STRATEGIC PROPOSALS, ENDOGENOUS COMMENTS, AND BIAS IN RULEMAKING

ABSTRACT. Agencies use notice-and-comment rulemaking to issue countless regulations with substantial economic stakes. The empirical literature has produced a complex set of descriptive findings, yet has struggled with informal concerns about selection bias. This paper characterizes notice-and-comment as a persuasion game played between regulators and outside interests. Analysis of this "stakeholder-balancing" model yields three key theoretical payoffs: an informational rationale for regulators to write rules with higher private and social costs, an explanation for strategic positioning by regulators even without oversight, and clarification that adverse priors are a more powerful mobilizing force than adverse policies. The model's two-sided selection dynamics reveal that well-established empirical regularities are inconsistent with extreme public-interest zealotry and strong capture, but fit a range of intermediate outcomes. To obtain deeper insights about bias in rulemaking, the model suggests focusing on the cost of rule revision, rule movement following abstention, and variation in stakeholder preferences.

Date: 12/3/18.

Ten days after assuming office, President Trump issued an executive order requiring agencies to remove two rules for every new one they issued (Executive Order 13771). Similar orders seeking to control agency rulemaking were issued by Reagan, Clinton, and Obama (Lubbers 2012). Yet the notice-and-comment process seems surprisingly resistant to external control. In spite of the executive order, rule output by federal agencies was not much different in 2017 than it was in 2009, the first year of the Obama presidency (Coglianese 2018). It is also unclear whether the policies emerging through rulemaking today are less important or even less "regulatory" than in past administrations. Salient policies issues being decided through notice-and-comment in the last two years include whether "bump stock" devices are machine guns (83 FR 13448), whether internet service providers should be regulated as utilities (83 FR 7852), and what protections consumers should have in the market for "payday" loans (82 FR 54472). Given the importance of rulemaking as a site of contemporary policymaking, which groups win and which groups lose through the process and why?

The growing empirical literature on rulemaking process suggests a complex fact-pattern that resists parismonious description. While some rules receive hundreds of thousands of comments or more, mass mobilization is rare. Instead, rulemaking usually involves conflict between a more limited set of stakeholders (de Figueiredo 2006; Kerwin and Furlong 2011). Mobilization typically favors regulated industries over industries not directly subject to regulation, favors businesses over consumers or public interest groups, and rarely involves citizens. Organized interests invest heavily in influencing rulemaking, most visibly by submitting comments on rules that agencies have proposed, but not yet finalized. A single comment letter can cost one hundred thousand dollars or more (Dash 2011). It is hard to know how much commenting costs in aggregate, because such activity falls outside of lobbying disclosure requirements. Nevertheless, the undisclosed costs associated with rulemaking advocacy conceivably exceed that of all reported lobbying expenditures. Even looking at reported lobbying expenses, evidence suggests that most of this may also be aimed at influencing agency decision-making during rulemaking (You 2017; Boehmke, Gailmard, and Patty 2013). Given the magnitude of these lobbying investments, one would imagine there

are returns. A significant body of evidence now shows that commenting appears to lead to favorable policy changes between proposed and final rules (Balla 1998; Yackee and Yackee 2006; Yackee 2006; Haeder and Yackee 2015). Troublingly, the gains from commenting do not appear to accrue equally to all commenters.

These and other findings are often regarded as descriptive evidence that rulemaking is biased (Yackee and Yackee 2006; Krawiec 2013). To be sure, the normative stakes of such empirical results are profound. Yet, there are doubts about how to interpret evidence of regulatory favoritism (Carpenter 2004). As Yackee (2019) writes, "most of the quantitative research on rulemaking thus far uses observational data to study patterns and correlates." Regulators and their stakeholders are highly strategic actors (Carpenter 2010). They make policy and advocacy decisions based not only on their own incentives, but also in anticipation of what their counter-party's incentives would be under various potential courses of action. Therefore, rulemaking data is likely susceptible to selection bias. How severe these issues are, in which direction they bias results, or how to control for them, are rarely discussed in detail. Complicating matters further, there is uncertainty not only about the effect of strategy on data, but also about which strategic actors are truly relevant. Congress, the President, and even courts rarely play a visible role in the typical or even most rulemakings. Even so, perhaps these powerful actors do exert a kind of gravitational force over rulemaking, which is nowhere seen yet everywhere felt. Without an understanding of how rulemaking would look in the absence of these actors, it is hard to know what would establish that they do play a significant role.

In light of such concerns, there is an acute need for new formal models of the rulemaking process, especially ones that focus on the distributive dynamics of rulemaking.[1] Here, I present one such model, where notice-and-comment is viewed as a persuasive signaling game

[1]The formal-theoretic literature on rulemaking specifically and administrative policymaking more generally has also grown in recent years (e.g., Stephenson 2007; Bueno De Mesquita and Stephenson 2007; Gailmard and Patty 2016; McCarty 2017; Turner 2017). These papers have largely focused on questions about how oversight effects bureaucratic capacity, however, rather than questions about how notice-and-comment structures the resolution of interest group competition.

played between outside actors who can reveal their preferred policies and an imperfect social planner who manages interest-group conflict in a possibly biased fashion. The "interest-balancing" framework that results is different from models in which commenting primarily produces expressive benefits (Coglianese 2006), where the gains of commenting stem from access to judicial review (Chubb 1983), or where comments are cheap-talk (Coleman 2016). While the latter viewpoints each capture some dynamics of regulatory advocacy at least some of the time, the persuasion framework is an especially important baseline for theoretical and empirical researchers to have available. The legislative history of the APA suggests that its authors made similar assumptions about how notice-and-comment rulemaking would work (Acheson Report, 1941). Additionally, the persuasion framework is useful for thinking about how interest groups would obtain influence from agencies with high degrees of bureaucratic autonomy (Carpenter 2010). On a theoretical level, this model provides three main payoffs for understanding the rulemaking process. It provides an informational rationale for regulators to write rules with private and social costs. It explains why strategic positioning by regulators during rulemaking might occur even without effective oversight by Courts or Congress. Finally, it presents a micro-founded account of what mobilizes outside actors during rulemaking. In particular, under this model, adverse regulatory priors are a more direct motivation for submitting comments than adverse regulatory policy. On an empirical level, the model also has important payoffs for the interpretation of empirical rulemaking data as it relates to regulatory bias. Bias in policy change and bias in mobilization are inconsistent with extreme forms of public-interest zealotry or regulatory capture, even given the two-sided selection issues the model assumes. Yet there is a vast observational equivalence between existing data and a range of intermediate forms of regulatory bias. To obtain deeper insights about bias in rulemaking, the model suggests focusing on the cost of rule revision, rule movement following abstention, and variation in stakeholder preferences.

Rulemaking as Learning and Adjusting to Public Preferences

The Model This game is played between a regulatory agency G and one of perhaps several outside actors i. One might think of i as an individual firm or interest group, or even as

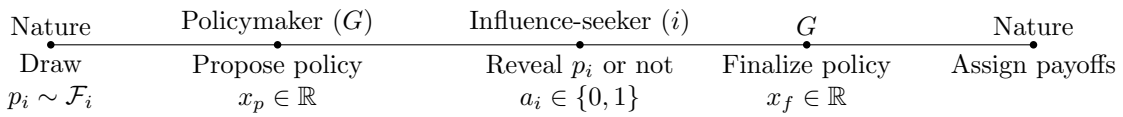

FIGURE 1. Order of play during the rulemaking game

a private citizen. The appendix provides proofs assuming an arbitrary number of outside actors, so I shall sometimes refer to i as a member of the "public." For presentational reasons, however, I shall begin by focusing on the simpler case where the "public" has a single actor, and then describe the wrinkles that are introduced by competition between influence-seekers. As Figure 1 shows, the game begins with regulator G proposing a policy $x_p \in \mathbb{R}$. The outside actor observes the proposal and chooses whether to send a comment by selecting $a_i \in \{0, 1\}$. If i comments ($a_i = 1$), then G can observe i's "bliss" policy choice, $p_i \in \mathbb{R}$. If i abstains ($a_i = 0$), G does not observe p_i. After observing i's signal, G selects a final policy $x_f \in \mathbb{R}$ and payoffs are realized.

