Hard Facts or Cheap Talk? Strategic Communication and Policy Change in Regulatory Rulemaking *

Brad Hackinen Giuseppe Carenini Jean Tirole Francesco Trebbi

November 2025

Abstract

This paper presents a theory and an empirical investigation of strategic communication in rulemaking, the process through which government agencies make regulations. It highlights the different empirical implications of models based on hard (verifiable), soft (cheap talk), and hybrid information in shaping agency policy. Using thousands of rules and millions of pages of public comments directed at regulatory agencies and recorded in a large U.S. government repository, we show that theories based on purely hard or purely soft information fail to match key moments in the data. Hybrid theories, based on combinations of hard and soft information, show promise.

JEL Codes: D72, P0, P16
Keywords: Government Regulation, Rulemaking, Notice and Comment, Cheap Talk vs Hard Information.

*Hackinen: Ivey School of Business, Western (bhackinen@ivey.ca). Carenini: Dept of Computer Science, University of British Columbia (carenini@cs.ubc.ca). Tirole: Toulouse School of Economics (jean.tirole@tse-fr.eu). Trebbi: Haas School of Business, UC Berkeley and National Bureau of Economic Research (ftrebbi@berkeley.edu). Early conversations with Marianne Bertrand, Matilde Bombardini, and Ray Fisman informed this project in many of its empirical elements. We are grateful to participants at seminars at University of Waterloo, and UC Berkeley for comments and valuable suggestions. Kourosh Khounsari, Linzi Xing and Amirhossein Abaskohi provided excellent research assistance. We thank Barbara Clark, Taylor Feltham, Dillon Gohil, Marie Liang, Nikolai Racicot, Rebekah Rodrigues, and Yuka Takami for their dedicated assistance with annotating rulemaking documents. An early version of the text was proofread with Refine.ink and we thank Ben Golub for his help. We acknowledge support from the Social Sciences and Humanities Research Council of Canada. Trebbi acknowledges support from the Smith Richardson Foundation grant #2022-2935.

1 Introduction

Policy making in modern democracies originates from executive and legislative branch activity and takes shape through the passage of statutes and implementation of associated rules. While statutes emanate from elected branches of government, regulations are directives and guidelines that have the force of law, but produced by agencies operating outside voter oversight. Precisely because regulation lacks the accountability and legitimacy derived from elections, but carries the force of law, United States statutes, such as the Administrative Procedure Act (APA) of 1946, aim at offering a compromise between safeguarding citizens against "*arbitrary and capricious*" government behavior and allowing agencies to apply subject-matter expertise to complex policymaking domains.[1] In U.S. federal rulemaking –the context that we study in our empirical analysis– consultation with the public is achieved through a process that can be reduced to three broad phases: (i) a notice of proposed rulemaking (NPRM), where a rule and its rationale are presented; (ii) a period of public comment; and (iii) rule finalization, which includes taking into consideration and responding to comments deemed relevant and substantive. Notwithstanding efforts to improve the transparency and accountability of the rulemaking process, regulatory choice is considered opaque in public opinion surveys and perceived as captive to special interest influence. In the words of Peltzman (2022), "*...ordinary citizens are skeptical and wary. They want less regulation and do not trust regulators to do what is right. The mistrust has become stronger over time.*" Questions about the nature and quality of the information exchanged in the regulatory process are debated and the role of strategic advocacy in rulemaking remains unclear: How responsive are regulators to the public or to special interests? How often do public comments result in policy change? Are there political influences/distortions evident in making rules? Are all commenters equally influential or do some have the ear of regulators more than others? An abundance of communication is generated between the public and the government throughout the *notice and comment* process, offering an ideal setting for the study of the economics of information among sophisticated agents. Yet, what the nature and the systematic consequences of this information exchange are is not firmly established (O'Connell, 2008).

This paper investigates, both theoretically and empirically, the nature of strategic communication and policy change in rulemaking: whether it is generally based on hard and verifi-

[1] The process of *notice and comment*, a mechanism of direct interaction between the public and regulatory agencies, is not unique to the American system. The European Commission of the European Union, for example, requires public consultation of stakeholders for directives and regulations (EU Commission, 2012), and it is required by law to organize "*broad consultations with parties concerned*". The government of Canada also offers a process of consultation with stakeholders.

able information or on soft and cheap talk, how issues of asymmetrical advocacy arise, who commands the attention of the regulator, and whether agencies are more likely to make policy changes in response to comments from larger organizations or from ideologically extreme advocates. Our objective is to contribute to the debate on regulation by presenting a theory and new facts about the nature of communication between interest groups and regulators. Ultimately, the combination of theoretical and empirical results in this paper paints a picture of the process of rulemaking where neither hard and verifiable information nor cheap talk and soft communication can alone explain all the data. Hybrid models, like the soft-hard information framework that we develop in Section 2.4, present more promise.

Our empirical analysis is possible because of the wealth of data assembled by U.S. regulatory agencies about the process of notice and comment regulation under the APA. We analyze all U.S. federal rules published between 2008 and 2022 on the government repository Regulations.gov and the majority of comments on these rules. This requires the analysis of millions of pages of text of regulations and public comments. We develop custom Natural Language Processing (NLP) tools for measuring policy changes related to individual comments and organizations that comment on rules, building on previous approaches to analyzing influence (see Xing et al. 2023; Bertrand et al. 2021) to measure commenting outcomes in detail. Careful NLP customization is necessary, as regulatory documents are among the most technical and impervious of text corpora. We parse rules to extract specific regulatory responses on narrow topics raised by commenters, parse response texts to detect mentions of policy changes and commenter stance in these responses, and identify which comments the agency is likely addressing in each response. The combined data allows us to measure influence in terms of the number of responses and each organization's success rate in obtaining a desired policy change over our sample period. It also allows us to quantify how responsive U.S. agencies are to public consultation in accordance to the APA, and regulators' heterogeneity in heeding to comments.[2]

We uncover a set of new facts that were previously undocumented or, in some cases, more narrowly quantified: *#1 Comments on regulations tend to be negative.* 75% of comment letters express opposition to the proposed rule and 90% want at least one change made to the proposed rule. *#2 A sizable share of comments contains detailed, verifiable information.* 35% of comment letters authored by organizations are at least 2 pages long, and 40% of these longer

[2] Let us clarify at the onset that, while most of the facts reported in this article are novel, some have appeared in more limited samples and/or for restricted time periods within the Law and/or Political Science literature. These fields are historically more attentive research areas to the subject matter of rulemaking and notice-and-comment processes (e.g. Yackee 2006; Yackee and Yackee 2006) than Economics, but also lack the set of large-scale regularities and methodologies that we propose. We try to integrate to the best of our ability these important interdisciplinary contributions and hope to amplify their impact within Economics.

comments contains citations or links to independent sources. *#3 Agencies receive requests for change by both sides of the political spectrum.* For groups for which political donations data is available, only 55% of comments requesting change come from organizations that are ideologically opposed to the administration. *#4 Politically aligned advocates have higher influence.* When organizations are politically aligned with the president, they comment 15% less often, but are twice as likely to express support for a proposed rule, and have 20% higher commenting success rate when they choose to comment. *#5 Ideologically moderate advocates have higher influence.* Left-leaning nonprofits are frequent commenters, but industry groups and other center-right organizations have 5-20% higher success rates when they choose to comment. *#6 Larger advocates comment more, but not with higher success rates.* Large public companies comment considerably more (both in terms of number of comment letters filed and in the length of each comment) than small public companies. However, large companies do not have a higher policy change success rate, conditional on sending a comment, relative to small companies. So, while large organizations drive the vast majority of changes in the data, this is determined by their much higher volume of commenting.

The theoretical approach in this paper focuses on presenting a flexible framework of strategic communication in rulemaking, with an emphasis on distinguishing between models of hard information and cheap talk.[3] We begin by presenting a core framework with two variants that differ only in whether the commenters can commit to truthful information (hard and verifiable information) or not (cheap talk), plus a third hybrid structure that mixes both forms of communication.

Our core framework is constructed as follows: An agency is designing a specific feature of a rule, which we model as a binary choice $a \in \{0, 1\}$. The agency cares both about its private value (bias, b) and the common good (ω) associated with choosing action $a = 1$.[4] The agency interacts with advocates –special interest groups with specific biases of their own and different stakes in this feature of the rule. Advocate i can privately collect information about ω and, if they so choose, disclose either a verifiable signal η of the state ω (co-collected with ω) or a soft recommendation $\hat{a}_i$, or both. The agency's bias is private information, leaving some residual uncertainty about the outcome of commenting for the advocate. Selection into commenting is endogenous in this model and the framework addresses head-on an important issue

[3] The portfolio of new facts that we document presents a useful set of empirical regularities necessary to discipline theories of persuasion and strategic advocacy. The literature is extremely vast and our model captures some salient, but by all means not all features of some of the classic models in this area (Crawford and Sobel, 1982; Dewatripont and Tirole, 1999; Krishna and Morgan, 2004; Kamenica and Gentzkow, 2011).

[4] A large class of common agency models applied to special interest politics and lobbying assume politicians care about a weighted average of aggregate welfare and private campaign contributions. See Grossman and Helpman (1994, 2001).

of the rulemaking environment: that not all players with a stake in a rule do in fact comment in equilibrium, and the selection into who comments is informative to the regulator –not just the content of the comment itself.

We study the hard information case first. We show that only certain organizations on a specific side of each issue will comment. Two forces are at work. First, there is the simple fact that in the absence of any comments, the agency will choose according to their (predictable) bias. Therefore advocates who want a policy change from the agency's default choice will be the ones who have the strongest incentive to comment –a feature strongly supported by our analysis of the comments and responses written by agencies (Fact #1). Second, there is a redundancy in collecting information from an advocate from the same side of the agency when an opponent also collects the same information about ω.

In support of the hard information model, Fact #2 shows that in the data a sizable share of comment letters, but by no means all, contain hard information.

Next, we consider the soft information case. In contrast to the hard information model, under cheap talk, a comment is most effective if the message is counter to the sender's bias. This produces a fundamentally different prediction about who will comment: Under hard information, comments will be "pro-attitudinal" (advocates with opposing biases request changes), while under cheap talk comments will be "counter-attitudinal" (advocates with similar bias to the agency request changes). On this dimension we find mixed empirical evidence (Fact #3): We can show that about 55% of comments are pro-attitudinal and the remaining 45% are counter-attitudinal. Importantly, the theory also highlights how the the models make different predictions for advocate's success in persuading agencies to make changes to the design of rules. With the soft information model, the identity of advocates is important and advocates with similar biases to the agency will have higher success rates. In contrast, with hard information, the presence of verifiable information makes the identity of the sender immaterial from the perspective of the decision by the agency. Here, our empirical results are more supportive of the soft information model. Facts #4 and #5 show that comments from advocates politically aligned with the agency and advocates that are more moderate in their ideological position have higher influence on changing regulation. In contrast, we do not find that large and small firms differ in their influence, controlling for the magnitude of their commenting activity (Fact #6).

Overall, we show that a mixed, semi-hard information environment in the spirit of Dewatripont and Tirole (2005) can reconcile several empirical regularities (Facts #3, #4, and #5, for example), pointing to this mode of communication as appropriate for rulemaking. This environment further allows to incorporate other empirically salient elements, like rubberstamping and agency review.[5]

Even under a consensus that rulemaking may be as consequential as lawmaking (Lawson, 1994), the Political Economy and Political Science's attention to the legislative branch has greatly surpassed that to regulation. From the Economics literature perspective especially, established theory of regulation is infrequently connected to its empirics, and aggregate studies on the costs and benefits of regulation are often less than complete (Bombardini et al., 2025). Few applications focus on rulemaking specifically,[6] while theory-to-data frameworks appear more frequently in the analysis of supervision and enforcement (among others, Garicano et al. 2016; Alesina et al. 2018; Aghion et al. 2021).

Early empirical analysis of regulatory rulemaking are limited to case studies or area-specific samples[7], because of the complexity, heterogeneity, and high specificity of the regulatory text and comments exchanged. More recent and related works on the quantitative analysis of regulation are more representative and range from the analysis of completion and quality of rules to their impact and costs for economic agents. In a recent contribution Bellodi et al. (2025), for example, shows, when looking at the US Unified Agenda of Federal Regulatory and Deregulatory Actions, rules that are overseen by regulators who are personally politically misaligned with the party of the US president take systematically longer to complete. Applications to the analysis of the compliance costs of regulation range from text-based methodologies (Kalmenovitz, 2023; Singla, 2023; Chang et al., 2023), to earning announcements and financial disclosures (Davis, 2017; Calomiris et al., 2020), to establishment-level survey evidence (Trebbi et al., 2026). Extant empirical research also rarely supplies a theoretical framework necessary to explicit mechanisms underlying the data generating process. This paper attempts to fill some of these gaps in the literature.

Our theoretical framework builds upon a substantial body of literature on the influence of interest groups in public decision-making. This influence can be indirect, operating through politicians' incentives to appeal to the electorate. The Ramsey literature, initiated by Lindbeck and Weibull (1987), demonstrates how the design of electoral platforms optimally reflects the relative elasticities of voting responses to platform characteristics. In the pandering or electioneering literature,[8] governments signal their empathy toward either the general population

[5]In Appendix A, Supplement C we show how to extend the model to incorporate pre-NPRM influence in shaping a rule proposal and post rule finalization judicial appeals, both important empirical elements of the rulemaking process. Heterogeneity in search technologies and alternative forms of information collection are also explored.

[6]See Furlong (1997); Yackee (2019) for discussions of the sparsity of empirical analysis in early work on rulemaking.

[7]Some examples include Yackee (2006); Magat et al. (2013); Lowande and Augustine Potter (2021).

[8]Battaglini and Coate (2007, 2008); Gavazza and Lizzeri (2009); Lizzeri and Persico (2005); Maskin and Tirole

or specific interest groups. Or influence can be direct, through information exchange. This approach—pioneered by Austen-Smith and Wright (1992, 1994) and most directly related to our work—examines the provision of information by well-informed and well-resourced SIGs to legislators and governments. Using a costly-state-verification framework, Austen-Smith and Wright show how lobbying is mainly performed by opponents and can improve legislative decisions; their 1994 paper introduces the notion of counteractive lobbying to explain why interest groups sometimes lobby their allies. This informational approach offers unique potential to address questions such as when public decision-makers are likely to be well informed[9] and, in the context of the NPRM process, who participates in commenting and with what degree of effectiveness. We argue that the NPRM process offers a comparatively "clean" environment for testing informational lobbying models: (i) alternative influence channels, such as vote buying (Dekel et al., 2008), are absent; (ii) access to officials is unrestricted; (iii) issue complexity and expertise are salient, making policy specialists, informational influence, and "quiet politics" more impactful; and (iv) the decision-making authority –the agency– is clearly identified[10].

The theoretical component of our paper extends the Austen-Smith and Wright agenda in several directions to inform our empirical analysis. First, we begin by providing a systematic analysis of how the characteristics of information affect its provision and its capacity to shape policy outcomes. While the persuasion literature typically studies either hard or soft information (or, less commonly, hybrid forms), it rarely compares or tests their respective predictions. To address this gap, we develop a unified model in which the sole variation lies in how easily the agency can verify the veracity of messages from SIGs. Second, we examine both the sending and receiving sides of informational lobbying. Indeed, the political economy literature offers competing perspectives on targeting by the sending side: (i) Fence-sitter lobbying: An elasticity-based view predicts that lobbying should focus on uncommitted legislators. (ii) Same-side lobbying: Classical political science emphasizes officials' receptiveness to "stimuli," suggesting that lobbyists should target like-minded policymakers. Relatedly, within the soft-information framework, the "access-buying" view posits that congruent lobbyists are willing to pay more for opportunities to engage sympathetic officials[11], for whom time is scarce,

(2004, 2019).

[9]The multi-advocate cheap talk literature finds conditions under which full revelation of the state of nature is possible (Krishna and Morgan, 2001a; Battaglini, 2002).

[10]As Schnakenberg (2017) notes, that most informational models involve only one legislator is limiting, as "interest groups do not lobby a single legislator in a vacuum but must instead build coalitions in favor of their preferred policy" (see his paper, as well as Alonso and Câmara (2016) and Caillaud and Tirole (2007) for contributions in which persuasion works differently in collective choice institutions). This is another reason why the NPRM process offers a simpler testing ground for informational theories of lobbying.

[11]Austen-Smith (1995); Grossman and Helpman (2001) (ch. 5).

thereby increasing their likelihood of influence[12]. (iii) Opposite-side lobbying: Austen-Smith and Wright's model predicts lobbying by groups predisposed to oppose the government's position.[13] Third, we offer a novel dimension: mode of argumentation. In addition to identifying who lobbies, we investigate how they seek to influence decisions. We define an argument as pro-attitudinal if it aligns with the position a lobby would adopt absent new information, and counter-attitudinal if it advocates the opposite. Like ideological distance, the "attitudinality" of an argument is observable, making it an analytically valuable dimension.

Finally, from a methodological perspective, we make advancements in the use of text analysis to measure outcomes in a real-world communication game. A general challenge in the literature on lobbying (broadly interpreted) is the difficulty of measuring outcomes related to communications with policymakers (Bombardini and Trebbi, 2020). The U.S. notice and comment procedure is appealing because the text of comments is publicly available and agencies are required to respond to comments in the published rule. This creates a paper trail recording interactions between advocates and policymakers that include both the advocate's preferred policy outcomes and the policymaker's decision. The challenge facing researchers is how to deal with the sheer scale and complexity of all the text recorded. Most influential early studies on influence rulemaking rely on painstaking work to hand-code commenter objectives and rulemaking outcomes (Yackee and Yackee, 2006; Libgober, 2020; Yackee, 2006; Naughton et al., 2009; Wagner et al., 2011). In all cases, the authors were forced to compromise in a variety of ways to make the task more feasible, such as coding a small number of rules or a subsample of the comments, simplifying the commenting positions to a pro- or anti-regulatory stance, or coding only a subset of the requests made in the comments. While very valuable, manual coding of commenting outcomes necessarily limits the types of analysis that is possible and its reproducibility. Several more recent studies have employed alternative approaches, using text similarity metrics to compare comment text to rules or rules to other features (Bertrand et al., 2021; Kalmenovitz and Chen, 2024; Dwidar, 2022; Carpenter et al., 2024) or large language models. These methods have the advantage of scaling up the analysis, which in turn enables researchers to use panel data methods to explore patterns of influence, as well as providing a more complete picture of the distribution of influence across commenters. However, ultimately textual similarity is at best only correlated with the actual policy outcomes of interest. Much of the response text written by agencies describes how the agency is rejecting various arguments

[12]Another theory predicting that lobbying is targeted to like-minded legislators is Hall and Deardorff (2006)'s "budget-based lobbying", in which lobbyists subsidize the work of legislators. McCarty and Poole (1998) observe that contributions to PACs are "same-side", i.e., inversely related to the ideological distance with the recipient.

[13]Also see Cukierman and Tommasi (1998) on how the identity of the policymaker affects the credibility of a policy message.

made in comments[14] –a strong correlation between comments and this kind of dismissive text is not a good indicator of influence. A contribution of this paper is that we extend previous Natural Language Processing work on comment-response linking (Xing et al., 2023) to develop a novel text-based outcome measure that captures granular comment-level outcomes, while also allowing large scale analysis of thousands of rules and millions of comments. This gives us a much more complete and detailed view of interactions between advocates and policymakers than previously available to scholars of lobbying and rulemaking.

The remaining of the paper is organized as follows. The next section presents our theoretical framework and its results. Section 3 presents the institutional details and examples useful for our understanding of the notice and comment regulatory process, the data and the details of the NLP processes that we use. The facts and results of our analysis appear in Section 4. Section 5 concludes.

2 Rulemaking: A theoretical framework

2.1 Model

The contribution of this section is to present a unified model useful to nest and compare the implications of hard, soft, and hybrid information in rulemaking. There are three players, the agency and two advocates representing two opposing interest groups. The agency takes action $a \in \{0, 1\}$, where $a = 1$ moves the public policy towards the right and $a = 0$ consists in not changing a particular aspect of the regulation[15]. A player's preferences reflect their own bias, their size or stake, and a common-good variable, ω, capturing the quality of the decision. The state of nature ω is initially unknown to all and drawn from distribution $F(\omega)$ with mean 0. We think of ω as an objective-quality, consensual parameter favoring one decision or the other[16]. The advocates have expertise, data, and budgets allowing them to learn a costly signal of ω . In contrast, the agency does not have resources that would enable it to acquire information (later, though, we will allow it to review arguments brought by the parties).

Advocates. There are two advocates, a left-wing advocate L ("she") and a right-wing one R ("he"). Advocate $i \in \{L, R\}$ has known bias b_i (such that $b_L < 0 < b_R$), stake (or size) s_i, and

[14] See Fact #7 in Appendix B.

[15] One may have in mind that the left-leaning party was in power, and now the right-leaning one has taken over.

[16] Alternatively, ω could reflect information drawing attention to the strong opposition or support of politically important interest groups. The analysis would be the same provided that the agency internalizes some of the interest groups' interests: The interest group's intensity of preferences would then be part of the "common good", as perceived not only by the relevant interest group but also by the agency.

net utility from action a in state of nature ω

$$U_i = (b_i + \omega)s_i a.$$

Agency. The agency ("it") similarly cares not only about its private valuation for the policy move, $b > 0$, but also about the common good, ω. Normalizing its stake to 1, the agency's utility function is:

$$V = (b + \omega)a.$$

The agency's ideological bias b may reflect its independence (the bias may be smaller for an independent agency), the extent to which the administration strives to have a good relationship with Congress, the country's polarization, and the agency's top brass's own preferences. It is common knowledge that the agency is right-biased, but the magnitude of the bias is private information. For the parties, the agency's bias admits smooth distribution $G(b)$, with mean $\bar{b}$ and full support on $(0, b^{max})$, where $b^{max} \leq +\infty$.[17] When comparing two distributions of partisanship for the agency, we will say that distribution $H(b)$ first-order stochastically dominates (FOSD) distribution $G(b)$ if $H(b) < G(b)$ for all $b > 0$: The agency is more partisan under distribution H.

The nature of information. An advocate's information-collection technology is invariant to the nature of information; so, the only moving part is how easily the resulting signal can be communicated. Each advocate i can engage in non-directed search to learn the state of nature[18]. When the search is successful, the advocate learns both the true state ω, which is soft information, and obtains a verifiable signal η of that state[19]. The signal η is hard information (verifiable by the agency at no cost)[20]: it can be concealed, but not falsified. Let $F(\omega \mid \eta)$ denote the posterior distribution of ω given η. It has full support for all η. The family $F(\omega \mid \eta)$ is ranked according to the monotone-likelihood-ratio property (MLRP): a higher signal is good news about the state of

[17]For much of the analysis, the asymmetry of information about b is irrelevant, and we will occasionally illustrate a result through a deterministic bias at $\bar{b}$. The one result that requires uncertainty about agency preferences is when we compare success rates, since with a known b, an advocate would always either comment successfully or not comment if there is an arbitrarily small cost of commenting. In contrast, when not knowing the agency's bias, advocate i in general achieves success rates between 0 and 1.

[18]The analysis can be extended to accommodate flexible information structures such as those associated with entropy and max-slope functionals.

[19]There is a single such signal for simplicity.

[20]One may have in mind for example that $\omega = \eta + \epsilon$, with $\epsilon \in \mathbb{R}$ and $E[\epsilon \mid \eta] = 0$, say, where η refers to the communicable part of knowledge (hard information) and ϵ stands for the less-communicable/tacit-knowledge part. The hard-information model is a special case of this one with $\epsilon \equiv 0$ and $\eta = \omega$. The soft-information model corresponds to $\eta \equiv 0$ and $\epsilon = \omega$.

nature.

