficit." Public participation during rulemaking is sporadic and unrepresentative. But what do high or low levels of public participation during rulemaking tell us about bias in regulators or regulation? This paper explores this question by modeling rulemaking as a persuasion game. The model shows that known patterns of public participation in rulemaking are inconsistent with total capture or absolute zealotry. Such patterns are consistent, however, with almost any other level of bias in the policy or the policymaker's decision-making, provided one fails to account for variability in outside actors' stakes in regulatory outcomes.

Date: 2/12/18.

Critics of adminsistrative policymaking are numerous and vociferous. Although they may blame the regulatory state for many ills, the crux of their argument frequently relates to representation. "Administrative legislation is unrepresentative," writes one author, "It therefore is without consent, without obligation, and without popular accountability. Indeed, it is a form of class power" (Hamburger 2014). Yet, one could make strong arguments that policymaking in *legislatures* is unrepresentative because of gerrymandering, disenfranchisement, lobbying, campaign finance issues, or other representational impediments. Further, some would argue that agency policymaking is *radically* democratic. Many contemporary national and subnational systems use "notice-and-comment" rulemaking procedures (Jensen and McGrath 2011). In this policymaking scheme, regulators must submit proposed regulations for public feedback. If the agency does not appropriately consider these comments, the submitter can sue and have a court invalidate the rule (see, e.g. Home Box Office, Inc v. FCC, 1977). Notice-and-comment procedures create an open-forum, whose openness is enforced through judicial sanction.

The quality of democratic representation and responsiveness in notice-and-comment procedures is an important topic both theoretically and practically, which a large and growing empirical literature addresses (e.g., Krawiec 2013; Michael 2013; Balla and Daniels 2007; Figueiredo 2006; Cuellar 2005; Golden 1998). Although this scholarship has produced profound insights into regulatory processes, it has also grappled with thorny theoretical issues clouding inference. Many related areas of empirical research suffer similar issues, and one often hears calls for additional formal theoretic work to address such difficulties (Hojnacki et al. 2012). Formal theorists have long worked on topics about or related to rulemaking (McCubbins and Schwartz 1984; McCubbins, Noll, and Weingast 1987; Epstein and O'Halloran 1999). In recent years, a number of notable formal-theoretic papers have addressed thematically similar questions about bureaucratic efficacy and legitimacy (Gailmard and Patty 2016; Bueno De Mesquita and Stephenson 2007; Stephenson 2007). Yet, none of these models has focused on the outcomes empirical scholars regularly report data about: public-participation rates and patterns of policy change. This is unfortunate because the data that empirical

scholars produce obviously relates to issues of representation and bias, although exactly how is often more complex than one might think.

Scholarly findings about public-participation rates illustrate the interpretive difficulties that motivate my model of rulemaking. The literature finds that most regulations receive no public comments and a few receive an overwhelming number of public comments. In more "typical" rulemakings that receive between one and several thousand comments, business interests are the most active participants. Arguably, these findings establish inequality in process and suggest inequality in outcomes. Yet, there are many reasons for doubt. A lack of participation may mean the proposal is uncontroversial or represented a favorable compromise to all interests. Similarly, a surfeit of participation may indicate public outrage; panic that agencies will ignore the preferences of many diverse interests; or that other nefarious dealings, such as "docket spamming," are in the works (see, for example, Hitlin, Omstead, and Toor 2017, discussing the FCC's repeal of net neutrality).

One could easily make similar informal arguments in regard to observations about how rules change from proposal to final stage. These alternative possibilities are sufficiently concrete that one could, in theory, investigate their merits. Other worries are more amorphous and do not lead to specific predictions. Rather, they cast a shadow of gloom over the entire research enterprise. For example, many would assume that actors that participate in notice-and-comment are strategic. As a result, the regulator and interested outside actors understand each other's incentives and respond accordingly. Possibly, they also respond to the strategic incentives of actors, such as Congress or the courts, who may only occassionally participate. The entire system is endogenous with respect to itself. Therefore, it is hard to know what to make of data that emerges from the process. Doubts like these are "productive" in the sense that they prevent scholars from overinterpreting their data. Yet, they are also deeply unproductive because they deter valuable data collection, provide limited resources for defining scope conditions, and suggest no robust or interesting research strategies for learning more about the quality and nature of representation during rulemaking.

This paper clarifies the interpretation of empirical data on rulemaking by modeling the notice-and-comment process as a two-sided persuasion game played between a regulator and interested outside actors. The model's framework distinguishes between three potential pitfalls in regulatory policymaking: *allocational bias*, *inferential error*, and *inferential bias*. "Allocational bias" relates to the weight the regulator assigns the preferences of outside actors. For example, a captured regulator might only care about the preferences of a single firm or industry, whereas a bureaucratic zealot might give no weight to the preferences of either concerned firms or citizens. "Inferential error" describes the difference between the regulator's assumptions about the preferences it cannot observe and what those preferences truly are. If inferential error is systematic, then it is "inferential bias." For example, a regulator might ignore participation incentives and assume that outside actors that comment have similar preferences to those that do not comment. If the regulator did that, it could lead to inferential bias. Inferential bias can increase the magnitude of inferential errors, but even an unbiased regulator makes inferential errors because it can only approximate preferences it does not observe. The less predictable the preferences of the abstaining outside actors, the bigger the expected inferential error, regardless of whether a regulator is inferentially biased.

The model informs us that two crucial levers affect outside actors' incentives to participate in the notice-and-comment process: inferential error and allocational bias. Each lever is also critically important in determining the resulting policy. As the empirical literature on rulemaking assumes, data about the level of participation in rulemaking *does* reveal information about regulatory bias under the model. For example, positive levels of public participation from multiple sectors is inconsistent with both absolute zealotry and single-sector capture, no matter how systematically mistaken the regulator is about outside actors' preferences. The reason is straightforward. If the regulator does not care about the preferences of outside actors, it would not respond to information it receives about those preferences. Therefore, outside actors would have no incentive to provide such information. Because scholars typically find some participation from many sectors during rulemaking, the model should leave us modestly encouraged.

However, stronger interpretations of known patterns are difficult to reach. If producers submit more comments than consumers, is that evidence that regulators favor producers over consumers? Not necessarily, says the model. Participation during rulemaking needs two ingredients: (1) sufficient concern for a potential participant's preferences and (2) sufficient anticipated error about those preferences. Provided that both ingredients are present, one cannot determine the ratio of these two substitutes purely from outputs. Producers might comment more than consumers because of regulatory favoritism *or* because its more difficult to predict the effect of regulation on producers. Similarly, if rules systematically change in a way that favors businesses over other groups (Yackee and Yackee 2006), that could occur because allocational bias favors business *or* because the regulator's inferential error with respect to other groups is smaller. If one knew that inferential error for business interests were smaller, then that would establish allocative bias toward business, at least under the model. In place of *generic* concerns about observational equivalency, the model develops a framework for understanding *specific* pathologies and provides scope conditions under which these pathologies would not cloud inference.

Although the model addresses rulemaking, it literally describes a participatory system in which a decision-maker proposes a policy, responds to costly signals from outside actors about their preferences, and then makes policy to respond to those preferences. Other policy fora besides rulemaking, including participatory budgeting, school board meetings, and party platform committees, arguably fit this mold. The bias of systems for participatory policymaking is a topic of perenial interest to scholars. Researchers in these areas may have their own concerns about how strategic interaction confounds their inferences. The framework this model develops may well prove useful.

Rulemaking: Examples, Processes, and Patterns

Rules, like laws or treaties, are one kind of official document that governments issue. For lawyers, the manner in which these documents relate to one another is a topic of great interest and – occassionally – confusion. For individuals and firms, the most pressing question

is, "where does it say what I am supposed to do?" Both laws and rules tell members of the public what to do. In most cases, these dictates are complementary. The law describes a requirement or a policy; the regulation clarifies or implements it. Conflicts nevertheless do occur. Theoretically, in cases of disagreement, the law triumphs. Practically speaking, policy-minded scholars should recognize that getting the law to triumph is easier said than done. Agencies zealously defend the rules they write in court. Litigating against a regulator is time-consuming and costly, and the standards of judicial review stack the deck in the agency's favor. Almost none of the individuals that theoretically might sue an agency will actually go on to seek judicial review (see, e.g., Stoll 2010, especially pp. 90-91).