Both G and i have utility functions whose values depend on (1) which policy is implemented and (2) actions undertaken during rulemaking. For i, indirect policy utility is a negative quadratic centered at p_i. Since G has ultimate discretion over final policy x_f, this aspect of i's utility function depends entirely on what the regulator chooses. i has no way of *forcing* the regulator to make a particular policy selection. Nevertheless, commenting indirectly influences policy selection by changing G's information set, usually for the benefit of both actors.[2] In order to comment, i must incur cost c. Formally, i's utility function is given as follows:

$$u_i(x_f, a_i) = -(x_f - p_i)^2 - ca_i$$

I model G as an imperfectly benevolent social planner, with an emphasis on "imperfectly." The indirect policy utility function of G is given by

$$v_G(x_f) = -(1 - \alpha)x_f^2 - \alpha(x_f - p_i)^2$$

[2]In equilibrium, commenting always brings policy benefits relative to abstention for both G and i, but if G is not fully rational then i can sometimes encourage G to compromise more than it should by abstaining, to i's benefit and G's loss.

Here, $\alpha \in [0,1]$ is the "allocative bias" parameter. It modulates the weight of outside actor i's preferences on G's decision-making. If $\alpha = 1$, then G has no preferences over policy independent of what the "public" wants: $v_G(x_f) = -(x_f - p_i)^2$. While one could conceive of policy utility in this case as welfarian, depending on the context one might imagine that the social optimum is not $\alpha = 1$. The contemporary bureaucratic politics literature has shown that agencies often do have competing concerns besides the preferences of the "public" (Gailmard and Patty 2007; Carpenter 2010; Clinton et al. 2012; Chen and Johnson 2015; Richardson, Clinton, and Lewis 2017), for better and for worse. The model can describe such dynamics as well. At the extreme, if $\alpha = 0$ then the public's welfare plays no role in G's policy selection: $v_G(x_f) = -x_f^2$. Following Gailmard and Patty (2007), I describe such a bureaucrat as a "zealot." Implicitly, this indirect utility function assumes that the bliss policy of the regulator is 0. This assumption is made without loss of generality and has no specific substantive content; it only means that G's independent preference is the reference point for all other actors preferences.[3] It is also worth noting that in the more complex case of n outside actors, the allocative bias parameter is a vector $\vec{\alpha}$, which describes not just an overall "public welfare" orientation but also the relative weight of each outside actor on G's decision-making.[4] In this way, $\vec{\alpha}$ can describe a regulator whose "public welfare" orientation tends to prioritize certain outside actors over others.

The regulator's utility does not only depend on policy preferences, but also on its choices during rulemaking. In particular, to modify its proposal, the regulator must pay a cost that increases with the size of the modification. The reputational, effort, and oversight costs associated with modifying policy are extremely high in the view of some legal scholars (Elliott 1992; McGarity 1992; Barron and Kagan 2001), although Yackee and Yackee (2010) provide empirical evidence that these costs are not so overwhelming as legal scholars had feared. In this model, moving policy from x_p to x_f imposes costs $k(x_f - x_p)^2$. The regulator's

[3]Since we make no restrictions on p_i or any of the other spatial parameters, 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.

[4]Formally, indirect policy utility is given by $v_G(x_f) = -(1 - \sum_{i=1}^{n} \alpha_i)x_f^2 - \sum_{i=1}^{n} \alpha_i(x_f - p_i)^2$.

overall utility is therefore given by

$$u_G(x_f, x_p) = -(1-\alpha)x_f^2 - \alpha(x_f - p_i)^2 - k(x_f - x_p)^2$$

The final ingredient in the model relates to the private information potentially transmitted during notice-and-comment. Fully capturing the richness of information transmission during rulemaking would be difficult for any parsimonious model, since different authorities tend to emphasize qualitatively different kinds of private information that comments communicate. Some sources suggest that comments must communicate technical or scientific information in order to persuade (Nelson and Yackee 2012). Others emphasize that comments serve as a signal of the strength of one's intentions to seek judicial review (Wagner 2012). Depending on the agency context, each of these factors might matter more or less. Regardless of domain, however, comments have the important function of signaling preferences. In this model, then, private information relates to i's bliss policy p_i, which i always learns at the start of the game, but G only observes if i comments. In the appendix, I characterize equilibrium assuming p_i comes from an arbitrary continuous distribution $\mathcal{F}_i$ that only has positive probability density within some bounded interval $(m_i - r_i, m_i + r_i)$. The prior distribution of i's type is common knowledge and so too are each actor's utility functions.[5]

In order to describe the Perfect Bayesian Equilibrium (PBE) of the rulemaking game, it will prove useful to introduce expressions for the location of several reference policies. These benchmarks are selected by the regulator in simplified versions of the complete rulemaking game. The first two reference policies, for example, each arise from eliminating a stage from the rulemaking game (see Figure 1).

[5]In the appendix, I describe how the strategic interaction changes if α is not common knowledge. If i has a prior on G's type α rather than knowing it surely, then x_p is an informative signal about α. Two-sided signaling creates additional strategic interpretation and selection problems, although in special cases x_p can fully reveal α. Whether α is fully revealed or not, i's best response is similar to the baseline, with $\mathbb{E}(\alpha \mid x_p)$ playing an analogous role to α. G's final policy choice is similar to the case where common-knowledge is assumed, although its problems in interpreting abstention are more complicated, as G will need some prior belief about what i believes about α.

Proposition 1. *If the regulator did not have access to notice and comment, then the regulator's best final policy selection would reflect its* prior expectations *of what the outside actor wants, given the weight the regulator deems appropriate.*

$$x_f^* = \alpha \mathbb{E}(p_i)$$

If the regulator received comments before issuing a proposal, the regulator's best final policy selection would reflect its posterior expectations *of what the outside actor wants, given the weight the regulator deems appropriate.*

$$x_f^* = \alpha \mathbb{E}(p_i \mid a_i)$$

Both expressions highlight the rulemaker's tendency to enact final policy that represents a "compromise" between competing values. The allocative bias α in G's utility function determines the balance. If $\alpha = 0$, then the regulator only cares about its independent preference, so the final policy outcome is the same in either simplified rulemaking game $x_f^* = 0$. If $\alpha = 1$, then the regulator only cares about the preferences of the outside interest, so it enacts the policy it expects the outside interest will most prefer. Depending on whether the regulator has access to commenting, the expectations will be more or less informed. For levels of allocative bias between 0 and 1, the regulator selects a final policy that is a weighted combination between the independent policy preference and its own expectations about outside interests' preferences.

Proposition 1 describes the resolution of preference conflict between one outside interest and an independent regulatory concern. There are analogous expressions that are useful for understanding the PBE given multiple influence-seekers. If there are n outside interests, $x_f^* = \sum_{i=1}^{n} \alpha_i \mathbb{E}(p_i)$ is the policy selected without access to comments, while $x_f^* = \sum_{i=1}^{n} \alpha_i \mathbb{E}(p_i \mid a_i)$ is the policy selected without proposals. Both final policies are weighted averages of the regulator's expectations about the bliss points of the outside interests and its own independent concern.[6] Clearly, balancing stakeholder interests is a key

[6]Implicity, in both expressions, G's own independent preference of 0 is given weight $1 - \sum_{i=1}^{n} \alpha_i$.

goal for the regulator in selecting final policy during rulemaking. Yet the regulator's own, paternalistic view of how to balance these competing interests greatly influences the policy outcomes. Going forward, I shall refer to the expression for final policy without access to comments as the *uninformed paternalistic compromise*. I will refer to the expression that comes from the regulator receiving comments before making policy as the *informed paternalistic compromise.*[7]

The last reference policy to consider is what would happen in the complete rulemaking game if nature exogenously decided whether outside interests comment. The resulting decision-problem is not really a "game" in the strict sense, because neither G or i makes strategic decisions in response to the other's incentives. Nevertheless, the outcomes are instructive, since they highlight how information provision incentives and strategy influence policy outcomes.