This representation accommodates the standard models of persuasion:

- *Hard information*: $\eta = \omega$, i.e. $F(\omega \mid \eta) = 0$ for $\omega < \eta$ and $F(\omega \mid \eta) = 1$ for $\omega \geq \eta$.
- *Soft information (cheap talk)*: $F(\omega \mid \eta) = F(\omega)$ for all η .
- *Hybrid information.* The agency, when receiving signal η from an advocate, can, at cost γ, learn ω.[21] Note that $\gamma = 0$ *de facto* implies hard information and that when $\gamma = +\infty$, any message beyond η is soft information. Hybrid information is therefore to be interpreted as an intermediate case, in which verifying the veracity of information is possible, but costly.

Acquisition and disclosure of information. Advocate i can privately learn $\{\omega, \eta\}$ with probability p_i at cost $C(p_i)$, where C is smoothly increasing and convex for positive values. Say, $C(0) = 0$, and, for $p_i > 0$, $C(p_i) = c_0 + c(p_i)$ where c_0 is a fixed cost, and $c(p_i)$, a differentiable and strictly increasing and convex function satisfies $c'(0) = 0$ and $c(1) = +\infty$ (guaranteeing an interior solution when the advocate acquires information[22]).

While ω in itself is soft information, the signal η can be disclosed to, and costlessly verified by the agency. The hard part of advocate i's comment is therefore either the absence of hard report, $r_i = \emptyset$, when uninformed, and $r_i \in \{\eta, \emptyset\}$ when informed. A disclosure by advocate i is more broadly a pair of hard and soft communications $\{r_i, \hat{a}_i\}$, where $\hat{a}_i \in \{0, \emptyset, 1\}$ is a recommended action ($\emptyset$ is then the absence of soft recommendation). "Not commenting" will refer to the complete absence of communication $\{r_i = \emptyset, \hat{a}_i = \emptyset\}$.

Definition (soft recommendations)

(i) We will restrict recommendation strategies when informed to be "consistent" or "plainspoken" in that agency obedience of the recommendation benefits the sender: $\hat{a}_i = 1$ if $b_i + \omega > 0$, and $\hat{a}_i = 0$ if $b_i + \omega < 0$.

(ii) A recommendation $\hat{a}_i \in \{0, 1\}$ is said to be pro-attitudinal (resp., counter-attitudinal) if $b_i(\hat{a}_i - \frac{1}{2}) > 0$ (resp., $b_i(\hat{a}_i - \frac{1}{2}) < 0$).

To clarify with respect to (ii), consider the case of a Republican-leaning agency evaluating

[21] For example by learning ϵ in the additive case $\omega = \eta + \epsilon$.

[22] For expositional convenience, we will more broadly assume that $c(1)$ is sufficiently high that the solution is interior even when $c(1) < +\infty$ (as is the case for, say, a quadratic cost function). Note also that this framework accommodates the standard binary-search model in which only two levels of information, 0 and p, are possible: let $c(\tilde{p}) = 0$ if $\tilde{p} \leq p$, and $c(\tilde{p}) = +\infty$ if $\tilde{p} > p$.

a rule change $a = 1$ and a Democratic advocate commenting on this. A pro-attitudinal recommendation on the part of the Democratic advocate would be a comment opposing the rule, $\hat{a}_L = 0$. A counter-attitudinal recommendation on the part of the Democratic advocate would be one supporting the rule change, $\hat{a}_L = 1$.

Timing. The timing goes as follows:

1. The agency issues an NPRM.
2. The advocates, $i \in \{L, R\}$, privately collect information, and may comment ($\{r_i, \hat{a}_i\} \neq \{\varnothing, \varnothing\}$) or not ($\{r_i, \hat{a}_i\} = \{\varnothing, \varnothing\}$).
3. Observing the comments (or their absence), the agency chooses $a \in \{0, 1\}$.

The model is thus explicitly designed around selection into commenting by firms. We assume that there is an arbitrarily small cost of commenting, that is of disclosing η or/and of making a recommendation $\hat{a}_i$ in $\{0, 1\}$. An advocate therefore does not comment unless this strictly increases the advocate's expected utility (in particular, unless the comment has strictly positive probability of changing the agency's decision).

Equilibrium. Regardless of the nature of information, we will look for a pure-strategy equilibrium in which: (i) recommendations, if any, are consistent, and (ii) the agency follows its initial instinct when receiving no comment: $a_\varnothing = 1$, where $a_\varnothing$ denotes the agency's default choice of action when none of the advocates comments.

We will show that such an equilibrium always exists; intuitively, the absence of comment by opponents to the reform should comfort the agency with pursuing its initial plan. This equilibrium may, but need not be unique. All supplemental material and proofs are available in Appendix A.

2.2 Hard information

2.2.1 Equilibrium

Assume hard information ($\eta = \omega$). We first describe an equilibrium in which: (i) the agency selects $a = 1$ when receiving no information (as it would were information acquisition very costly, e.g., c_0 high), (ii) only advocate L, namely the side that objects to the agency's decision, reports hard information, and (iii) there is no recommendation (soft communication). This implies that efforts to acquire information satisfy:

$$p_L^* > p_R^* = 0.$$

Because the agency selects $a_\varnothing = 1$ when receiving no comment, advocate L comments when potentially pivotal, i.e., when learning state $\omega < 0$, implying that: (i) she wants to alter the agency's decision ($b_L + \omega < 0$), and (ii) reporting alters the decision with strictly positive probability (namely, when b is small enough that $b + \omega < 0$). Commenting is then pro-attitudinal. The optimal search by advocate L is such that her marginal cost of information acquisition is equal to her marginal benefit, namely the advocate's net gain when her information overturns the agency's decision:

$$c'(p_L^*) = v_L \equiv E_b\left[\int_{-\infty}^{-b}[-(b_L + \omega)s_L]dF(\omega)\right], \tag{1}$$

where v_L denotes L's value of information[23]. To complete the demonstration that this indeed is an equilibrium, we need to check that $a_\varnothing = 1$, that there is no scope for cheap talk, and that advocate R does not collect information:

Default decision. Given that $p_L^* > 0$, receiving no comment from advocate L shifts the agency's posterior distribution of ω to the right[24], reinforcing its preference for $a = 1$ (that is, $b + \omega_\varnothing$, where $\omega_\varnothing \equiv E[\omega \mid \varnothing] > 0$). So $a_\varnothing = 1$ is indeed the default option. The game exhibits "(unilateral) expectation conformity"[25], meaning that being expected to acquire information raises a player's incentive to acquire information.

No scope for cheap talk. Let us, first, show that the possibility of hard communication crowds out soft communication. Suppose that advocate L has learned the state of nature. She reports ω when favorable ($\omega < 0$). She would also want to recommend $\hat{a}_L = 0$, although not report

[23] The elasticity of commenting to the left-side's stake in this one-sided-commenting equilibrium is therefore equal to:

$$\frac{dp_L^*/p_L^*}{ds_L/s_L} = \frac{C'(p_L^*)}{p_L^* C''(p_L^*)}.$$

For instance, the elasticity is equal to 1 in the quadratic-cost case. For the elasticity to be less than 1, the cost function must be more convex that the quadratic one.

[24] More formally, letting $M^+(\omega) \equiv E[\tilde{\omega} \mid \tilde{\omega} > \omega]$ denote the lower-tail-truncated mean, the posterior mean of ω conditional on not receiving information, $\omega_\varnothing$, is given by:

$$\omega_\varnothing = \frac{p_L^*[1 - F(0]M^+(0)}{1 - p_L^* F(0)} > 0$$

and so $b + \omega_\varnothing > 0$ for all b.

[25] See Pavan and Tirole (2024).

ω, when $\omega > 0$ and $b_L + \omega < 0$. But in the absence of hard evidence, the agency infers that ω is positive if advocate L indeed has information, and sticks with the default decision. Thus, there is no scope for soft communication by advocate L. A similar reasoning would also apply to advocate R. Further, there is no scope either for a cheap-talk announcement by the agency of the realized value of b [26]. Intuitively, the agency, regardless its of ideology, gains from the advocate collecting more information; thus, there cannot be any signaling of the agency's type.

One-sided commenting. For this one-sided-commenting equilibrium to exist, advocate L, but not advocate R, must invest in information acquisition. Define (by analogy with v_L) R's value of information as:

$$v_R \equiv E_b\left[\int_{-\infty}^{\min\{-b,-b_R\}}[-(b_R+\omega)s_R]dF(\omega)\right],$$

implying that v_R is a decreasing function of b_R and is smaller than v_L provided that $s_R \leq s_L$. Existence of the "L-advocacy-only" equilibrium requires when $C(p_i) = c_0 + c(p_i)$:

$$max_{p_L}\{p_L v_L - c(p_L)\} \geq c_0 \geq max_{p_R}\{p_R(1-p_L^*)v_R - c(p_R)\}. \tag{2}$$

In a symmetric environment (i.e., with equal stakes $s_L = s_R = s$, intensities of bias $b_L + b_R = 0$, and identical cost functions, as we assumed), the left term in (2) exceeds the right term for different reasons:

- *Status-quo bias.* First, there is a status-quo bias in favor of advocate R. The agency will go ahead with the proposed rule in the absence of information ($a_\emptyset = 1$). So, acquiring information serves to reverse the decision from $a = 1$ to $a = 0$ and because $b_R > b_L$, the right-side advocate has less incentive to search for information. In particular, if the distribution of b were to put all the weight on values above b_R, then $v_L - v_R = E_b[F(-b)](b_R - b_L)s]$.

[26]Consider a message m giving rise to posterior beliefs $\hat{G}(m)$ and thereby search intensity $p_L^*(m)$. When of type b, the agency would choose message m so as to solve

$$max_m\{(1-p_L^*(m))b + p_L^*(m)\int_{-b}^{+\infty}(b+\omega)dF(\omega)\}$$

The derivative of the maximand with respect to $p_L^*(m)$ is equal to $\int_{-\infty}^{-b}[-(b+\omega)]dF(\omega) > 0$: As is natural, the agency is better off with more information. The cross-derivative of the maximand with respect to $p_L^*(m)$ and b is equal to 0. And so there is no sorting condition. Intuitively, all types would like to pretend they have the lowest possible bias (0 here), signaling that they are most receptive to left-side comments. The low types gain more from more information than the high types, as they will make more use of that information. However, the common preference for pretending to be receptive precludes any separation.

- *Different intensities of commenting*. Second, when advocate R turns out to be more extremist than the agency ($b_R > b$), he makes little use of the information he receives. Because the lower bound of the support of G is 0, we have $v_L - v_R > [E_b F(-b)](b_R - b_L)s$.[27]

- *Redundancy.* Third, and endogenously, advocate R is disincentivized from searching, as his search may duplicate advocate L's own search (this redundancy is captured by the term $(1 - p_L^*)$). Despite the divergence of preferences between advocates, the common-good parameter ω confers in part a public-good flavor to the information-acquisition game. Unlike the first two effects, which unambiguously favor information acquisition by the L advocate, expectation conformity is an expectational feature and depends on who is expected to acquire information in equilibrium.

More generally, as long as $s_L \geq s_R$, there exists an equilibrium with $p_L^* > p_R^* \geq 0$. We therefore make from now on an assumption so that the natural proclivity of the opposing group to disclose more (or be the only one to disclose) hard information not be reversed by a much higher stake for the same-side group:

Assumption. $s_L \geq s_R$.

Success rate. The success rate of advocate L's comments is advocate's bias- and stake-independent:

$$\sigma = \frac{p_L^* \int_0^{b^{max}} F(-b) dG(b)}{p_L^* F(0)} \in (0, 1).$$

Because $F(-b)$ is a decreasing function of b, a FOSD increase in the agency's bias (i.e., a more partisan agency) reduces the success rate.

Equilibrium uniqueness. The case in which c_0 exceeds L's upper bound on her benefit of advocacy ($p_L^* v_L - c(p_L^*)$) is uninteresting: There is no information collection at all and no commenting, and the agency accordingly selects according to its prior: $a \equiv 1$. We will henceforth assume that $p_L^* v_L - c(p_L^*) \geq c_0$. When c_0 is small in contrast, the right inequality in condition (2) is violated, and there can be an equilibrium in which R (also) acquires information.

Note, though, that if

$$c_0 > max_{p_R}\{p_R v_R - c(p_R)\}, \tag{3}$$

a condition that is stronger than the right-hand-side condition in (2), then, given $a_\emptyset = 1$, searching for information is a dominated strategy for advocate R, and so an equilibrium with $a_\emptyset = 1$

[27] Note that $v_L - v_R = [E_b F(-b)](b_R - b_L)s + \int_0^{b_R} dG(b) \int_{-b_R}^{-b} [-(b_R + \omega)s] dF(\omega)$.

and advocate R commenting does not exist. Appendix A Supplement A characterizes the various equilibria, including the possibility that $a_{\emptyset} = 0$, if the equilibrium is not unique.

Proposition 2.1. ***(Hard information)***

Under hard information:

(i) Equilibrium. *If* $p_L^* v_L - c(p_L^*) \geq c_0 \geq max_{p_R}\{p_R(1-p_L^*)v_R - c(p_R)\}$*, there exists an equilibrium in which only the opposite side comments (*$p_L^* > p_R^* = 0$*). That is, there is no jousting for influence between the majority and the opposition. Comments are pro-attitudinal.*

(ii) Success rate. *In this equilibrium, the success rate* σ *is lower than 1 and independent of the advocate's stake,* s_i*, and bias,* b_i*. It decreases with the agency's bias (in the sense of FOSD).*

(iii) Uniqueness. *If furthermore* $c_0 > max_{p_R}\{p_R v_R - c(p_R)\}$*, the equilibrium described in (i) is the unique equilibrium with* $a_{\emptyset} = 1$*. More generally, when the marginal cost of information acquisition is large, so that the probabilities of discovering hard evidence are small, and/or the agency is very partisan (in the sense of FOSD shifts in the distribution* $G(b)$*), an equilibrium with comments from advocate* R *does not exist.*

Remark. The analysis accommodates the presence of multiple advocates on the same side, with different degrees or partisanship and different stakes. Appendix A Supplement B solves for equilibrium and shows in particular that Proposition 1 (ii) (equal success rates) carries over to this more general environment.

Finally, Appendix A, Supplement C explores in detail the implications for our framework of endogenizing the proposed rule and of explicitly modeling pre-NPRM influence. It assumes that the agency incurs an opportunity cost for each NPRM, and therefore may be loath to engage in a process if its expected payoff is too small. Advocate R then gains by nudging the agency, i.e., by bringing some information, making it worthwhile for the agency to initiate the process. Advocate L in contrast comments on the proposed rule. We direct to this theoretical supplement the reader interested in modeling this antecedent phase of the rulemaking process.

2.3 Soft information

Suppose now that the information on ω is entirely soft: The signal η cannot be verified by the agency (equivalently it has no informational content), and so advocates can only recommend an action. To illustrate the working of the model in this cheap talk environment, we focus on a deterministic b for expositional conciseness, setting $b = \bar{b} > 0$.

We look for a pure-strategy equilibrium. Advocate $i \in \{L, R\}$ acquires information (with the same technology as in the hard-information model), and can either recommend an action $\hat{a}_i \in \{0, 1\}$, or not comment ($\varnothing$). As we will see, the outcome is qualitatively very different under cheap talk. Intuitively, under hard information, an advocate brings evidence to support her case. Under soft information, the message is much more likely to be effective if counter-attitudinal, i.e., if it "surprises" the agency: either advocate L recommends $\hat{a}_L = 1$ or advocate R recommends $\hat{a}_R = 0$.

2.3.1 The single-advocate case

The case of a single advocate will underlie the analysis of the multiple-sender case. We investigate whether the advocate's message can be informative. For that, we look for an informative (non-babbling) pure-strategy equilibrium with default action $a_\varnothing = 1$. Suppose therefore that there are two recommendations, leading to actions $a = 0$ and $a = 1$ respectively. Throughout the entire section, we will posit that c_0 is not too large, so that a player will collect some information when having a stake in doing so.

L-advocate. Suppose that communication is effective and so consider the two messages, $\hat{a}_L = 0$ and Ø, that lead the agency to choose $a = 0$ and $a = 1$, respectively[28]. Because the default action is $a_\varnothing = 1$ and the only information collector is advocate L, the latter does not comment if $b_L + \omega > 0$, and recommends $\hat{a}_L = 0$ if either $b_L + \omega < 0$ or she has not acquired any information (as $b_L < 0$). Because advocate L's recommendation of action $a = 0$ is rubber-stamped by the agency, her incentive to get informed is to avoid inaction when she actually would benefit from reform. Hence, advocate $L's$ equilibrium probability of information acquisition, $p_L = \hat{p}_L$, solves:

$$\max_{p_L} \Big\{ p_L \Big[\int_{-b_L}^{+\infty} (b_L + \omega) s_L dF(\omega) \Big] - C(p_L) \Big\}. \tag{4}$$

An informative communication requires message $\hat{a}_L = 0$ to be obeyed by the agency, or

$$\bar{b} + \frac{\hat{p}_L \int_{-\infty}^{-b_L} \omega dF(\omega)}{1 - \hat{p}_L[1 - F(-b_L)]} < 0. \tag{5}$$

From (4), advocate L's incentive to acquire information is meager, as the advocate gains only if $\omega > -b_L > 0$. Condition (5) and $\bar{b} > 0$ imply that informative messages from advocate L are

[28]Recall that under a vanishingly small cost of commenting, the advocate selects not to comment rather than sending the non-pivotal recommendation $\hat{a}_L = 1$.

unlikely to be feasible, especially if the marginal cost of information acquisition is large ($\hat{p}_L$ low) or if advocate L is quite partisan (b_L low).

R-advocate. Let us now perform the same exercise with advocate R as the only advocate, and an informative equilibrium with two messages, $\hat{a}_R = 0$ and Ø. The previous reasoning in this case has advocate R not comment when either he is uniformed (as $b_R > 0$) or he is informed that $b_R + \omega > 0$, and, provided that the recommendation overturns the planned decision, to recommend $\hat{a}_R = 0$ when informed that $b_R + \omega < 0$. The agency selects $a = 1$ in the absence of comment (as $\bar{b} + p_R \int_{-b_R}^{+\infty} \omega dF(\omega)/[1 - p_R F(-b_R)] > 0$ for any p_R). When $\hat{a}_R = 0$, obedience by the agency requires that, letting $M^{-}(\omega) \equiv E[\tilde{\omega} \mid \tilde{\omega} \leq \omega]$ denote the upper-tail-truncated mean,

$$\bar{b} + M^{-}(-b_R) < 0. \tag{6}$$

$p_R = \hat{p}_R$ then solves

$$\max_{p_R} \left\{ p_R \int_{-\infty}^{-b_R} (b_R + \omega) s_R dF(\omega) - C(p_R) \right\}. \tag{7}$$

Condition (6), unlike condition (5) is a relatively weak condition; in particular, it is always satisfied when the agency is not more partisan than the interest group: $\bar{b} \leq b_R$. The comparison between (5) and (6) is instructive: The presence of $\hat{p}_L$ in the first and the absence of $\hat{p}_R$ in the second comes from the fact that advocate R has no incentive to mislead the agency when uninformed. Hence, advocate R is trusted more (he recommends against the status-quo only when informed) and so is more likely to inform the agency. This suggests that counter-attitudinal messages by the same-side interest group are more likely in the two-advocate case, to which we now turn.

2.3.2 Multiple advocates: The location of activism

Assume now that both advocates can collect information. After acquiring information (with probability p_i for advocate $i \in \{L, R\}$) advocates simultaneously send messages $\hat{a}_i \in \{0, \varnothing, 1\}$ and the agency then selects action $a(\hat{a}_L, \hat{a}_R) \in \{0, 1\}$[29].

Proposition 2.2. ***(Soft information)***

The single-advocate informative equilibria are equilibria of the two-advocate cheap-talk game under the following conditions:

- *Counter-attitudinal communication by advocate R exists under a weak condition, (6), that*

[29]The babbling equilibrium has $a(\hat{a}_L, \hat{a}_R) = 1$ for all $\{\hat{a}_L, \hat{a}_R\}$, $p_L = p_R = 0$ and $\hat{a}_L = \hat{a}_R = \varnothing$.

is always satisfied for $b_R \geq \bar{b}$ *or, when* $b_R < \bar{b}$, *for* $\bar{b} + M^{-}(-b_R) < 0$ *, provided that (5) is violated.*

- *Counter-attitudinal communication by advocate* L *requires a strong condition, (5), that cannot be satisfied for instance with a high marginal cost of information acquisition or when the left advocate is very partisan. Such communication may not exist even if (5) is satisfied, as information collection by advocate* R *may reduce her incentives to collect information and thereby damage her credibility.*

Proposition 2.2 demonstrates the difference in nature between hard and soft information. Information is naturally collected by the opposition advocate under hard information and by the majority advocate under soft information. Ignoring the babbling equilibrium, are there other equilibria? Appendix A Supplement D shows there may exist another possible non-babbling equilibrium. Like for the equilibrium involving counter-attitudinal communication by advocate L, very strong conditions must prevail for this equilibrium to exist, though.

Foreshadowing our empirical analysis, when comparing the implications of the two classes of models, empirical discriminating moments in the data arise from who is more likely to comment and who is successful in changing policy. The theory identifies several such moments. First, in hard information environments, commenters with a bias opposite to the agency's should communicate, while in cheap talk aligned commenters should communicate more. Second, in hard information environments, comments should be pro-attitudinal, while in cheap talk comments should be counter-attitudinal. Third, under hard information, the identity of the sender should not matter for success (hard information speaks for itself, irrespective of who brings it to the agency), while trust and identity matter under soft information (for example, commenters with bias similar to the agency are trusted more). This is illustrated in Table 1.

Table 1: Some Discriminating Theoretical Predictions

	Hard information	**Soft information**
Who comments?	Opposite-bias advocate	Same-bias advocate
Nature of comment	Pro-attitudinal	Counter-attitudinal
Influence of congruence in bias on success rate	Congruence-independent success rate	Congruence-dependent success rate

Finally, we can show that other implications of our setting, including an equilibrium success rate less than one, independence of success rates from the stake of the sender, and absence of jousting for influence in both hard (per Proposition 2.1) and soft information models.

2.4 Hybrid information

We have seen that hard information overwhelmingly generates pro-attitudinal comments from the parties, but has little scope for trust to matter (facts speak for themselves). In contrast, soft information is conducive to counter-attitudinal comments and emphasizes the role of trust, as measured by the congruence in bias between advocate and agency. In either case, comments do not call for assessments by the agency. Under hard information, a comment may fail to trigger a new decision because the latter would go against the agency's preferences, not because the agency questions the validity of the comment. Soft information is by definition unverifiable. In this sense, there is no scope for agency review of the comments under either paradigm, negating the widespread concern for agency-ex-post moral hazard, which figures prominently in NPRM institutions. This suggests considering the more general framework of hybrid information. As stated in Subsection 2.1, let $F(\omega \mid \eta)$ denote the agency's distribution over the state of nature when receiving hard signal η, and γ denote the agency's cost of learning the full state of nature when already knowing η (in contrast, the agency's cost of collecting information from scratch is prohibitive).