Rules have consequences for virtually every policy domain. Indeed, rules often address the core policy trade-offs in a much deeper and more practical way than the laws that authorize them. For example, no provision of the U.S. Code specifically instructs the Consufer Financial Protection Bureau (CFPB) to regulate payday loans. Instead, section 1031 of the Dodd-Frank Act asks the CFPB to make rules preventing "unfair" or "abusive" consumer-finance practices. The 2-page section describes how the agency should determine which practices are unfair or abusive. The payday-loan rule is, by contrast, a 1,600-page behemoth. At a high level, the rule represents a balancing act (Cordray 2017). On the one hand, the rule aims to discourage lenders from leading borrowers into viscious debt spirals. On the other hand, it tries to preserve consumer access to this market. Without access to payday loans, desperate consumers may switch to potentially worse sources of finance, such as pawn-brokers, loan-sharks, or high-interest credit cards (Bhutta, Goldin, and Homonoff 2016). The regulation forces lenders to conduct a "full-payment test" before issuing a loan. The test ensures that borrowers will not need another high-interest loan to pay back the first loan. The rule describes what the full-payment test must include, which loans are subject to it, which loans are exempt, and various provisions relating to recordkeeping and disclosures (Kully and Robinson 2017). The payday loan is just *one* example of the regulations that section 1031 makes possible.

There is a large scholarly literature exploring various aspects of the notice-and-comment" process by which most rules are produced.[1] Recent empirical work on rulemaking in political science focuses largely on two crucial questions relating to public influence over policymaking: "Who participates in rulemaking" and "How do rules changes from proposal to final form?" This literature suggests several stylized facts:

- Defining the typical rule or comment is difficult, since there is so much variation in the number of comments that rules receive (Golden 1998; West 2004; Yackee 2005; Balla and Daniels 2007). An analysis of data from regulations.gov suggests that most rules receive no comments and most comments were submitted on a small number of rules (see Appendix, Figure 1).
- Participation during notice-and-comment is pluralistic, in the sense that firms from many sectors, associations, and even ordinary citizens at times participate (Golden 1998; Balla 1998; Cuellar 2005; Furlong and Kerwin 2005; Yackee and Yackee 2006).
- Business interests participate more consistently in rulemaking, although even firms from within a given sector sometimes oppose one another (Golden 1998; West 2004; Yackee and Yackee 2006).
- Some policy change almost always occurs between a proposed rule and a final rule. Most of these change are not very substantial, but they are occassionally significant. (Golden 1998; West 2004; Yackee and Yackee 2006). For example, the CFPB amended the payday loan rule to also regulate "balloon loans."
- The position that business interests advocate predicts the direction of policy change from proposed rule to final rule, but the same is not true for other kinds of commenters (Yackee and Yackee 2006; Haeder and Yackee 2015).

Several formal theoretic papers have explored rulemaking, but none have sought to make sense of the theoretical significance of these particular empirical patterns. Rather, existing work on rulemaking emerges from a larger literature on the role of information in bureaucratic politics (Gailmard and Patty 2012). The starting point for this literature is the puzzle

[1]Excellent overviews are (Kerwin and Furlong 2011) and (Lubbers 2012).

of delegation. Why would Congress grant an agent like the CFPB so much discretion to enact policies like the payday loan rule? The answer, some have argued, has to do with informational asymmetries (e.g., Epstein and O'Halloran 1999; Bawn 1995). Regulatory agencies have informational advantages that helps them respond to the fundamental uncertainties of policymaking. If policymaking is uncertain, then Congress can do better by delegating authority than it can acting on its own, even if Congress and the agency are not that closely aligned (Bendor and Meirowitz 2004). This answer begs the question, however, of when and why agencies possesss the information that makes delegation advantageous. Models exploring this second question come in two flavors. One kind of model focuses on the circumstances affecting agency investment in bureaucratic capacity (e.g., McCarty 2017; Gailmard and Patty 2007). In these models, the ideal policy is uncertain, but costly efforts at acquiring expertise can decrease or eliminate that uncertainty. The willingness of the regulator to invest in expertise depends on salient features of the political environment.

The second kind of model is premised on the idea that agencies invest not only in expertise, but also in procedures for receiving and channeling information from the public. These procedures present the regulator with a "learning problem" (Carpenter 2004). How does the regulator transform a mix of present and absents signals into beliefs about the underlying state of the world? The manner in which policymakers resolve learning problems can have wide-ranging political and policy implications (Callander 2011; Volden, Ting, and Carpenter 2008; Carpenter and Ting 2007).

One can describe the notice-and-comment process as either an investment problem or a learning problem. Recent papers have taken the investment approach (Gailmard and Patty 2016; Bueno De Mesquita and Stephenson 2007; Stephenson 2007). Yet, to the author's knowledge none has tried casting rulemaking as learning problem, which is surprising. The authors of the APA explicitly argued that the purpose of the notice-and-comment process was to create a space for learning that could *not* come about by one expert actor's studying an issue on its own.

> The reason for [an administrative agency's] existence is that it is expected to bring to its task greater familiarity with the subject than legislators, dealing with many subjects, can have. But its knowledge is rarely complete, and it must always learn the frequently clashing viewpoints of whom its regulations will affect.
>
> These differences are and should be reflected in its procedures, which should be adapted to giving adequate opportunity to all persons affected to present their views, the facts within their knowledge, and the dangers and benefits of alternative courses. They should also be adapted to eliciting, far more systematically and specifically than a legislature can achieve, the information, facts, and probabilities which are necessary to fair and intelligent action. (Attorney General's Committee on Administrative Procedure 1941, p.101-102)

The report writers seem to have envisioned rulemaking as a forum for agencies to learn about the consequences of various policy choices on interested parties. The agency would then engineer a compromise between what it wanted to do on the basis of its own expert opinion and what the public would want. In other words, the rulemaker should aggregate preferences. When political scientists describe preference aggregation, they typically refer to algorithms for choosing from rank orders over alternatives. But rulemaking seems to involve a form of preference aggregation that is fundamentally unlike voting in elections. In rulemaking, a rational actor rather than an algorithm combines preferences, in response to public signals it may or may not observe, from members of the public it may or may not care about.[2]

[2]A further complication is that the signals that the public sends about its preferences may be dishonest or exaggerated. For reasons discussed in the following section, this paper adopts the persuasion-game framework of (Milgrom and Roberts 1986; Milgrom 2008), which makes it difficult to address dishonesty or exaggeration. Deviations from perfect credibility may result in the regulator's discounting the value of information it observes and preferring its prior beliefs. Thus, different degrees of credibility in public signals can be partially described in the model by modulating the degree of concern for public signals.

Although this formal model receives its inspiration from an informal account that predates the APA, it is consistent with each of the stylized facts mentioned above. Perhaps more importantly, the model serves as a useful tool for thinking about rulemaking and the data that it generates. It shows that patterns of participation and policy change can provide misleading signals about the underlying bias of the regulator or the regulation.

Rulemaking as Learning and Adjusting to Public Preferences

The Model. This game has two actors: a regulatory agency G and a member of the public i. One might think of i as an individual firm or interest group, or even as a private citizen. The technical Appendix extends the analysis to the case where the public has n individual members, however the dynamics this article emphasises are are visible even if the public has a single member. Actors in this model have utility functions that are dependent on (1) which policy is implemented and (2) actions during rulemaking. I first discuss how policy utility is modeled and then describe how actions taken during rulemaking affect utility.

I understand policy utility for the regulator and the public using a single-dimensional spatial model. Member of the public i has negative quadratic preferences centered around an ideal policy p_i, which is the public's type. Formally, $v_i(x) = -(x - p_i)^2$. Whereas p_i is known to i, the regulator knows only the distribution from which p_i is drawn. In order to obtain closed-form solutions, I assume that the distribution of types is uniform, in particular $p_i \sim \mathcal{U}(m_i - r_i, m_i + r_i)$.[3] This parameterization implies that prior to notice and comment, the expectation is that the public will most prefer policy m_i.

I model the policymaker as an imperfectly benevolent social planner, with an emphasis on the word "imperfectly." In particular, the Bernoulli utility function of G is given by

$$u_G(x_f) = -(1 - \alpha)x_f^2 + \alpha v_i(x_f)$$

[3]As we shall see, the model developed here requires forming the posterior from partially censored data. While closed form solutions may still be obtained if types are uniformly distributed, for most other distributions posterior inference from censored data leads to irreducible integrals.