Proposition 2. *If the commenting decision of the outside interest is exogenously fixed by nature, then the best proposal for the regulator is the uninformed paternalistic compromise*

$$x_p^* = \alpha \mathbb{E}(p_i)$$

while the best final policy is a convex combination of the informed and uninformed paternalistic compromise, where the weight of each policy is determined by the amendment costs.

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

[7]It is worth mentioning two important caveats. First, an "informed" compromise does not necessarily mean "fully" informed. It is possible that the only piece of information G will acquire is that i declined to comment, in which case p_i is still unknown to G. Second, it only makes sense to speak of "the" informed paternalistic after nature assigns the outside interest a bliss point. Prior to that, $\alpha \mathbb{E}(p_i \mid a_i)$ is a random variable, so many informed paternalistic compromises are possible *ex ante*. By contrast, it always makes sense to speak of "the" uninformed paternalistic compromise because it is fixed by the beliefs and allocative bias of the regulator.

As the proposition shows, if the regulator did not have to consider the effect of its decisions on the strategic incentives of the outside interests, final policy would lie somewhere between the informed and the uninformed paternalistic compromise with outside interests. If modifying proposals were costless ($k = 0$), the regulator would surely implement the informed paternalistic compromise of $\alpha\mathbb{E}(p_i \mid a_i)$. The regulator always prefers to implement the informed compromise to one that relies only on prior beliefs. Yet the regulator's desire to enact this better policy is tempered by the the amendment costs. As $k \to \infty$, the best final policy converges on the initial proposal, which was the uninformed paternalistic compromise of $\alpha\mathbb{E}(p_i)$. For moderate values of k, the regulator will adjust policy from its prior toward an informed compromise as much as it can given the modifications costs. Therefore, I will refer to the outcome to the rulemaking game with exogenous commenting as the *affordable paternalistic compromise.*

These benchmakrs clarify the significance of the mathematical expressions that appear in the equilibrium conditions for the full rulemaking game. Detailed proofs that these conditions are necessary for equilibrium, and also that these conditions can always be satisfied, are left for the appendix

Proposition 3. *A PBE of the rulemaking game consists of a proposed policy x_p^*, a commenting decision rule a^*, a best guess z_i^*, and a final policy x_f^* satisfying the following four conditions:*

(1) Regulator proposes policy that deviates from the uninformed paternalistic compromise it would make without access to notice-and-comment. The greater the allocative bias in favor of outside interests, the more aggressively the regulator engages in such "strategic positioning." Conversly, the greater the modification costs, the smaller its strategic deviations:

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

(2) Regulator forms expectations about what preferences would lead an outside actor to abstain:

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

(3) Outside actor comments if the error in the regulator's expectation about abstention is worth correcting:

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

(4) Regulator adjust its proposal toward the informed paternalistic compromise, with the amount of adjustment depending on the modification costs.

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

This proposition is best explored from last to first, following the logic of backward induction. After commenting ends, the regulator chooses a policy based on some new information about public preferences. Condition (4) implies that the final proposal will be somewhere between (a) the policy it initially proposed and (b) the paternalistic compromise the regulator would most like to make given the information it has learned through notice and comment. As in Proposition 2, the amendment cost parameter k controls where the policy falls between the two alternatives. Yet the final policy selected in equilibrium is *not* the same as the affordable paternalistic compromise that arises if commenting is exogneous. In particular, if commenting is endogenous, the regulator's first move is to strategically deviate from the compromise it would make if it did not have access to commenting at all. As I elaborate in much greater detail in the next section, this adjustment helps make commenting more likely and abstention more informative. While the adjustment has fully served its purpose after comments are received, the regulator cannot walk it back because the proposal is sticky. In equilibrium, threfore, final policy represents a complex balancing act between the regulator's preconceptions about what outside interests want, its paternalistic view of how

important their preferences are, what it learns about those preferences during commenting, and strategic policy sacrifices that help the regulator learn more through the commenting process.

Before exploring the tradeoffs in commenting, it is worth considering what happens to equilibrium policy selection as modification costs become very large or very small. As amendent costs go to infinity, the equilibrium final (and proposed) policy converges on the expression for the uninformed paternalistic compromise. The reemergence of this expression suggests that the policy outcomes from rulemaking with very high modification costs should resemble the choices the policymaker would make without access to notice-and-comment at all. Conversly, as amendment costs vanish, the equilibrium final policy choice converges on the informed paternalistic compromise (and any proposal is admissable). In other words, the policy outcomes from the rulemaking game with low modification costs resembles the ones that would emerge if the regulator received commentary without making an initial proposal. The exact same limiting behavior is observed in Propostion 2, where commenting is assumed to be exogenous.

What is the significance of the observation that Propositions 2 and 3 are the same in the limit? A hasty interpretation might be that endogeneity does not matter for policy outcomes when amendment costs are extreme. This claim is subtly incorrect, because the posterior expectation $\mathbb{E}(p_i \mid a_i)$ stands for a different calculation in the two propositions. In particular, nature's assignment mechanism could give the regulator different information than would emerge endogenously if i followed its own incentives. Nevertheless, the formal similarity in the limits between the two propositions is telling. Namely, it shows that as modification costs are pushed to their ends, strategic positioning ceases to play a role in equilibrium policy selection. The regulator avoids strategically positioning in the case of $k = 0$ and $k = \infty$ for different reasons. If modification costs are infinitely large, then it is too costly for G to actually use information acquired during notice-and-comment. Therefore, the policymaker has no incentive to manipulate i's information provision incentives. Instead, G must follow through on its uninformed expectations to the bitter end. By contrast, if

modification costs are vanishing, then the regulator is not bound tight enough for strategic positioning to matter. Outside interests know that the regulator is unable to prevent itself from enacting the informed paternalistic compromise after receiving comments. No proposal can credibly change outside interests incentives. Only if modification costs are between the two extremes does the proposal have commitment value that the regulator can actually use.

Having elaborated the drivers of final policy selection under various permutations of the model, I turn to the incentives facing outside interest i in commenting. In particular, i has very limited means of influencing its policy utility. The costliness of policy modification k, allocative bias α, and the initial policy proposal x_p are all important factors affecting i, yet they are all beyond i's control. What i is able to influence is G's posterior expectations. If influence-seeker i comments ($a_i = 1$), then the regulator learns i's type immediately and with certainty: $\mathbb{E}(p_i \mid a_i = 1) = p_i$. If i does not comment, then the regulator will have to guess i's type. Denote this guess by z_i. A crucial observation is that the marginal incentive for i to comment depends entirely on the difference between z_i and p_i, which I will sometimes call the "inferential error." If individual i's type p_i is close to z_i so that the inferential error is small, then individual i will have little incentive to comment. Similar policy outcomes would follow from commenting or abstaining, so the marginal difference in policy cannot justify the costs. Conversely, if p_i and z_i are different enough, then there might be great benefits (or costs) to commenting as opposed to abstaining.