We look for an equilibrium in which: (i) as in the hard-information case, only advocate L searches for information in equilibrium, (ii) the agency picks $a_{\emptyset} = 1$ when no hard signal η is disclosed. Despite the presence of soft information, we offer no scope for pure cheap talk; the idea is related to the standard unraveling argument: An advocate cannot have influence unless it can convince the agency that it has information, which requires disclosing η.[30]

While an advocate knows the state of nature when acquiring it, the agency can, when confronted with signal η:

1. *Rubber-stamp* without checking, in the sense of doing what advocate L desires ($a = 0$),
2. *Review* at cost γ and then choose a according to the realization of the state of nature ($b + \omega \gtrless 0$),
3. *Fail to be receptive* to advocate L's attempt to influence policy (chooses the status quo $a = 1$).

[30] The reader may be concerned that the hybrid model does not nest cheap talk as a limit case. Note, though, that even when η is hardly informative, disclosing η may stand for the recommendation that the agency disposes of the reform (saying implicitly $b_i + \omega < 0$), while the absence of comment suggests going ahead with it. In a sense, the advocate in equilibrium must provide at least a narrative, even if this narrative has limited informational content.

Advocate L does not comment (here, reveal the realization of η) unless:

$$b_L + \omega < 0.$$

A comment from advocate L thus reveals that $\omega < -b_L$. Realizations of ω between $-b$ and $-b_L$ represent the extent of the conflict of interest between the agency and L (agency prefers $a = 1$, while L prefers $a = 0$[31]). Interestingly, what matters for the agency is not only the signal that is being transmitted (as under hard information), but also who transmitted it (as under soft information); in that sense, the hybrid model combines features of the hard and soft information paradigms. Let

$$\hat{F}(\omega \mid \eta, b_L) \equiv \frac{F(\omega \mid \eta)}{F(-b_L \mid \eta)}$$

denote the agency's posterior distribution, with support $(-\infty, -b_L)$, conditional on signal η and the fact that advocate L revealed η. The function $\hat{F}(\omega \mid \eta, b_L)$ is decreasing in η (due to MLRP) and increasing in b_L. Because the distribution of ω is right-tail-truncated, the posterior mean of ω in the absence of a comment satisfies $\omega_\varnothing > 0$ and so $b + \omega_\varnothing > 0$ for all b; hence $a_\varnothing(b) = 1$ for all b.

Suppose that advocate L reveals η. The agency's payoffs to its three strategies, given distribution $\hat{F}$, are:

- Nonreceptivity (agency does not review/evaluate and chooses $a = 1$):

$$V_1(\eta; b_L, b) = \int_{-\infty}^{-b_L} (b + \omega) d\hat{F}(\omega \mid \eta, b_L).$$

- Rubber-stamping following the comment (agency does not review and picks $a = 0$):

$$V_0(\eta; b_L, b) = 0.$$

- Review (agency pays γ and chooses $a = 1$ iff $b + \omega > 0$):

$$V_e(\eta; b_L, b) = \int_{-b}^{-b_L} (b + \omega) d\hat{F}(\omega \mid \eta, b_L) - \gamma = b - b_L - \int_{-b}^{-b_L} \hat{F}(\omega \mid \eta, b_L) d\omega - \gamma.$$

[31] For $\omega < -b$ instead both agency and L prefer $a = 0$.

To reduce the number of cases under consideration, we make weak assumptions on the impact of extreme values of the signal on these payoffs: for all $\{b_L, b\}$,(i) $lim_{\eta\to-\infty}\{V_1(\eta;b_L,b) - V_0(\eta;b_L,b)\} < 0$ (for very low signals, rubber-stamping dominates reforming); (ii) $lim_{\eta\to+\infty}\{V_1(\eta;b_L,b) - V_0(\eta;b_L,b)\} > 0$ (for very high signals, reforming dominates rubber-stamping); (iii) $lim_{\eta\to-\infty}\hat{F}(-b \mid \eta, b_L) = 1$ (almost all of the weight in the distribution of ω is on low values for very low signals); (iv) $lim_{\eta\to+\infty}\frac{F(\omega|\eta,b_L)}{F(-b_L|\eta,b_L)} = 0$ for $\omega < -b_L$ (for very high signals, the posterior distribution puts the weight close to the highest possible state of nature $\omega = -b_L$).

Proposition 2.3. ***(Hybrid information)***

The net benefits from skepticism, $V_1 - V_e$ (ignoring rather than reviewing), and $V_e - V_0$ (reviewing rather than rubber-stamping), decrease in b_L and increase in b and η.

(i) Strength of the narrative. *There exist $\eta_1(b_L,b) \le \eta_2(b_L,b)$ such that (a) the agency rubber-stamps and picks $a = 0$ for $\eta < \eta_1(b_L,b)$, (b) the agency reviews if $\eta_1(b_L,b) < \eta < \eta_2(b_L,b)$, (c) ignores the comment for $\eta > \eta_2(b_L,b)$. The review region (b) exists ($\eta_1(b_L,b) < \eta_2(b_L,b)$) if and only if $\gamma < \bar{\gamma} \equiv \int_{-b}^{-b_L}(b+\omega)d\hat{F}(\omega \mid \eta^\sharp, b_L)$, where $\eta^\sharp$ is defined by $V_1(\eta^\sharp;b_L,b) = V_0(\eta^\sharp;b_L,b)$.*

(ii) Advocate's partisanship. *The thresholds $\eta_1(b_L,b)$ and $\eta_2(b_L,b)$ increase as advocate L becomes more moderate (as b_L increases).*

(iii) Agency's partisanship. *The thresholds $\eta_1(b_L,b)$ and $\eta_2(b_L,b)$ decrease as the agency becomes more partisan (b increases).*

(iv) Success rate. *The L advocate's comment success rate decreases with η (the signal is less favorable to the L cause).*

Proposition 2.3 first observes that the agency is more friendly to a comment by advocate L, the more moderate the advocate and the agency, and the more favorable the signal is to the advocate's cause. The subsequent results draw the implications of these observations. Part (i) states that a stronger narrative (a lower η) is more likely to convince the agency, or at least to induce it to investigate. Part (ii) implies that a more moderate advocate is more likely to induce rubber-stamping when a thorough review would have taken place, or to induce such a review when the narrative would have been dismissed without further ado. Similarly, an increase in agency partisanship reduces the probability of rubber-stamping; part (iii) shows that a more partisan agency is less prone to rubber-stamp and more prone to just ignore the signal. Finally, part (iv) notes that a higher signal hurts the L cause[32].

[32]We conjecture that, like in the soft information case, advocate L's expected success rate is an increasing function of b_L: a more moderate advocate's comments are accepted more often. Starting from distribution $H(\eta)$, let $\hat{H}(\eta, b_L)$ denote the distribution conditional on b_L ; it is increasing in b_L. The overall success rate for an advocate of

The rest of the analysis is very similar to that in the hard-information case. First, we must verify that $a_\emptyset = 1$. This is slightly more involved than in the hard-information case. Here advocate L may not comment even though $\omega < 0$. Indeed, we have made an assumption guaranteeing that $E[\omega \mid \omega < -b_L, \eta] > 0$ for η large enough. So, the absence of comment may coincide with a state in which advocate L would prefer $a = 0$, but the narrative η is so favorable to R that L has no hope to bring the agency on board. Nonetheless, advocate L never comments when $b_L + \omega > 0$, and may comment (for lower η) when $b_L + \omega < 0$. Thus, conditionally on being informed and not commenting, the expectation of ω is strictly positive. As this expectation is 0 when uninformed, then the expectation of $b + \omega$ is strictly positive in the absence of comment, for all $b \geq 0$. We thus have verified that $a_\emptyset = 1$.

Second, we can again follow the analysis of the hard-information case and find conditions under which advocate R does not acquire information, and (stricter) conditions under which the equilibrium above is the unique equilibrium. Letting v_L and v_R denote advocates L and R's values of information[33], the condition for commenting by advocate L only is: $max_{p_L}\{p_L v_L - c(p_L)\} \geq c_0 \geq max_{p_R}\{p_R v_R - c(p_R)\}$.

To conclude, in terms of empirical predictions, the hybrid model accommodates a mixture of hard and soft information predictions as originally listed in Table 1. The effectiveness of messages is determined by the couple $\{b_i, \eta\}$ and therefore the hybrid model allows for hard information being part of the comment, but it also introduces a role for advocate identity and trust in determining which messages change agency policy. The hybrid variant of the framework further allows for both pro-attitudinal and for counter-attitudinal messages and for success rates of comments depending directly on the bias (moderate or extreme) and type (opposite or same

type b_L is therefore: $\sigma(b_L) = E_b[\hat{H}(\eta_1(b_L, b), b_L) + \int_{\eta_1(b_L,b)}^{\eta_2(b_L,b)} \sigma_e(b_L, b \mid \eta) d\hat{H}(\eta, b_L)]$, where $\sigma_e(b_L, b \mid \eta) = Pr(b + \omega < 0 \mid b_L + \omega < 0)$. The functions $\hat{H}(\eta, b_L)$, $\eta_1(b_L, b)$ and $\eta_2(b_L, b)$ are increasing in b_L, and that $\sigma_e(b_L, b \mid \eta)$ is decreasing in b_L and decreases with η . An increase in b_L indices a shift in composition (the rubberstamping region -success rate 1- expands to the detriment of the examination region -success rate between 0 and 1- and the latter expands to the detriment of the nonreceptivity region-success rate 0-), that by itself increases the overall success rate.

[33]Letting $E_\eta[.]$ denote the expectation relative to the prior distribution of η (if H denotes this distribution of η, $F(\omega) = \int F(\omega \mid \eta)] dH(\eta)$), and generalizing the previous expression,

$$v_L \equiv E_b[E_\eta[\int_{-\infty}^{\eta_1} \int_{-\infty}^{-b_L} [-(b_L + \omega)s_L] dF(\omega \mid \eta)] + \int_{\eta_1}^{\eta_2} \int_{-\infty}^{-b} [-(b_L + \omega)s_L] dF(\omega \mid \eta)]].$$

Similarly,

$$v_R \equiv E_b[E_\eta[\int_{-\infty}^{\eta_1} \int_{-\infty}^{-b_R} [-(b_R + \omega)s_R] dF(\omega \mid \eta)] + \int_{\eta_1}^{\eta_2} \int_{-\infty}^{min\{-b,-b_R\}} [-(b_R + \omega)s_R] dF(\omega \mid \eta)]].$$

Again, $v_R < v_L$ provided that $s_R \leq s_L$.

bias as the agency) of the advocate. In Section 4 we show that the empirical regularities found in the US rulemaking process most closely align with these predictions.

3 Institutional setting, data, and measurement

3.1 Institutional setting

We begin by providing a short overview of the process of rulemaking in the United States. While our theory is broader than this specific application, all our evidence will come from the US environment and therefore by necessity we focus on it. Notice that this subsection is not designed to be an exhaustive review from the Administrative Law perspective[34], but rather is an incomplete primer, useful to the reader lacking immediate familiarity with this important area of government function.

In the US, substantial amount of legislative activity is not performed directly by Congress, but delegated via the APA of 1946 to agencies which perform two fundamental quasi-legislative functions[35]: formal and informal rulemaking.[36] The terms "formal" and "informal" are somewhat of a misnomer, as both are highly formalized processes. Formal rulemaking is a form of rule making involving a quasi-judicial procedure, including trial-like hearings, testimony, cross examination, etc. It is costly and typically used by agencies only if Congress explicitly specifies that this type of process is needed for a specific rule.

Informal rulemaking, which is the subject of our empirical analysis, is the modal form of rulemaking across US federal agencies. It is based on a notice-and-comment structure, where agencies formulate a proposed rulemaking, publish such proposal on the Federal Register (including sufficient details and a rationale for the rule), open a comment period to receive comments from the public, and then decide (without obligation of incorporating comments) which comments to accept (changing the rule/policy accordingly). A revised and final version of the rule appears on the Federal Register (FR) and updates the Code of Federal Regulations (CFR).[37] The final rule includes a preamble, where significant and relevant comments raised by the public must be discussed, as per APA requirement.

[34] Mashaw and Merrill (1985).

[35] In the words of Croley (1995), *"rulemaking is the most important device federal agencies use to specify, clarify, and refine Congress' work-product –in short to finish the task of legislating."* (p.1512).

[36] Such agencies perform also functions of formal and informal adjudication under the APA, but we do not address these quasi-judicial activities here.

[37] The Code of Federal Regulations is a codification of the rules and regulations issued by all US federal agencies. It contains the rules implementing laws passed by Congress.

Importantly, differently from US federal lobbying of Congress or the Executive branch(Bombardini and Trebbi, 2020), rulemaking allows the researcher to retrace the subject matter of the communication, the issues/arguments raised by the commenter, whether the request made in the comment is accepted or not, and the rationale for the rule change or its rejection.

3.2 Data sources

Our raw data comes from two sources: official rulemaking documents published in the FR, and public comments submitted to federal regulators on the government website Regulations.gov.

The FR is the official publication for US federal rulemaking. All rules, proposed rules, and related notices are published in daily issues. Our sample of rules consists of all rules published from the beginning of 2008 to the end of 2022, a total of 16,079 documents. We collect the full text of these rules from XML format files[38] along with cleaned document metadata for all federal register documents available on the FR website API [39]. We also collect Unified Agenda data in XML format from reginfo.gov. The Unified Agenda records agencies planned rulemaking activities each year in a format that groups related proposed and final rules. Our code for downloading, parsing, and combining rulemaking documents from these sources is available as an open-source Python package[40].

We collect comments from Regulations.gov, the main site used by US federal agencies to manage their public comments. Regulations.gov was launched in January 2003 as a common platform for comment management. Agencies gradually adopted the platform over several years. By 2008, 86% of all calls for comments published in the FR included a link to Regulations.gov. We use the Regulations.gov API to collect all comments submitted from 2008-2022. In our analysis, we focus on comments submitted to the set of agencies that were using Regulations.gov by 2010. This gives us an initial sample of 8.5 million comment documents submitted to 87 agencies and covering 82% of all calls for comments published in the FR from 2008-2022.

On Regulations.gov, a comment consists of a text-entry field, optional fields where comment submitters can enter authorship information, any number of attached files, and some simple metadata tracking basic information, such as the submission date and docket ID. Most comments can be broken into two broad types: short comments submitted by individuals who write a message to the agency using the text entry field, and sophisticated comments written as formal letters and submitted as an attachment (usually PDF, but sometimes Word documents or

[38]https://www.govinfo.gov/bulkdata/FR
[39]see: https://www.federalregister.gov/reader-aids/developer-resources
[40]https://github.com/bradhackinen/frdocs

other file formats). We use the Tika Parser with OCR to extract text from all attachments, skipping the occasional attachment with an unusual file type. We further process both the text from the submission field and all attachments by splitting it into paragraphs, and using simple pattern matching rules to separate body text from extraneous text such as addresses, salutations, page numbers, and signature blocks. Our code for collecting comments, extracting, and preprocessing the raw text is available as another open-source Python package[41].

The linking of comments to rules requires two stages. In the first stage, we match FR documents to their copies on Regulations.gov based on available metadata (title, FR document number, citation), and full text when necessary. This gives us a mapping from Regulations.gov comment IDs to consistent FR document IDs for the documents containing the calls for comments. The second stage is to link FR documents together, so that comments linked to a proposed rule can be inferred to be relevant for a rule published at a later date. We use a mix of techniques, relying on Unified Agenda identifiers where possible, and then matching based on docket ID, title, topic similarity, and the sections of the CFR each rule is expected to modify. In our final linked data we find that more than 90% of comments linked to a rule are indeed relevant for that rule.

3.3 Measuring agency responses

Under the APA, US federal regulators are required to respond to public comments submitted during the rulemaking process. These responses must appear in the preamble of the final rule published in the FR [42]. We parse all rules published from 2008-2022 to construct a complete record of US federal regulator's responses to comments.

The main challenge with parsing these responses is that they can appear in a variety of formats. Agencies have a responsibility to make the responses clear, and many rules have an explicitly labeled section where comments are discussed, with a list of point-by-point replies to each narrow issue raised by commenters. A typical "response" consists of a brief summary of the issue (e.g. "Several commenters were concerned that....") and then the agency's response in the next paragraph (e.g. "We agree....and therefore ..."). However, the exact formatting of these responses varies from rule to rule, and in some rules responses are written in a much less structured way, or even integrated into the general plain language explanation of the rule. To overcome this challenge, we trained several flexible natural language classifiers to identify re-

[41] https://github.com/bradhackinen/regcomments

[42] "After consideration of the relevant matter presented, the agency shall incorporate in the rules adopted a concise general statement of their basis and purpose." (5 U.S.C. §§ 553(c))

sponse boundaries and contents using manually annotated paragraphs from rule documents.

We began by sampling short sequences of paragraphs (up to 11 paragraphs long) from the preambles of 2,269 rules, selected to maximize the diversity of agencies and years represented in the sample. The full training set contains 21,989 paragraphs. We hired a small team of law students from the University of Western Ontario to read the sampled text from each rule and provide paragraph-level annotations indicating whether each paragraph is part of a response, which paragraphs should be grouped together in the case of responses that extended over multiple paragraphs, and whether the agency indicated making a policy change in the response. The precise instructions given to the research assistants are presented in Appendix E. For the variable indicating whether the agency made a policy change we had multiple law students review each paragraph, and then used GPT-4 to review the same data, allowing an experienced law student to choose between GPT-4 and the initial manual annotations in the cases where the two disagreed. We also experimented with using GPT-4 to generate additional labels to extend our training data and found that adding GPT-4 labeled data was helpful in increasing the accuracy of the classifier for this task. For a subset of the rule paragraph sample, we also had a law student indicate if the response addressed a single or multiple commenters and the degree of disagreement between commenters that the agency was addressing (according to the agency's description in the response). The degree of opposition was initially coded on a 5-point Likert scale, but we simplified the final labels to three categories to make the task easier for the automatic classifier.

With manually annotated training data in hand, we fine-tuned several RoBERTa-base models for text classification (Liu et al. (2019); Raffel et al. (2020)).[43] Our first model is trained to detect the boundaries of responses in the preamble text of rules. The model takes a single rule paragraph as input, plus the nearest preceding non-bullet/non-list paragraph additional context. We train the model to classify each paragraph as the beginning of a response, a continuation of a response, or a non-response paragraph (the most frequent category). To extract responses at scale, we use the trained model to classify every paragraph in our sample of rules and extract responses using the predicted labels.

Table 3 gives several measures of accuracy for our response extraction procedure measured

[43]These models are small predecessors of today's LLMs. They are much more computationally efficient, but are still capable of giving very good performance on simple classification tasks with appropriate training data. The main limitation of these models is their relatively short context length of 512 tokens, or about 380 words, meaning that our classification models need to make their inferences based on fairly short chunks of input text, based on a rule of thumb of 4 characters per token and 4 tokens per 3 words of English text suggested by OpenAI. See https://platform.openai.com/tokenizer .https://platform.openai.com/tokenizer For cases where the input text is too long, we truncate from the middle, retaining the beginning and end of the text.

against a manually coded test set. The first row reports the F1 score –a metric indicating how well a model identifies true positives while minimizing both false positives and false negatives– for individual paragraphs extracted from the rule, the second row gives a “fuzzy” F1 score computed at the response level, which awards partial precision and recall scores according to the fraction of the paragraphs that overlap between the true and predicted responses, and the third row gives a “strict” F1 score where extractions are only counted as true positives if the extracted response contains exactly the same paragraphs as the manually coded response. The second column gives the equivalent measures, but only counting the “R” paragraphs where the actual response (change or no change) occurs. The test results show that our algorithm performs well at extracting response paragraphs accurately (F1=0.88), but sometimes groups paragraphs incorrectly leading to a lower strict F1 score. The intermediate fuzzy F1 scores show that the responses are mostly correct if we account for partial matches, particularly when we focus on the most important paragraphs in the response.

Table 2: Response Extraction Accuracy

Measure	**All Paragraphs**	**"R" Paragraphs Only**
Paragraph F1	0.88	0.86
Fuzzy Response F1	0.73	0.83
Strict Response F1	0.57	0.74

After responses are extracted, we use three additional classifiers to determine response-level variables indicating whether there was a policy change, the number of commenters (single or multiple), and the degree of alignment between commenters indicated in the response (single side, multiple sides, opposing sides). These classifiers take an entire response as input and produce a prediction about the response characteristics. This is a straightforward text classification task, and our large training samples give us very high accuracy, particularly on identifying policy changes (F1=0.93) and distinguishing between single and multiple commenters (F1=0.96). These scores are similar to individual human RAs (who make occasional errors).

Table 3: Response Classification Accuracy

Classifier	**Labels**	N_{Gold}	N_{GPT}	N_{Val}	**Macro F1**	**Accuracy**
Any Change	Y/N	3157	5940	631	0.93	0.94
Multiple Commenters	Y/N	3902	0	780	0.96	0.96
Multiple Sides	single/multiple/opposing	1935	0	387	0.75	0.79

3.4 Scoring comment success

We quantify outcomes for individual commenters by linking each comment to specific responses. The goal is to identify, for each response, the set of comments that the agency is responding too. This could be a single comment, many similar comments, or many different comments that all raise the same issue in different ways somewhere in the comment text. It is therefore a challenging text retrieval task. We developed a hybrid approach. First, we split comments into roughly paragraph-length chunks of text and estimate comment chunk-response level similarity scores using another RoBERTa model that is trained to place chunks of comment text and response text into a common embedding space (Xing et al. (2023)). We then take these similarity scores, along with key features of the comment, rule, and response (including the total number of comments linked to the rule, the mean and max similarity scores for that rule, and whether the response addresses a single commenter or the number of sides), and use a sample of manually labeled comment-response pairs to fit a penalized logistic regression that predicts the probability that a given comment chunk should be matched to a given response. Comment-response pairs are then assigned the maximum match probability across the text chunks within the comment. These comment-response match probabilities work well for rules with a relatively small number of comments. However, for rules with thousands of comments the accuracy falls and the probabilities tend to be less differentiated and less informative. So as a final step, we augmented our data by using Open AI's GPT-4o-mini model to detect if a comment is being discussed in a response. The GPT-4o-mini context size is large enough that most comments can be included in their entirety without chunking, and while the model still makes some mistakes, the error rate is much lower than the match probabilities when the match probabilities are close to 0.5. We collected GPT-4o-mini data for all comment-response pairs with match probabilities between 0.3 and 0.7, asking the model to assign a score between 1-5 indicating the quality of the match, which we normalize to another 0-1 score. We compute a final matching score putting 2/3rds weight on the GPT-4o-mini data where available, and 1/3rd weight on the original match probability (which still contains useful information), and link a comment to a response if the final score is greater than 0.5.

Figure 1 shows the results of an end-to-end accuracy test for three key comment-level outcome measures. To run this test we selected 100 random rules with 1-20 responses, and selected one of the comments linked to each rule. We had RAs manually replicate our entire scoring pipeline by reading through the original comment and rule documents and identifying which responses addressed the comment, and which of those responses included a policy change. The x-axes of the three plots show our main predicted measures: the number of responses linked to a comment, the number of changes linked to a comment, and the fraction of linked responses with a policy change. The y-axis of each plot shows the average value coded by the RAs for each comment. Accuracy is generally high.

