

OPTIMAL AGENCY BIAS AND REGULATORY REVIEW

RYAN BUBB* AND PATRICK WARREN**

ABSTRACT. Why do bureaucratic principals appoint agents who hold different policy views from the principal and subject their decisions to review by a more aligned agent? We posit an explanation based on the interplay between two types of agency costs: shirking on information production and policy bias. Principals employ biased agents because they shirk less, and principals institute review of the agents' decisions to mitigate the resulting bias in those decisions. We apply the theory to explain various features of the administrative state. In contrast to existing accounts, in our model the use by the president of ideological bureaucrats at the regulatory agencies and centralized regulatory review are complements. The use of bias to mitigate shirking results in an amplification of the swings of regulatory policy and heightens the role of regulatory policy in partisan politics.

Preliminary and incomplete.

Date: November 21, 2012.

We are grateful to Barry Adler, Jennifer Arlen, Stephen Choi, Oren Bar-Gill, Rachel Barkow, Charles Cameron, Adam Cox, Kevin Davis, John Ferejohn, Jody Freeman, Douglas Ginsburg, Jeffrey Gordon, Sanford Gordon, Oliver Hart, Sam Issacharoff, Marcel Kahan, Louis Kaplow, Lewis Kornhauser, Michael Levine, Daryl Levinson, Michael Livermore, Rick Pildes, Ricky Revesz, Catherine Sharkey, Steve Shavell, Matthew Stephenson, Charles J. Thomas, and workshop participants at NYU School of Law, Clemson University, and Harvard Law School for helpful discussions and comments.

*New York University School of Law. Email: ryan.bubb@nyu.edu.

**Clemson University. Email: patrick.lee.warren@gmail.com.

1. INTRODUCTION

In January 2010 the Environmental Protection Agency (EPA) proposed a rule that would have substantially tightened the standard for ozone under the Clean Air Act.¹ But in September 2011 the EPA's ozone proposal was quashed following a review of the policy by the White House Office of Information and Regulatory Affairs (OIRA). In his letter to EPA Administrator Lisa Jackson announcing the decision, Cass Sunstein, the Administrator of OIRA, explained that President Obama had "directed [him] to give careful scrutiny to all regulations that impose significant costs on the private sector" and "does not support finalizing the rule at this time."² The President, of course, had appointed both Jackson and Sunstein to their respective posts, choosing a self-described environmentalist to lead EPA and a proponent of cost benefit analysis to head OIRA. Moreover, President Obama issued executive orders requiring that significant rules issued by EPA (and other executive agencies) be subject to OIRA review, continuing a practice that dates back to the Nixon administration.³

In this paper we provide a rationale for a bureaucratic principal appointing agents who hold different policy views from the principal and instituting a review process led by a bureaucrat with views more in line with the principal's. Our explanation is based on the interplay between two types of agency costs that stem from delegation: *shirking* and *bias*.

First, information is a key input into policymaking, and generating information is costly. When responsibility for information production is delegated to an agent, the agent might shirk rather than exert the optimal amount of costly effort to generate information (Stephenson, 2011).

Second, agents may have policy preferences that differ from those of the principal. These could be intrinsic policy preferences or alternatively preferences that are induced by some implicit incentive scheme. Biased policy preferences may skew both agents' willingness to divulge information and the policy choices they make.

¹75 Fed. Reg. 2,938 (Jan. 10, 2010).

²Letter from Cass Sunstein to Lisa Jackson, Sept. 2, 2011, available at http://www.whitehouse.gov/sites/default/files/ozone_national_ambient_air_quality_standards_letter.pdf.

³Exec. Order 13,563, 76 Fed. Reg. 3,821 (Jan. 18, 2011); Exec. Order 13,610, 77 Fed. Reg. 28,469 (May 14, 2012).

We show how policy bias can be harnessed to mitigate the problem of shirking. Suppose that the principal has to choose an agent responsible for regulating in some domain, say the environment. Should the principal appoint an agent whose intrinsic policy preferences correspond to the principal's? While such an agent would choose the principal's preferred rules, conditional on their information, they have suboptimal incentives to exert effort to generate information.

Accordingly, we show that the principal should choose an agent who is more pro-regulatory than the principal. For example, the principal might appoint someone who values environmental quality to a greater extent than she does. A person who places greater value on clean air, say, is willing to work harder to find regulatory opportunities to improve air quality. Hence, appointing such a person can help on the extensive margin of regulation—more harmful pollutants, say, are identified and brought under control. This incentive effect comes at a cost, of course. A biased agent (relative to the principal) will generally get the intensive margin of regulation wrong. That is, conditional on the information the agent has generated, the agent will not choose the rule (e.g., stringency) preferred by the principal. We show that this tradeoff generally results in the principal preferring a relatively biased agent.

How does this use of agency bias interact with other tools for controlling agents? We focus in particular on review of the agent's decisions by a more aligned bureaucrat. One might worry that such review would nullify the agency bias approach to incentivizing information production by the agent. With the ultimate policy decisions made according to the preferences of the principal, a biased agent would be getting policies that could be quite far from his ideal point and thus have less incentive to generate information about regulatory opportunities. However, regulatory review also reduces the cost of agency bias, and we show that regulatory review and agency bias are thus complements. An important (and heretofore unconsidered) effect of regulatory review is to encourage the appointment of more extreme bureaucrats. We show that the principal can generally do better with a biased agent and regulatory review than she can by delegating complete authority to the agent.

An important limitation of this approach to mitigating shirking by agents is the possibility of strategic information disclosure by the agent. If the preferences of the agent and the reviewer

are too far apart, the agent may hide information from the reviewer, which results in regulations that are less tailored to the specific circumstances of the rule (e.g., the level of harm done by a pollutant). This problem can inhibit the use of agency bias as a motivational instrument.

Another consideration that can influence the principal's use of agency bias is a stock of existing rules. If the principal wants to *deregulate*, by reducing the stringency of an existing regulation, that too will require effort by the agent. Accordingly, the principal may appoint an agent who is even more anti-regulatory than the principal is in order to motivate deregulatory effort.

Our theory helps explain a range of features of the institutions, politics, and policy outcomes of administrative decisionmaking. Our main application is to presidential appointments decisions and centralized regulatory review. To mitigate shirking by a regulatory agency, the president may appoint an official to head the agency who is relatively pro-regulatory over the agency's domain and subject the agency's decisions to review by a more aligned bureaucrat. This provides a new explanation for why the president would institute centralized regulatory review to control agencies when he could instead simply appoint loyalists at the agencies. Moreover, the use of agency bias as a motivational instrument can result in an amplification in the swings of regulatory policy and heighten the role of regulatory policy in partisan politics.

Another application of our analysis is to the internal organization of agencies. Regulatory agencies are large organizations made of many, typically thousands of, bureaucrats, most of whom are career civil servants and some of whom are political appointees. The senior political appointees have a range of levers by which they can shape the preferences of subordinate staff, such as through reorganizations and, of course, hiring decisions. Our analysis shows why the head of an agency will employ relatively pro-regulatory staff to generate information about regulatory opportunities.

Our analysis also applies to the legislative branch. For example, the chairs of congressional committees may appoint biased committee staff to motivate effort at investigating legislative opportunities in the committee's purview. Similarly, committee seats may be assigned to members of Congress who are preference outliers with respect to policies in the committee's jurisdiction in order to incentivize information production (*pace* Gilligan and Krehbiel, 1987). Moreover, Congress plays a role as principal vis-à-vis the regulatory agencies both in confirming appointments to

the agencies and ex post in its oversight function. Congressional actors may also prefer relatively pro-regulatory agency staff, particularly when they intend to engage in robust oversight.

Of course, bureaucratic appointment decisions involve many other political and managerial considerations, and there are alternative potential explanations for appointments of biased agents. For example, it may be that a principal can generate support from a particular political constituency by appointing an ideologue to the agency, but can temper the effect on policy using a review process that is not well-understood by the constituency. Or a principal might choose a biased agent to counterbalance the expected effect of lobbying by an interest group. Our goal in this paper is not to empirically test competing theories but rather to analyze the implications of an explanation based on the principal's desire to mitigate shirking.

While we focus on public bureaucracies, our analysis applies to other types of organizations as well. For example, a CEO of a private firm may prefer to hire division heads who are biased towards the activities of their division and prone to “empire building” because such agents exert more effort. But the CEO may want to subject the investment decisions of those division heads to review by a more aligned agent, such as the CFO of the company.

In this paper we present a “control rights” view of how organizations address incentive problems (Grossman and Hart, 1986; Tirole, 1994). That is, rather than consider explicit incentive schemes, we consider how allocating various decision rights to particular agents can allow the principal to best achieve her goals.⁴ While (to the best of our knowledge) our analysis is new, it uses building blocks introduced in prior work on delegation in organizational economics and political science.

One key building block of our analysis is the idea that biased policy preferences can induce effort by otherwise weakly-motivated agents. Prendergast (2007) applies this insight to the problem of inducing street-level bureaucrats to identify the proper recipients for some “treatment” (e.g., a drivers license or a prison sentence)—a setting quite different from the policymaking bureaucracy considered here. Moreover, he models delegation of only search effort, not of the ultimate allocation decision, and hence does not analyze the policy-effort tradeoff that is central to our analysis.

⁴Previous work that has taken a similar approach to modeling organizational behavior includes Gilligan and Krehbiel (1987), Aghion and Tirole (1997), Dewatripont and Tirole (1999), Dessein (2002), Besley and Ghatak (2005), Bueno de Mesquita and Stephenson (2007), Gailmard and Patty (2007), Prendergast (2007), and Stephenson (2007).

Gailmard and Patty (2007) provide a model in which only bureaucrats with strong policy preferences are willing to put in effort. While this produces an implicit policy-effort tradeoff, in their model the principal cannot directly control the policy preferences of bureaucrats. Instead, they focus on how civil service tenure and the nature of Congressional delegation to agencies affect the incentives of bureaucrats to make up-front investments in expertise. In contrast, we model the appointment and delegation decisions of the principal and show why the principal will actively appoint biased bureaucrats to regulatory agencies but subject their decisions to review by more aligned bureaucrats.

Another important feature of our analysis that has roots in the existing literature is that giving the agent a policy outcome closer to his ideal policy will generate more effort by the agent. This idea is the driver of Aghion and Tirole (1997)'s analysis of the delegation of authority. Gilligan and Krehbiel (1987) similarly show that legislatures can incentivize committees to generate information by delegating policy authority to the committee. However, in contrast to our model, in those models the principal would prefer an agent who is perfectly aligned with the principal.

The idea that principals prefer agents who share their policy preferences—referred to as the “ally principle”—is a key result in the extensive formal literature on delegation in political science (Bendor, Glazer, and Hammond, 2001). Bendor and Meirowitz (2004) provide an illuminating theoretical framework that synthesizes much of this literature and note a theoretical exception to the ally principle: when information acquisition is endogenous, the principal may prefer to delegate to a non-ally if that agent faces lower information acquisition costs, or receives greater benefits from information acquisition, than an ally does. However, they do not indicate conditions where that might be the case. We posit a particular (and we think plausible) structure to policy preferences—agents vary in the weight they place on the gross benefits of regulation—and systematically analyze its implications for an important set of potential institutional designs. Another difference between our analysis and their framework is that they (and much of the literature) consider an already-identified regulatory opportunity and focus on the production of additional information relevant to that policy decision, whereas we focus on the agent's effort to identify regulatory opportunities. We show that using a biased agent can help on the extensive margin but hurt on the intensive margin.