A corollary of the observation that inferential error determines the commenter's marginal incentive is that, potentially, abstention is informative about i's type. If the regulator wants to select the right policy in the final period, it must reverse-engineer i's decision-problem to understand the probability that i has a particular type p_i. Yet, somewhat circularly, the regulator's own expectations determine what those probabilities are. The Perfect Bayesian solution concept offers one take about how this "observer paradox" should resolve itself. Namely, the regulator and outside interest will converge on probabilities that are strategically stable. Indeed, Condition (2) of Proposition 3 asserts that in a PBE the regulator guesses correctly on average the type of the outside interest that abstains, while

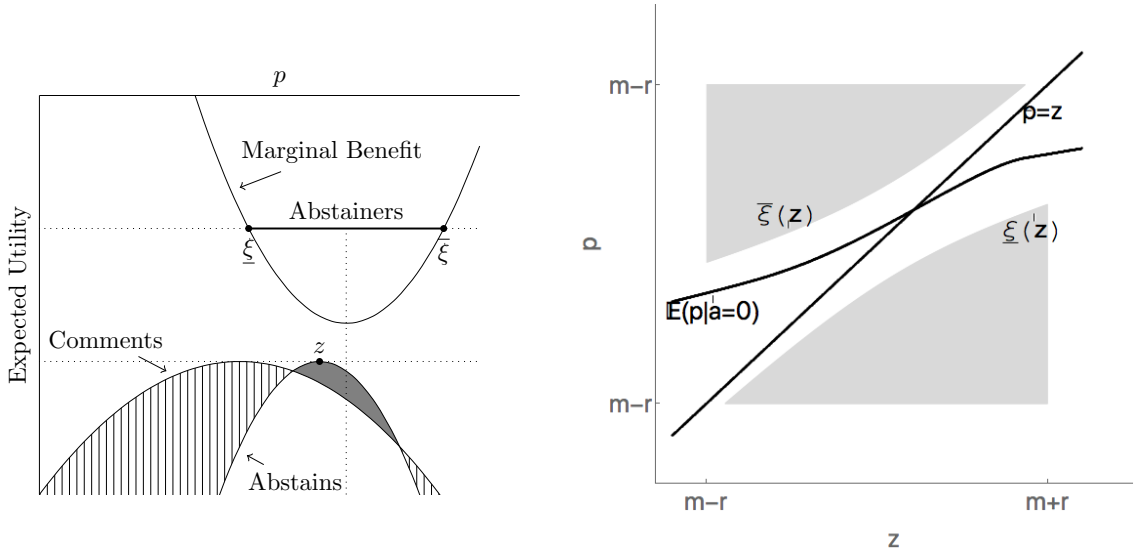

(A) i's decision problem portrayed for arbitrary p_i and fixed z_i.

(B) i's decision problem for arbitrary p_i and arbitrary z_i

FIGURE 2. Two ways of looking at the consistency of $\mathbb{E}(p_i \mid a_i = 0)$ and the commenting decision function $a_i(p_i)$.

Condition (3) says that i only comments if the marginal benefit of commenting exceed the cost, given what the regulator assumes about outside interests that abstain.

The trickiest part of establishing equilibrium existence in this model is showing that regulator's beliefs and i's commenting decision-rule can actually coincide, as they do in Figure 2(b). In this context, a crucial observation is that for any particular z_i, the marginal incentives are well-behaved with respect to p_i. Indeed, they are quadratic, as an informal argument shows. G's utility function implies that it wants to make a compromise between its independently preferred alternative and the competing preferences in the public. The more i's preferences diverge from G's independent concern of 0, the further away from 0 is the best compromise in G's view. For i, the "goldilocks" z_i is the one that leads G to think the appropriate compromise is p_i. If z_i were even more extreme than that, then G would over-compromise even for i's taste. Conversely, if z_i is not sufficiently extreme, then i will believe that the compromise does not go far enough. If p_i is sufficiently far away from z_i,

then, the benefit of commenting will eventually be greater than any fixed cost c.[8] Figure 2(a) provides a graphical illustration.

A corollary of the previous discussion is that if i's private type is between two cutoffs, call them $\underline{\xi}$ and $\overline{\xi}$, then the gains from correcting inferential error are too small to be worth the trouble of commenting. I shall refer to the space between $\underline{\xi}$ and $\overline{\xi}$ as the "abstention interval," since i abstains if its bliss point is between these cutoffs. If i's type is outside the abstention interval, then it will see commenting as cost-beneficial. Given a_i is such a well-behaved decision-rule that varies continuously with z_i, it is easy to see that the conditional expectation $\mathbb{E}(p_i \mid a_i = 0)$ will also be well-behaved and continuous with respect to z_i. The assumption that $p_i \in (m_i - r_i, m_i + r_i)$ allows for the application of powerful fixed point theorems, which establish that the intersection shown Figure 2(b) must always occur. At the same time, they generate one additional complication: boundary constraints. For example, the abstention interval may subsume the entire support of the distribution from which bliss points are drawn. In this case, one could hypothetically imagine an outside interest that would have the incentive to comment, but such an extreme p_i cannot actually occur in the game. Even so, there are only a few ways that the probability distribution's support $(m_i - r_i, m_i + r_i)$ can intersect with the abstention interval $(\underline{\xi}, \overline{\xi})$, as Figure 3 shows. Either i abstains for all possible types (two active constraints), i comments if its type is one side of the distribution (one active constraint), the other side of the distribution comments (the other active constraint), or abstainers have preferences between the commenters (unconstrained). These possibilities give the model flexibility to explain different patterns of participation in rulemaking, although they also show why some segment of the public will fail to participate despite their interest in the policymaking outcome.

Having addressed the joint selection-inference problem, it is possible to turn to G's initial policy choice. 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

[8]The logic is unchanged if we assume heterogenous costs c_i, which might reflect homogenous costs c but varying expressive benefits D_i to commenting, or heterogenous costs and benefits.

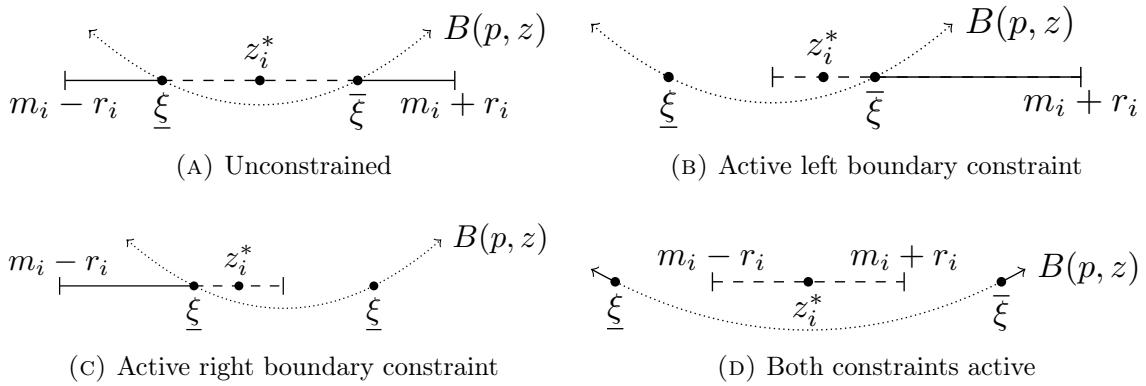

(A) Unconstrained

(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^*

competing considerations. On the one hand, it wants to avoid incurring future modification costs. Therefore, the regulator is inclined to propose the uninformed paternalistic compromise it would make if it did not have access to the notice-and-comment process. On the other hand, the regulator is aware of its dependency on the public for information. Final policy will necessarily reflect to some degree G's own expectations about the public's type. This creates a quandary for G. G can only create the "right" policy with certainty if the public fully communicates its preferences. Yet if p_i is close to z_i, then i will have limited incentive to comment and so G may have to take a guess in making final policy. If G could commit to inferring z_i very far from the support of the distribution from which p_i is drawn, then it could ensure full participation. Essentially, G would be threatening to punish abstention with extremely adverse policy. Yet under the model's assumptions, G has no way to commit to actually punishing abstention. Instead, G's own incentives will lead it to select z_i based on a rational analysis of i's incentives. Because i's incentives depend on the proposal x_p, however, the proposed rule gives G an indirect way of shaping particiaption and abstention for G's own benefit. In particular, Condition (1) describes when G would deviate in proposal from its natural inclination to make an uninformed paternalistic compromise. It will do so when deviating would bring G's posterior expectations about i's type closer to reality. How and why different choices of x_p bring about this better match, and what the implications

are for rulemaking as a whole, are easier seen through the examples presented in the next section.