Figure 1: End-to-end Comment Outcome Scoring Accuracy

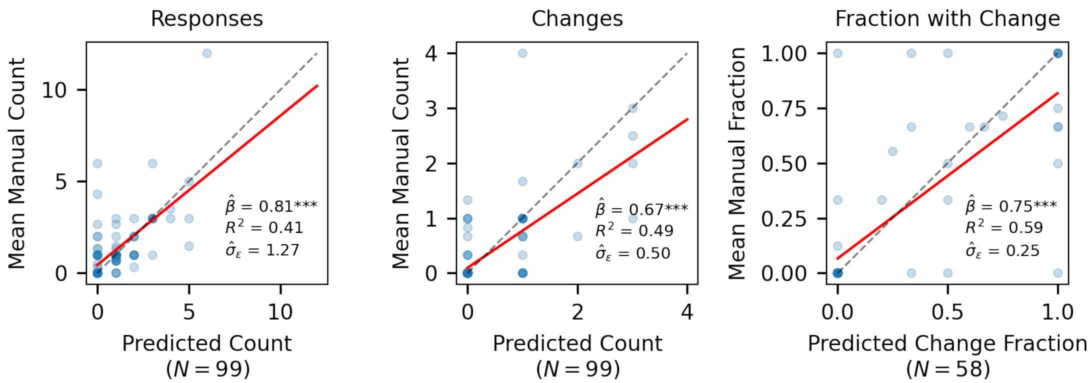

3.5 Identifying commenters

We further track individual commenting organizations across multiple rules as well as link organization commenters to other data including lobbying records, Compustat financial information, and DIME ideology scores.[44] To do so, we need to extract authorship information from comments and then link organizations across comments and to organization identifiers. Authorship information on Regulations.gov is often incomplete. Some commenting organizations are identified in an explicit metadata field, but this field is often blank or misleading (for example, identifying only one of several authors, or listing an affiliated organization, such as a school or place of work that we do not want to attribute as an author). Organization names can also appear in comment title fields (depending on how the agency formats these), though this is not consistent either. In general, the only way to be certain of authorship is to look at the content of the comment itself, which may require parsing text extracted from a PDF. Commenters often identify themselves in the first paragraph of the comment, or include organization information in the signature block at the end of the comment. In our experience, modern LLMs can easily identify comment authors with fairly high accuracy (comparable to human RAs), but using these models to extract authorship information for millions of comments would be prohibitively expensive. Instead, we rely on light-weight RoBERTa model-based classifiers to efficiently extract, classify and link organization authors. Specifically, we break the task of iden-

[44]We focus on identifying organizations that can be considered authors of the comment either because the comment is explicitly written on behalf of the organization or because the one of the comment authors has a senior/leadership role in the organization.

tifying authors in each comment into four steps: 1) Extract organization names, 2) Identify the role of organizations as authors, affiliations or other reasons for being mentioned, 3) Classify the type of organization that wrote the comment. At each stage, we used GPT-4 to generate thousands of labeled examples to use as training data, and then fine-tuned a RoBERTa model to complete the narrow task. After training, our light-weight classifiers achieve similar accuracy to GPT-4, with F1 scores of 0.95 on name extraction, 0.96 on role classification, and a macro F1 score of 0.83 on organization type classification on held out test sets of GPT-4 labeled data. The prompts used to generate the training and test data are available in Appendix F. In the final data, we identify organization comments as those that include organization names where the organizations are classified as authors. We attribute organization types based on the modal comment-level organization type classification.

Linking authorship information across comments and linking comment authors to external data requires one final linking step. We use a custom name-matching algorithm that is trained on a combination of names from lobbying filings cleaned by the Center for Responsive Politics (opensecrets.org), a nonpartisan non-profit focused on electoral transparency, with additional manual cleaning and some GPT-4 labeled data[45]. The algorithm takes all raw name strings from all our organization comments and external data files, embeds them into a common space, and uses a variant of agglomerative clustering to find groups of names that refer to the same organization. The most common string in each group (by occurrences) is assigned as the unique organization identifier.

4 Rulemaking: Six Stylized Facts

This section explores the content of the comments, the overall responsiveness of regulatory agencies in terms of policy change (i.e. adoption of a comment or rejection), the degree of explicit opposition between commenters on a rule, the degree of concentration of influence among small groups of organizations, the role of alignment with the agency in determining commenting behavior and outcomes, of ideology/bias, and of organization size. At least to the degree of generality at which we report them, the empirical regularities in this section are novel to the literature and useful to provide a set of relevant stylized facts to the reader interested in rulemaking.

A further goal is to build this set of regularities with an eye to our theoretical framework and ultimately to help separate among hard, soft and hybrid information models. While not

[45]See https://github.com/bradhackinen/nama

all results reported are individually discriminant across theories, and we make no attempt at structurally estimating the model, overall our results are consistent with a mixture of both hard and soft information in US rulemaking. Additional ancillary stylized facts, of lesser pertinence to the theory, are available in Appendix B.

Fact 1. Comments on regulations tend to be negative. *75% of comment letters express opposition to the proposed rule and 90% want at least one change made to the proposed rule.*

We begin with an important feature of our equilibrium. Recall that the agency follows its initial instinct when receiving no comment and it chooses action $a_{\emptyset} = 1$, with $a_{\emptyset}$ being the agency's default choice of action, when none of the advocates comments. In our theory most comments should therefore aim at inducing action $a = 0$, countering the proposed rule. This is also what we see in the data.

Most comment letters present some common features. First, comment letters are predominantly submitted by individuals and organizations who are critical of the proposed rule. Using our automatic classifiers, we find that 75% of comments contain at least one paragraph where the advocate expresses opposition to the proposed rule, while only 29% of comment letters contain an expression of support. Sophisticated comment letters give more detailed feedback and argumentation, sometimes supporting certain features of the proposed rule and expressing opposition to others. Second, almost all commenters –even those broadly supportive of the proposed rule– request at least one specific policy change or raise a concern that would require making changes to the proposed rule to address. With our automatic classifiers, we find that only about 10% of comment letters express support without also raising a single concern or requesting a change.

Figure 2: Comment Letter Counts by Sentiment

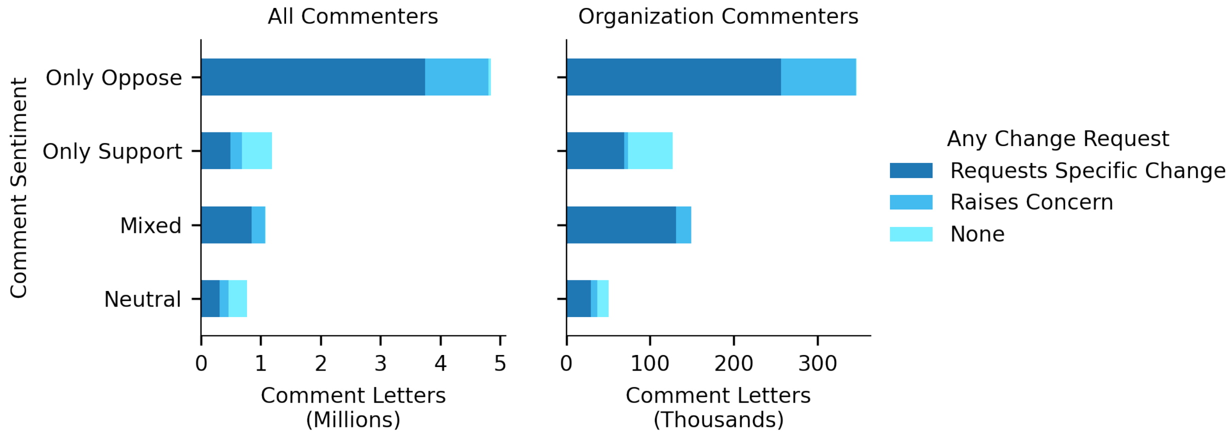

Figure 2 shows the number of comment letters in our linked sample broken down according to their general sentiment and whether they request a specific change or raise a more general concern about the proposed rule. Sentiment detection is done at the paragraph-level and then aggregated up to the whole letter: "Only Oppose" means that the comment letter contains at least one text paragraph expressing opposition to the proposed rule and no paragraphs expressing support. "Only Support" means that the comment letter contains at least one paragraph expressing support for the proposed rule and no paragraphs expressing opposition. "Mixed" means that the letter contains both paragraphs expressing support and paragraphs expressing opposition. "Neutral" means that no clear supportive or opposition language was detected anywhere in the letter.

Fact 2. A sizable share of comments contains detailed, verifiable information. *35% of comment letters authored by organizations are at least 2 pages long, and 40% of these longer comments contain citations/links to independent sources.*

Comment letters vary widely in length and sophistication. At one end of the spectrum are short statements where citizens express support –or more often opposition– to the proposed rule with little additional information, very much alike the $\hat{a}_i \in \{0, 1\}$ recommendations discussed in Section 2.3. At the other end of the spectrum are long, carefully researched comment letters presenting exhaustive scientific and legal analysis, closely resembling the hard information, verifiable η messages discussed in Section 2.2. These more sophisticated letters often contain multiple separate comments about the proposed rule, each with different supporting arguments and evidence.

We measure two simple characteristics of comment letters to quantify this pattern. Panel A of figure 3 shows the quantity of organization comments in different length categories. Even among organization-authored comments, most comments are short –less than a page in length, where each "page" is measured as 3000 characters of body text. Roughly 35% of organization comments are two or more pages long. Very long comments (50 pages or more) exist, but are infrequent. Panel B shows the fraction of organization-authored comment letters in each category where we detect at least one citation to an independent and verifiable source. Examples of these verifiable sources are published academic articles, privately sponsored scientific studies, and white papers produced by think tank organizations. The fraction of comments including one such a citation is very strongly related to comment length: very short comments rarely contain verifiable citations, but very long comments almost always do. The rate at which comment include some variant of hard information in the sense discussed in Section 2.2 therefore varies, but for comments above 2-5 pages of length this share is above 25% (per Panel B of figure 3). Data for comment letters written by individuals, as opposed to organizations, look similar, except that these documents are much shorter on average[46].

Figure 3: Comment Letter Lengths and Citation Rates

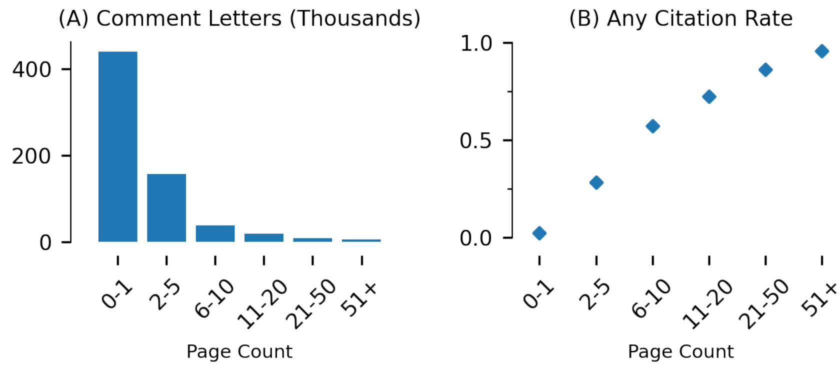

Figure 3 panel A shows the number of comments with organization authors broken down by number of pages, where one "page" is measured as 3000 characters of body text. Body text paragraphs are de-duplicated before measuring total character count. Panel B shows the fraction of comments in each bin where we detect at least one citation or link to an independent, verifiable source.

Fact 3. Agencies receive requests for change by both sides of the political spectrum. *For organizations with ideological data, only 55% of comments requesting change come from organizations that are ideologically opposed to the administration.*

In the US, agency heads are appointed by the president and the president has the power to set the regulatory agenda. This means that the ideological congruence between an organization and the current administration can be an important factor in determining who participates and how much influence an advocate has over the rulemaking process.[47] Recall further

[46]95% of comments from individuals are less than one page long. However, citation rates are similar between comments from organizations and individuals for comments of the same length.

[47]Regulated "insiders" are seen as generally opposed to more regulation and also highly influential in the pre-proposal stage (before comments are solicited, , consistent with Appendix A, Supplement C), while public interest "outsiders" benefit from the transparency of the notice and comment process–if they can overcome the challenge of becoming sufficiently informed to comment effectively (Wagner et al. 2011). However, who is an "insider" and

that, as discussed in Table 1, the hard information model predicts opposite-bias advocates, i.e. political opponents sending comments to agencies, while soft information predicts same-bias advocates. A hybrid information model, in principle, allows for both types, but Section 2.4 also shows that conditions for opposite-bias advocates (as in the hard information case) are more easily satisfied.

In order to map the data to our theoretical model, we link each commenting organization to the DIME database Bonica (2016), which provides "CF" ideology scores (Campaign Finance scores) for a wide range of organizations. CF scores are constructed by inferring the political alignment of an organization based on its campaign contributions to political candidates. An organization that donates exclusively to mainstream Democratic candidates would have a CF score of -1, while an organization that donates exclusively to mainstream Republican candidates would have a CF score of 1. Organizations that donate to both parties have a scores closer to 0, while organizations that donate to extreme candidates can have scores outside the range of -1 to 1. We report this information in Figure 4.

We find that commenting organizations are ideologically diverse, with a range of scores from far left (less than -1) to far right (greater than 1). The distribution of advocates' CF scores is bimodal, with a large peak around 0.5 (the "center-right") and a smaller peak closer to -1 (the "left")[48]. Importantly this mixture holds even within administration, suggesting that neither pure hard information nor pure soft information frameworks can rationalize the entirety of the data.

who is an "outsider" may depend on the administration in power.

[48]At the risk of over-simplification, organizations on the center-right tend to be business interests, while the left peak corresponds to a cluster of left-leaning non-profit organizations (See figure 7).

Figure 4: Distribution of CF Scores Associated with Linked Comments

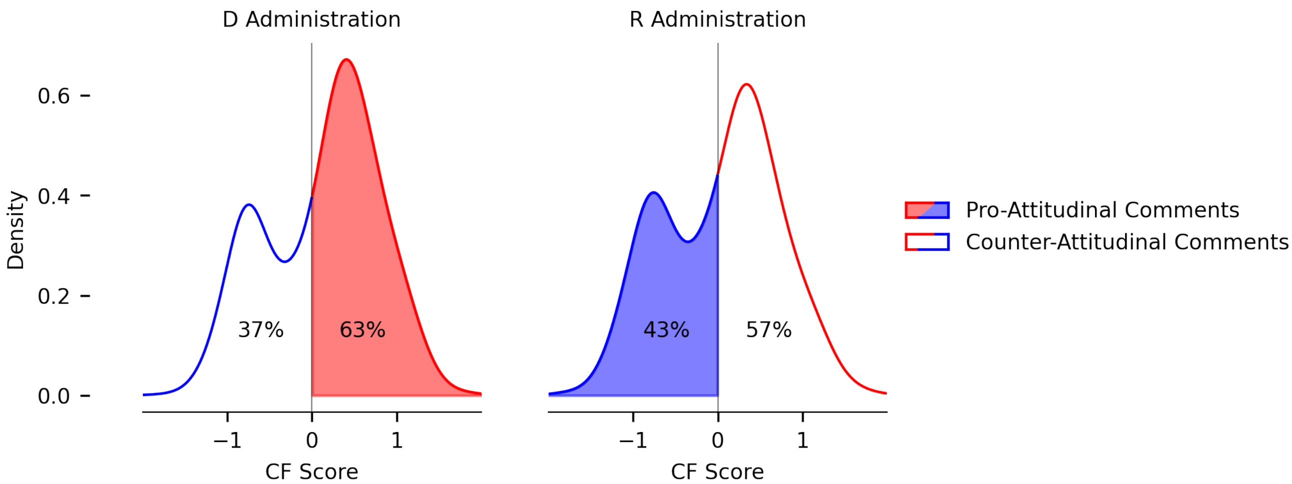

Figure 4 shows the distribution of CF scores associated with linked comments under Democrat and Republican administrations. In these distributions, each observation corresponds to a single comment, identified as a unique organization-response pair where the text of the response implies the commenters to want policy change. The CF scores come from the DIME database (Bonica, 2016), linked to comment authors by organization name. Comments are labeled "pro-attitudinal" if the organization's CF score has the same opposite sign as the part of the president and "counter-attitudinal" otherwise.

Table 1 is predictive of the fraction of comments that are "pro-attitudinal" versus "counter-attitudinal" under different informational theories. An example of a pro-attitudinal comment is one from an advocate on the left asking for change to a rule proposed under a Republican administration, or an advocate on the right asking for a change under a Democratic administration. These types of comments are "expected" (specifically by the agency receiving them), while "counter-attitudinal" cases where an advocate asks for a change in a rule proposed under an aligned administration might be somewhat more surprising. Table 1 lists several predictions. Under a Republican president, hard information models predict pro-attitudinal comments and opposite-bias advocates (that is, Democratic commenters advocating against the rule change), while under a Democratic president, hard information models would predict Republican commenters advocating against the rule change. Under a Republican president, soft information models would predict counter-attitudinal comments and same-bias advocates (that is, Republican commenters advocating against rule changes), while under Democratic administrations, soft information models would predict Democratic commenters advocating against rule changes.

Figure 4 illustrates the distribution of organization ideologies associated with comments requesting change under Democratic and under Republican administrations separately. It shows that slightly more than half (55%) of comments are pro-attitudinal in nature, while 45% of comments are counter-attitudinal –a mixture of types. As discussed in Section 2.4, these mixtures are plausible under the hybrid information model, where both verifiable messages and soft communication happen in equilibrium (and communication can be verified). Pure hard and soft information models are only consistent with one of the two subsets of the data, but not both.

Fact 4. Politically aligned advocates have higher influence. *When organizations are politically aligned with the president, they comment 15% less often, they are twice as likely to express support for a proposed rule, and have 20% higher commenting success rate when they choose to comment.*

We now explore how commenting outcomes vary according to ideology, and in particular, to the ideological congruence between commenting organization and the presidential administration. Recall that Section 2.3 discusses explicitly how information is naturally collected by opposition advocates (i.e. not aligned with the president) under hard information and by the majority advocate (i.e. aligned with the president) under soft information.

To examine the relationship between commenting outcomes and alignment, we perform a regression-based decomposition exercise. Our analysis starts from the fact that the total number of changes linked to a commenter is proportional to the number of rules for which it submits comments (the commenting rate), times the number of responses per rule (the response rate), times the fraction of those responses that include a policy change conditional on receiving a response (the success rate):

$$\text{Total Changes} = \underbrace{(\text{Rule Count})}_{\text{"Commenting Rate"}} \times \underbrace{\left(\frac{\text{Response Count}}{\text{Rule Count}}\right)}_{\text{"Response Rate"}} \times \underbrace{\left(\frac{\text{Change Count}}{\text{Response Count}}\right)}_{\text{"Success Rate"}} \tag{8}$$

We conduct this decomposition exercise to break down the variation in linked changes into differences in the number of rules an organization comments on, amount of text submitted per rule, the number of responses an organization receives per unit of text, and the fraction of responses that include a change, focusing on the responses where commenters want a change to the rule. We interpret the last fraction as a “success rate” that corresponds to the probability of a regulator changing their action in response to a comment in our theoretical model.[49] Appendix

[49]This interpretation rests on two assumptions. First, we assume that any policy change described in a linked response is an indication that the organization achieved its goal. We believe this is reasonable in the vast majority of cases. Exceptions include cases where responses address multiple commenters with different goals, or where the response is incorrectly classified as making a change. However, both of these cases are relatively uncommon in the data.

Second, and more challenging, we must assume that the count of responses linked to the organization where

B shows in its additional stylized facts that success rates hover on average around 20%, but with substantial heterogeneity depending on agency.

We use four separate Poisson regressions[50] to estimate the elasticity of each outcome with respect to the alignment between the organization and the administration, which we measure as the interaction between the CF score and a dummy variable indicating that the presidential administration is Republican:

$$\log(\text{Total Changes}_{it}) = \beta_1 \text{CF}_i \times R_t + \mu_i + \mu_t + \varepsilon_{it} \tag{9}$$

$$\log(\text{Rule Count}_{it}) = \beta_2 \text{CF}_i \times R_t + \mu_i + \mu_t + \varepsilon_{it} \tag{10}$$

$$\log(\text{Response Count}_{itr}) = \beta_3 \text{CF}_i \times R_t + \mu_i + \mu_t + \varepsilon_{itr} \tag{11}$$

$$\log(\text{Change Count}_{itr}) = \beta_4 \text{CF}_i \times R_t + \log(\text{Response Count}_{ir}) + \mu_i + \mu_t + \varepsilon_{itr} \tag{12}$$

Total Changes_{it} is the total number of changes linked to firm i in year t, and Rule Count_{it} is the number of rules that firm i commented on in year t. CF_i is the CF score of i, R_t is a dummy variable indicating that the presidential administration is Republican. We include organization and year fixed effects in all regressions, μ_i, μ_t. The structure of these regressions has some features in common with a difference-in-differences model where right-leaning organizations under a Republican president are the treatment group. Constant differences between organizations are absorbed by the organization fixed effects. Differences between administrations are absorbed by the year fixed effects. The remaining variation in outcomes is driven by differences in the differences between outcomes for left and right-leaning organizations as the party of the administration changes. The specification makes the assumption that the treatment effect is constant over time, but we believe this is a reasonable compromise given the small number of complete presidential terms available in our data (for our sample covering 2008-2022, we are missing the first three years of the Bush administration and the last two years of the Biden administration). Asymmetries between increasing and decreasing alignment cannot be identified

we detect any commenter wanting change is a reasonable estimate of the number of requests for change in the focal organization's comment letter. This is assumption is reasonable if a) we have linked comments to responses with high enough accuracy, and b) agencies reliably address all requests for change in the comments. In terms of accuracy, we believe we have achieved high enough linking accuracy to extract meaningful differences between outcomes for different organizations. This is based partly on results on test data (see section 3.4), and partly based on the fact that we see sensible results for regressions where we have a strong prior about the expected relationships (such as seeing more responses linked to larger firms). The question of whether agencies address every request for change is harder to verify in the raw data and would involve counting requests in the comment text itself, which is a difficult NLP task.

[50] We choose Poisson regression because it can estimate accurate log-log relationships in count data with zeros.

because of the year fixed effects. Therefore, each β can be interpreted as the (multiplicative) effect on that occurs when a right-leaning organization with a CF score of 1 experiences a switch from an Democratic to a Republican administration, or equivalently, the effect that occurs when a left-leaning organization with a CF score of -1 experiences a switch from a Republican administration to a Democrat administration.

Figure 5 plots the estimated coefficients for each outcome.

Figure 5: Effect of Political Congruence on Commenting Outcomes

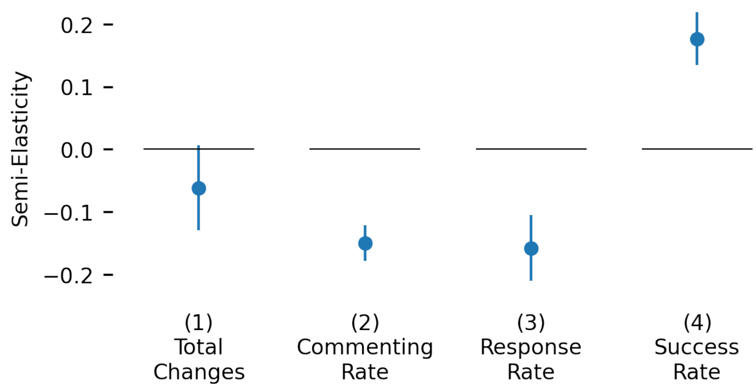

Figure 5 shows the results of a decomposition exercise that breaks down the number of changes linked to an organization into three multiplicative factors: rules per organization (the "Commenting Rate"), responses per rule (the "Response Rate"), and changes per response (the "Success Rate"). Only responses where commenters want the rule to change are included in the analysis. Each point shows the estimated coefficient (a semi-elasticity) of this factor with respect to CF score when the administration changes from a Democrat to Republican president. Positive values indicate that organizations increase in this factor under an aligned president. The estimation model assumes linearity in the log effect size as a function of CF score and symmetry between switching from Democrat to Republican and vice versa. All points are estimated conditional on year and organization fixed effects. The whiskers on each point indicate 95% confidence intervals. Standard errors are clustered by organization in columns 1-2, and by organization and rule in columns 3-4.