In our application to presidential appointments, we build on a substantial existing literature on presidents' use of ideologically motivated appointees and centralized review to control agencies. These two tools of presidential control are referred to in the literature as "politicization" and "centralization," respectively. Moe and Wilson (1994) argue that while presidents can improve their control of agencies by appointing "loyal, ideologically compatible people in pivotal positions" at the agencies, such a politicization strategy will be imperfect. Political appointees at the agencies remain at an informational disadvantage vis-à-vis career civil servants, and moreover are influenced by the career staff to take the perspective of the agency. Consequently, presidents also centralize decisionmaking authority to further rein in agencies' residual noncompliance with presidential policy objectives. On this standard account, then, politicization and centralization are substitutes.

Calvert, McCubbins, and Weingast (1989)'s classic formal model of political control of agencies takes a similar approach, with centralized control only useful because of uncertainty about the preferences of appointees *ex ante*. Many subsequent formal models of regulatory review take the policy preferences of agencies and the centralized reviewer as exogenous (e.g., Bueno de Mesquita and Stephenson, 2007; Acs and Cameron, 2012). In contrast, we incorporate the appointments power and centralized review into a single model and show that the agency shirking problem can lead to *complementarity* between politicization and centralization.

The paper is organized as follows. In section 2 we provide our baseline model of the use of agency bias and regulatory review to control a policymaking agent. In section 3 we consider two extensions of our baseline model: (1) asymmetric information between the agent and the reviewer; and (2) an existing regulation that the principal wants to revise. We also discuss more generally the types of information production to which our basic results apply. In section 4 we illustrate the application of our analysis to administrative decisionmaking institutions using two historical examples from the Nixon administration: (1) the revitalization of the Federal Trade Commission; and (2) the creation of the Environmental Protection Agency and parallel emergence of centralized regulatory review. In section 5 we conclude by suggesting some implications of our analysis for the debate over the normative desirability of centralized regulatory review.

2. THE BASELINE MODEL

2.1. **Setup.** We consider a setting where there are potential regulatory opportunities in some domain, but they are initially unknown. To be concrete, consider environmental regulation and think of a regulatory opportunity as, for example, a pollutant that can be controlled. Suppose Congress has delegated authority to a regulatory agency to generate rules in this domain. Taking this delegation by Congress as exogenous, we model the institutional design problem of a Principal who wants to control the agency to further certain policy objectives. We focus on two design issues: the type of bureaucrats the Principal will appoint and whether to appoint a separate bureaucrat to review rules proposed by the agency. For now we suppose that there are no extant rules in this domain. The baseline model most directly describes the design of a new regulatory agency. We consider the revision of existing regulations in an extension to the model in section 3.2 below.

A bureaucrat at the agency can generate information about regulatory opportunities within its purview by exerting costly effort to *search*. We will refer to this bureaucrat as simply the *Agent*. In particular, to generate a probability e of finding a new regulatory opportunity, the Agent must bear a cost $\psi(e)$. For simplicity, we assume that the Agent's effort-cost function takes a quadratic form, $\psi(e) = A\frac{e^2}{2}$, with A sufficiently large so as to guarantee an interior solution.

If the Agent finds a regulatory opportunity, he can then create a regulation. A regulation is defined by its *stringency* $s \geq 0$. Think of stringency as how tightly the regulation controls the pollutant. A higher stringency would correspond to a lower parts per million regulatory standard, for example.

We assume that the Principal and the Agent are policy-motivated. In particular, we assume that stringency has gross benefit Bs and gross cost $c(s)$ to the Principal and Agent. Think of these policy payoffs as a form of social preferences. A natural interpretation is that Bs and $c(s)$ are each a fraction of the *social* benefits and costs of the regulation.⁵ The benefits that the Principal and Agent care about include, for example, a reduction in respiratory disease in society, while the costs

⁵If you prefer, that fraction could be explicit in the utility function, for example $U_A(e, s) = \gamma[Bs - \frac{s^2}{2}] - \psi(e)$, where $\gamma < 1$ is the fraction of the regulation's net benefits that the Agent internalizes via his social preferences. Including such a γ would not change any of the analysis that follows.

include the cost of installation of equipment at power plants to control the levels of the pollutant. For simplicity, we will assume that $c(s)$ takes a quadratic form so that $c(s) = \frac{s^2}{2}$.

The Principal faces two incentive problems posed by delegation to the Agent. First, the Agent bears all of the costs of his search effort. Hence the Principal faces a problem in motivating the Agent to exert effort to search. We assume that incentive pay and the like cannot implement first best effort levels, perhaps due to difficulty in measuring bureaucratic effort and output.⁶

Second, the Agent and the Principal may put a different relative weight on the benefits and costs of regulation. The Agent weights the gross costs of regulation by 1 but weights the gross benefits by $k_A \in [0, k^{max}]$. The greater is k_A , the more the Agent cares about the benefits relative to the costs of the regulation. Think of k_A as measuring how “pro-regulatory” or “mission-oriented” (where the mission is defined in terms of regulatory benefits, e.g., environmental protection) the Agent is. Similarly, the Principal weights the gross costs of regulation by 1 but weights the gross benefits by $k_P \in [k_P^{min}, k_P^{max}]$. We assume that $k_P^{min} > 0$ and $k_P^{max} < k^{max}$ so that it is always possible for the Agent to be strictly less or strictly more pro-regulatory than the Principal.

Together, these assumptions imply that the Agent’s and the Principal’s ex post payoffs from the ultimate policy decision and the Agent’s search effort are given by the utility functions,

$$(1) \quad U_A(e, s) = k_A B s - \frac{s^2}{2} - A \frac{e^2}{2},$$

and,

$$(2) \quad U_P(e, s) = k_P B s - \frac{s^2}{2},$$

respectively, where the case of not finding a regulatory opportunity corresponds to $s = 0$.

Suppose that the Principal has authority to appoint the Agent and that k_A is observable. What type of Agent will the Principal appoint? We consider this question under three alternative institutional designs that the Principal might employ: (1) full delegation of authority to the Agent

⁶Of course, in general extrinsic motivations like the desire for promotion produce some effort by bureaucrats independent of any social preferences. Our focus in this model is on the residual shirking that such extrinsic motivations do not eliminate.

to search for regulatory opportunities and set stringency; (2) delegation to the Agent of the job of searching for regulatory opportunities but institution of regulatory review by a separately appointed bureaucrat, a *Reviewer*; and (3) the case in which the Principal cannot commit to delegating authority to set stringency.

2.2. Full delegation to the Agent. We begin with full delegation. The sequence of moves in the model is:

1. Principal appoints Agent by choosing k_A .
2. Agent chooses search effort e .
3. With probability e , Agent finds a regulatory opportunity (if not, game ends).
4. Agent chooses stringency s .

Note that if $k_A \neq k_P$, then after the Agent finds a regulatory opportunity, the Principal would want to intervene and set stringency according to her preferences. We assume in this section and the next that the Principal can commit to delegating policy authority and defer the case without commitment to section 2.4 below.

2.2.1. Stringency and search effort. We find the equilibrium of the model by starting with the Agent's choice of stringency. The Agent chooses s to solve,

$$(3) \quad \max_{s \geq 0} \left\{ k_A B s - \frac{s^2}{2} \right\}.$$

Denote the solution to this problem, as a function of k_A , as $s^*(k_A)$. Given our assumptions, the solution is defined by the first order condition,

$$(4) \quad s^*(k_A) = k_A B.$$

Note that the Agent's choice of stringency is strictly increasing in k_A , since $s^{*'}(k_A) = B > 0$.

It is easy to see that the Principal, in contrast, prefers the stringency $k_P B$. Thus, the Agent chooses the Principal's preferred stringency for any regulatory opportunity he finds if and only if $k_A = k_P$.

Denote the value of the maximal policy payoff to the Agent in (3) by $V(k_A) = \frac{1}{2}k_A^2B^2$.

Turning now to the Agent's search effort, the Agent chooses effort level e to solve,

$$(5) \quad \max_{e \in [0,1]} \left\{ eV(k_A) - A\frac{e^2}{2} \right\}.$$

Denote the solution, as a function of k_A , by $e^*(k_A)$. Our assumptions guarantee that it is defined by the first order condition,

$$(6) \quad e^*(k_A) = \frac{V(k_A)}{A} = \frac{k_A^2B^2}{2A}.$$

Note that e^* increases in k_A . The Principal thus faces a tradeoff between incentives for effort provision and stringency choice. An Agent who shares the Principal's policy preferences, $k_A = k_P$, will choose the Principal's preferred stringency. But the Principal can get more search effort from an Agent who places greater weight on the benefits of regulation, $k_A > k_P$, at a cost of biased stringency. Agent bias helps the Principal on the extensive margin of regulation—more regulatory opportunities are identified—but hurts the Principal on the intensive margin of regulation—stringency is set too high. The key reason this tradeoff exists is that increasing the weight the Agent puts on the gross benefits of regulation increases both the marginal benefit of stringency (which increases the Agent's optimal stringency choice) as well as the level of the Agent's payoff from finding a regulatory opportunity (which increases his effort).

2.2.2. *Agent bias.* Consider now the Principal's optimal choice of Agent bias given this tradeoff. Denote the policy payoff to the Principal, when stringency is chosen according to preferences k_A , by $V(k_P, k_A)$, which is given by,

$$(7) \quad V(k_P, k_A) = k_P B s^*(k_A) - \frac{s^*(k_A)^2}{2} = \left[k_P k_A - \frac{k_A^2}{2} \right] B^2.$$

The Principal thus solves the problem,

$$(8) \quad \max_{k_A \in [0, k^{max}]} \left\{ e^*(k_A) V(k_P, k_A) \right\}.$$

Denote the solution k^* .

Proposition 1. *Equilibrium under full delegation.*

The Principal chooses a relatively pro-regulatory Agent, $k^ = \min\{\frac{3}{2}k_P, k^{max}\} > k_P$.*

Proposition 1 characterizes how principals will use their appointments power in the case of full delegation to the Agent. The Principal does not want to appoint an “ally” in the sense of someone who shares the Principal’s policy preferences. Rather, the Principal prefers an Agent who is biased towards the mission of the agency, despite the consequent bias to policy.

The reason it is optimal to have a relatively pro-regulatory Agent is that increasing k_A above $k_A = k_P$ causes a first order increase in the Principal’s utility by inducing the Agent to work harder to find regulatory opportunities. While this causes a distortion on the choice of stringency, by the envelope theorem the utility loss from that effect as you move away from the Principal’s ideal point is only second order. But the cost of the bias to policy gets larger as k_A increases. If k^{max} is sufficiently high ($> \frac{3}{2}k_P$), then the Principal’s optimal choice of Agent bias is in the interior (i.e., $< k^{max}$).

Finally, note that the proposition establishes that k^* is increasing in k_P , strictly so for $k^{max} > \frac{3}{2}k_P$. Hence Principals who are more pro-regulatory over the agency’s domain appoint more pro-regulatory Agents, as one would expect.