Interpreting the Model

Strategic Positioning and Costly Rules It is widely believed that regulator's proposals during rulemaking are strategic. Some rulemaking scholars have suggested the primary audience for such "strategic" policymaking are the courts, or perhaps other actors with oversight responsibility such as the President or Congress. Yet even without oversight, strategic positioning can occur for purely informational reasons, as the model shows. What role strategic positioning plays in determining final policy is also debatable. According to one influential account (Elliott 1992), notice-and-comment is analogous to "kabuki theater," an exaggerated performance of policy change that inevitably results in final policy being exactly what the agency always wanted. For those that endorse this view, notice-and-comment only imposes transaction costs upon the agency. Notice-and-comment does not, in itself, meaningfully influence policy selection. The stakeholder-balancing model described here disputes this conclusion. Strategic positioning effects who participates and which final policies are selected. It has costs for the regulator and, potentially, the collective public. At the same time, strategic positioning does in expectation improve outcomes from the regulator's perspective, so such strategic positioning has no obvious normative interpretation.

Figure 4 illustrates some dimensions of the strategic positioning problem. In the left panel, the regulator's expected utility is graphed for various policy proposals assuming G and i subsequently play their best responses. The y-axis reflects the "natural benchmark" of the uninformed paternalistic compromise G would reach without access to notice and comment. The maximum of the expected utility curve does not occur on this reference line. Rather, there is a set of policy proposals that are in expectation as good or better for G. Why are these proposals better? The right panel delivers the key insight. The Figure plots G's posterior belief about i's bliss point against the truth assuming each follows their equilibrium strategy. For the moment, assume that the regulator gives no weight to its independent preference and modification cost are zero. Under these assumption, the right

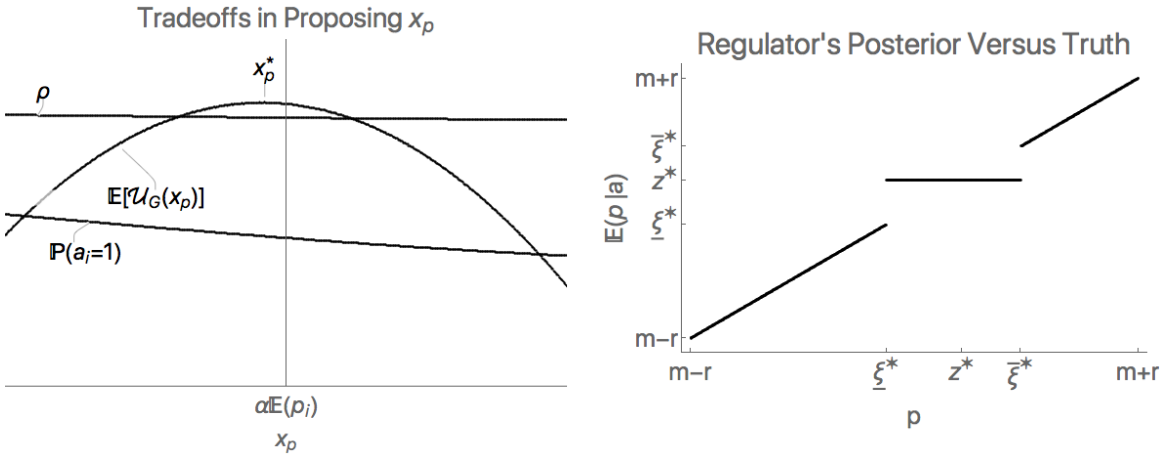

(A) (B)

FIGURE 4. The optimal proposal deviates from the ex ante expected rule in order to increase the covariance between the posterior and truth. Covariance is improved when the abstention interval is narrower and less probable.

panel is also a graph of final policy x_f against the bliss policy of i. Outside of the abstention interval, G's posterior/final policy are perfectly correlated with the true bliss point. Within the interval, they are perfectly uncorrelated. Perhaps counter-intuitively, this pattern of perfect correlation of final policy and i outside the abstention interval, no correlation within the abstention interval, holds under very weak assumptions about allocative bias. Provided G is willing to give even a smidgen of responsiveness ($\alpha > 0$) to the preferences of the outside interest, policy will always move in the direction G expects i's preferences to lie. Yet if i abstains, then G can only respond to its best approximation. Inevitably, this approximation results in no correlation between final policy and i's bliss point.

The lack of agreement between final policy and i's preferences lead to utility losses for the regulator, and open up the possibility that alternative proposals may lead to utility gains on average. Figure 4(a) shows not only the expected utility for G from various proposals, but also shows how different proposed policies impact the correlation between the regulator's posterior expectations and reality. Although the figure does not show it, the expected bliss point in 4(a) is somewhere to the right of the y-axis and the regulator's independent object

of concern is somewhere to the left. The downward slope of the correlation line shows that the agreement between posteiror and reality gets as the proposal approaches i's bliss policy. Conversely, the correlation is better the more slanted the proposal is toward G's independent preferencs. There are two reasons why this occurs. First, more adverse policies increase the probability that i will choose to reveal its preferences, as Figure 4(a) also shows. Second, adversarial postures may decrease the *width* of the abstention interval.[9] In most familiar probability distributions, shorter intervals tend to have lower probabilities, so these factors go together. Potentially, however, these factors might oppose. If one were to give the regulator the choice of two differently sized intervals which gave the same overall expected level of commenting, the regulator will always prefer the narrower one. The reason is that if the interval is narrower, the best guess z_i will more closely approximate p_i, so the overall correlation across all possible values of p_i will be tighter.

The analysis becomes more subtle when more outside actors are introduced. In that case, the optimal policy proposal is

$$\sum_{i=1}^{n} \alpha_i \mathbb{E}(p_i) + \frac{1}{2k} \sum_{i=1}^{n} \alpha_i^2 \left\{ \frac{\partial}{\partial x_p} \mathrm{Cov}(p_i, \mathbb{E}(p_i \mid a)) \right\} \tag{1}$$

As mentioned in the last section, the optimal proposal is defined with reference to the multi-actor version of the uninformed paternalistic compromise: $\sum_{i=1}^{n} \alpha_i \mathbb{E}(p_i)$. In the n outside actor case, hoewver, strategic positioning is qualitatively different because G is cross-pressured. It cannot take an adversarial posture with respect to all outside interests at once. Which of the many outside interests will have the most influence in driving the deviation from the uninformed paternalistic compromise? Paradoxically, those with high α_i, which is to say those that are most favored by G's allocative bias. Perhaps surprisingly, notice-and-comment *does* bias outcomes, on average. All else equal, however, it biases policy *against* the interests the regulator weighs the most heavily. As we shall see, however, notice-and-comment creates a surplus by decreasing the mismatch between the regulator's expectations and reality. Those with higher α tend to reap more of this surplus.

[9]Figure 2(b) shows, relatedly, that different choices of z_i would result in different width intervals.

The remarks of the last paragraph require two caveats. First, the prior distributions of bliss points $\{\mathcal{F}_i\}_{i=1}^n$ are an important confounder. Even if allocative bias completely favors one interest, strategic positioning is not observed unless it can actually increase the agreement between the regulator's posterior beliefs about public preferences and the reality. Conversely, suppose the inferences that the regulator can make about policy can be improved by taking an adversarial position with respect to one group about which the regulator cares only a little. If these gains are great enough, the regulator will strategically position with respect to that group despite limited concern for that actor's preferences. Secondly, strategic-position does not depend on oversight in this model, but it does require that G be risk averse. The regulator is willing to permanently give up some amount of utility through strategic positioning that will help it avoid bigger losses in final policy. A risk neutral regulator would not be inclined to make this tradeoff.

Despite these caveats, Equation (1) suggests an important comparative static that clarifies intuitions about how strategic positioning relates to the costs of modifying rules. If amendment costs are decreased, then the policy proposal becomes a less effective commitment device. Yet G views proposed policies as a valuable instrument for committing to make different, and more privately useful inferences when it comes time to make final policy. As a result, the regulator compensates for lowered modification costs by making increasingly extreme initial policy selections. Decreasing the unit costs of modifying rules does not bring a corresponding decrease in the modification costs that the regulator incurs. Conversely, if modification costs are increased, one should see less strategic positioning. The reason is that the policy choice x_p is a strategic investment that is harder to walk back.