We find that increased political alignment with the administration has a small negative (but statistically insignificant) effect on the total number of changes linked to an organization. This small effect masks larger shifts in commenting outcomes that mostly cancel each other out. Aligned organizations are significantly less likely to comment, and receive fewer responses when they do. However, they also have higher success rates, obtaining 20% more changes per response than organizations with neutral CF scores. This is consistent with soft information per the discussion in Section 2.3.1. In the case of a Republican administration, advocate R has no incentive to mislead the agency when uninformed and R is trusted more (he recommends against the status-quo only when informed). As we see in the data, R is more likely to inform the agency and induce a change.

In Figure 6, we can see further that there are differences in the content and sentiment of comment letters submitted by organizations who are aligned with the administration. Aligned organizations submit shorter letters, are twice as likely to express only support, and are half as likely to express only opposition to the rule. Nevertheless, both aligned and non-aligned organizations have similar rates of requesting at least one change to the rule.

Figure 6: Effect of Political Congruence on Comment Letter Characteristics

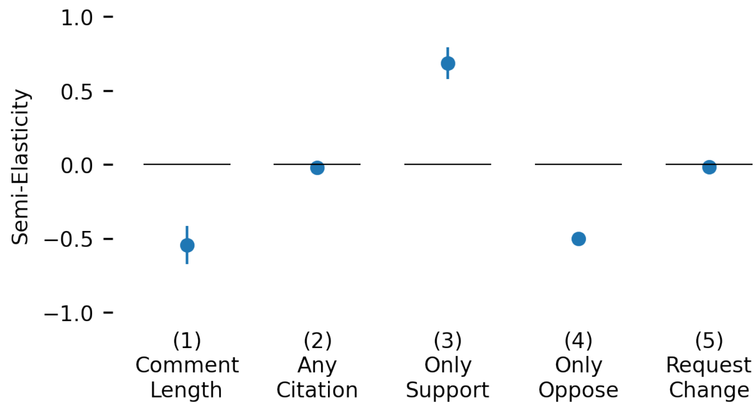

Figure 6 shows the results of five separate Poisson regressions of different variables describing comment characteristics on an interaction term between organization CF score and an indicator for whether the current president is a Republican. Comment length is measured by summing the length (in characters) of all unique body paragraphs in documents submitted by one organization for one rule. The outcome variables in columns 2-5 are binary indicators for whether the comment letter contains any citations to external documents, contains an expression of support for the proposed rule and no expressions of opposition, contains an expression of opposition to the proposed rule but no expressions of support, or requests a specific policy change. All regressions use organization-rule level data and include organization and year fixed effects. The whiskers on each point indicate 95% confidence intervals. Standard errors are clustered by organization and rule.

Fact 5. Ideologically moderate advocates have higher influence. *Left-leaning nonprofits are frequent commenters, but industry groups and other center-right organizations have 5-20% higher success rates when they choose to comment.*

While the influence of commenting organizations appears to depend on ideological congruence with the administration, there are also differences in the average level of influence between organizations across the ideological spectrum. Figure 7 shows the organizations in our commenting data that have linked CF scores.

Figure 7: Commenting Success and Ideology

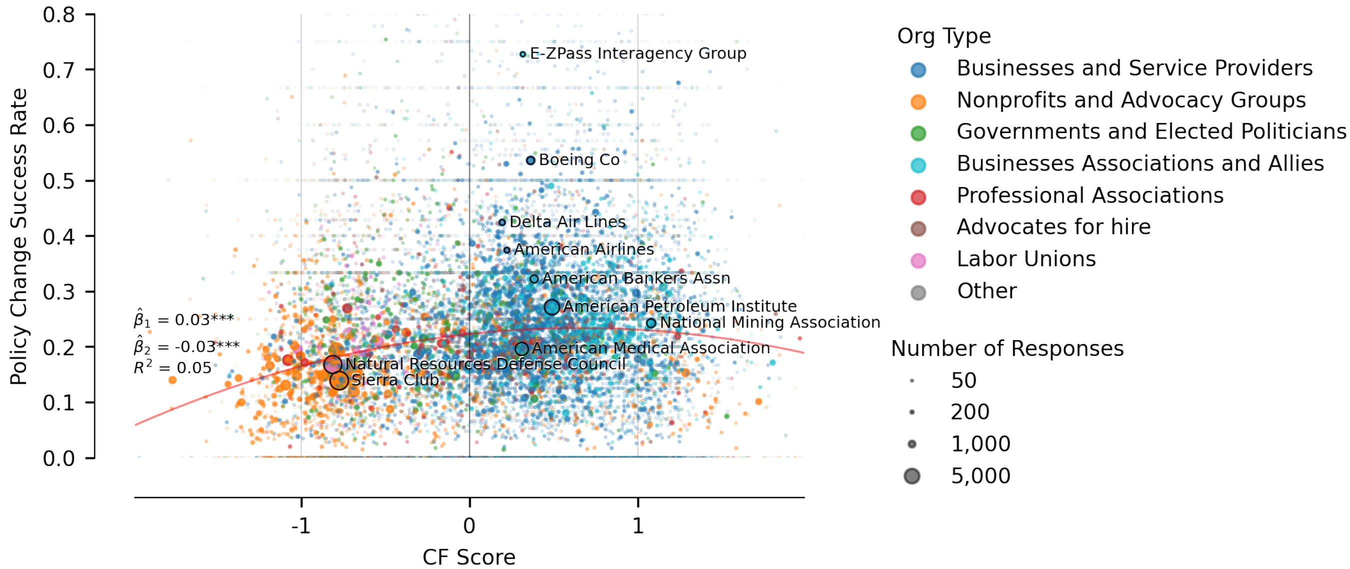

Figure 7 shows all organizations in our data that also appear the DIME database Bonica (2016). Each point represents one organization, which may be associated with multiple DIME ids if the names are assigned to the same organization during name matching. The x-axis indicates DIME CF score, a measure of ideology inferred from patterns of political donations (when multiple DIME scores are linked to an organization we take the mean of all linked scores). A score of -1 roughly corresponds to the average congressional Democrat, while a score of 1 corresponds to the average congressional Republican. The size of each point indicates the number of responses linked to each organization through its comments, and the y-axis indicates the policy change "success rate", defined as the fraction of responses that include a policy change among all responses linked to the organization where commenters want a change. Points are colored according to our own classification of the organizations. The line of best fit is calculated using a quadratic WLS regression where each organization-level observation is weighted by its response count. *** indicates statistical significance at the $\alpha = 0.01$ level using heteroskedasticity-robust (HC3) standard errors.

Commenting organizations are ideologically diverse, with a range of scores from far left (less than -1) to far right (greater than 1). Businesses and business groups are frequent commenters and tend to have moderate CF scores between 0 and 1. Nonprofit advocacy groups are also active commenters and tend to have more extreme CF scores, with a large cluster on the left and a smaller number on the right. The vertical axis indicates each organization's policy change "success rate", a measure of each organization's ability to persuade agencies to make desired policy changes (as discussed in section 6). The quadratic line of best fit shows that moderate and center-right organizations tend to have the highest success rates, while more extreme organizations have lower success rates.

This finding is supported by our hybrid model. Recall that Proposition 2.3 first observes that the agency is more friendly to a comment by an advocate, the more moderate the advocate. Specifically, part (ii) of the proposition implies that a more moderate advocate is more likely to induce rubber-stamping when a thorough review would have taken place, or to induce such a review when the narrative would have been dismissed.

We can also estimate parametric relationships between CF score and the same four influence measures shown for ideological alignment, using separate quadratic Poisson regressions of the form:

$$\log(\text{Total Changes}_{it}) = \beta_1 \text{CF}_i + \beta_2 \text{CF}_i^2 + \mu_t + \varepsilon_{it} \tag{13}$$

$$\log(\text{Rule Count}_{it}) = \beta_2 \text{CF}_i + \beta_2 \text{CF}_i^2 + \mu_t + \varepsilon_{it} \tag{14}$$

$$\log(\text{Response Count}_{itr}) = \beta_3 \text{CF}_i + \beta_2 \text{CF}_i^2 + \mu_r + \varepsilon_{itr} \tag{15}$$

$$\log(\text{Change Count}_{itr}) = \beta_4 \text{CF}_i + \beta_2 \text{CF}_i^2 + \log(\text{Response Count}_{ir}) + \mu_r + \varepsilon_{itr} \tag{16}$$

Here the third regression uses firm-rule level observations, conditional on the firm commenting on the rule. Response Count$_{itr}$ is the number of responses in rule r (published in year t) linked to firm i where the commenters are described as wanting a change. The final regression is also at the firm-rule level, but it conditions on the firm receiving at least one response. These last two regression also include rule-level fixed effects, μ_r. Figure 8 illustrates the predicted curves implied by the estimated coefficients in the four equations (see Table 12 in Appendix for the full regression results).

Figure 8: Ideological Alignment Relative Influence Decomposition

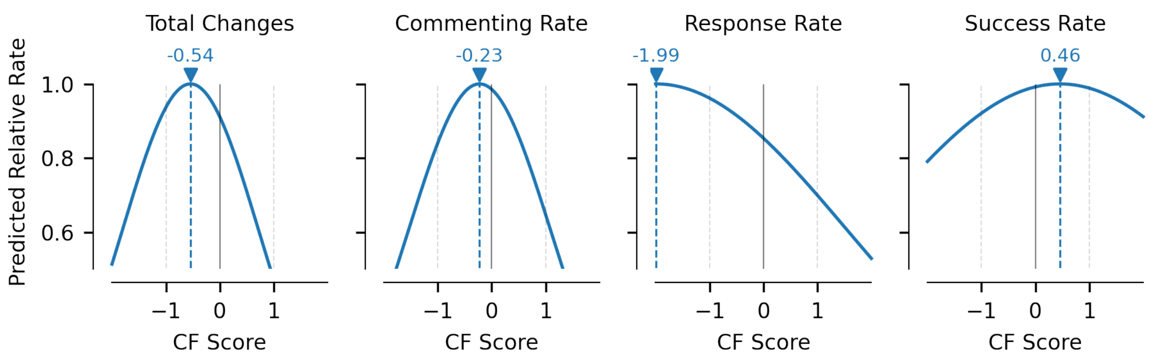

Figure 8 shows predicted Total Changes, Commenting Rates, Response Rates and Success Rates as a function of organization CF score using separate quadratic Poisson regressions. The peak of each curve is normalized to 1. The first two regressions include year fixed effects, while the latter two include rule fixed effects. The sample includes all organizations with linked DIME data and at least one comment letter in our sample.

The results show that left-leaning organizations comment on the most rules (per-organization) and receive the most responses[51]. However, center-right organizations tend to have the highest success rates. As can bee seen visually in Figure 7, organizations in this range tend to be businesses and business allies. Meanwhile, the cluster of non-profit organizations on the left tends to have 5-20% lower success rates, depending on how extreme their CF ideology scores are.

Fact 6. Larger advocates comment more, but not with higher success rates. *Large public companies comment considerably more (both in terms of number of comment letters filed and in the length of each comment) than small public companies. However, large companies do not have a higher policy change success rate, conditional on sending a comment, relative to small companies.*

We explore now whether large organizations are more influential using Compustat data to compare commenting outcomes for large and small publicly traded firms. Our data allows us to distinguish between several different stories about how size might related to influence. For size of a firm, we use various proxies, including number of employees, market value, as well as total lobbying expenditures by the firm over the 2008-2022 sample period, which is known to correlate with organization size but also captures additional elements of political influence (Bombardini and Trebbi, 2020)[52]. We again separately measure how frequently an organization submits comment letters, the number of linked agency responses involving a request for change, and the fraction of these responses that actually include a policy change.

Figure 9 plots our linked Compustat firms showing the number of responses, overall success rate (total number of changes divided by total number or responses), and number of employees. Its clear that large firms comment more and receive more responses. However, the relationship between firm size and policy change success rate is weak. Firms have a wide range of success rates, and size explains little of the variation. For example, WPX Energy and Entergy Arkansas, Inc are of comparable size, but have rates of change of 62% and 4%, respectively. Importantly, Section 2.2.1 can help rationalize these facts quite naturally. If the agency expects large firms to comment, then unilateral expectation conformity requires large firms to comment in equilibrium, even if their success rate happens to be not higher than for smaller firms. This is simply because otherwise the agency would interpret their silence as support for the rule (action $a = 1$). Further, as information acquisition is costly, small firms simply free ride on the large ones. As it can be seen, this appears a more benevolent (but surely not exclusive) interpretation relative to

[51] Note that the vertical position of the curves is not identified because of the rule fixed effects. We normalize the curves by setting the peak to one, so that they indicate predicted rates relative to the peak predicted rate

[52] We do not wish to argue that lobbying amount is exogenous to commenting activity, but we include it to check if firms that invest in more lobbying (usually larger, highly regulated firms) have systematically different commenting outcomes relative to firms that spend little on lobbyists.

Stiglerian capture theory with respect to large firms.

Figure 9: Commenting Success and Firm Size

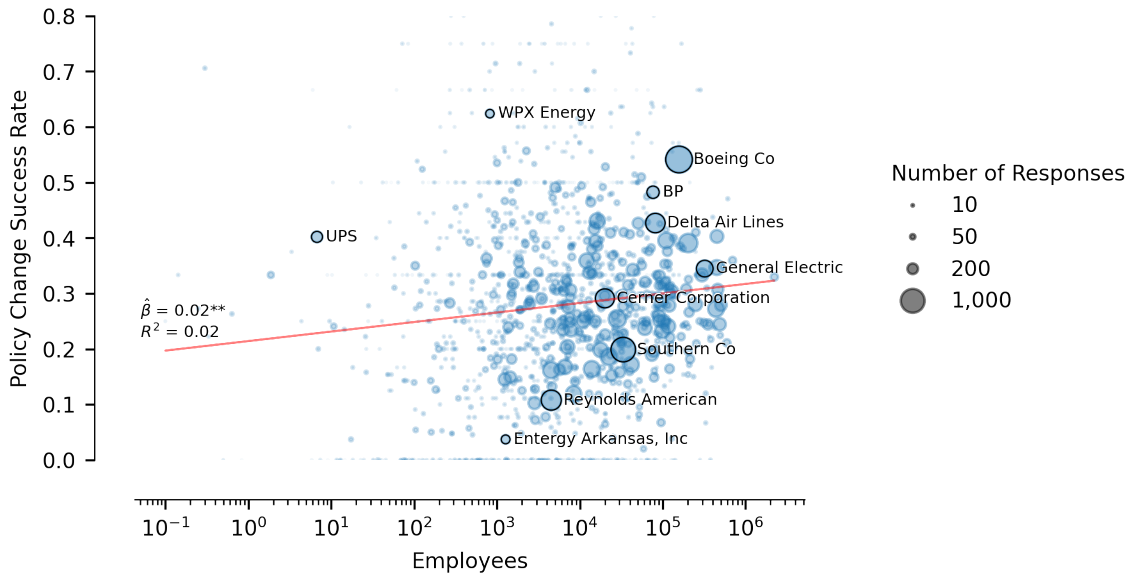

Figure 9 shows all business organizations in our data that also appear in Compustat data on publicly traded firms. Each point represents one organization, which may aggregate multiple Compustat GVKEYs if the compustat firms are assigned to the same organization during name matching. The x-axis indicates the average number of employees for the 2008-2022 period. The size of each point indicates the number of responses linked to each organization through its comments, and the y-axis indicates the fraction of linked responses that include a policy change, focusing on the set of responses where commenters want a policy change. The line of best fit is calculated using a linear WLS regression where each organization-level observation is weighted by its response count. ** indicates statistical significance at the $\alpha = 0.05$ level using heteroskedasticity-robust (HC3) standard errors.

We are particularly interested in how response and success rates vary with size after accounting for selection into commenting on certain rules. In other words, we wish to measure whether, on the same rule, a small firm and a large firm tend to achieve the same outcomes when they comment. Size s_i also plays a role in our model, for example as explicitly discussed in in Proposition 2.1, part (ii). In the hard information model, characteristics of the advocate like s_i do not influence the agency beyond the verifiable information η.

We estimate conditional elasticities corresponding to each of these terms with respect to firm size using four separate Poisson regressions:

$$\log(\text{Total Changes}_{it}) = \beta_1 \log(\text{Firm Size}_{it}) + \mu_t + \varepsilon_{it} \tag{17}$$

$$\log(\text{Rule Count}_{it}) = \beta_2 \log(\text{Firm Size}_{it}) + \mu_t + \varepsilon_{it} \tag{18}$$

$$\log(\text{Response Count}_{itr}) = \beta_3 \log(\text{Firm Size}_{it}) + \mu_r + \varepsilon_{itr} \tag{19}$$

$$\log(\text{Change Count}_{itr}) = \beta_4 \log(\text{Firm Size}_{it}) + \log(\text{Response Count}_{itr}) + \mu_r + \varepsilon_{itr} \tag{20}$$

The first two regressions use data that has been aggregated to the firm-year level. We omit firm-year observations for which no firm size information is available (for example, if the firm did not exist). μ_t indicates year fixed effects. The coefficient on the $\log(\text{Response Count}_{itr})$ in the last regression is constrained to one, so that β_4 measures the elasticity of the ratio of changes per response with respect to firm size (a Poisson rate regression).

Figure 10: Firm Size Influence Decomposition

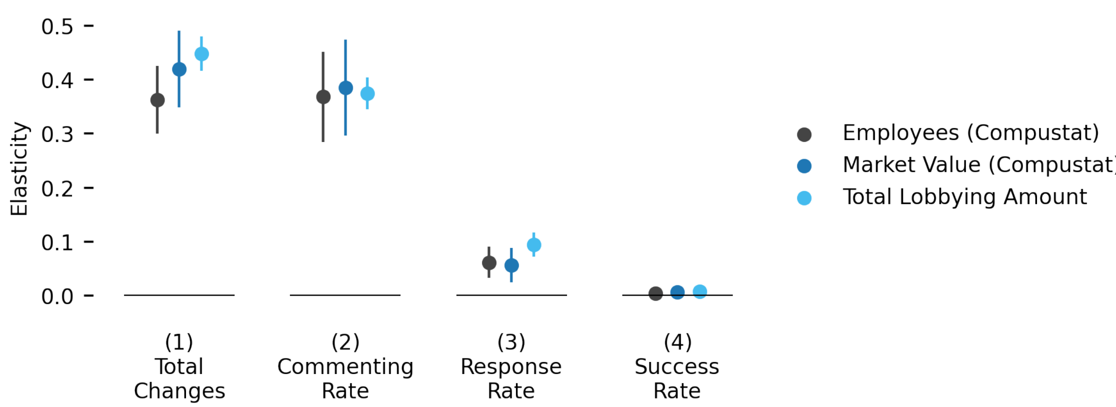

Figure 10 shows the results of a decomposition exercise that breaks down the number of changes linked to a firm into three multiplicative factors: rules per firm (the "Commenting Rate"), responses per rule (the "Response Rate"), and changes per response (the "Success Rate"). Only responses where commenters want the rule to change are included in the analysis. Each point shows the estimated elasticity of this factor with respect to differences in firm size, as measured by employees, market value, and total expenditures on lobbying firms over the 2008-2022 period. In the cases where multiple compustat GVKEYs are linked to the same organization name, we compute organization-level firm size as the sum of the sizes associated with the linked GVKEYs. Columns 1-2 are estimated using aggregated firm-year level counts with year fixed effects, while columns 3-4 are estimated at the firm-rule level conditioning on rule fixed effects. The whiskers on each point indicate 95% confidence intervals. Standard errors are clustered by firm in columns 1-2, and by firm and rule in columns 3-5.

Figure 10 shows the estimated $\beta_1, ..., \beta_4$ elasticities using three different measures of firm size (full regression tables are available in Appendix C). Large firms are linked to many more policy changes than small firms. On average, a firm that has 10% higher market value with have 4% more linked changes. However, differences between small and large firms are driven almost entirely by differences in the number of comments submitted. We see this in the high elasticity of the commenting rate (indicating more comment letters submitted by large firms) and the positive elasticity on the response rate, indicating that each comment letter contains more substantive requests that the agency needs to address. Another striking fact in Figure 10 is that when conditioning on rule fixed effects, the elasticity of the success rate with respect to firm size is less than 0.01 for all measures of firm size, including lobbing expenditures. This is, however, not inconsistent with our theory. As mentioned above, lack of sensitivity of success rates to firm size is discussed in our Proposition 2.1, part (ii).

We finally explore how firm size correlates with comment characteristics, including its sentiment. We run organization-rule level Poisson regressions and estimate the relationship between organization size and comment length, whether the comment contains any citations to verifiable sources, whether the letter contains only support for the proposed rule, or only expresses opposition, and whether the letter contains any requests for specific changes.[53] Each coefficient is estimated with a separate regression of the form:

$$\log(y_{itr}) = \beta \log(\text{Firm Size}_{it}) + \mu_r + \varepsilon_{itr} \tag{21}$$

where y_{itr} is the given outcome variable for firm i commenting on rule r that was published in year t.

We find that large firms write longer letters and are more likely to include citations to verifiable sources (a proxy for hard information). Large firms are also less likely to express only opposition to the proposed rule, though they have similar rates of expressing only support, and similar probability of requesting at least one specific change. These results suggest that large firms write more detailed and sophisticated comments than small firms –making it all the more surprising that they have similar success as smaller firms on a per-response basis.

[53]Here we define a firm's "comment letter" as the set of all unique paragraphs linked to that firm by authorship. This means that comments co-authored by two firms will appear twice, but duplicated comment letters linked to the same firm will only appear once.

Figure 11: Firm Size and Comment Letter Characteristics

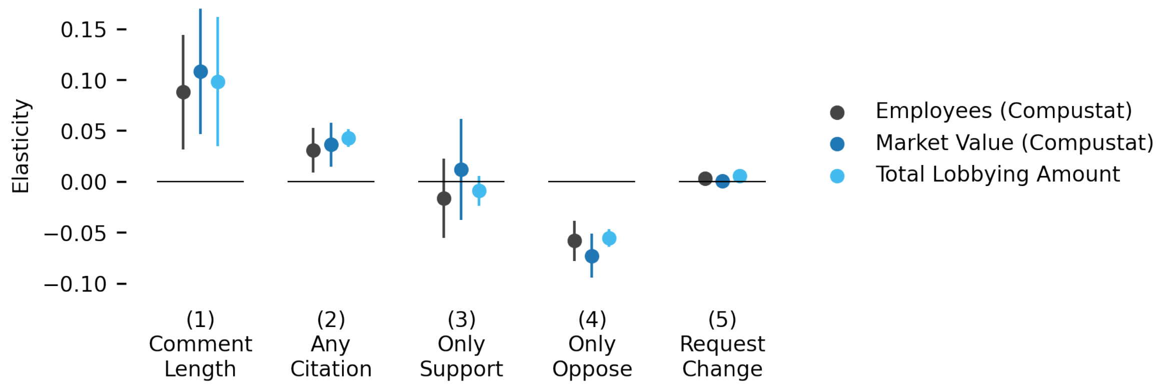

Figure 11 shows the results of 15 separate Poisson regressions of five different variables describing comment letter characteristics on three different measures of firm size. Comment length is measured by summing the length (in characters) of all unique body paragraphs in documents submitted by one organization for one rule. The outcome variables in columns 2-5 are binary indicators for whether the comment letter contains any citations to external documents, contains an expression of support for the proposed rule and no expressions of opposition, contains an expression of opposition to the proposed rule but no expressions of support, or requests a specific policy change. All regressions use firm-rule level data and include rule fixed effects. The whiskers on each point indicate 95% confidence intervals. Standard errors are clustered by firm and rule.