2.3. Regulatory review. Consider now how the institution of regulatory review affects the use of agency bias as a motivational instrument. In particular, suppose that rather than choosing the stringency, the Agent can only “propose a rule” to a Reviewer. If the Agent does not propose a rule, the game ends and the players get policy payoffs from the status quo, i.e., $s = 0$. If the Agent does propose a rule, then the Reviewer gets to choose the rule’s stringency.⁷

The Reviewer’s policy preferences have the same basic structure as those of the Principal’s. In particular, the Reviewer weights the gross benefits of regulation by k_R , with $k_R \in [0, k^{max}]$, so the

⁷Our assumption that the Reviewer gets to set stringency following a proposal is not necessary for our basic results. The alternative assumption that the Reviewer simply has a veto right would yield similar results, as discussed in more detail in footnote 8 below.

Reviewer's preferences over the ultimate policy decision can be represented by the utility function,

$$(9) \quad U_R(s) = k_R B s - \frac{s^2}{2}.$$

The sequence of moves in the model is now:

1. Principal appoints Agent and Reviewer by choosing k_A and k_R .
2. Agent chooses search effort e .
3. With probability e , Agent finds a regulatory opportunity (if not, game ends).
4. Agent chooses whether to propose a rule to the Reviewer (if not, game ends).
5. Reviewer chooses stringency s .

2.3.1. *Stringency and search effort.* We begin with the subgame in which the Reviewer chooses stringency. The Reviewer's optimal choice of stringency is given by the now familiar function, $s^*(k_R) = k_R B$. Given this equilibrium strategy of the Reviewer, the policy payoff to the Agent, when stringency is chosen according to preferences k_R , is given by $V(k_A, k_R)$ (the same function as in (7) but evaluated at (k_A, k_R) instead of (k_P, k_A)). Note that $V(k_A, k_R) > 0$ if and only if $k_A > k_R/2$. If this condition holds, the Agent chooses search effort $e = \frac{V(k_A, k_R)}{A}$. Otherwise, the Agent will exert no effort and find no regulatory opportunities. Denote the Agent's optimal search effort by $e^*(k_A, k_R)$.⁸

2.3.2. *Agent and Reviewer bias.* Consider now the choice of the Principal. Her choice problem is to choose the Agent and Reviewer in order to affect both policy and effort. Formally, she solves,

$$(10) \quad \max_{k_A, k_R} \left\{ e^*(k_A, k_R) V(k_P, k_R) \right\}.$$

Denote the choices of k_A and k_R that solve this problem by k_A^* and k_R^* .

⁸ If we instead assume that the Reviewer only has the right to veto the Agent's proposal, and not the right to set stringency, then the Agent will propose the policy closest to his ideal point that the Reviewer will accept and the Reviewer will accept it. The Agent will anticipate that outcome and choose his effort accordingly. The key insight in connecting this alternative model to our baseline model is that in both cases the Principal essentially completely chooses the stringency to be implemented by choosing the Reviewer. In the baseline model she implements stringency s^* by choosing $k_R = B/s^*$ while in the alternative model she does it by choosing the k_R such that $k_R B s^* - (s^*)^2/2 = 0$. The key results, that she will choose an extreme Agent and a Reviewer that leads to an implemented policy that is more stringent than her own ideal policy in order to induce extra effort by the Agent, is unchanged.

Proposition 2. *Equilibrium under regulatory review.*

- (1) *The Principal chooses a maximally pro-regulatory Agent, $k_A^* = k^{max}$.*
- (2) *The Principal chooses a Reviewer who is more pro-regulatory than she is, but not maximally pro-regulatory, $k_P < k_R^* < k^{max}$, and more pro-regulatory Principals appoint more pro-regulatory Reviewers, $\frac{\partial k_R^*}{\partial k_P} > 0$.*
- (3) *The Principal's choice of Reviewer is less pro-regulatory than her choice of Agent under full delegation, $k_R^* < k^*$.*
- (4) *The Principal strictly prefers a regulatory process that includes regulatory review to full delegation to the Agent.*

The Principal, whatever her policy preferences, prefers a maximally pro-regulatory Agent because the tasks of search and policymaking are separated. The Agent's search effort increases in his bias, and with the tasks separated, there is no policy cost from having a biased Agent. Nevertheless, the tradeoff between policy and effort remains at a different level. The Agent is willing to work even harder as the Reviewer becomes more similar to him, and the Principal will use that extra incentive to induce even higher effort, to some extent. The Principal prefers a relatively pro-regulatory Reviewer because the consequent improvement in the Agent's effort outweighs the cost from the resulting bias to policy.

Finally, the Principal always strictly prefers to have regulatory review rather than to delegate fully to the Agent. There are two reasons for this. First, regulatory review allows the Principal to choose a *more extreme* Agent. Holding the preferences of the Reviewer (and, therefore, policy) fixed, this move alone increases Agent effort, resulting in more regulatory opportunities discovered, and increases the Principal's payoff. An implication of the model, then, is that the availability of regulatory review encourages greater use of agency bias as a motivational instrument. For example, all else equal, we expect presidents to appoint more extreme agency staff for agencies subject to OIRA review than for independent agencies.

Second, the presence of regulatory review reduces equilibrium policy bias. With the additional effort provided by an extreme Agent, the Principal prefers to choose a Reviewer who is more

aligned with her policy preferences than the optimal Agent under full delegation because the benefit in terms of reduced policy bias is greater than the cost in terms of less effort from the Agent when the Agent is set optimally at $k_A = k^{max}$.

2.4. Delegation without policy commitment. In both the full delegation case and the regulatory review case considered above, we assumed that the Principal can commit to delegating authority over stringency to either the Agent or the Reviewer. This is a reasonable way to model regulatory review in some contexts. For example, in the case of centralized regulatory review by the White House, the many competing demands on the president's time and attention provide a commitment device that at a minimum limits the number of rules in which the president can personally intervene.⁹ However, in other settings the Principal may not be able to commit to not intervening ex post. Hence we consider here the case in which the Principal delegates search to the Agent but retains authority over stringency. This case is equivalent to a special case of the regulatory review model considered above with k_R constrained to be equal to k_P .

The sequence of moves in the model is now:

1. Principal appoints Agent by choosing k_A .
2. Agent chooses search effort e .
3. With probability e , Agent finds a regulatory opportunity (if not, game ends).
4. Agent chooses whether to propose a rule to the Principal (if not, game ends).
5. Principal chooses stringency s .

Proposition 3. *Equilibrium under delegation without policy commitment.*

- (1) *The Principal chooses a maximally pro-regulatory Agent, $k_A^* = k^{max}$.*
- (2) *The Principal strictly prefers committing to regulatory review by a separate Reviewer to delegation without policy commitment, which in turn she strictly prefers to full delegation.*

⁹Aghion and Tirole (1997) show that giving the principal a broad span of control can be a useful device to commit to delegation, leading to greater initiative by the agent. Moreover, analyses of the practice of regulatory review show that the president indeed only rarely personally intervenes in rulemakings under OIRA review (Livermore and Revesz, 2012).

The results here are similar to the case with regulatory review by a separate Reviewer. In particular, the Principal still prefers a maximally pro-regulatory Agent even when she cannot commit to not intervening herself *ex post*. But part (2) of the proposition makes clear that the Principal would do better if she could commit to delegating policy authority to a separate Reviewer. The reason is that the Principal could then get more effort out of the Agent by tilting the Reviewer's preferences towards the Agent's. In turn the Principal always prefers delegation without policy commitment to full delegation to the Agent. The reason is that in this model the benefit to the Principal of retaining policy authority in terms of reduced policy bias is always greater than the cost in terms of reduced Agent effort.

3. EXTENSIONS

In our baseline model above, we abstracted from a number of important issues. We turn now to two issues that can moderate or reverse the prediction of our baseline model that bureaucratic principals will employ relatively pro-regulatory agents. First, we consider the possibility that a biased Agent will manipulate the information available to the Reviewer. Second, we consider the case in which the Principal wants to identify opportunities to change *existing* regulations. Finally we discuss the extent to which our results extend to other types of information production by agencies.

3.1. Asymmetric information and regulatory review. In our baseline model, we assumed that there is no uncertainty about the benefits of a regulatory opportunity once the Agent discovers it and proposes a rule. Suppose now that the Agent has private information about the marginal benefit of a regulatory opportunity. The idea that bureaucrats have an informational advantage over those who delegate to them is central in the study of bureaucratic politics (Niskanen, 1975; McCubbins, Noll, and Weingast, 1987; Stephenson, 2011). Relative to this literature, our contribution is in considering how this information asymmetry affects the principal's use of agency bias.¹⁰ We show

¹⁰Dessein (2002) considers a different, but related, problem, in which the principal takes the agent's ideology as given, but can choose the ideology of the supervisor to induce revelation. In his model, allowing the principal control over both ideologies would be uninteresting, since there is no effort dimension, so the principal would always simply choose someone like himself. We allow the principal control over both, but with an effort-policy tradeoff.

that the risk that the agent will strategically withhold information curbs the principal's willingness to appoint extreme agents.

3.1.1. *Setup.* To be concrete, suppose that if the Agent finds a regulatory opportunity, it is one of two types: it is a high value opportunity with marginal benefit B_H with probability q and a low-value opportunity with marginal benefit B_L with probability $1 - q$, with $0 < B_L < B_H$. However, the Agent only learns what type the opportunity is with probability p . Moreover, suppose that, if the Agent learns what the marginal benefit is, he can choose whether to disclose this information to the Reviewer. We assume that the Agent can credibly disclose the true B , but that if the Agent hides B , the Reviewer does not know whether the Agent knows B .

The sequence of moves in the model is now:

1. Agent chooses search effort e .
2. With probability e , Agent finds a regulatory opportunity (if not, game ends).
3. Nature chooses marginal benefit B of the regulatory opportunity.
4. With probability p , Agent learns B .
5. Agent chooses whether to propose a rule to the Reviewer (if not, game ends).
6. If Agent knows B , he chooses whether to disclose B to Reviewer.
7. Reviewer chooses stringency s .

Our equilibrium concept is perfect Bayesian equilibrium (PBE).

3.1.2. *Stringency, disclosure, and search effort.* To find the equilibrium, we start with the Reviewer's choice of stringency. If the Agent has conveyed the marginal benefit of the regulatory opportunity to the Reviewer, then the Reviewer's choice problem is,

$$(11) \quad \max_{s \geq 0} \left\{ k_R B s - \frac{s^2}{2} \right\},$$

where the marginal benefit is denoted $B \in \{B_L, B_H\}$. Note that the Reviewer's objective function is linear in B so that, if B is uncertain, all that matters to the Reviewer is the expected value of B . If the Agent has not disclosed the marginal benefit, then the Reviewer will form beliefs about the marginal benefit that are, in equilibrium, consistent with the Agent's disclosure strategy.

Hence, (11) is also the Reviewer's problem in the subgame in which the Agent has not disclosed the marginal benefit, but now B denotes the expected marginal benefit based on the Reviewer's beliefs. The solution to this problem is $s^*(k_R, B) = k_R B$.