Here, α is a variable controlling the regulator's allocational bias, or the weight attributed to the public welfare. It is comparable with the taste for political contributions parameter a in Grossman and Helpman (1994). I view bureaucrats as potentially policy-minded. If $\alpha = 1$, then G has no preferences over policy independent of what the public wants. Indeed, its utility function $u_G(x) = v_i(x)$ would be identical with the public's indirect utility over policy. Even if there is no preference conflict between the public and the agency, there is nevertheless an informational problem because G does not know v_i. Unlike the Grossman-Helpman model, no auction-style mechanism exists for members of the public to reveal the location and intensity of their preferences. Instead, the notice-and-comment process may offer G an opportunity to learn p_i provided i is willing to incur some signaling costs. If p_i remains unknown, then G would select policy only on the basis of its prior beliefs. Hence, the utility-function representations of the public and the agency already highlights the presence of both inferential and allocative problems in crafting regulation. Further description of the inferential aspects of the regulator's decision problem will follow some additional discussion of allocative bias.

Although early political science scholarship sometimes described agencies as benevolent social planners, and might therefore be willing to assume $\alpha = 1$, the contemporary bureaucratic politics literature supports the notion that agencies have policy concerns independent of public welfare. Theoretically, these concerns might relate to administrability and budget (Niskanen 1968), organizational mission (Wilson 1989), bureaucratic reputation (Carpenter 2010), or political ideology (Gailmard and Patty 2007). Indeed, recent empirical work has provided strong evidence that agencies have ideological biases that are likely related to their policymaking choices (Richardson, Clinton, and Lewis 2017; Chen and Johnson 2015; Clinton et al. 2012). In order to illustrate how the modeling technology might describe a non-welfarian social planner, I consider the extreme case where $\alpha = 0$. Following Gailmard and Patty (2007), I describe such a bureaucrat as a "zealot." If $\alpha = 0$, then $u_G(x) = -x^2$, so the regulator has a utility function quadratically decreasing from 0. It is important to emphasize that one can select 0 without loss of generality, and, in principle, 0

could reflect any substantive policy choice. If the ideal policy of the agency were $p_G \neq 0$, one could simply shift other quantities in the model horizontally by p_G and derive formally identical conclusions. However, using the regulator's independent preferences as a reference point does have important interpretive consequences. If $p_i > 0$ then it means that i's ideal policy is on one side of the regulator in this policy space, whereas if $p_i < 0$, it means that i's ideal policy is on the other side.

Any account of allocative bias in rulemaking should be able to describe some notion of regulatory capture. Different scholars might use the term "capture" to refer to different phenomena (Carpenter and Moss 2013). In this model, capture refers to governmental favoritism affecting the distribution of policy benefits (Stigler 1971). To give some idea of how capture is formalized, suppose that the public had two members. In this case, the regulator's utility function would be given as

$$u_G(x) = -\alpha_0 x^2 + \alpha_1 v_1(x) + \alpha_2 v_2(x)$$

The allocative bias parameter α is now a vector of weights corresponding to each component of the regulator's welfare function. If $\alpha_1 = 1$, then G cares only about the utility of the first member of the public: it gives no weight to either its independent preference or the preferences of the second member of the public. Hence, the regulator is "completely captured." Larger α_i entails greater policy favoritism for a given part of the public.

Having specified utility over policies, I now turn to describing the strategic interaction driving this model of rulemaking. At the beginning of the game, regulator G proposes a policy, x_p. The public observes the proposal and must choose to send a comment by selecting $a \in \{0, 1\}$.[4] If $a = 1$, then regulator observes p_i, but the member of the public incurs cost c. Alternatively, i can sends no signal and incur no cost. As in other persuasion

[4]In the Appendix, each member of the public makes this decision independently and without any ability to communicate with other potential participants. Coalitional commenting is an important part of the politics of rulemaking (Nelson and Yackee 2012), although how coalitions form and why they are effective remains poorly understood. While not strictly necessary to establish this article point, future modeling should try to allow for a degree of coordination between potential rulemaking participants.

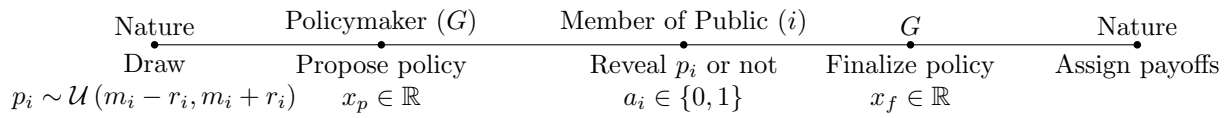

FIGURE 1. Order of play during the rulemaking game

games, the information-sender cannot provide the receiver false or dishonest signals. After observing the public signals, G must select a final policy x_f. Amending policy is costly in this model. Moving from policy x_p to x_f imposes costs $k(x_f - x_p)^2$ on the regulator. Figure 1 shows the order of play in this rulemaking game.

Persuasion games describe situations where a decision-maker "need[s] to rely on suggestions and information provided by individuals who are affected by their decisions" (Milgrom and Roberts 1986). The APA authors portray the notice-and-comment process in such terms. The core econoimc assumption in such models is that communication must be honest, but need not be given. From a rational-choice perspective, honesty is justified if information is easily verified or if the cost for reporting false information are sufficiently large. Frequently in rulemaking, both are true. Regulators have independent investigatory capacity they can use to double-check the information that firms and interest groups provide. Firms and interest groups do not wish to be caught lying to their regulator because doing so would hurt their reputations. Those who write comment letters within firms frequently consider future employment with the regulator, or have social connections with employees of the regulator. Career-concerns may also raise the costs of attaching one's name to a misleading comment letter. In cases where verification and reputation do not constrian actors, I would expect the regulator to *discount* the value of communication these actors provide and *increase* the weight of their prior. Modulating the α parameter can therefore capture some, if perhaps not all, of the effects of exaggerated or dishonest communication.

I assume that modifying rules is costly to the regulator. While analysis of agency's reptuational and managerial incentives can support this assumption, I wish to emphasize the strategic incentives a regulator faces in surviving judicial review. "If the agency is to state the detailed basis for its actions in such a way that its actions will survive judicial review," wrote

one EPA general counsel, "public input through formal notice-and-comment rulemaking must come relatively close to the end of the agency's process, when the proposed rule has 'jelled' into something fairly close to its final form." (Elliott 1992) The author cites three examples in which courts forced agencies to repeat the notice-and-comment process becasue they made changes that were too "fundamental." These sentiments are substantially echoed by Barron and Kagan (2001); they write that "the more work agencies put into their proposals, the less flexibility they show during rulemaking to respond to the concerns of affected parties." Indeed, some authors go so far as to argue that aggressive judicial oversight has gone so far as to "ossify" rulemaking, vitiating the policymaking forum of its intended dynamic character (McGarity 1992). Although quantitative evidence suggests ossification is overstated (Yackee and Yackee 2010), it is nonetheless reasonable to assume that agencies would rather not make changes to proposed rules and that costs to changes are increasing at an increasing rate. The quadratic form $k(x_f - x_p)^2$ captures these dynamics and has a free parameter $k \in \mathbb{R}^+$ that controls the importance of these costs relative to the other parameters in the model.

As a dynamic game of incomplete information, it is appropriate to use the Perfect Bayesian Equilibrium (PBE) solution concept. A PBE is a set of strategies and beliefs such that, during any stage in the game, "strategies are optimal given beliefs, and beliefs are obtained from equilibriums strategies and observed action using Bayes rule" (Fudenberg and Tirole 1991). The following proposition summarizes the conditions for a pure-strategy equilibrium. Detailed proofs that these conditions are necssary for equilibrium, and also that these conditions can always be satisfied, are left for the Appendix.