Adverse Policies and Adverse Priors An important insight of the formal analysis is that the incentive to comment primarily derives from correcting some mistaken belief on the part of the regulator. The likelihood that the regulator will issue an adverse policy is not sufficient, on its own, to mobilize commenters. This is not to say that adverse policy does not matter. The larger the degree of conflict between G and i, the more attractive even marginal improvements will be to i. If what i wants is already pretty well known by the regulator,

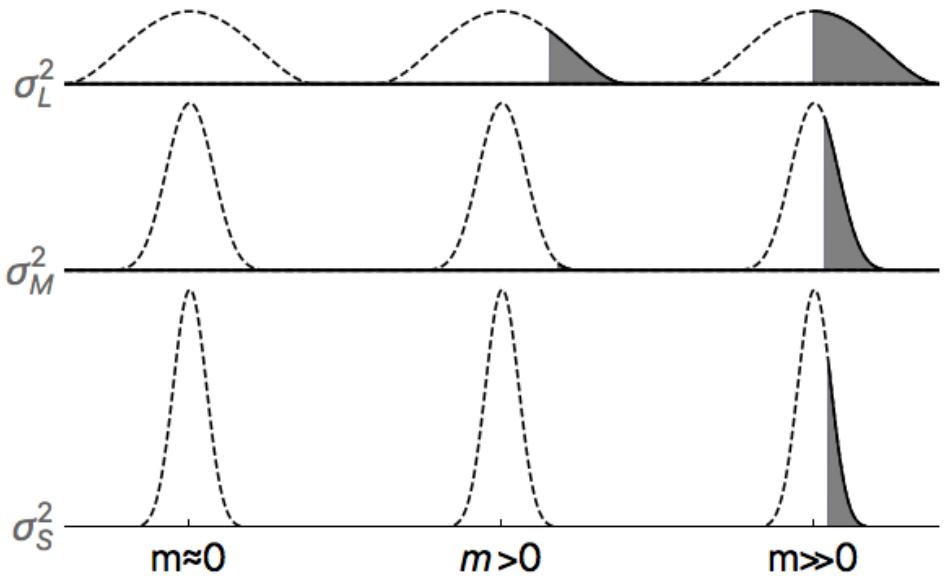

FIGURE 5. Prior preferences are increasingly unpredictable along the y-axis ($\sigma_L^2 > \sigma_M^2 > \sigma_S^2$), while preferential conflict between i and G is greater along the x-axis. Shading indicates types p_i which, if drawn, would comment.

however, then the marginal benefit of commenting is necessarily small. Predictability of preferences is a key obstacle to mobilization in a two-sided model where preferences are the only information of value that stakeholders can share.

Figure 5 illustrates the probability of observing a comment in the model's equilibrium for various distributions of i's bliss point. Going upward, the distributions are increasingly diffuse and unpredictable. Moving rightward, the degree of preferential conflict between the outside interest and the regulator are increasingly large. Each of the distributions only has positive probability within an interval of identical width. As the Figure shows, if preference conflict between the regulator and the outside interest is negligible, then mobilization cannot occur in equilibrium regardless of the unpredictability of preferences. If preference conflict is moderate, however, then the cases are distinguishable. Given this level of conflict, there is no chance of commenting when predictability is high, a small chance of commenting when predictability is moderate, and a fair chance of commenting when predictability is low. If preference conflict is great, then the odds of commenting go from fair to probable as unpredictability increases. Generally speaking, for any level of preference conflict betweent

the outside interest and the policymaker, then if what i wants is sufficiently predictable, participation will be infeasible. The converse is also true.

On a theoretical level, this observation about "concentrated" and "diffuse" priors recalls in some ways classic intuitions from the public choice literature (Olson 1971). Olson argues that goods with diffuse benefits and concentrated costs will tend to be under provided by governments, because concentrated groups will have less temptation to free-ride. Similar, large n free-riding dynamics are observable in variations on this model.[10] Even if $n = 1$, however, the model shows an additional concern. Potential influence-seekers are not only tempted to free-ride on one another's costly investments in changing policy, they are also tempted to free-ride on the policymaker's beliefs. "Epistemic" free-riding is easier to justify for the influence-seeker whose preferences are predictable than it is for the one whose preferences are hard to know ex ante. "Concentrated" in this sense is bad for mobilization, it turns out, while "diffuse" is good. In the regulatory context, larger groups such as consumers are often affected more homogeneously and predictably than smaller groups such as regulated industries. As a result, the classical analysis and this model would tend to coincide about which groups have mobilization advantages in rulemaking and which do not.

At the same time, this model provides an alternative view to the classical public choice literature on the policy causes and consequences of mobilization failure. For Olson, mobilization failures provide observable evidence that policy selection deviates significantly from the social optimum. As the next section shows, however, mobilization failure in this model is observationally equivalent with a range of policy outcomes and directions of regulatory bias. While mobilization surely brings some policy benefits to those who manage it, these benefits are highly contingent on the regulator's allocative bias. Moreover, this allocative bias can lead to policy that favors certain outside interests even if they do not make the same costly investments in influencing policy.

[10]For example, let $\Gamma(i)$ assign each outside actor i to one of j "groups" and suppose nature imbues the regulator with a group-specific allocative bias α_j. Suppose further that G's concern for each group member is equal. Then this is equivalent to the baseline model where $\alpha_i = \frac{\alpha_j}{n_j}$ and $n_j = \sum_{\Gamma(i)=j} 1$. If α_j is fixed, then the larger n_j is, the less mobilization will occur, effectively because it implies a smaller α_i.

Interpreting Rulemaking Data

One of the unique strengths of the stakeholder-balancing model described here is that it readily lends itself to interpreting descriptive facts about who gets what, when, and how during the rulemaking process. Yackee (2019) describes dozens of such findings on the notice-and-comment process, a substantial proportion of them published in just the last several years. The sub-literature appears especially poised to benefit from the increasing availability of sophisticated computational tools for analyzing rulemaking data (Libgober and Carpenter 2017) Yet theoretical progress has been slower, which also hinders empirical research. Selection issues are often acknowledged, but seldom addressed. Similar questions are explored in different contexts without clear theoretical expectations of why the results should vary or what the significance of such differences would be. In this context, the stakeholder-balancing model provides resources for explaining and organizing research findings. As I argue here, empirical regularities documented by the literature are consistent with only *some* equilibria of this model, which gives the model explanatory and interpretive power. At the same time, it also suggests reasons for caution. With regard to the key allocative bias quantity α, the existing facts are observationally equivalent. New, theoretically-informed data collection strategies are necessary to push the empirical literature forward.