Now consider the Agent's disclosure strategy, for now taking the preference parameters k_A and k_R as exogenous. We begin by considering the conditions under which there is a *full-disclosure* equilibrium in which the Agent always discloses any information he obtains. A full-disclosure equilibrium exists if and only if the Agent receives a higher payoff when he discloses than when he hides, both when he observes B_L and when he observes B_H .

Consider for example a full-disclosure equilibrium for the case in which $k_A > k_R$. Suppose that the Agent observes B_L . If he discloses, then the Reviewer will select stringency $k_R B_L$, which is lower than the Agent's preferred stringency, $k_A B_L$. If instead he deviates by hiding this information, then the Reviewer would believe that the Agent did not observe the marginal benefit and hence that the opportunity has expected marginal benefit \bar{B} . Thus the Reviewer would choose a higher stringency, $k_R \bar{B} > k_R B_L$. For this deviation not to be attractive to the Agent, hiding must result in the Reviewer overshooting by selecting a stringency that is *above* the Agent's ideal stringency for B_L . In fact, this implemented stringency must be so far above the Agent's ideal that the Agent prefers the too-low stringency chosen when he tells the truth. Hence it is easier to maintain disclosure as $B_H - B_L$ grows, since hiding B_L would induce a bigger jump in stringency, leading to overshooting for a larger set of preference parameters k_A and k_R . Of course, with $k_A > k_R$ the Agent never has incentive to hide B_H . The case in which $k_A < k_R$ is similar but the relevant incentive constraint applies to hiding B_H instead of B_L .

Analysis of these two incentive constraints yields constraints on how far apart the policy preferences of the Agent and the Reviewer can be for a full-disclosure equilibrium to exist, which are summarized in the following lemma.

Lemma 1. *A full-disclosure equilibrium exists if and only if $-\frac{k_R(1-q)}{2} \frac{B_H - B_L}{B_H} \leq k_A - k_R \leq \frac{k_R q}{2} \frac{B_H - B_L}{B_L}$.*

Let $\bar{k}(k_R) \equiv k_R + \frac{k_R q}{2} \frac{B_H - B_L}{B_L}$ represent the highest value of k_A , given k_R , for which the full-disclosure equilibrium exists, given by Lemma 1.

The value to the Agent of finding a regulatory opportunity in a full-disclosure equilibrium is given by,

$$\begin{aligned}
(12) \quad V^{disc}(k_A, k_R) &= qp[k_A B_H s^*(k_R, B_H) - \frac{s^*(k_R, B_H)^2}{2}] + \\
&\quad (1 - q)p[k_A B_L s^*(k_R, B_L) - \frac{s^*(k_R, B_L)^2}{2}] + \\
&\quad (1 - p)[k_A \bar{B} s^*(k_R, \bar{B}) - \frac{s^*(k_R, \bar{B})^2}{2}] \\
&= [k_A k_R - \frac{k_R^2}{2}] [p[qB_H^2 + (1 - q)B_L^2] + (1 - p)\bar{B}^2] \\
&= [k_A k_R - \frac{k_R^2}{2}] [\bar{B}^2 + pq(1 - q)(B_H - B_L)^2].
\end{aligned}$$

The Agent thus chooses search effort to satisfy the first-order condition $e = \frac{V^{disc}(k_A, k_R)}{A}$.

If a full-disclosure equilibrium does not exist, then the only equilibrium that exists is a *hiding* equilibrium in which the Agent hides information about one of the two states. Since (as we show below) the Principal will choose an Agent with relatively high k_A , much as in the baseline model, the relevant case is $k_A > k_R$, in which case the only possible hiding equilibrium is one in which the Agent hides B_L . Because they are not reached on the equilibrium path, for brevity we omit discussion of subgames with $k_A < k_R$.

Denote the Reviewer's beliefs in such an equilibrium about the probability that $B = B_H$ in the subgame in which the Agent does not disclose by \hat{q} , with corresponding expected marginal benefit \hat{B} . The Reviewer's equilibrium beliefs are given by Bayes' Rule:

$$(13) \quad \hat{q} \equiv \frac{q(1 - p)}{(1 - p) + p(1 - q)} < q.$$

The Agent's expected policy payoff from finding a regulatory opportunity is thus given by,

$$\begin{aligned}
V^{hide}(k_A, k_R) &= qp[k_A B_H s^*(k_R, B_H) - \frac{s^*(k_R, B_H)^2}{2}] + \\
&\quad (1 - qp)[k_A \hat{B} s^*(k_R, \hat{B}) - \frac{s^*(k_R, \hat{B})^2}{2}] \\
(14) \qquad &= \left[k_A k_R - \frac{k_R^2}{2} \right] \left[qp B_H^2 + (1 - qp) \hat{B}^2 \right] \\
&= \left[k_A k_R - \frac{k_R^2}{2} \right] \left[\bar{B}^2 + \frac{qp(1 - q)^2}{1 - qp} (B_H - B_L)^2 \right].
\end{aligned}$$

The Agent chooses search effort $e = \frac{V^{hide}(k_A, k_R)}{A}$.

For parameter values for which both the full-disclosure equilibrium and this hiding equilibrium exist, the full-disclosure equilibrium Pareto dominates the hiding equilibrium for the Agent, Reviewer, and Principal. First, observe that $V^{disc}(k_A, k_R) > V^{hide}(k_A, k_R)$. This implies that the Agent's search effort is higher in the full-disclosure equilibrium, and that the Agent is better off in the full-disclosure equilibrium. Furthermore, the policy payoffs to the Reviewer and to the Principal from the Agent finding a regulatory opportunity are the same functions $V^{disc}(k, k_R)$ and $V^{hide}(k, k_R)$ but with k_R and k_P , respectively, substituted for the first argument of the functions. Therefore this implies that the Reviewer and the Principal are also better off in the full-disclosure equilibrium than in the hiding equilibrium.¹¹ Because the full-disclosure equilibrium Pareto dominates the hiding equilibrium, we will assume that the full-disclosure equilibrium is played if it exists.¹²

3.1.3. Agent and Reviewer bias. With this characterization of the equilibrium play in the Reviewer-Agent subgames in hand, let us turn finally to the Principal's choice of k_A and k_R . Let k_A^{**} and k_R^{**} denote the Principal's equilibrium choice of the type of Agent and Reviewer, respectively. The following proposition summarizes the equilibrium under asymmetric information.

¹¹It is also the case that for $k_A < k_R$ (a case we omit because it is not on the equilibrium path), the full-disclosure equilibrium Pareto dominates the hiding equilibrium (in this case hiding B_H).

¹²Mixed-strategy equilibria, in which the Agent mixes between hiding and disclosing for one value of the marginal benefit, exist for some parameter values. But whenever a mixed strategy equilibrium exists, so does a full-disclosure equilibrium, which Pareto dominates it. We similarly assume that in such cases the Agent and Reviewer play the full-disclosure equilibrium.

Proposition 4. *Equilibrium under asymmetric-information.*

(1) *Holding \bar{B} fixed, there exists a unique threshold $T > 0$, such that:*

(a) *(Hiding) If $B_H - B_L < T$, then the Principal appoints a maximally biased Agent, $k_A^{**} = k_A^* = k^{max}$, and a Reviewer who is more pro-regulatory than she is, but not maximally pro-regulatory, $k_P < k_R^{**} < k^{max}$. The Agent hides the marginal benefit when it is equal to B_L .*

(b) *(Full Disclosure) If $B_H - B_L > T$, then the Principal appoints an Agent who is indifferent about whether to disclose the marginal benefit for low marginal benefit opportunities, $k_A^{**} = \bar{k}(k_R^{**})$, and a Reviewer who is more pro-regulatory than she is, but less pro-regulatory than the Agent, $k_P < k_R^{**} < k_A^{**}$. The Agent always discloses B .*

(2) *This T is decreasing in k_P and increasing in \bar{B} .*

(3) *The Principal strictly prefers a regulatory process that includes regulatory review to full delegation to the Agent.*

Part (1)(a) of the proposition shows that when there is relatively little information asymmetry between the Agent and the Reviewer about the benefits of regulation after the Agent proposes a rule ($B_H - B_L$ is small), then the Principal chooses an extreme pro-regulatory Agent. But unlike in the baseline model with regulatory review, the incentive effect of Agent bias comes at a cost. This cost is different than the cost in the full delegation case, namely, here appointing an extreme Agent results in a loss of information that would help to fine-tune regulation. Conflict and deception in the Agent and Reviewer relationship are avoidable in this model, but in this area of the parameter space the Principal (second-best) optimally chooses not to avoid it.

The Principal only chooses bureaucrats who will fail to communicate when the cost associated with the loss of information is small. If the asymmetric information is important ($B_H - B_L$ is large), then part (1)(b) of the proposition shows that the Principal will forgo the extra effort he might get by appointing a very extreme Agent and instead choose an Agent more in line with the Reviewer in order to guarantee disclosure. Large $B_H - B_L$ makes inducing disclosure more attractive to the

Principal for two reasons. First, the asymmetric information about the true state is more valuable as that information has a bigger effect on the preferred stringency. When $B_H - B_L$ is large, there is a large gap between the preferred stringency in each state. Second, as $B_H - B_L$ grows it actually becomes easier to induce full disclosure (in the sense that the maximum gap between the Agent's and Reviewer's preferences under which full disclosure is an equilibrium grows).

Part (2) of the proposition implies that there will be less conflict between the Reviewer and the Agent, and in particular less manipulation of information by the Agent, when the Principal is more pro-regulatory over the agency's domain. Thus, one might think that a Republican president will get more disclosure from national security agencies, while Democrats will get more disclosure from the EPA.

Furthermore, T increasing in \bar{B} , along with part (1) of the proposition, implies that we will observe more extreme appointments to agencies, and more conflict between the Agent and the Reviewer, when finding regulatory opportunities is particularly important relative to the importance of the asymmetric information about the regulatory opportunity that remains after the Agent proposes a rule. A range of institutions serve to reduce this information asymmetry and thus, according to our model, encourage the use of agency bias as a motivational instrument. For example, the Administrative Procedure Act provides interested parties notice of rulemaking and an opportunity to comment on proposed rules and requires agencies to state the "basis and purpose" for their decisions¹³ and thereby reduces the information asymmetry between the agencies and OIRA. Similarly, one way of understanding the reason that OIRA requires agencies to provide cost-benefit analysis with their rules is as a way of forcing the agency to disclose information (Posner, 2001). For domains in which these institutions are effective, we expect to see presidents appointing particularly biased agency staff and subjecting their decisions to review by more centrist bureaucrats at OIRA. In contrast, for domains in which information asymmetry is more important, presidents can be expected to appoint less biased agency staff. These may include domains in which regulatory issues are highly technical and for which there are not many competing interest groups who can reduce the information asymmetry by providing information.

¹³5 U.S.C. § 551 *et seq.*

Finally, the introduction of information asymmetry does not change the fact that the Principal prefers to employ regulatory review. To see this, note that the Principal could still recreate the full-delegation outcome here by choosing $k_A = k_R$, since the Agent gains nothing from hiding information when the Reviewer exactly shares his policy preferences. The Principal never makes that choice, and actually strictly prefers not to. The reasons regulatory review is useful are the same as in the baseline model.