5 Conclusions

This paper investigates theoretically and empirically the process of rulemaking. The goal of the article is provide a new perspective on this important area of quasi-legislative government intervention, where both robust theoretical guidance and general stylized facts are still sparse.

We first present a theoretical model of rulemaking. Based on the predictions of this framework, we investigate whether theories of hard versus soft information have higher explanatory power when compared to a set of empirical regularities that we uncover, using data from the United States. We show that theories which present an hybrid (soft-hard) structure align better with the data. Both pure hard information and pure soft information models fail to capture a number of important stylized facts about the process of notice-and-comment regulation. Hard information is part of the process of public commenting: comments contain verifiable information and messages are often pro-attitudinal. Yet advocate influence also depends on political position and alignment with the administration of the advocate, suggesting also a role for soft information and cheap talk.

Future research should build on the analysis presented in this paper to assess the welfare consequences of public commenting and to identify if our results have external validity in other regulatory contexts, such as, for example, the European Commission public consultation process.

References

Aghion, P., Bergeaud, A., and Van Reenen, J. (2021). The impact of regulation on innovation. Technical report, National Bureau of Economic Research.

Alesina, A., Battisti, M., and Zeira, J. (2018). Technology and labor regulations: theory and evidence. *Journal of Economic Growth*, 23:41–78.

Alonso, R. and Câmara, O. (2016). Persuading voters. *American Economic Review*, 106(11):3590–3605.

Austen-Smith, D. (1995). Campaign contributions and access. *American Political Science Review*, 89:566–581.

Austen-Smith, D. and Wright, J. (1992). Competitive lobbying for a legislator's vote. *Social Choice and Welfare*, 9:229–257.

Austen-Smith, D. and Wright, J. (1994). Counteractive lobbying. *American Journal of Political Science*, 38(1):25–44.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401.

Battaglini, M. and Coate, S. (2007). Inefficiency in legislative policy-making: A dynamic analysis. *American Economic Review*, 97(1):118–149.

Battaglini, M. and Coate, S. (2008). A dynamic theory of public spending, taxation, and debt. *American Economic Review*, 98(1):201–236.

Becker, G. S. (1983). A theory of competition among pressure groups for political influence. *The quarterly journal of economics*, 98(3):371–400.

Bellodi, L., Morelli, M., Spenkuch, J. L., Teso, E., Vannoni, M., and Xu, G. (2025). Personnel is policy: Delegation and political misalignment in the rulemaking process. *Mimeo UC Berkeley*.

Bertrand, M., Bombardini, M., Fisman, R., Hackinen, B., and Trebbi, F. (2021). Hall of mirrors: Corporate philanthropy and strategic advocacy. *The Quarterly Journal of Economics*, 136(4):2413–2465.

Bombardini, M. and Trebbi, F. (2020). Empirical models of lobbying. *Annual Review of Economics*, (12):391–413.

Bombardini, M., Trebbi, F., and Zhang, M. B. (2025). Measuring the costs and benefits of regulation. *Annual Review of Economics*, 17.

Bonica, A. (2016). Database on ideology, money in politics, and elections: Public version 2.0. *Computer file, Stanford University Libraries. data. stanford. edu/dime (accessed June 30, 2020)*.

Caillaud, B. and Tirole, J. (2007). Consensus building: How to persuade a group. *American Economic Review*, 97(5):1877–1900.

Calomiris, C. W., Mamaysky, H., and Yang, R. (2020). Measuring the cost of regulation: A text-based approach. Technical report, National Bureau of Economic Research.

Carpenter, D. P., Dagonel, A., Judge-Lord, D., Kenny, C. T., Libgober, B., Rashin, S., Waggoner, J., and Yackee, S. W. (2024). Inequality in agency rulemaking.

Chang, S., Kalmenovitz, J., and Lopez-Lira, A. (2023). Follow the pipeline: Anticipatory effects of proposed regulations. *Available at SSRN 4360231*.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451.

Croley, S. (1995). Making rules: An introduction.

Cukierman, A. and Tommasi, M. (1998). When does it take a nixon to go to china? *American Economic Review*, 88(1):180–197.

Davis, S. J. (2017). Regulatory complexity and policy uncertainty: headwinds of our own making. *Becker Friedman Institute for Research in economics working paper*, (2723980).

Dekel, E., Jackson, M. O., and Wolinsky, A. (2008). Vote buying: General elections. *Journal of Political Economy*, 116(2):351–380.

Dewatripont, M. and Tirole, J. (1999). Advocates. *Journal of political economy*, 107(1):1–39.

Dewatripont, M. and Tirole, J. (2005). Modes of communication. *Journal of political economy*, 113(6):1217–1238.

Dwidar, M. A. (2022). Coalitional lobbying and intersectional representation in american rulemaking. *American Political Science Review*, 116(1):301–321.

Furlong, S. R. (1997). Interest group influence on rule making. *Administration & Society*, 29(3):325–347.

Garicano, L., Lelarge, C., and Van Reenen, J. (2016). Firm size distortions and the productivity distribution: Evidence from france. *American Economic Review*, 106(11):3439–79.

Gavazza, A. and Lizzeri, A. (2009). Transparency and economic policy. *Review of Economic Studies*, 76(3):1023–1048.

Grossman, G. and Helpman, E. (1994). Protection for sale. *American Economic Review*, 84(4):833–850.

Grossman, G. M. and Helpman, E. (2001). *Special interest politics*. MIT press.

Hall, R. and Deardorff, A. (2006). Lobbying as legislative subsidy. *American Political Science Review*, 100(1):69–84.

Huneeus, F. and Kim, I. S. (2018). The effects of firms' lobbying on resource misallocation.

Kalmenovitz, J. (2023). Regulatory intensity and firm-specific exposure. *The review of financial studies*, 36(8):3311–3347.

Kalmenovitz, J. and Chen, J. (2024). Regulatory Similarity. *Journal of Law and Economics, forthcoming.*

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kang, K. (2016). Policy influence and private returns from lobbying in the energy sector. *The Review of Economic Studies*, 83(1):269–305.

Kerr, W. R., Lincoln, W. F., and Mishra, P. (2014). The dynamics of firm lobbying. *American Economic Journal: Economic Policy*, 6(4):343–379.

Krishna, V. and Morgan, J. (2001a). A model of expertise. *Quarterly Journal of Economics*, 116(2):747–775.

Krishna, V. and Morgan, J. (2001b). A model of expertise. *The Quarterly Journal of Economics*, 116(2):747–775.

Krishna, V. and Morgan, J. (2004). The art of conversation: eliciting information from experts through multi-stage communication. *Journal of Economic theory*, 117(2):147–179.

Lawson, G. (1994). The rise and rise of the administrative state. *Harvard Law Review*, 107(6):1231–1254.

Libgober, B. (2020). Meetings, comments, and the distributive politics of rulemaking. *Quarterly Journal of Political Science*, 15:449–481.

Lindbeck, A. and Weibull, J. (1987). Balanced-budget redistribution as the outcome of political competition. *Public Choice*, 52(3):273–297.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach.

Lizzeri, A. and Persico, N. (2005). A drawback of electoral competition. *Journal of the European Economic Association*, 3(6):1318–1348.

Lowande, K. and Augustine Potter, R. (2021). Congressional oversight revisited: Politics and procedure in agency rulemaking. *The Journal of Politics*, 83(1):401–408.

Magat, W., Krupnick, A. J., and Harrington, W. (2013). *Rules in the making: A statistical analysis of regulatory agency behavior*. RFF Press.

Mashaw, J. L. and Merrill, R. A. (1985). *Administrative Law: The American Public Law System. Cases and Materials*. St. Paul MN: West Publishing.

Maskin, E. and Tirole, J. (2004). The politician and the judge: Accountability in government. *American Economic Review*, 94:1034–1054.

Maskin, E. and Tirole, J. (2019). Pandering and pork-barrel politics. *Journal of Public Economics*, 176:79–93.

McCarty, N. and Poole, K. (1998). An empirical spatial model of congressional campaigns. *Political Analysis*, 7(1):1–30.

Naughton, K., Schmid, C., Yackee, S. W., and Zhan, X. (2009). Understanding commenter influence during agency rule development. *Journal of Policy Analysis and Management*, 28(2):258–277.

O'Connell, A. J. (2008). Political cycles of rulemaking: An empirical portrait of the modern administrative state. *Virginia Law Review*, pages 889–986.

Pavan, A. and Tirole, J. (2024). Expectation conformity in strategic cognition. Mimeo.

Peltzman, S. (2022). Public opinion about regulation. *The Journal of Law and Economics*, 65(S2):S327–S353.

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.

Schnakenberg, K. E. (2017). Informational lobbying and legislative voting. *American Political Science Review*, 61(1):129–145.

Singla, S. (2023). Regulatory costs and market power. *mimeo LBS*.

Trebbi, F., Zhang, M. B., and Simkovic, M. (2026). The cost of regulatory compliance in the united states. *Review of FInancial Studies*.

Wagner, W., Barnes, K., and Peters, L. (2011). Rulemaking in the shade: An empirical study of epa's air toxic emission standards. *Administrative Law Review*, pages 99–158.

Xing, L., Hackinen, B., and Carenini, G. (2023). Tracing influence at scale: A contrastive learning approach to linking public comments and regulator responses. *arXiv preprint arXiv:2311.14871*.

Yackee, J. W. and Yackee, S. W. (2006). A bias towards business? assessing interest group influence on the us bureaucracy. *The Journal of Politics*, 68(1):128–139.

Yackee, S. W. (2006). Sweet-talking the fourth branch: The influence of interest group comments on federal agency rulemaking. *Journal of Public Administration Research and Theory*, 16(1):103–124.

Yackee, S. W. (2019). The politics of rulemaking in the united states. *Annual Review of Political Science*, 22(1):37–55.

A Supplements to the theory

Supplement A. Characterization of equilibria in hard-information game

Suppose $a_\emptyset = 1$. Because $\min\{-b, -b_R\} < -b$ for all b, the sets of states of nature for which the two advocates' collections of information impact the agency's decision are nested. Let

$$\alpha \equiv \frac{1}{v_L} E_b[\int_{-\infty}^{\min\{-b,-b_R\}} [-(b_L + \omega)s_R] dF(\omega)].$$

Note that $\alpha \in (0, 1]$ and $\alpha = 1$ iff $b_R \leq b$ for all b. The gain from collecting information is

$$\left(1 - \alpha p_R\right) p_L v_L - c(p_L) - c_0,$$

for advocate L and

$$(1 - p_L) p_R v_R - c(p_R) - c_0,$$

for advocate R. There are three possible equilibria:

1. ***L* advocacy**: $p_L^* = (c')^{-1}(v_L) > 0$ and $p_R^* = 0$, when

$$p_L^* v_L - c(p_L^*) \geq c_0 \geq \max_{p_R} (1 - p_L^*) p_R v_R - c(p_R).$$

2. **Two-sided advocacy**: $p_L^*, p_R^* > 0$ solve:

$$c'(p_L) = \left(1 - \alpha p_R\right) v_L,$$
$$c'(p_R) = (1 - p_L) v_R,$$

   and satisfy

$$c_0 \leq \left(1 - \alpha p_R^*\right) p_L^* v_L - c(p_L^*),$$
$$c_0 \leq (1 - p_L^*) p_R^* v_R - c(p_R^*).$$

   Multiple pairs of admissible (p_L^*, p_R^*) may exist in this case, always including one with $p_L^* > p_R^*$ whenever $s_L \geq s_R$.

3. ***R* advocacy**: $p_R^* = (c')^{-1}(v_R) > 0$ and $p_L^* = 0$, when

$$p_R^* v_R - c(p_R^*) \geq c_0 \geq \max_{p_L}(1 - \alpha p_R^*) p_L v_L - c(p_L).$$

Supplement B. Multiple advocates on the same side.

Collective action problem. There are n_L interest groups on the left and n_R interest groups on the right. Interest groups/advocates are indexed by i, that determines both the size $\{s_i\}_{i \in \{1,\ldots,n_L\}}$ and the bias $\{b_i\}_{i \in \{1,\ldots,n_L\}}$ of the group[54]. Preferences of group i are given by:

$$U_i = (b_i + \omega) s_i a.$$

All interest groups have the same cost function $C(p_i)$. We let I and J denote the subset of interest groups that are active in searching and commenting, on the left- and right- sides, respectively; $i \in I \cup J$ if and only if $p_i > 0$. We assume that the interest groups act independently: they search and comment simultaneously[55]. This implies that multiple interest groups may issue the same (successful or unsuccessful) comment concomitantly. We must therefore distinguish a "successful" comment (the agency's decision goes in the direction of the comment) from a "pivotal" one, which coincides with a successful one only if only one advocate on the same side comments[56].

We look for an equilibrium in which only left-leaning groups comment ($I \neq \emptyset$ while $J = \emptyset$); having also information acquisition by some large right-wing interest groups would not alter the equal-success-rate result below. All comment when there is a positive probability that the comment will be impactful, i.e, when $\omega < 0$. As earlier, the agency selects $a_\emptyset = 1$ when receiving no information: $b + \omega_\emptyset > 0$ as the expected value $\omega_\emptyset$ in the absence of comment is (strictly) positive when only left-side interest groups search. While the right-side advocates do not search

[54]For instance, suppose there is a unit carbon tax τ , that the agency is thinking of removing (or decreasing or not enforcing). Industry i 's stake is then τs_i where s_i captures the production size, or the industry's emissions intensity, or else the farms' crop choice.

[55]Whether interest groups lobby independently or in coordination depends on the extent of their common interest (Bombardini-Trebbi 2012). We here analyze the subgame that unfolds once the structure of lobbying has been determined.

[56]This distinction already arises with one interest group on each side, when advocates L and R both collect information.

(under an assumption that generalizes (2)), left advocate i, when active[57], solves:

$$max_{p_i}\{p_i(1-p_{-i})v_i - c(p_i)\} \Rightarrow c'(p_i) = (1-p_{-i})v_i, \tag{22}$$

where p_{-i} denotes the probability that the other active left-side interest groups find the information:

$$1-p_{-i} \equiv X_{i'\neq i}(1-p_{i'}),$$

and v_i is i's expected partisanship-adjusted value from issuing a successful comment. Let

$$v_i \equiv \int_0^{+\infty}\int_{-\infty}^{-b}[-(b_i+\omega)s_i]dF(\omega)dG(b).$$

Then for $i \in I$,

$$max_{p_i}\{p_i(1-p_{-i})v_i - c(p_i)\} \geq c_0, \tag{23}$$

while the weak inequality hold reverse if $i \notin I$. Moreover, advocates $j \in \{1,...,n_R\}$ do not acquire information if

$$c_0 \geq max_{j\in\{1,...,n_R\}} max_{p_j}\{p_j(1-p_L)v_j - c(p_j)\}, \tag{24}$$

where

$$v_j \equiv \int_0^{+\infty}\int_{-\infty}^{-min\{-b,-b_j\}}[-(b_j+\omega)s_j]dF(\omega)]dG(b).$$

Proposition A.1. ***(collective action).***

Suppose that information is hard and advocates differ solely through their size and their partisanship; assume further that (23) and (24) are satisfied, so that some left-leaning advocates comment while right-leaning advocates do not. Then,

(i) the success rates of same-side active advocates are identical,

(ii) the common success rate is smaller, the more biased the agency (in the sense of FOSD).

Proof of Proposition A.1

(i) The probability of a comment by advocate i is equal to the equilibrium probability p_i of

[57]We take an arbitrary equilibrium. If the s_i and the $\mid b_i \mid$ are both ranked in decreasing order, it may be natural to look for an equilibrium in which only the left-side, highest-stakes interest groups ($i \in \{1,...,m\}$ with $m \leq n_L$) are active, i.e., search and possibly comment- there then always exists such an equilibrium. But the properties below are independent of equilibrium selection.

finding information times the probability $F(0)$ that the information is favorable to the left-side cause for a positive measure of bs. The probability that the comment is made and accepted is $p_i Prob(b+\omega<0) = p_i \int_0^{+\infty} F(-b)dG(b)$

Advocate i's success rate, $[p_i \int_0^{+\infty} F(-b)dG(b)]/[p_i F(0)]$, is therefore independent of the interest group's stake s_i.

(ii) Because the probability of being pivotal when informed, $F(-b)$, is decreasing in b, if H FOSD G, then $\int_0^{+\infty} F(-b)dH(b) < \int_0^{+\infty} F(-b)dG(b)$. ■

Finally, consider the case of two interest groups, i and i', and suppose that the equilibrium $\{p_i, p_{i'}\}$ is stable. Then, when s_i grows, $p_{i'}$ has to decrease: with two groups, the intuitive property that an increase in one's stake crowds out the other interest group's effort is verified.

Remark: An environment in which success rates are not equalized is when interest groups have access to different search technologies. Suppose that at cost $c(\rho_i)$, interest group i receives a signal ω, that we can take to be the posterior mean, according to symmetric distribution $F(\omega;\rho_i)$ where ρ_i is a rotation parameter. The rotation point is 0[58]: $F_\rho(\omega;\rho) > 0$ for $\omega < 0$ and $F_\rho(\omega;\rho) < 0$ for $\omega > 0$ (and $F_\rho(0;\rho) \equiv 1/2$). Maximum information corresponds to $\rho = +\infty$ and no information (a spike at 0) to $\rho = 0$. To fix ideas, consider two interest groups with precisions $\rho_H > \rho_L$. In comparative statics (there is a single advocate with this precision), $2\int_0^{+\infty} F(-b;\rho_H)dG(b) > 2\int_0^{+\infty} F(-b;\rho_L)dG(b)$, so the interest group has more influence (a higher success rate) when better informed.

Supplement C. Endogenous NPRMs and pre-NPRM nudges

The new feature in this Section is the existence of an agency opportunity cost $\kappa > 0$ of issuing an NPRM. If κ is not sunk by the agency, the policy remains the initial policy $a = 0$. If there is an opportunity cost of issuing an NPRM, one would expect:

- Right-wing advocacy prior to an NPRM, where the right-wing interest group tries to nudge the agency to issue an NPRM.
- Left-wing advocacy post NPRM as above.

If so, there is de facto two-sided advocacy, but *observationally* advocacy is one-sided. Intuitively, the issuing of an NPRM changes the status-quo from $a = 0$ to $a = 1$; the implicit burden of proof prior to an NPRM is on the right-leaning side. Advocate L economizes by delaying her

[58]Thus, ρ indexes a mean-preserving spread.

search process. Once there is an NPRM, the new status-quo is $a = 1$, shifting the implicit burden of proof to the left-leaning side. For conciseness only, we assume a known agency bias $\bar{b} > 0$. The case for unknown b and the concomitant self-selection into an NPRM is addressed below.

Suppose (for conciseness) that in the absence of pre-NPRM information, the equilibrium with left-advocate-commenting-only prevails. An NPRM, once issued, elicits information collection from advocate L given by $c'(p_L^*) = v_L$, where $v_L = \int_{-\infty}^{-\bar{b}}[-(b_L + \omega)s_L]dF(\omega)$. The agency issues no NPRM if its opportunity cost exceeds its benefit of eliciting the information through an NPRM:

$$(1 - p_L^*)\bar{b} + p_L^* \int_{-b}^{+\infty} (\bar{b} + \omega)dF(\omega) < \kappa. \tag{25}$$

To introduce a "nudge" by advocate R, suppose for instance that the agency's payoff is $V = (\bar{b} + \omega + \xi)a$, where $\xi = \xi^+ > 0$ with probability x, and $\xi = \xi^-$ with probability $1 - x$; without loss of generality assume that $x\xi^+ + (1 - x)\xi^- = 0$. Discovering ξ may or may not be costly. As earlier, this information can be acquired only by the advocates, L and R, and is, like ω, hard.

Nudge by advocate R. We first look for an equilibrium in which only advocate R discloses signal ξ^+ prior to an NPRM[59]. Let $\bar{b}^+ \equiv \bar{b} + \xi^+$, $b_L^+ \equiv b_L + \xi^+$, $b_R^+ \equiv b_R + \xi^+$, $v_L^+ \equiv \int_{-\infty}^{-\bar{b}^+} [-(b_L^+ + \omega)s_L]dF(\omega)$, and $c'(p_L^{**}) \equiv v_L^+$[60]. If

$$(1 - p_L^{**})\bar{b}^+ + p_L^{**} \int_{-b^+}^{+\infty} (\bar{b}^+ + \omega)dF(\omega) > \kappa, \tag{26}$$

then advocate R strictly benefits from communicating that $\xi = \xi^+$ prior to the NPRM. The right-side group then gains $(1 - p_L^{**})b_R^+ s_R + p_L^{**} \int_{-b^+}^{+\infty} (b_R^+ + \omega)s_R dF(\omega) > 0$ from nudging the agency.

Nudge by advocate L. In contrast, advocate L may lose from the disclosure of $\xi = \xi^+$. First, the cost of information acquisition $C(p_L^{**})$ is borne entirely by the left advocate: A nudge puts the burden of proof on the left advocate. Second, let us compare the advocates' payoffs gross of information acquisition with and without nudge ξ^+. Without nudge the choice is always $a = 0$. Whenever the nudge leads to $a = 1$, advocate R gains $b_R - b_L > 0$ more (or loses $b_R - b_L$ less) than advocate L when the reform is undertaken. Thus advocate R benefits much more from

[59] Advocate L then has no incentive to search for information about ξ : either the signal is ξ^+ and then it duplicates advocate R's disclosure, or the signal is ξ^- and disclosing it further discourages the issuance of an NPRM.

[60] Good news about the common good component ω partly reduces the incentive to search for evidence against the reform: $p_L^{**} < p_L^*$. If p_L^{**} does not vary too much with ξ^+, as when the cost function is very convex- in the extreme has a kink at p_L^*-, the agency benefits from issuing an NPRM whenever it is informed that $\xi = \xi^+$.

the nudge on two fronts: He does not incur any search cost and it benefits more from induced policy changes. This yields:

Proposition A.2. ***(pre-NPRM Nudges).***

(i) Advocate R always gains by nudging the agency. Nudging is effective if conditions (4) and (5) are satisfied.