Table 1 presents some important stylized facts about rulemaking as described by Yackee (2019), Kerwin, Furlong, and West (2011), and others. This scholarship has shown that rulemaking is a dynamic and highly variable site of interest group conflict, where important policy gains are won and lost. These patterns do not fit easily with earlier, informal descriptions of notice-and-comment as beleaguered or only superficially a site of significant policy contestation (Harter 1982; Elliott 1992). Yet they do fit naturally with the persuasive signaling framework. The stakes of rulemaking arise from the presence of preferential conflict between the regulator and external influence-seekers. Depending on the issue, preferential conflict could be small or large. Policy change from proposal to final rule occurs in the model because of learning about stakeholder preferences, which can occur with or without actual commenter participation. Indeed, why rules might change given no commenter participation

Fact	Source	Model-consistent Explanations
Policy selection responds, at least marginally, to comments.	(Golden 1998; Balla 1998; Cuellar 2005; Furlong and Kerwin 2005; Yackee 2006)	Agencies are not completely zealous.
No group completely dominates participation or gets everything they ask for	(Fritschler 1975; Chubb 1983; Nixon, Howard, and DeWitt 2002; Kamieniecki 2006)	Agencies are not completely captured.
Many rules receive no comments, some receive many comments	(Golden 1998; West 2004; Coglianese 2006; Balla and Daniels 2007)	• The policy stakes for rules with more comments are higher. • The preferences of stakeholders for rules with more comments are less predictable. • Agencies have higher concern for stakeholder preferences in rules with more comments.
Firms and interest groups participate more than citizens.	(Golden 1998; Yackee and Yackee 2006; Kerwin and Furlong 2011)	• Allocative bias favors firms and interests. • Citizens have preferences that are more predictable.
Major policy change is rare during rulemaking.	(Golden 1998; West 2004; Yackee 2006)	• Modification costs are high. • Agencies are zealous. • Allocative bias cancels out inequalities in mobilization. • Agencies are reacting to abstention.
Policy change is biased in favor of business interests, especially when unopposed.	(Yackee and Yackee 2006; Haeder and Yackee 2015)	• Agencies are biased in favored of business. • Agencies preferences are aligned with non-business interests. • Non-business interests have predictable preferences.

TABLE 1. Stylized facts and non-exclusive model consistent explanations.

is hard to square theoretically with other accounts. Moreover, heterogeneity in patterns of participation and policy change across rules arises, under the model, as a result of variation in the underlying stakeholder preferences regarding the issues at stake. Potentially, this

issue-specific variation is also increased by variation in the allocative bias α. Allocative bias surely varies across agencies. It is also likely to vary within agencies across rules.

Under the model, these stylized facts eliminate the possibility of two significant edge cases: agencies are not total zealots ($\alpha \neq 0$), nor are they completely captured ($\alpha \neq 1$). Intuitively, if allocative bias were at either extreme, one could not observe the patterns of commenter participation or policy change that rulemaking scholars often do. If the agency is a total zealot, then it proposes and finalizes its most preferred policy. The reason is that the regulator sees no benefit in satisfying the preferences of any of its other stakeholders. As a result, influence-seekers should see no marginal incentive to invest in commenting, which is costly. Similarly, if the regulator is completely beholden to one outside interest i's preferences, then there is no marginal incentive for any other influence-seeker j to submit a comment. From j's perspective, the completely captured and totally zealous regulator look alike. In the case of complete capture, policy might move in a direction more or less favorable to j through rulemaking, but j should realize that its own actions have nothing to do with it. A second problem with the strong capture equilibrium is that, in this equilibrium, policy should always change in i's favor whenever i comments. The existing literature suggests that on average certain interests might win more than others, but policy change during rulemaking has never been found to favor one interest or firm exclusively.

Other well-known empirical regularities also arise under the model, but the observed evidence is potentially multi-causal. Take, for example, studies demonstrating that mobilization is biased during rulemaking and that certain groups tend to win the lion's share of policy change. In their exemplary study, Yackee and Yackee (2006) argue that their findings raise questions about whether "the interests of the broader public are furthered in agency rulemaking." The question is particularly distressing given the vastly important policies being decided in this venue. The model confirms the reason for concern, but suggests the possibility that there are more benign explanations (Carpenter 2004). If one influence-seeker is more favored by an agency during rulemaking than another, then the favored actor is more likely to mobilize and receive better policy outcomes, assuming all else is equal between the two

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

TABLE 2. Numerical Example. The vertical line separates quantities that empirical scholars have reproducibly measured (right) from those that are unobservable using existing methods (left).

outside actors. Of course it is not actually clear that all else *is* equal. The prior distribution of preferences can make it so that the disfavored actor actually outperforms the favored one in terms of mobilization and policy change. Yet policy will, unsurprisingly, remain largely biased toward whichever actor the regulator favors more.

A numerical example helps illustrate the point. [11] Suppose there are two stakeholders, a supplier of a good s and a buyer b. Each interest's bliss points come from their own prior distributions, in particular $p_s \sim \text{Uniform}(m_s - r_s, m_s + r_s)$ and $p_b \sim \text{Uniform}(m_b - r_b, m_b + r_b)$. To avoid clouding the example with issues relating to preference alignment, which also can lead to observational equivalency, suppose that $m_b = -1$ and $m_s = 1$, so that the position of the firms with respect to G is symmetrical. Clearly, negative policies typically favor buyers while positive policies typically favor suppliers. Key quantities of interest are collected in Table 2. These include the expected participation rates for s and b, as well as the position where policy is proposed, finalized, and how policy changes, for various choices of $\vec{\alpha} =< \alpha_0, \alpha_s, \alpha_b >$. I focus on two cases, one where G heavily favors suppliers ($\vec{\alpha}_s =< 0, 0.1, 0.9 >$) and one where G heavily favors buyers ($\vec{\alpha}_b =< 0, 0.9, 0.1 >$). Viewing the derived formulas for the participation rate, it becomes clear that if the preferences of buyers are sufficiently predictable ex ante (e.g., in this case, $r_b < 1$), then the probability of participation by buyers is actually 0 for both, very different levels of allocative bias. Depending on whether G has $\vec{\alpha}_s$ or $\vec{\alpha}_b$, one can also find a level of predictability in supplier preferences so that the probability s participates is ≈ 0.4. Let us assume that these levels of predictability are suitably chosen. Then consider what happens if G is biased in favor of suppliers ($\vec{\alpha}_s$). Policy typically starts out positive and becomes more so, favoring s. If G is biased in favor of buyers ($\vec{\alpha}_b$), then the

[11] The appendix presents more details on these calculations.

proposed and final policies are in expectation both negative. Nevertheless, policy usually changes in the suppliers direction. For scholars who only observe the direction of policy change and mobilization rates, these two cases look identical, despite the radical difference in regulatory bias being represented.

Mobilization and policy change are observationally equivalent for the key allocative bias parameter. Yet these quantities represent something important on their own. In the example, the suppliers did better through notice-and-comment than they would have without it, while the buyers did worse. Importantly, no matter G's bias, the policy surplus notice-and-comment creates is reaped by the ones that can potentially mobilize. The word "potentially" in the last sentence is important. Indeed, *commenting does not actually need to occur* for s to reap some of the benefits of notice-and-comment, which b never is able to do. The reason is that if G does not believe commenting makes sense for any buyer ex ante, G regards a buyer's failure to comment as uninformative. Rationally, G cannot justify any changes from its prior view of the best compromise. By contrast, if G considers participation by s feasible ex ante, then s's failure to comment *is* informative. Rationally, G will need to make an accommodation. The patterns of policy change described in the example can occur during rulemaking without parties needing to take *any* observable, costly actions.

Importantly, this analysis is descriptive rather than normative, although the normative implications are intuitive. If the regulator's allocative bias is socially optimal way, then G only moves policy if it enhances overall welfare. In this case, the gains notice-and-comment make possible for s but not b are welfare improving. Similarly, if the regulator's allocative bias is such that it places less weight on s than it normatively should, then inordinate opportunities for s to make gains through notice-and-comment can compensate for the regulator's inherent bias. Potentially, this is also to society's benefit. On the other hand, if the regulator is inappropriately biased toward the influence-seekers that can potentially mobilize in the first place, then notice-and-comment will further exacerbate bias.

Other regularities described by the literature also have multiple explanations consistent with the model, however the more robust description it provides helps make sense of

certain puzzles. For example, some rules receive few comments, while others receive many comments. According to regulations.gov, the majority of rules receive no comments.[12] One interpretation of this fact is that some rules are "significant" and "controversial" while others are more "technical." While some rules surely raise few distributive issues, commenting volume and importance do not move in lock-step. Congress has seldom exercised its powers over rulemaking under the Congressional Review Act, however one of the few rules controversial enough to achieve this distinction had only sixteen comments ("Unemployment Insurance Drug Testing Rule," 81 FR 50298). The model informs us that "significance" or "controversy," understood as differences in preferences between influence-seekers and the regulator G, are but one ingredient necessary for commenter mobilization. Predictability of preferences are also important, as is the degree of allocative bias possed by the regulator. Rather than view rules without comments as settling policy matters without much consequence, one should instead believe that the level of controversy was insufficient given the regulator's bias, the potential information value of comments from stakeholders, and the costs of commenting.[13] Similarly, the finding that major policy change is rare during rulemaking does not necessarily imply that agencies are "overzealous" in their preferences. Rather, it could also mean that the costs of modification are relatively high or that the agency learned conflicting information about stakeholder preferences through the notice-and-comment period, which forced the agency to remain closely tethered to its initial proposal. Alternatively, it could mean that interest group preferences are usually quite predictable.