3.2. Revising existing regulations. In our baseline model, the Principal is born into a regulatory vacuum and thus can move regulatory policy in only one direction—increased regulation. Suppose instead that there is a stock of extant regulations and the Principal would like to revise them.

3.2.1. Setup. For simplicity, assume that there is a single existing regulation with marginal benefit B , which was set according to the preference parameter k_O , so that its current stringency is $s^*(k_O) = k_O B$ (the subscript O stands for “old”). We assume that $k_P \neq k_O$, so the Principal knows she would like to change existing regulations. But suppose that changing regulations requires additional information. For example, the precise regulations that can be usefully revised may be unknown, or changing stringency may require additional information.

The Principal must delegate to an Agent the task of searching for a revision opportunity. In particular, to generate a probability e of identifying an appropriate opportunity to revise an existing regulation, the Agent must bear a cost $\psi(e) = A\frac{e^2}{2}$. If the Agent finds an opportunity, he can propose a rule to the Reviewer, who then chooses stringency, s .¹⁴ This model nests the baseline regulatory review model from section 2.3 above in the particular case where $k_O = 0$.

3.2.2. Stringency and search effort. Beginning with the stringency choice of the Reviewer, if the Agent proposes a revision of an existing regulation, the Reviewer sets stringency at $s^*(k_R) = k_R B$.

Given this stringency setting strategy, the Agency receives the following incremental payoff over his policy payoff from the status quo from finding and proposing a revision,

$$(15) \quad V(k_A, k_R) = B[(k_A k_R - \frac{k_R^2}{2}) - (k_A k_O - \frac{k_O^2}{2})].$$

¹⁴For brevity, we only explicitly model the case with regulatory review; the full-delegation case is a straightforward extension.

The Agent's policy payoff from uncovering an opportunity to revise regulation is non-negative as long as the reviewer's preference is closer to his own than the old regulation was, i.e. $|k_R - k_O| \leq |k_A - k_O|$. If that condition fails, the agent will exert no effort and propose no changes. Otherwise, he will choose effort to satisfy the first-order condition $e = \frac{V(k_A, k_R)}{A}$ and propose any regulation he finds. Denote the agent's optimal effort, as a function of k_A and k_R , by $e^*(k_A, k_R)$.

3.2.3. *Agent and Reviewer bias.* Finally, consider the Principal's choice of k_A and k_R . She solves the following problem:

$$(16) \quad \max_{(k_A, k_R) \in [0, k^{max}]^2} \left\{ e^*(k_A, k_R) V(k_P, k_R) \right\}.$$

As before, let (k_A^*, k_R^*) denote the solution to this problem. The following proposition characterizes the equilibrium in this extension with regulatory revision.

Proposition 5. *Equilibrium with revision of existing regulations.*

- (1) *If $k_O < k_P$, the equilibrium Agent preference is $k_A^* = k^{max}$. If $k_O > k_P$, the equilibrium Agent preference is $k_A^* = 0$.*
- (2) *The Principal appoints a Reviewer with preferences strictly between her own and the Agent's.*
- (3) *A more pro-regulatory Reviewer is selected as the Principal becomes more pro-regulatory or the old regulation gets less extreme, i.e., $\frac{\partial k_R^*}{\partial k_P} > 0$, $\frac{\partial k_R^*}{\partial k_O} < 0$.*

The key factors guiding the Principal's choice are similar to those in the baseline model. The Principal wants the Agent to work hard to uncover opportunities to revise regulation, and choosing an extreme Agent gives the strongest such incentives. The key difference is what "extreme" means in this context. Here, the Principal wants an Agent who is extreme in the direction that she wants to move regulation. If $k_O < k_P$, the Principal wants to tighten regulation, and the optimal Agent is exactly like the baseline case ($k_A^* = k^{max}$). In contrast, when $k_O > k_P$, the Principal wants to deregulate, and the optimal Agent is extremely anti-regulation ($k_A^* = 0$). In either case, the Principal will choose a Reviewer with preferences between the Principal's and the Agent's, since

moving the Reviewer closer to the Agent will again entail an effort-versus-policy trade-off. Part (3) shows that the bias of the Reviewer will respond to changes in the relative preferences of the Principal and the extant regulation in intuitive ways.

This extension of the model suggests several empirical implications. First, the shirking problem results in an amplification of both partisan conflict over regulation and the cycling of regulatory policy. To see this, suppose that a president enters office with relatively pro-regulatory preferences in some domain with an extant stock of regulation such that the president wants to move policy towards more stringent regulation. Because the president faces an agency that will shirk, under our theory the president will want to appoint a very pro-regulatory agency, since such an agency will shirk less. Moreover, the president will appoint a reviewer who is more aligned with the president but still biased towards regulation. The result of this strategy, however, is that the policies set by the agency and reviewer will be even more stringent than the president prefers.

Given this extremism, regulatory policy becomes a salient political issue, and leaders of the competing party decry the “over-regulation” of the incumbent administration. Suppose this competing party wins the next election, installing a relatively anti-regulatory president. The anti-regulatory president wants to move regulatory policy to become less stringent. But this president also faces a problem of shirking at the agency. So to motivate the agency to deregulate, he appoints a relatively anti-regulatory agency, who will work hard to deregulate but, together with an optimally chosen reviewer, will set policy even looser than the president prefers. This makes regulatory policy a salient issue in the next round of electoral politics. And so forth.

If presidents did not face this shirking problem—or did not use biased preferences of agencies to mitigate it—the swings of regulatory policy would be smaller in amplitude. The use of zealots as a way to mitigate shirking results in higher variance in regulatory policy and heightens the role of regulatory policy in electoral politics. Interestingly, effective regulatory review can dampen the swings of policy, but will *amplify* the swings of preferences at the agencies. While reviewers are chosen to be more aligned than are agencies under full delegation, the agencies are chosen to be more extreme when regulatory review is employed.

For simplicity, throughout the paper we have assumed that the Agent works on a single regulation. In a environment with multiple regulations (or regulatory opportunities), little changes as long as the Principal wants to move the stringency of all regulations in the same direction. The Principal still prefers an extremist Agent and a Reviewer with preferences between her own and the Agent's. However, if there are both opportunities to revise regulations that are too tight and opportunities to tighten regulations that are too loose (or to find new regulatory opportunities) that must be delegated to the same Agent, the result is a multi-task principal-agent problem (Holmstrom and Milgrom, 1991). The Principal's choice of Agent will balance incentives to work on both directions, and could result in extremist Agents in either direction or even in relatively centrist Agents depending on the relative importance of the regulatory tasks and the deregulatory tasks. The Reviewer, of course, would continue to be located between the Principal and the Agent.¹⁵

This multi-task perspective sheds new light on the dynamics of agency bias. Downs (1967, p. 5) argues that new agencies are "initially dominated either by advocates or zealots." In our theory, new agencies will also be led by zealots, since identifying new regulatory opportunities is the goal. However, over time the agency creates a stock of rules that are more stringent than the appointing president would prefer (on average). As this stock grows, the same president may begin appointing less extreme agents in order to encourage some effort on the deregulatory tasks without abandoning all effort on the search for new regulations. Finally, a more extreme deregulatory motive can be triggered when a president from the party that puts less weight on the agency's regulatory objective is elected. Such a president may then appoint a "deregulatory zealot" in order to incentivize the agency to identify deregulatory opportunities and reset the stringency of the stock of agency rules. These dynamics implied by our theory may explain some historical episodes of deregulation, for example at the beginning of the Reagan administration.

¹⁵One might wonder why the regulatory and deregulatory tasks must be assigned to the same Agent. Our model suggests that creating a specialized regulatory agency to find opportunities to increase regulation and a separate specialized deregulatory agency to find opportunities to loosen regulation may result in more efficient incentives. In practice, regulatory and deregulatory tasks are typically assigned to the same agency. One potential explanation for this is that the two tasks are complementary in the production function. Searching for regulatory opportunities may as a byproduct also produce information about deregulatory opportunities, and vice-versa.

3.3. Other types of information production. In the baseline model, we focused on a particular type of information production by agencies: the search for “regulatory opportunities.” A natural question is whether our results extend to other types of information production by agencies.

To explore how general our results are, it is helpful to provide a somewhat more general formulation of the basic model. Consider the case with regulatory review and suppose that the Agent can generate “information” with probability e at a cost of $\psi(e) = A\frac{e^2}{2}$. For simplicity, suppose that if the Agent finds information, the Reviewer learns it as well (that is, abstracting from the asymmetric information issues considered above). The Reviewer will choose policy to maximize his policy payoff given the information generated by the Agent. Denote the policy payoff to the Agent when the Agent has found information and when not by $V^I(k_A, k_R)$ and $V^N(k_A, k_R)$, respectively. The Agent chooses effort e to solve,

$$(17) \quad \max_{e \in [0,1]} \left\{ V^N(k_A, k_R) + e[V^I(k_A, k_R) - V^N(k_A, k_R)] - \psi(e) \right\}.$$

The solution is $e^*(k_A, k_R) = V^I(k_A, k_R) - V^N(k_A, k_R)$. Thus, the Agent’s effort is increasing in k_A if and only if $V_1^I(k_A, k_R) - V_1^N(k_A, k_R) > 0$. In economic terms, this condition simply requires that Agents value information more the greater is their k_A . In our baseline model, we had $V^N(k_A, k_R) = 0$, so this condition amounted to $V_1^I(k_A, k_R) > 0$, which is true in that model.

It is also true in other settings. For example, suppose that the agency has already identified a regulatory opportunity but the marginal benefit of the opportunity B is unknown. As in section 3.1, suppose that it is a high value opportunity with marginal benefit B_H with probability q and a low-value opportunity with marginal benefit B_L with probability $1 - q$. But now suppose that the Agent chooses effort e to investigate the marginal benefit at a cost of $\psi(e)$ and discovers the true marginal benefit with probability e . It can be seen from inspection of (12) above that,

$$(18) \quad V^I(k_A, k_R) - V^N(k_A, k_R) = (k_A k_R - \frac{k_R^2}{2})q(1 - q)(B_H - B_L)^2,$$

and,

$$(19) \quad V_1^I(k_A, k_R) - V_1^N(k_A, k_R) = k_R q(1 - q)(B_H - B_L)^2 > 0.$$

The Agent’s effort thus increases in k_A in this general setting as well. Moreover, it can be shown that all of our results in Proposition 2 from the baseline model hold here as well.

However, there are other environments in which more pro-regulatory agents put *lower* value on information and hence work *less* hard to generate information. One obvious and important example consider above is for information required to *deregulate*. To give another example, consider the model of investigating the value of B just introduced, but suppose now that there is an upper bound on stringency, $s \leq \bar{s}$. Suppose that if the Principal were not informed about B , the Principal would prefer $s = \bar{s}$, but if the Principal knew that $B = B_L$, then the Principal would prefer $s = k_P B_L < \bar{s}$. It can easily be shown that in this model, $V_1^I(k_A, k_R^*) - V_1^N(k_A, k_R^*) < 0$, so that Agents who put lower weight on the gross benefits of regulation will work harder to generate information about B . The intuition for why is similar to the intuition in the case of deregulatory opportunities considered above. If the Agent does not discover the marginal benefit, policy will be set at \bar{s} . If the Agent discovers that $B = B_L$, the policy will be set at $k_R B_L < \bar{s}$. Agents value more highly this opportunity to set stringency lower the lower the weight they put on the gross benefits of regulation (i.e., the lower their k_A).