Proposition 1. *A PBE of the rulemaking game consists of a policy* $x_p^* \in \mathbb{R}$*, a commenting decision rule* $a^* : [m_i - r_i, m_i + r_i] \to \{0, 1\}$*, an imputed type* $z_i^* \in [m_i - r_i, m_i + r_i]$ *and a best response function* $x_f^* : \mathbb{R} \times [m_i - r_i, m_i + r_i] \times \{0, 1\}$ *satisfying the following four conditions.*

(1) $x_p^* = \alpha m_i + \left(\frac{\alpha^2}{2k}\right)\left(\frac{\partial}{\partial x_p} Cov(p_i, \mathbb{E}(p_i \mid a))\right)$

$$(2)\ a_i^*(p_i) = \begin{cases} 1 & \left(\frac{k}{k+1}x_p + \frac{\alpha}{k+1}z_i^* - p_i\right)^2 - \left(\frac{k}{k+1}x_p + \frac{\alpha p_i}{k+1} - p_i\right)^2 > c \\ 0 & \textit{otherwise} \end{cases}$$

$$(3)\ z_i^* = \mathbb{E}\left[p_i \mid a = 0\right]$$

$$(4)\ x_f^*(x_p, p_i, a_i) = \tfrac{k}{k+1}x_p^* + \tfrac{\alpha}{k+1}\mathbb{E}\left(p_i \mid a^*\right)$$

To unpack this proposition, I proceed in reverse order, following the logic of backward induction. Condition (4) states that the regulator enacts final policy according to the following formula.

$$x_f^*(x_p, p_i, a_i) = \frac{k}{k+1}x_p^* + \frac{\alpha}{k+1}\mathbb{E}\left(p_i \mid a_i^*\right)$$

The final proposal is a weighted average of the policy the regulator initially chose and its posterior expectation of i's type. The relative weight of the initial proposal x_p is a function of k, the costliness of amending policy. The suggested relationship is intuitive: the more expensive it is for the regulator to change the policy, the more similar the proposal and the final policy will be. In the limit, $\lim_{k\to\infty} x_f^* = x_p^*$. By contrast, the less expensive it is to change policy, the less the initial proposal matters for the regulator's final choice. Formally, if amending policy is costless, the term x_p indiciating the location of initial proposal drops:

$$\lim_{k\to 0} x_f^* = \alpha E\left(p_i \mid a_i^*\right)$$

This expression makes it easier to see that my initial description of final policy was a simplification. It is not the posterior expectation *per se* that matters for final policy, but rather the posterior expectation after adjusting downward for the regulator's allocative bias α. Continuing under the assumption that amending policy is costless, if $\alpha = 0$ then the regulator attributes no weight to public preferences and so its posterior about i's type is irrelevant to final policy: $x_f^* = 0$. If $\alpha = 1$, then final policy is identical to the regulator's posterior expectation of the public's type. For levels of allocational bias between 0 and 1, the regulator selects a final policy that is a convex combination between the posterior expectation and its independently preferred policy of 0.

The posterior expectation $E\left(p_i \mid a_i^*\right)$ makes the notice-and-comment process a "learning problem." If the member of the public i comments ($a_i = 1$), then the regulator learns i's type immediately and with certainty: $\mathbb{E}\left(p_i \mid a_i = 1\right) = p_i$. If i does not comment, it might seem that the regulator can learn nothing about the public's type. If abstention were uninformative to the regulator, then the posterior expectation is equal to the prior expectation: $\mathbb{E}\left(p_i \mid a_i = 0\right) = \mathbb{E}(p_i) = m_i$. Yet a fully rational regulator will understand that strategic-interaction may create selection issues. Commenters and abstainers may differ in ways that relate to the policy outcomes they anticipate. If i abstains, it means that i anticipated gains of participating that were too small to justify the costs. Because the benefits of commenting depend on i's type, abstention provides information about that type. If the regulator wants to make the right policy, it must reverse-engineer the public's selection bias to understand which types had sufficient incentive to comment.

To understand how selection-bias works in the model, consider i's decision problem. The public knows that the bureaucrat will enact policy according to Condition (4). After abstention, i does not necessarily know the regulator's expectations regarding p_i. Yet the public knows that if it abstains, the regulator will expect *something*. Call this arbitrary estimate z_i. If the public knows – or thinks it knows it – z_i, the decision calculus for individual i then becomes simple. "If I comment, the regulator assumes that my type is p_i, which leads to policy x. If I abstain, the regulator will assume that my type is z_i, which leads to final policy x'. Therefore, I should comment if and only if the difference in expected utility between x and x' exceeds the commenting cost." Formally, it is possible to formulate the best response of the public *regardless* of what we assume about how the regulator makes inferences:

$$a_i^*(p_i) = \begin{cases} 1 & \left(\frac{k}{k+1}x_p + \frac{\alpha}{k+1}z_i^* - p_i\right)^2 - \left(\frac{k}{k+1}x_p + \frac{\alpha p_i}{k+1} - p_i\right)^2 > c \\ 0 & \text{otherwise} \end{cases}$$

The above expression is the same as Condition (3) of the main proposition. The lefthand side of the inequality is the policy benefit of commenting, the right hand side is the commenting cost.

These remarks show that the regulator's expectations about the public depend on the public's incentives. Yet the public's incentives depend on the regulator's expectations. Escaping this circularity requires assumptions about bureaucratic rationality – for example, that agencies ignore selection bias or assume the worst. Here, I use the Perfect Bayesian Equilibrium concept, which requires the regulator's beliefs and the public's incentives to reinforce each other. In particular, we have condition (2)

$$z_i^* = \mathbb{E}\left[p_i \mid a = 0\right]$$

As expansion of Condition (3) confirms, the public's marginal benefit of participation depends on the inferential error it anticipates: $p_i - z_i$. When inferential error is small, the difference in policy outcomes following participation must also be small. If the difference is small enough, participation cannot be worth the costs. Conversely, if $p_i - z_i$ is larger, the marginal difference in outcomes is also larger. If $p_i - z_i$ is large enough, the marginal differences in outcomes must be positive. The informal argument is that, as long as the regulator cares an iota about the public's preferences, the final policy accommodates public preferences as the regulator understands them. The further p_i is from 0, the more G is willing to concede following commenting. As p_i approaches ∞, the size of the concession that follows participation is arbitrarily large. Yet z_i fixes the concession that follows abstention. Hence, as p_i approaches ∞, the policy concession that follows commenting becomes marginally more attractive, whereas the concession that follows abstaining becomes increasingly inadequate. The argument is analogous for $p_i \to -\infty$. Figure 2 provides a graphical illustration.

Before the public participates, the regulator proposes a policy x_p to create the most favorable circumstances for issuing its final policy. The regulator has two competing considerations. On the one hand, it wants to avoid incurring future modification costs. Therefore, the regulator is inclined to propose the policy it would finalize if it did not have access to the

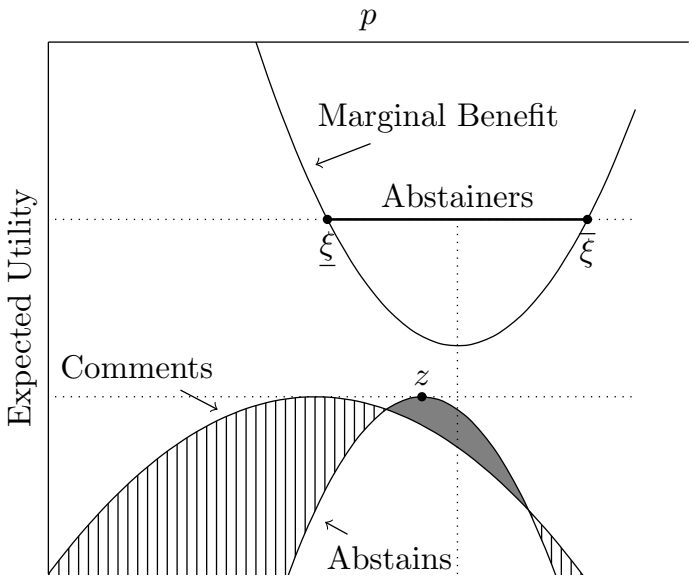

FIGURE 2. i's decision problem portrayed for arbitrary z. Along the x-axis are types p_i that i might possess. The y-axis shows expected utility following each action and also the marginal benefit of commenting. Vertical lines indicate values of p_i where there are policy gains from commenting, though perhaps too small to justify the participation costs. The shaded area shows area where i would experience policy losses. The horizontal line through $\underline{\xi}, \overline{\xi}$ indicates the cost of commenting c. If $p_i \in (\underline{\xi}, \overline{\xi})$ then i abstains, otherwise i comments.

notice-and-comment process: αm_i. On the other hand, the regulator knows that final policy will reflect to some degree its own expectations about the public's type. These expectations depend on the public's participation decision. The regulator will threrefore wish to propose policy so as to minimize the inferential error implicit in final policy. If the public comments, then the regulator becomes perfectly informed, so no inferential error occurs. If the public abstains, some error occurs, but the regulator wants that error to be small. Condition (1) shows that the initial policy responds to both the desire to avoid modification costs and the desire to get the most informative mix of present and absent signals:

$$x_p^* = \alpha m_i + \left(\frac{\alpha^2}{2k}\right)\left(\frac{\partial}{\partial x_p}\text{Cov}(p_i, \mathbb{E}(p_i \mid a))\right)$$

Note that this condition also shows that a foresighted regulator may *strategically position* itself as to encourage more commenting and more informative abstention.[5]

One important wrinkle is that this model has boundary constraints. Recall that the justifiability of participation depends on the inferential error $p_i - z_i$ being sufficiently grave. The model assumes, however, that types p_i are distributed uniformly on $(m_i - r_i, m_i + r_i)$. This assumption creates the possibility that one or both of the distributional constraints would be "active." It may be impossible for i to have a type so far below $m_i - r_i$ that the regulator would make a sufficiently large inferential error about i's preferences. The boundary constraints can determine how the interval $(m_i - r_i, m_i + r_i)$ is divided between participants and non-participants (Figure 3). There are four possibilities: everyone abstains (two active constraints), one side of the distribution comments (one active constraint), the other side of the distribution comments (the other active constraint), or abstainers have preferences between the commenters (unconstrained). Although all these possibilities give the model flexibility to explain diferent patterns of participation, they also create presentational difficulties. The expectation in Condition (3) has a different formula in each case (see Appendix). The remainder of the paper mainly focuses on the unconstrained case, in which the closed-form solutions are the simplest.