(ii) The left advocate's payoff from the NPRM is equal to the right advocate's payoff

$$\Delta \equiv C(p_L^{**}) + [(1 - p_L^{**}) + p_L^{**}[1 - F(-\bar{b}^{+})]][b_R s_R - b_L s_L] + p_L^{**}[\int_{-\bar{b}^{+}}^{+\infty} \omega dF(\omega)](s_R - s_L).$$

Going back to the issue of self-selection into an NPRM, consider a distribution $G(b)$. Let $\hat{G}(b \mid NPRM)$ denote the posterior distribution of b given that the agency (which has private information about its partisanship) elects to incur opportunity cost κ. An NPRM elicits information collection from advocate L given by $c'(p_L^{*}) = v_L$, where

$$v_L = \int_0^{+\infty} [\int_{-\infty}^{-b} [-(b_L + \omega)s_L]dF(\omega)]d\hat{G}(b \mid NPRM)$$

The advocate exerts more effort to collect information if it believes that the agency will be a "good listener", i.e., that it will not be too partisan. In turn, the agency issues an NPRM if and only if it is partisan enough to promote change: $b \geq b^{*}$ (and so $\hat{G}(b \mid NPRM) = \frac{G(b)-G(b^{*})}{1-G(b^{*})}$ for $b \geq b^{*}$, and $= 0$ otherwise), where

$$(1 - p_L^{*})b^{*} + p_L^{*} \int_{-b^{*}}^{+\infty} (b^{*} + \omega)dF(\omega) = \kappa$$

Suppose a wide-enough support for b (say, $G(\kappa) < 1$), so that there exists an interior cutoff b^{*}

satisfying:

$$
\begin{aligned}
\kappa &= (1-p_L^*)b^* + p_L^* \int_{-b^*}^{\infty} (b^*+\omega) dF(\omega) \\
&= b^* + p_L^* \int_{-\infty}^{\infty} \max\{-b^*, \omega\} dF(\omega) \equiv \Lambda(p_L^*, b^*) > 0 \\
\implies \frac{\partial \Lambda}{\partial p_L^*} &= \int_{-\infty}^{\infty} \max\{-b^*, \omega\} dF(\omega) > 0, \quad \text{and} \\
\frac{\partial \Lambda}{\partial b^*} &= 1 - p_L^* \int_{-\infty}^{-b^*} dF(\omega) = 1 - p_L^* F(-b^*) > 0 \\
\implies \frac{dp_L^*}{db^*} &= -\frac{1-p_L^* F(-b^*)}{\int_{-\infty}^{\infty} \max\{-b^*, \omega\} dF(\omega)} < 0 \\
\implies \frac{d^2 p_L^*}{d(b^*)^2} &= \Big(\frac{2F(-b^*)}{\int_{-\infty}^{\infty} \max\{-b^*, \omega\} dF(\omega)} + \frac{p_L^* f(-b^*)}{1-p_L^* F(-b^*)} \Big) \frac{dp_L^*}{db^*} < 0.
\end{aligned}
$$

This defines a decreasing and concave function $p_L^*(b^*)$, with $p_L^*(\kappa) = 0$ and $p_L^*(0) = \frac{\kappa}{(1-F(0))M^+(0)}$. On the other hand, we have

$$
p_L^* = (c')^{-1}(v_L) \text{ if this implies } c_0 \le p_L^* v_L - c(p_L^*), \quad = 0 \text{ otherwise.}
$$

Moreover,

$$
\begin{aligned}
v_L &= \int [-(b_L+\omega)s_L] dF(\omega)\, dG(b \big| b \ge b^*) \\
&= \frac{1}{1-G(b^*)} \int_{b^*}^{\infty} \int_{-\infty}^{-b} [-(b_L+\omega)s_L] dF(\omega) dG(b) \\
&= \int_{-\infty}^{-b^*} \frac{G(-\omega)-G(b^*)}{1-G(b^*)} [-(b_L+\omega)s_L] dF(\omega).
\end{aligned}
$$

v_L is decreasing in b^*, because what is inside the expectation is decreasing in b^*. This gives $p_L^*(b^*) < 1$, weakly positive, decreasing with a jump to zero at the value b^* for which the corresponding v_L gives

$$
\max_p \{p v_L - c(p)\} = c_0.
$$

Therefore, we have two decreasing functions that may have several intersections. We can also compute the derivative of the second function using:

$$
\frac{dp_L^*}{db^*} = \frac{s_L}{c''(p_L^*)} \frac{dv_L}{db^*} \le 0,
$$

where

$$\begin{aligned}\frac{dv_L}{db^*} &= \frac{g(b^*)}{1-G(b^*)}\int_{-\infty}^{-b^*}\frac{1-G(-\omega)}{1-G(b^*)}(b_L+\omega)dF(\omega)\\ &= \frac{g(b^*)}{1-G(b^*)}\Big(v_L+\int_{-\infty}^{-b^*}(b_L+\omega)dF(\omega)\Big).\end{aligned}$$

Supplement D. Two-comment equilibrium under soft information

The following can be an equilibrium of the soft information game if c_0 is small enough: (a) Advocate L recommends $\hat{a}_L = 0$ if $\omega < -b_L$, but does not comment when she is uninformed or informed that $\omega > -b_L$. (b) Advocate R recommends $\hat{a}_R = 0$ if $\omega < -b_R$ and this recommendation is always followed; otherwise advocate R does not comment. (c) The agency listens to R if R recommends $\hat{a}_R = 0$ or if R does not comment and L recommends $\hat{a}_L = 0$. (d) $a = 1$ is selected if and only if none of the advocates comments.

Suppose that the agency's type is known and equal to $\bar{b}$. Let $a_\emptyset = 1$ be the default action. Advocate L , when not knowing the state of nature, is not willing to recommend $\hat{a}_L = 0$ if:

$$b_L + p_R(1-F(-b_R))M^+(-b_R) \geq 0.$$

This holds true if

$$|b_L| \leq \frac{M^+(-b_R)}{1+\frac{1-p_R}{p_R F(-b_R)}}.$$

The agency listens to the L advocate's (pro-attitudinal) recommendation that $\hat{a}_L = 0$ if:

$$\bar{b} + p_R E[\omega \mid -b_R < \omega < -b_L] + (1-p_R)F(-b_L)M^-(-b_L) \leq 0.$$

These are strong conditions. The first requires that p_R be large. As for the second condition, suppose e.g. equal biases ($b_R + b_L = 0$); the second term is equal to 0. The second condition requires that p_R be small (and this is not sufficient).

Proofs of Propositions

Proof of Proposition 2.2

Suppose that (6) is satisfied. Does advocate L have an incentive to invest in information acquisition when advocate R acquires information with intensity $\hat{p}_R$ and has real authority when recommending $\hat{a}_R = 0$? When $\omega > -b_L$, advocate L would not comment upon learning the state of nature, as this would not alter the decision (advocate R then does not comment either, whether informed or uninformed, and so $a = 1$). When $\omega < -b_L$, advocate L may discover redundant information when collecting information (this happens with probability $\hat{p}_R$ when $\omega < -b_R$). So L's incentive for information acquisition (and thus $\hat{p}_L$) is smaller than the value given by (4). And so advocate L *a fortiori* has no real authority if (5) is violated. Having no real authority, advocate L collects no information. Next, suppose that (5) is satisfied. Even so, provided that advocate R collects information, there is a possible redundancy when $\omega < -b_R$ and so L's incentive to collect information is smaller than as described in equation (6), so that (5) need no longer be satisfied.■

Proof of Proposition 2.3

Note that (after integration by parts):

$$V_1(\eta; b_L, b) - V_e(\eta; b_L, b) = \int_{-\infty}^{-b} (b+\omega) d\hat{F}(\omega \mid \eta, b_L) + \gamma = -\int_{-\infty}^{-b} \hat{F}(\omega \mid \eta, b_L) d\omega + \gamma$$

is decreasing in b_L and increasing in η and b, and so, as the advocate becomes more moderate, the review region gains at the expense of the region in which the comment is ignored. Note further (again integrating by parts and using $\hat{F}(-b_L \mid \eta, b_L) = 1$) that

$$V_e(\eta; b_L, b) - V_0(\eta; b_L, b) = \int_{-b}^{-b_L} (b+\omega) d\hat{F}(\omega \mid \eta, b_L) - \gamma = (b - b_L) - \int_{-b}^{-b_L} \hat{F}(\omega \mid \eta, b_L) d\omega - \gamma.$$

Together with MLRP (which, recall, implies that $\partial \hat{F}/\partial \eta < 0$), this implies that $V_e(\eta; b_L, b) - V_0(\eta; b_L, b)$ and $V_1(\eta; b_L, b) - V_e(\eta; b_L, b)$ are increasing in η and b[61]. Furthermore,

$$\frac{\partial (V_e(\eta; b_L, b) - V_0(\eta; b_L, b))}{\partial b_L} = -\int_{-b}^{-b_L} \frac{\partial \hat{F}(\omega \mid \eta, b_L)}{\partial b_L} < 0$$

We made assumptions guaranteeing that the rubber-stamping and nonreceptivity regions always exist (for η very low or very high, respectively). In contrast, the review region exists if and only if $\gamma < \bar{\gamma}$ for some $\bar{\gamma} > 0$. Namely, V_1 is increasing in η while V_0 is decreasing; they intersect at $\eta = \eta^\sharp$ such that $b + E[\omega \mid \omega < -b_L, \eta^\sharp] = 0$, and so $\bar{\gamma} \equiv V_1(\eta^\sharp; b_L, b) = V_0(\eta^\sharp; b_L, b)$. This yields the value of $\bar{\gamma}$ given in the Proposition part *(i)*.

[61] $\frac{\partial}{\partial b} V_e(\eta; b_L, b) = 1 - \hat{F}(-b \mid \eta, b_L) \geq 0.$

(iv) Success rate. The success rate is equal to $\sigma_0 = 1$ when $\eta < \eta_1(b_L, b)$. It is equal to $\sigma_e(b_L, b \mid \eta) = \frac{F(-b|\eta)}{F(-b_L|\eta)}$ if $\eta_1(b_L, b) < \eta < \eta_2(b_L, b)$, and to $\sigma_1 = 0$ for $\eta > \eta_2(b_L, b)$. Note that σ_e increases with b_L, decreases with b, and, under the monotone-likelihood-ratio property, decreases with η. ■

B Supplement to the empirics: Additional stylized facts, tables, and figures

Fact 7. Agencies are responsive to public comments. *On rules with less than a thousand comments, 75% of comment letters receive at least one detailed response. On rules with tens of thousands of comments, the agency response rate falls to 50% for letters from organizations and to 20% for letters from individuals.*

Our theory is predicated on the strategic transfer of information from advocates to agencies in equilibrium. It is uncommon to find direct empirical validation that informational transfer between the agent and the principal has indeed occurred. For example, in the case of federal lobbying, the Lobbying Disclosure Act requires no disclosure of the content of messages or of the response by targeted politicians. In the case of rulemaking, we have direct metrics.

About half of all proposed rules receive a positive number of comments from the public. The total number is, however, heterogeneous [62] and the distribution of comment letter counts is skewed: A few salient rules receive hundreds of thousands of comments, with the top 20 rules accounting for 50% of all comments linked to a rule in our sample. In turn, the number of agency responses to these comments is much less skewed, with 50% of all responses in our data occurring in the case of rules that receive fewer than 100 comments. Perhaps unsurprisingly given these patterns, the probability that a comment receives at least one detailed response from an agency depends on the total number of comment letters received by the regulator in the first place. For rule proposals with only a handful of comments, roughly 75% of comments receive at least one response. Organizations are more likely to receive a response than individuals, and this gap widens on proposed rules that receive many comments. Overall, one can interpret the fact that agencies acknowledge comments and engage in a response to advocates, as corroborating evidence of informational exchange within the NPRM process.

[62] While roughly half of the rules proposed receives no comment, those are rules of trivial impact or published as matter of pure procedure –for example, the U.S. Coast Guard must publish a rule every time they establish a temporary safety zone that regulates activities on the water for a special event.

Figure 12: Comments and Responses

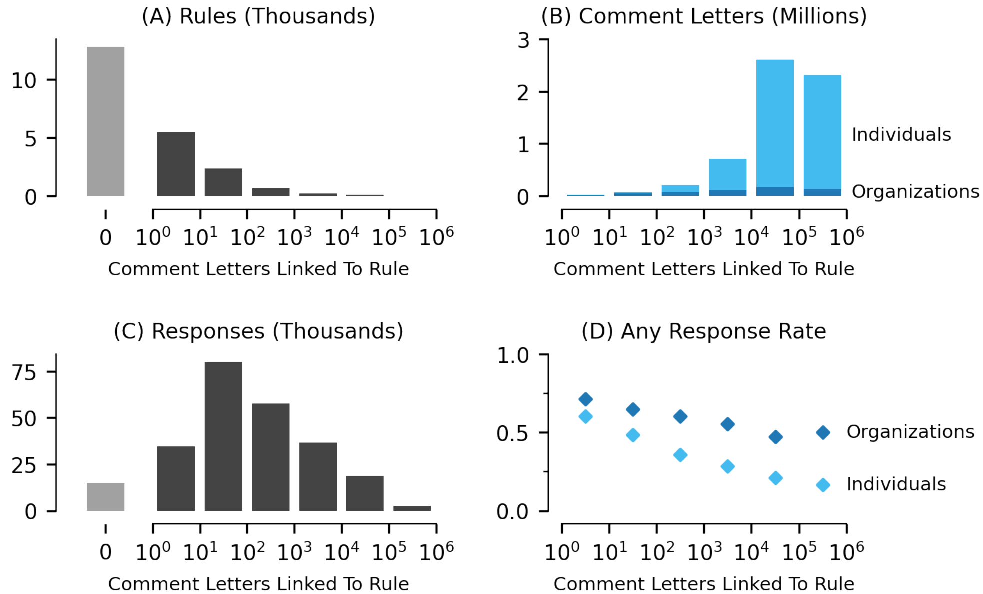

Figure 12, panel A shows the number of rules in our data broken down by the number of comment letters linked to the rule. Panel B shows the number of comments linked to rules in each comment letter count bin. Comments are categorized as coming from organizations if we detect at least one organization author, and the remaining comments are assumed to be authored by individuals. Panel (C) shows the total number of responses in our data broken down according to the number of comments linked to the rule. 6% of responses occur in rules with no linked comments, which reflects failures to find and link the appropriate comments to the rule. Panel D shows the fraction of comments that are linked to at least one response according to the number of comments linked to the rule.

Fact 8. Most comments are rejected. *Agencies agree to make policy changes in 22% of responses to comments, with vast heterogeneity across agencies and rules.*

Agencies frequently make policy changes in response to comments. Most of these changes are small, but some are substantial. Common examples of changes include adding, modifying or removing provisions from the proposed rule, changing key definitions, or delaying implementation dates.[63] In our data, agencies describe making a policy change in 22% of all responses. Given the number comments received and responses written by agencies this strikes as a fairly large rate of change.

The fraction of responses including a policy change varies substantially across rules and agencies. Figure 13 shows number of comment letters, responses, and fraction of responses

[63]Estimating the economic significance of each policy change is extremely difficult, and we do not attempt to do so in this paper. Instead we focus primarily on the count of responses to comments that are included in each rule and the fraction of these responses where the agency describes making a policy change to address commenter's concerns. See (Bombardini et al., 2025) for a discussion.

with a policy change for all rules published from 2008-2022 that are linked to at least one comment and contain at least one response. The skewed distribution of comments and responses discussed earlier is clearly visible. Roughly half of rules that receive thousands of comments are classified as "economically significant" under Executive Order 12866[64]. We find that rules that receive many comment letters tend to have a lower fraction of responses with a policy change. This pattern is robust to controlling for economic significance, the number of proposed rules published before the final rule, and agency-by-year fixed effects–however, it is primarily driven by comments from individuals, not comments from organizations (see Table 6 for full regression results).

Figure 13: Rule Comments and Responses

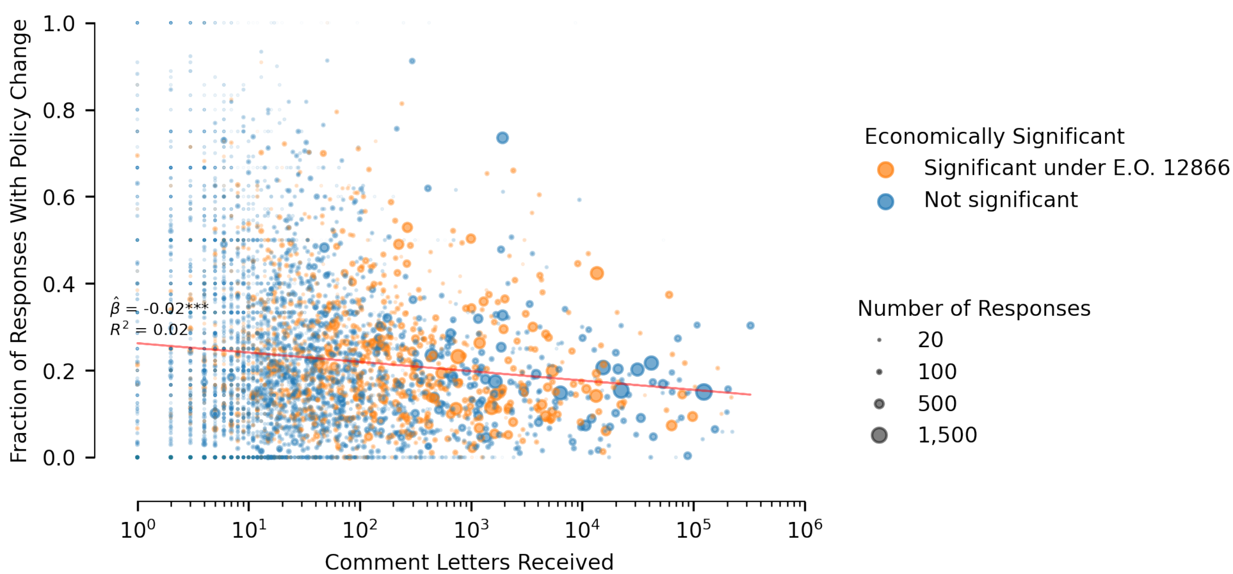

Figure 15 shows all rules published in the Federal Register from 2008-2022 that are linked to at least one comment in our data and contain at least one response. Each point represents one rule. The x-axis indicates the number of comment letters (unique Regulations.gov comment IDs) linked to each rule, and the size and opacity of the point is scaled to indicate the number of responses extracted from the rule using our automated approach. The y-axis indicates the fraction of responses in the rule for which our automatic classifier detects a policy change. Rules are classified as "economically significant" under Executive Order 12866 if they "an annual effect on the economy of $100 million or more" or result in other adverse material effects, or raise important legal issues. The line of best fit is calculated using a WLS regression where each rule observation is weighted by its response count. *** indicates statistical significance at the $\alpha = 0.01$ level using heteroskedasticity-robust (HC3) standard errors.

Figure 14 shows that similar variation exists across agencies. Some key agencies, such as the Federal Aviation Administration (FAA) and Environmental Protection Agency (EPA), are extremely active in rulemaking, each publishing more than 10,000 rules and writing tens of thou-

[64]This classification initiates an additional process of review by the Office of Information and Regulatory Affairs and can be triggered by several broadly defined economic or legal impacts, including "an annual effect on the economy of $100 million or more".

sands of responses to comments. The U.S. federal government also has many other smaller agencies like the Administration on Aging (AOA) or National Archives and Records Administration (NARA) that only occasionally engage in rulemaking or publish less salient regulations and write fewer responses to comments. Agencies also vary widely in the fraction of responses that contain a change. Notable extremes include the Federal Crop Insurance Program (FCIC), which makes a change in 49% of its responses, and the Federal Energy Regulatory Commission (FERC), which makes a change in only 10% of its responses.

Figure 14: Agency Rules and Responses

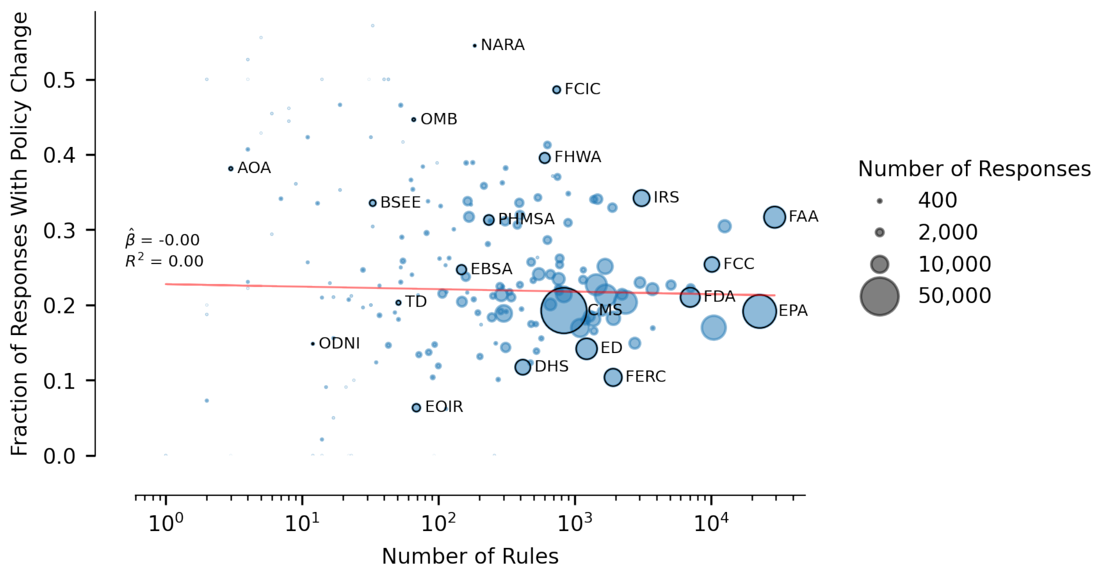

Figure 14 shows all agencies that appear as "primary authors" of rules published from 2008-2022. The x-axis indicates the number of rules authored by each agency. Many rules list multiple agencies as authors, but these often reflect nested organizational hierarchies. We assign primary authorship by removing parent agencies and departments from the author list when a child agency is also listed. If rules still have multiple authors after removing parent agencies these are assigned to multiple points in the plot. The size of the point indicates the total number of responses in all the rules assigned to each agency, while the y-axis indicates the fraction of those responses where we detect a policy change. The line of best fit is calculated using a WLS regression where each rule observation is weighted by its response count.

Fact 9. There is little direct opposition among advocates within the same rule. *Fewer than 10% of agency responses address comments on opposing sides of an issue in a proposed regulation.*

Within our theory, multi-sided commenting is possible, but unlikely. Proposition 2.1 part (i) discusses the one-sided equilibrium under hard information. Under the restrictive conditions in Proposition 2.2, opposite-sided commenting should be rare. However, many alternative models of advocacy and strategic communication (e.g. Dewatripont and Tirole, 1999; Krishna and Morgan, 2001b; Battaglini, 2002) and influence (Becker, 1983), as well as in empirical contest function models, such as Kang (2016), make a case for competition among multiple special interest groups as potentially welfare increasing. Theoretically, political principals benefit from the tug-of-war among informed agents and their jousting for influence over key policy decisions.

Empirically, we can show that this tug-of-war is infrequent. An analysis of the responses to comments written by agencies suggests that competition for influence among multiple advocates is rare. The majority of responses addresses a single advocate or a single group of commenters, who all make a similar comment on a specific issue or request the same specific change in the rule. To establish this pattern, we trained a classifier to distinguish between four degrees of opposition between the comments addressed in each response. In most cases it easy to tell whether a response is addressing a single commenter or multiple commenters from the way the agency uses either singular or plural nouns when referring to the commenters (our automatic classifier achieves an F1 score of 0.96 on this task). For responses addressing multiple comments, it is also possible to infer the number of "sides" being addressed. We classify responses into three categories: *Single Side*, indicating that all comments take the same position, *Multiple Sides*, indicating that there are multiple views represented without necessarily being strongly opposed to other, and *Opposing Sides*, indicating that at least two commenters want opposing outcomes[65]. Figure 15 shows the distribution of the number of sides addressed by agency responses using our automatic classifiers. We find that 74% of all responses address either single advocate or a single side. Only 5% of responses address opposing sides[66].

[65] Distinguishing between these cases is harder than simply detecting whether a response is addressing a single or multiple commenters, but our classifier still achieves a Macro F1 score of 0.75.