4. APPLICATIONS

Our simple model predicts that, absent a strong deregulatory motive, bureaucratic principals will typically employ relatively pro-regulatory agents in order to better motivate the agents to generate information about regulatory opportunities but subject their decisions to review by a more aligned bureaucrat. The basic structure that we model—a policy-motivated principal who can choose the agents responsible for regulating in some domain and can structure their decisionmaking process—appears in many contexts across public bureaucracies. We focus our discussion of applications on two related issues: (1) presidential appointments of the heads of regulatory agencies and review of those agencies’ decisions; and (2) the appointment of subordinate agency staff by the heads of regulatory agencies. To make our discussion more concrete, we consider two historical examples from the 1970s that can be interpreted through the lens of our theory: the revitalization of the

Federal Trade Commission (FTC) and the creation of the EPA with a corresponding increase in centralized regulatory review.

4.1. The Federal Trade Commission. Created by Congress in 1914, the FTC is responsible for administering both antitrust laws and more general consumer protection laws. We can apply our model to understand the strategies that three different principals used to control the FTC: the president, Congress, and the FTC chairman. By statute the FTC is an independent agency outside of the direct control of the president, who simply gets to appoint commissioners and to choose which commissioner will serve as chairman. From the perspective of the president, then, this case should be thought of as entailing full delegation to the agency to set policy. Congress is another important principal. The president's appointments to the FTC are subject to the advice and consent of the Senate. Moreover, Congressional oversight *ex post* can be thought of as a form of the regulatory review that we model. Finally, the FTC chairman has broad authority to shape the personnel of the FTC and to structure its internal organization. While the staff of the FTC do much of the agency's actual work, by statute the commissioners themselves retain authority to make final regulatory decisions. So under this interpretation, the chairman is a principal who can appoint agents but cannot commit to delegating policymaking authority to the agents. As we recount below, in the 1970s each of these principals worked to appoint relatively pro-regulatory bureaucrats at the FTC to generate more regulatory effort at the agency. And when that strategy resulted in overreach by the FTC, Congress stepped in *ex post* to limit the bias of FTC policy.

By 1969, the FTC was widely regarded as a moribund agency. Spurred by two reports critical of the FTC, one by the Nader organization¹⁶ and the other by the American Bar Association,¹⁷ in the fall of that year President Nixon announced that “the time has now come for the reactivation and

¹⁶Cox, Fellmeth, and Schulz (1969). The Nader Report revealed an organization that was passive and ineffective, relying on reports from consumers to detect violations and rarely bringing enforcement actions, and used by congressmen as a source of patronage employment.

¹⁷American Bar Association (1969).

revitalization of the FTC.”¹⁸ Nixon cited in particular the need for more effort to identify regulatory opportunities, stating that the FTC “should seek out new information on consumer problems through more energetic field investigations.”¹⁹

Key to Nixon’s approach was the appointment of Casper Weinberger as chairman of the FTC. Nixon reported that Weinberger had “assured me that he intends to initiate a new era of vigorous action.”²⁰ Personnel was a major focus during Weinberger’s brief six-month chairmanship, during which he discharged 18 of the top 31 staff member at the FTC.²¹ Nixon appointed Miles Kirkpatrick to succeed Weinberger, who continued the focus on personnel, replacing about a third of the mid- and lower-level staff with new people who had a “strong commitment to consumer protection.”²² As a result, the ideology of the FTC staff became dramatically more pro-consumer protection, with many FTC staff members decidedly more pro-consumer protection than Nixon and the FTC chairman. For example, Kirkpatrick, with the consent of the Nixon White House, appointed two “card-carrying activist Democrats” to the high-level FTC posts of Director of the Bureau of Consumer Protection and Director of the Bureau of Economics.²³

The revamped FTC dramatically increased its enforcement and regulatory activities. The first stage of this expansion, from 1970 to 1975, entailed more aggressive case-by-case enforcement against unfair and deceptive business practices.²⁴ The second stage began in 1975, when Congress expressly delegated to the FTC the authority to issue industry-wide rules to regulate unfair and deceptive industry practices.²⁵ Prior to this legislation, the FTC’s legal authority to promulgate industry-wide rules was widely doubted and only sparingly asserted.²⁶ The FTC used the authority

¹⁸President Nixon, Special Message to the Congress on Consumer Protection, Oct. 30, 1969.

¹⁹*Id.*

²⁰*Id.*

²¹Clarkson and Muris (1981, p. 4).

²²Harris and Milkis (1996, p. 167).

²³Senate Committee on Government Operations, Study on Federal Regulation: The Regulatory Appointments Process, vol. 1, p. 217, S. Doc. No. 25, 95th Cong., 1st Sess. (1977).

²⁴Harris and Milkis (1996, p. 181)

²⁵Magnuson-Moss Warranty - FTC Improvement Act of 1975, Jan. 4, 1975, Pub. L. 93-637, 88 Stat. 2193.

²⁶The first judicial recognition of the FTC’s authority to issue industry-wide rules came in *National Petroleum Refiners Ass’n v. FTC*, 482 F. 2d 672 (D.C. Cir. 1973), which upheld an FTC rule issued in 1971 requiring octane ratings to be posted on gasoline pumps.

granted it under the statute to propose rules regulating numerous industries, including eyeglasses, franchising, used cars, mobile homes, and vocational schools.

The congressional record supports the interpretation that the appointment of more pro-regulatory agency personnel was meant to spur more regulatory effort at the FTC. At the confirmation hearing of Lewis Engman, Kirkpatrick's successor as chairman of the FTC, Republican Senator Norris Cotton said that the FTC "has had a need for some kind of injection to pep it up so it would fulfill its mission."²⁷ Ted Stevens, Republican Senator from Alaska, told Engman, "I am really hopeful that ... you will become a real zealot in terms of consumer affairs and some of these big business people will complain to us that you are going too far. That would be the day, as far as I am concerned."²⁸

But while the FTC was certainly more activist than prior to the overhaul initiated by Nixon, consumer advocates criticized the FTC for failing to become "a real zealot in terms of consumer affairs."²⁹ That changed in 1977, when President Carter, at the recommendation of Ralph Nader, appointed an even more pro-regulatory FTC chairman in Michael Pertschuk, prompting *Fortune* magazine to report, "Nader's invaders were inside the gates."³⁰ Under Pertschuk, the FTC put greater emphasis on hiring committed consumer advocates onto the staff. For example, Harris and Milkis (1996, p. 178) quotes an FTC attorney as saying,

[W]ho is better, a 4.0 graduate from Harvard who engaged in no "public service" programs, or a 3.85 graduate from Harvard who ran the legal aid program or was otherwise actively involved in proconsumer programs. ... Under Chairman Pertschuk, I am confident the latter would have been chosen.

Many of the regulatory actions taken under Pertschuk were initiated under previous Republican appointed chairmen, but FTC policy under Pertschuk was decidedly more activist than under his predecessors.³¹ The most controversial rulemaking proceeding in the period entailed a proposal to

²⁷Nomination of Lewis A. Engman, To Be a Commissioner, Federal Trade Commission: Hearings Before the Senate Comm. On Commerce, 93d Cong., 1st Sess. 4 (1973).

²⁸*Id.*

²⁹U.S. Congress, Senate, Committee on Commerce, 1974. Federal Trade Commission Oversight. 93rd Cong., 2d sess., March 1, 7, 14; May 9.

³⁰Harris and Milkis (1996, p. 155).

³¹Harris and Milkis (1996, p. 177)

restrict television advertisements aimed at children. The proposal prompted a backlash against the FTC, ultimately leading to Congressional action to curtail the agency.

4.2. The Environmental Protection Agency and regulatory review. The creation of the EPA and the parallel development of centralized regulatory review by the White House further illustrate our theory. In 1969 a White House task force recommended to President Nixon that responsibility for environmental protection and natural resources management be combined in a new Department of Natural Resources.³² In response, Nixon charged his Advisory Council on Executive Organization, better known as the Ash Council, to come up with a detailed proposal. The staff of the Ash Council believed that combining environmental regulation with natural resources management would result in less vigorous protection of the environment, since the concerns of environmentalists would be overwhelmed by better organized natural resource developers. In contrast, a single-mission agency would be a more single-minded advocate for pollution control.³³ Such a concern is consistent with our theory—agency bias is most effective at mitigating shirking when the agency has a relatively focused mission and not multiple missions that are frequently at odds. Persuaded in part by this concern,³⁴ Nixon adopted the Ash Council’s proposal and created the EPA by executive order on July 9, 1970.³⁵

Nixon appointed William Ruckelshaus, an attorney from the Department of Justice, to be the first administrator of the EPA. Ruckelshaus had worked on environmental actions at the state level earlier in his career, but was largely unknown to both industry and environmentalists at the time of his appointment. At his confirmation hearing, Ruckelshaus was asked whether he would resolve statutory ambiguities in favor of “the environmental view.” He replied, “that is ... precisely what

³²Marcus (1980, p. 31).

³³Marcus (1980, pp. 34-37). This view was echoed by Senator Edmund Muskie, who argued, “If the control of pollution is assigned to those responsible for the promotion of polluting activities at the same time, we compromise our goal of environmental protection. ... The agency which sets environmental quality standards must have only one goal—protection of this and future generations against changes in the natural environment which adversely affect the quality of life.” Statement of Senator Edmund Muskie, Reorganization Plan No. 3, Creating the Environmental Protection Agency, Before the Subcommittee on Government Organization, Senate Committee on Government Operations, July 28, 1970.

³⁴Whitaker (1976, p. 55).

³⁵Reorganization Plan No. 3 (35 F.R. 15623, 84 Stat. 2086).

I would intend to do.”³⁶ Ruckleshaus quickly established his environmentalist bona fides with an aggressive campaign of enforcement of existing environmental laws against polluters,³⁷ leading the *New York Times* to refer to him as the “house liberal” in the Nixon Administration.³⁸

Ruckleshaus’s aggressive policies soon brought him into conflict with the White House. One of the most controversial policy areas was the EPA’s implementation of the 1970 Clean Air Act. The law was designed to force the EPA to take aggressive steps to reduce air pollution and achieve “healthy air” by 1975 by mandating specific short-term deadlines for the agency to issue rules and by prohibiting the EPA from considering economic costs in its rule setting. In April 1971 the EPA circulated draft guidelines for states in formulating their implementation plans for the air quality standards set by the EPA under the Act. To force the EPA to consider economic costs in its policymaking under the Act, in May 1971 the director of the Office of Management and Budget (OMB) George Shultz sent Ruckleshaus a letter informing him that EPA regulations had to be cleared through OMB and other agencies before being issued.³⁹ The following month, OMB asserted that authority by preventing EPA from publishing its guidelines in the Federal Register. Following completion of OMB review, the final guidelines published in August gave states more flexibility in implementing the air quality standards and directed states to consider the economic impact of their implementation plans.⁴⁰

Soon after Nixon’s reelection in 1972, Nixon decided to retain Ruckleshaus as administrator, but Ruckleshaus stipulated as a condition of staying that Nixon revise the system of OMB review to clarify that EPA had final authority over regulatory decisions.⁴¹ While Nixon agreed, the OMB

³⁶Nomination of William D. Ruckleshaus, Hearings Before the Committee on Public Works, United States Senate, 91st Cong. Second Session, Dec. 1 and 2, 1970.