An important takeaway from the last paragraph is that uncertainty about preferences and participation are opposed. If r_i is sufficiently big, the distributional constraints will not bind and participation is possible. Indeed, as we shall emphasize in the next section, the larger r_i is, the more probable participation is. Conversly, if r_i is too small, both constraints bind and abstention is inevitable. The fact that r_i can strongly influence patterrns of participation and responsiveness is what leads to many interpetive pathologies, and explains why it is undesirable to leave public preferences as an unmeasured latent variable.

[5]Strategic positioning is possible but not inevitable under the model. The reason is that the covariance term can be flat as a function of x_p in the neighrborhood of αm_i, in which case $x_p = \alpha m_i$. Strategic positioning is only observed when it can shrink the abstention interval (see Figures 2 and 3).

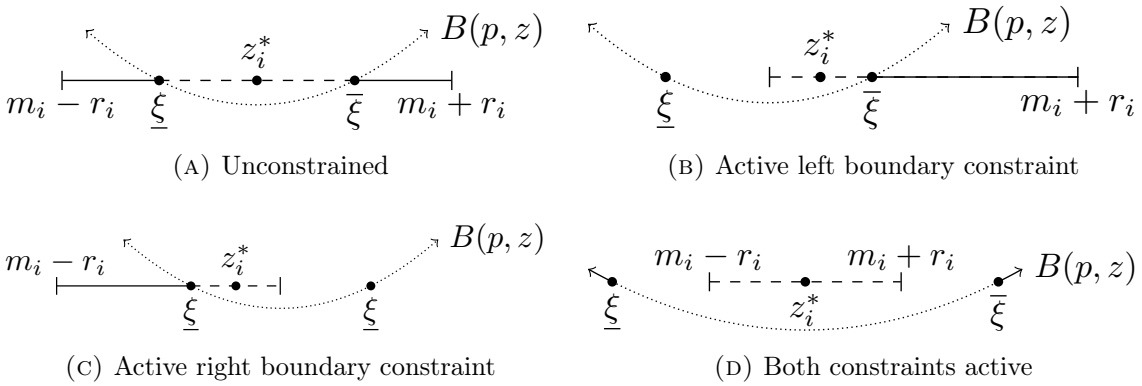

(A) Unconstrained (B) Active left boundary constraint

(C) Active right boundary constraint (D) Both constraints active

FIGURE 3. A typology of equilibrium participation. Solid lines indicate types that would participate according to a^*, while dashed lines indicate types that would abstain. The midpoint of the abstention region is indicated by z_i^*

INTERPRETING RULEMAKING DATA

The previous section developed a model of rulemaking as learning and responding to public preferences. In the equilibrium of this game, public participation and policy choice reflect strategic positioning by policymakers and endogenous selection by influence-seekers. For empirical applications, these confounding factors are important. The model provides structure for thinking about what role they might play.

Yet the model suggests there are even more fundamental issues than endogeneity for studies that relate participation and policy change to bias. Encouragingly, the model shows that public participation and policy change are strongly related to regulatory bias. One can derive predicted participation propensities and policy changes from the equilibrium conditions. Because allocative bias α occurs in every condition, data about these key quantities must reflect regulatory bias.

On the other hand, inferential error also matters. As the model shows, inferential error and allocative bias must be sufficiently large to allow for participation and policy change. If the regulator cares too little about outside actors' preferences, or if those preferences are too predictable, then participation rates will be low, and policy will be static. If both quantities are sufficient, however, the two can work as economic substitutes. The substitutability of allocative bias and inferential error makes interpretation of rulemaking data difficult.

Unless one accounts for the distribution of preferences among the public, inferential error can potentially do all the work of explaining inequality in participation and policy change.

The rest of this section develops these points in greater detail. The first subsection emphasizes how participation propensities and policy changes relate to allocative bias α under the model. The second subsection presents a numerical example that helps illustrate the points this paper makes. For the purposes of presentational simplicity, I adopt the assumption that the regulator is "myopic" in the sense that it ignores the effects of its policy proposal on the public's informational incentives. Formally, this means I replace Condition (4) of Proposition 1 with the assumption that $x_p^* = \alpha m_i$.

Zealots and Near-Sighted Compromisers. To highlight the role of allocative bias α in shaping participation patterns, I contrast two different types of regulators. The first is a "zealot" that strongly values an independent idea of what public policy should be. For the zealot, $\alpha = 0$, so the regulator derives utility only from making close to the policy it would independently prefer, without loss of generality normalized to 0. The second regulator has $\alpha > 0$. As the previous section showed, this assumption implies that the regulator will seek some compromise between its independent vision of policy and the policy that the public would most prefer. Hence, the second regulator is a "compromiser."

How different are zealots and compromisers on the basis of the data that empirical scholars collect? The key quantities of interest are individual i's participation propensity,

$$\Pr(a(p_i) = 1)$$

and the policies that the regulator select's

$$x_f - x_p$$

If one assumes types are independent, one does not need to consider the broader case of an arbitrarily large finite population because

$$\mathbb{E}(\sum_{i=1}^{n} a_i(p_i)) = n \Pr(a(p_i) = 1)$$

The dynamics prove similar. Note that the results are also the same if the n individuals come from one of k different distributions, provided that the regulator knows the distribution from which individual i 's type arises.

I begin by characterizing the participation propensity, policy selection, and policy change for the case where the regulator is a zealot.

Proposition 2. *If the regulator is a zealot ($\alpha = 0$), then*

(1) $\Pr(a(p_i) = 1) = 0$

(2) $x_f - x_p = 0$

These results follow readily from Proposition 1, but they are also intuitive. A regulator that is a zealot only derives utility from making policy at its independent preference of 0. As a result, the zealot proposes and finalize policy at its ideal pont. Since the zealot will not change policy in response to public comments, the public has no reason to provide them. Note that the result does not depend upon what one assumes about z.

Predicting participation and policy change for the comrpomiser is more subtle. I shall take each in turn.

Proposition 3. *If the regulator is a myopic compromiser ($\alpha > 0$,$x_p = \alpha m_i$) and the distributional constraints do not bind,then*

$$\Pr(a(p_i) = 1) = 1 - \frac{(k+1)}{r_i}\sqrt{\frac{c}{\alpha(2k+2-\alpha)}}$$

The Proposition reveals a direct, if constrained relationship between α and the probability of participation. The maximum value that α may obtain is 1. Provided k, c, and r_i are suitably chosen, the probability of participation can be strictly greater than zero, but it need

not be.[6] It is unsurprising that participation might not satisfy the cost-benefit analysis even if α were at its maximum. Policy benefits are necessarily finite in this model and the costs of signaling, while also finite, are theoretically unbounded. Given exogenous parameters such that a positive probability of participation exists when $\alpha = 1$, there is also a value $\tilde{\alpha}$ so that $\Pr(a(p_i) = 1) = 0$ for the same parameters. Only if $\alpha \in (\tilde{\alpha}, 1)$ is there a direct relationship between allocative bias and participation probability. These conclusions still hold if just one of the distributional constriants is active, as Proposition 4 summarizes.