Conclusion: Directions for Future Research on Rulemaking

Who wins and who loses in the making of regulation is a significant question for scholars in political science, law, and public administration. Learning about the mechanisms driving the process on the basis of observable evidence is hard, however, because regulators

[12]The appendix shows the distribution of comments on rules, as well as a table of comments on rules vetoed via the Congressional Review Act (CRA).

[13]If costs are heterogenous across actors, for example because of varying expressive benefits, the problems the model describes are only exacerbated

and influence-seekers are strategic. In particular, regulators anticipate what information outside actors can provide. In turn, outside actors anticipate the costs and benefits of abstention versus participation. Correctly reading outputs of rulemaking as evidence of regulatory bias require a comprehensive view of the "chessboard" that notice-and-comment represents. Nevertheless, the literature offers few formal characterizations of what that board looks like. This model of notice-and-comment as a persuasion game shows that mobilization rates and directions of policy change are directly informative about regulatory bias. Yet present data collection strategies have only produced evidence that can reject two edge cases: strong capture and strong public-interest zealotry. Additional progress for the empirical literature will require new researcher strategies to overcome or circumnavigate this observational equivalency.

Under the model presented here, the major confounding variable is uncertainty about the shape of prior preferences of the influence-seekers, especially the "unpredictability" or variance in these preferences. Measuring the preferences of potential influence seekers is not an impossible research task. For example, surveys of commenters have proven a useful source of data for empirical rulemaking scholars (Furlong 1997; Furlong and Kerwin 2005; Yackee 2012). These methods have not been used to define the scope of preferential conflict at stake in particular rulemakings. Survey-based strategies will generally require planning and foresight, because survey respondents may forget what their positions on an issue were. Another plausible research strategy is to use ideal-point measures (Chen and Johnson 2015; Richardson, Clinton, and Lewis 2017). Generic application of partisan measures is ill-advised, since interest group conflict during rulemaking is not always partisan. Nevertheless, in certain cases one would expect partisan alignment to coincide with the true dimension of interest group conflict. The regulation of "bump stocks" is an example where preferential conflict during rulemaking might conform to partisan conflict. Conjointly positioning regulators and outside interests in the policy space would enable two key contributions. First, it would test the model's predictions regarding the tendency of "extremists" to comment and those with

more central preferences to abstain. Second, one could use these measures to control for the two-sided selection bias problems with regard to mobilization and direction of policy change.

Analysis of the model also suggested other empirical tasks that would have smaller identification issues. The model predicts that aggregate commenting levels will respond to changes in modification costs, in particular when costs of modification increase then commenting volume goes down. Modification costs are plausibly related to technical complexity or duration of rule-development, and so observationally one could look for the model's expected relationship between modification costs and commenting or degree of policy change. In general, one would have a hard time justifying the necessary assumptions about prior preferences to discount confounding. Yet in some policy domains, one might plausibly assume preference conflict is homogenous from year to year. Alternatively, one could collect pairs of rules where there is good reason to think that the conflict between stakeholders is balanced, for example because the same statutory provisions are addressed. Regression discontinuity designs to test the model's implications might focus on shocks to modification costs that are plausibly exogenous to sampled rules, for example the introduction of new procedural requirements on rulemaking throughout the entire regulatory state, or changes in the staffing of oversight bodies due to unforeseen retirements. The existence of policy change following complete abstention by all stakeholders is hard to fathom from most informal models of the notice-and-comment process. Among all possible theories of the rulemaking process, this model uniquely rationalizes policy movement in the absence of commenting, so exploration of these rules could yield a falsification test of the model.

Besides more empirical work, more theoretical models of the rulemaking process are also necessary. Position-taking is an important aspect of what makes influence-seekers persuasive, however other factors are also important. The literature on bureaucratic politics informs us that many factors influence agency's choices, including regulatory networks, reputations, staffing, insulation, dynamics of sub-delegation, and oversight by courts and Congress. How each of these factors influence the dynamics of policy selection during notice-and-comment is hard to know, and has not received much formal treatment. Similarly, the role

of ad hoc coalitions and organized trade associations appears significant during rulemaking. The theoretical terrain is open. There remains much to explore.

References

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.

Boehmke, Frederick J., Sean Gailmard, and John W. Patty. 2013. "Business as usual: interest group access and representation across policy-making venues." *Journal of Public Policy* 33:3–33.

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

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.

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.

Chubb, John E. 1983. *Interest Groups and the Bureaucracy*. Stanford, CA: Stanford University Press.

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.

Coglianese, Cary. 2006. "Citizen Participation in Rulemaking: Past, Present, and Future." *Duke Law Journal* 55 (5): 943–968.

— . 2018. *Let's Be Real About Trump's First Year in Regulation.*

Coleman, W. 2016. "How Cheap Is Corporate Talk? Comparing Companies' Comments on Regulations With Their Securities Disclosures." *Harvard Environmental Law Review* 40 (1): 10681–10688.

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

Dash, Eric. 2011. Feasting on Paperwork.

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

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

Fritschler, A. Lee. 1975. *Smoking and Politics*. Englewood Cliffs, NJ: Prentice-Hall.

Furlong, S. R. 1997. "Interest Group Influence on Rule Making." *Administration & Society* 29, no. 3 (): 325–347.

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. 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." *American Journal of Political Science* 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.

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.

Harter, Philip J. 1982. "Negotiating Regulations: A Cure for Malaise." *Georgetown Law Journal* 71 (1): 1–118.

Kamieniecki, Sheldon. 2006. *Corporate America and Environmental Policy*. Stanford, CA: Stanford Law / Politics.

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

Kerwin, Cornelius, Scott R. Furlong, and William West. 2011. "Interest Groups, Rulemaking, and American Bureaucracy." *The Oxford Handbook of American Bureaucracy*, no. December 2017: 1–24.

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

Libgober, Brian, and Daniel P. Carpenter. 2017. "Lawyers and Lobbyists."

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.

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

Nelson, David, and Susan Webb Yackee. 2012. "Lobbying Coalitions and Government Policy Change: An Analysis of Federal Agency Rulemaking." *The Journal of Politics* 74 (2): 339–353.

Nixon, David C., Robert M. Howard, and Jeff R. DeWitt. 2002. "With Friends Like These: Rule-Making Comment Submissions to the Securities and Exchange Commission." *Journal of Public Administration Research and Theory* 12 (1): 59–76.

Olson, Mancur. 1971. *The Logic of Collective Action*. Cambridge: Harvard University Press.

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.

Turner, Ian R. 2017. "Working smart and hard? Agency effort, judicial review, and policy precision." *Journal of Theoretical Politics* 29 (1): 69–96.

Wagner, Wendy. 2012. "Revisiting the Impact of Judicial Review of Agency Rulemakings: An Empirical Investigation." *William & Mary Law Review* 53:1717.

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

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 rulemaking "ossified"?" *Journal of Public Administration Research and Theory* 20 (2): 261–282.

Yackee, Susan Webb. 2006. "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.

— . 2012. "The politics of ex parte lobbying: Pre-proposal agenda building and blocking during agency rulemaking." *Journal of Public Administration Research and Theory.*

— . 2019. "The Politics of Rulemaking." *Annual Review of Political Science* 22.

You, Hye Young. 2017. "Ex Post Lobbying." *Journal of Politics* 79 (4): 1162–1176.