³⁷Marcus (1980, pp. 88-90).

³⁸*New York Times*, “Government’s Pollution Fighter,” April 12, 1973.

³⁹Marcus (1980, p. 125).

⁴⁰This sequence of events was recounted at a Senate hearing called over the controversy, at which Ruckleshaus insisted that he had made the final call on the modifications. Implementation of the Clean air act amendments of 1970 (title I). Hearings before the Subcommittee on Air and Water Pollution of the Committee on Public Works, U.S. Senate, Ninety-second Congress, second session, Feb. 16, 17, 18, and 23, 1972. U.S. Government Printing Office, Washington, DC.

⁴¹Quarles (1976, p. 118)

review process continued and OMB successfully pressured EPA to modify rules in response to White House concerns.⁴²

Subsequent presidents continued the practice of OMB review of executive agencies' proposed rules. President Reagan further formalized the process by executive order and expressly authorized OMB to block regulations by directing agencies to "refrain from publishing" proposed rules until OMB's review was concluded.⁴³ Responsibility for coordinating the regulatory review process was lodged in the Office of Information and Regulatory Affairs (OIRA) within OMB. When President Clinton took office in 1993, many expected him to terminate the OIRA review process. Instead, he continued the regulatory review regime with a few minor changes, and the institution of OIRA review continues to this day.

Why do presidents from both parties find OIRA review useful when they could instead simply appoint agency heads who share the president's policy preferences? In the case of Ruckleshaus, for example, Nixon could have used a different strategy for controlling the EPA: appoint an EPA administrator who shared Nixon's policy preferences.

Our theory offers an explanation: presidents prefer to appoint a relatively biased agency head in order to mitigate shirking by the agency, and this use of agency bias creates a role for regulatory review in reducing the consequent bias to the intensive margin of policy. In the case of Ruckleshaus, Nixon wanted energetic effort out of the EPA and achieved it by appointing an agency head who put relatively large weight on the benefits of environmental regulation. The OMB review process helped keep in check the resulting bias in EPA's specific rule proposals.

5. CONCLUSION

Our goal in this paper has been to provide an account of why bureaucratic principals appoint agents with policy preferences more extreme than the principal's and subject their policy decisions to review by a more aligned bureaucrat. While our primary goal is descriptive, we conclude by

⁴²For example, Quarles (1976, pp. 117-142) recounts an episode in which OMB officials successfully pressured the EPA to delay the target date of new regulations restricting the lead content of gasoline in 1973.

⁴³Exec. Order 12,291 § 3(f), 46 Fed. Reg. 13,193 (Feb. 17, 1981).

suggesting a few implications of our analysis for the debate over the normative desirability of regulatory review.

The traditional justification for centralized regulatory review is that it keeps in check the inherent bias of agencies toward their mission. On this account, “an agency succeeds by accomplishing the goals Congress set for it as thoroughly as possible — not by balancing its goals against other, equally worthy goals” (DeMuth and Ginsburg, 1986, p. 1081). Relatedly, “capture” of regulatory agencies by special interest groups creates another source of bias, and OIRA is viewed as less vulnerable to such influence. In a recent paper, Livermore and Revesz (2012) provide a detailed analysis of the features of OIRA that insulate it from capture, pointing in particular to its generalist jurisdiction.⁴⁴

A second justification for OIRA review focuses on the value of presidential control rather than on specifically either a need to check overzealous agencies or to correct for agency capture. Kagan (2000) argues that OIRA review facilitates presidential control over the administrative state. Under her view, presidential control of the bureaucracy enhances the democratic legitimacy of bureaucratic decisionmaking because the president is elected by a national constituency. Moreover, because the president is a unitary actor in a central position within the regulatory state, presidential control results in more rational, cost-effective, and consistent bureaucratic decisionmaking.

But critics of regulatory review argue that the institution in practice has been systematically biased against regulation.⁴⁵ For example, they observe that changes made to rules during OIRA review tend to reduce the stringency of regulations,⁴⁶ and that OIRA only reviews agency decisions to act, not agency decisions not to act.⁴⁷ Moreover, critics argue that there is no compelling theory or evidence to support the hypothesis that agencies are overzealous in pursuit of their missions and therefore need to be checked via centralized review.⁴⁸

⁴⁴But the authors argue that OIRA review has in practice been biased against regulation and propose reforms that would make OIRA play an anti-capture role that would correct rules that favor special interest groups (whether they be pro-regulatory or anti-regulatory) at the expense of the broader public interest.

⁴⁵*See, e.g.*, Morrison (1985); Percival (1991).

⁴⁶Bressman and Vandenberg (2006).

⁴⁷Bagley and Revesz (2010).

⁴⁸*Id.*

Furthermore, critics of the presidential control justification for OIRA review point out that the appointments power enables presidents to choose loyalists to head agencies, leaving little role for OIRA review in furthering presidential control.⁴⁹ In addition, presidents and their senior staff only rarely directly intervene in agency rules under OIRA review.⁵⁰

In our view, centralized regulatory review is fundamentally about presidential control over the administrative state. As influentially argued by Moe (1985), presidents have strong political incentives to assert control over agencies and have instituted OIRA review as part of a set of strategies to achieve some measure of such control. The fact that every president since Nixon has retained centralized regulatory review is strong evidence that such review furthers presidential objectives. Other social benefits of OIRA review, such as checking overzealous or captured agencies, are a product of the ways that presidents have employed regulatory review to achieve their policy and political objectives.

Moreover, our analysis highlights difficulties with some of the arguments made against the presidential control view. The primary channels of presidential control over agency rulemaking can be grouped into two categories: (1) direct presidential policy decisions on specific rules; and (2) presidential decisions on the structure of the administrative state and its personnel. In our model of regulatory review, the president makes no direct policy decisions. Rather, the president achieves her policy objectives by choosing what bureaucrats to appoint and by structuring their decision-making process. As the model makes clear, in terms of achievement of presidential policy objectives, a system with regulatory review strictly dominates full delegation. Hence the frequency with which the president personally intervenes in rulemakings under OIRA review is not an accurate measure of the degree to which OIRA review facilitates presidential control. In fact, we show that a president who has the opportunity to personally intervene actually does strictly worse than one who can commit to delegating the policy decision to a reviewer.

⁴⁹Livermore and Revesz (2012, p. 13).

⁵⁰Livermore and Revesz (2012, p. 14) argue that “for the vast majority of OIRA’s work—the bulk of the regulatory iceberg that is submerged below the gaze of the President or other senior political officials—the presidential power justification for OIRA review is largely irrelevant.”

Furthermore, our analysis shows that ideological conflicts between OIRA and the regulatory agencies can be useful to, and indeed are sometimes intentionally created by, the president. When the president does not have too strong a deregulatory motive, the president has an incentive to appoint pro-regulatory biased agency staff and subject them to review by a more aligned bureaucrat at OIRA. This helps explain why in practice OIRA review more often leads to a reduction rather than an increase in the stringency of agencies' proposed rules. Moreover, it also explains why OIRA is structured to review only agency action, not agency decisions not to act. In our theory, the optimal division of labor is to give more ideologically motivated staff in the agencies the responsibility to research and propose rules. The comparative advantage of more centrist staff at the center is in reviewing proposed agency action, not in researching the many ways in which the agency has chosen not to act. Because OIRA is fundamentally about presidential control, the normative debate about the desirability of OIRA review should focus on the normative attractiveness of presidential control—an important and immensely complicated issue implicating, *inter alia*, separation of powers concerns and the nature of electoral politics—not on whether OIRA review is systematically biased.

REFERENCES

- ACS, A., AND C. CAMERON (2012): "Regulatory Auditing at the Office of Information and Regulatory Affairs," Unpublished Manuscript.
- AGHION, P., AND J. TIROLE (1997): "Formal and real authority in organizations," *Journal of Political Economy*, pp. 1–29.
- AMERICAN BAR ASSOCIATION (1969): *Report of the ABA Commission To Study the Federal Trade Commission*.
- BAGLEY, N., AND R. REVESZ (2010): "Centralized oversight of the regulatory state," *Columbia Law Review*, 106, 1260.
- BENDOR, J., A. GLAZER, AND T. HAMMOND (2001): "Theories of delegation," *Annual review of political science*, 4(1), 235–269.
- BENDOR, J., AND A. MEIROWITZ (2004): "Spatial models of delegation," *American Political Science Review*, 98(2), 293–310.
- BESLEY, T., AND M. GHATAK (2005): "Competition and incentives with motivated agents," *The American Economic Review*, 95(3), 616–636.
- BRESSMAN, L., AND M. VANDENBERGH (2006): "Inside the Administrative State: A Critical Look at the Practice of Presidential Control," *Michigan Law Review*, 105, 47.

- BUENO DE MESQUITA, E., AND M. STEPHENSON (2007): "Regulatory quality under imperfect oversight," *American Political Science Review*, 101(03), 605–620.
- CALVERT, R., M. MCCUBBINS, AND B. WEINGAST (1989): "A theory of political control and agency discretion," *American Journal of Political Science*, pp. 588–611.
- CLARKSON, K. W., AND T. J. MURIS (1981): *The Federal Trade Commission since 1970: economic regulation and bureaucratic behavior*. Cambridge University Press, Cambridge.
- COX, E. F., R. C. FELLMETH, AND J. E. SCHULZ (1969): *The Nader report on the Federal Trade Commission*. R. W. Baron, New York.
- DEMUTH, C., AND D. GINSBURG (1986): "White House review of agency rulemaking," *Harvard Law Review*, 99(5), 1075–1088.
- DESSEIN, W. (2002): "Authority and communication in organizations," *The Review of Economic Studies*, 69(4), 811–838.
- DEWATRIPONT, M., AND J. TIROLE (1999): "Advocates," *Journal of Political Economy*, pp. 1–39.
- DOWNES, A. (1967): *Inside bureaucracy*. Little, Brown, Boston.
- EDLIN, A., AND C. SHANNON (1998): "Strict monotonicity in comparative statics," *Journal of Economic Theory*, 81(1), 201–219.
- GAILMARD, S., AND J. PATTY (2007): "Slackers and zealots: Civil service, policy discretion, and bureaucratic expertise," *American Journal of Political Science*, 51(4), 873–889.
- GILLIGAN, T., AND K. KREHBIEL (1987): "Collective decisionmaking and standing committees: An informational rationale for restrictive amendment procedures," *Journal of Law, Economics, & Organization*, 3, 287.
- GROSSMAN, S., AND O. HART (1986): "The costs and benefits of ownership: A theory of vertical and lateral integration," *The Journal of Political Economy*, pp. 691–719.
- HARRIS, R. A., AND S. M. MILKIS (1996): *The politics of regulatory change: a tale of two agencies*. Oxford University Press, New York, 2nd ed edn.
- HOLMSTROM, B., AND P. MILGROM (1991): "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design," *Journal of Law, Economics, & Organization*, 7, 24.
- KAGAN, E. (2000): "Presidential administration," *Harvard Law Review*, 114, 2245.
- LIVERMORE, M., AND R. REVESZ (2012): "OIRA, Regulatory Review, and Capture," Unpublished Manuscript.
- MARCUS, A. A. (1980): *Promise and performance: choosing and implementing an environmental policy*, vol. no. 39. Greenwood Press, Westport, Conn.
- MCCUBBINS, M., R. NOLL, AND B. WEINGAST (1987): "Administrative procedures as instruments of political control," *Journal of Law, Economics, & Organization*, 3(2), 243–277.
- MILGROM, P., AND C. SHANNON (1994): "Monotone comparative statics," *Econometrica*, pp. 157–180.
- MOE, T. (1985): "The Politicized Presidency," in *The New Direction in American Politics*, ed. by J. E. Chubb, and P. E. Peterson. Brookings Institution, Washington, D.C.
- MOE, T., AND S. WILSON (1994): "Presidents and the Politics of Structure," *Law & Contemporary Problems*, 57, 1.
- MORRISON, A. (1985): "OMB interference with agency rulemaking: the wrong way to write a regulation," *Harvard Law Review*, 99, 1059.
- NISKANEN, W. (1975): "Bureaucrats and politicians," *Journal of Law and Economics*, 18(3), 617–643.