Proposition 4. *If the exogenous parameters k, c, m_i, and r_i are suitably chosen, the probability of i participating is 0 even when $\alpha = 1$. If the exogenous parameters are such that there is some chance of participation when $\alpha = 1$, there exists a critical level of allocative bias $\tilde{\alpha}$ with the following properties. If $0 < \alpha < \tilde{\alpha}$, the expected amount of participation from the public is 0. For $\alpha \in (\tilde{\alpha}, 1)$, the probability of participation increases with α.*

The model shows that there is a significant interpretive asymmetry in approaching public participation during rulemaking. If the probability of public participation is greater than 0, then that implies that $\alpha > 0$. However, if the probability of public participation is observed to be 0, then it does not follow that $\alpha = 0$. Indeed, it could be the case that the costs of commenting, c; the costs of amending, k; and the unpredictability of the public's preferences, r_i, are such that participation would not be justified for *any* level of concern. One might also observe some positive participation if only α were higher.

A similar interpretive asymmetry affects data about rule changes. Proposition 1.4 implies that if the regulator is myopic,

$$x_f^* - x_p = \frac{\alpha}{k+1} \left(\mathbb{E}(p_i \mid a_i^*) - \mathbb{E}(p_i) \right)$$

Myopic compromisers change their policy position, according to the formula, whenever the regulator updates its prior beliefs after observing the public's participation decision.

[6]Note that the assumption of lax distributional constraints is also necessary, as certain combinations of α, k, c, and $r_i$can make the implied participation propensities negative. Such non-sensical results only occur because one or both of the boundary constraints are active. See the Appendix for the alternative formulae.

Researcher Observers	Rule Change	No Rule Change
Comments	$\alpha > 0$	Impossible Under Model
No Comments	$\alpha > 0$	$\alpha = 0$ or $\alpha > 0$

TABLE 1. Possible degrees of allocative bias under the model

If the exogenous parameters were such that participation was hypothetically possible ex ante, then even abstention is informative. Thus, policy change *can* occur under the model despite a lack of comments. If the exogenous parameters were such that participation was always impossible ex ante, then commenting does not occur and abstention is not informative. In this case, the regulator's posterior is its prior and so policy change does not occur. No change is also consistent with the regulator having an obsequious willingness to compromise, which is to say even if $\alpha = 1$.

If a researcher observes rules changing or public commenting, it establishes that agencies are compromisers, according to the model. Yet, when a researcher observes that no one has commented and the rules have not changed, it does not establish that the rulemaker is a zealot. Rather, it could be the case that participation was strategically infeasible: thus, the compromising regulator rationally did not adjust its prior belief. All of the parameters of the model are potential explanations for why participation was infeasible. More comments might have been submitted if the regulator had more concern for the public, or if commenting were less costly, or even if the costs of amending policy were less costly. These interpretive asymmetries suggest that rule change and participation rate data are ambivalent evidence about α. Nevertheless, the combination of these data do yield a falsifiable prediction of the model. Rules should always change following commenting. If rules did not change following commenting, that would constitute evidence of zealotry. Yet zealotry makes commenting irrational, so the model would be false. Table 1 summarizes the interpetive conclusions of this section.

Favoritism By Examples. This section provides several numerical examples that highlight the insights of the previous sections. The first goal of this section is to more concretely

illustrate how participation rates and policy changes can mislead researchers about the direction of regulatory bias. It also aims to show how the model gives insights into how the regulatory capture phenomenon might manifest during rulemaking. The fact-pattern for this example is based on Golden (1998), a well-regarded and often-cited study in the rulemaking literature. In that paper, Golden randomly selects eleven rules from the Federal Register. She observes and analyzes who comments and how the rules change between proposal and final stage. She finds that "the bias in participation rates of business and citizen groups and the virtual absence of actual citizen participation demonstrate conclusively that the accent of the heavenly choir in the rule-making process is off key."

Suppose that the public has two sectors, suppliers of a product and buyers of a product. As a further simplification, each sector has one member, s and b, respectively. To analyze rulemaking given a multi-member public, I model allocative bias as a vector $\alpha = \langle \alpha_0, \alpha_b, \alpha_s \rangle$. Each sector has its own distribution from which nature draws types. Again suppose that $p_i \sim \mathcal{U}(m_i - r_i, m_i + r_i)$ for $i \in \{s, b\}$. For further concreteness, suppose that $m_b = -1$ and $m_s = 1$, so that negative policies favor buyers and positive policies favor suppliers. Ignoring the public, the regulator would want a policy squarely between the two, so that the expected level of preference conflict between the regulator and each segment of the public is identical.

Projecting onto Propensities. A researcher collects data on public participation rates by sector. The propensity of suppliers and buyers to submit comments is estimated as $\langle \hat{e}_b, \hat{e}_s \rangle = \langle 0, 0.1 \rangle$. It turns out that the estimated participation probabilities are consistent with the regulator strongly favoring *either group*, depending on r_b and r_s. These two parameters control the fundamental uncertainty about preferences in the environment and, therefore, the inferential error that members of the public may expect. Following the derivations in the appendix, the ex ante probability of participation by s, assuming the boundary constraints do not bind is

$$\Pr(a_s = 1) = e_s = 1 - \left(\frac{k+1}{r_s} \right) \sqrt{\frac{c}{\alpha_s(2k + 2 - \alpha_s)}}$$

The formula is entirely analogous for the case of consumer b, and is virtually identical to Proposition (3). If one assumes modification costs are 0 and commenting costs are 1,

$$\langle e_b, e_s \rangle = \left\langle 1 - \left(\frac{1}{r_b}\right)\sqrt{\frac{1}{\alpha_b(2-\alpha_b)}}, 1 - \left(\frac{1}{r_s}\right)\sqrt{\frac{1}{\alpha_s(2-\alpha_c)}} \right\rangle$$

The expression shows that the participation propensities are separable. This fact follows even though allocative bias is zero-sum and the policy outcomes are jointly dependent on each member of the public's participation and abstention decisions. In particular, virtually any value of e_b and e_s are possible, regardless of the allocative bias, so long as one suitably adjusts r_s and r_b. If $\vec{\alpha} = \vec{\alpha}_s = \langle 0, 0.1, 0.9 \rangle$, so that the regulator greatly favors suppliers, we can have $r_b \approx 1.12$ and $r_s \approx 2.3$. According to the formula, it follows that the true participation propensities would be approximately $\langle 0, 0.1 \rangle$. Suppose instead that $\vec{\alpha} = \vec{\alpha}_b = \langle 0, 0.9, 0.1 \rangle$, so that the regulator greatly favors buyers. If $r_b \approx 1$ and $r_s \approx 2.55$, the true participation propensities are also approximately $\langle 0, 0.1 \rangle$. In other words, strong favoritism toward one group or *the other* can equally lead to a participation propensity that favor suppliers.

The key omitted variables that allow for this mischief, r_b and r_s, control the average level of inferential error in the environment. A 10% participation propensity is preserved despite a decrease in α_s from 0.9 to 0.1 by increasing r_s from 2.3 to 2.55. A 0% participation propensity is preserved despite an increase in α_b from 0.1 to 0.9 by decreasingr_b from 1.12 to 1. If one knew that supplier's preferences were *more predictable* than the buyer's preferences, so that $r_s < r_b$, none of these examples would be possible. If the sector with more predictable preferences participates more, the only interpretation under the model is that the regulator desires to create policy more favorable to that sector.

Shifty Changes. A scholar now considers data collection on policy changes as well as participation. Unfortunately, policy change also yield ambiguous evidence of allocative bias. To see why, suppose for simplicity that the regulator make a "myopic" policy proposal:

$$x_p = \alpha_b m_b + \alpha_s m_s$$

To ensure that this proposal matters for final policy, let $k = 1$. I will identify *numerically* the position of the initial policy x_p; the final policy x_f; and the policy change $x_f - x_p$. I will do this twice, once with $\vec{\alpha}_s = \langle 0, 0.1, 0.9\rangle$ and once with $\vec{\alpha}_b = \langle 0, 0.9, 0.1\rangle$. As an aide the reader, Table 2 summarizes the two sets of calculations.