[66] A weakness of our automatic classification approach is that it considers each response in isolation. It is possible that agencies could split opposing sides into separate responses and address them separately. To address this concern we had an RA manually review all the responses in a sample of 29 rules with responses to commenters, carefully checking for opposing sides addressed in separate responses. In this small test set, 99/108 (92%) of responses address issues raised by commenters with with no direct opposition. Similarly, 23/25 (92%) of policy changes occur in responses to these unopposed comments.

Figure 15: Response Counts by Commenter Sides

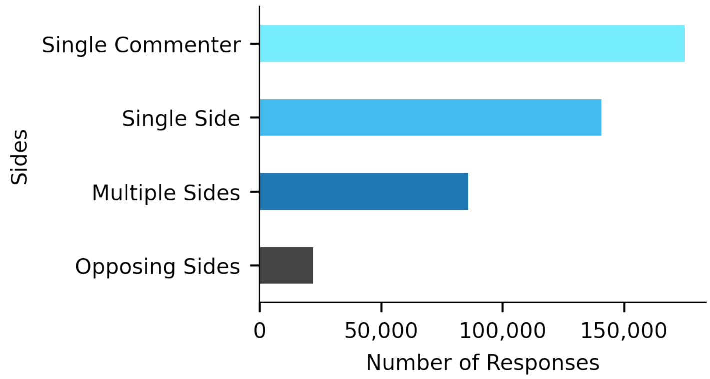

Figure 12 shows the number of responses in our sample, broken down according how many "sides" were described in the text of the regulators response and the inferred commenter stance. See section 3.3 for a discussion of the classification methodology. 74% of responses address either a single commenter or a group of commenters taking the same side. Only 5% of responses address opposing sides.

Overall, it appears that opposition between interest groups in rulemaking is less common, or at least more indirect, than what models of multilateral advocacy and communication may postulate. Of course, it is likely that many commenters have opposing goals in some very broad sense. But, when it comes to writing comments, they do not seem to argue over the same detailed issues of a rule. Instead, advocates with different objectives seem to "talk past each other", raising many separate, mostly independent points.

Fact 10. Influence is highly concentrated. *Up to 50% of all policy changes can be attributed to the top 1000 commenting organizations (30% to the top 100 organizations).*

Lobbying in Washington is notoriously concentrated within a small number of large businesses and industry groups (Bombardini and Trebbi, 2020). Most businesses do not lobby at all, but those that do lobby a lot (Huneeus and Kim, 2018; Kerr et al., 2014). A key question about the notice-and-comment process is whether a wider range of organizations participate, and whether participation leads to equitable outcomes. One way to quantify this is to estimate the fraction of all policy changes that can be attributed to a small number of organizations, where each "change" is identified as a response containing a policy change.

Figure 16 shows the result of such an exercise, using two different methods to attribute changes to organizations. The first, "greedy" attribution, starts by finding the organization with the most linked changes and attributing every change linked to this organization exclusively to this top organization. Then the second-most influential organization is identified by finding the organization with the most linked changes among the unattributed changes, and these linked changes are attributing to the second organization. The process is iterated until no more changes can be attributed to any organization. Thus, the greedy attribution measure gives the maximum number of changes that could potentially be attributed to the n most influential organizations, which we interpret as an upper bound on the concentration of influence. The second attribution measure is the "equal" attribution: when multiple organizations are linked to the same change, each organization is attributed an equal share of the change. Given the large disparity between different commenting organizations, this measure is likely a lower-bound on the concentration of influence.

Figure 16: The Cumulative Influence of Top-n Organizations (Policy Changes)

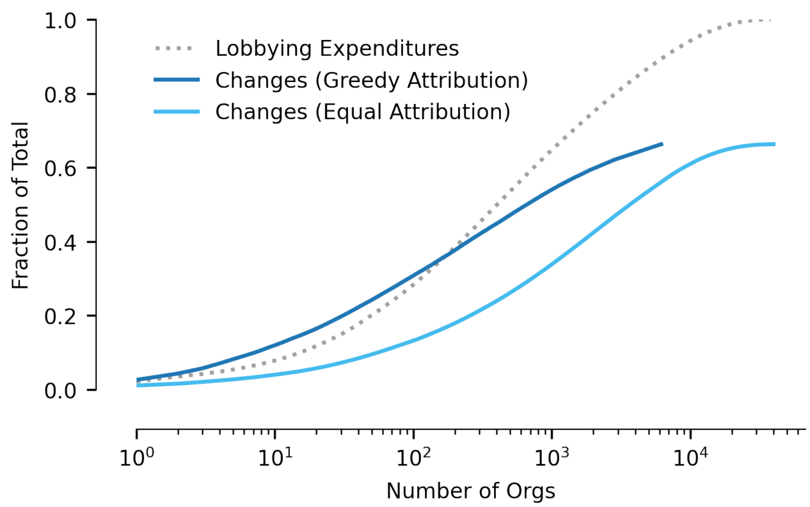

Figure 16 shows the fraction of changes that can be attributed to the n-most influential organizations in our data using two different algorithms for assigning credit for changes. Equal attribution divides the credit evenly among all organizations linked to the response. A point on the curve then indicate what fraction of the total credit for all responses can be attributed to the top n organizations ranked by total credit assigned. The greedy attribution algorithm finds the single organization that is linked to the largest number of changes, assigns all credit for those changes to that organization, and removes that organization and all its assigned changes from the pool of changes. This assignment process is repeated until no un-assigned changes remain. A point on this curve indicates the fraction of all responses that are linked to at least one of the top n organizations ranked in this manner. The lobbying expenditures curve is presented for reference. It shows the fraction of all lobbying expenditures recorded under the Lobbying Disclosure Act from 2008-2022 that can be explained by the top n organizations with the highest lobbying spending.

Our data suggests that commenting influence is highly concentrated. The 100 most influential organizations are responsible for between 13% and 31% of all changes in our linked sample, while the top 1000 organizations are responsible for 34% to 54% of all changes. For reference, over this period, the top 100 lobbying clients were responsible for 28% of all registered lobbying spending, while the top 1000 lobbying clients were responsible for 65% of all lobbying spending. Therefore, while the total number of organizations commenting is much higher, the share of changes that can be attributed to the most active participants is similar to the fraction of lobbying spending that can be attributed to top spenders.

Fact 11. Top commenting organizations are diverse. *The top 100 most influential commenters include a mix of business associations, individual businesses, non-profits, professional associations, and government agencies.*

A notable difference between commenting and congressional lobbying is that the top commenting organizations are more diverse. The top 100 lobbying clients are consist almost exclusively of businesses (68%) and business associations (22%). In contrast, for the top 100 commenting organizations (ranked by a simple count of linked changes), business associations play a more important role (35%) while businesses only account for 18% of the top 100 commenting organizations. There are also a significant number of nonprofits (22%), professional associations (13%), and various government agencies (7%).

Table 4: Top commenting organizations by linked change count

Organization Name	Rules	Linked Responses	Linked Changes	Change Fraction
American Petroleum Institute	351	5,205	1,380	0.27
Natural Resources Defense Council	519	7,402	1,207	0.16
Earthjustice	453	7,342	1,127	0.15
Center for Biological Diversity	569	7,353	1,107	0.15
Sierra Club	601	7,801	1,055	0.14
American Medical Association	95	3,970	745	0.19
Environmental Defense Fund	245	3,854	726	0.19
Boeing Co	847	1,359	716	0.53
US Chamber of Commerce	292	3,054	636	0.21
American Chemistry Council	219	2,646	537	0.20
Federation of American Hospitals	58	2,652	470	0.18
American Hospital Assn	85	2,662	454	0.17
American Bankers Assn	119	1,395	434	0.31
National Association of Manufacturers	183	2,035	422	0.21
National Mining Association	126	1,760	420	0.24
American Fuel & Petrochemical Manufacturers	126	1,984	407	0.21
Defenders of Wildlife	180	2,358	390	0.17
American Bar Association	95	1,445	382	0.26
American Academy of Family Physicians	71	1,866	374	0.20
AFL-CIO	179	1,830	362	0.20
American Forest & Paper Association	135	1,590	361	0.23
Association of American Medical Colleges	58	2,105	352	0.17
American College of Physicians	51	1,597	351	0.22
Independent Petroleum Association of America	87	1,229	344	0.28
Hunton & Williams LLP	121	1,870	341	0.18
Securities Industry and Financial Markets Association	63	1,209	336	0.28
Alliance of Automobile Manufacturers	89	1,787	335	0.19
Delta Air Lines	333	766	322	0.42
American College of Surgeons	27	1,643	319	0.19
Western Energy Alliance	76	1,049	316	0.30

Table 4 shows the top 30 commenting organizations ranked by linked change count. By this measure, the most influential commenting organization is the American Petroleum Institute, with 1,380 linked changes. Five of the top ten organizations are environmental nonprofits who comment together on many rules (and should be seen as a coalition that is linked to roughly a thousand changes rather than five separate organizations with a thousand changes each). Boeing stands out as the most influential individual business, with 716 linked changes. This is partly due to the way airlines are regulated, as, for example, the Federal Airline Administration produces thousands of rules regulating aircraft safety, and many of these rules only affect specific airplane models. The majority of the remaining top organizations are professional organizations and trade groups representing specific industries. Only one union, the AFL-CIO, appears on the list with 362 linked changes. [67]

Table 5: Rule Characteristics and Response Counts

	Responses (Poisson)				
	(1)	(2)	(3)	(4)	(5)
Significant under EO 12866	1.377***				0.325***
	(0.101)				(0.074)
Multiple Proposals		1.207***			0.355***
		(0.064)			(0.060)
Log(1 + Comments)			0.462***		-0.055**
			(0.011)		(0.022)
Log(1 + Org Commenters)				0.731***	0.747***
				(0.016)	(0.033)
Agency×Year FE	x	x	x	x	x
Observations	21451	21451	21451	21451	21451

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Rule Characteristics and Change Fraction

	Change Fraction (OLS)				
	(1)	(2)	(3)	(4)	(5)
Significant under EO 12866	0.019**				0.021**
	(0.008)				(0.008)
Multiple Proposals		0.031***			0.038***
		(0.007)			(0.007)
Log(1 + Comments)			-0.010***		-0.041***
			(0.002)		(0.003)
Log(1 + Org Commenters)				-0.003*	0.038***
				(0.002)	(0.004)
Agency×Year FE	x	x	x	x	x
Observations	10268	10268	10268	10268	10268
R^2	0.203	0.204	0.205	0.203	0.213

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

[67] In our linking process we include local chapters as part of the main organization

Table 7: Response by Commenter Sides

	Any Change (OLS)	
	(1)	(2)
Single Side	0.054***	0.062***
	(0.003)	(0.002)
Multiple Sides	-0.004	0.012***
	(0.003)	(0.003)
Opposing Sides	-0.011**	-0.011***
	(0.005)	(0.004)
Intercept	0.201***	
	(0.003)	
Rule FE	-	x
Observations	424156	424156

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Outcomes are measured at the individual response level. The outcome variable indicates whether there was any policy change detected in the response text. "Sides" indicator coefficients are measured relative to the "Single Commenter" category. Standard errors are clustered by rule.

Figure 17: Distribution of Top 100 Org Types for Lobbying and Commenting

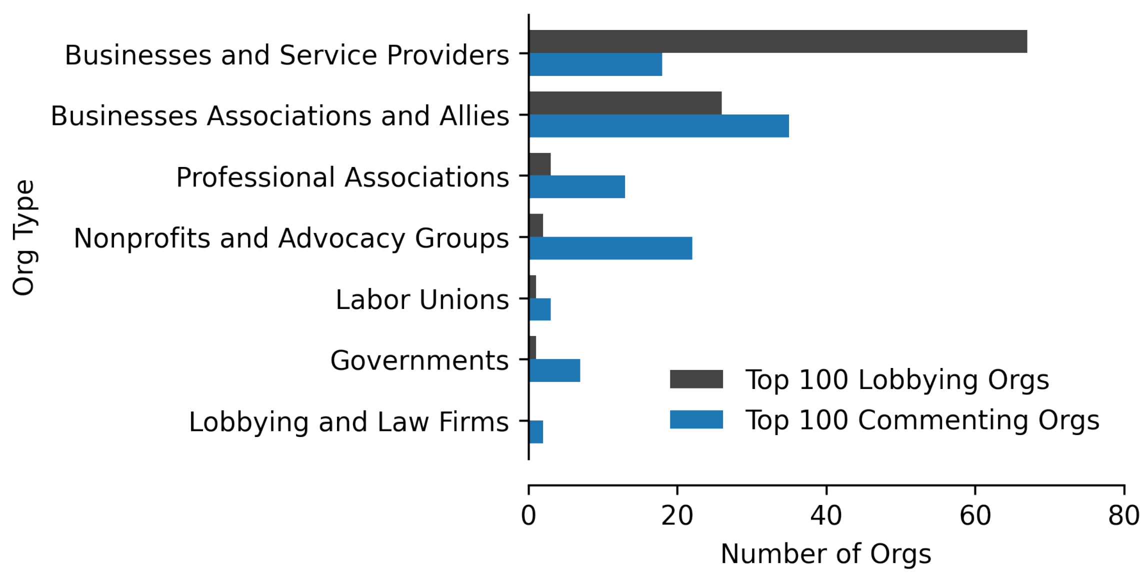

C Firm size regression tables

Table 8: Influence Decomposition by Compustat Firm Employment

	Success Rate (1)	Commenting Rate (2)	Response Rate (3)	Success Rate (4)
Log(Employees)	0.363*** (0.032)	0.368*** (0.043)	0.061*** (0.015)	0.003 (0.006)
Year FE	x	x	-	-
Rule FE	-	-	x	x
Observations	22960	22960	11239	7422
R^2	-	-	-	-

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are clustered by Organization in columns 1-2 and Organization + Rule in columns 3-4. Estimates in columns 3-4 are conditional on commenting on the rule and receiving at least one response respectively.

Table 9: Influence Decomposition by Compustat Firm Market Value

	Success Rate (1)	Commenting Rate (2)	Response Rate (3)	Success Rate (4)
Log(Market Value)	0.419*** (0.036)	0.385*** (0.045)	0.055*** (0.016)	0.006 (0.006)
Year FE	x	x	-	-
Rule FE	-	-	x	x
Observations	20531	20531	9216	6047
R^2	-	-	-	-

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are clustered by Organization in columns 1-2 and Organization + Rule in columns 3-4. Estimates in columns 3-4 are conditional on commenting on the rule and receiving at least one response respectively.

Table 10: Influence Decomposition by Total Lobbying Expenditures

	Success Rate (1)	Commenting Rate (2)	Response Rate (3)	Success Rate (4)
Log(Lobbying Amount)	0.448***	0.374***	0.094***	0.008***
	(0.016)	(0.015)	(0.011)	(0.002)
Year FE	x	x	-	-
Rule FE	-	-	x	x
Observations	183765	183765	84398	57369
R^2	-	-	-	-

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are clustered by Organization in columns 1-2 and Organization + Rule in columns 3-4. Estimates in columns 3-4 are conditional on commenting on the rule and receiving at least one response respectively.

D Ideology regression tables

Table 11: Effects of Ideological Congruence

	Total Changes (1)	Commenting Rate (2)	Response Rate (3)	Success Rate (4)
CF $\times$ Republican President	-0.061*	-0.150***	-0.158***	0.177***
	(0.035)	(0.015)	(0.027)	(0.021)
Year FE	x	x	x	x
Organization FE	x	x	x	x
Observations	196755	474450	99674	65931
R^2	-	-	-	-

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are clustered by Organization + Rule.

Table 12: Influence Decomposition by Ideology

	Total Change Count	Commenting Rate	Response Rate	Success Rate
	(1)	(2)	(3)	(4)
CF	-0.343***	-0.131***	-0.159***	0.036***
	(0.059)	(0.035)	(0.033)	(0.006)
CF^2	-0.315***	-0.290***	-0.040	-0.039***
	(0.055)	(0.036)	(0.026)	(0.007)
Rule FE	-	-	x	x
Year FE	x	x	-	-
Observations	474450	474450	110570	68003

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are clustered by Organization in columns (1-2) and Organization + Rule in columns (3-4).

E Annotation instructions

Paragraph Labeling Scheme

Each paragraph should be labeled with as “S”, “R”, “SR”, or “O” in the “label” column:

- “**S**” indicates that the paragraph summarizes one or more comments
- “**R**” indicates that the paragraph responds to one or more comments
- “**SR**” indicates that the paragraph contains both a summary and a response
- “**O**” (the letter, not the number) indicates that the paragraph is neither a summary nor response

Any Change Labeling Scheme

Each response paragraph ("R" or "SR") should be coded as "Y", "N" or "YN" in the "any_change" column:

- "**Y**" indicates that the regulator describes making a change to the policy environment in response to the comments. For example, the following cases should all be coded "Y":
    1. Any change made by the regulator, even if it is not asked for by the commenter.
    2. A change in the regulation's effective date or date for compliance.
    3. Withdrawing a rule or a part of a rule.
    4. Amending a regulation, even if the change is seemingly trivial.
    5. A commitment to a change in enforcement or interpretation of the law, even if the letter of the law is not changed.
- "**N**" indicates that there was no change made to the policy environment described in this response paragraph. For example, the following cases should all be coded "N":
    1. The regulator provides additional information, explanation, or clarification, but does not make a change.
    2. The regulator's response is something to the effect of "we will consider the commentator's request in the future."
    3. The regulator expresses agreement with the comment, but does not say they are making a change.
    4. The regulator's response is something to the effect of "we will endeavor/try to use the information provided."
- "**YN**" indicates that the regulator is responding to multiple commenters with different views and making a change in response to some but not others
- "**?**" Can be used to indicate uncertainty, either on its own, or combined with the best guess (e.g. "N?")

Comment-Response Labeling Scheme

For each pair of comment-response texts, enter a number between 1 and 5 in the "match_quality" column. The number should indicate degree of overlap between the topics discussed in the two texts and how likely it is that the agency's response text is intended as a response to the selected comment text:

1. **Incorrect match**. Comment and response text are clearly discussing very different issues. The agency is definitely not responding to this comment text in the response text.
2. **Poor match**. Comment and response text are somewhat related, but appear to be discussing different specific issues. It is unlikely that the agency is responding to this comment text in the response text.
3. **Partial Match**. Comment and response text are discussing related issues but the degree of overlap is either imperfect or somewhat ambiguous.
4. **Good match**. Comment text appears closely related to the agency's response. It is likely that the agency is responding to this comment text.
5. **Perfect match**. Comment text contains the exact argument or information that the agency is responding to in the response text. The agency is definitely responding to this specific comment text.

Notes:

- The response text could also be addressing other comments as well. This should not detract from the score. For example, if the regulator is clearly responding to two different comments A and B, and the selected comment text appears to exactly match the summary of comment A, then enter a "5".
- Sometimes there is a tension between recognizing that the comment is likely the one being discussed, and whether there is a good topic match. For example, both the comment and response might identify the commenter by name making it clear that this is the correct comment. However, if the topics do not match, the score should still be low (keep in mind this is only a sample of the comment text - it is likely that there is another omitted sample of the comment text that would be a better match).

F Prompts for LLM annotations

The following text comes from a comment letter submitted to a government regulator as part of the note-and-comment process for U.S. federal rulemaking.

INSTRUCTIONS:

First, please identify whether the comment is authored by an organization or an individual.

- Respond with "organization" if the letter is written on behalf of one or more organizations
- Respond with "individual" if the letter is written on behalf of an individual
- Only print one word.

Next, list all organizations affiliated with the author of the comment on separate lines.

- Use the most complete and unambigous version of each name.
- Do not include departments or divisions within an organization.
- IMPORTANT: Do not include CC'd organizations.

COMMENT TEXT:

< “organization” metadata field >

< “title” metadata field >

< first paragraph of comment body >

< closing text (e.g. sincerely....) >

< comment description accompanying attachments (if any) >

The following text comes from a comment letter submitted to a government regulator as part of the notice-and-comment process for U.S. federal rulemaking.
INSTRUCTIONS:
First, describe who authored the comment in 10 words or less.
Then, classify the comment as one of the following types:
- "organization" if the comment is written on behalf of one or more organizations such as a business, nonprofit, or government agency
- "member" if the comment is written by someone who represents themselves a rank and file member of an organization (for example, a union member, student at a school, nurse at a hospital, etc.)
- "politician" if the comment is written by a politician or political candidate (for example, a member of congress, state legislator, mayor, or candidate running in an election)
- "expert" if the comment is written by an independent expert in the field being regulated (for example, an academic, scientist, engineer, or doctor)
- "individual" if the comment is written by someone who is does not have a strong affiliation with an organization (sometimes described as a citizen or taxpayer)
- "unknown" if the comment type cannot be determined from the text
No further explanation is required.
COMMENT TEXT:
< comment text >

The following text comes from a comment letter submitted to a government regulator as part of the note-and-comment process for U.S. federal rulemaking.
INSTRUCTIONS:
First, please identify whether the comment is authored by an organization or an individual.
- Respond with "organization" if the letter is written on behalf of one or more organizations
- Respond with "individual" if the letter is written on behalf of an individual
- Only print one word.
Next, list all organizations affiliated with the author of the comment on separate lines.
- Use the most complete and unambigous version of each name.
- Do not include departments or divisions within an organization.
- IMPORTANT: Do not include CC'd organizations.
COMMENT TEXT:
< “organization” metadata field >
< “title” metadata field >
< first paragraph of comment body >
< closing text (e.g. sincerely....) >
< comment description accompanying attachments (if any) >

The goal of this task is to review selected names to identify organizations that authored comment letters submitted to a U.S. federal agency.
The name below should be classified into one of the following categories:
- "author": The comment can be inferred to represent the official views of this organization. This includes comments:
- written by this organization
- written on behalf of this organization
- written by a leader or senior member of this organization
- written by a subsidiary or part of this organization
- "affiliated": One of the authors is affiliated with this organization as a rank and file member or lower-level employee.
- "mentioned": This organization is only mentioned in the letter. This includes cases where the author endorses a separate comment written by the organization.
- "recipient": This organization is a recipient of the comment letter (including CC'd organizations).
- "invalid": The name is not a valid organization. This includes individual names, addresses, or uninterpretable strings.
COMMENT TEXT:
<comment text>
ORGANIZATION NAME: <org_name>
INSTRUCTIONS:
First, use the comment text and organization name above to learn about the organization.
Next, Describe the role of the organization in the comment in 10 words or less.
Finally, classify the organization into one of the categories listed above. Print only the word associated with the category.
No further explanation is required.

The goal of this task is to classify an organization that has submitted a comment to a U.S. federal agency.

The organization should be classified into one of the following categories:

CATEGORIES:

B: "Businesses and Service Providers", including:

- Corporations
- Small Businesses
- Banks
- Utility Companies
- Universities
- Hospitals
- Transit Systems

A: "Businesses Associations and Allies", including:

- Trade Associations
- Industry funded Think Tanks and Institutes
- Business advocacy groups
- Chambers of Commerce
- Industry funded astroturf campaigns

U: "Labor Unions", including:

- Labor Unions
- Employee Associations

P: "Professional Associations", including:

- Medical Associations
- Bar Associations
- Engineering Associations
- Other Professional Associations

N: "Nonprofit Foundations and Advocacy Groups", including:

- Independent Philanthropic and Charitable Foundations
- Non-business Think Tanks and Research Organizations
- Religious Organizations
- Scientific and Academic Organizations
- Museums and Cultural Institutions
- Ideological Advocacy Groups
- Consumer Advocacy Groups
- Civil Rights Organizations
- Social Justice Organizations
- Advocacy Groups for marginalized communities
- Environmental nonprofits
- Conservation Organizations
- Animal Rights Organizations