- PERCIVAL, R. (1991): “Checks Without Balance: Executive Office Oversight of the Environmental Protection Agency,” *Law & Contemporary Problems*, 54, 127.
- POSNER, E. (2001): “Controlling agencies with cost-benefit analysis: A positive political theory perspective,” *University of Chicago Law Review*, 68, 1137.
- PRENDERGAST, C. (2007): “The motivation and bias of bureaucrats,” *The American Economic Review*, pp. 180–196.
- QUARLES, J. (1976): *Cleaning up America: an insider’s view of the Environmental Protection Agency*. Houghton Mifflin, Boston.
- STEPHENSON, M. (2007): “Bureaucratic Decision Costs and Endogenous Agency Expertise,” *Journal of Law, Economics, & Organization*, 23(2), 469–498.
- (2011): “Information Acquisition and Institutional Design,” *Harvard Law Review*, 124, 1422–1829.
- TIROLE, J. (1994): “The internal organization of government,” *Oxford Economic Papers*, pp. 1–29.
- WHITAKER, J. C. (1976): *Striking a balance: environment and natural resources policy in the Nixon-Ford years*. American Enterprise Institute for Public Policy Research, Washington.

APPENDIX

Proof of Proposition 1. We begin by taking the derivative of the Principal’s objective function in (8) with respect to k_A , ignoring the $k_A \leq k^{max}$ constraint, which yields the first-order condition,

$$(20) \quad \left[e^{*'}(k^*)(k_P k^* - \frac{k^{*2}}{2}) + e^*(k^*)(k_P - k^*) \right] B^2 = 0.$$

Replacing for the $e^*(k^*)$ and $e^{*'}(k^*)$, this becomes,

$$(21) \quad \left[\frac{k^* B^2}{A} (k_P k^* - \frac{k^{*2}}{2}) + \frac{k^{*2} B^2}{2A} (k_P - k^*) \right] B^2 = 0,$$

which simplifies to,

$$(22) \quad k_P - \frac{k^*}{2} + \frac{k_P}{2} - \frac{k^*}{2} = 0.$$

Solving for k^* gives the result that $k^* = \frac{3}{2}k_P$. The second-order condition is everywhere negative, so the Principal will prefer the boundary if the unconstrained optimum exceeds it. \square

Proof of Proposition 2.

(1) The derivative of the Principal’s objective function in (10) with respect to k_A is,

$$(23) \quad e_1^*(k_A, k_R) V(k_P, k_R)$$

$e^* = \frac{\max\{0, V(k_A, k_R)\}}{A}$, and $V(k, k_R) > 0$ whenever $k_R > k/2$, so for any k_R the Principal would consider selecting, this expression is equal to,

$$(24) \quad \frac{k_R B^2}{A} V(k_P, k_R).$$

It is, therefore, strictly positive, which implies that $k_A^* = k^{max}$.

(2) The derivative of the Principal’s objective function in (10) with respect to k_R is,

$$e_2^*(k_A, k_R) V(k_P, k_R) + e^*(k_A, k_R) V_2(k_P, k_R).$$

Substituting in for these functions and for $k_A = k^{max}$, this derivative becomes,

$$\begin{aligned}
(25) \quad & [(k^{max} - k_R)(k_P k_R - \frac{k_R^2}{2}) + (k_P - k_R)(k^{max} k_R - \frac{k_R^2}{2})] \frac{B^4}{A} = \\
& [(k^{max} - k_R)(k_P - \frac{k_R}{2}) + (k_P - k_R)(k^{max} - \frac{k_R}{2})] \frac{k_R B^4}{A} = \\
& [k^{max} k_P - \frac{k^{max} k_R}{2} - k_R k_P + \frac{k_R^2}{2} + k^{max} k_P - \frac{k_R k_P}{2} - k^{max} k_R + \frac{k_R^2}{2}] \frac{k_R B^4}{A} = \\
& [2k^{max} k_P - \frac{3k^{max} k_R}{2} - \frac{3k_R k_P}{2} + k_R^2] \frac{k_R B^4}{A}.
\end{aligned}$$

Evaluating this at $k_R = k_P$, we get,

$$\begin{aligned}
(26) \quad & [2k^{max} k_P - \frac{3k^{max} k_P}{2} - \frac{3k_P^2}{2} + k_P^2] \frac{k_P B^4}{A} \\
& = [\frac{k^{max} k_P}{2} - \frac{k_P^2}{2}] \frac{k_P B^4}{A},
\end{aligned}$$

which is strictly greater than 0, so $k_R^* > k_P$.

Evaluating this derivative at $k_R = k^{max}$ yields,

$$\begin{aligned}
(27) \quad & [2k^{max} k_P - \frac{3k^{max2}}{2} - \frac{3k^{max} k_P}{2} + k^{max2}] \frac{k^{max} B^4}{A} \\
& = [\frac{k^{max} k_P}{2} - \frac{k^{max2}}{2}] \frac{k^{max} B^4}{A},
\end{aligned}$$

which is strictly less than 0, so $k_R^* < k^{max}$.

To show that $\frac{\partial k_R^*}{\partial k_P} > 0$, we apply the implicit function theorem. The first-order condition for the Principal's choice of k_R is,

$$(28) \quad g(k_R; k_P) \equiv [2k^{max} k_P - \frac{3k^{max} k_R}{2} - \frac{3k_R k_P}{2} + k_R^2] \frac{k_R B^4}{A} = 0$$

Because the second-order condition is true, we need only show that $\frac{\partial g(k_R; k_P)}{\partial k_P} > 0$. We have that,

$$(29) \quad \frac{\partial g(k_R; k_P)}{\partial k_P} = [2k^{max} - \frac{3k_R}{2}] \frac{k_R B^4}{A} > 0.$$

(3) If $k^* = k^{max}$ this result follow from part (2). So consider the case when $k^* = \frac{3}{2}k_P$. The derivative of the Principal's objective function with respect to k_R , provided in (30), evaluated at $k_R = \frac{3}{2}k_P$, is,

$$(30) \quad [2k^{max} k_P - \frac{9k^{max} k_P}{4} - \frac{9k_P^2}{4} + \frac{9k_P^2}{4}] \frac{3k_P B^4}{2A} = [-\frac{k^{max} k_P}{4}] \frac{3k_P B^4}{2A},$$

which is strictly negative, which implies that $k_R^* < \frac{3}{2}k_P = k^*$.

(4) Observe that setting $k_A = k_R = k^*$ would replicate the outcome from full delegation, but the characterization above shows it is always possible to improve on that choice by either increasing k_A (if $k^* < k^{max}$) or decreasing k_R (if $k^* = k^{max}$), or both.

□

Proof of Proposition 3.

- (1) The proof is identical to the proof of part (1) of Proposition 2
- (2) In the case in which the Principal can commit to delegating policy authority to a separate Reviewer analyzed above, the Principal could choose the no commitment outcome by setting $k_R = k_P$. But as we proved above, the Principal never chooses to do so, and so she must be strictly better off in that case than without policy commitment.
- (3) The Principal's payoff in partial delegation without policy commitment is given by

$$W^{partial} = \frac{B^4}{A} \left(k^{max} k_P - \frac{k_P^2}{2} \right) \frac{k_P^2}{2}.$$

The Principal's payoff in the case of full delegation is $e^*(k^*)V(k_P, k^*)$.

Consider first the case where $k^* = k^{max}$, in which case full delegation yields the principal an expected payoff of

$$W_1^{full} = \frac{B^4}{A} \frac{(k^{max})^2}{2} \left(k^{max} k_P - \frac{(k^{max})^2}{2} \right).$$

It's useful, here, to replace for $k^{max} = Mk_P$, where $M \geq 1$. With this replacements, we can write

$$\begin{aligned} (31) \quad W^{partial} - W_1^{full} &= \frac{B^4 k_P^{*4}}{4A} [2M - 1 - 2M^3 + M^4] \\ &= \frac{B^4 k_P^{*4}}{4A} [(M - 1)^2 (M - 1) (M + 1)] > 0 \end{aligned}$$

Consider now the case where $k^* = \frac{3}{2}k_P < k^{max}$, in which case full delegation yields the principal an expected payoff of

$$W_2^{full} = \frac{B^4}{A} \frac{9k_P^2}{8} \left(\frac{3}{2}k_P^2 - \frac{9}{8}k_P^2 \right)$$

Making the same substitution as before,

$$\begin{aligned} (32) \quad W^{partial} - W_2^{full} &= \frac{B^4 k_P^{*4}}{4A} \left[2M - 1 - \frac{9}{2} \left(\frac{3}{2} - \frac{9}{8} \right) \right] \\ &= \frac{B^4 k_P^{*4}}{4A} \left[2M - \frac{43}{16} \right]. \end{aligned}$$

So this difference is positive whenever $M > \frac{43}{32}$. But to ever be in this case, it must be that $k^{max} > \frac{3}{2}k_P$, and so $M > \frac{3}{2} > \frac{43}{32}$. Thus the difference is always positive, and the result holds unambiguously for both cases.

□

Proof of Lemma 1. Full disclosure is an equilibrium iff

$$\left[k_A k_R - \frac{k_R^2}{2} \right] B_j^2 > \left[k_A B_j k_R \bar{B} - \frac{k_R^2 \bar{B}^2}{2} \right]$$

for $j \in \{H, L\}$. We can rewrite this condition as $k_A k_R B_j [B_j - \bar{B}] > \frac{k_R^2}{2} [B_j^2 - \bar{B}^2]$. Since $B_H > \bar{B}$, for B_H this simplifies to

$$k_A > k_R \frac{B_H + \bar{B}}{2B_H},$$

which can be rewritten as the lower bound on k_A in the lemma. expression in the text. Since $B_L < \bar{B}$, for B_L this simplifies to

$$k_A < k_R \frac{B_L + \bar{B}}{2B_L},$$

which can similarly be rewritten as the upper bound on k_A in the lemma.

Proof of Proposition 4. PROOF IN PROGRESS.

Proof of Proposition 5. PROOF IN PROGRESS.