Supplier favoritism. If $\vec{\alpha} = \vec{\alpha}_s = \langle 0, 0.1, 0.9\rangle$, the regulator favors suppliers and proposes policy at 0.8. The formulas in the appendix show how to calculate the public's best-response function. If the distributional constraints do not bind, the following formalas apply:

$$z_i = \left(\frac{kx_p + \sum_{j\neq i}\alpha_j m_j}{k+1-\alpha_i}\right) \quad a_i(p_i) = \begin{cases} 1 & p_i \in (z_i - \delta_i, z_i + \delta_i) \\ 0 & \text{o/w} \end{cases}$$

where $\delta_i = (k+1)\sqrt{\frac{c}{\alpha_i(2k+2-\alpha_i)}}$. I have already shown that manipulating r_s and r_b can make participation impossible or ensure that neither boundary constraint binds. Assume $r_s = 2$, which is sufficient to ensure the boundary constraints are slack for s. Let r_b be sufficiently small that participation is infeasible for b. This implies that

$$z_s \approx 0.6 \qquad a_s(p_s) \approx \begin{cases} 1 & p_s \in (1, -0.6) \cup (1.8, 3) \\ 0 & p_s \in [-0.6, 1.8] \end{cases}$$

$$z_b = -1 \qquad a_b(p_b) = 0$$

According to the Appendix, final policy is given by the following equation

$$\left(\frac{k}{k+1}\right) x_p + \left(\frac{1}{k+1}\right) \sum \alpha_i \mathbb{E}(p_i \mid a_i)$$

The final policy that results is therefore a random variable that depends on p_s. Importantly, it does not depend on p_b since abstention is certain. Consider what happens when s has typical preferences for a commenter ($p_s = 1.6$) . Then final policy is

$$x_f^1 = \frac{1}{2}(0.8) + \frac{1}{2}(0.1 * (-1) + 0.9 * 1.6) = 1.07$$

Thus, policy has changed from 0.8 to 1.07, a shift of +0.27.

These calculations reveal all the components of the data generating process, but the researcher only observes a limited set of these quantities. In particular, she notes that the supplier sector participated, but the buyer sector did not. Fruther, the rule moved in a positive, or "supplier," direction. The researcher argues that the regulator is likely biased in favor of the suppliers, and she is correct.

Buyer favoritism. Consider what the data-generating process would have been like if $\alpha = \alpha_b = \langle 0, 0.9, 0.1 \rangle$, $r_s = 5.5$, and r_b is still too small for any buyer to justify participation. Then the initial policy is $x_p = -0.8$, favoring the buyers. The following formula summarizes the decision and imputation strategies

$$z_s = -0.9 \qquad a_s(p_s) = \begin{cases} 1 & p_i \in (-4.5, -4.1) \cup (2.3, 6.5) \\ 0 & p_i \in [-4.1, 2.3] \end{cases}$$

$$z_b = -1 \qquad a_s(b) = 0$$

Consider how policy changes when s has typical preferences for a commenter given these parameters (p_s =4.08). Final policy is given by

$$x_f^1 = \frac{1}{2}(-0.8) + \frac{1}{2}(0.9 * (-1) + 0.1 * 4.08) = -0.646$$

Thus policy changes from -0.8 to -0.646, a shift of +0.154.

In this circusmtance, a researcher would see that suppliers had participated participated but the buyers did not. The policy has also shifted in a direction that favors the suppliers. In others words, this researcher observes the same evidence that the other researcher did. If the researcher concludes that the regulator and regulation are biased in favor of the supplier, however, he is wrong.

Discussion. The key driver of the observational equivalency between strong favoritism toward the suppliers and strong favoritism toward the buyers is, as before, manipulation of the predictability of preferences. In both examples, the policy benefits the suppliers because

Regulator Favors	$\vec{\alpha}$	r_s	x_p	$\Pr(s \text{ comments})$	$\mathbb{E}(x_f \mid s \text{ comments})$	$\mathbb{E}(x_f - x_p \mid s \text{ comments})$
Suppliers	$\langle 0, 0.1, 0.9\rangle$	2	0.8	0.4	1.07	0.27
Buyers	$\langle 0, 0.9, 0.1\rangle$	5.5	-0.8	0.42	-0.646	0.154

TABLE 2. Calculated quantities. By assumption, r_b is small enough that participation by b is impossible. The probability that s comments is similar in both instances, and so is the direction of policy change.

they are the ones that comment. Interestingly, one can show that policy would have shifted in favor of s even if its best response were to abstain. The reason is that, so long as participation is feasible, abstention is meaningful; hence the regulator responds to it. If participation is unfeasible, the regulator learns nothing from its absence and so it does not respond. Additional allocative concern does not change the basic learning problem.

The example of this subsection assumed that participation was infeasible for the buyers. Yet, the same conclusions do follow even if the regulator learns about both sectors simultaneously through the comment process. If the change in expectations about s were larger than the change in expectation about b, the overall policy shift could favor s even if the regulator were biased in favor of b. The only way to prevent such pathologies under the model is to control the fundametal uncertainty, r_i. This technique is effective because it limits the possible inferential error $p_i - z_i$, which is the primary alternative to allocative bias as a driver of participation and responsiveness under the model.

CONCLUSIONS

Critics of administrative policy-making have long argued that Congress delegates too much policy-making authority to agencies. The question of who wins and who loses in the making of regulation is significant for scholars in political science, law, and public administration. Researchers working in these disciplines have adduced facts are in some ways consistent with our worst fears about rulemaking's "democratic deficit," yet in others ways not. Much of the reason for continued uncertainty about what these data reveal is current scholarly frameworks provide insufficient theoretical resources to guide interpretation.

I address this topic by developing a formal model of rulemaking as learning and responding to public preferences. The task of the regulator during rulemaking are to elicit public preferences and to engineer a compromise between these demands and the agency's own view of what policy should be. In certain respects, this account is similar to the model that the authors of the APA seem to have used. Through analysis of this formalization, I show that the model is consistent with evidence adduced by empirical researchers. Moreover, pluralistic participation and ubiquitous policy change is inconsistent with strong capture and bureaucratic zealtory. Yet stronger conclusions are diffucult to reach, because the preferences of outside actor's over policy – their "stakes" – can do all the work explaining who participates and who gets the benefits of policy change. In particular, bias toward business in policy change is consistent with regulator hostility toward business so long as the effect of regulation on businesses is sufficiently difficult for the regulator to predict.

Analysis of the model also provides an interpretive framework for distinguishing between three potential pitfalls in regulatory policymaking. Allocational bias relates to the weight the regulator assigns to the concerns of each stakeholder in rulemaking. Inferential error describes the difference between what the regulator assumes about public preferences it cannot observe and what they truly are. Systematic inferential error is inferential bias. Participation and responsiveness depend on sufficient allocational bias and inferential error. If enough of each is present, then these two quantities can act as economic substitutes. The substitutibility of these two quantities makes it hard to determine how much of each is present, as this paper shows formally through proof and numerically by examples.

The model presented here offers significant challenges to the literature, yet it also provides opportunities. Advances are continuously made in terms of ideal point estimation of executive agencies (Richardson, Clinton, and Lewis 2017; Chen and Johnson 2015; Clinton et al. 2012). It may soon be possible to measure the position of public actors and agencies on the per rule level. Even if measuring public preferences remains difficult, the model also raises many questions about the strength and validity of its core assumptions. What proportion of commenting activity can be characterized as informing regulators about public preferences

versus attempting to attract attention of other political principals? How do regulators think about the interests of the silent public? These and other questions the model raises should be engaged quantitiatively, qualitatively, and with additional formal theory.

References

Attorney General's Committee on Administrative Procedure. 1941. *Final Report*. Tech. rep. Washington, DC.

Balla, Steven J. 1998. "Administrative Procedures and Political Control of the Bureaucracy." *American Political Science Review* 92 (3): 663–673.

Balla, Steven J., and Benjamin M. Daniels. 2007. "Information technology and public commenting on agency regulations." *Regulation & Governance* 1 (1): 46–67.

Barron, David J., and Elena Kagan. 2001. "Chevron's Nondelegation Doctrine." *The Supreme Court Review* 2001:201–265.

Bawn, Kathleen. 1995. "Political Control Versus Expertise: Congressional Choices about Administrative Procedures." *American Political Science Review* 89 (1): 62–73.

Bendor, Jonathan, and Adam Meirowitz. 2004. "Spatial models of delegation." *American Political Science Review* 98 (2): 1–19.

Bhutta, Neil, Jacob Goldin, and Tatiana Homonoff. 2016. "Consumer Borrowing after Payday Loan Bans." *The Journal of Law and Economics* 59 (1): 225–259.

Bueno De Mesquita, Ethan, and Matthew C. Stephenson. 2007. "Regulatory Quality Under Imperfect Oversight." *American Political Science Review* 101 (3): 605–621.

Callander, Steven. 2011. "Searching for Good Policies." *American Political Science Review* 105 (4): 643–662.

Carpenter, Daniel P. 2004. "Protection without Capture: Product Approval by a Politically Responsive Learning Regulator." *American Political Science Review* 98 (4): 613–631.

— . 2010. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton: Princeton University Press.

Carpenter, Daniel P., and David A. Moss. 2013. *Preventing Regulatory Capture*. New York: Cambridge University Press.

Carpenter, Daniel P., and Michael M. Ting. 2007. "Regulatory Errors with Endogenous Agendas." *American Journal of Political Science* 51 (4): 835–852.

Chen, Jowei, and Tim Johnson. 2015. "Federal employee unionization and presidential control of the bureaucracy: Estimating and explaining ideological change in executive agencies." *Journal of Theoretical Politics* 27 (1): 151–174.

Clinton, Joshua D., et al. 2012. "Separated Powers in the United States: The Ideology of Agencies, Presidents, and Congress." *American Journal of Political Science* 56 (2): 341–354.

Cordray, Richard. 2017. *Prepared Remarks of CFPB Director Richard Cordray on the Payday Rule Press Call*. Washington, D.C.

Cuellar, Mariano-Florentino. 2005. "Rethinking Regulatory Democracy." *Administrative Law Review* 57 (2): 411–499.

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

Epstein, David, and Sharyn O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. New York: Cambridge University Press.

Figueiredo, John M. de. 2006. "E-Rulemaking: Bringing Data to Theory at the Federal Communications Commission." *Duke Law Journal* 55:969.

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. 603. Cambridge: MIT Press.

Furlong, Scott R., and Cornelius M. Kerwin. 2005. "Interest Group Participation in Rule Making: A Decade of Change." *Journal of Public Administration Research and Theory* 15 (3): 353–370.

Gailmard, Sean, and John W. Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15 (1): 353–377.

Gailmard, Sean, and John W Patty. 2016. "Participation, process and policy: the informational value of politicised judicial review." *Journal of Public Policy* 1 (28).

— . 2007. "Slackers and Zealots: Civil Policy and Bureaucratic Expertise." 51 (4): 873–889.

Golden, Marissa Martino. 1998. "Interest Groups in the Rule-Making Process: Who Participates? Whose Voices Get Heard?" *Journal of Public Administration Research and Theory* 8 (2): 245–270.

Grossman, Gene M., and Elhanan Helpman. 1994. "Protection for Sale." *American Economic Review* 84 (4): 833–850.

Haeder, Simon F., and Susan Webb Yackee. 2015. "Influence and the Administrative Process: Lobbying the U.S. President's Office of Management and Budget." *American Political Science Review* 109 (3): 507–522.

Hamburger, Philip. 2014. *Is Administrative Law Unlawful?* Chicago: University of Chicago Press.

Hitlin, Paul, Kenneth Omstead, and Skye Toor. 2017. *Public Comments ot the Federal Communications Commission About Net Neutrality Contain Many Inaccuracies and Duplicates*. Tech. rep. Pew Research Center.

Hojnacki, Marie, et al. 2012. "Studying Organizational Advocacy and Influence: Reexamining Interest Group Research." *Annual Review of Political Science* 15, no. 1 (): 379–399.

Jensen, Christian B., and Robert J. McGrath. 2011. "Making Rules about Rulemaking: A Comparison of Presidential and Parliamentary Systems." *Political Research Quarterly* 64 (3): 656–667.

Kerwin, Cornelius M., and Scott R. Furlong. 2011. *Rulemaking: How Government Agencies Write Law and Make Policy*. 4th ed. Washington, DC: CQ Press.

Krawiec, Kimberly D. 2013. "Don't 'Screw Joe the Plummer': The Sausage-Making of Financial Reform." *Arizona Law Review* 55:53–103.

Kully, Kris D, and Stephanie C Robinson. 2017. *The CFPB's Payday Lending Rule*. Tech. rep. Mayer Brown.

Lubbers, Jeffrey S. 2012. *A Guide to Federal Agency Rulemaking*. 5th ed. Chicago, IL: American Bar Association.

McCarty, Nolan M. 2017. "The Regulation and Self-Regulation of a Complex Industry." *Journal of Politics* 79 (4): 1220–1236.

McCubbins, Mathew D., and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *Political Science* 28 (1): 165–179.

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): 243–277.

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

Michael, Gabriel J. 2013. "Politics and Rulemaking at the Copyright Office." *Journal of Information Technology & Politics* 1681 (January 2014): 131213092240006.

Milgrom, Paul. 2008. "What the Seller Won't Tell You: Persuasion and Disclosure in Markets." *Journal of Economic Perspectives* 22 (2): 115–131.

Milgrom, Paul, and John Roberts. 1986. "Relying on the Information of Interested Parties." *RAND Journal of Economics* 17 (1): 18–32.

Nelson, David, and Susan Webb Yackee. 2012. "Lobbying Coalitions and Government Policy Change: An Analysis of Federal Agency Rulemaking." *The Journal of Politics* 74 (02): 339–353.

Niskanen, William A. 1968. "The Peculiar Economics of Bureaucracy." *American Economic Review* 58 (2): 293–305.

Richardson, Mark D., Joshua D. Clinton, and David E. Lewis. 2017. "Elite Perceptions of Agency Ideology and Workforce Skill." *The Journal of Politics* 80 (1): 303–308.

Stephenson, Matthew C. 2007. "Bureaucratic Decision Costs and Endogenous Agency Expertise." *Journal of Law, Economics, and Organization* 23 (2): 469–498.

Stigler, George J. 1971. "The theory of economic regulation." *The Bell Journal of Economics and Management Science* 2 (1): 3–21.

Stoll, Richard G. 2010. *Effective EPA Advocacy*. Oxford: Oxford University Press.

Volden, Craig, Michael M. Ting, and Daniel P. Carpenter. 2008. "A Formal Model of Learning and Policy Diffusion." *American Political Science Review* 102 (03): 319–332.

West, F. 2004. "Procedures, Informal Processes, Accountability, and Responsiveness in Bureaucratic Policy Making: An Institutional Policy Analysis." *Public Administration* 64 (1): 66–80.

Wilson, James. 1989. *Bureaucracy: What Government Agencies Do and Why They Do It.* New York: Basic Books.

Yackee, Jason Webb, and Susan Webb Yackee. 2006. "A Bias Towards Business? Assessing Interest Group Influence on the U.S. Bureaucracy." *The Journal of Politics* 68 (1): 128–139.

— . 2010. "Administrative procedures and bureaucratic performance: Is federal rule-making "ossified"?" *Journal of Public Administration Research and Theory* 20 (2): 261–282.

Yackee, Susan Webb. 2005. "Sweet-talking the Fourth Branch: The Influence of Interest Group Comments on Federal Agency Rulemaking." *Journal of Public Administration Research and Theory* 16:103–124.

Appendix

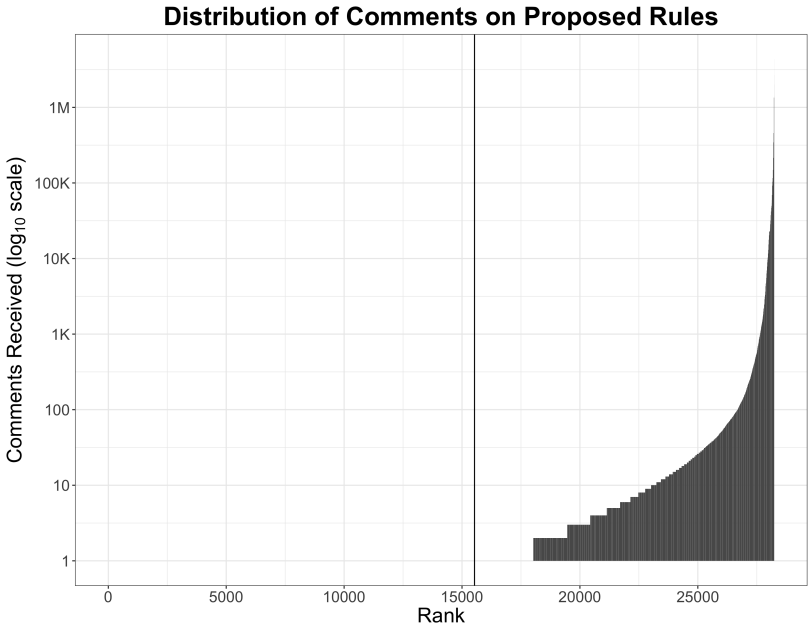

Figure A1. Distribution of comments on proposed rule dockets (source: regulations.gov)